



**The Anh Han**  
Mestre em Lógica Computacional

## **Intention Recognition, Commitment and Their Roles in the Evolution of Cooperation**

Dissertação para obtenção do Grau de Doutor em  
Informática

Orientador: Prof. Doutor Luís Moniz Pereira  
Professor Catedrático  
Universidade Nova de Lisboa

Júri:

Presidente: Prof. Doutor Pedro Manuel Corrêa Calvente Barahona

Arguentes: Prof<sup>a</sup>. Doutora Fariba Sadri  
Prof. Doutor Tom Lenaerts

Vogais: Prof. Doutor Carlos Augusto Isaac Piló Viegas Damásio  
Prof. Doutor Francisco João Duarte Cordeiro Correia dos Santos  
Prof. Doutor Luís Manuel Sancho Moniz Pereira



FACULDADE DE  
CIÊNCIAS E TECNOLOGIA  
UNIVERSIDADE NOVA DE LISBOA

**Maio de 2012**





**The Anh Han**  
Mestre em Lógica Computacional

## **Intention Recognition, Commitment and Their Roles in the Evolution of Cooperation**

Dissertação para obtenção do Grau de Doutor em  
Informática

Orientador: Prof. Doutor Luís Moniz Pereira  
Professor Catedrático  
Universidade Nova de Lisboa

Júri:

Presidente: Prof. Doutor Pedro Manuel Corrêa Calvente Barahona

Arguentes: Prof<sup>a</sup>. Doutora Fariba Sadri  
Prof. Doutor Tom Lenaerts

Vogais: Prof. Doutor Carlos Augusto Isaac Piló Viegas Damásio  
Prof. Doutor Francisco João Duarte Cordeiro Correia dos Santos  
Prof. Doutor Luís Manuel Sancho Moniz Pereira

 FACULDADE DE  
CIÊNCIAS E TECNOLOGIA  
UNIVERSIDADE NOVA DE LISBOA

**Maio de 2012**



**Copyright** © 2012 The Anh Han, FCT-UNL, UNL

A Faculdade de Ciências e Tecnologia e a Universidade Nova de Lisboa tem o direito, perpétuo e sem limites geográficos, de arquivar e publicar esta dissertação através de exemplares impressos reproduzidos em papel ou de forma digital, ou por qualquer outro meio conhecido ou que venha a ser inventado, e de a divulgar através de repositórios científicos e de admitir a sua cópia e distribuição com objectivos educacionais ou de investigação, não comerciais, desde que seja dado crédito ao autor e editor.



*Dành tặng bố mẹ kính yêu*

*For my parents*



# Acknowledgements

First of all, I would like to express my deepest gratitude and appreciation to my supervisor, Prof. Luís Moniz Pereira, who has worked closely with me all these years. For he is knowledgeable in so many research areas, I cannot recall how much I have learnt from him through an enormous number of discussions (not only in this PhD, but also during my Master study). His genuine enthusiasm for new research ideas and his endless desire to sort them out show me an example of an extraordinarily dedicated scientist, which, I believe, will strongly influence my research approach and attitudes for the rest of my career. I truly enjoy and appreciate the many long discussions that he always generously arranges whenever I wish. Furthermore, I am profoundly indebted to his constant (and always speedily) support in all matters, not only in my studies, but also in the everyday life ones.

Besides my supervisor, I am most grateful to Francisco Santos, who has been, not only providing me with the significant guidance for the evolutionary part of my PhD research, but also exhibiting essential support and warm friendship. I truly appreciate all the precious scientific connections Francisco has helped to establish, which are undoubtedly crucial for my career, and our several interesting discussions and collaborations during these years.

I am very grateful to the members of the Thesis Plan Advisory Committee, Carlos Damásio, Fariba Sadri, and Francisco Santos, who have provided me with essential suggestions to converge into an appropriate plan for my PhD research. I also thank them, plus Tom Lenaerts, for useful suggestions as members of my thesis program examination committee. I thank the course lecturers during my PhD studies, Nuno Marques, Carlos Damásio, Luís Caires, José Alferes, João Leite, for providing me with essential scientific background. I also thank Luís Caires, Nuno Correia, Pedro Barahona, and José Alferes, who were always of great help with respect to the administration issues.

I acknowledge the Portuguese national foundation of science and technologies (FCT-MCTES) for giving me the PhD grant (ref. SFRH/BD/62373/2009), which enables me to focus on doing research. I acknowledge the Department of Informatics, Faculty of Sciences and Technologies, New University of Lisbon and the Center for Artificial Intelligence (CENTRIA) for providing me with a place and tools to work.

I thank Sandra Raínha, Filipa Reis and Anabela Duarte for helping me with the several bureaucratic issues and with arranging my travels during these years.

I would like to thank Arne Traulsen for providing me with a nice place for a research visit in his group at the Max Planck institute for Evolutionary Biology. I had a great time there, enabled by the great hospitality of all the group members, whom I truly thank. I thank Francisco Santos and Jorge Pacheco for their help with making this connection possible. I also truly thank Jorge (and his ATP group) for allowing me to use their cluster to run several simulations during my PhD study.

I am very thankful to my friend Martin Slota, for being a wonderful apartment-mate and companion in this PhD journey. I enjoy the many discussions, scientific as well as regarding many other things, and joyful activities we had, which help to balance my life. I thank all my wonderful friends at FCT/UNL, João Moura, Ana Sofia Gomes, Ari Saptawijaya, Alexandre Pinto, Venil Elangovan, who are always very supportive whenever I need help. I also thank my Vietnamese friends, Dung, Toan, Trang, Xiem, Cuong, Giang, Khoi, Nghia, Trinh, who make me less homesick via several activities, especially those during the Vietnamese national holidays.

When writing up this thesis, beside my supervisor Luís, who always rapidly and diligently provided me with his valuable comments and suggestions, chapter by chapter, whom I heartedly appreciate, a final draft of this thesis has been read by Francisco Santos, Ari Saptawijaya, Julián García, and Chaitanya Gokhale. I appreciate their important feedbacks, which have turned this work into a better one. I truly thank Ana Sofia Gomes for her kind help with translating the thesis abstract into Portuguese. Moreover, during my PhD research, several manuscripts have been improved with the useful comments and suggestions from (when not being a co-author) Fariba Sadri, Martin Slota, Ana Sofia Gomes and Ari Saptawijaya.

Finally, none of this would have been possible without the love and support of my family. My wife, Ly, has always been by my side, loving and understanding me, and constantly supporting me through this challenging period. My wife and my son, Vinh, who was born during my PhD study, have always been my greatest inspiration and motivation to become stronger and work harder, so that I can be with them soon. And above all, I am most indebted to my parents, who, throughout my life, have always been of the greatest and unconditional support in all matters; and who have always encouraged me to pursue higher studies and a career in science. This thesis is dedicated to my parents.

# Abstract

---

The goal of this thesis is twofold. First, intention recognition is studied from an Artificial Intelligence (AI) modeling perspective. We present a novel and efficient intention recognition method that possesses several important features: (i) The method is context-dependent and incremental, enabled by incrementally constructing a three-layer Bayesian network model as more actions are observed, and in a context-dependent manner, relying on a logic programming knowledge base concerning the context; (ii) The Bayesian network is composed from a knowledge base of readily specified and readily maintained Bayesian network fragments with simple structures, enabling an efficient acquisition of the corresponding knowledge base (either from domain experts or else automatically from a plan corpus); and, (iii) The method addresses the issue of intention change and abandonment, and can appropriately resolve the issue of multiple intentions recognition. Several aspects of the method are evaluated experimentally, achieving some definite success. Furthermore, on top of the intention recognition method, a novel framework for intention-based decision making is provided, illustrating several ways in which an ability to recognize intentions of others can enhance a decision making process.

A second subgoal of the thesis concerns that, whereas intention recognition has been extensively studied in small scale interactive settings, there is a major shortage of modeling research with respect to large scale social contexts, namely evolutionary roles and aspects of intention recognition. Employing our intention recognition method and the tools of evolutionary game theory, this thesis explicitly addresses the roles played by intention recognition in the final outcome of cooperation in large populations of self-regarding individuals. By equipping individuals with the capacity for assessing intentions of others in the course of social dilemmas, we show how intention recognition is selected by natural selection, opening a window of opportunity for cooperation to thrive, even in hard cooperation prone games like the Prisoner's Dilemma.

In addition, there are cases where it is difficult, if not impossible, to recognize the intentions of another agent. In such cases, the strategy of proposing commitment, or of intention manifestation, can help to impose or clarify the intentions of others.

Again using the tools of evolutionary game theory, we show that a simple form of commitment can lead to the emergence of cooperation; furthermore, the combination of commitment with intention recognition leads to a strategy better than either one by itself.

**How the thesis should be read?** We recommend that the thesis be read sequentially, chapter by chapter [**1-2-3-4-5-6-7-8**].

However, for those more interested in intention recognition from the AI modeling perspective, i.e. the first subgoal of the thesis, Chapters 6 and 7 can be omitted and Chapters 4 and 5 are optional [**1-2-3-(4)-(5)-8**].

In addition, for those more keen on the problem of the evolution of cooperation, i.e. the second subgoal of thesis, Chapter 3 and even Chapter 2, can be omitted [**1-(2)-4-5-6-7-8**].

**Keywords:** Intention Recognition, Commitment, Evolution of Cooperation, Evolutionary Game Theory, Prisoner's Dilemma, Bayesian Network, Logic Programming, Evolution Prospection, Decision Making.

---

# Sumário

---

Nesta tese é estudado o Reconhecimento de Intenções definido numa perspectiva de modelação da Inteligência Artificial (IA). Neste contexto, esta tese é composta por dois objectivos principais. Como primeiro objectivo, apresentamos um novo e eficiente método de Reconhecimento de Intenções que possui propriedades importantes: (i) O método é dependente do contexto e incremental, permitindo a construção de uma rede Bayesiana de três camadas. Assim, esta rede é construída incrementalmente à medida que mais acções são observadas, e dependente do contexto, já que depende de uma base de conhecimento definida em programação lógica, integrando a informação sobre contexto; (ii) A rede Bayesiana é construída através de uma base de conhecimento composta por vários fragmentos de rede Bayesiana, fragmentos estes com uma estrutura simples, permitindo assim um processo eficiente de construção desta base de conhecimento (tanto através de informação fornecida por especialistas da área ou automaticamente, através de um *plan corpus*); (iii) O método, para além de considerar a questão da mudança e abandono de intenções, consegue resolver correctamente o problema do reconhecimento de múltiplas intenções. Por outro lado, diversos aspectos deste método são avaliados experimentalmente, demonstrando resultados claros do sucesso da solução proposta. Adicionalmente, em cima deste método de reconhecimento de intenções, propomos uma nova framework para processos de decisão baseados em intenções, ilustrando diversos exemplos nos quais a capacidade de reconhecer as intenções de outrem pode melhorar o processo de decisão.

O segundo objectivo desta tese está relacionado com o estudo dos vários aspectos de reconhecimento de intenções no contexto da teoria de jogos evolucionária. Se, por um lado, o reconhecimento de intenções tem sido extensivamente estudado em contextos pequenos e interactivos, existe uma escassez de trabalho relacionado de modelação de contextos sociais de larga escala, nomeadamente no que diz respeito ao papel do reconhecimento de intenções nos modelos de evolução destes contextos. Utilizando o nosso método de reconhecimento de intenções em conjunto com ferramentas de teoria de jogos evolucionária, esta tese permite estudar o papel que este reconhecimento de intenções possui no resultado final da cooperação de agentes

individualistas em grandes populações. Deste modo, ao equiparmos indivíduos com a capacidade de avaliar as intenções de outrem na resolução de dilemas sociais, mostramos como o reconhecimento de intenções é utilizado no processo de selecção natural, abrindo assim uma nova janela de oportunidades para o aparecimento de cooperação entre indivíduos, mesmo em jogos de difícil cooperação como o Dilema do Prisioneiro.

No entanto, existem diversos casos onde é bastante difícil, se não impossível, de reconhecer as intenções de outro agente. Nestes casos, a estratégia de propor um compromisso ou manifestação de intenções pode ajudar a impor, ou a clarificar as intenções dos restantes agentes. Mais uma vez, utilizando as técnicas da teoria de jogos evolucionária, demonstramos como uma simples proposta de compromisso pode levar à emergência de cooperação; e, ainda, que a combinação de compromissos com o reconhecimento de intenções resulta numa melhor estratégia que qualquer uma delas em separado.

**Como é que a tese deve ser lida?** A nossa recomendação é que a tese deve ser lida sequencialmente, capítulo a capítulo [~~1-2-3-4-5-6-7-8~~].

No entanto, para os leitores particularmente interessados no reconhecimento de intenções de uma perspectiva de modelação de IA, i.e. o primeiro objectivo desta tese, os Capítulos 6-7 podem ser omitidos e os Capítulos 4-5 são opcionais [~~1-2-3-(4)-(5)-8~~].

Por último, para os leitores especialmente focados no problema da evolução da cooperação, i.e. no segundo objectivo desta tese, o Capítulo 3, e até o Capítulo 2, podem ser omitidos [~~1-(2)-4-5-6-7-8~~].

**Palavras-chave:** Reconhecimento de Intenções, Compromisso, Teoria de Jogos Evolucionária, Emergência da Cooperação, Dilema do Prisioneiro, Rede Bayesiana, Programação em Lógica, Prospecção da Evolução, Processos de Decisão.

---

# Contents

<b>I</b>	<b>Intention Recognition</b>	<b>1</b>
<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Motivations . . . . .	3
1.1.1	Intention Recognition in AI . . . . .	4
1.1.2	Evolutionary, Psychological, and Philosophical Aspects of Intention Recognition . . . . .	6
1.1.3	Commitment and Intention Recognition . . . . .	9
1.2	Aim and Scope . . . . .	10
1.3	Thesis Contributions . . . . .	11
1.4	Outline of the Thesis . . . . .	12
<b>2</b>	<b>Intention Recognition Methods</b>	<b>15</b>
2.1	Introduction . . . . .	16
2.2	Related Work . . . . .	19
2.3	Bayesian Networks . . . . .	21
2.4	Incremental Intention Recognition . . . . .	22
2.4.1	Operators for constructing IRBNs . . . . .	27
2.5	Relation Among Intentions . . . . .	32
2.6	Single Intention Recognition . . . . .	33
2.6.1	The Model . . . . .	33
2.6.2	Experiments with Linux Plan Corpus . . . . .	35
2.7	IPD Plan Corpora . . . . .	40
2.7.1	Iterated Prisoner’s Dilemma . . . . .	40
2.7.2	Corpus Description . . . . .	41
2.7.3	Plan Corpora Generation . . . . .	43

## CONTENTS

---

2.7.4	Social Learning . . . . .	44
2.7.5	Results on IPD Plan Corpora . . . . .	48
2.7.6	Scope Enlargement Discussion . . . . .	49
2.8	More on Situation Sensitiveness . . . . .	51
2.8.1	Situation-sensitive Bayesian Networks . . . . .	51
2.8.2	Situation-sensitive selection of unit fragments . . . . .	55
2.9	Conclusions and Future Work . . . . .	58
<b>3</b>	<b>Intention-based Decision Making and Applications</b>	<b>61</b>
3.1	Introduction . . . . .	62
3.2	Evolution Prospection . . . . .	63
3.2.1	Preliminary . . . . .	63
3.3	Evolution Prospection with Intention Recognition . . . . .	67
3.3.1	Intentions Triggering Active Goals . . . . .	68
3.3.2	Intention Triggering Preferences . . . . .	69
3.3.3	Hostile setting . . . . .	73
3.4	Ambient Intelligence in the Home Environment and Elder Care . . . . .	75
3.4.1	Proactive Support . . . . .	76
3.4.2	Security and Emergency . . . . .	80
3.4.3	Discussion of Other AmI Issues . . . . .	84
3.5	Conclusions and Future Work . . . . .	85
<b>II</b>	<b>Intention Recognition, Commitment, and the Evolution of Cooperation</b>	<b>87</b>
<b>4</b>	<b>Intention Recognition Promotes the Emergence of Cooperation: A Bayesian Network Model</b>	<b>97</b>
4.1	Introduction . . . . .	98
4.2	Bayesian Network for Intention Recognition in Social Dilemmas . . . . .	99
4.3	Results . . . . .	102
4.3.1	Evolution of short-memory intention recognizers . . . . .	103
4.3.2	The role of memory size . . . . .	107
4.3.3	The role of IRs optimism and their tolerance to noise . . . . .	108
4.4	Concluding Remarks . . . . .	110
4.5	Appendices . . . . .	112

4.5.1	Decision making with intention recognition . . . . .	112
4.5.2	Memory-Two Intention Recognizers . . . . .	113
4.5.3	Comparing IR, WSLS, and TFT . . . . .	116
<b>5</b>	<b>Intention Recognition Promotes The Emergence of Cooperation: A Corpus-based Model</b>	<b>119</b>
5.1	Introduction . . . . .	120
5.2	Materials and Methods . . . . .	122
5.2.1	Plan Corpus Description . . . . .	122
5.2.2	Corpus-based Intention Recognition Model . . . . .	123
5.2.3	Intention-based Decision Making Model . . . . .	124
5.3	Models Acquisition and Evaluation . . . . .	126
5.3.1	Plan Corpus Generation . . . . .	126
5.3.2	Intention Recognition Model . . . . .	127
5.3.3	Decision Making Model . . . . .	128
5.4	Experiments and Results . . . . .	130
5.4.1	Analysis . . . . .	130
5.4.2	Evolutionary Simulations . . . . .	132
5.4.3	Intensities of Selection . . . . .	132
5.5	Discussion – two models of intention recognition for the evolution of cooperation . . . . .	134
5.6	Concluding Remarks . . . . .	136
<b>6</b>	<b>The Emergence of Commitments and Cooperation</b>	<b>139</b>
6.1	Introduction . . . . .	140
6.2	Related Work . . . . .	143
6.3	Models and Methods . . . . .	144
6.3.1	Model . . . . .	144
6.3.2	Methods . . . . .	146
6.4	Results . . . . .	148
6.5	Discussion . . . . .	154
<b>7</b>	<b>Intention Recognition, Commitment and the Evolution of Cooper- ation</b>	<b>157</b>
7.1	Introduction . . . . .	158

## CONTENTS

---

7.2	Models . . . . .	160
7.2.1	Commitment variant of the Prisoner's Dilemma . . . . .	161
7.2.2	Combination of intention recognition and commitment . . . . .	162
7.3	Results . . . . .	167
7.4	Discussion . . . . .	169
7.5	Conclusion . . . . .	170
<b>8</b>	<b>Conclusions and Future Work</b>	<b>173</b>
8.1	Summary and Conclusions . . . . .	173
8.2	Future works . . . . .	176
8.2.1	Part I: Intention Recognition Methods . . . . .	176
8.2.2	Part II: The Evolution of Cooperation . . . . .	177

# List of Tables

2.1	Intention Recognition Results on the Linux Plan Corpus . . . . .	39
5.1	Intention recognition results for each strategy and for the totality. .	127



# List of Figures

2.1	General structure of a Bayesian network for intention recognition. . .	23
2.2	Elder Intentions Recognition IRBN . . . . .	24
2.3	Two Types of Unit Fragments. . . . .	26
2.4	Noisy-OR Combination Method . . . . .	28
2.5	Noisy-OR Combination for the node <i>Looking</i> . . . . .	29
2.6	Plot of the incremental intention recognition method’s precision and convergence for the Linux Plan Corpus. . . . .	38
2.7	Bayesian Network for intention recognition (IRBN) in the context of iterated Prisoner’s Dilemma, where players can change their initial intention by imitating successful others. . . . .	45
2.8	Precision and convergence of the incremental intention recognition on <code>Testset-IRFIX</code> dataset. . . . .	47
2.9	Plot of the precision of the context-dependent intention recognition for three different levels of contextual information with respect to <code>Testset-IRCHANGE</code> dataset. . . . .	48
3.1	Tea-Coffee Considering Intentions: A <i>Priori</i> Preferences . . . . .	70
3.2	Tea-Coffee Considering Intentions: A <i>Posteriori</i> Preferences . . . . .	70
3.3	Example: Saving or Leaving . . . . .	71
3.4	Fox’s intentions recognition Bayesian Network. . . . .	72
3.5	Elder Care Example: Suggestion for a Drink . . . . .	78
3.6	Elder’s intentions recognition Bayesian Network . . . . .	83
4.1	Bayesian Network for Intention Recognition in Social Dilemmas . . .	99
4.2	Transition probabilities and stationary distributions in a population of <i>AllCs</i> , <i>AllDs</i> and either <i>WSLSs</i> , <i>TFTs</i> or <i>IRs</i> . . . . .	104

LIST OF FIGURES

---

4.3	Percentage of time spent at <i>AUC</i> , <i>AUD</i> and either <i>WSLS</i> , <i>TFT</i> or <i>IR</i> .	105
4.4	Stationary distribution in a population of five strategies <i>AUC</i> , <i>AUD</i> , <i>WSLS</i> , <i>TFT</i> and <i>IR</i> .	106
4.5	Probabilities of reactions computed numerically for <i>IRs</i> with different memory sizes and analogous quantities for <i>WSLS</i> and <i>TFT</i> strategists, as a function of the level of noise.	108
5.1	Intention recognizers' decision making model for different values of the discount factor $\alpha$ .	128
5.2	Interactions of <i>IR</i> with <i>AUC</i> , <i>AUD</i> , <i>TFT</i> , <i>WSLS</i> and another <i>IR</i> , in the absence of noise and $\alpha = 1$ .	131
5.3	Simulation results for Donation game in different scenarios, including populations of three strategies, <i>AUC</i> , <i>AUD</i> and either <i>IR</i> , <i>TFT</i> or <i>WSLS</i> , and populations of all those strategies together.	133
5.4	Stationary distribution (in percentage) of each strategy depending on the intensity of selection, in a population consisting of five strategies <i>AUC</i> , <i>AUD</i> , <i>TFT</i> , <i>WSLS</i> , and <i>IR</i> .	134
6.1	Stationary distribution and fixation probabilities, considering a population of COM_C, C, D, and D_COM.	147
6.2	Contour plot of the frequency of COM_C as a function of $\epsilon$ and $\delta$ .	150
6.3	Frequency of each strategy as a function of benefit-to-cost ratio for the Donation game, in a population of COM_C, COM_D, C, and D individuals.	151
6.4	Frequency of each strategy as a function of mutation rates, in a population consisting of COM_C, COM_D, C, and D individuals.	152
7.1	Frequency of each strategy as a function of confidence threshold, in a population of IRCOM, COM_D, C, and D individuals,	165
7.2	Frequency of IRCOM as a function of the confidence threshold, for different values of management cost, $\epsilon$ , and different values of punishment cost, $\delta$ .	166
7.3	Frequency of IRCOM as a function of confidence threshold, $\theta$ , for different probability distributions of the confidence level, reflecting the efficiency or precision of the intention recognition model at work.	168

Part I

# Intention Recognition



# Chapter 1

---

## INTRODUCTION

---

*All that counts in life is intention.*

Andrea Bocelli

*Our intention creates our reality.*

Wayne Dyer

### 1.1 Motivations

ONCE you recognize it, intention recognition is everywhere. Driving on the street, we recognize intentions of other drivers for our safety. Do they want to turn right, turn left, or just go straight? Talking to friends, for a smooth conversation we usually need to recognize what they mean as messages are not always explicitly conveyed. To secure a successful collaboration with others, we recognize what they want and intend to do. One might risk our life or the life of our beloved ones if one cannot recognize intentions of a hostile enemy. Intention recognition frequently occurs even when we are enjoying our favorite fairy tale stories or watching a comedy episode.

### 1.1.1 Intention Recognition in AI

In AI, because of its important role, research on intention and plan recognition has been carried out over more than three decades. Generally, *intention recognition* (also called *goal recognition*) is defined as the process of becoming aware of the intention of another agent and, more technically, as the problem of inferring an agent’s intention through its actions and their effects on the environment (Charniak and Goldman, 1993; Tahboub, 2006; Heinze, 2003; Armentano and Amandi, 2007; Sadri, 2011b). For the recognition task, several issues can be raised grounded on the eventual distinction between the model an agent creates about himself and the one used to describe others, often addressed in the context of the “Theory of Mind” theory, which neurologically reposes in part on “mirror neurons”, at several cortical levels, as supporting evidence (Iacoboni et al., 2005; Rizzolatti and Craighero, 2004; Nakahara and Miyashita, 2005). *Plan recognition* is closely related to intention recognition, extending it to also recognize the plan the observed agent is following in order to achieve his intention (Armentano and Amandi, 2007; Sadri, 2011b). Intention recognition is performed in domains in which it is better to have a fast detection of just the user goal or intention than a more precise but time consuming detection of the user’s complete plan (Armentano and Amandi, 2007; Sadri, 2011b). Generally, the input to both intention and plan recognition systems is a set of conceivable intentions and a set of plans achieving each intention, given in terms of a plan library (Charniak and Goldman, 1993; Geib and Goldman, 2009) or a plan corpus (Lesh, 1998; Blaylock and Allen, 2003, 2004; Armentano and Amandi, 2009; Han and Pereira, 2011b).

Intention and plan recognition have been applied and shown to be useful in a wide range of application domains (Sadri, 2011b), including story understanding (Charniak and Goldman, 1990), human-computer interaction and interface-agents systems (Lesh, 1998; Hong, 2001; Armentano and Amandi, 2007), traffic monitoring (Pynadath and Wellman, 1995), assistive living (e.g. Elder Care, Ambient Intelligence) (Geib, 2002; Haigh et al., 2004; Tahboub, 2006; Roy et al., 2007; Pereira and Han, 2011a,b; Han and Pereira, 2010c) and military settings (Mao and Gratch, 2004; Heinze, 2003).

Cohen et al. (1981) distinguish between two kinds of intention recognition, *intended* and *keyhole* intention recognition. In the intended case, the observed agent (intending agent) is cooperative; he wants his intentions to be understood by the

observing agent (recognizing agent) and intentionally signals his true intentions. An example would be in case of language understanding where the speaker wants to convey his intentions (Cohen et al., 1981; Pinker et al., 2008). In the keyhole case, the observed agent is unaware of, or indifferent to, being observed and recognized; it is as if the recognizing agent is looking through a keyhole. In this case, partial observability of actions of the intending agent might need to be considered.

Sadri (2011b) further classifies intention recognition as either *adversarial* or *diversionary*. In the former case, the observed agent is hostile to his actions being observed, for instance when the actions are aimed at intrusion in a computer network system (Geib and Goldman, 2001). The latter case points to the circumstances in which the observed agent is in fact attempting to conceal his intentions by performing misleading actions or signaling false information. For example, in the Fox-Crow fable (Aesop), the Fox tried to hide his intention of deceiving the Crow for a piece of cheese from her, pretending that he wants to listen to her singing and praising her—so that she sings and drops the cheese (Pereira and Han, 2009c).

Intention recognition is a rich and challenging field. Often multiple competing hypotheses are possible regarding the intentions and plans of an observed agent. The choice between these hypotheses for the efficiency of the envisaged method so as not to lose its generality is a challenge. For instance, when can we assume the complete information of observed actions, or partial observability needs instead to be considered? When can we assume that the observed agent follows only a single intention at a time, or multiple intentions recognition needs to be performed? Usually, the assumptions are made depending on the nature of the application domains which the envisaged intention recognition method aims at.

Moreover, there are many challenging issues that have attracted much effort from the community. Some most typical and popular ones are the issues of handling incompleteness of observations (e.g. in (Geib and Goldman, 2009; Pereira and Han, 2009c)) and of incompleteness of prior plan library (e.g. in (Sadri, 2010)); the observed agent may follow multiple intentions or interleaved plans simultaneously, or multiple plans toward the same intention (Geib and Goldman, 2009); agents may change or abandon their original intentions (Geib and Goldman, 2003; Han and Pereira, 2011b); incremental (online) intention and plan recognition (see, e.g. (Bui et al., 2002; Han and Pereira, 2011b)). Also, intention recognition becomes more difficult when the uncertainty of the observed actions or even their interpretation

needs to be taken into account, e.g. when dealing with people suffering from cognitive impairments (Roy et al., 2007).

Many methods have been proposed to address these challenging issues, which we will look more closely at in Chapter 2. But usually each of them can handle only a small subset of the issues. That is partly because of their lack in taking into account the aspect of situation-dependent modelling—mostly, only the available observed actions are used for the recognition task. Part of this PhD research is devoted to provide intention recognition methods that take into account contextual information, and show how they can appropriately tackle several unavoidable issues in intention recognition.

### 1.1.2 Evolutionary, Psychological, and Philosophical Aspects of Intention Recognition

*(...) because intention-reading is so critical for human social functioning and the development of key human abilities, such as language and culture, it is reasonable to assume that it has been shaped by natural selection.*

Woodward et al. (2009)

Intention recognition can be found abundantly in many kinds of interactions and communications, not only in Human but also many other species (Tomasello, 1999; Woodward et al., 2009; Cheney and Seyfarth, 2007; Meltzoff, 2007; Tomasello, 2008). Tomasello documented the evidence concerning nonhuman primate understanding of the intentionality of other animate beings, which comes from both experimental and naturalistic studies (Tomasello, 1999, pages 19-22) (Tomasello, 2008). Furthermore, Woodward et al. (2009) provide an extensive review on experimental evidence showing the development of humans' capacity for understanding intentions of others, which trace its roots to the first year of life. Undoubtedly, the knowledge about intentions of others in a situation could enable to plan in advance, either to secure a successful cooperation or to deal with potential hostile behaviors (Geib and Goldman, 2001; van Hees and Roy, 2008; Roy, 2009a; Han and Pereira, 2011c). Given the advantage of knowing the intentions of others and the abundance of intention recognition among different species, undoubtedly intention recognition should be taken into account when studying or modeling collective behavior.

Intention recognition, and the recognition process or heuristics in general, are important mechanisms used by human bounded rationality to cope with real-life complexity (Simon, 1957, 1990; Arthur, 1994; Selten, 2001). In contrast to the unrealistic assumption of perfect rationality, we usually have to make decision in complicated and ill-defined situations, with incomplete information and time constraints (Todd, 2001; Simon, 1990), under computational processing limitations, including limited memory capacity (Simon, 1990; Kareev, 1995; Johnson-Laird, 2010). The recognition processes, based on stored knowledge in terms of generic patterns, have been known to play a major role in that respect. In problems with such complications, we look for patterns, based on them we simplify the problem by using these to construct temporary internal models as working hypotheses (Arthur, 1994; Simon, 1990). It becomes even more relevant when considering interactive settings where the achievement of a goal by an agent does not depend solely on its own actions, but also on the decisions and actions of others, especially when the possibility of communication is limited or includes uncertainty (Kraus, 1997; Heinze, 2003; Pinker et al., 2008; Van Segbroeck et al., 2010). The agents cannot rely on others to behave under perfect or improved rationality, and therefore need to be able to recognize their behaviors and even predict the intention beyond the surface behaviors.

In population-based artificial intelligence applications (Bonabeau et al., 1999; Ampatzis et al., 2008; Gutierrez et al., 2009), such as collective robotics and others, the inherent problem of lack of intention recognition due to the simplicity of the agents is often solved by assuming homogeneous populations, in which each agent has a perfect image of the other as a copy of their own self. Yet, the problem remains in heterogeneous agent systems where it is likely that agents speak different languages, have different designs or different levels of intelligence; hence, intention recognition may be the only way agents understand each other to secure successful cooperation or coordination among heterogeneous agents (Heinze, 2003). Moreover, in more realistic settings where deceiving may offer additional profits, individuals often attempt to hide their real intentions and make others believe in faked ones (Robson, 1990; Tomasello, 1999, 2008; Skyrms, 2010; Pereira and Han, 2009c, 2011b; Santos et al., 2011; Trivers, 2011).

The problem of intention recognition has been paid much attention in AI, Philosophy and Psychology for several decades (Kautz and Allen, 1986a; Charniak and Goldman, 1993; Bratman, 1987, 1999; Geib and Goldman, 2009; Searle, 2010).

Whereas intention recognition has been extensively studied in small scale interactive settings, there is a significant shortage of modelling research with respect to large scale social contexts; namely the evolutionary roles and aspects of intention recognition.

In this thesis, we study the role of intention recognition for one of the most challenging but intriguing issues, traversing areas as diverse as Biology, Economics, Artificial Intelligence, Political Science, or Psychology: the problem of *evolution of cooperation* (Hardin, 1968; Axelrod, 1984; Nowak, 2006b; Sigmund, 2010). In his “The Origin of Species”, Darwin already recognized it as a very fundamental explanatory difficulty (Darwin, 1859). Humans are the champions of cooperation, but cooperative acts are also pervasive in many other species. Insects forgo reproduction to help others. Primates gang up against predators by protecting each other at personal cost. Cooperation also occurs in low levels of organization, from genomes, cells to multicellular organisms (Nowak, 2006a), but we envisage the scope of the roles of intention recognition in the evolution of cooperation for rather high level species, including humans and primates.

The problem of the evolution of cooperation usually entails a social dilemma—situations in which collective interests are at odds with private interests. The most popular and well studied one is the Prisoner’s Dilemma (PD) (Axelrod, 1984; Hofbauer and Sigmund, 1998; Nowak, 2006b; Sigmund, 2010)<sup>1</sup>. In an interaction, each player has two options, cooperate (C) or defect (D), and defect is the dominant option – it is always better to defect in a one-shot interaction. Rationally, both players should choose to defect, while they would be better off by choosing to cooperate instead, thus leading to the destruction of social welfare and individuals’ fitness. In an evolutionary setting, we may ask why would natural selection equip selfish individuals with altruistic tendencies while it incites competition between individuals and thus apparently rewards only selfish behavior? Several mechanisms responsible for promoting cooperative behavior have been recently identified (Sigmund, 2010; Nowak, 2006b). From kin and group ties (Hamilton, 1964; West et al., 2002; Traulsen and Nowak, 2006), to different forms of reciprocity (Nowak and Sigmund, 1992; Imhof et al., 2005; Trivers, 1971; Pacheco et al., 2006; Nowak and Sigmund, 2005) and networked populations (Santos and Pacheco, 2005; Santos et al., 2006a;

---

<sup>1</sup>There are other social dilemmas such as the Stag Hunt and Chicken Game (for an overview see (Sigmund, 2010)), but the Prisoner’s Dilemma is known to represent one of the most difficult or fierce environments for cooperation to emerge.

Szabó and Fáth, 2007; Santos et al., 2008; Lindgren and Nordahl, 1994), several aspects have been shown to play an important role in the emergence of cooperation. In contradistinction, in this thesis we shall describe how cooperation may emerge from the interplay between population dynamics and individuals' cognitive abilities, namely the ability to perform intention recognition.

We provide computational models showing that intention recognition can promote the emergence of improved cooperation, i.e. it leads to higher levels of cooperation in the population and/or a facilitating condition for cooperation to emerge. The study will be carried out within the framework of Evolutionary Game Theory (EGT) (Maynard-Smith, 1982; Hofbauer and Sigmund, 1998; Nowak, 2006a; Sigmund et al., 2010).

### 1.1.3 Commitment and Intention Recognition

There are cases where it is difficult, if not impossible, to recognize the intentions of another agent. It might be your first interaction with someone in your life, and you have no information about him/her which can be used for intention recognition. You also might know someone well, but you still might have very little relevant information in a given situation to predict the intentions with high enough confidence. Furthermore, you might also have abundance of relevant observations about him/her, but he/she is so unpredictable that you have rarely managed to predict his/her true intention in the past. In all such cases, the strategy of proposing commitment, or intention manifestation <sup>2</sup>, can help to impose or clarify the intentions of others.

Agents make commitments towards others when they give up options in order to influence others in a certain way. Most commitments depend on some incentive that is necessary to ensure that an action (or even an intention) is in the agent's interest and thus will be carried out (Gintis, 2001). The capacity for using commitment strategies effectively is so important that natural selection may have shaped specialized capacities to make this possible (Skyrms, 2010; Robson, 1990; Santos et al., 2011; Ruse, 2001; Nesse, 2001b; de Vos et al., 2001; Back and Flache, 2008).

---

<sup>2</sup>*Intention is choice with commitment* (Cohen and Levesque, 1990; Bratman, 1987; Roy, 2009b). Once an agent intends to do something, it must settle on some state of affairs for which to aim, because of its resource limitation and in order to coordinate its future actions. Deciding what to do established a form of commitment (Cohen and Levesque, 1990; Roy, 2009b). Proposing a commitment deal to another agent consists in asking it to express or clarify its intentions.

The commitment we all know is marriage. By giving up the option to leave someone else, spouses gain security and an opportunity for a much deeper relationship than would otherwise be possible (Nesse, 2001b). It might be risky to assume a partner’s intention of staying faithful without the commitment of marriage. A contract is another popular kind of commitment, e.g. for an apartment lease (Frank, 2001). When it is risky to assume another agent’s intention of being cooperative, arranging an appropriate contract provides incentives for cooperation. However, for instance in accommodation rental, a contract is not necessary when the cooperative intention is of high certainty, e.g. when the business affair is between closed friends or relatives. It said, arranging a commitment deal can be useful to encourage cooperation in case intention recognition is difficult, or cannot be performed with sufficiently high confidence. On the other hand, arranging commitments is not for free, and requires a specific capacity to set it up within a reasonable cost (for the agent to actually benefit from it) (Nesse, 2001b,a)—therefor it should be avoided when opportune. In the case of marriage, partners sometimes choose to stay together without an official commitment when it might be too costly (e.g., it could be against parents’ or families’ wish, or it may need to be in secret because of their jobs) and/or they strongly trust each other’s faithfulness (e.g., because of emotional attachment (Frank, 1988)). In short, a combination of the two strategies, those of commitment and intention recognition, seems unavoidable.

In this thesis, again resorting to EGT tools, we show how the most simple of commitment strategies work, and how they can give rise to the emergence of cooperation. Furthermore, we show that combining the strategies of intention recognition and commitment in a reasonable way can lead to the emergence of improved cooperation, not solely achievable by either strategy. Our study seeks what can be a reasonable combination of such strategies.

## 1.2 Aim and Scope

The aim of the research is twofold. First, we aim at developing a *context-dependent* and *incremental* intention recognition method, based on which we will address and clarify different challenging issues in the field, and namely those mentioned above (Section 1.1.1). Second, resorting to this intention recognition method we shall study the roles of intention recognition in the evolution of cooperation, providing

computational models showing that intention recognition promotes the emergence of cooperation. Furthermore, we study how commitment, or intention manifestation, in combination with intention recognition, can lead to the emergence of improved cooperation.

We will use Bayesian networks as a basis structure of our intention recognition model, and sometimes couple them with Logic Programming techniques for modelling the changing world and reasoning about it (Baral, 2003). The flexibility of Bayesian networks for representing probabilistic dependencies and the efficiency of inference methods for Bayesian networks have made them an extremely powerful tool for problem solving under uncertainty (Pearl, 1988, 2000).

Furthermore, for decades, Logic Programming has proven to be a very powerful tool for knowledge representation (monotonic and non-monotonic logics, and different kinds of negation, etc.) (Baral and Gelfond, 1994; Lloyd, 1987; Baral, 2003; Kowalski, 2011) and cover a wide range of reasoning problems, such as deduction (Hayes, 1973), abduction, integrity constraint handling (Kakas et al., 1993; Alferes et al., 2004), preferences modelling (Brewka and Eiter, 1999; Pereira et al., 2011a; Pereira and Han, 2009b), updates and revision (Alferes et al., 2000, 2005). Logic Programming representation and reasoning techniques are appropriate for our aim of modelling the changing world and reasoning about it—in order to extract relevant information for intention recognition as well as fine tuning the model.

In order to study the roles of intention recognition and commitment in the evolution of cooperation, we will resort to Evolutionary Game Theory as analytical and simulation tools, and to the Prisoner’s Dilemma game as the interacting model of agents.

## 1.3 Thesis Contributions

There are two primary contributions of this PhD thesis, corresponding to the two deployed composite aims described above.

- We provide a novel and efficient method for incremental and context-dependent intention recognition. Our method can appropriately address several important issues in intention recognition, including intention change/abandonment and multiple intentions recognition (Han and Pereira, 2011a, 2010a, 2011b). The method is applicable in particular in the context of evolutionary game

theory, which is important for the second aim of the thesis.

- We provide computational models showing the roles of intention recognition in the evolution of cooperation (Han et al., 2011a,b). The result of our research provides new evidence for the evolutionary advantages of having high cognitive skills in the problems where cooperation is beneficial, and putting forward a new approach to using AI cognitive modelling methods to tackle the problem of explaining the evolution of cooperation. We also provide computational models showing that commitment can lead to emergence of reinforced cooperation (Han et al., 2012b), and even more when combining with the intention recognition strategy (Han et al., 2012c).

Additionally, the thesis contains two secondary contributions.

- We provide an intention-based decision making framework, where we demonstrate several ways intention recognition can be useful for decision making (Han and Pereira, 2011c).
- We describe how intention recognition can be useful for different application domains, including Elder Care and Ambient Intelligence in the home environment (Han and Pereira, 2010c; Pereira and Han, 2011a,b).

## 1.4 Outline of the Thesis

The thesis is split into two parts. In Part I, we study techniques for intention recognition from the AI modelling perspective, which spans two chapters (Chapters 2 and 3).

Chapter 2 presents a corpus-based method for intention recognition, which is context-dependent and incremental. The method is evaluated with several benchmarks. We show that taking into account contextual information appropriately influences the intention recognition process, and can significantly improve the performance (Han and Pereira, 2010a, 2011b,a).

Chapter 3 describes a framework for intention-based decision making on top of our previously implemented Evolution Prospection system (Pereira and Han, 2009a,b; Han, 2009). It serves to bridge intention recognition to real application domains, where it is important and useful to take into account intentions of other

agents in the environment. Its illustration on the Elder Care and Ambient Intelligence domains will be described (Han and Pereira, 2011c, 2010c; Pereira and Han, 2011b,a; Han and Pereira, 2011a).

In Part II, we study the roles of intention recognition and commitment in the evolution of cooperation, including four chapters (Chapters 4-7).

In Chapter 4, based on a general Bayesian network for intention recognition provided in Part I, we derive a Bayesian network model for intention recognition in the context of social dilemmas. It takes into account mutual trusts between the intention recognizer and the co-player, accumulated throughout past interactions. We show, analytically and by computer simulations, that the intention recognition strategy outperforms the most famous successful strategies in the iterated Prisoner's dilemma (Han et al., 2011a).

Chapter 5 offers a corpus-based intention recognition model in the course of iterated Prisoner's Dilemma, which is based on the general corpus-based incremental intention method described in Part I. It results in an intention recognition strategy that can recognize the strategy of its co-player, and make appropriate decisions based on the recognized strategy. We show that the strategy prevails against the most famous successful strategies of the iterated Prisoner's Dilemma (Han et al., 2011b).

Chapter 6 provides a computational model showing that commitment of the simplest form can lead to the emergence of cooperation, even in non-iterated settings of the one-shot Prisoner's Dilemma (Han et al., 2012b).

Chapter 7 describes a model showing that an appropriate combination of the two previous strategies, those of intention recognition and commitment, can lead to a strategy performing better than either one solely (Han et al., 2012c).

The thesis ends with conclusions and discussion of future work, in Chapter 8.



# Chapter 2

---

## INTENTION RECOGNITION METHODS

---

*We cannot perceive the thought of another person at all, we can only infer it from his behavior.*

Charles Darwin (1809-1882): *Old and useless notes, 1840*

*We call an intention good which is right in itself,  
but the action is good, not because it contains within it some good, but because it issues from a good intention.*

Peter Abelard (1079-1142)

**Abstract.** In this chapter, we present a novel method for context-dependent and incremental intention recognition. It is performed by incrementally constructing a three-layer Bayesian network model as more actions are observed, and in a context-dependent manner. The Bayesian

---

<sup>0</sup>This chapter encompasses the publications (Han and Pereira, 2010a), (Han and Pereira, 2011b), (part of) (Pereira and Han, 2011b), and a journal submission (Han and Pereira, 2011a).

network is composed from a knowledge base of readily specified and readily maintained Bayesian network fragments. The simple structure of the fragments enables us to easily and efficiently acquire the corresponding knowledge base, either from domain experts or else automatically from a plan corpus. We illustrate the method in the Elder Care domain, showing how contextual information can appropriately influence the recognition results of an observed elder’s intention. Furthermore, we experimentally show the improved performance of the method in the Linux plan corpus and the so-called IPD plan corpora (that is, in the framework of the Iterated Prisoner’s Dilemma). We show that taking into account contextual information can considerably improve the intention recognition performance.

## 2.1 Introduction

We propose a method for intention recognition in a dynamic, real-world environment. An important aspect of intentions is *future-directedness*, i.e., if we intend something now, we mean to execute a course of actions to achieve something in the future (Bratman, 1987; Roy, 2009b). Most actions may be executed only at a far distance in time. During that period, the world is changing, and the initial intention may be changed to a more appropriate one or even abandoned (Singh, 1991; Bratman, 1992; Geib and Goldman, 2009). An intention recognition method should take into account these changes, and, when necessary, be able to reevaluate the intention recognition model depending on some time limit. In addition, as new actions are observed, the model should be reconfigurable to incorporate them. In other words, the model is incremental and, furthermore, the intention recognition prediction is available at anytime (Han and Pereira, 2010a).

Generally, *intention recognition* (also called *goal recognition*) is defined as the process of becoming aware of the intention of another agent and, more technically, as the problem of inferring an agent’s intention through its actions and their effects on the environment (Charniak and Goldman, 1993; Tahboub, 2006; Heinze, 2003; Armentano and Amandi, 2007). *Plan recognition* is closely related to intention recognition, extending it to also recognize the plan the observed agent is following in order to achieve his intention (Sadri, 2011b; Armentano and Amandi, 2007).

Mere intention recognition is performed in domains in which it is preferred to have a fast detection of just the user’s goal/intention rather than a more precise but time consuming detection of the user’s complete plan, e.g., in the interface agents domain (Armentano and Amandi, 2007; Horvitz et al., 1998; Madani et al., 2009). Like most other intention or plan recognition work, we assume that the actions to be used for the recognition task are given, and are observed with certainty. Usually, the actions are given by an *action* or *activity recognition* system (see, e.g. (Turaga et al., 2008; Duong et al., 2009)). Dealing with uncertainty in action recognition is beyond the scope of our work here <sup>1</sup>. Generally, the inputs to both intention and plan recognition systems are a set of conceivable intentions and a set of plans for achieving each intention, given in terms of a plan library (Charniak and Goldman, 1993; Geib and Goldman, 2009) or a plan corpus (Blaylock and Allen, 2003, 2004; Armentano and Amandi, 2009)). There are also generative approaches based on planning algorithms, which do not require plan library/corpus (e.g., see (Ramírez and Geffner, 2010)).

The future-directedness of intentions also means that once an agent intends something, he has settled on a particular course of action (Bratman, 1987; Cohen and Levesque, 1990; Singh, 1991; Roy, 2009b). This makes the intentions relatively stable, pending new information. An agent who made the decision to act in a certain way commits to sticking to this decision by the reasons which led to it, unless counterbalancing reasons meanwhile appear and trigger further deliberations (for more discussion see (Bratman, 1987, p. 16-17) and (Singh, 1991)). In other words, intentions are relatively resistant to reconsideration unless there are pondered reasons to do so (Bratman, 1992, 1987; Singh, 1991; Roy, 2009b). Following this, any attempt to tackle the issues of intention change or abandonment cannot solely be based on observable actions. The reasons why the intention is changed or abandoned must be taken into account. The reasons can be changes in the environment (possibly made by other agents) which do not support or force the observed agent to keep following his initial intention anymore. And here the context-dependent modeling appears to be unavoidable (Section 2.7).

In this work, we resort to Bayesian Networks (BNs) as the intention recognition model. The flexibility of BNs for representing probabilistic dependencies and the

---

<sup>1</sup>Though in our paper “Moral Reasoning with Uncertainty” (Han et al., 2012d), we address the case of judging under uncertainty of actions performed, and the giving of a verdict by a jury, illustrated with an example.

efficiency of inference methods for BNs have made them an extremely powerful and natural tool for problem solving under uncertainty (Pearl, 1988, 2000). To perform intention recognition, we construct a three-layer BN (Pereira and Han, 2009c, 2011b)—justified based on Heinze’s causal intentional model (Heinze, 2003; Tahboub, 2006)—and use it for evidential reasoning from observations to intention hypothesis.

We surmise a knowledge representation method to support incremental BN model construction for performing intention recognition during runtime, from an initially given domain knowledge base. As more actions are observed, a new BN is constructed from the previous one reinforcing some intentions whilst ruling out others. This incremental method allows domain experts to specify knowledge in terms of small and simple BN fragments, which can be easily maintained and changed, and which are used to compose the situated ongoing BN model. Alternatively, these fragments can be easily learned from data. We also propose a method to represent relationships among intentions, when considering the case of agents that may pursue multiple intentions simultaneously (Section 2.5). It is an indispensable aspect, but mostly omitted in prior work, which moreover allows us to sometimes significantly decrease the complexity of the BN inference (Gogate and Dechter, 2011).

Our method is generally inspired in that knowledge experts often consider a related set of variables together, and organize domain knowledge in larger chunks. An ability to represent conceptually meaningful groupings of variables and their inter-relationships facilitates both knowledge elicitation and knowledge base maintenance (Laskey and Mahoney, 1997). To this end, there have been several methods proposed for Bayesian Network construction from small and easily maintained network fragments (Pearl, 1988; Pfeffer et al., 1999; Mahoney and Laskey, 1996; Laskey and Mahoney, 1997; Xiang and Poh, 2005; Natarajan et al., 2008; Laskey, 2008). In essence, a combination of BNs is a graph that includes all nodes and links of the networks, where nodes with the same name are combined into a common node. The main issue for a combination method is how the influence of different parents of the common node can be combined in the new network, given the partial influence of each parent in the corresponding fragment. The most extensively used and popular combination method is Noisy-Or, firstly proposed by (Pearl, 1988) for BNs of Boolean variables, and generalized by (Srinivas, 1993; Diez, 1993) for the general case of arbitrary domains. The Noisy-OR method is discussed in Section 2.4.

The rest of this chapter has the following structure. The next section describes a short review of prior work, pointing out those limitations of that work which we address in our method (Section 2.2). Section 2.3 recalls some background of BNs that is necessary for further discussion of the intention recognition method, which is described in detail in Section 2.4. Section 2.5 describes a method for expressing relationships amongst intention variables in a BN model for intention recognition. Section 2.6 discusses terminologies used for the evaluation of our method, including the evaluation metrics, and the first set of experimental results on the Linux plan corpus. Section 2.7 presents our own, so-called IPD plan corpora benchmarks based on the iterated Prisoner’s Dilemma, and show our experimental results for it. We also describe how to incorporate contextual information in our model, and how it helps to improve the intention recognition performance. Section 2.8 proffers some extensions of our method to take into account contextual information. Further developments and concluding remarks, in Section 2.9, end the chapter.

## 2.2 Related Work

Work on intention and plan recognition has been paid much of attention for more than thirty years, and a large number of methods have been applied. They can be roughly categorized into two main groups: *Consistency* and *Probabilistic* approaches (Armentano and Amandi, 2007; Singla and Mooney, 2011; Geib and Goldman, 2009; Sadri, 2011b).

Consistency approaches face the problem by determining which intention is consistent with the observed actions, i.e. whether the observed actions match with at least a plan achieving the intention. The earliest work on plan recognition belongs to this group (Schmidt et al., 1978; Wilensky, 1983; Kautz and Allen, 1986a; Hong, 2001). More recent work can be found in a rather comprehensive survey by Sadri (2011b). The problem with the consistency approaches is that they cannot handle well the case where the current observed actions enable more than one intention—they cannot directly select between those intentions.

Probabilistic approaches, on the other hand, are mainly based on Bayesian network and (Hidden) Markov models (Charniak and Goldman, 1993; Pynadath and Wellman, 1995; Forbes et al., 1995; Albrecht et al., 1998; Forbes et al., 1995; Conati et al., 1997; Albrecht et al., 1998; Bui et al., 2002; Huber and Simpson, 2004; Tah-

boub, 2006; Schrempf et al., 2007; Geib and Goldman, 2009; Pereira and Han, 2009c, 2011b; Armentano and Amandi, 2009). A significant advantage of the probabilistic approaches is that they can directly address the above issue of the consistency approaches—by finding the most probable intentions given the current observations, on the basis of accumulated statistical evidence or simply subjective beliefs encoded in a Bayesian network or Markov model.

Bayesian approaches have exhibited the most successful models applied to intention/plan recognition (Charniak and Goldman, 1993; Pynadath and Wellman, 1995; Goldman et al., 1999; Geib, 2004; Geib and Goldman, 2009). The first model was built by Charniak and Goldman (1991, 1993). Depending on the structure of plan libraries, a knowledge-based model construction is employed to build BNs from the library—which is then used to infer the posterior probability of explanations (for the set of observed actions). This approach, mostly advanced by Goldman et al. (1999) and especially in the more recent work (Geib and Goldman, 2009)<sup>2</sup>, addresses a number of issues in intention/plan recognition, e.g., when the observed agent follows multiple intentions or interleaved plans simultaneously; fails to observe actions; addresses partially ordered plans. However, there are some important aspects not yet explored therein, partially for the sake of computational efficiency. First, prior probabilities of intentions are assumed to be fixed. This assumption is not always reasonable because those prior probabilities should in general depend on the situation at hand (Bratman, 1992, 1987; Pynadath and Wellman, 1995; Brown, 1998), and can justifiably be captured by causes/reasons of the intentions, as in our method (Pereira and Han, 2011b; Han et al., 2011a; Tahboub, 2006; Heinze, 2003). Indeed, Geib and Goldman (2009) also highlighted the need to account for contextual information or state of the world as a potential extension to their plan recognizer. In (Pynadath and Wellman, 1995), a similar context-dependent Bayesian approach is used, though the model therein is not incremental. The authors demonstrated that taking into account contextual information is important to appropriately recognize drivers’ intention in the traffic monitoring domain (Pynadath and Wellman, 1995).

Second, intentions are assumed to be independent of each other. This is not generally the case since the intentions may support or exclude one another, leading to the need to reconfigure the model. Those works hence might not appropriately address multiple intentions recognition. Pynadath and Wellman (1995) proposed

---

<sup>2</sup>Note that this work is based on Bayesian inference, though they do not build Bayesian networks as in (Charniak and Goldman, 1991, 1993).

to combine, in their BN model for plan recognition, the mutually exclusive plan nodes into a single variable. As a step further, we formally define how that can be done appropriately, so as to guarantee the consistency in the obtained BN (Section 2.5). This latter assumption must always, explicitly or implicitly, be made by the approaches based on (Hidden) Markov models, e.g. (Armentano and Amandi, 2009; Bui, 2003), or statistical corpus-based machine learning (Blaylock and Allen, 2003, 2004). Generally, in those approaches, a separate model is built for each intention; thus no relations amongst the intentions are expressed or can be expressed. These works were restricted to the single intention case. The method in this chapter attempts to tackle the multiple case more appropriately.

Different from most above mentioned works, our model is *context-dependent*, which is partly achieved by including in it causes/reasons of intentions. This way, our model can appropriately deal with the abandonment/changes of intentions—when the causes/reasons do not support or force the intending agent to hold those intentions anymore—in an *integrated* manner. In contrast, in (Geib and Goldman, 2003), the authors build a *separate* model to recognize when the observed agent abandons its current intention, which may then trigger revision of the intention recognition model. To the best of our knowledge, this is the only work addressing the abandonment issue. However, the system therein is only evaluated with a rather small benchmark (with three intentions), and only for the accuracy of the abandonment recognition itself. The benefit from having this additional intention abandonment recognition module for enhancing intention/plan recognition performance has not been studied, as the authors themselves mention in their recent study (Geib and Goldman, 2009). We address this issue in Section 2.7.

## 2.3 Bayesian Networks

**Definition 2.3.1** *A Bayesian Network (BN) is a pair consisting of a directed acyclic graph (DAG) whose nodes represent variables and missing edges encode conditional independencies between the variables, and an associated probability distribution satisfying the Markov assumption of conditional independence, saying that variables are independent of non-descendants given their parents in the graph (Pearl, 1988, 2000).*

In a BN, associated with each node of its DAG is a specification of the dis-

tribution of its variable, say  $A$ , conditioned on its parents in the graph (denoted by  $pa(A)$ )—i.e.,  $P(A|pa(A))$  is specified. If  $pa(A) = \emptyset$  ( $A$  is called *root* node), its unconditional probability distribution,  $P(A)$ , is specified. These distributions are called Conditional Probability Distribution (CPD) of the BN.

The joint distribution of all node values can be determined as the product of conditional probabilities of the value of each node on its parents

$$P(X_1, \dots, X_N) = \prod_{i=1}^N P(X_i|pa(X_i)) \quad (2.1)$$

where  $V = \{X_i | 1 \leq i \leq N\}$  is the set of nodes of the DAG.

Suppose there is a set of evidence nodes (i.e. their values are observed) in the DAG, say  $O = \{O_1, \dots, O_m\} \subset V$ . We can determine the conditional probability distribution of a variable  $X$  given the observed value of evidence nodes by using the conditional probability formula

$$P(X|O) = \frac{P(X, O)}{P(O)} = \frac{P(X, O_1, \dots, O_m)}{P(O_1, \dots, O_m)} \quad (2.2)$$

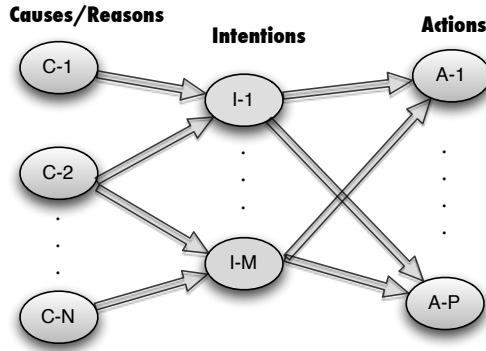
where the numerator and denominator are computed by summing the joint probabilities over all absent variables with respect to  $V$  as follows

$$\begin{aligned} P(X = x, O = o) &= \sum_{av \in ASG(AV_1)} P(X = x, O = o, AV_1 = av) \\ P(O = o) &= \sum_{av \in ASG(AV_2)} P(O = o, AV_2 = av) \end{aligned} \quad (2.3)$$

where  $o = \{o_1, \dots, o_m\}$  with  $o_1, \dots, o_m$  being the observed values of  $O_1, \dots, O_m$ , respectively;  $ASG(Vt)$  denotes the set of all assignments of vector  $Vt$  (with components are variables in  $V$ );  $AV_1, AV_2$  are vectors components of which are corresponding absent variables, i.e. variables in  $V \setminus (O \cup \{X\})$  and  $V \setminus O$ , respectively.

## 2.4 Incremental Intention Recognition

In (Pereira and Han, 2009c, 2011b), a general BN model for intention recognition is presented and justified based on Heinze’s causal intentional model (Heinze, 2003; Tahboub, 2006). Basically, the BN consists of three layers: cause/reason nodes in



**Figure 2.1: General structure of a Bayesian network for intention recognition.** The Bayesian network consists of three layers. The pre-intentional layer consists of cause/reason nodes, connecting to intention nodes in the intentional layer, which in turn connect to action nodes in the activity layer.

the first layer (called *pre-intentional*), connecting to intention nodes in the second one (called *intentional*), in turn connecting to action nodes in the third (called *activity*) (Figure 2.1).

In general, it is possible to build a single BN containing all the relevant factors. But in a real application domain, it can be envisaged that such a network is very large, which clearly leads to high complexity for the BN inference (Pearl, 2000, 1988). To address this problem, in this work we present a method for incrementally constructing a BN model with the purpose of performing *incremental* intention recognition. The idea is that, given the current observations, only their relevant factors are incorporated for the moment into the network.

**Definition 2.4.1 (Intention Recognition BN – IRBN)** A BN for intention recognition (IRBN)  $W$  is a triple  $\langle \{Cs, Is, As\}, pa, P_W \rangle$  where

- $Cs$ ,  $Is$  and  $As$  are the sets of cause/reason nodes, intention nodes and action nodes, respectively. They stand for binary random variables (i.e. their value is either true (T) or false (F)).
- $pa$  is a mapping which maps a node to the set of its parent nodes such that:  $pa(C) = \emptyset \quad \forall C \in Cs$ ;  $pa(I) \subseteq Cs \quad \forall I \in Is$ ; and  $\emptyset \neq pa(A) \subseteq Is \quad \forall A \in As$ .
- CPD tables are given by the probability distribution  $P_W$ , i.e.  $P_W(X|pa(X))$  defines the probability of  $X$  conditional on  $pa(X)$  in  $W$ ,  $\forall X \in Cs \cup Is \cup As$ .

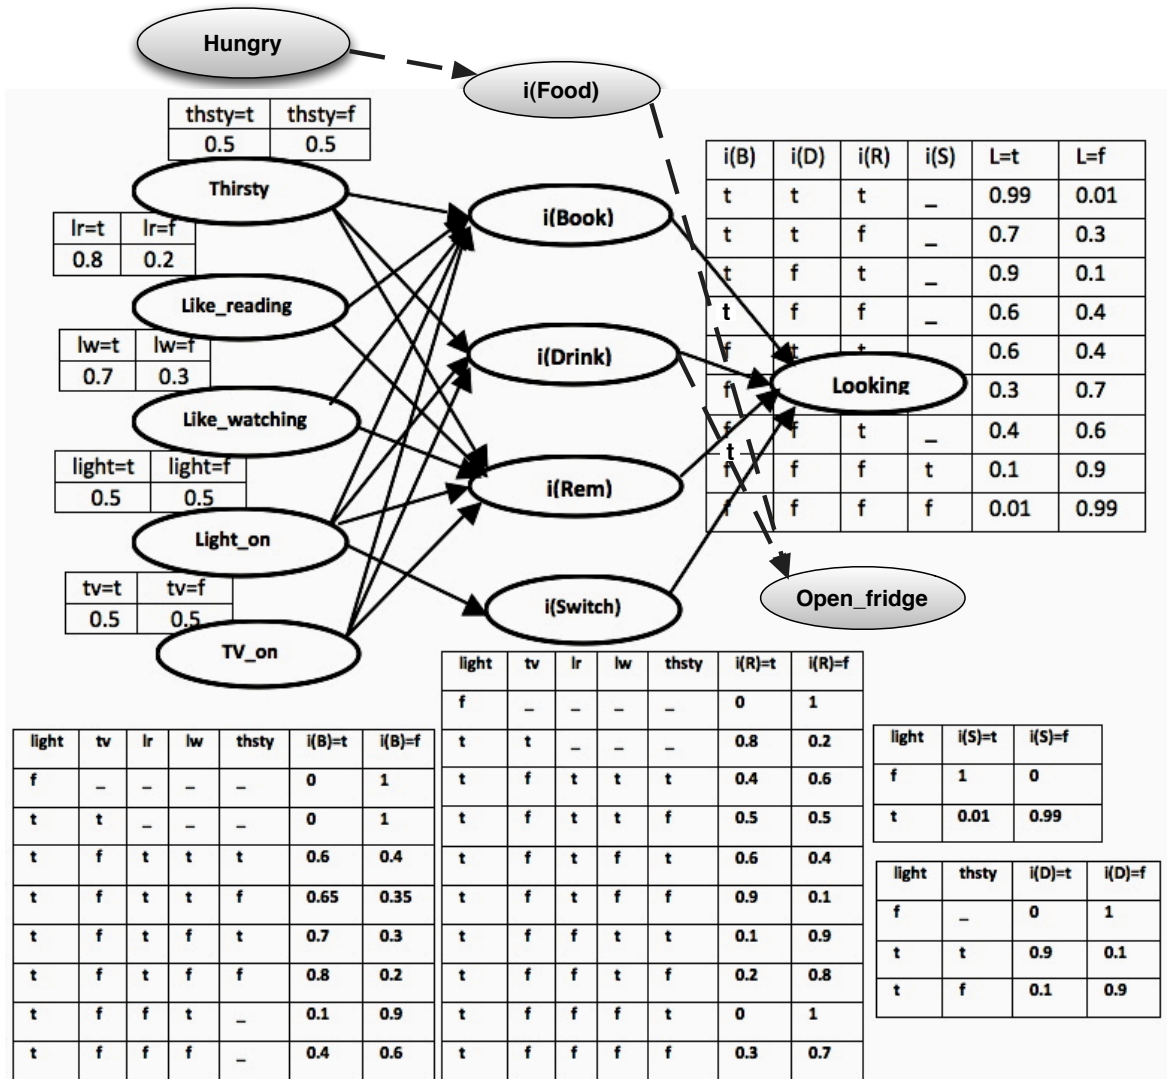


Figure 2.2: Elder Intentions Recognition IRBN

Note that the set of cause/reason nodes of an IRBN,  $Cs$ , can be empty, as in the case of the Linux plan corpus we shall see later on (Section 2.6).

For illustration of the concepts and the intention recognition method presented in this section, we consider an extended example from the Elder Care domain (Pereira and Han, 2011a,b).

**Example 2.4.1 (Elder’s Care with Intention Recognition)** *An elder stays alone in his apartment. The assistant system (with the capability of intention recognition) observes that he is looking for something in the living room. In order to assist him, the system needs to figure out what he intends to find. The possible things are: something to read (Book –  $i(B)$ ); something to drink (Drink –  $i(D)$ ); the TV remote control (Rem –  $i(R)$ ); and the light switch (Switch –  $i(S)$ ). The IRBN representing this scenario is that of Figure 2.2, but without the grey-filled nodes for now.*

*There are five nodes in the first layer: Thirsty ( $thsty$ ), Like\_reading ( $lr$ ), Like\_watching ( $lw$ ), Light\_on ( $light$ ) and TV\_on ( $tv$ ). We have,  $Cs = \{thsty, lr, lw, light, tv\}$ . Intention nodes in the middle layer are,  $Is = \{i(B), i(D), i(R), i(S)\}$ . Action nodes are,  $As = \{Looking\}$ .*


*The mapping  $pa$  is defined by the arrows in the IRBN, e.g.,  $pa(i(B)) = \{thsty, lr, lw, light, tv\}$  and  $pa(Looking) = \{i(B), i(D), i(R), i(S)\}$ .*

The intention recognition method will be performed by incrementally constructing an IRBN as more actions are observed. The construction is based on a prior knowledge base consisting of unit fragments of BN (Figure 2.3).

**Definition 2.4.2 (Unit Fragments)** *There are just two types of unit fragments issued for IRBN model construction (Figure 2.3):*

1. *A unit fragment for an action  $A$  consists of an intention  $I$  connecting to (i.e. causally affecting)  $A$ , and is denoted by  $UF_{\mathfrak{A}}(I, A)$ .*
2. *A unit fragment for an intention  $I$  consists of a context-independent and fixed over time set of causes/reasons  $Cs$  connecting to (i.e. causally affecting)  $I$ , and is denoted by  $UF_{\mathfrak{I}}(Cs, I)$ .*

Note that a unit fragment for an action contains a single action, which thus enables to easily handle the set of conceivable intentions in the IRBN, in a context-dependent manner, as implemented with operators described in the sequel. In contradistinction,



**Figure 2.3:** Two Types of Unit Fragments.

a unit fragment for an intention includes a fixed set of causes/reasons rather than a single one. This simplified modification can be made because the prior probability of each cause/reason node can be updated in a context-dependent manner (as we shall see in Section 2.8), allowing us to easily switch on and off the effect of a cause/reason node. Differently, prior probabilities of the intention nodes depend in general on the cause/reason nodes.

**Example 2.4.2 (Unit Fragments)** *Here are some unit fragments for actions and intentions that will be used later for constructing the IRBN in Figure 2.2.*

*Some unit fragments for the action Looking:  $UF_{\Delta}(i(B), L)$ ,  $UF_{\Delta}(i(D), L)$ ,  $UF_{\Delta}(i(R), L)$  and  $UF_{\Delta}(i(S), L)$ .*

*Some unit fragments for intentions:  $UF_{\Upsilon}(\{thsty, lr, lw, light, tv\}, i(B))$ ,  $UF_{\Upsilon}(\{thsty, light, tv\}, i(D))$ ,  $UF_{\Upsilon}(\{thsty, lr, lw, light, tv\}, i(R))$  and  $UF_{\Upsilon}(\{light\}, i(S))$ .*

We next stipulate some conditions to guarantee the consistency of the knowledge base.

**Definition 2.4.3 (Knowledge Base)** *The domain knowledge base KB consists of a set of actions  $\Delta$ , a set of intentions  $\Upsilon$ , a set of unit fragments for each action in  $\Delta$  and a single unit fragment for each intention in  $\Upsilon$ , satisfying that*

- *An intention  $I$  has a unique unit fragment in KB. The set of its parents (causes/reasons) and the CPD table associated with it are fixed. Let  $\mathfrak{C}(I)$  denote the set of the parents of  $I$  and  $P_{KB}(I|\mathfrak{C}(I))$  define its CPD table.*
- *A cause  $C$  has the same prior probability distribution in all the unit fragments (for intentions) that it belongs to, denoted by  $P_{KB}(C)$ .*

The simple structures of unit fragments enable domain experts to easily construct and maintain the knowledge base. The fragments can also be learnt from appropriate datasets, as we shall see later for the Linux and IPD plan corpora.

Before presenting the intention recognition algorithm, let us define some (original) operators for handling CPD tables and IRBNs.

### 2.4.1 Operators for constructing IRBNs

As a new action  $A$  is observed, we need to incorporate it into the current IRBN. Firstly, the appropriate unit fragments for  $A$  are selected from the prior domain knowledge base. In Section 2.8.2, we will discuss methods for selecting the appropriate fragments in a situation-sensitive manner. They are based on the intuition that whether an intention may give rise to an action depends on the situation in which the action is observed. That enables to reduce the size of the BN model, which otherwise could be very large.


For now, let us assume that the operator  $select(A, SIT)$  provides the (context-dependent) set of unit fragments for action  $A$  given the situation at hand,  $SIT$ . If  $SIT$  is empty,  $select(A, SIT)$  is the set of all unit fragments for action  $A$  from the knowledge base.

Then, after having obtained the appropriate fragments, we combine them using the Noisy-OR method (Pearl, 1988; Srinivas, 1993; Cozman, 2004), thereby obtaining a BN with a single action (Figure 2.4). We then add into it appropriate causes/reasons for each intention.

**Definition 2.4.4 (Unit IRBN via Noisy-OR)** *The unit IRBN  $W$  for action  $A$  in a given situation  $SIT$  is an IRBN with a single action, denoted by  $irBN(A) = \langle \{Cs, Is, \{A\}\}, pa, P_W \rangle$ . It is obtained via the Noisy-OR method as follows.*

*Let  $select(A, SIT) = \{UF_{\mathfrak{A}}(I_1, A), \dots, UF_{\mathfrak{A}}(I_N, A)\}$  and for  $1 \leq i \leq N$ ,  $P(A = T | I_i = T) = q_i$  (defined in fragment  $UF_{\mathfrak{A}}(I_i, A)$ ). Then,*

- $Is = \{I_1, \dots, I_N\}$ ;  $Cs = \bigcup_{I \in Is} \mathfrak{C}(I)$ ;
- $pa(I) = \mathfrak{C}(I) \quad \forall I \in Is$ ;  $pa(A) = Is$ ;
- $P_W(C) = P_{KB}(C) \quad \forall C \in Cs$ ;  $P_W(I|pa(I)) = P_{KB}(I|\mathfrak{C}(I)) \quad \forall I \in Is$ ; and,



**Figure 2.4:** Noisy-OR Combination Method (Pearl, 1988; Srinivas, 1993; Diez, 1993):  $P(A = T|I_1, \dots, I_N) = 1 - \prod_{i:I_i=T} (1 - q_i)$  where  $P(A = T|I_i = T) = q_i$ .

according to the Noisy-OR method

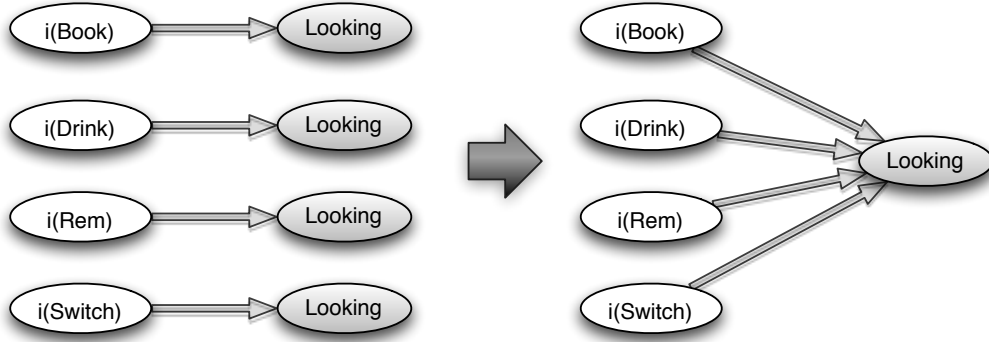
$$P_W(A = T|pa(A)) = 1 - \prod_{1 \leq i \leq N: I_i=T} (1 - q_i). \quad (2.4)$$

The rationale and appropriateness of the application of the Noisy-OR method here for combining unit fragments for an action is based on the intuition that each intention can be interpreted as “cause” of the observed action (Bratman, 1987; Pynadath and Wellman, 1995); and the action occurs when one or more of such intentions are active. More detailed arguments for this can be found in (Cozman, 2004; Pearl, 1988).

**Example 2.4.3 (Noisy-OR)** Consider the action node *Looking* (i.e.,  $L$ ), with four parent intention nodes  $i(B)$ ,  $i(D)$ ,  $i(R)$ , and  $i(S)$ . In the unit fragments for  $L$ ,  $UF_{\mathfrak{A}}(i(B), L)$ ,  $UF_{\mathfrak{A}}(i(D), L)$ ,  $UF_{\mathfrak{A}}(i(R), L)$ , and  $UF_{\mathfrak{A}}(i(S), L)$ , we define  $P(L = T|i(B) = T) = 0.9$ ,  $P(L = T|i(D) = T) = 0.7$ ,  $P(L = T|i(R) = T) = 0.8$ ,  $P(L = T|i(S) = T) = 0.2$ , respectively.

The combination of these unit fragments using the Noisy-OR method, in Figure 2.5, and the CPD table for the node  $L$  in the obtained BN is defined following Eq. (2.4). The CPD for the node *Looking* ( $L$ ) in Figure 2.2, now instead of being fully constructed beforehand, can be defined by this Noisy-OR combination from the simple unit fragments.

Obviously, from the design point of view, it is easier and usually much cheaper to construct the small fragments (and then combine them) than to construct the



**Figure 2.5:** Noisy-OR Combination for the node *Looking*.

whole BN (Pearl, 1988; Laskey, 2008) (see also the discussion in this chapter's introduction, Section 2.1).

Now we need to combine the obtained unit IRBN,  $\text{irBN}(\mathbf{A})$ , with the current IRBN. For that, in the sequel we define how to combine two IRBNs. Intuitively, we simply add up all the new nodes and links of the new IRBN to the current IRBN, keeping the CPD tables from the two original IRBNs.

**Definition 2.4.5 (Combination of IRBNs)** Let  $W_1 = \langle \{Cs_1, Is_1, As_1\}, pa_1, P_1 \rangle$  and  $W_2 = \langle \{Cs_2, Is_2, As_2\}, pa_2, P_2 \rangle$  be two IRBNs, such that  $As_1 \cap As_2 = \emptyset$  (any actions in  $As_2$  which are already so named in  $As_1$  are renamed). The combination of these two IRBNs is an IRBN, denoted by  $\text{comb}(W_1, W_2) = \langle \{Cs, Is, As\}, pa, P_W \rangle$ , where

- $As = As_1 \cup As_2$ ;  $Is = Is_1 \cup Is_2$ ;  $Cs = Cs_1 \cup Cs_2$ ;
- $pa(I) = \mathfrak{C}(I) \quad \forall I \in Is$ ;  $pa(A) = pa_1(A) \cup pa_2(A) \quad \forall A \in As$ ;
- $P_W(C) = P_{KB}(C) \quad \forall C \in Cs$ ;  $P_W(I|pa(I)) = P_{KB}(I|\mathfrak{C}(I)) \quad \forall I \in Is$ ;  
 $P_W(A|pa(A)) = P_k(A|pa_k(A)) \quad \text{if } A \in As_k \quad (\text{with } k = 1, 2).$

Note that here it is allowed the possibility that the observed agent follows multiple intentions simultaneously.

When some intentions are found irrelevant, e.g., because they are much unlikely<sup>3</sup>, those intentions should be removed from the IRBN. These intentions might resurface

<sup>3</sup>One intention is much less likely than the other if the fraction of its likelihood and that of the most likely intention is less than some small threshold. It is up to the KB designer to provide it.

in the IRBN later if there appear newly observed actions triggering them. This could happen frequently when one considers the case where agents might abandon/change their initial intentions—especially for the multiple intention recognition case.

The removal is enacted by considering the intentions (to be removed) as completely false and employing a *project* operator, described below.

**Definition 2.4.6 (Project of CPD Table)** *Let  $Tb$  be a CPD table defining  $P(X|V)$ , the probability of a random variable  $X$  conditional on a set of random binary variables  $V$ . Considering a strict subset  $V' \subsetneq V$ , the project of  $Tb$  on  $V'$ , denoted by  $\text{proj}(Tb, V')$ , is the part of  $Tb$  corresponding to all variables in  $V \setminus V'$  being false.*

Note that this approach guarantees the consistency of the probability distribution over the set of the remaining intentions. It requires neither normalization over the set nor creating a new CPD: any computation can be done directly with the original CPD.

**Definition 2.4.7 (Remove Intentions from IRBN)** *Let  $W = \langle \{Cs, Is, As\}, pa, P_W \rangle$  be an IRBN and  $R \subset Is$  be a strict subset of  $Is$ . The result of removing the set of intentions  $R$  from  $W$  is an IRBN, denoted by  $\text{remove}(W, R) = \langle \{Cs_R, Is_R, As_R\}, pa_R, P_R \rangle$ , where*

- $As_R = As$ ;  $Is_R = Is \setminus R$ ;  $Cs_R = \bigcup_{I \in Is_R} \mathfrak{C}(I)$ ;
- $pa_R(I) = \mathfrak{C}(I) \quad \forall I \in Is_R$ ;  $pa_R(A) = pa(A) \setminus R \quad \forall A \in As_R$ ;
- $P_R(C) = P_{KB}(C) \quad \forall C \in Cs_R$ ;  $P_R(I|pa_R(I)) = P_{KB}(I|\mathfrak{C}(I)) \quad \forall I \in Is_R$ ;  
and for each  $A \in As_R$ ,  $P_R(A|pa_R(A))$  is defined by the CPD table  $\text{proj}(Tb, pa_R(A))$  where  $Tb$  is the CPD table for  $A$  in  $W$ , i.e. defined by  $P_W(A|pa(A))$ .

Based on these operators, we now describe an algorithm for incremental intention recognition in a real-time manner.

**Algorithm 1 (Incremental Intention Recognition)** *Repeat the following steps until some given time limit is reached. The most likely intention in the previous cycle of the repeat loop is the final result <sup>4</sup>.*

---

<sup>4</sup>Obviously, one can easily modify the algorithm to find the set of  $N$  most likely intentions. In the next section we shall see that in more detail.

1. Let  $A$  be a new observed action. Combine the current IRBN  $W$  with  $\text{irBN}(A)$  to obtain  $W' = \text{comb}(W, \text{irBN}(A))$ . If  $A$  is the initially observed action, let  $W' = \text{irBN}(A)$ .
2. Compute the probability of each intention in  $W'$ , conditional on the set of current observations in  $W'$ . Remove the intentions which are much less likely than the others (following Definition 2.4.7).

**Example 2.4.4 (Elder’s Incremental Intention Recognition)** When observing the action *Looking*, the IRBN in Figure 2.2 (without action node *Open\_Fridge*) is created. We compute the probabilities<sup>5</sup>  $V_1, V_2, V_3, V_4$  of each intention  $i(B), i(D), i(R),$  and  $i(S)$ , respectively, conditional on the observations, including action node *Looking* and the two cause/reason nodes, *tv\_on* and *light\_on*. Let us consider the possible cases.

- If light is off, then  $V_1 = V_2 = V_3 = 0$  and  $V_4 = 1$ , regardless of the observed state of the TV.
- If light is on and tv is off, then  $V_1 = 0.71, V_2 = 0.49, V_3 = 0.50,$  and  $V_4 = 0.011$ .
- If light is on and tv is on, then  $V_1 = 0, V_2 = 0.50, V_3 = 0.80,$  and  $V_4 = 0.01$ .

Hence, if one observes that the light is off, the elder is definitely looking for the light switch, given that he is looking around. Otherwise, if one observes the light is on, and that the TV is on too, the intention of looking for the remote control,  $i(R)$ , is most probable; but, if the TV is off, then looking for something to read,  $i(B)$ , is most probable. They are the output of the algorithm whenever the decision needs to be made immediately after observing the first action.

Now suppose we are in the second case (light is on and tv is off), and the decision has not had to be made. Another action, open the fridge – *Open\_fridge* (*OpenF*), is observed. Suppose there are two unit fragments for this action,  $UF_{\mathfrak{A}}(i(D), \text{OpenF})$ , and  $UF_{\mathfrak{A}}(i(F), \text{OpenF})$ , and one unit fragment for the intention  $i(F)$ ,  $UF_{\mathfrak{J}}(\{Hungry\}, i(F))$ , in the knowledge base (where  $i(F)$  stands for the intention of looking for some food). Respectively, CPD tables in these unit fragments

---

<sup>5</sup>In this work, for Bayesian network reasoning and inference, we utilize the Bayesian reasoning engine SMILE running in a MAC-OS-X computer, publicly available at: <http://genie.sis.pitt.edu/>.

are given by,  $P(\text{OpenF} = T | i(D) = T) = 0.3$ ,  $P(\text{OpenF} = T | i(F) = T) = 0.8$ , and  $P(i(F) = T | \text{Hungry} = T) = 0.9$ . A new IRBN with the new action node is created (Figure 2.2 in full, with all grey-filled nodes). Note that now the intention node  $i(\text{Switch})$ , and with all its in- and out- connections, is being removed, since it is very unlikely compared with others (see above,  $V_4 = 0.011$ ).

Now the conditional probabilities of intentions in the IRBN are:  $V_1 = 0.57$ ,  $V_2 = 0.55$ ,  $V_3 = 0.47$ , and  $V_5 = 0.63$  ( $V_5$  is that of the intention  $i(\text{Food})$ ). Looking for food becomes the most probable intention.

## 2.5 Relation Among Intentions

When considering the case in which the observed agent may pursue multiple intentions simultaneously, it is undoubtedly indispensable to take into account and express the relations amongst the intentions in the model. Pursuing one intention may exclude some other intention to be pursued (Pynadath and Wellman, 1995; Bratman, 1987; Roy, 2009b). It may be so because of some resource limitation, e.g., allowance time is not enough for accomplishing both intentions at the same time (Keng and Yun, 1989). It also may be because of the nature or restriction of the observed agent’s task: the agent is restricted to pursuing a single intention (e.g., in constructing Linux and Unix plan corpora, a user is given one task at a time to complete) (Blaylock and Allen, 2003, 2004; Lesh, 1998).

We introduce a so-called *exclusive relation*  $e$ —a binary relation on the set of intention nodes—representing that if one intention is pursued, then the other intention cannot be pursued. It is usually, although perhaps not always, the case that intentions exclusiveness is symmetric. It holds for the resource limitation case: one intention excludes the other intention because there is not enough resource for accomplishing both, which in turn implies that the latter intention excludes the former one too. It also clearly holds for the case where the agent is restricted to pursuing a single intention. Here we assume that  $e$  is symmetric; it can be renamed *mutually exclusive relation*.

Intentions  $I_1$  and  $I_2$  are mutually exclusive iff they cannot be pursued simultaneously, i.e.  $P(I_1 = T, I_2 = T) = 0$ . Thus, for any action  $A$ , if  $I_1, I_2 \in pa(A)$  then the CPD table for  $A$  is undefined. Hence, the BN needs to be restructured. The mutually exclusive intentions must be combined into a single intention node since they cannot

co-exist as parents of a node. Each intention represents a possible value of the new combined node. Namely, let  $I_1, \dots, I_t$  be such that  $e(I_i, I_j), \forall i, j : 1 \leq i < j \leq t$ . The new combined node,  $I$ , stands for a random variable whose possible outcomes are either  $I_i, 1 \leq i \leq t$ , or  $\tilde{I}$ —the outcome corresponding to the state that none of the  $I_i = T$ . Note that if the intentions are exhaustive,  $\tilde{I}$  can be omitted. Next,  $I$  is linked to all the action nodes that have a link from one of  $I_i, 1 \leq i \leq t$ .

There remains to re-define CPD tables in the new BN. They are kept the same for action  $A$  where  $I \notin pa(A)$ . For  $A$  such that  $I \in pa(A)$ , the new CPD table at  $I = I_i$  corresponds to the CPD table in the original BN at  $I_i = T$  and  $I_j = F \forall j \neq i$ , i.e.  $P(A|I = I_i, \dots) = P(A|I_0 = F, \dots, I_{i-1} = F, \mathbf{I}_i = \mathbf{T}, I_{i+1} = F, \dots, I_t = F, \dots)$ . Note that the left hand side is defined in the new BN, and the right hand side is defined in the original BN. Similarly, the new CPD table at  $I = \tilde{I}$  corresponds to  $I_i = F$  for all  $1 \leq i \leq t$ . In addition, prior probability  $P(I = I_i) = P(I_i = T)$  and  $P(I = \tilde{I}) = \prod_{i=1}^t P(I_i = F)$  (and then being normalized).

We now specify the CPD table of  $I$ . In the new BN, the causes/reasons of each intention are connected to the combined node, i.e.  $pa(I) = \bigcup_{i=1}^t \mathfrak{C}(I_i)$ . Applying the Markov assumption (see Def. 2.3.1) we have  $P(I = I_i|pa(I)) = P_i(I_i = T|\mathfrak{C}(I_i))$  and  $P(I = \tilde{I}|pa(I)) = \prod_{i=1}^t P_i(I_i = F|\mathfrak{C}(I_i))$ , where  $P_i$  is the probability distribution of the unit fragment for  $I_i$ .

In the next section we focus on the single intention recognition case, showing how the approach to representing relationships amongst several intentions can significantly decrease the complexity of the probability inference therein. We then present experimental results on the Linux plan corpus. After that, in Section 2.7, we provide further experimentation on our novel so-called IPD plan corpora.

## 2.6 Single Intention Recognition

### 2.6.1 The Model

Suppose the observed agent pursues a single intention at a time. In this case, all intentions are mutually exclusive, and they can be combined into a single node. The IRBN then has a single intention node, linking to all action nodes. All cause/reason nodes are connected to the intention node.

Let  $I_1, \dots, I_n$  be the intentions in the original IRBN. As usual, they are assumed to be exhaustive, i.e. the observed agent is assigned an intention from them. The

combined node  $I$  thus has  $n$  possible outcomes  $I_i$ ,  $1 \leq i \leq n$ . Let  $As = \{A_1, \dots, A_m\}$  be the set of current observed actions, all linked to the single intention  $I$ . The set of all cause/reason nodes are  $Cs = \cup_{i=1}^n \mathfrak{C}(I_i)$ . Suppose  $C_e \subseteq Cs$  is the set of cause/reason nodes which are observed (evidence nodes). For instance, in the Elder Care examples presented above, the state of TV and the state of the light are observed cause/reason nodes.

Let  $C_{ne} = Cs \setminus C_e$ . Applying Eq. (2.2), we obtain the probability of each intention conditional on the current observations

$$P(I = I_j | C_e, As) = \frac{P(I_j, C_e, As)}{\sum_{i=1}^n P(I_i, C_e, As)} \quad \text{for all } 1 \leq j \leq n \quad (2.5)$$

where, by applying the joint probability formulas (2.1) and (2.3), we obtain for all  $j$  that

$$P(I_j, C_e, As) = \prod_{i=1}^m P(A_i | I_j) \left( \sum_{C_{ne}} P(I_j | Cs) \prod_{C \in Cs} P(C) \right). \quad (2.6)$$

This implies that, when not including causes/reasons for intentions (i.e.,  $Cs = \emptyset$ ) as in the case of Linux plan corpus below, our intention recognizer has a linear complexity on the number of intentions  $O(|n|)$ .

If no cause/reason nodes are observed, i.e.  $C_{ne} = Cs$  (as in the case of the Linux and IPD plan corpora we examine in the next subsection), we obtain

$$P(I_j, C_e, As) = P(I_j) \prod_{i=1}^m P(A_i | I_j). \quad (2.7)$$

If all of them are observed, i.e.  $C_{ne} = \emptyset$  (as we shall see in the IPD Plan corpora), the term  $\prod_{C \in Cs} P(C)$  is simplified in the fraction of Eq. (2.5), since it appears and is the same in both numerator and denominator. Thus, in these two cases, we do not need to define prior probabilities distribution of the root nodes in  $Cs$ , as to be applied to obtain experimental results for the Linux and IPD plan corpora in the following sections. Note that in the latter case we still need to compute the conditional probabilities  $P(I_j | Cs)$ , which will be explicitly defined when needed.

## 2.6.2 Experiments with Linux Plan Corpus

### 2.6.2.1 The Linux Plan Corpus

Plan corpus is the term used to describe a set of plan sessions and consists of a list of goals/intentions and the actions a user executed to achieve them (Armentano and Amandi, 2009). Although there are many corpora available for testing machine learning algorithms in other domains, just a few are available for training and testing plan/intention recognizers; furthermore, each of the plan/intention recognizers using plan corpora usually has its own datasets—which leads to a difficult comparison amongst each other. For that important reason, we chose the Linux plan corpus (Blaylock and Allen, 2004)—one of the rare regularly used plan corpora—which was kindly made publicly available by Nate Blaylock—in order to test our system. It also enables a better comparison with other systems using this corpus (Blaylock and Allen, 2005, 2004; Armentano and Amandi, 2009; Singla and Mooney, 2011).

The Linux plan corpus is modeled after Lesh’s Unix plan corpus (Lesh, 1998). It was gathered from 56 human users (graduate and undergraduate students, faculty, and staff) from the University of Rochester Department of Computer Science. The users have different levels of expertise in the use of Linux, and they were allowed to perform as many times as they wished, in order to contribute more plan sessions. The sessions, consisting in sequences of commands performed by the users to achieve a given goal/intention, were automatically recorded. For example, a goal is to find a file with a given name or copy some files to a given folder, and the users can use the Linux commands such as “find”, “cp”, “cd”, “ls”, etc. At the end of each session, the users were asked to indicate whether they succeeded in achieving their goal/intention. In total, there are 547 sessions, 457 of which were indicated as successfully completing the goal, 19 goals and 43 actions (commands).

The Linux plan corpus is an important and hard benchmark for intention/goal recognition. First, data is collected from real humans and thus noisy. Second, involved humans expertise is varied, and they sometimes used wrong commands due to limited knowledge about the domain (Blaylock and Allen, 2004; Linux-Plan-Corpus). Furthermore, we observe that plan sessions’ lengths in the corpus are quite varied. The minimum, maximum, and mean number of actions in a plan session are 1, 60, and 6.124, respectively.

### 2.6.2.2 Learning Unit Fragments from Data

For unit fragment  $UF_{\mathfrak{A}}(I, A)$ , the conditional probability of  $A$  given  $I$  is defined by the frequency of  $A$  in a plan session for achieving the goal/intention  $I$  divided by the frequency of any action for achieving  $I$ :

$$P(A = T|I = T) = \frac{freq(A_I)}{freq(I)}. \quad (2.8)$$

For better understanding, in the plan corpus each action is marked with the intention which the action is aiming at. Then,  $freq(A_I)$  is the frequency of  $A$  being marked by  $I$ , and  $freq(I)$  is the frequency of seeing the mark  $I$ .

Note that prior probabilities of all the intentions in the corpus are given initially, and used for generating tasks for users (Linux-Plan-Corpus; Blaylock and Allen, 2004).

### 2.6.2.3 Making Predictions

Similar to (Blaylock and Allen, 2004; Armentano and Amandi, 2009), instead of letting the recognizer make a prediction after each observed action, we set a *confidence* threshold  $\tau$  ( $0 \leq \tau \leq 1$ ), which allows the recognizer to decide whether or not it is confident enough to make a prediction; the recognizer only makes a prediction if the likelihood of the most likely intention in the model is greater than  $\tau$ . Otherwise, it predicts “don’t know”.

In addition, instead of only predicting the most likely intention, the recognizer provides a set of  $N$  most likely ones (*N-best prediction*).

### 2.6.2.4 Evaluation Metrics

For evaluating our system and comparing with the previous ones (Blaylock and Allen, 2004; Armentano and Amandi, 2009), we use three different metrics. *Precision* and *recall* report the number of correct predictions divided by total predictions (predicts when confident enough) and total prediction opportunities (always predicts, whatever the confidence is), respectively. More formally (Armentano and Amandi, 2009), let  $Seq = a_1, \dots, a_n$  be a sequence of actions for achieving intention

$I$  (a session for testing). Considering the  $N$ -best prediction case, let

$$correct(A) = \begin{cases} 1 & \text{if } I \text{ is one of the } N \text{ most likely intentions when observing } A \\ 0 & \text{otherwise} \end{cases}$$

Then, precision and recall for  $Seq$  are defined as follows, respectively,

$$precision(Seq) = \frac{\sum_{i=1}^n correct(a_i)}{z},$$

$$recall(Seq) = \frac{\sum_{i=1}^n correct(a_i)}{Z}$$

where  $z$  and  $Z$  are the number of predictions made (when the recognizer is confident enough) and the total number of prediction opportunities, respectively.


On the other hand, *convergence* is a metric that indicates how much time the recognizer took to converge on what the current user goal/intention was. Let  $t$  be such that  $correct_i = 0$  for  $0 \leq i \leq t - 1$  and  $1$  for  $t \leq i \leq n$  (i.e.  $t$  is the first time point which from there on the system always correctly predicts), convergence for sequence  $Seq$  is

$$convergence(Seq) = \frac{z - t + 1}{z}.$$

Finally, the overall precision, recall and convergence are obtained by taking averages over all testing sessions.

### 2.6.2.5 Experiments and Results

Because of the small size of the Linux corpus, and similar to previous work, we ran experiments using the one-out cross validation method (Armentano and Amandi, 2009). Just one at a time, one plan session in the whole corpus is left out. The rest of the corpus is used for training the model, which is then evaluated against the left out plan session. We study the effect of confidence level  $\tau$  with respect to the precision and convergence (for recall, it clearly is a decreasing function of  $\tau$ ) (Figure 2.6). As expected, the greater  $N$  the better the precision and convergence scores. The difference in the precision and convergence between two different values of  $N$  is large when  $\tau$  is small, and gets smaller for greater  $\tau$ . Most interestingly, we observe that the precision and convergence are not monotonic increasing on  $\tau$ . There are *critical values* of  $\tau$  at which the measures have the maximal value, and those



**Figure 2.6:** Plot of our method’s precision and convergence for  $\tau \in [0, 1]$  and for different values of  $N$  ( $N = 1, 2, 3, 4$ ) in Linux plan Corpus. The greater  $N$ , the better precision and convergence. The difference in precision and convergence between two different values of  $N$  is large when  $\tau$  is small, and gets smaller for greater  $\tau$ . Most interestingly, we observe that precision and convergence are not monotonic increasing on  $\tau$ . There are *critical values* of  $\tau$  at which the measures have maximal value, and those values are smaller for greater  $N$ . This observation suggests that in plan/intention recognition task, the more precise (i.e. the smaller  $N$ ) the decision needed to make is, the greater confidence level the recognizer should gain to make a good (enough) decision. On the other hand, the recognizer should not be too cautious to make a prediction, leading to refuse to make a prediction when it would have been able to make a correct one. In short, it suggests the important need to study (experimentally) the confidence threshold  $\tau$  carefully for particular application domains.

values are smaller for greater  $N$ . This observation suggests that in a plan/intention recognition task, the more precise (i.e. the smaller  $N$ ) the decision needed to make is, the greater confidence level the recognizer should gain to make a good (enough) decision. On the other hand, the recognizer should not be too cautious, possibly leading to refuse to make a prediction when it would have been able to make a correct one. In short, these experiments suggest an important need to (experimentally) study the confidence threshold  $\tau$  carefully, for particular application domains and particular values of  $N$ . Using the same  $\tau$  for all values of  $N$  could decrease the recognizer’s performance.

Table 2.1 shows some of the results for different values of  $N$  (and the corresponding value of  $\tau$ ). Similar to the previous works on the same Linux corpus (Blaylock and Allen, 2004; Armentano and Amandi, 2009), we keep the best results of each

**Table 2.1:** Intention Recognition Results on the Linux Plan Corpus

<b>N-best</b>	<b>1-best</b>	<b>2-best</b>	<b>3-best</b>	<b>4-best</b>
$\tau$	0.95	0.5	0.45	0.42
<b>Precision</b>	0.786	0.847	0.870	0.883
<b>Recall</b>	0.308	0.469	0.518	0.612
<b>Converg.</b>	0.722	0.799	0.822	0.824

case with respect to  $\tau$  for the comparison. For example, we obtained a precision of 78.6% for 1-best that is increased to 87.0% for 3-best prediction and 88.3% for 4-best one. Convergence is increased from 72.2% for 1-best to 82.2% for 3-best and 82.4% for 4-best prediction.

The best performance on the Linux plan corpus so far, in terms of precision and convergence (recall is not referred), has been reported in (Armentano and Amandi, 2009), where the authors use a variable Markov model with exponential moving average. In this respect we obtained an increment of 14% better precision and 13.3% better convergence for 1-best prediction, 8.2% better precision and 9.3% better convergence for 2-best prediction, and 7.5% better precision and 7.7% better convergence for 3-best prediction. We also obtained better recalls in all cases and for all metrics of 4-best prediction compared with (Blaylock and Allen, 2004), the only work reporting these scores. For the Linux plan corpus, Singla and Mooney (2011) reports only the precision score, which is worse than that of Armentano and Amandi (2009).

Note that in (Armentano and Amandi, 2009) the authors use a more fine-grained preprocessing method for their work, but we suspect it will have increased their performance. To fairly compare with all previous work we use the original corpus.

The Linux plan corpus allows an appropriate comparison with prior work. However, it does not include contextual information (reasons/causes of intentions), and there is no intention change/abandonment occurrences (users follow a single intention throughout entire plan sessions). To evaluate the context-dependent aspect as well as the capability of dealing with intention change/abandonment, we next present new plan corpora.

## 2.7 IPD Plan Corpora

We present new plan corpora in the context of the iterated Prisoner’s Dilemma (IPD) <sup>6</sup> (Sigmund, 2010) and provide experimental results for them. The intentions/goals to be recognized are the (known) strategies in IPD (see below). Strategies in the (evolutionary) game theory context are usually considered representative types of agents’ behavior in a social or biological setting. An agent that adopts a given strategy can be interpreted as if it intends to follow the corresponding type of behavior, in the given social or biological setting. This interpretation satisfies the generally accepted definition of intention, which is ‘choice with commitment’ (Bratman, 1987; Cohen and Levesque, 1990). Plan sessions are sequences of moves played by such strategies. This way of interpreting game strategies was explored in (Han et al., 2011b) for studying the evolutionary roles of intention recognition.

### 2.7.1 Iterated Prisoner’s Dilemma

Prisoner’s Dilemma (PD) is a symmetric two-player non-zero game defined by the payoff matrix

$$\begin{array}{cc} & \begin{array}{cc} C & D \end{array} \\ \begin{array}{c} C \\ D \end{array} & \left( \begin{array}{cc} R, R & S, T \\ T, S & P, P \end{array} \right). \end{array}$$

Each player has two options in each round, cooperates (C) or defects (D). A player who chooses to cooperate with a player who defects receives the sucker’s payoff  $S$ , whereas the defecting player gains the temptation to defect,  $T$ . Mutual cooperation (resp., defection) yields the reward  $R$  (resp., punishment  $P$ ) for both players. PD is characterized by the payoff ranking  $T > R > P > S$  (and, in addition,  $2R > S + T$  for IPD). Thus, in a single round, it is always best to defect, but cooperation may be rewarded if the game is iterated. Let  $r$  denote the average number of rounds the game is iterated.

IPD is usually known as a story of tit-for-tat (TFT), which won both Axelrod’s tournaments (Axelrod, 1984). *TFT* starts by cooperating, and does whatever the opponent did in the previous round. It will cooperate if the opponent cooperated,

---

<sup>6</sup>The approach also applies readily to other social dilemmas such as Snow Drift and Stag Hunt (Sigmund, 2010). A social dilemma represents a particular type of situation in a social or biological context.

and will defect if the opponent defected. But if there are erroneous moves due to noise (i.e. an intended move is wrongly performed with a given execution error), the performance of *TFT* declines: it cannot correct errors or mistakes. Tit-for-tat is then replaced by generous tit-for-tat (*GTFT*), a strategy that cooperates if the opponent cooperated in the previous round, but sometimes cooperates even if the opponent defected (with a fixed “forgiveness” probability  $p > 0$ ) (Nowak and Sigmund, 1992; Sigmund, 2010). *GTFT* can correct mistakes. Subsequently, *TFT* and *GTFT* were replaced by win-stay-lose-shift (*WSLS*) as the winning strategy chosen by evolution (Nowak and Sigmund, 1993; Sigmund, 2010). *WSLS* repeats the previous move whenever it did well, but changes otherwise.

Some other less famous strategies which we are going to use later are *GRIM* – a grim version of *TFT*, prescribing to defect except after a round of mutual cooperation, and Firm-But-Fair (*FBF*) – known as a tolerant brother of *TFT*, prescribing to defect only if getting a sucker’s payoff  $S$  in the previous round. Details of all strategies summarily depicted above can be found in (Sigmund, 2010) (Chapter 3).

### 2.7.2 Corpus Description

In the following, we prescribe how to create plan corpora for training and testing the described intention recognition method, for a given set of strategies. We start by making an assumption that all strategies to be recognized have the memory size bounded by  $M \geq 0$ , i.e. their decision at the current round is independent of the past rounds that are at a time distance greater than  $M$ . Note that the above mentioned strategies have memory bounded by  $M = 1$ .

For clarity of representation, abusing notations,  $R$ ,  $S$ ,  $T$  and  $P$  are henceforth also referred to as elementary game states, in a single round of interaction. Additionally,  $E$  (standing for *empty*) is used to refer to a game state having had no interaction. The most basic element in a plan corpus is the corpus action, having the following representation.

**Definition 2.7.1 (Corpus Action)** *An action in a plan corpus is of the form  $s_1 \dots s_M \xi$ , where  $s_i \in \{E, R, T, S, P\}$ ,  $1 \leq i \leq M$ , are the states of the  $M$  last interactions, and  $\xi \in \{C, D\}$  is the current move <sup>7</sup>.*

---

<sup>7</sup>From now on, this notion of an action is used, which is different from the notion of a move (either C or D).

**Definition 2.7.2 (Plan Session)** *A plan session of a strategy is a sequence of corpus actions played by that strategy (more precisely, of a player using that strategy) against an arbitrary player.*

We denote by  $\Sigma_M$  the set of all possible types of action for memory size  $M$ . Clearly,  $|\Sigma_M| = 2 \times 5^M$ . For example,

$$\Sigma_1 = \{EC, RC, TC, SC, PC, ED, RD, TD, SD, PD\}.$$

Now, for an example of a plan session, let us consider the strategy *TFT* and the following sequence of its interactions with some other player (denoted by  $X$ ), in the presence of noise

<b>round :</b>	0	1	2	3	4	5
<b>TFT :</b>	–	<i>C</i>	<i>C</i>	<i>D</i>	<i>D</i>	<i>D</i>
<b>X :</b>	–	<i>C</i>	<i>D</i>	<i>D</i>	<i>C</i>	<i>D</i>
<i>TFT-states :</i>	<i>E</i>	<i>R</i>	<i>S</i>	<i>P</i>	<i>T</i>	<i>P</i>

The corresponding plan session for *TFT* is  $[EC, RC, SD, PD, TD]$ . At the 0-th round, there is no interaction, thus the game state is *E*. *TFT* starts by cooperating (1-st round), hence the first action of the plan session is *EC*. Since player  $X$  also cooperates in the 1-st round, the game state at this round is *R*. *TFT* reciprocates in the 2-nd round by cooperating, hence the second action of the plan session is *RC*. Similarly for the third and the fourth actions. Now, at the 5-th round, *TFT* should cooperate since  $X$  cooperated in 4-th round, but because of noise, it makes an error to defect. Therefore, the 5-th action is *TD*.

This way of encoding actions and the assumption about the players' memory size lead to the equivalent assumption that the action in the current round is *independent* of the ones in previous rounds, regardless of the memory size. The independence of actions will allow to derive a convenient and efficient intention recognition model, discussed in the next subsection. Furthermore, it enables to save us the game states without having to save the co-player's moves, thus simplifying the representation of plan corpora.

**Definition 2.7.3 (Plan Corpus)** *Let  $\mathcal{S}$  be a set of strategies to be recognized. A plan corpus for  $\mathcal{S}$  is a set of plan sessions generated for each strategy in the set.*

For a given set of strategies, different plan corpora can be generated for different purposes. In the sequel, for example, we generate different plan corpora for training and for testing the intention recognition method.

### 2.7.3 Plan Corpora Generation

Let us start by generating a plan corpus for seven most popular strategies within the IPD framework: *AllC* (always cooperate), *AllD* (always defect), *TFT*, *GTFT* (probability of forgiving a defect is  $p = 0.5$ ), *WSLS*, *GRIM* and *FBF* (described above).

We collect plan sessions of each strategy by playing a random move (C or D) with it in each round. To be more thorough, we can also play all possible combinations for each given number of rounds  $r$ . E.g, if  $r = 10$ , there are  $2^{10} = 1024$  combinations: C or D in each round. When noise is present, each combination is played repeatedly a number of times, since each time one might obtain different reaction moves from the simulated co-player.

The training corpus to be used here is generated by playing with each strategy all the possible combinations 10 times, for each number of rounds  $r$  from 5 to 10. The testing dataset is generated by playing a random move with each strategy in each round, also for  $r$  from 5 to 10<sup>8</sup>. We continue until obtaining the same number of plan sessions as for the training dataset (corpus). Both datasets are generated in the presence of noise. Namely, an intended move is wrongly performed, with respect to the definition of the strategy, with probability 0.05.

In this testing dataset, changes/abandonment of players' initial intention (strategy) are not allowed. The players use the same strategy in all the rounds. We refer to this testing dataset as **Testset-IRFIX**.

For testing the context-dependent aspect of our intention recognition method, as well as taking into account intention changes/abandonment, we next introduce the concept of *social learning* within the framework of evolutionary game theory (Sigmund et al., 2010; Hofbauer and Sigmund, 1998; Tuyls and Parsons, 2007).

---

<sup>8</sup>It might seem unnatural to interact with a strategy randomly, but it enables us to collect a highly diverse set of move sequences of the strategy, for both training and testing. It allows to thoroughly evaluate our intention recognizer, even if our additional experiments confirm that its relative performance (shown in Subsection 2.7.5) would not change much when the strategy interacts with a big enough set of strategies.

### 2.7.4 Social Learning


We consider it necessary for an agent to acquire knowledge from other agents, i.e. learn “by being told” instead of learning only by experience. Indeed, this social learning is a fairly practical and economical way of increasing abilities, widely used by human beings, as widely studied in evolutionary biology and economics (Richardson and Boyd, 2006; Rendell et al., 2010; Sigmund et al., 2010). Let us see how social learning can be modeled in Evolutionary Game Theory (EGT) (Hofbauer and Sigmund, 1998; Sigmund, 2010) given a fixed set of strategies. Agent strategies can change, say through mutation or learning from experience, but we shall not consider that issue here.

In social learning, agents in a population can observe the behavior of others and the outcomes of that behavior. They may copy the behavior of others whenever these appear to be more successful (Sigmund, 2010; Sigmund et al., 2010; Rendell et al., 2010). The accumulated payoff from all interactions emulates the agents’ *fitness* or social *success* and the most successful agents will tend to be imitated by others. There are many ways to model social learning (Hofbauer and Sigmund, 1998; Sigmund et al., 2010). The most popular one is implemented using the so-called pairwise comparison rule (Traulsen et al., 2006; Sigmund, 2010): an agent **A** with fitness  $f_A$  will adopt the strategy of a randomly chosen agent **B** with fitness  $f_B$  with a probability given by

$$p(f_A, f_B) = \left(1 + e^{-\beta(f_B - f_A)}\right)^{-1}, \quad (2.9)$$

where the quantity  $\beta$  controls the “imitation strength”, i.e. how strongly the players are basing the decision to imitate on payoff comparisons. Henceforth, **A** and **B** are referred to as imitating and imitated agents, respectively. For simplicity, we use  $\beta = 1$  in the remainder of this chapter: the imitation depends on comparing the exact payoffs.

It is now allowed the possibility that a player can change his/her strategy (intention) by imitating the randomly met player’s strategy (intention), depending on how much the latter player is more successful. The two players’ ongoing success difference (SD) causally affects the imitating player’s current intention. In addition, this intention is causally affected by the so-called imitation event (IE), stating whether the player is meeting some other player for learning/imitating.



**Figure 2.7: Bayesian Network for intention recognition (IRBN) in the context of iterated Prisoner’s Dilemma**, where players can change their initial intention (strategy) by imitating successful others, based on a social learning model (Hofbauer and Sigmund, 1998; Sigmund et al., 2010; Rendell et al., 2010). The IRBN have two cause/reason nodes in the first layer ( $IE$  and  $SD$ ), connecting to a single intention node in the second layer ( $I$ ), in turn connecting to action nodes in the third layer.

In an everyday life situation, an imitation event,  $IE$ , can be more generally thought of as a contact of the intending agent with some source of information that might influence its current intention to change. For example, an intention of going out might be influenced to change/abandon by information about weather, traffic, etc., whether the source of that information is a person, a newspaper, TV or Radio. In turn, the success difference,  $SD$ , measures how strong is the influence on the change, based on the credibility or strength of the information. The above mentioned intention of going out is undoubtedly influenced more by the forecast information provided by a trustworthy TV channel than an untrustworthy one or unreliable person. Another example is that a suggestion from a doctor, after some tests, apparently has more effect on changing/abandoning the intention of drinking alcohol than that from a friend, or that from a daily newspaper.

Now we can specify a Bayesian network for intention recognition IRBN with two cause/reason nodes, a single intention node, and observed action nodes (Figure 2.7). In the following we define the CPD tables of the IRBN.

**Bayesian Network definition.** Recall that the intention node  $I$  has  $n$  possible outcomes  $I_i$ ,  $1 \leq i \leq n$  (Subsection 2.6.1). We define the conditional probability distribution  $P(I_i|IE, SD)$ . If the player does not meet any other player for imitation (i.e.,  $IE = F$ ),  $I_i$  is independent of the success difference  $SD$ . Hence,  $P(I_i|IE = F, SD) = P(I_i|IE = F)$ . Now, let us consider the case  $IE = T$ . If the

successes are also observable (thus,  $SD$  is observed, say, equal  $\chi$ )<sup>9</sup>, but the strategy of the imitated player is not, we have

$$P(I_i|IE = T, SD = \chi) = (1 - u)p_i + \frac{u}{S - 1} \sum_{j \neq i} p_j \quad (2.10)$$

where  $u = (1 + e^{-\chi})^{-1}$ ,  $p_i$  is the probability that  $I_i$  was the player's intention in the last prediction, and  $S$  is the number of strategies in the corpus. The equation is explained as follows. With probability  $(1 - u)p_i$  the imitating player's strategy remains  $I_i$ . Moreover, when not being observed, the probability that  $I_i$  was the imitated player's strategy is (assumed) equal  $1/(S - 1)$ . The second term expresses the probability that the player adopts the new strategy  $I_i$  by imitation.

In the case the imitated player's strategy is also observed, denoted by  $I_{i^*}$ , similarly we have


$$\begin{aligned} P(I_{i^*}|IE = T, SD = \chi) &= (1 - u)p_{i^*} + u \sum_{j \neq i^*} p_j, \\ P(I_i|IE = T, SD = \chi) &= (1 - u)p_i \quad \forall i \neq i^*. \end{aligned} \quad (2.11)$$

**Testing Dataset Generation (IRCHANGE).** The testing dataset in this setting is generated by using a simplified evolutionary simulation involving the seven strategies described in Subsection 2.7.3. Similarly to the previous dataset, we play a random choice sequence with each of the seven above mentioned strategies for 10 rounds. The payoff of each strategy is accumulated over these rounds. Then, for each strategy, another strategy is randomly chosen from the other six, for possible imitation. The imitation decision is guided by the pairwise comparison rule (see Eq. (2.9)), which is based on a comparison of the performance of the two strategies, in terms of the difference of their accumulated payoffs in these 10 rounds.

After all the seven strategies are given the chance to change their strategy by imitating another, the interactions are repeated for 10 more rounds. At the 10-th round, we save the accumulated payoff values of the imitating and imitated strategies. We experiment until obtaining the same number of plan sessions as in the training

---

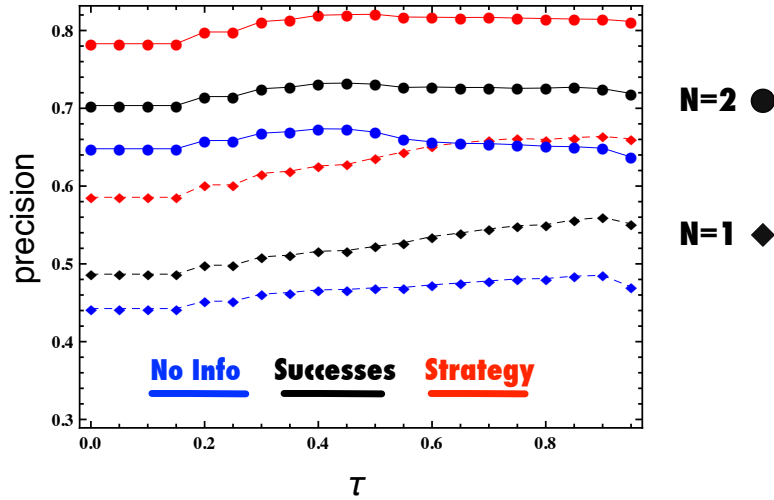
<sup>9</sup>There may be noise in the evaluation of the successes. The observed value  $\chi$  of  $SD$  is randomly taken in the range  $((1 - \epsilon)\chi_1, (1 + \epsilon)\chi_1)$ , where  $\epsilon$  is a small positive number (here we use  $\epsilon = 0.01$ ) and  $\chi_1$  is the exact value of the difference.



**Figure 2.8: Precision and convergence for  $\tau \in [0, 1]$  and for different values of  $N$  ( $N = 1, 2, 3$ ) with respect to Testset-IRFIX dataset, where  $N$  is the number of most likely considered intentions. The greater  $N$ , the better the precision and convergence. For high enough confidence  $\tau$ , the intention recognizer obtained quite good scores.**

dataset. The PD payoff matrix being used is:  $T = 20$ ,  $R = 15$ ,  $P = 10$ ,  $S = 5$ ; and  $noise = 0.05$ . This testing dataset is referred to as **Testset-IRCHANGE**.

Clearly, each plan session in the dataset consists of 20 actions. The first 10 actions are played by the initial strategy. The last 10 ones are played by a strategy which might be either the initial one or a different one, as the outcome of imitation (or social learning). The intention recognition system must identify both the first and the second strategies. While identifying the first strategy can be done efficiently without taking into account contextual information, as we have seen in the previous datasets, identifying the second strategy is more difficult due to the stochastic change of the strategies (resulting from the stochasticity from the comparison rule). In the course of observing the first 10 actions generated by the initial strategy, the system (on average) converges to a correct prediction of it. If the strategy is changed to a different one, and the system has no clue about that change, it would poorly predict the new strategy, since it would still consider the old actions generated by the first strategy as belonging to the new strategy. Having (contextual) clues about the potential effects leading to strategy change could significantly enhance the prediction of the second strategy, as we shall see from our experimental results in the sequel.



**Figure 2.9:** Precision for three different levels of contextual information, for  $\tau \in [0, 1]$ , with respect to `Testset-IRCHANGE` dataset. We plot for  $N = 1$  (dashed diamond) and  $N = 2$  (circle).

### 2.7.5 Results on IPD Plan Corpora

The intention recognition model is acquired using the training corpus. Figure 2.8 shows the precision and convergence of the model with respect to the `Testset-IRFIX`, where again  $N$  is the number of most likely considered intentions. Given that the training as well as the testing datasets are generated in the presence of noise, the achieved performance is quite good. Namely, for a big enough  $\tau$ , both precision and convergence scores are greater than 0.9, even for the 1-best case. It is noteworthy that, in Chapter 5 (see also (Han et al., 2011b)), a similar good performance is obtained for a smaller set of four strategies (intentions). Therein, based on the intention recognition model above presented, we implement a novel strategy capable of intention recognition that outperforms other strategies in an evolutionary setting.

In Figure 2.9 we show the effects of having different levels of contextual information on the intention recognition performance, using `Testset-IRCHANGE` dataset, where we consider  $N = 1$  (dashed diamond) and  $N = 2$  (circle). Namely, in the first case (blue curves), there is no information about the imitation event (IE) – it is not known if the recognized player may imitate and adopt another strategy. In the second case (black curves), *IE* and the successes are observable. In the third case

(red curves), the strategy of the imitated player is also observable (see a discussion below for how these three levels of contextual information can be mapped into real world situations). It is clearly shown that the performance is considerably increased as more contextual information is available. Namely, comparing with the first case where no contextual information is taken into account, an increase of about 5% and 15% precision is achieved in the second and third cases, respectively.

These results show that, besides observable actions from the recognizing agent, it is important to take into account contextual information for intention recognition. That is what most previous work on intention/plan recognition has omitted, with the exception of Pynadath and Wellman (1995), though the authors did not provide experimental results to show its importance (see Related Work in Section 2.2).

Note that our three-layer Bayesian network model described in Figure 2.7 is not restricted to context-dependent intention recognition in the (evolutionary) game theory setting. The two cause/reason nodes in the BN, imitation event (*IE*) and success difference (*SD*), account for key relevant contextual information for intention recognition. The first one stands for relevant factors that might influence the stability of current intentions in the model, while the second measures how strong are such influences. Conceivably, those are all the relevant contextual information needed for the context-dependent intention recognition task in a given application domain.

We next further discuss the rationale of using (evolutionary) game-theoretic methods and concepts for modeling real world situations.

### 2.7.6 Scope Enlargement Discussion

The task of intention/plan recognition is mostly concerned with actions or behavior. In previous work, for a given application domain, these actions are supposed to be supplied by an *activity* recognition system. In contrast, here we aim at providing a general framework (and benchmarks) for studying and evaluating intention/plan recognition models. For that, it is crucial to have a general model of agent behavior and its changes.

We adopted here one of the most popular frameworks for modelling agents' behavior (including those amongst humans), the social dilemmas (not only the Prisoner's dilemma) and social learning within the framework of evolutionary game theory (Hofbauer and Sigmund, 1998; Sigmund, 2010; Rendell et al., 2010). They

have been widely used in Economics, Biology, Psychology, Artificial Intelligence, etc. to describe human (and other species) behavior (Axelrod, 1984; Sigmund et al., 2010; Traulsen et al., 2010; Sigmund, 2010; Rendell et al., 2010; Tuyls and Parsons, 2007). Although they do not exhibit all aspects of human behavior in the real world, they have been, experimentally as well as analytically, shown to reflect its core nature (most famously, in the context of the evolution of cooperation study (Axelrod, 1984; Traulsen et al., 2010; Morgan et al., 2011; Santos and Pacheco, 2011; Sigmund, 2010)). For example, the notorious strategies described above (including *TFT*, *WSLS*, *GTFT*) frequently show up in many real-world experimental settings, ranging from Global Warming issues to Psychology and Economics behavioral researches (Milinski et al., 2006; Falk et al., 2008; Morgan et al., 2011). They are also widely observed in a broad range of real human social behavior. For a review of social learning and evolutionary games in the real world the readers are referred to the seminal experimental studies in (Traulsen et al., 2010; Morgan et al., 2011).

Using game-theoretic concepts allows a higher level of abstraction of domain knowledge. That is, it enables us to account for and generalize similar situations in different application domains. As a result, our IPD benchmarks above presented are deemed domain-independent. Let us consider a simple scenario in the Elder Care domain to illustrate how it is mapped to game theoretic concepts, namely the three levels of contextual information, “No Info”, “Successes”, and “Strategy”, alluded above (Figure 2.9). Suppose an elder intends to go fishing the next morning. But ‘being told’ that evening that it is going to rain then, it was suggested (he intend) to do something else. Whether he changes (or even abandons) his intention strongly depends on by whom he was ‘being told’. Was it his 3-year old granddaughter or the TV/Radio forecast? That is, whether he adopts another person’s suggested intention (of not going fishing the next morning) depends on that person’s expertise or past success in the matter, or the credibility of the information provided. It seems almost sure that he would drop the intention of going fishing if it was said by the TV/Radio forecast. In order to appropriately recognize the elder’s intention, it is crucial to account for such contextual information, and that it was encountered. Further, if he was also ‘being told’ by the forecast that, although it is going to rain in the morning, the weather in the afternoon should be perfect, one can predict that he might likely to adopt the intention of going fishing in the afternoon instead. Briefly, it corresponds to the first case, “No Info”, if the system cannot observe (despite its

actual occurrence) that the elder was ‘being told’ by another person. The second case, “Successes”, corresponds to when the system can observe who passed him the information. The third case, “Strategy”, corresponds to when the system also observes what he was advised by another.

In short, using the models of evolutionary game theory and social learning, one can generate general-purpose or domain-independent benchmarks for intention/plan recognition systems. The applicability of a system to such benchmarks may serve to show its applicability to a wider range of application domains.

## 2.8 More on Situation Sensitiveness

In the following we discuss some extensions to the above intention recognition model, which enable to take into account contextual information in several manners. We mainly focus on exploiting the powerful knowledge representation tool of Logic Programming (LP) (Baral, 2003)<sup>10</sup>.

### 2.8.1 Situation-sensitive Bayesian Networks

Undoubtedly, BNs should be situation-sensitive since using a general BN for all specific situations of a problem domain is unrealistic and most likely imprecise (Xiang and Poh, 2005; Laskey, 2008). For example, in the Elder Care domain, different elders might have different living conditions and habits that need to be taken into account to recognize their intentions. Also, place, time of day, temperature, etc. need to be considered (Pereira and Han, 2011b; Han and Pereira, 2010b). However, consulting the domain expert to manually change the BN with respect to each situation is costly or unfeasible.

In (Pereira and Han, 2009c, 2011b), we described a method to construct situation-sensitive BNs, i.e. ones that change according to the given situation. It uses LP techniques to compute situation specific probabilistic information that is then updated into the BN. The BNs themselves are also encoded with P-log,

---

<sup>10</sup>Here we consider solely normal logic programs, which consist of a set of rules of the form

$$A \leftarrow B_1, \dots, B_m, \text{ not } C_1, \dots, \text{ not } C_n \quad (m, n \geq 0)$$

where  $A, B_1, \dots, B_m, C_1, \dots, C_n$  are domain atoms, and  $\text{not } C_i$  denotes the default negation of atom  $C_i$ , i.e. being false by default. The rule reads, if all the atoms and default negations in the body are true, then the head  $A$  is true. If the body is empty, that is,  $m = n = 0$ , the head is true unconditionally. In this case, it has the form,  $A \leftarrow$  or simply  $A$ , and is also called a *fact*.

a probabilistic logic implemented system (Baral et al., 2009; Han et al., 2008, 2009), which supports coherent updates. The LP techniques employed for this are deduction with top-down procedure (XSB-Prolog) (XSB, 2009) (to deduce situation-specific probabilistic information) plus integrity constraints and abduction (to abduce probabilistic information needed to support and explain observations in the given situation). However, we can employ too various other types of LP based reasoning, e.g., constraint satisfaction, contradiction removal, preferences, or inductive learning, whose results can be compiled (in part) into an evolving BN.

Consider three-layer BN for intention recognition (IRBN), we assume that only prior probabilities of the top nodes (causes/reasons) are subject to changes along with the changing situation. The CPD tables for intention and action nodes reflect the internal state of the recognizing agent about world (Heinze, 2003; Tahboub, 2006). They may change in the long-term periods, but just occasionally. In short, for a given situation at hand, an IRBN can be situated in it by recomputing the prior probabilities of the cause/reason nodes in accordance with the situation.

Based on this idea, we now revise the Algorithm 1, making the IRBN model situation-sensitive. For that, we first describe a new operator. In a given situation, an IRBN can be situated by recomputing the prior probabilities of the top nodes.

**Definition 2.8.1 (Situate IRBN)** *Let  $W = \langle \{Cs, Is, As\}, pa, P_W \rangle$  be an IRBN. We say that  $W$  is situated into a situation  $SIT$  if the prior probabilities of the top nodes of  $W$ , i.e.  $P_W(C)$  ( $C \in Cs$ ), are recomputed according to  $SIT$ . In this work, a situation is encoded by a logic program and the prior probabilities of the top nodes are computed using LP techniques, as described above (Pereira and Han, 2009c, 2011b). Formally, the situate operator is defined by  $\text{situate}(W, SIT) = \langle \{Cs, Is, As\}, pa, P_S \rangle$ , where*

- $P_S(C)$ , for all  $C \in Cs$ , are the new prior probabilities of top nodes, resulted from the re-computation according to  $SIT$ .
- $P_S(X|pa(X)) = P_W(X|pa(X)) \forall X \in Is \cup As$ .

It might be too costly to always change the IRBN as the situation relevant factors could constantly evolve. Here, we propose a simple criterion allowing to decide when the IRBN should be reconfigured to account for the situation at hand. It says, if there is a conceivable “salient” intention currently figuring in the IRBN, the IRBN should be reconfigured (after observing a new action), that is, be situated according

to the latest information about the situation at hand. Otherwise, the IRBN remains the same. The property “salient” reads differently in different application domains, and it is up to the domain experts to design its specification. For example, in the Elder Care domain, “salient” may read *dangerous*, e.g., when suspecting an intrusion intention to elders’ house, or the elders’ committal suicide intention. It may also read *urgent*, e.g., when detecting the elders’ intention to find something to eat or something to drink.

Accordingly, a simple change to the Algorithm 1 needs to be made.

**Algorithm 2** *After the first step (item 1) in Algorithm 1, insert the following item:*

- *If there is a “salient” intention in the current IRBN  $W'$ , situate it according to the situation at hand  $SIT$ , i.e.  $\text{situate}(W', SIT)$ . Otherwise, the IRBN remains the same.*

For illustration, let us extend the previous example in the Elder Care domain.

**Example 2.8.1 (Elder Care (cont’d))** *In the scenario provided in Example 2.4.4, the IRBN model (after observing and including both actions) may vary depending on some factors such as the time of day, of the elders’ last drink or last meal, his interest in football, etc. For illustration, we design a simple logical component for the IRBN to take into account those factors.*

```

pa_rule(pa(hg(T),d_(0,1)),[])    :- time(X), last_eating(X1), X-X1 < 1.
pa_rule(pa(hg(T),d_(9,10)),[])  :- time(X), last_eating(X1), X-X1 > 3.

pa_rule(pa(thsty(T),d_(1,10)),[]) :- time(X), last_drink(X1), X1-X < 1.
pa_rule(pa(thsty(T),d_(9,10)),[]) :- time(X), last_drink(X1), X1-X > 3.

pa_rule(pa(lr(T),d_(1,100)),[])  :- time(X), X > 0, X < 5.

pa_rule(pa(lw(T),d_(9,10)),[])   :- time(X), schedule(X1,football), X1-X < 0.25, !.
pa_rule(pa(lw(T),d_(1,100)),[])  :- time(X), (X > 23; X < 5).
    
```

Basically, in P-log, probabilistic information is given by  $pa/2$  rules (Han et al., 2008; Baral et al., 2009). For example, the rule  $(pa(hg(T), d_(9, 10)) \leftarrow Body)$  means that the probability of being hungry (i.e.  $hg(T)$ ) is 9/10 if the precondition in the body of the rule,  $Body$ , holds. We provide a reserved  $pa\_rule/2$  predicate which takes the head and body of some  $pa/2$  rule as its first and second arguments, respectively,

and includes preconditions for its activation in its own body. Thus, e.g., the second *pa\_rule* above means, if the elder’s last eating is more than 3 hours ago, it is quite probably that he is being hungry (with probability 9/10) (note that here the *pa* rule body, representing the precondition, is empty—represented by an empty list []). Now, a situation is given by asserted facts representing it and, in order to find the probabilistic information specific to the given situation, we simply use the XSB Prolog built-in *findall/3* predicate (XSB, 2009) to find all true *pa/2* literals expressed by the *pa\_rule/2* rules with true bodies in the situation.

There are several predicates specific to the domain knowledge base. For example, *time/1* provides the current time of the day (in hours). The predicates *last\_drink/1* and *last\_eating/1*, with an argument in time, hold if the elder’s last drink and last meal were at the given time, respectively.

Now let us exhibit how some different situations can be encoded in our framework, and how they appropriately affect the intention recognition results, in a context-dependent manner.

- If the current time is 18 (i.e. *time(18)*) and the last time the elders ate was half an hour before (i.e. *last\_eating(17.5)*). But they did not have any drink for 3 hours (e.g., *last\_drink(14)*). Those three facts are asserted. Hence, the following two *pa\_rule/2* literals are true, and are updated into the general IRBN

```
pa_rule(pa(hg(T),d_(0,1)), []).
```

```
pa_rule(pa(thty(T),d_(9,10)), []).
```

Now the result is,  $V_1 = 0.55$ ;  $V_2 = 0.77$ ;  $V_3 = 0.41$ ;  $V_5 = 0.21$ . It is now the case that looking for something to drink, *i(Drink)*, is the most likely intention, instead of looking for food as previously.

- If the elder also had the last drink just half an hour before (i.e. *last\_drink(17.5)*), then we have the following two rules instead

```
pa_rule(pa(hg(T),d_(0,1)), []).
```

```
pa_rule(pa(thty(T),d_(1,10)), []).
```

We obtain,  $V_1 = 0.60$ ;  $V_2 = 0.14$ ;  $V_3 = 0.54$ ;  $V_5 = 0.12$ . In other words, the intention of looking for something to read, *i(Book)*, becomes the most likely intention.

- If the situation remains the same, but the time is changed to *1 a.m.*, that is, the facts *time(1)*, *last\_eating(0.5)*, and *last\_drink(0.5)* hold. And if there is no football scheduled around that time, we have two additional rules

```
pa_rule(pa(lr(T),d_(1,100)), []).
```

```
pa_rule(pa(lw(T),d_(1,100)), []).
```

We thus obtain,  $V_1 = 0.40$ ;  $V_2 = 0.14$ ;  $V_3 = 0.29$ ;  $V_5 = 0.121$ , i.e., the intention of looking for something to read, *i(Book)*, remains the most likely intention. But, if there is football scheduled around that time, say, *schedule(1.15, football)*, then we have additionally the following two rules instead

```
pa_rule(pa(lr(T),d_(1,100)), []).
```

```
pa_rule(pa(lw(T),d_(9,10)), []).
```

Now the result is,  $V_1 = 0.13$ ;  $V_2 = 0.14$ ;  $V_3 = 0.21$ ;  $V_5 = 0.12$ . In other words, the intention of looking for the remote control now becomes the most likely one.

In short, we have demonstrated the flexibility of our technique by means of representing situations as logic programs, and used them to update the current IRBN. We show how the context information can appropriately influence the interpretation of the observed agent's intention.

### 2.8.2 Situation-sensitive selection of unit fragments

In all the plan corpora benchmarks we have been using in previous sections, no context-dependent selection of unit fragments for actions is needed. In operator *select(A, SIT)* (recall from Subsection 2.4.1), the parameter *SIT* is always empty. However, the selected set of intentions for an action should in general be context dependent, that is, which intentions conceivably give rise to the action should depend on the situation in which the action is observed (Bratman, 1992, 1987; Roy, 2009b). Commonsense reasoning can be employed for this purpose, which will be enacted by LP reasoning techniques (Baral, 2003) in the following.

The prior domain knowledge base KB is accompanied by a logic program  $P_{KB}$  in order to help decide which unit fragments for an observed action are selected in

a given situation, i.e., which intentions are conceivable. Let  $\Delta = \{A_1, \dots, A_N\}$  be the set of actions of KB and  $\Upsilon_i = \{I \mid UF_{\mathfrak{A}}(I, A_i) \in KB\}$  the set of intentions that belong to a unit fragment for action  $A_i$  in KB.

We say that an intention  $I$  is conceivable when observing action  $A$  if it is expected in the given situation and there is no expectation to the contrary (a similar technique was used in (Pereira and Han, 2009b, 2011b; Han and Pereira, 2011c)). Thus, for  $1 \leq i \leq N$  and  $I \in \Upsilon_i$ ,  $P_{KB}$  contains the following rule:

$$\text{conceivable}(I) \leftarrow A_i, \text{expect}(I), \text{not expect\_not}(I).$$

Furthermore, for each  $I \in \bigcup_{i=1}^N \Upsilon_i$ ,  $P_{KB}$  contains two rules:

$$\begin{aligned} \text{expect}(I) &\leftarrow \text{Cond}_1, \\ \text{expect\_not}(I) &\leftarrow \text{Cond}_2. \end{aligned} \tag{2.12}$$

The rules about expectations are domain-specific knowledge used to constrain the conceivable intentions in a situation. Counter-expectation rules supplement expectation rules for representing defeasible conditions or exceptions.

Now suppose that an action  $A$ ,  $A \in \Delta$ , is observed. The current situation is encoded by a logic program  $SIT$ , consisting of LP facts describing the situation. In order to compute the set of conceivable intentions that may give rise to  $A_i$ , we simply use the XSB Prolog built-in *findall/3* predicate (XSB, 2009) to find all true *conceivable/1* atoms of the program  $P_{KB} \cup SIT \cup \{A \leftarrow\}$ . This provides the operational definition of the *select(A, SIT)* operator discussed in Subsection 2.4.1.

For illustration, we modified the previous example in Elder Care domain, as follows.

**Example 2.8.2 (Elder Care variant)** *An elderly person stays alone in his apartment. An intention recognition system is set up to support his activities in the living room. At the moment the system observes that the elder is looking around for something (look). The knowledge base KB of the system has a unit IRBN for this action. For illustration, consider a small set of conceivable intentions,  $Is = \{\text{book}, \text{water}, \text{weapon}, \text{lightSwitch}\}$ .*

The accompanying logic program  $P_{KB}$  contains the following rules, for each  $I \in Is$ :

$$\text{conceivable}(I) \leftarrow \text{look}, \text{expect}(I), \text{not expect\_not}(I).$$

Suppose in  $P_{KB}$  the expectation and counter-expectation rules for these intentions are

1.  $expect(book)$ .  
 $expect\_not(book) \leftarrow light\_off$ .  
 $expect\_not(book) \leftarrow burglar\_alarm\_ring$ .
2.  $expect(water)$ .  
 $expect\_not(water) \leftarrow light\_off$ .  
 $expect\_not(water) \leftarrow burglar\_alarm\_ring$ .
3.  $expect(weapon) \leftarrow burglar\_alarm\_ring$ .  
 $expect\_not(weapon) \leftarrow light\_off$ .  
 $expect\_not(weapon) \leftarrow no\_weapon\_availabe$ .
4.  $expect(lightSwitch)$ .  
 $expect\_not(lightSwitch) \leftarrow light\_on, tv\_on$ .

For example, the rules in line 1 say the intention of looking for a book is always expected except when the light is off or the burglar alarm is ringing. Let us consider some cases. If at the moment the light is off, i.e.  $SIT = \{light\_off \leftarrow\}$ , then  $conceivable(light\_switch)$  is the only true *conceivable/1* atom of the program  $P_{KB} \cup SIT \cup \{look \leftarrow\}$ . In this case, since there is only one conceivable intention, we can conclude immediately that looking for the light switch is the elder's intention, without having to construct an IRBN.

Now suppose the light is on, the tv is not on, and the burglar alarm is not ringing. That is,  $SIT = \{light\_on \leftarrow\}$ . Recall that what we consider here are normal logic programs, with default negation. It means that if something is not stated true as a fact, its negation is true by default. There are three conceivable intentions, *book*, *water*, and *lightSwitch*, since they are expected in this situation, and there are no expectations to the contrary (lines 1, 2, and 4). Hence, they are selected for constructing the unit IRBN for action *look*, given the current situation. The intention *weapon* is irrelevant, or not conceivable, and can be ruled out immediately at this situation-dependent selection stage.

If light is on, tv is not on, but now the burglar alarm is ringing, i.e.  $SIT = \{light\_on \leftarrow, burglar\_alarm\_ring \leftarrow\}$ . *Weapon* is accordingly expected, and there is no expectation to the contrary (line 3). Furthermore, *water* and *book*

are not expected (line 2). Hence, there are two conceivable intentions: *weapon* and *lightSwitch*.

If, additionally, the tv is on, that is,  $SIT = \{light\_on \leftarrow, tv\_on \leftarrow, burglar\_alarm\_ring \leftarrow\}$ , then the intention of looking for the light switch is not expected (line 4). Hence, looking for *weapon* is the only conceivable intention recognizable in this situation.

Despite its simplicity, in considering only a small set of possible conceivable intentions, this example has demonstrated how the LP techniques can help to (significantly) reduce the size of the result IRBN, and sometimes reach the solution immediately. We envisage that the approach can provide even greater advantage, in terms of time complexity, when the set of conceivable intentions is large.

In this section we have presented two LP techniques which aim at enabling and enhancing context-dependent intention recognition. Taking into account contextual information can help significantly reduce the size of the IRBN in which the BN inference is performed. Although much work remains to be done—including gathering real data or generating appropriate plan corpora to evaluate them—the proposed techniques are quite promising. We admit that all these might be done using BNs solely, it however appears that their combination can provide a more efficient solution. The BN inference is much costlier than LP inference. For example, consider a (definite) logic program that encodes all the causal relations of a BN. In the general case, while the inference, e.g. of a conditional probability in the BN, is exponential, the corresponding deductive inference in the logic program can be performed in polynomial time.

## 2.9 Conclusions and Future Work

We have presented a novel method for context-dependent and incremental intention recognition. The method is performed by dynamically constructing a three-layer BN model for intention recognition (IRBN), from a prior knowledge base consisting of readily maintained and constructed fragments of BN. Their simple structures allow easy maintenance by domain experts or automatically building from available plan corpora. The three-layer IRBN follows a *causal* intentional structure (Heinze, 2003; Tahboub, 2006), from causes/reasons (of intentions) to intentions (as causes

of actions), then to (actually observed) actions. This causal structure enabled us to appropriately capture different aspects of the context-dependent intention recognition, from the context-dependent selection of intentions for the model construction to sensitizing prior probabilities of the intentions already in the model.

The presented above examples in the Elder Care domain illustrate several aspects of our intention recognition method, highlighting the importance of taking into account contextual information – an aspect usually omitted in the previous work (Pynadath and Wellman, 1995; Geib and Goldman, 2009). In addition, for the first time, we have shown experimentally that accounting for contextual information is crucial for the recognition accuracy when the observed agents might change or abandon their initial goals/intentions. It is apparently an unavoidable aspect of real agents, clearly pointed out in (Geib and Goldman, 2003, 2009). Taking into account such information enables us to account for the reasons why the agents change or abandon their goals/intentions.

Our method performs particularly well on the Linux plan corpus, showing its applicability to the important interface-agents domain (Horvitz et al., 1998; Armentano and Amandi, 2007; Madani et al., 2009). It outperforms the existent methods that make use of the corpus. For further experimentation, we have presented the so-called IPD plan corpora for the famous strategies in the context of the iterated Prisoner’s Dilemma. We employed the popular model of human behavior enacted by means of social learning and evolutionary game theory to simulate intention changes/abandonment—thus enabling us to evaluate the context-dependent aspect of our intention recognizer as well as its capability for dealing with intention changes/abandonment. Given that this modelling approach has been widely adopted and employed in several fields, as diverse as Economics, Psychology and Biology (Hofbauer and Sigmund, 1998; Sigmund, 2010), the good performance of the method for the IPD corpora makes it highly applicable for a wide range of application domains therein. Particularly in Part II of the thesis (see also (Han et al., 2011b,a)), we employ this intention recognition method to implement strategies that outperform the most successful known strategies in the iterated Prisoner’s Dilemma.

Herein we have also attempted to tackle the problem where an observed agent may follow multiple intentions simultaneously in a more appropriate manner. We formally described how to represent relationships amongst intentions in the Bayesian network for intention recognition, particularly in order to maintain its consistency

when one needs to combine mutually exclusive intentions (Pynadath and Wellman, 1995). This aspect is indispensable in multiple intentions recognition, but mostly omitted in previous work. However, the scalability of our method remains to be seen. For its evaluation we need to gather an appropriate plan corpus allowing for the possibility that users might pursue multiple intentions simultaneously. An idea is that instead of giving users one task at a time (in the case of the Linux or Unix domain (Lesh, 1998; Blaylock and Allen, 2004)), they are given several, so that they can complete them separately or in parallel. Another method is to use a planner to generate interleaved plans for multiple intentions.

A limitation of the current formalization in the multiple intentions recognition case is that we need to assume that the intentions to be combined are perfectly mutually exclusive. This assumption can be relaxed by utilizing a latent variable for any subset of perfectly mutually exclusive intention nodes. The latent variable figures in the BN either as a child or parent of the nodes, whichever works better for inference. We are exploring this direction to provide a more general method for representing relationships amongst intention nodes.

Another limitation of our current method is that it did not explicitly take into account temporal evolution of domain variables in the BN. It is usually done using Dynamic Bayesian Networks (DBNs) (Pynadath and Wellman, 1995; Forbes et al., 1995; Albrecht et al., 1998), where the state of each variable is represented by a series of nodes. In our method, the time evolution, to some degree, is implemented by means of updating the IRBN from time to time. For example, the states of the cause/reason nodes are updated using an external logic program which represents the evolving world. In this way not only can one significantly reduce the size of the BN, thereby its inference complexity (as the LP inference is significantly less expensive), but the declarative representation and reasoning of the LP techniques (Baral, 2003) could be important when the states of a node cannot be easily represented in an explicit manner. However, in domains where the states of the nodes might constantly change, the explicit time series representation of DBNs is apparently necessary. We envisage to bring in, to some degree, the time series representation of DBNs (Dean and Kanazawa, 1992; Forbes et al., 1995; Albrecht et al., 1998), to improve our method.

# Chapter 3

---

## INTENTION BASED DECISION MAKING AND APPLICATIONS

---

*Each decision we make, each action we take,  
is born out of an intention.*

Sharon Salzberg, *O Magazine*,  
*The Power of Intention, January 2004*

**Abstract.** In this chapter we present an intention-based decision making system. We exhibit a coherent combination of two Logic Programming based implemented systems, Evolution Prospection and Intention Recognition. The Evolution Prospection system has proven to be a powerful system for decision making, designing and implementing several kinds of preferences and useful environment-triggering constructs. It is here enhanced with an ability to recognize intentions of other agents—an important aspect not well explored so far. The usage and usefulness of the combined system are illustrated with several extended examples in different application domains, including Ambient Intelligence and Elder Care.

---

<sup>0</sup>This chapter is based on the publications: (Han and Pereira, 2011c), (Han and Pereira, 2010c), (Pereira and Han, 2011b), and (Pereira and Han, 2011a).

### 3.1 Introduction

Given the important role that intentions play in the way we make decisions (Bratman, 1987; Roy, 2009b; Searle, 2010), one would expect intentions to occupy a substantial place in any theory of action. Surprisingly enough, in what is perhaps the most influential theory of action—rational choice theory (Russell and Norvig, 2003; Roy, 2009b)—which includes the theory of decision making—explicit reference is made to actions, strategies, information, outcomes and preferences but not to intentions.

This is not to say that no attention has been paid to the relationship between rational choice and intentions. Quite the contrary, a rich philosophical and AI literature has developed on the relation between rationality and intentions (Bratman, 1987; Cohen and Levesque, 1990; Singh, 1991; van Hees and Roy, 2008). Some philosophers, e.g. in (Bratman, 1987, 1999; Roy, 2009b), have been concerned with the role that intention plays in directing rational decision making and guiding future actions. In addition, many agent researchers have recognized the importance of intentions in developing useful agent theories, architectures, and languages, such as Rao and Georgeff with their BDI model (Rao and Georgeff, 1995), which has led to the commercialization of several high-level agent languages, e.g. in (Burmeister et al., 2008; Wooldridge, 2000, 2002). However, to the best of our knowledge, there has been no real attempt to model and implement the role of intentions in decision making, within a rational choice framework. Intentions of other relevant agents are always assumed to be given as the input of a decision making process; no system that integrates a real intention recognition system into a decision making system has been implemented so far.

In this chapter, we present a coherent Logic Programming (LP) based framework for decision making—which extends our previous work on Evolution Prospection for decision making (Pereira and Han, 2009a,b; Han, 2009)—but taking into consideration now the intentions of other agents. Obviously, when being immersed in a multi-agent system, knowing the intentions of other agents can benefit the recognizing agents in a number of ways. It enables the recognizing agents to predict what other agents will do next or might have done before—thereby being able to plan in advance to take the best advantage from the prediction, or to act so as to take remedial action. In addition, an important role of recognizing intentions is to enable coordination of your own actions and in collaborating with others (Bratman, 1999,

1987; Roy, 2009b; Searle, 1995, 2010; Kaminka et al., 2002). We have also recently shown the role of intention recognition in promoting improved cooperative behavior in populations or societies of self-interested agents (Han et al., 2011a,b) (this will be described in Part II of the thesis).

The Evolution Propection (EP) system is an implemented LP-based system for decision making (Pereira and Han, 2009a,b)—which will be described in Section 3.2. An EP agent can prospectively look ahead a number of steps into the future to choose the best course of evolution that satisfies a goal. This is achieved by designing and implementing several kinds of prior and post preferences (Pereira et al., 2011a; Pereira and Lopes, 2009) and several useful environment-triggering constructs for decision making.

In order to take into account intentions of other agents in decision making processes, we employ the intention recognition system described in the previous chapter as an external module of the EP system. For an easy integration, the Bayesian network inference of the intention recognition system is performed by P-log (Baral et al., 2009; Han et al., 2008, 2009), a probabilistic logic system <sup>1</sup>.

## 3.2 Evolution Propection

### 3.2.1 Preliminary

The implemented EP system <sup>2</sup> has proven useful for decision making (Pereira and Han, 2009a,b; Pereira and Lopes, 2009; Han et al., 2012d). It is implemented on top of ABDUAL <sup>3</sup>, a preliminary implementation of (Alferes et al., 2004), using XSB Prolog (XSB, 2009). We next describe the constructs of EP, to the extent we use them here. A full account can be found in (Pereira and Han, 2009b; Han, 2009).

**Language** Let  $\mathcal{L}$  be a first order language. A domain literal in  $\mathcal{L}$  is a domain atom  $A$  or its default negation *not*  $A$ . The latter is used to express that the atom is false

---

<sup>1</sup>The implementation of P-log systems described in (Baral et al., 2009) can be found in:  
<http://www.cs.ttu.edu/~wezhu/>

<sup>2</sup>The implementation of the Evolution Propection system can be downloaded at:  
<http://centria.di.fct.unl.pt/~lmp/software/epa.zip>

<sup>3</sup>The implementation of ABDUAL system can be downloaded at:  
<http://centria.di.fct.unl.pt/~lmp/software/contrNeg.rar>

by default (Closed World Assumption). A domain rule in  $\mathcal{L}$  is a rule of the form:

$$A \leftarrow L_1, \dots, L_t \quad (t \geq 0)$$

where  $A$  is a domain atom and  $L_1, \dots, L_t$  are domain literals. An integrity constraint in  $\mathcal{L}$  is a rule with an empty head. A (logic) program  $P$  over  $\mathcal{L}$  is a set of domain rules and integrity constraints, standing for all their ground instances.

Here we consider solely Normal Logic Programs (NLPs), those whose heads of rules are positive literals, or empty (Baral, 2003). We focus furthermore on abductive logic programs (Kakas et al., 1993; Alferes et al., 2004), i.e. NLPs allowing for abducibles – user-specified positive literals without rules, whose truth-value is not fixed. Abducibles instances or their default negations may appear in bodies of rules, like any other literal. They stand for hypotheses, each of which may independently be assumed true, in positive literal or default negation form, as the case may be, in order to produce an abductive solution to a query.

**Definition 3.2.1 (Abductive Solution)** *An abductive solution is a consistent collection of abducible instances or their negations that, when replaced by true everywhere in  $P$ , affords a model of  $P$  (for the specific semantics used on  $P$ ) which satisfies the query and the ICs – a so-called abductive model.*

**Active Goals** In each cycle of its evolution the agent has a set of active goals or desires. We introduce the *on\_observe/1* predicate, which we consider as representing active goals or desires that, once triggered by the observations figuring in its rule bodies, cause the agent to attempt their satisfaction by launching all the queries standing for them, or using preferences to select them. The rule for an active goal AG is of the form:

$$on\_observe(AG) \leftarrow L_1, \dots, L_t \quad (t \geq 0)$$

where  $L_1, \dots, L_t$  are domain literals. During evolution, an active goal may be triggered by some events, previous commitments or some history-related information. When starting a cycle, the agent collects its active goals by finding all the *on\_observe(AG)* that hold under the initial theory without performing any abduction, then finds abductive solutions for their conjunction.

**Preferring Abducibles** An abducible  $A$  can be assumed only if it is a considered one, i.e. if it is expected in the given situation, and, moreover, there is no expectation to the contrary

$$\text{consider}(A) \leftarrow \text{expect}(A), \text{not expect\_not}(A), A.$$

The rules about expectations are domain-specific knowledge contained in the theory of the program, and effectively constrain the hypotheses available in a situation.

Handling preferences over abductive logic programs has several advantages, and allows for easier and more concise translation into NLPs than those prescribed by more general and complex rule preference frameworks. The advantages of so proceeding stem largely from avoiding combinatory explosions of abductive solutions, by filtering irrelevant as well as less preferred abducibles (Pereira et al., 2011b,a).

To express preference criteria among abducibles, we envisage an extended language  $\mathcal{L}^*$ . A preference atom in  $\mathcal{L}^*$  is of the form  $a \triangleleft b$ , where  $a$  and  $b$  are abducibles. It means that if  $b$  can be assumed (i.e. considered), then  $a \triangleleft b$  forces  $a$  to be assumed too if it can. A preference rule in  $\mathcal{L}^*$  is of the form:

$$a \triangleleft b \leftarrow L_1, \dots, L_t \quad (t \geq 0)$$

where  $L_1, \dots, L_t$  are domain literals over  $\mathcal{L}^*$ . The semantics of this preference rule can be found in (Pereira et al., 2011b,a).

*A priori* preferences are used to produce the most interesting or relevant conjectures about possible future states. They are taken into account when generating possible scenarios (abductive solutions), which will subsequently be preferred amongst each other *a posteriori*.

**Example 3.2.1** Consider a situation where I need to choose to drink either tea or coffee (but not both). I prefer coffee over tea when sleepy, and do not drink coffee when I have high blood pressure. This situation is described by the program with abducibles `coffee` and `tee`:

```
abds[tea/0, coffee/0].
on_observe(drink).
drink ← tea.      drink ← coffee.
expect(tea).      expect(coffee).
```

*expect\_not(coffee) ← blood\_high\_pressure.*  
*← tea, coffee.*  
*coffee ≺ tea ← sleepy.*

This program has two abductive solutions, one with *tea* and the other with *coffee*. Adding literal *sleepy* triggers the only *a priori* preference in the program, which defeats the solution where only *tea* is present (due to the impossibility of simultaneously abducing coffee). If later we add *blood\_pressure\_high*, coffee is no longer expected, and the transformed preference rule no longer defeats the abduction of *tea* which then becomes the single abductive solution, despite the presence of *sleepy*.

**A Posteriori Preferences** Having computed possible scenarios, represented by abductive solutions, more favorable scenarios can be preferred a posteriori. Typically, *a posteriori* preferences are performed by evaluating consequences of abducibles in abductive solutions. An *a posteriori* preference has the form:

$$A_i \ll A_j \leftarrow \text{holds\_given}(L_i, A_i), \text{holds\_given}(L_j, A_j)$$

where  $A_i, A_j$  are abductive solutions and  $L_i, L_j$  are domain literals. This means that  $A_i$  is preferred to  $A_j$  a posteriori if  $L_i$  and  $L_j$  are true as the side-effects of abductive solutions  $A_i$  and  $A_j$ , respectively, without any further abduction when testing for the side-effects. Optionally, in the body of the preference rule there can be any Prolog predicate used to quantitatively compare the consequences of the two abductive solutions.

**Evolution Result A Posteriori Preference** While looking ahead a number of steps into the future, the agent is confronted with the problem of having several different possible courses of evolution. It needs to be able to prefer amongst them to determine the best courses from its present state (and any state in general). The *a posteriori* preferences are no longer appropriate, since they can be used to evaluate only one-step-far consequences of a commitment. The agent should be able to also declaratively specify preference amongst evolutions through quantitatively or qualitatively evaluating the consequences or side-effects of each evolution choice.

A *a posteriori* preference is generalized to prefer between two evolutions. An *evolution result a posteriori* preference is performed by evaluating consequences of

following some evolutions. The agent must use the imagination (look-ahead capability) and present knowledge to evaluate the consequences of evolving according to a particular course of evolution. An *evolution result a posteriori preference* rule has the form:

$$E_i \lll E_j \leftarrow \text{holds\_in\_evol}(L_i, E_i), \text{holds\_in\_evol}(L_j, E_j)$$

where  $E_i, E_j$  are possible evolutions and  $L_i, L_j$  are domain literals. This preference implies that  $E_i$  is preferred to  $E_j$  if  $L_i$  and  $L_j$  are true as evolution history side-effects when evolving according to  $E_i$  or  $E_j$ , respectively, without making further abductions when just checking for the side-effects. Optionally, in the body of the preference rule there can be recourse to any Prolog predicate, used to quantitatively compare the consequences of the two evolutions for decision making.

### 3.3 Evolution Propection with Intention Recognition

There are several ways an EP agent can benefit from the ability to recognize intentions of other agents, both in friendly and hostile settings. Knowing the intention of an agent is a means to predict what he will do next or might have done before. The recognizing agent can then plan in advance to take the best advantage of the prediction, or act to take remedial action. Technically, in the EP system, this new kind of knowledge may impinge on the body of any EP constructs, such as active goals, expectation and counter-expectation rules, preference rules, integrity constraints, etc., providing a new kind of trigger.

In order to account for intentions of other agents in decision making with EP, we provide a built-in predicate,  $\text{has\_intention}(Ag, I)$ , stating that an agent  $Ag$  has the intention  $I$ . The truth-value of this predicate is evaluated by the intention recognition system. Whenever this predicate is called in an EP program, the intention recognition system is employed to check if  $Ag$  has intention  $I$ , i.e.  $I$  is the most likely conceivable intention at that moment. We also provide predicate  $\text{has\_intention}(Ag, I, Pr)$ , stating that agent  $Ag$  has intention  $I$  with probability  $Pr$ . Hence, one can express, for example, the situation where one needs to be more, or less, cautious.

One can also generalize to consider the  $N$ -best intention recognition approach (see Chapter 2), that is, to assess whether the intention of the agent is amongst

the  $N$  most likely intentions. It has been shown that by increasing  $N$ , the recognition accuracy is significantly improved (Blaylock and Allen, 2004; Armentano and Amandi, 2007; Han and Pereira, 2011b).

In the sequel we draw closer attention to some EP constructs, illustrating with several examples how to take into account intentions of other agents for enhancement of decision making.

### 3.3.1 Intentions Triggering Active Goals

Recall that an active goal has the form

$$on\_observe(AG) \leftarrow L_1, \dots, L_t \quad (t \geq 0)$$

where  $L_1, \dots, L_t$  are domain literals. At the beginning of each cycle of evolution, those literals are checked with respect to the current evolving knowledge base and trigger the active goal if they all hold. For intention triggering active goals, the domain literals in the body can be in the form of *has\_intention* predicates, taking into account intentions of other agents.

This way, any intention recognition system can be used as the goal producer for decision making systems, the inputs of which are (active) goals to be solved (see for instance (Pereira and Han, 2011a,b; Han and Pereira, 2010c)).

It is easily seen that intention triggering active goals are ubiquitous. New goals often appear when one recognizes some intentions in others. In a friendly setting, one might want to help others to achieve their intention, which is generally represented as follows

```
on_observe(help_achieve_goal(G)) <-
    friend(P), has_intention(P,G),
```

while in a hostile setting, we probably want to prevent the opponents from achieving their goals

```
on_observe(prevent_achieve_goal(G)) <-
    opponent(P), has_intention(P,G).
```

Or, perhaps we simply want to plan in advance to take advantage of the hypothetical future obtained when the intending agent employs the plan that achieves his intention

```
on_observe(take_advantage(F)) <- agent(P),
        has_intention(P,G), future(employ(G),F).
```

Let us look a little closer at each setting, providing some ideas how they can be enacted. When helping someone to achieve an intention, what we need to do is to help him/her with executing a plan achieving that intention successfully, i.e., all the actions involved in that plan can be executed. This usually occurs in multi-agent collaborative tasks (see for example (Kaminka et al., 2002)), wherein the agents need to be able to recognize their partners' intention to secure an efficient collaboration.

In contrast, in order to prevent an intention from being achieved, we need to guarantee that all possible plans achieving that intention cannot be executed successfully. To that effect, at least one action in each plan must be prevented if the plan is conformant (i.e., a sequence of actions (Tu et al., 2011)). If the plan is conditional (see for example (Pereira and Han, 2009c; Tu et al., 2007)), each branch is considered a conformant plan and must be prevented.

We shall see several examples in the next sections.

### 3.3.2 Intention Triggering Preferences

Having recognized an intention of another agent, the recognizing agent may either favor or disfavor an abducible (*a priori* preferences), an abductive solution (*a posteriori* preferences) or an evolution (*evolution result a posteriori* preferences) with respect to another, respectively, depending on the setting they are in. If they are in a friendly setting, the one which provides more support to achieve the intention is more favored; in contrast, in a hostile setting, the one providing more support is disfavored. The recognizing agent may also favor the one which takes better advantage of the recognized intention.

To illustrate the usage of intention triggering *a priori* preferences, we revise the Example 3.2.1.

**Example 3.3.1 (Tea-Coffee with Intention Recognition)** *Being thirsty, I consider making tea or coffee. I realize that my roommate, John, also wants to have a drink. To be friendly, I want to take into account his intention when making my choice. This scenario is represented with the following EP program.*

```

1. abds([coffee/0, tea/0]).
2. expect(coffee).    expect(tea).
3. on_observed(drink) <- thirsty.
   drink <- tea.      drink <- coffee.
   <- tea, coffee.
4. expect_not(coffee) <- blood_high_pressure.
5. tea <| coffee <- has_intention(john,tea).
   coffee <| tea <- has_intention(john,coffee).

```

**Figure 3.1:** Tea-Coffee Considering Intentions: A *Priori* Preferences

It is enacted by the preference rules in line 5. The first rule says that *tea* is preferable, a priori, to *coffee* if John intends to drink *tea*; and vice versa, the second rule says that if John intends to drink *coffee*, *coffee* is preferable. Note that the recognition of what John intends is performed by the intention recognition system—which is triggered when a reserved predicate *has\_intention/2* is called.

This scenario also can be encoded using intention triggering *a posteriori* preferences. As a good friend of John, I prefer an abductive solution with a side-effect of John being happy to the one with a side-effect of John being unhappy. This can be coded as in Figure 3.2.

```

unhappy <- coffee, has_intention(john,tea).
happy <- coffee, has_intention(john,coffee).
unhappy <- tea, has_intention(john,coffee).
unhappy <- tea, has_intention(john,tea).
Ai << Aj <- holds_given(happy,Ai),
             holds_given(unhappy,Aj).

```

**Figure 3.2:** Tea-Coffee Considering Intentions: A *Posteriori* Preferences

Despite its simplicity, the example demonstrates how to solve a class of collaborative situations, where one would like to take into account the intentions and the need of others when deriving relevant hypothetical solutions of our current goals.

Next, to illustrate other kinds of preferences, we consider the following revised extended version of the saving city example, presented in (Pereira and Han, 2009b).

**Example 3.3.2 (Saving cities by means of intention recognition)** *During war time, agent David, a general, needs to decide to save a city from his enemy's attack or leave it to keep the military resource, which might be important for some future purpose. David has recognized that a third party is intending to make an attack on the enemy on the next day. David will have a good chance to defeat the enemy if he has enough military resource to coordinate with the third party. The described scenario is coded with the following EP program.*

```

1. abds([save/0, leave/0]).
2. expect(save).          expect(leave).
3. on_observe(choose) <- has_intention(enemy,attack_my_city).
   choose <- save.          choose <- leave.
4. save_men(5000) <- save.      save_men(0) <- leave.
   lose_resource <- save.      save_resource <- leave.
5. Ai << Aj <- holds_given(save_men(Ni), Ai),
   holds_given(save_men(Nj), Aj), Ni > Nj.

6. on_observe(decide) <- decide_strategy.
   decide <- stay_still.
   decide <- counter_attack.
7. good_opportunity <- has_intention(third_party,attack).
   expect(counter_attack) <- good_opportunity, save_resource.
   expect(stay_still).
8. pr(win,0.9) <- counter_attack.
   pr(win,0.01) <- stay_still.
9. Ei <<< Ej <- holds_in_evol(pr(win,Pi), Ei),
   holds_in_evol(pr(win,Pj), Ej), Pi > Pj.
```

**Figure 3.3:** Example: Saving or Leaving

In the first cycle of evolution, there are two abducibles, *save* and *leave*, declared in line 1, to solve the active goal *choose*. The active goal is triggered when David recognizes the intention of the enemy to attack his city (line 3).

Similar to the original version in (Pereira and Han, 2009b), in the case of being a bad general who just sees the situation at hand, David would choose to *save the city* since it would save more people (5000 vs. 0, line 4), i.e. the *a posteriori* preference in line 5 is taken into account immediately, to rule out the case of leaving the city since it would save less people. Then, next day, he would not be able to attack since

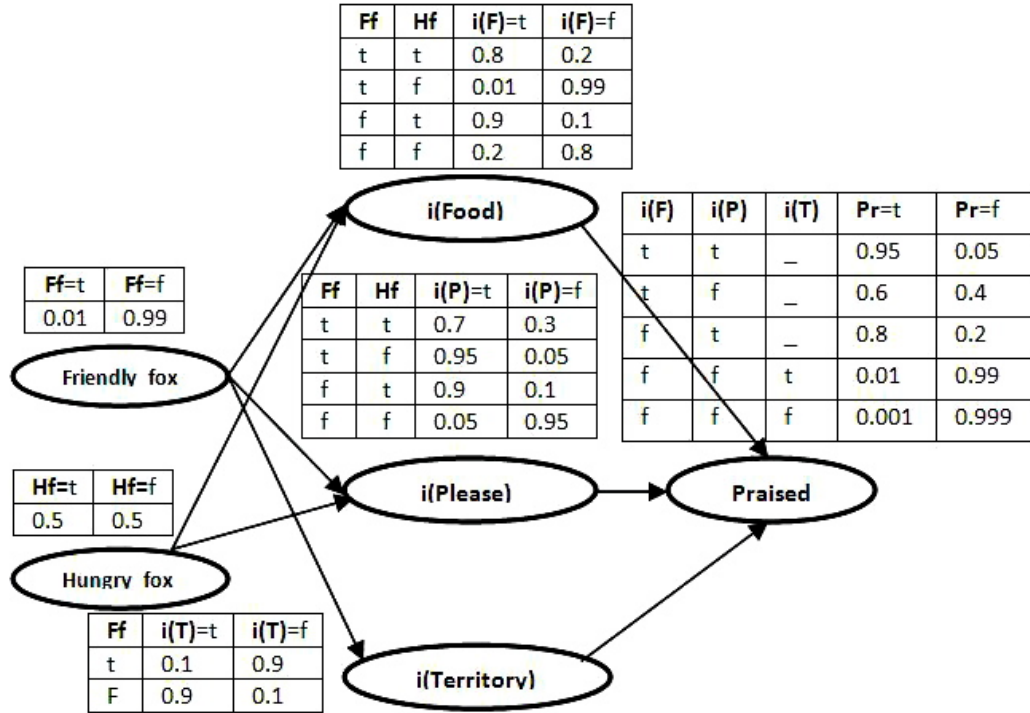


Figure 3.4: Fox's Intentions IRBN

the military resource is not saved (line 7), and that leads to the outcome with very small probability of winning the whole war (line 8).

But, fortunately, being able to look ahead plus to do intention recognition, David can see that on the next day, if he has enough military resources, he has a good opportunity to make a counter-attack on his enemy (line 7), by coordinating with a third party who exhibits the intention to attack the enemy on that day as well; and a successful counter-attack would lead to a very much higher probability of winning the conflict as a whole (line 8). The *evolution result a posteriori* preference is employed in line 9 to prefer the evolution with higher probability of winning the whole conflict.

In this example we can see, in line 7, how a detected intention of another agent can be used to enhance the decision making process. It is achieved by providing an (indirect) trigger for an abducible expectation which affects the *evolution result a posteriori* preference in line 9.

### 3.3.3 Hostile setting

In this hostile setting, having confirmed the intention (and possibly also the plans achieving that intention being carried out by the intending agent), the recognizing agent might act to prevent the intention from being achieved, that is, prevent at least one action of each intention achieving plan from being successfully executed; and, in case of impossibility to doing so, act to minimize losses as much as possible.

**Example 3.3.3 (Fox-Crow)** *Consider Fox-Crow story, adapted from Aesop's fable (Aesop). There is a crow, holding a cheese. A fox, being hungry, approaches the crow and praises her, hoping that the crow will sing and the cheese will fall down near him. Unfortunately for the fox, the crow is very intelligent, having the ability of intention recognition.*

The Bayesian network for recognizing Fox's intentions is depicted in the Figure 3.4. Fox's conceivable intentions (that Crow comes up with) are: Food -  $i(F)$ , Please -  $i(P)$  and Territory -  $i(T)$ . The facts that might give rise to those intentions are how friendly the Fox is (*Friendly\_fox*) and how hungry he is (*Hungry\_fox*). Currently, there is only one observed actions: Fox praised Crow (*Praised*). More details and examples can be found in (Pereira and Han, 2011b).

**Example 3.3.4 (Fox-Crow, cont'd)** *Suppose in Example 3.3.3, the final confirmed Fox's intention is that of getting food (details can be found in (Pereira and Han, 2009c)). That is, the predicate `has_intention(fox,food)` holds. Having recognized Fox's intention, what should Crow do to prevent Fox from achieving it? The following EP program helps Crow with that.*

There are two possible ways so as not to lose the Food to Fox, either simply decline to sing (but thereby missing the pleasure of singing) or hide or eat the cheese before singing.

Line 1 is the declaration of program abducibles (the last two abducibles are for the usage in the second phase, starting from line 9). All of them are always expected (line 2). The counter-expectation rule in line 4 states that an animal is not expected to eat if he is full. The integrity constraints in line 5 say that Crow cannot decline to sing and sing, hide and eat the cheese, at the same time. The *a priori* preference in line 6 states that eating the cheese is always preferred to hiding it (since it may

```
1. abds([decline/0, sing/0, hide/2, eat/2, has_food/0, find_new_food/0]).
2. expect(decline).      expect(sing).
   expect(hide(_,_)).    expect(eat(_,_)).
3. on_observe(not_losing_cheese) <- has_intention(fox, food).
   not_losing_cheese <- decline.
   not_losing_cheese <- hide(crow,cheese), sing.
   not_losing_cheese <- eat(crow,cheese), sing.
4. expect_not(eat(A,cheese)) <- animal(A), full(A).
   animal(crow).
5. <- decline, sing.
   <- hide(crow,cheese), eat(crow,cheese).
6. eat(crow,cheese) <| hide(crow,cheese).
7. no_pleasure <- decline.  has_pleasure <- sing.
8. Ai << Aj <- holds_given(has_pleasure,Ai),
   holds_given(no_pleasure,Aj).

9. on_observe(feed_children) <- hungry(children).
   feed_children <- has_food.  feed_children <- find_new_food.
   <- has_food, find_new_food.
10. expect(has_food) <- decline, not eat(crow,cheese).
    expect(has_food) <- hide(crow,cheese), not stolen(cheese).
    expect(find_new_food).
11. Ei <<< Ej <- hungry(children), holds_in_evol(has_food,Ei),
    holds_in_evol(find_new_food,Ej).
12. Ei <<< Ej <- holds_in_evol(has_pleasure,Ei),
    holds_in_evol(no_pleasure,Ej).
```

be stolen), of course, just in case eating is a possible solution (this is assured in our semantics of *a priori* preferences (Pereira and Han, 2009b)).

Suppose Crow is not full. Then, the counter expectation in line 4 does not hold. Thus, there are two possible abductive solutions:  $[decline]$  and  $[eat(crow,cheese), sing]$  (since the *a priori* preference prevents the choice containing *hiding*).

Next, the *a posteriori* preference in line 8 is taken into account and rules out the abductive solution containing *decline* since it leads to having *no pleasure* which is less preferred to *has pleasure*—the consequence of the second solution that contains *sing* (line 7). In short, the final solution is that Crow eats the cheese then sings, without losing the cheese to Fox and having the pleasure of singing.

Now, let us consider a smarter Crow who is capable of looking further ahead into the future in order to solve longer term goals. Suppose that Crow knows that her children will be hungry later on, in the next stage of evolution (line 9); eating the cheese right now would make her have to find new food for the hungry children. Finding new food may take long, and is always less favourable than having food ready to feed them right away (*evolution result a posteriori* preference in line 11). Crow can see three possible evolutions:  $[[decline], [has\_food]]$ ;  $[[hide(crow, cheese), sing], [has\_food]]$  and  $[[eat(crow, cheese), sing], [find\_new\_food]]$ . Note that in looking ahead at least two steps into the future, local preferences are taken into account only after all evolution-level ones have been applied (a full discussion can be found in (Pereira and Han, 2009b)).

Now the two *evolution result a posteriori* preferences in lines 11-12 are taken into account. The first one rules out the evolution including *finding new food* since it is less preferred than the other two which includes *has\\_food*. The second one rules out the one including *decline*. In short, Crow will hide the food to keep for her hungry children, and still take pleasure from singing.

### 3.4 Ambient Intelligence in the Home Environment and Elder Care

Ambient Intelligence (AmI) is the vision of a future in which environments support people inhabiting in them. The envisaged environment is unobtrusive, interconnected, adaptable, dynamic, embedded and intelligent. It should be sensitive to the needs of inhabitants, and capable of anticipating their needs and behaviour. It

should be aware of their personal requirements and preferences, and interact with people in a user-friendly way (Sadri, 2011a).

From this view of AmI, we can see a number of issues where individual intention recognition techniques can step in, providing help and enabling improvement. For example, in order to provide appropriate support, the environment should be able to proactively recognize the inhabitants' intention—to see whether they need help to accomplish what they intend to do—or to warn them (or their carers) in case they intend something dangerous. Undoubtedly, an ability to recognize intentions of assisted people as well as other relevant concerns such as intruders, would enable to deal with a combination of several issues, e.g. pro-activeness (either agonistic or antagonistic), security, emergency, etc. in a more integrated and timely manner (Roy et al., 2007; Han and Pereira, 2010c).

One of the key issues of Ambient Intelligence (AmI), which has not been well studied yet, and reported as an ongoing challenge (Cook et al., 2009), is that AmI systems need to be aware of users' preferences, intentions and needs. Undoubtedly, respecting users' preferences and needs in decision making processes would increase their degree of acceptance with respect to the systems, making them more friendly and thoughtful. Furthermore, an ability to recognize intentions of assisted people as well as other relevant concerns such as intruders, would enable to deal with a combination of several issues—including pro-activeness, security, and emergency—in a more integrated and timely manner. We shall discuss these issues in the sequel.

### 3.4.1 Proactive Support

A important feature of AmI, particularly desirable in the Elder Care domain, is that the assisting system should take initiative to help the people it assists. To this end, the system must be capable of recognizing their intentions on the basis of their observable actions, then provide suggestions or help achieve the recognized intentions (Pereira and Han, 2011a). A suggestion can be, for example, what are the appropriate kinds of drink for the elder, considering the current time, temperature, or even future scheduled events such as going to have a medical test on the next day, upon having recognized that he has an intention to drink something. Or, a suggestion can simply be telling the elder where he put his book yesterday, having recognized that he might be looking for it. This feature is especially desirable and important when the assisted people are elderly or individuals with disabilities or

suffering from mental difficulties (Roy et al., 2007). The need for technology in this area is obvious looking at the fact that in the last twenty years there has been a significant increase of the average age of the population in most western countries and the number of elderly people has been and will be constantly growing (Cook et al., 2009; Haigh et al., 2004; Cesta and Pecora, 2004; Geib and Goldman, 2001; Geib, 2002; Roy et al., 2007; Sadri, 2011a).

The EP system can be engaged to provide appropriate suggestions for the elders, taking into account the external environment, elders' preferences and already scheduled future events. Expectation rules and *a priori* preferences cater for the physical state (health reports) information of the elders to guarantee that only contextually safe healthy choices are generated; then, information such as the elders' pleasure, interests, etc. are then considered by *a posteriori* preferences and the like.

In the Elder Care domain, assisting systems should be able to provide contextually appropriate suggestions for the elders based on their recognized intentions. The assisting system is supposed to be better aware of the environment, the elders' physical states, mental states as well as their scheduled events, so that it can provide good and safe suggestions, or simply warnings. We continue with the Elder Care example from a previous chapter for illustration.

**Example 3.4.1 (Elder Care, cont'd)** *Suppose in the Elder care examples in Chapter 2 (see Figure 2.2), the final confirmed intention is that of looking for a drink. The possibilities are: natural pure water, tea, coffee and juice. The EP system now is employed to help the elder with choosing an appropriate drink. The scenario is coded with the EP program in Figure 3.5 below.*

*The elder's physical states are utilized in a priori preferences and expectation rules to guarantee that just choices which are contextually safe for the elder are generated. Only after that other aspects, for example the elder's pleasure with respect to each kind of drink, are taken into account, with the a posteriori preferences.*

The information regarding the environment (current time, current temperature) and the physical states of the elder is coded in the Prolog part of the program (lines 9-11). The assisting system is supposed to be aware of this information in order to provide good suggestions.

Line 1 is the declaration of program abducibles: *water, coffee, tea, and juice*. All of them are always expected (line 2). Line 3 picks up a recognized intention verified by the intention recognition module. The counter-expectation rules in line 4 state

```

1. abds([water/0, coffee/0, tea/0, juice/0,
        precise_result/0, imprecise_result/0]).
2. expect(coffee). expect(tea). expect(water). expect(juice).
3. on_observe(drink) <- has_intention(elder,drink).
   drink <- tea. drink <- coffee. drink <- water. drink <- juice.
4. expect_not(coffee) <- prolog(blood_high_pressure).
   expect_not(coffee) <- prolog(sleep_difficulty).
   expect_not(coffee) <- prolog(late).
   expect_not(juice) <- prolog(late).
5. <- tea, coffee. <- coffee, juice.
   <- tea, juice. <- tea, water.
6. coffee <| tea <- prolog(morning_time).
   coffee <| water <- prolog(morning_time).
   coffee <| juice <- prolog(morning_time).
7. juice <| coffee <- prolog(hot). juice <| tea <- prolog(hot).
   juice <| water <- prolog(hot). water <| coffee <- prolog(hot).
   water <| tea <- prolog(hot).
8. tea <| coffee <- prolog(cold). tea <| juice <- prolog(cold).
   tea <| water <- prolog(cold).
9. pleasure_level(3) <- coffee. pleasure_level(2) <- tea.
   pleasure_level(1) <- juice. pleasure_level(0) <- water.
10. sugar_level(1) <- coffee. sugar_level(1) <- tea.
    sugar_level(5) <- juice. sugar_level(0) <- water.
11. caffein_level(5) <- coffee. caffein_level(0) <- tea.
    caffein_level(0) <- juice. caffein_level(0) <- water.
12. Ai << Aj <- holds_given(pleasure_level(V1), Ai),
    holds_given(pleasure_level(V2), Aj), V1 > V2.

13. on_observe(health_check) <- time_for_health_check.
    health_check <- precise_result.
    health_check <- imprecise_result.
14. expect(precise_result) <- no_high_sugar, no_high_caffein.
    expect(imprecise_result).
    no_high_sugar <- sugar_level(L), prolog(L < 2).
    no_high_caffein <- caffein_level(L), prolog(L < 2).
15. Ei <<< Ej <- holds_in_evol(precise_result, Ei),
    holds_in_evol(imprecise_result, Ej).

beginProlog.
 :- assert(scheduled_events(1, [has_intention(elder,drink)])),
    assert(scheduled_events(2, [time_for_health_check])).
late :- time(T), (T > 23; T < 5).
morning_time :- time(T), T > 7, T < 10.
hot :- temperature(TM), TM > 32.
cold :- temperature(TM), TM < 10.
blood_high_pressure :- physical_state(blood_high_pressure).
sleep_difficulty :- physical_state(sleep_difficulty).
endProlog.

```

Figure 3.5: Elder Care: Suggestion for a Drink

that *coffee* is not expected if the elder has high blood pressure, experiences difficulty to sleep or it is late; and juice is not expected if it is late. Note that the reserved predicate *prolog/1* is used to allow embedding Prolog code in an EP program. More details can be found in (Pereira and Han, 2009a,b). The integrity constraints in line 5 say that is is not allowed to have at the same time the following pairs of drink: tea and coffee, tea and juice, coffee and juice, and tea and water. However, it is the case that the elder can have coffee or juice together with water at the same time.

The *a priori* preferences in line 6 say in the morning coffee is preferred to tea, water and juice. And if it is hot, juice is preferred to all other kinds of drink and water is preferred to tea and coffee (line 7). In addition, the *a priori* preferences in line 8 state if the weather is cold, tea is the most favorable, i.e. preferred to all other kinds of drink.

Now let us look at the suggestions provided by the Elder Care assisting system modelled by this EP program, considering some cases:

1. time(24) (*late*); temperature(16) (*not hot, not cold*); no high blood pressure; no sleep difficulty: there are two a priori abductive solutions: [*tea*], [*water*]. Final solution: [*tea*] (since it has greater level of pleasure than water, which is ruled out by the *a posteriori* preference in line 12).
2. time(8) (*morning time*); temperature(16) (*not hot, not cold*); no high blood pressure; no sleep difficulty: there are two abductive solutions: [*coffee*], [*coffee, water*]. Final: [*coffee*], [*coffee, water*].
3. time(18) (*not late, not morning time*); temperature(16) (*not cold, not hot*); no high blood pressure; no sleep difficulty: there are six abductive solutions: [*coffee*], [*coffee, water*], [*juice*], [*juice, water*], [*tea*], and [*water*]. Final: [*coffee*], [*coffee, water*].
4. time(18) (*not late, not morning time*); temperature(16) (*not cold, not hot*); high blood pressure; no sleep difficulty: there are four abductive solutions: [*juice*], [*juice, water*], [*tea*], and [*water*]. Final: [*tea*].
5. time(18) (*not late, not morning time*); temperature(16) (*not cold, not hot*); no high blood pressure; sleep difficulty: there are four abductive solutions: [*juice*], [*juice, water*], [*tea*], and [*water*]. Final: [*tea*].

6. time(18) (*not late, not morning time*); temperature(8) (*cold*); no high blood pressure; no sleep difficulty: there is only one abductive solution: *[tea]*.
7. time(18) (*not late, not morning time*); temperature(35) (*hot*); no high blood pressure; no sleep difficulty: there are two abductive solutions: *[juice]*, *[juice,water]*. Final: *[juice]*, *[juice,water]*.

If the *evolution result a posteriori preference* in line 15 is taken into account and the elder is scheduled to go to the hospital for health check in the second day: the first and the second cases do not change. In the third case: the suggestions are *[tea]* and *[water]* since the ones that have *coffee* or *juice* would cause high caffeine and sugar levels, respectively, which can make the checking result (health) imprecise (lines 13-15). Similarly for all the other cases.

Note future events can be asserted as Prolog code using the reserved predicate *scheduled\_events/2*. For more details of its use see (Pereira and Han, 2009a,b).

As one can gather, the suggestions provided by this assisting system are quite contextually appropriate. We might elaborate current factors (time, temperature, physical states) and even consider more factors to provide more appropriate suggestions if the situation ever gets more complicated.

### 3.4.2 Security and Emergency

**Security in AmI.** Security is one of the key issues for AmI success (Friedewald et al., 2007), and particularly important in home environments (Friedewald et al., 2005). It comprises two important categories : security in terms of Burglary Alarm systems and security in terms of health and well-being of the residents (prevention, monitoring) (Friedewald et al., 2005).

So far Burglary Alarm technology has been mainly based on sensing and recognizing the very last action of an intrusion plan, such as “breaking the door” (Friedewald et al., 2005; Wikipedia). However, it may be too late to provide an appropriate protection. Burglary Alarm systems need to be able to guess in advance the possibility of an intrusion on the basis of the very first observable actions of potential intruders. For example, it would be useful to find out how likely a stranger constantly staring at your house has an intrusion intention, taking into account the particular situation, e.g. if he has weapon or if it is night time. This information can be sent to the carer (or the elders if there is no carer available) to get prepared

(e.g. turn on the light or sounders to scare off burglars or call relatives, police, etc.). Our incremental intention recognition system (Chapter 2) appears to be appropriate to deal with this scenario. Given any currently observed actions, one can compute the probability of the on-going conceivable intentions, and if the one of the intrusion intention is large enough or is among (some of) the most likely intentions, the carer should be informed of a potential intrusion. To be more certain about the possibility of an intrusion, more observations may need to be made, but at least now the carer is ready to handle any potentially negative forthcoming situations. Waiting until being sure to get ready can be too late to take appropriate actions. The EP system can thus be used to provide suggestions on the appropriate course of actions to take. For illustration let us consider the following example.

**Example 3.4.2 (Solving Intrusion)** *Consider a situation where the intention recognition system recognized an intention of intrusion at night. The system must either warn the elders who are sleeping, automatically call the nearest police, or activate the embedded burglary alarm. If the elders are sleeping and ill, they do not expect to be warned, but prefer other solutions. Due to potential disturbance, the elders prefer simply activating the burglary system to calling the police as long as no weapon is detected and there is a single intruder.*

The situation is described by the program with three abducibles: *call\_police*, *warn\_persons*, *activate\_alarm*, and can be coded in EP as follows

1.  $on\_observe(solve\_intrusion) \leftarrow at\_night, has\_intention(stranger, intrusion).$
2.  $solve\_intrusion \leftarrow call\_police.$   
 $solve\_intrusion \leftarrow warn\_persons.$   
 $solve\_intrusion \leftarrow activate\_alarm.$
3.  $expect(call\_police). \quad expect(warn\_persons). \quad expect(activate\_alarm).$
4.  $expect\_not(warn\_persons) \leftarrow ill, sleeping.$
5.  $activate\_alarms \triangleleft call\_police \leftarrow no\_weapon\_detected, individual.$
6.  $call\_police \triangleleft activate\_alarms \leftarrow weapon\_detected.$

Suppose it is night-time and an intrusion intention is recognized, then the active goal *solve\_intrusion* (line 1) is triggered, and the EP system starts reasoning to find the most appropriate solutions.

This program has three abductive solutions: [*call\_police*], [*warn\_persons*], and [*activate\_alarm*] since all the abducibles are expected and there is no expectation to

their contrary. Note that for each abducible a *consider*-rule is added automatically (Pereira and Han, 2009b). Suppose it detects that the elders are sleeping and known to be ill, i.e. literals *ill* and *sleeping* are factual. In this case, the elders do not expect to be warned (line 4), thus ruling out the second solution [*warn\_persons*]. And if no weapon is detected and there is only a single intruder, the *a priori* preference in line 5 is triggered, which defeats the solution where only *call\_police* is present (due to the impossibility of simultaneously abducting *activate\_alarm*). Hence, the only solution is to activate the burglary alarm. However, if weapons were detected, the preference in line 6 is triggered and defeats the [*activate\_alarm*] solution. The only solution left is to call the police (*call\_police*).

Regarding Burglary Alarm systems, in the following example we consider a simple scenario of recognizing an elder's intentions.

**Example 3.4.3 (Elder Intentions)** *An elder stays alone in his apartment. One day the Burglary Alarm is ringing, and the assisting system observes that the elder is looking for something. In order to assist him, the system needs to figure out what he intends to find. Possible things are: Alarm button (AlarmB); Contact Device (ContDev); Defensible Weapons (Weapon); and light switch (Switch). The IRBN representing this scenario is in Figure 3.6.*

The nodes representing the conceivable intentions are:  $i(\text{AlarmB})$ ,  $i(\text{ContDev})$ ,  $i(\text{Weapon})$ , and  $i(\text{Switch})$ . The Bayesian network for intention recognition (IRBN) has three top nodes in the pre-intentional level representing the causes/reasons of the intentions, which are *Alarm\_On*, *Defensible* and *Light\_on*. The first and last nodes are evidence nodes, i.e. their values are observable. There is only one observable action, represented by the node *Looking* in the last layer. It is a direct child of the intention nodes. The conditional probability tables (CPD) of each nodes in the BN are given. For example, the table of the node *Defensible* says that the elder is able to defense himself (with weapons) with probability of 0.3 and not able to do so with probability 0.7. The table in the top-right corner provides the probability of the elder looking around for something conditional on the intentions. Based on this IRBN one can now compute the conditional probability of each intention given the observed action.

Another security issue concerns health and well-being of the residents. AmI systems need to be able to prevent hazardous situations, which usually come from dangerous ideas or intentions (e.g. take a bath when drunk, drink alcohol while

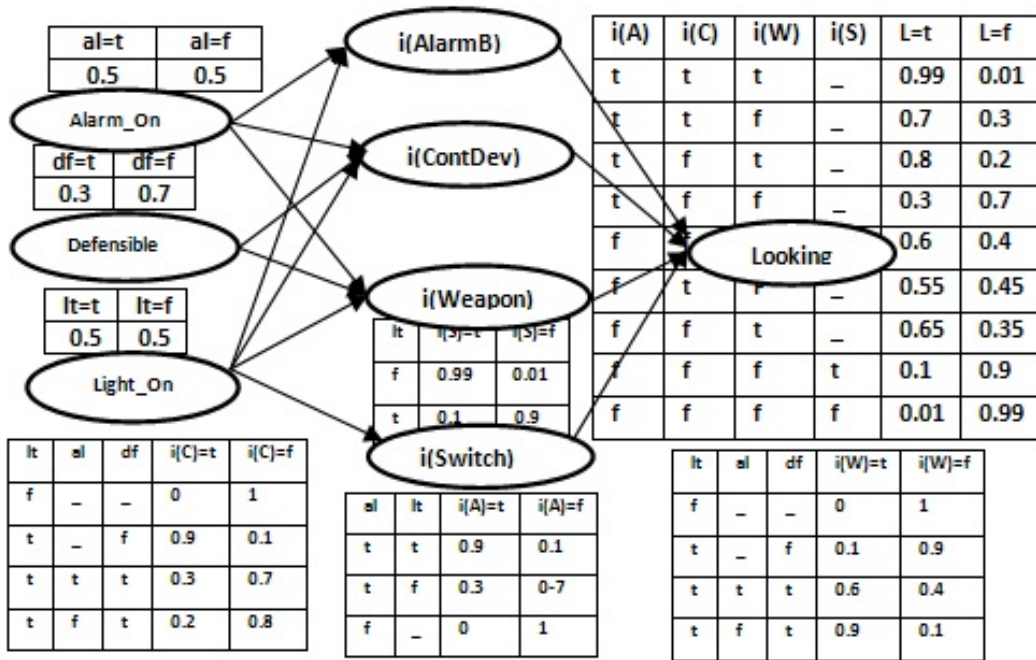


Figure 3.6: Bayesian network for recognizing the elder's intentions

not permitted, or even commit suicide) of the assisted persons, especially those with mental impairments (Roy et al., 2007). To this end, guessing their intentions from the very first relevant actions is indispensable to take timely actions. In our incremental intention recognition method, an IRBN will be built to compute how likely there is a dangerous intention, with respect to any currently observed actions, and carers would be informed in case it is likely enough, in order to get prepared in time.

**Emergency in AmI.** Handling emergency situations is another important issue in AmI. There is a wide range of emergency situations, e.g. in security, when recognizing intrusion intention of a stranger or dangerous intentions of the assisted person. They also can occur when detecting fire, unconsciousness or unusualness in regular activities (e.g. sleep for too long), etc. Emergency handling in the EP system can be done by having an active goal rule for whichever emergency situation. For solving the goal, a list of possible actions, all represented by abducible enablers, are available to form solutions. Then, users' preferences are encoded using all kinds of

preference of EP: *a priori* ones for preferring amongst available actions, *a posteriori* ones for comparing solutions taking into account their consequences and utility, and *a posteriori evolution result* ones for comparing more-than-one-step consequences. Moreover, the expectation and counter expectations rules are used to encode pros and cons of the users towards each available action, or towards any abducible in general.

### 3.4.3 Discussion of Other AmI Issues

We have shown how our intention-based decision making framework can enable the provision of proactive support for assisted people, and the tackling of the AmI security and emergency issues. We now briefly sketch how it can be utilized to address yet other important issues in AmI.

First of all, it is known that intention recognition plays a central role in human communication (Pinker et al., 2008; Heinze, 2003; Tomasello, 2008). In addition, an important aspect of intentions is future-directedness, i.e. if we intend something now, it means we intend to execute a course of actions to achieve it in the future (Bratman, 1987; Singh, 1991; Roy, 2009b). Most actions may be executed only at a far distance in time. Thus, we usually need to guess others' intentions from the very first clues, such as their actions or spoken sentences, in order to secure a smooth conversation or collaboration. Perhaps we guess a wrong intention, but we need to be able to react in a timely manner; and that is also part of the conversation. We can simply attempt to confirm by asking, e.g. "is this (...) what you mean?". Our intention-based decision making framework can be used to design better and more friendly human-computer interaction devices that can react to human behavior and speech, communicate with them to confirm their intentions so as to provide appropriate help when necessary, after having guessed their likely intentions using an intention recognition system.

Another issue is that, in order to be highly accepted by the users, an assistive system should be able to proffer explanations for the suggestions it provides. In EP, that can be easily done by keeping all the preferences, integrity constraints, expectation and counter expectation rules that were used both to consider and to rule out abductive solutions.

## 3.5 Conclusions and Future Work

We have summarized our previous work on Evolution Prospection (EP) (Han, 2009; Pereira and Han, 2009b) and shown its coherent combination with the intention recognition system described in Chapter 2, for intention-based decision making. The EP system has proven useful for decision making (Han, 2009; Pereira and Han, 2009b; Han et al., 2012d; Pereira and Han, 2011a), and has now been empowered to take into account the intentions of other agents—an important aspect that has not been well explored so far. The fact that both systems are Logic Programming based enabled their easy integration. We have described and exemplified several ways in which an EP agent can benefit from having an ability to recognize intentions of other agents.

Nonetheless, the approach we used here is not restricted to Logic Programming based systems. In general, any intention recognition system, and indeed, any decision making system, can be considered. The ideas of integration described here can be adopted in other decision making systems to account for intentions.

We have addressed the need of intention-based decision making in different application domains, including Ambient Intelligence (Han and Pereira, 2010c; Sadri, 2011a) and Elder Care (Pereira and Han, 2011a; Han and Pereira, 2010b), where decision making techniques as well as intention recognition abilities are of increasing importance (Sadri, 2011b; Geib, 2002; Sadri, 2011a; Han and Pereira, 2010c). However, the scalability of our system remains to be studied.

In future work, we consider to apply our combined system to other application domains, including story understanding (Charniak and Goldman, 1990), human-computer and interface-agents systems (Lesh, 1998; Hong, 2001; Armentano and Amandi, 2007), traffic monitoring (Pynadath and Wellman, 1995), assistive living (Geib, 2002; Haigh et al., 2004; Tahboub, 2006; Roy et al., 2007; Pereira and Han, 2011a), military settings (Mao and Gratch, 2004; Heinze, 2003), and moral reasoning (Han et al., 2012d), where intention recognition has proven useful and of great practicality. Another area of future development is to extend our system to enable collective or group intention recognition (Sukthankar and Sycara, 2008; Sukthankar, 2007) in a decision making process. In this regard, we have made some initial attempt in the Elder Care domain (Han and Pereira, 2010c,b).



## Part II

# Intention Recognition, Commitment, and the Evolution of Cooperation



---

The main concern of Part I was to develop efficient intention recognition methods, and application of intention recognition for different AI problems. In Part II, based on these methods, we study the role of intention recognition in the evolution of cooperation. In addition, we study how a simple commitment strategy, in combination with intention recognition, can lead to improvement in the evolution of cooperation.

---

*(...) because intention-reading is so critical for human social functioning and the development of key human abilities, such as language and culture, it is reasonable to assume that it has been shaped by natural selection.*

Woodward et al. (2009).

*(...) The capacity for using commitment strategies effectively is so important that natural selection may have shaped specialized signaling capacities to make this possible.*

Nesse (2001b).

The problem of intention recognition has been paid much attention in AI, Philosophy and Psychology for several decades (Kautz and Allen, 1986a; Charniak and Goldman, 1993; Bratman, 1987, 1999; Bui, 2003; Heinze, 2003; Armentano and Amandi, 2007; Geib and Goldman, 2009). However, it has mostly been restricted to small scale interactive settings, focusing on how to efficiently and appropriately recognize intentions of an individual or a group of agents. There is a significant shortage of

---

modelling research with respect to large scale social contexts, namely to address the evolutionary roles and aspects of intention recognition.

Intention recognition can be found abundantly in many kinds of interactions and communications, not only in Human but also many other species (Tomasello, 1999; Woodward et al., 2009; Meltzoff, 2007; Tomasello, 2008; Cheney and Seyfarth, 2007). The knowledge about intention of others in a situation could enable to plan in advance, either to secure a successful cooperation or to deal with potential hostile behaviours (van Hees and Roy, 2008; Roy, 2009a; Pereira and Han, 2011b; Han and Pereira, 2011c). Given the advantage of knowing the intentions of others and the abundance of intention recognition among different species, intention recognition undoubtedly should be taken into account when studying or modeling collective behavior. This issue becomes even more relevant when the achievement of a goal by an individual does not depend uniquely on its own actions, but also on the decisions and actions of others, namely when individuals cooperate or have to coordinate their actions to achieve a task—especially when the possibility of communication is limited (Heinze, 2003; Kraus, 1997; Pinker et al., 2008; Van Segbroeck et al., 2010; Tomasello, 2008). Additionally, in more realistic settings where deceit may offer additional profits, individuals often attempt to hide their real intentions and make others believe in faked ones (Robson, 1990; Tomasello, 2008; Skyrms, 2010; Pereira and Han, 2011b; Santos et al., 2011; Trivers, 2011). Hence, undoubtedly, a capability of intention recognition would confer on its holder great evolutionary benefits.

In this part of the thesis we study the role of intention recognition for one of the most challenging but intriguing issues, that traverses areas as diverse as Biology, Economics, Artificial Intelligence, Political Science, and Psychology: the problem of *evolution of cooperation* (Hardin, 1968; Axelrod, 1984; West et al., 2007; Nowak, 2006b,a; Sigmund, 2010). In its simplest form, a cooperative act is metaphorically described as the act of paying a cost to convey a benefit to someone else. If two players simultaneously decide to cooperate or not, the best possible response will be to try to receive the benefit without paying the cost. In an evolutionary setting, we may also wonder why would natural selection equip selfish individuals with altruistic tendencies while it incites competition between individuals and thus apparently rewards only selfish behavior? Several mechanisms responsible for promoting cooperative behavior have been recently identified (Sigmund, 2010; Nowak, 2006b). From kin and group ties (Hamilton, 1964; West et al., 2002; Traulsen and Nowak,

---

2006), to distinct forms of reciprocity (Nowak and Sigmund, 1992; Imhof et al., 2005; Trivers, 1971; Pacheco et al., 2006; Nowak and Sigmund, 2005), networked populations (Santos and Pacheco, 2005; Santos et al., 2006a; Szabó and Fáth, 2007; Santos et al., 2008; Lindgren and Nordahl, 1994), and different forms of punishment (Hauert et al., 2007; Fehr and Gächter, 2002; Dreber et al., 2008; Ohtsuki et al., 2009; Rand and Nowak, 2011), several aspects have been shown to play an important role in the emergence of cooperation. Differently from these, here we shall describe how cooperation may emerge from the interplay between population dynamics and individuals' cognitive abilities, namely the basic and important ability to perform intention recognition.

Additionally, as a natural extension and supplement to intention recognition, an ability to arrange commitment might enable to clarify intentions of others when it is difficult to perform intention recognition with sufficient precision and/or efficacy (see thesis introduction for a brief discussion, Section 1.1.3). Hence, in this thesis we shall study how a simple form of commitment provides a pathway to the evolution of cooperative behavior, and furthermore, how it supplements intention recognition, thereby leading to a better combined strategy.

Our study is carried out within the framework of Evolutionary Game Theory (EGT) (Maynard-Smith, 1982; Hofbauer and Sigmund, 1998; Nowak, 2006a; Sigmund et al., 2010). Here, individuals' success or fitness is expressed in terms of the outcome of a two-player game, which, in turn, is used by the individuals to copy others whenever these appear to be more successful. Comparative accumulated pay-offs are used to update the population: more successful individuals produce more offspring, which inherit their strategy. Equivalently, the same process can be seen as if, instead of inheriting strategies, new individuals adapt by copying strategies from acquaintances that did better. Overall, this type of dynamics can be conveniently described as an ordinary differential equation – the replicator equation (Hofbauer and Sigmund, 1998)–, which nicely describes any simple evolutionary process. This framework is however more general one could initially foresee, as the ensuing dynamics may be also shown to be equivalent to finite-action learning automata (Börgers et al., 1997; Van Segbroeck et al., 2010), in which agents revise their strategies by means of incipient reinforcement learning techniques (Narendra and Thathachar, 1989).

We start this second part of the thesis by describing the EGT methods based on

---

which the models in the next four chapters will be built and analyzed.

**Interaction between Individuals** Interactions are modeled as symmetric two-player games defined by the payoff matrix

$$\begin{array}{cc} & C & D \\ C & (R, R) & (S, T) \\ D & (T, S) & (P, P) \end{array}.$$

A player who chooses to cooperate (C) with someone who defects (D) receives the sucker’s payoff  $S$ , whereas the defecting player gains the temptation to defect,  $T$ . Mutual cooperation (resp., defection) yields the reward  $R$  (resp., punishment  $P$ ) for both players. Depending on the ordering of these four payoffs, different social dilemmas arise (Macy and Flache, 2002; Santos et al., 2006a; Sigmund, 2010). In this thesis we are only concerned with the Prisoner’s Dilemma (PD), which is characterized by the ordering  $T > R > P > S$ .

In a single round of the PD, it is always best to defect, but cooperation may be rewarded if the game is iterated. In the iterated PD, it is also required that mutual cooperation is preferred over an equal probability of unilateral cooperation and defection ( $2R > T + S$ ); otherwise alternating between cooperation and defection would lead to a higher payoff than mutual cooperation. For convenience and a clear representation of results, we sometimes use the Donation game (Sigmund, 2010)—a famous special case of the PD—where  $T = b$ ,  $R = b - c$ ,  $P = 0$ ,  $S = -c$ , satisfying that  $b > c > 0$ , where  $b$  and  $c$  stand respectively for “benefit” and “cost” (of cooperation).

In a population of  $N$  individuals interacting via an iterated PD, whenever two specific strategies are present in the population, say **A** and **B**, the fitness of an individual with a strategy **A** (resp., **B**) in a population with  $k$  **A**s and  $(N - k)$  **B**s can be written as

$$\begin{aligned} \Pi_A(k) &= \frac{1}{r(N-1)} \sum_{j=1}^r [(k-1)\pi_{A,A}(j) + (N-k)\pi_{A,B}(j)], \\ \Pi_B(k) &= \frac{1}{r(N-1)} \sum_{j=1}^r [k\pi_{B,A}(j) + (N-k-1)\pi_{B,B}(j)], \end{aligned} \tag{3.1}$$

---

where  $\pi_{X,Y}(j)$ , with  $X, Y \in \{A, B\}$ , stands for the payoff obtained from a round  $j$  as a result of their mutual behavior of an  $X$  strategist in an interaction with a  $Y$  strategist (as specified by the payoff matrix above), and  $r$  is the total number of rounds of the iterated PD. As usual, instead of considering a fixed number of rounds, upon completion of each round, there is a probability  $w$  that yet another round of the game will take place, resulting in an average number of  $\langle r \rangle = (1-w)^{-1}$  rounds per interaction (Hofbauer and Sigmund, 1998; Sigmund, 2010).

The iterated PD will be used as the interacting environment in Chapters 4 and 5, while the non-iterated or one-shot PD will be used in Chapters 6 and 7. For the one-shot PD, we can simply substitute  $r = 1$  in Eq. (3.1). In all analyses and simulations, values of  $\Pi$  will be derived analytically, and when this is not possible, we shall use numerical simulations.

We now recall some of the most important strategies of the iterated PD, needed for the studies in Chapters 4 and 5.

**Strategies in Iterated Prisoner’s Dilemma** The iterated PD is usually known as a story of tit-for-tat (TFT), which won both Axelrod’s tournaments (Axelrod, 1984; Axelrod and Hamilton, 1981). *TFT* starts by cooperating, and does whatever the opponent did in the previous round. It will cooperate if the opponent cooperated, and will defect if the opponent defected. But if there are erroneous moves because of noise (i.e. an intended move is wrongly performed with a given execution error, referred here as “noise”), the performance of *TFT* declines, in two ways: (i) it cannot correct errors (e.g., when two *TFT*s playing with one another, an erroneous defection by one player leads to a sequence of unilateral cooperation and defection) and (ii) a population of *TFT* players is undermined by random drift when the pure cooperator *AllC* mutants appear (which allows exploiters to grow). Tit-for-tat is then replaced by generous tit-for-tat (GTFT), a strategy that cooperates if the opponent cooperated in the previous round, but sometimes cooperates even if the opponent defected (with a fixed probability  $p > 0$ ) (Nowak and Sigmund, 1992). *GTFT* can correct mistakes, but remains suffering the random drift.

Subsequently, *TFT* and *GTFT* were replaced by win-stay-lose-shift (WSLS) as the winning strategy chosen by evolution (Nowak and Sigmund, 1993). *WSLS* starts by cooperating, and repeats the previous move whenever it did well, but changes

---

otherwise. *WLSL* corrects mistakes better than *GTFT* and does not suffer random drift. However, it is severely exploited by the pure defector *AllD* players.

Intention recognition strategies implemented in Chapters 4 and 5 will be compared to these well-known strategies of the iterated PD.

**Evolutionary Dynamics** In the framework of Evolutionary Game Theory (EGT), the accumulated payoff from all interactions (see Eq. (3.1)) emulates individuals' *fitness* or social *success* and the most successful individuals will tend to be imitated by others, implementing a simple form of social learning (Szabó and Toke, 1998; Traulsen et al., 2006; Sigmund, 2010; Rendell et al., 2010). Any player can change its strategy by adopting another player's strategy with a probability defined by the Fermi probability distribution below. If a strategy has a higher (average) payoff or fitness than another, it tends to be imitated more by the other.

A strategy update event is defined in the following way, corresponding to the so-called pairwise comparison (Szabó and Toke, 1998; Traulsen et al., 2006). At each time-step, one individual  $i$  with a fitness  $f_i$  is randomly chosen for behavioral revision. Individual  $i$  will adopt the strategy of a randomly chosen individual  $j$  with fitness  $f_j$  with a probability given by the Fermi function (from statistical physics)

$$p(f_i, f_j) = \left(1 + e^{-\beta(f_j - f_i)}\right)^{-1},$$

where the quantity  $\beta$ , which in physics corresponds to an inverse temperature, controls the 'intensity of selection' or 'imitation strength'. When  $\beta = 0$  we obtain the limit of neutral drift, and with the increasing of  $\beta$  one enhances the role played by the game payoff in the individual fitness, and behavioral evolution (Traulsen et al., 2006, 2007).

In the absence of mutations, the end states of evolution are inevitably monomorphic, as a result of the stochastic nature of the evolutionary dynamics and update rule. As we are interested in a global analysis of the population dynamics with multiple strategies, we further assume that with a small probability  $\mu$  individuals switch to a randomly chosen strategy, freely exploring the space of possible behaviors. By introducing a small probability of mutation or exploration, the eventual appearance of a single mutant in a monomorphic population, this mutant will fixate or will become extinct long before the occurrence of another mutation and, for this reason, the population will spend all of its time with a maximum of two strategies present

---

simultaneously (Fudenberg and Imhof, 2005; Imhof et al., 2005; Hauert et al., 2007; Traulsen et al., 2009; Santos et al., 2011). This allows one to describe the evolutionary dynamics of our population in terms of a reduced Markov Chain of a size equal to the number of different strategies, where each state represents a possible monomorphic end-state of the population associated with a given strategy, and the transitions between states are defined by the fixation probabilities of a single mutant of one strategy in a population of individuals who adopt another strategy. The resulting stationary distribution characterizes the average time the population spends in each of these monomorphic states, and can be computed analytically (Karlin and Taylor, 1975; Fudenberg and Imhof, 2005; Imhof et al., 2005; Hauert et al., 2007; Santos et al., 2011) (see below).

In the presence of two strategies the payoffs of each are given by Eq. (3.1), whereas the probability to change the number  $k$  of individuals with a strategy **A** (by  $\pm$  one in each time step) in a population of  $(N - k)$  **B**-strategists is

$$T^\pm(k) = \frac{N - k}{N} \frac{k}{N} \left[ 1 + e^{\mp\beta[\Pi_A(k) - \Pi_B(k)]} \right]^{-1}.$$

The fixation probability of a single mutant with a strategy **A** in a population of  $(N - 1)$  **B**s is given by (Traulsen et al., 2006)

$$\rho_{B,A} = \left( \sum_{i=0}^{N-1} \prod_{j=1}^i \lambda_j \right)^{-1} \quad (3.2)$$

where  $\lambda_j = T^-(j)/T^+(j)$ .

In the limit of neutral selection ( $\beta = 0$ ),  $\lambda_j = 1$ . Thus,  $\rho_{B,A} = 1/N$ . Considering a set  $\{1, \dots, n_S\}$  of different strategies, these fixation probabilities determine a transition matrix  $[T_{ij}]_{i,j=1,\dots,n_S}$ , with  $T_{ii} = 1 - \sum_{k=1, k \neq i}^{n_S} \rho_{k,i}/(n_S - 1)$  and  $T_{ij, j \neq i} = \rho_{ji}/(n_S - 1)$ , of a Markov Chain. The normalized eigenvector associated with the eigenvalue 1 of the transposed of  $M$  provides the stationary distribution described above (Karlin and Taylor, 1975; Fudenberg and Imhof, 2005; Imhof et al., 2005; Hauert et al., 2007; Santos et al., 2011), describing the relative time the population spends adopting each of the strategies.

---

**Structure of the remainder of Part II** The remainder of this part is organized as follows. In the next two chapters, Chapters 4 and 5, we provide two different models showing that intention recognition promotes the emergence of cooperation, within the framework of the iterated PD. In Chapter 6, we study the role of commitment in the evolution of cooperation. Then, in Chapter 7, we provide a model combining intention recognition and commitment strategies, leading to improved cooperation, even in the framework of non-repeated interactions (the one-shot PD).

# Chapter 4

---

## INTENTION RECOGNITION PROMOTES THE EMERGENCE OF COOPERATION: A BAYESIAN NETWORK MODEL

---

*There is a good side and a bad side to most people, and  
in accordance with your own character and disposition  
you will bring out one of them  
and the other will remain a sealed book to you.*

Mark Twain (1835-1910).

**Abstract.** We derive a three-layer Bayesian Network model for intention recognition in the context of repeated social dilemmas and evolutionary game theory, by assessing the internal dynamics of trust between intention recognizers and their opponents. On the basis of past direct interactions, intention recognizers are able to some extent to predict the next move of their opponents (cooperate or defect). That in turn enables them to prevail over the most famous strategies of repeated dilemmas of cooperation, even in the presence of noise.

---

<sup>0</sup>This chapter is based on the publication (Han et al., 2011a).

## 4.1 Introduction

We model intention recognition within the framework of repeated interactions. In the context of direct reciprocity (Trivers, 1971; Nowak, 2006a; Sigmund, 2010) intention recognition is being performed using the information about past *direct* interactions. We study this issue using the iterated Prisoner’s Dilemma (PD), i.e., intentions are inferred from past individual experiences.

Contrary to other approaches dealing with the integration of (direct or indirect) information about the past in individual decisions, e.g. in (Masuda and Ohtsuki, 2009; Ohtsuki and Iwasa, 2006; Wang et al., 2008; Vukov et al., 2011), intention recognition is performed using a Bayesian Network (BN) model. BNs have proven to be one of the most successful approaches for intention recognition (Charniak and Goldman, 1993; Bui et al., 2002; Tahboub, 2006; Pereira and Han, 2011b). Their flexibility for representing probabilistic dependencies as well as causal relations, and the efficiency of inference methods have made them an extremely powerful tool for problem solving under uncertainty (Pearl, 1988, 2000; Spirtes et al., 2000; Glymour, 2001), and appropriate to deal with several probabilistic as well as causal dependencies occurring in intention recognition. There has been mounting evidence that the behavior of the human perceptual, motor, and cognitive systems are captured remarkably well by BN models, which provide a framework for understanding information processing in the brain (Wolpert and Ghahramani, 2004; Kording and Wolpert, 2004; Tenenbaum and Griffiths, 2001).

We derive a BN model for intention recognition in the context of social dilemmas, taking into account mutual trusts between the intention recognizer and his opponent. Trusts are accumulated through past interactions, assuming that intention recognizers have a memory. Greater memory sizes enable to build longer-term mutual trusts, and therefore allow better tolerance to the errors of intended actions. We study analytically the case of small memory size, and experimentally the effect of having greater memory sizes. In addition, we compare the performance of intention recognizers with the most famous strategies of the iterated PD (see above).

Here we show that our innovative intention recognition strategy (*IR*) can correct mistakes even better than win-stay-lose-shift (*WSLS*), and not be exploited by pure defectors. We compare the performance of tit-for-tat (*TFT*), *WSLS*, and *IR*, using EGT methods in finite populations (Imhof et al., 2005; Nowak et al., 2004), in a well-mixed population of pure cooperators (AllC), pure defectors (AllD) and,



**Figure 4.1: Bayesian Network for Intention Recognition in Social Dilemmas.** Pre-intentional level has one node, *oTrust (Tr)*, receives Boolean values, *t (true)* or *f (false)*, representing the other’s trust on us (the intention recognizers). Intentional level has one node, *Intention (I)*, receiving value *C* or *D*, corresponding to more cooperative and more defective, respectively, in the past. It is causally affected by *oTrust*. Activity level has one node, *pastObs (O)*, causally affected by *Intention* node. Its value is a pair  $(n_C, n_D)$  where  $n_C$  and  $n_D$  are the number of times the recognized player cooperated and defected, respectively, in the recent  $M$  (memory size) steps. The node *pastObs* is the only observed (evidence) node.

additionally, of individuals of either of the above other three strategies, as well as all these strategies together. The results show that *IR* performs best, in the sense that populations spend more time in a homogenous state of *IRs*, even in the presence of noise.

## 4.2 Bayesian Network for Intention Recognition in Social Dilemmas

In Chapter 2 (also in (Pereira and Han, 2009c, 2011b)), a general BN model for intention recognition is presented and justified based on Heinze’s intentional model (Heinze, 2003; Tahboub, 2006). Basically, the BN consists of three layers: cause/reason nodes in the first layer (called *pre-intentional*), connecting to intention nodes in the second one (called *intentional*), in turn connecting to action nodes in the third (called *activity*). Intuitively, the observed actions of an agent are causally affected by his/her intentions, which are in turn causally affected by the causes/reasons for which he committed to the intentions (Bratman, 1987, 1999; Cohen and Levesque, 1990; Singh, 1991).

Based on this general model, we present an intention recognition model in the context of the social dilemmas, taking into account the past *direct* interactions (Figure 4.1). The model is described from the view of an intention recognizer (denoted by  $\mathcal{I}$ ) with respect to a co-player (denoted by  $\mathcal{J}$ ), whose intention (C or D) is to be recognized. A player’s intentions here can be understood as the characters or

types of the player: how cooperative or defective he is in general when playing with me. Saying that the co-player has intention C (resp., D) means that, in general, he intends to cooperate with me (resp., exploit or defect towards me). Thus, if he has been cooperative in the past, it is likely he will continue to cooperate in the current interaction.

$\mathcal{J}$ 's intention in a given interaction is causally affected by the trust he holds towards his opponent ( $\mathcal{I}$ ), which is accumulated over their past (observed) interactions.  $\mathcal{J}$ 's intention in turn has given rise to his past actions. Let  $M > 0$  be the memory size of intention recognizers, i.e. they can remember their moves and their opponents' moves in the last  $M$  rounds of interaction with any specific players.

For this BN, we need to determine the prior probability of the node *oTrust*, i.e.  $P(Tr)$ ; the conditional probability table (CPD) of node *Intention*—specifying the conditional probability of  $\mathcal{J}$  having an intention (C or D) given the trust he holds towards his opponent ( $\mathcal{I}$ ), i.e.  $P(I|Tr)$ ; and the CPD table of the node *pastObs*—specifying the conditional probability of the past observations given  $\mathcal{J}$ 's intention (C or D), i.e.  $P(O|I)$ .

To begin with, let  $n_C(\mathcal{X}, \mathcal{Y})$  and  $n_D(\mathcal{X}, \mathcal{Y})$  denote the numbers of times a player  $\mathcal{X}$  cooperated and defected, respectively, in the last  $M$  interactions with another player  $\mathcal{Y}$ . Note that  $n_C(\mathcal{X}, \mathcal{Y}) + n_D(\mathcal{X}, \mathcal{Y}) \leq M$ , and the equality occurs only when the two players have interacted with each other at least  $M$  times.

**Trust Distribution.** The probability that  $\mathcal{J}$  trusts  $\mathcal{I}$  is given by how often  $\mathcal{I}$  cooperated with  $\mathcal{J}$ . This can be written as

$$P(Tr = t) = \frac{1}{2} + \frac{n_C(\mathcal{I}, \mathcal{J}) - n_D(\mathcal{I}, \mathcal{J})}{2M}. \quad (4.1)$$

It is easily seen that  $0 \leq P(Tr = t) \leq 1$ , and  $P(Tr = t) = 0$  (resp., 1) if  $\mathcal{I}$  defected (resp., cooperated) in all recent  $M$  interactions. These correspond to the extremes that  $\mathcal{I}$  lost all his/her trust (resp., gained complete trust) concerning  $\mathcal{J}$ . We further assume that, in the first interaction, the trust level is neutral:  $P(Tr = t) = \frac{1}{2}$ .

**Definition of  $\mathbf{P(I|Tr)}$ .** We use the following CPD table

$$P(I = C|Tr = t) = P(I = D|Tr = f) = h,$$

$$P(I = C|Tr = f) = P(I = D|Tr = t) = 1 - h,$$

where  $h$  is the probability the intention recognizer ( $\mathcal{I}$ ) thinks the co-player ( $\mathcal{J}$ ) has intention  $C$  given that he ( $\mathcal{J}$ ) completely trusts me ( $\mathcal{I}$ ).

As the intention recognizers are cooperative, i.e. they seek the cooperators to cooperate with and generously start by cooperating with everybody, we assume  $h \geq 0.5$ . This probability reflects the intention recognizers' **optimistic level**—assumed fixed for their entire life cycle (generation).

**Definition of P(O|I).** The conditional probability of the past observations about the co-player given his intention. It can be interpreted as how trustful or cooperative the intention recognizer ( $\mathcal{I}$ ) thinks his co-player ( $\mathcal{J}$ ) is, and can be defined as how often  $\mathcal{J}$  cooperated with  $\mathcal{I}$  in an interaction. It can be given as follows

$$\begin{aligned} P(O = (n_1, n_2)|I = C) &= \frac{1}{2} + \frac{n_1 - n_2}{2M}, \\ P(O = (n_1, n_2)|I = D) &= \frac{1}{2} + \frac{n_2 - n_1}{2M}, \end{aligned} \tag{4.2}$$

where  $n_1 = n_C(\mathcal{J}, \mathcal{I})$  and  $n_2 = n_D(\mathcal{J}, \mathcal{I})$ .

In a nutshell, the intention recognition model presented takes into account the past direct interactions, in terms of mutual trusts, which are encoded into a Bayesian Network.

**Intention Recognizer.** In an interaction, the probabilities of the co-player having intention  $C$  and having intention  $D$ , given his  $M$  recent past actions  $o = (n_C, n_D)$ , are computed using Eq. (2.2)

$$\begin{aligned} p(I = C|O = o) &= \frac{p(I = C, O = o)}{p(O = o)}, \\ p(I = D|O = o) &= \frac{p(I = D, O = o)}{p(O = o)}. \end{aligned} \tag{4.3}$$

The intention recognizer plays  $C$  if he recognizes that the co-player is more cooperative (thus more likely to play  $C$  than to play  $D$ ), i.e.  $p(I = C|O = o) \geq p(I = D|O = o)$ . Otherwise, he plays  $D$ . Players using this intention recognition strat-

egy are henceforth referred to as *IR* players<sup>1</sup>. With the trust functions given in Eq. (4.1) and (4.2), we obtain that the intention recognizer ( $\mathcal{I}$ ) cooperates with his co-player ( $\mathcal{J}$ ) iff (see Appendix 4.5.1 for details)

$$\Delta = sP + Q \geq 0, \quad (4.4)$$

where  $s = 2h - 1$  ( $0 \leq s \leq 1$ );  $P = n_C(\mathcal{I}, \mathcal{J}) - n_D(\mathcal{I}, \mathcal{J})$  and  $Q = n_C(\mathcal{J}, \mathcal{I}) - n_D(\mathcal{J}, \mathcal{I})$ .

### 4.3 Results

We consider a population of constant size  $N$ . At each evolution step, a random pair of players are chosen to play with each other. The population consists of pure cooperators (*AllCs*), pure defectors (*AllDs*) plus either of *TFTs* or of *WSLSs* or of intention recognizers who, being capable of recognizing another individual's intention based on the past interactions, seek the cooperators to cooperate with and to defect toward detected defectors.

We shall start by considering a finite population consisting of *AllC*, *AllD* and *IR* players. It is easily seen that in the absence of noise, a player adopting a *IR* strategy performs similarly to a *TFT* player, i.e. always cooperates with an *AllC*, always defects with *AllD* after cooperating in the first round, and always cooperate amongst themselves. In the sequel we study the performance of *IR* in the presence

---

<sup>1</sup>For those who are interested in Logic Programming modeling (other readers can just ignore this footnote), in the following we provide a declarative representation of this *IR* strategy using the intention-based decision making framework in Chapter 3.

1. *abds*([*move*/1]).
2. *on\_observed*(*decide*)  $\leftarrow$  *new\_round*.
3. *decide*  $\leftarrow$  *move*(*c*).      *decide*  $\leftarrow$  *move*(*d*).  
      $\leftarrow$  *move*(*c*), *move*(*d*).
4. *expect*(*move*(*X*)).
5. *move*(*c*)  $\triangleleft$  *move*(*d*)  $\leftarrow$  *has\_intention*(*co\_player*, *c*).  
     *move*(*d*)  $\triangleleft$  *move*(*c*)  $\leftarrow$  *has\_intention*(*co\_player*, *d*).

Starting a new interaction, an *IR* needs to choose a move, either cooperate (C) or defect (D) (lines 2-3). Both options are expected, and there are no expectation to the contrary (line 4). There are two *a priori* preferences in line 5 stating that an *IR* prefers to cooperate if the co-player's recognized intention is to cooperate, and prefers to defect otherwise. The built-in predicate *has\_intention*/2 in the body of the preferences triggers the intention recognition module to validate if the co-player is more likely to have the intention given in the second argument.

of noise, and compare with *TFT* and *WSLS*. First we study analytically the case of memory two, i.e.  $M = 2$ .

#### 4.3.1 Evolution of short-memory intention recognizers

In the presence of noise, let us assume that an intended action ( $C$  or  $D$ ) can fail with probability  $\epsilon \in [0, 1]$ . We obtain the following payoff matrix for *AllC*, *AllD* and *IR*, where all terms of order  $O(\epsilon^2)$  have been ignored (see Appendix 4.5.2.1)


$$\begin{pmatrix} (b-c)(1-\epsilon) & -c+(b+c)\epsilon & \frac{c(\omega-1+\epsilon(1-4\omega+\omega^2))}{1-\omega+3\epsilon\omega-\epsilon\omega^2} \\ b-b\epsilon-c\epsilon & (b-c)\epsilon & \frac{b(\omega-1+2\epsilon(1-3\omega+\omega^2))}{1-\omega+3\epsilon\omega-\epsilon\omega^2} \\ \frac{b(1-\epsilon-\epsilon\omega^2)-c(1-\epsilon(2-\omega+\omega^2))}{1-\epsilon\omega^2} & \frac{\epsilon(b-c(2+\omega-\omega^2-\omega^3))}{1-\epsilon\omega^2(1+\omega)} & \frac{\epsilon(b(2+\omega-\omega^2)-c)}{1-\epsilon\omega} \end{pmatrix}.$$

Let  $A(X, Y)$  be the payoff of strategist  $X$  when playing against strategist  $Y$ . We can show that (Appendix 4.5.3)

$$\begin{aligned} A(IR, IR) &> A(WSLS, WSLS) > A(TFT, TFT), \\ A(AllD, IR) &< A(AllD, TFT) < A(AllD, WSLS), \\ A(IR, AllD) &> A(TFT, AllD) > A(WSLS, AllD), \\ A(AllC, WSLS) &< A(AllC, TFT) < A(AllC, IR), \\ A(WSLS, AllC) &> A(TFT, AllC) > A(IR, AllC). \end{aligned}$$

The first inequality implies that *IR* deals with noise better than *TFT* and *WSLS*, when interacting with individuals alike. As a result, a homogeneous population of *IRs* has a higher level of cooperation (thus, greater average payoff) than the ones of *WSLSs* and *TFTs*<sup>2</sup>. The next two inequalities imply that *IR* deals with *AllD* better than *TFT*, which is in turn better than *WSLS*. The fixation probability of an *IR* taking over a population of *AllD* is greater than those of *TFT* and *WSLS* (Figure 4.2). Finally, the last two inequalities imply that *IR* is more cooperative to

<sup>2</sup>Recall that a homogeneous population or homogeneous state is a population with just a single strategy, i.e., all individuals of the population adopt that strategy.



**Figure 4.2: Transition probabilities and stationary distributions (in percentage)**, computed analytically for two distinct sets of parameters. We consider a population of *AIIc*s, *AIID*s and either *WSLS*s, *TFT*s or *IR*s ( $M = 2$ ). The black arrows are only shown for the transition directions that are rather more likely than neutral. The strongest transition is from *AIIc* to *AIID*. The transition from *AIID* to *IR* is stronger than to *TFT*. This is reversed in case of *WSLS*, where the most probable transition occurs from *WSLS* to *AIID*. For slow intensity of selection  $\beta$  (panel a), the transitions between *AIIc* and *IR* and *AIIc* and *TFT* are near neutral, and for strong selection (panel b), there is a transition from *TFT* to *AIIc*, which is stronger than from *IR* to *AIIc*. Also, the greater  $\beta$ , the stronger the transition from *AIIc* to *AIID* and to *WSLS*. The calculations in both cases are made with  $N = 100$  and  $\omega = 0.9$ ;  $\rho_N = 1/N$  denotes the neutral fixation.

### 4.3. Results

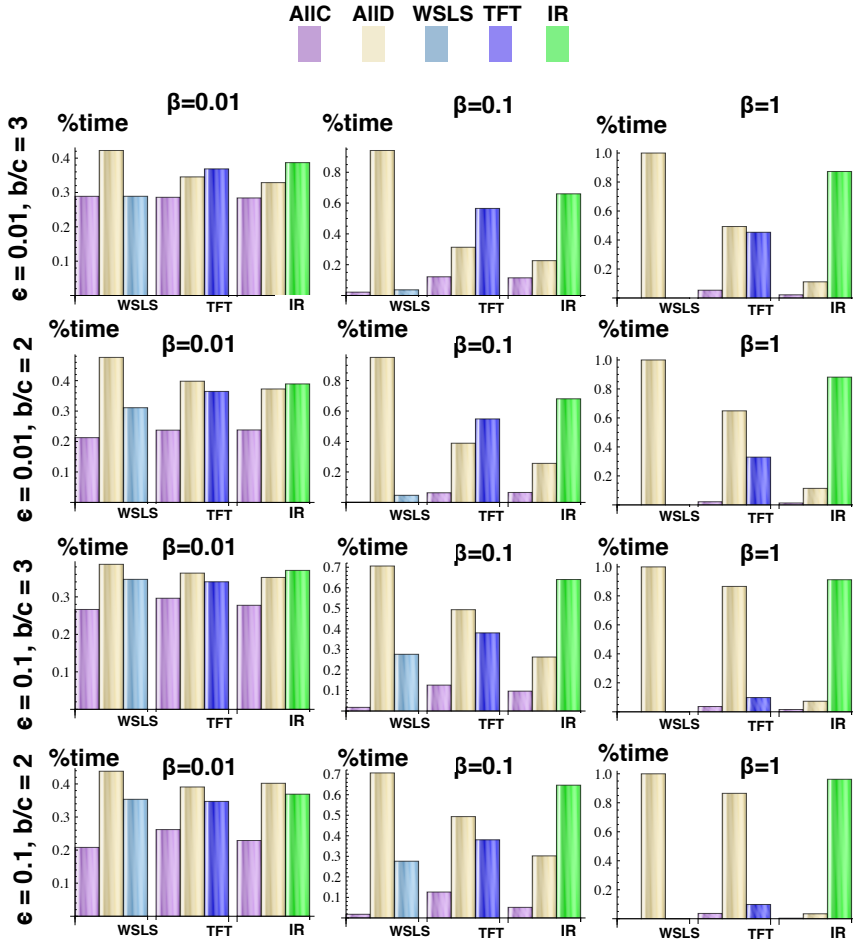


Figure 4.3: Percentage of time spent at *AUC*, *AID* and either *WSLS*, *TFT* or *IR*. We compare *IR* with *WSLS* and *TFT* when interacting with *AUCs* and *AIDs*. We consider different selection intensities  $\beta$  (0.01, 0.1 and 1), noise levels  $\epsilon$  (0.01 and 0.1) and PD benefit-to-cost ratios  $b/c$  (2 and 3). In all cases, *IR* is better than *TFT* and *WSLS*. In a population of *AUCs*, *AIDs* and *IRs*, the system spends more time in the homogeneous state of *IRs*, especially when selection is strong. *TFT* and *WSLS* perform poorly at strong selection intensities, as the population spends most of the time at *AIDs*. The calculations in all cases are made with  $N = 100$  and  $\omega = 0.9$ .




Figure 4.4: Stationary distribution in a population of five possible strategies *AIC*, *AID*, *WSL*, *TFT* and *IR* ( $M = 2$ ). The population spends most of the time in the homogeneous state of *IRs*. *WSL* also performs well in this 5-strategy setting, reconfirming that it needs other catalyzers such as *TFT* to perform well. On the contrary, these results show that *IR* performs well in either case (the other setting in Figure 4.3). The calculations in all cases are made with  $N = 100$ ,  $b/c = 3$ ,  $\omega = 0.9$  and  $h = 0.6$ . The average fitness of *IR*, *WSL* and *TFT* players interacting with each other are obtained by averaging over  $10^7$  simulated interactions.

*AllC* (i.e. it is more tolerant of noise originated by *AllC*).

Figures 4.2 and 4.3 show a comparison among *WSLS*, *TFT* and *IR* in terms of the percentage of time the population spends at their homogenous state, in a setting of either strategy and *AllC* and *AllD* individuals. In all cases, *IR* is better than *TFT* and *WSLS*. In a population of *AllC*s, *AllD*s and *IR*s, the population spends more time in the homogeneous state of *IR*s, especially when the intensity of selection is strong. *TFT* and *WSLS* poorly perform at strong selection intensities—the population spends most of the time at *AllD*s. The poor performance of *WSLS* is not surprising, as *WSLS* needs *TFT* players as a catalyst to perform well, as discussed in (Sigmund, 2010).

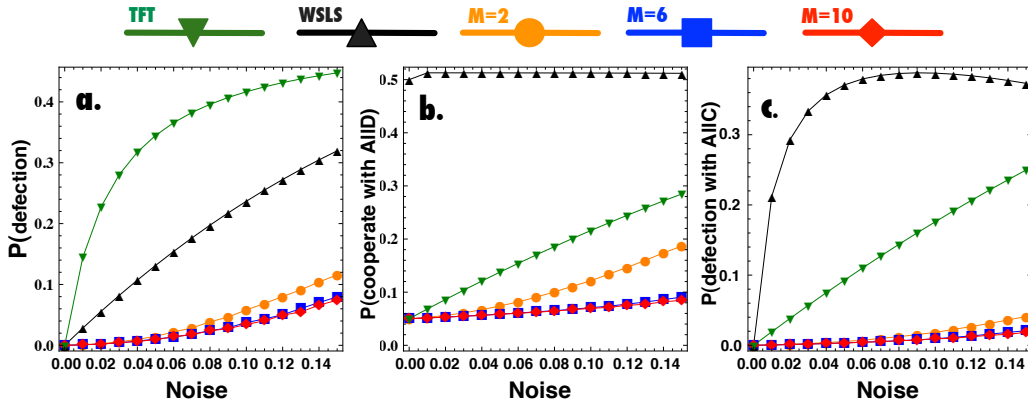
Figure 4.4 shows the results for the setting where all the five strategies *AllC*, *AllD*, *TFT*, *WSLS* and *IR* are in the population. Again, the population will be most likely found in the homogeneous state of *IR*s. *WSLS* also performs well in this setting (since *TFT*s are present).

In short, in these two commonly used settings, *IR* always outperforms *TFT* and *WSLS*. The population spends more time in the homogeneous state of *IR*s. Furthermore, since a population of *IR*s is highly cooperative, it is clear that the introduction of intention recognition significantly increases the cooperation level of the population, leading to a greater social welfare.

### 4.3.2 The role of memory size

The impressive results obtained in the previous section addressed the evolutionary chances of very short-memory intention recognizers. Yet, it is reasonable to suppose that individuals may record a larger number of rounds and use it at their profit. As shown in Figure 4.5a, a greater memory size enables a better tolerance to noise. It allows the build-up of long-term mutual trusts, which enables a better assessment of errors. The intention recognizers become more generous (see Fig. 4.5a) and, above all, more tolerant to other players' errors. As a result, a homogenous population of large memory *IR*s can reach closer to the state of perfect cooperation, even in the presence of noise.

In addition, they are more resilient to changing opinion about untrustworthy players. Namely, the greater  $M$  the smaller probability that an *IR* defects with another *IR* and with an *AllC* (see Fig. 4.5c), and the smaller the probability of cooperating with *AllD* (see Fig. 4.5b).



**Figure 4.5:** Probabilities of reactions computed numerically for *IRs* of different memory sizes ( $M = 2, 6, 10$ ), and analogous quantities for *WSLS* and *TFT* strategists, as a function of the level of noise. a) Probability of defection of *IR*, *TFT* and *WSLS* strategists when playing with themselves. The probability of defection between *IRs* decreases with the memory size. b) Probability that each of the three strategies—*IR*, *TFT* and *WSLS*—cooperates with a *AID*. The probability that *IR* cooperates with *AID* decreases when memory size increases. c) Probability that a *IR*, *TFT* and *WSLS* player defects with a *AIC*. The probability that *IR* defects with *AIC* decreases when memory size increases. In all cases,  $\omega = 0.9$ ,  $h = 0.6$  and the probabilities are obtained by averaging over  $10^7$  simulated interactions.

### 4.3.3 The role of *IRs* optimism and their tolerance to noise

In the following, we study the relation of the optimistic level  $h$  of an *IR* with his/her noise tolerance, towards a cooperative and towards a defective co-player. For simplicity, we consider the following two cases.

First of all, suppose that at the  $m$ -th round, we have  $P = P_0 > 0$  and  $Q = Q_0 > 0$  (i.e. the players were cooperative), and the co-player will constantly defect (with probability  $1 - \epsilon$ ). Let  $k$  be the expected number of rounds *IR* is tolerant to his co-player's defection. Clearly,  $k \leq M$  (see Appendix 4.5.1).

If  $m \geq M$ , from Eq. (4.4) and the fact that *IR*'s memory size is  $M$  we have that  $k$  must satisfy the following equation

$$sP_k + Q_k = 0, \quad (4.5)$$

where  $P_{i+1} = P_i + (1 - \epsilon - P_i/M)$  and  $Q_{i+1} = Q_i - (1 - \epsilon + Q_i/M)$  for all  $1 \leq i \leq k-1$ . This can be explained as follows. With probability  $(1 - \epsilon)$ , *IR* cooperates, thus the value of  $P_i$  is increased by  $(1 - \epsilon)$ . But since  $m \geq M$ , the value of  $P_i$  is also decreased,

on average, by  $P_i/M$  (one round is “forgotten”). A similar argument can be used for  $Q_i$ .

By a simple computation<sup>3</sup> we obtain that Eq. (4.5) is equivalent to

$$\left(\frac{M-1}{M}\right)^k = \frac{M(1-s)(1-\epsilon)}{M(1-s)(1-\epsilon) + sP_0 + Q_0} = \frac{1}{1 + \frac{sP_0 + Q_0}{M(1-s)(1-\epsilon)}}.$$

The right-hand side is clearly a decreasing function of  $s$ . Thus,  $k$  is an increasing function of  $s$ . Now if  $m < M$ , by putting  $k = k_1 + M - m$  and applying the same method we can show that  $k_1$  is also an increasing function of  $s$ .

In short, in both cases  $k$  is an increasing function of  $s$  (hence, also of  $h$ ). It means that the more optimistic the *IR* is (i.e. the greater  $h$ ), the more tolerant/generous he is towards a cooperative co-player. It also means that the more optimistic an *IR*, he might become more generous to defective players if errors occur more frequently at the beginning of their interaction (defective players show up to be cooperative).

Similarly, let us consider the case where at some round, both players were defective, i.e.  $P_0 < 0$  and  $Q_0 < 0$ , and the co-player starts cooperating (either by mistakes or because the co-player is actually a cooperative one). The number of rounds that *IR* keeps on defecting while the co-player cooperates, by using the same method, can be shown to be an increasing function of  $h$ . Hence, the more optimistic a *IR* is, the less generous he is towards a defective player. However, it also means that the more optimistic a *IR* is, the less generous he might become to cooperative players if errors occur more frequently at the beginning of their interaction (cooperative players show up to be defective).

In short, when errors are not frequent at the beginning, a more optimistic *IR* becomes more tolerant to noise, since he/she is more generous towards cooperative players and less so to defective ones. On the other hand, if errors are frequent at the beginning (cooperative players show up to be defective, and vice-versa, defective players show up to be cooperative), the more optimistic a *IR* is, more rounds he/she would take to recognize correctly the opponents' intentions. Overall, this suggests that, ideally, a *IR* should not be optimistic at the beginning of an interaction since otherwise an error could create a wrong bad impression which is hard to recover from. When more interactions are made, a higher optimism increases the tolerance

---

<sup>3</sup>Consider the sequence  $X_i$ , where  $X_i = sP_i + Q_i$ . It satisfies the recursive equation  $X_{i+1} = \frac{M-1}{M}X_i + (1-s)(1-\epsilon)$ .

to noise of an *IR* strategist. As a result, it suggests that  $h$  can be expressed as an increasing function of time.

## 4.4 Concluding Remarks

Using the tools of EGT, in this chapter we have shown the role of intention recognition for the emergence of cooperation within the framework of the iterated Prisoner's Dilemma. Intention recognition is performed using a Bayesian Network model via computing mutual trusts between the intention recognizers and their opponents. Given the broad spectrum of problems which are addressed using this cooperative metaphor, our result indicates how intention recognition can be pivotal in social dynamics. We have shown that the intention recognition strategy prevails over the most successful existent strategies (*TFT*, *WSLS*) of the iterated PD, even when players have a very limited memory. *IR* deals with *AUD* better than *TFT* – the best known defector-dealer, and correct mistake better than *WSLS* – the best known mistake-corrector (Nowak, 2006b; Sigmund, 2010). As a result, a homogenous population of *IRs* has a higher level of cooperation than the ones of *WSLSs* and *TFTs*, resisting the invasion of other strategies.

In (Imhof et al., 2005), it has been shown that in the absence of noise, in a population of *AUCs*, *AUDs* and *TFTs*, the population spends most of the time in a homogeneous state of *TFTs*. However, as we have shown here, that it is not the case if noise is present, especially under strong selection. In the absence of noise, *IR* behaves as well as *TFT*. Moreover, *IRs* are selected by evolution in the latter case where noise is present. We have shown that in a population of *AUCs*, *AUDs* and *IRs*, the population spends most of the time in the homogeneous state of *IRs*—in a broad range of scenarios and parameters—especially when the intensity of selection is strong. We have also exhibited experimentally that in a population where all the five strategies *AUC*, *AUD*, *TFT*, *WSLS* and *IR* are present, *IRs* still prevail most of the time. Therefore, together with the fact that *IRs* can correct mistakes better than *WSLSs* and *TFTs*, the presence of *IRs* would significantly increase the overall level of cooperation of the population.

Additionally, we have shown the role of a large memory size in recognizing/correcting errors. Having a greater memory size allows intention recognizers to build longer-term mutual trusts/distrusts, and hence enables them to better

recognize erroneous moves. It then further enables them to better tolerate a selfish act made by cooperative trustful individuals, and to refuse to cooperate after an erroneous cooperation made by a defective untrustworthy ones. Indeed, intention recognition gives rise to an incipient mechanism of commitment formation, from which future behaviors may be assessed and trust bonds established. Further discussions on the relation between commitment and intention recognition will be found in Chapter 7.

Overall, our work provides new insights into the complexity and beauty of behavioral evolution driven by elementary forms of cognition.

Notwithstanding, intentions are inferred from past individual experiences only, within the framework the iterated Prisoner's Dilemma (PD). Naturally, the same principles could be extended to cope with indirect information, as in indirect reciprocity (Nowak and Sigmund, 2005; Pacheco et al., 2006; Ohtsuki and Iwasa, 2006). This eventually introduces moral judgment and concern for individual reputation, which constitutes "per se" an important area where intention recognition may play a pivotal role. In subsequent work we aim to extend our model to incorporate this indirect information to further enhance the performance of intention recognizers.

## 4.5 Appendices

### 4.5.1 Decision making with intention recognition

Here we derive a simplified expression for intention recognizers' decision making, i.e., when to cooperate and when to defect. Recall that an *IR* cooperates iff he recognizes that the co-player is more likely to cooperate than to defect, i.e.  $p(I = C|O = o) \geq p(I = D|O = o)$ . From Eq. (4.3) (main text), we have that the inequality holds if and only if

$$\begin{aligned}
 & p(I = C, O = o) \geq p(I = D, O = o) \\
 \iff & p(I = C, O = o, Tr = t) + p(I = C, O = o, Tr = f) \geq \\
 & p(I = D, O = o, Tr = t) + p(I = D, O = o, Tr = f) \\
 \iff & p(O = o|I = C)p(I = C|Tr = t)p(Tr = t) + \\
 & p(O = o|I = C)p(I = C|Tr = f)p(Tr = f) \geq \\
 & p(O = o|I = D)p(I = D|Tr = t)p(Tr = t) + \\
 & p(O = o|I = D)p(I = D|Tr = f)p(Tr = f) \\
 \iff & mtr [otr.h + (1 - otr)(1 - h)] \geq (1 - mtr) [otr(1 - h) + (1 - otr)h],
 \end{aligned}$$

where  $otr = P(Tr = t)$  and  $mtr = P(O = o|I = C)$ .

Simplifying both sides we obtain

$$h(2.otr - 1) \geq otr - mtr. \quad (4.6)$$

From Eqs. (4.1) and (4.2) (main text), (4.6) can be rewritten as

$$\Delta = sP + Q \geq 0, \quad (4.7)$$

where  $s = 2h - 1$  ( $0 \leq s \leq 1$ ),  $P = n_C(\mathcal{I}, \mathcal{J}) - n_D(\mathcal{I}, \mathcal{J})$ , and  $Q = n_C(\mathcal{J}, \mathcal{I}) - n_D(\mathcal{J}, \mathcal{I})$ .

In short, (4.7) provides a simple decision making model for the intention recognizer, taking into account the co-player's  $M$  recent moves as well as the intention recognizer own's  $M$  recent moves, linking by a factor of  $(2h - 1)$  where  $h$  is the optimistic level of the intention recognizer.

We have that  $-M \leq P, Q \leq M$ . Let us look at some extreme cases.

- If  $Q = M$ , i.e. the co-player cooperates in all  $M$  recent steps, then  $\Delta > 0$ . The *IR* cooperates in the next round.
- If  $Q = -M$ , i.e. the co-player defects in all  $M$  recent steps, then  $\Delta < 0$ . The *IR* defects in the next round.
- If  $h = 1/2$ , then  $\Delta = Q$ . The *IR* only considers the co-player's past actions to decide his next move: if the co-player cooperated at least as much as defected in the  $M$  recent rounds, then *IR* cooperates in the next round, and defects otherwise. We henceforth consider  $h > 1/2$ : the *IR* also takes its own moves into account.

### 4.5.2 Memory-Two Intention Recognizers

To provide a simple mathematical analysis, let us consider the simplest case where *IR* players have a very short memory  $M = 2$ . By following a similar method as described in (Sigmund, 2010), we derive the analytical payoff matrix for *AllC*, *AllD* and *IR* in the presence of noise. Based on that, we then compare analytically *IR* with other famous strategies, including *TFT* and *WSLS*.

To begin with, since  $s > 0$ , a memory-two *IR* decides his current move depending on the state of his own and his co-player's last two moves. There are 16 possible states, forming by all the combinations of four possible game situations ( $R, S, T, P$ ) in each of the last two encounters. We enumerate these states by  $state_i$ , with  $1 \leq i \leq 16$ .

We consider stochastic strategies  $(f, l, q_1, q_2, \dots, q_{16}) \in [0, 1]^{18}$  where  $f$  and  $l$  are the propensities to play  $C$  in the initial and second rounds, respectively, and  $q_i$  are propensities to play  $C$  after having been at  $state_i$ ,  $1 \leq i \leq 16$ .

Let us assume that player 1 using  $(f_1, s_1, p_1, p_2, \dots, p_{16})$  encounters a co-player 2 using  $(f_2, s_2, q_1, q_2, \dots, q_{16})$ . We have a Markov chain in the state space  $\{state_1, \dots, state_{16}\}$ . The transition probabilities are given by the stochastic matrix  $Q$  below. Note that one player's  $S$  is the other player's  $T$ .

$$Q = \begin{pmatrix} p_1q_1 & \dots & (1-p_1)(1-q_1) & 0 & \dots & 0 & & 0 \\ 0 & \dots & 0 & p_2q_3 & \dots & 0 & \dots & 0 \\ \dots & & & \dots & & \dots & & \\ p_5q_9 & \dots & (1-p_5)(1-q_9) & 0 & \dots & 0 & & 0 \\ \cdot & & & & & & & \\ \cdot & & & & & & & \\ \cdot & & & & & & & \\ 0 & \dots & 0 & 0 & \dots & 0 & \dots & (1-p_{16})(1-q_{16}) \end{pmatrix}.$$

There are several zeros in the matrix. The state with the second component  $X \in \{R, S, T, P\}$  can only go to the states with the first component being  $X$ . For example, in the first row,  $state_1$ , i.e.  $(R, R)$ , can only go to states with the first component being  $R$ , i.e.  $(R, R), (R, T), (R, S), (R, P)$  (i.e.  $state_i, 1 \leq i \leq 4$ ).

The initial probabilities for the sixteen states are given by the vector

$$\mathbf{f} = \{f_1f_2s_1s_2, f_1f_2s_1(1-s_2), f_1f_2(1-s_1)s_2, f_1f_2(1-s_1)(1-s_2), \\ f_1(1-f_2)s_1s_2, f_1(1-f_2)s_1(1-s_2), f_1(1-f_2)(1-s_1)s_2, \\ f_1(1-f_2)(1-s_1)(1-s_2), (1-f_1)f_2s_1s_2, (1-f_1)f_2s_1(1-s_2), \\ (1-f_1)f_2(1-s_1)s_2, (1-f_1)f_2(1-s_1)(1-s_2), \\ (1-f_1)(1-f_2)s_1s_2, (1-f_1)(1-f_2)s_1(1-s_2), \\ (1-f_1)(1-f_2)(1-s_1)s_2, (1-f_1)(1-f_2)(1-s_1)(1-s_2)\}.$$

In the next round, these probabilities are given by  $\mathbf{f}Q$ , and in the round  $n$  by  $\mathbf{f}Q^n$ . We denote by  $\mathbf{g}$  the vector  $\{R, S, T, P, R, S, T, P, R, S, T, P, R, S, T, P\}$ , then the payoff for player 1 in round  $n$  is given by

$$A(n) = \mathbf{g} \cdot \mathbf{f}Q^n. \tag{4.8}$$

For  $\omega < 1$  the average payoff per round is  $(1-\omega) \sum w^n A(n)$  (Sigmund, 2010), i.e.,

$$(1-\omega)\mathbf{g} \cdot \mathbf{f}(Id - \omega Q)^{-1}, \tag{4.9}$$

where  $Id$  is the identity matrix of size 16.

#### 4.5.2.1 Payoff Matrices in Presence of Noise

We now derive the payoff matrix for the 3-strategy game involving  $AUC$ ,  $AllD$  and  $IR$  in the presence of noise, i.e., an intended action (C or D) can fail with probability

$\epsilon \in [0, 1]$ . The strategies *AllC* and *AllD* are given by  $\{1 - \epsilon, 1 - \epsilon, 1 - \epsilon, \dots, 1 - \epsilon\}$ ,  $\{\epsilon, \epsilon, \epsilon, \dots, \epsilon\}$ , respectively. Strategy *IR* is given by  $\{1 - \epsilon, l, q_1, \dots, q_{16}\}$  where  $l = 2\epsilon(1 - \epsilon)$  when playing with *AllD* and  $\epsilon^2 + (1 - \epsilon)^2$  when playing with *AllC* or *IR*<sup>4</sup>; and  $q_i = 1 - \epsilon$  for  $i \in \{1..5, 7, 9..11, 13\}$  and  $\epsilon$  for  $i \in \{6, 8, 12, 14..16\}$ . Note that  $X..Y$ , where  $X \leq Y$  are two natural numbers, denote the sequence  $X, X + 1, \dots, Y$ . Basically, *IR* will cooperate in the next round (with probability  $1 - \epsilon$ ) iff there is at least a mutual cooperation (i.e. *R*) in the last two steps (i.e.  $i \in \{1..5, 13\}$ ) or there is at least one *T* and no *P* (i.e.  $i \in \{7, 9..11\}$ ).

Considering the PD game with  $T = b, R = b - c, P = 0, S = -c$ , the payoff matrix for *AllC*, *AllD* and *IR*, applying Eq. (4.9) for each pair of strategies, is approximately given (where all terms of order  $O(\epsilon^2)$  have been ignored),

$$\left( \begin{array}{ccc} (b-c)(1-\epsilon) & -c + (b+c)\epsilon & \Lambda \\ b - b\epsilon - c\epsilon & (b-c)\epsilon & \frac{\epsilon(b(2+\omega-\omega^2)-c)}{1-\epsilon\omega} \\ \frac{b(1-\epsilon-\epsilon\omega^2)-c(1-\epsilon(2-\omega+\omega^2))}{1-\epsilon\omega^2} & \frac{\epsilon(b-c(2+\omega-\omega^2-\omega^3))}{1-\epsilon\omega^2(1+\omega)} & \frac{(b-c)(1+\epsilon(-2+4\omega+\omega^2+\omega^3))}{1+\epsilon\omega(3+\omega+\omega^2)} \end{array} \right),$$

where  $\Lambda = \frac{c(\omega-1+\epsilon(1-4\omega+\omega^2))+b(\omega-1+2\epsilon(1-3\omega+\omega^2))}{1-\omega+3\epsilon\omega-\epsilon\omega^2}$ <sup>5</sup>.

By a similar method, we derive the payoff matrixes *AllC*, *AllD* and either *TFT* or *WSLS* (the results for the general case of Prisoner's Dilemma can be found in (Imhof et al., 2007)). First, for *AllC*, *AllD* and *WSLS*

$$\left( \begin{array}{ccc} (b-c)(1-\epsilon) & b\epsilon - c(1-\epsilon) & c(-1+\epsilon) - \frac{b(1-\epsilon)(1-(1-2\epsilon)\omega)}{-1+(1-2\epsilon)^2\omega} \\ b(1-\epsilon) - c\epsilon & (b-c)\epsilon & -c\epsilon + \frac{b(1+2\epsilon^2\omega-\epsilon(1+\omega))}{1+(1-2\epsilon)^2\omega} \\ \Lambda' & \frac{b\epsilon(1+(1-2\epsilon)^2\omega)-c(1-\epsilon-\epsilon\omega+2\epsilon^2\omega)}{1+(1-2\epsilon)^2\omega} & (b-c)(1-\epsilon)(1-2\epsilon\omega+4\epsilon^2\omega) \end{array} \right),$$

where  $\Lambda' = b(1-\epsilon) + \frac{c(1-\epsilon)(1-(1-2\epsilon)\omega)}{-1+(1-2\epsilon)^2\omega}$ .

Now, for *AllC*, *AllD* and *TFT*, the payoff matrix reads

$$\left( \begin{array}{ccc} (b-c)(1-\epsilon) & b\epsilon - c(1-\epsilon) & c(-1+\epsilon) + b(1+2\epsilon^2\omega - \epsilon(1+\omega)) \\ b(1-\epsilon) - c\epsilon & (b-c)\epsilon & -c\epsilon - b(-1+\epsilon)(1+(-1+2\epsilon)\omega) \\ \Lambda'' & b\epsilon - c(1-\epsilon)(1-(1-2\epsilon)\omega) & \frac{(b-c)(1-\omega+\epsilon(-1+2\omega))}{1+(-1+2\epsilon)\omega} \end{array} \right),$$

where  $\Lambda'' = b(1-\epsilon) - c(1-\epsilon-\epsilon\omega+2\epsilon^2\omega)$ .

<sup>4</sup>These can be easily seen from the fact that in the second round *IR*'s decision only depends on what the co-player did in the first round (cf. (4.7)).

<sup>5</sup>This notation is added only for a better alignment.

### 4.5.3 Comparing IR, WSLS, and TFT

Let  $A(X, Y)$  be the payoff of strategist X when playing with strategist Y (from the payoff matrices). In the sequel we will show that, for  $b \geq (1 + \omega)c$ ,  $\omega > 3/4$  and  $\epsilon < 1/3$ :

$$\begin{aligned} A(IR, IR) &> A(WSLS, WSLS) > A(TFT, TFT), \\ A(AllD, IR) &< A(AllD, TFT) < A(AllD, WSLS), \\ A(IR, AllD) &> A(TFT, AllD) > A(WSLS, AllD), \\ A(AllC, WSLS) &< A(AllC, TFT) < A(AllC, IR), \\ A(WSLS, AllC) &> A(TFT, AllC) > A(IR, AllC). \end{aligned}$$

The first condition always holds by the usual assumption in Donation game  $b \geq 2c$ ; the second condition means the game is repeated at least 4 rounds.

#### 4.5.3.1 TFT, WSLS, IR: Competing Amongst Themselves

Comparing the bottom rightmost element of each payoff matrix we have

$$\begin{aligned} A(WSLS, WSLS) - A(TFT, TFT) &= \\ \frac{(b - c)(1 - 2\epsilon)^2 \epsilon \omega (2\omega - 2\epsilon\omega - 1)}{1 + (-1 + 2\epsilon)\omega} &> 0 \end{aligned}$$

(for small enough  $\epsilon$  and big enough  $\omega$ , namely  $\epsilon < 1 - \frac{1}{2\omega}$ ),

$$\begin{aligned} A(IR, IR) - A(WSLS, WSLS) &= \\ \frac{(b - c)\epsilon (-1 + 3\omega + \epsilon (-3\omega + 7\omega^2 + 3\omega^3 + 2\omega^4))}{1 + \epsilon\omega (3 + \omega + \omega^2)} &> 0. \end{aligned}$$

(all the terms of order  $O(\epsilon^3)$  in the numerator have been ignored).  
In short, for  $\epsilon < 1 - \frac{1}{2\omega}$  we have

$$A(IR, IR) > A(WSLS, WSLS) > A(TFT, TFT).$$

**4.5.3.2 TFT, WSLS, IR: With AllD**

$$A(\text{AllD}, \text{WSLS}) - A(\text{AllD}, \text{TFT}) = \frac{b(1-2\epsilon)^2\omega(\epsilon + \omega - 3\epsilon\omega + 2\epsilon^2\omega)}{1 + (1-2\epsilon)^2\omega} > 0,$$

$$A(\text{AllD}, \text{TFT}) - A(\text{AllD}, \text{IR}) = \frac{b(1-\omega)(1-3\epsilon-2\epsilon\omega)}{1-\epsilon\omega} > 0$$

(all the terms of order  $O(\epsilon^2)$  in the numerator have been ignored),

$$A(\text{TFT}, \text{AllD}) - A(\text{WSLS}, \text{AllD}) = \frac{c(1-2\epsilon)^2\omega(\epsilon + \omega - 3\epsilon\omega + 2\epsilon^2\omega)}{1 + (1-2\epsilon)^2\omega} > 0,$$

$$A(\text{IR}, \text{AllD}) - A(\text{TFT}, \text{AllD}) = \frac{c\epsilon\omega^3 + c(1-\omega)(1-3\epsilon-\epsilon\omega-\epsilon\omega^2)}{1-\epsilon\omega^2(1+\omega)} > 0$$

(all the terms of order  $O(\epsilon^2)$  in the numerator have been ignored).

In short, it always holds that

$$A(\text{AllD}, \text{WSLS}) > A(\text{AllD}, \text{TFT}) > A(\text{AllD}, \text{IR}),$$

$$A(\text{IR}, \text{AllD}) > A(\text{TFT}, \text{AllD}) > A(\text{WSLS}, \text{AllD}).$$

**4.5.3.3 TFT, WSLS, IR: With AllC**

$$A(\text{AllC}, \text{WSLS}) - A(\text{AllC}, \text{TFT}) = \frac{-b(1-2\epsilon)^2\epsilon\omega(1+(1-2\epsilon)\omega)}{1-(1-2\epsilon)^2\omega} < 0,$$

$$A(\text{AllC}, \text{TFT}) - A(\text{AllC}, \text{IR}) =$$

$$\frac{-\epsilon(b(1-\omega)(2\omega-1) + \epsilon\omega(b - b\omega^2 + 4b\omega - 3c + c\omega))}{1-\omega + 3\epsilon\omega - \epsilon\omega^2} < 0$$

(all the terms of order  $O(\epsilon^2)$  in the numerator have been ignored;  $\omega > 3/4$ ),

$$A(\text{WSLS}, \text{AllC}) - A(\text{TFT}, \text{AllC}) = \frac{c(1-2\epsilon)^2\epsilon\omega(1+\omega-2\epsilon\omega)}{1-(1-2\epsilon)^2\omega} > 0,$$

$$A(TFT, AllC) - A(IR, AllC) =$$

$$\frac{\epsilon (2c\epsilon^2\omega^3 + c(2\omega(1 - \epsilon) - 1) + \epsilon (b\omega^2 - c\omega^2 - c\omega^3))}{1 - \epsilon\omega^2} > 0$$

(for big enough  $\omega$  and small enough  $\epsilon$ , namely  $\epsilon < 1 - \frac{1}{2\omega}$ , and  $b \geq (1 + \omega)c$ ).  
 In short, for  $b \geq (1 + \omega)c$ ,  $\omega > 3/4$ ; and  $\epsilon < 1/3$  we have

$$A(AllC, WSLS) < A(AllC, TFT) < A(AllC, IR),$$

$$A(WSLS, AllC) > A(TFT, AllC) > A(IR, AllC).$$

# Chapter 5

---

## INTENTION RECOGNITION PROMOTES THE EMERGENCE OF COOPERATION: A CORPUS-BASED MODEL

---

*His plans are calm and deeply hidden, so no one can figure them out. He changes his actions and revises his plans, so that people will not recognize them. He changes his abode and goes by a circuitous route so people cannot anticipate him. When people never understand what your intention is, then you win.*

SUN TZU (544 BC - 496 BC): *THE ART OF WAR*

**Abstract.** We introduce a new strategy that is able to assign an intention (strategy) to the actions of opponents, on the basis of an acquired corpus consisting of possible plans achieving that intention, as well as to then make decisions on the basis of such recognized intentions. The success of these intention recognizers is grounded on the unshameful beneficial exploitation of pure cooperators whilst remaining robust against defection strategies. In addition, we show how intention recognizers do indeed prevail against the most famous successful strategies of iterated dilemmas of cooperation, even in the presence of noise.

---

<sup>0</sup>This chapter is based on the publication (Han et al., 2011b) and a journal submission (Han et al., 2012a).

## 5.1 Introduction

Similarly to the previous chapter, by resorting to the tools of Evolutionary Game Theory (EGT) we study the role of intention recognition in the evolution of cooperation, within the framework of iterated Prisoner’s Dilemma (IPD). But in contrast to the intention recognition model therein, herein we describe a model to recognize a larger set of conceivable intentions, in terms of more sophisticated patterns of behavioral strategies for the IPD.

As usual, the inputs of an intention recognition system are a set of conceivable intentions and a set of plans achieving each intention—given in terms of either a plan library (Geib and Goldman, 2009; Charniak and Goldman, 1993) or a plan corpus (Blaylock and Allen, 2003, 2004; Armentano and Amandi, 2009; Han and Pereira, 2011b) (for more details see Related Work of Chapter 2). In the EGT context here, the conceivable intentions are the strategies already known to the intention recognizer, whose recognition model is learnt from a plan corpus consisting of sequences of moves in successive game rounds (called plan sessions) for each of the different strategies whilst playing IPD. For the corpus-based intention recognition, we use the method described in Chapter 2.

The rationale of the corpus-based approach in the evolutionary context relies firstly on the idea of nature-nurture co-evolution or experience inheritance (Richardson and Boyd, 2006; Shennan, 2002): the corpus represents ancestors’ given knowledge or acquired experience in interacting with known strategies. Additionally, intention recognizers can use themselves as a framework for learning and understanding those strategies by self-experimenting with them (Woodward et al., 2009)—as suggested by the famous ‘like-me’ framework (Meltzoff, 2007, 2005), i.e. hypothetically assuming that I am like the other or that the other is like me. This is often addressed in the context of the “Theory of Mind” theory (ToM) (Premack and Woodruff, 1978; Whiten, 1991; Cheney and Seyfarth, 2007), neurologically relying in part on “mirror neurons”, at several cortical levels, as supporting evidence (Iacoboni et al., 2005; Rizzolatti and Craighero, 2004; Nakahara and Miyashita, 2005). For a recent general account of the neuronal basis of mind see (Gazzaniga, 2011). Indeed, intention recognition can be considered as an elementary component of the ToM. The other components of the ToM include attributions of other mental states (to other agents) such as beliefs, desires, emotions, attention, and knowledge (Premack and Woodruff, 1978; Whiten, 1991). Closely related to this second point, notice that

contrary to the formal Evolutionary Game Theory framework being used here, there have been several works showing the evolution of mind, which are based on genetic algorithm simulations, e.g. (Takano and Arita, 2006; Zanolungo, 2007; Sayama et al., 2011; Kim and Lipson, 2009).

In addition, to the best of our knowledge, although there is a large body of literature on learning in games (Fudenberg and Levine, 1998; Macy and Flache, 2002), e.g. reinforcement learning (Sandholm and Crites, 1995; Van Segbroeck et al., 2010), very little attention has been paid to studying how some cognitive ability that requires learning (as the ability of intention recognition in this work) fares in an evolutionary setting, particularly within the framework of Evolutionary Game Theory. On the other hand, there have been some efforts to study the effects of increased memory size in evolutionary settings, e.g. see (Hauert and Schuster, 1997), though individual learning is not considered. Differently from this literature, our aim is to provide a computational model showing how the cognitive ability of intention recognition, which is so-critical and ubiquitous in humans' activities (Woodward et al., 2009; Tomasello, 2008; Meltzoff, 2007), is a viable possibility that might have been retained by natural selection.

We offer a method to acquire an intention-based decision making model from the plan corpus, stating what to play with a given co-player based on the recognized intention and the game's current state. The intention-based decision maker attempts to achieve the greatest expected benefit for itself, taking advantage of the knowledge about the co-player's intention (or strategy). The model is discussed in Subsection 5.2.3.

We show that our intention recognizers prevail against the most famous successful strategies of repeated dilemmas of cooperation, including tit-for-tat (TFT), generous tit-for-tat (GTFT), and win-stay-lose-shift (WSLS) (see above), even in the presence of noise.

For easier following by the readers, we next recall the corpus-based intention recognition method described in Chapter 2, including the description of the plan corpora in the context of the iterated Prisoner's Dilemma and a simplified description of the intention recognition model.

## 5.2 Materials and Methods

### 5.2.1 Plan Corpus Description

We describe how to create plan corpora for training and testing the corpus-based intention recognition models, for a given set of strategies. We start by making an assumption that all strategies to be recognized have memory size bounded-up by  $M$  ( $M \geq 0$ )—i.e. their decision at the current round is independent of the past rounds that are at a time distance greater than  $M$ . Note that the strategies described above all have memory size bounded by  $M = 1$ . Nevertheless, our method is for general  $M$ .

For clarity of representation, abusing notations,  $R$ ,  $S$ ,  $T$ , and  $P$  are henceforth also referred to as (elementary) game states, in a single round of interaction. Additionally,  $E$  (standing for *empty*) is used to refer to the game state having had no interaction yet. The most basic element in a plan corpus is corpus actions, having the following representation.

**Definition 5.2.1 (Corpus Action)** *An action in a plan corpus is of the form  $s_1 \dots s_M \xi$ , where  $s_i \in \{E, R, T, S, P\}$ ,  $1 \leq i \leq M$ , are the states of the  $M$  last interactions, and  $\xi \in \{C, D\}$  is the current move<sup>1</sup>.*

**Definition 5.2.2 (Plan Session)** *A plan session of a strategy is a sequence of corpus actions played by that strategy (more precisely, a player using that strategy) against an arbitrary player.*

We denote by  $\Sigma_M$  the set of all possible types of action. Clearly,  $|\Sigma_M| = 2 \times 5^M$ . For example,

$$\Sigma_1 = \{EC, RC, TC, SC, PC, ED, RD, TD, SD, PD\}$$

This way of encoding actions and the assumption about the players' bounded memory size lead to the equivalent assumption that the action in the current round is *independent* of the ones in previous rounds, regardless of the memory size. The independence of actions will allow us to derive a convenient and efficient intention recognition model, discussed in the next subsection. Furthermore, it enables to save

<sup>1</sup>From now on, this notion of an action is used, which is different from the notion of a move (which is either C or D).

the game states without having to save the co-player’s moves, thus simplifying the representation of plan corpora.

As an example, let us consider *TFT* and the following sequence of its interactions with some other player (denoted by *X*), in the presence of noise

<b>round :</b>	0	1	2	3	4	5
<b>TFT :</b>	–	<i>C</i>	<i>C</i>	<i>D</i>	<i>D</i>	<i>D</i>
<b>X :</b>	–	<i>C</i>	<i>D</i>	<i>D</i>	<i>C</i>	<i>D</i>
<i>TFT-states :</i>	<i>E</i>	<i>R</i>	<i>S</i>	<i>P</i>	<i>T</i>	<i>P</i>

The corresponding plan session for *TFT* is [*EC, RC, SD, PD, TD*]. At the 0-th round, there is no interaction, thus the game state is *E*. *TFT* starts by cooperating (1-st round), hence the first action of the plan session is *EC*. Since player *X* also cooperates in the 1-st round, the game state at this round is *R*. *TFT* reciprocates in the 2-nd round by cooperating, hence the second action of the plan session is *RC*. Similarly for the third and the fourth actions. Now, at the 5-th round, *TFT* should cooperate since *X* cooperated in 4-th round, but because of noise, it makes an error to defect. Therefore, the 5-th action is *TD*.

**Definition 5.2.3 (Plan Corpus)** *Let  $\mathcal{S}$  be a set of strategies to be recognized. A plan corpus for  $\mathcal{S}$  is a set of plan sessions generated for each strategy in the set.*

For a given set of strategies, different plan corpora can be generated for different purposes. In Subsection 5.3.1, for example, we generate plan corpora for training and testing intention recognition models.

### 5.2.2 Corpus-based Intention Recognition Model

We could use any corpus-based intention recognition model in the literature for this work (see Related Work in Chapter 2). Here we deploy our corpus-based intention recognition model that was described in Chapter 2. For easy following, in the sequel we provide a simplified description of the intention recognition model.

Let  $I_i, 1 \leq i \leq n$ , be the intentions to be recognized, and  $O = \{A_1, \dots, A_m\}$  the set of current observed actions. The intention recognition task is to find the most

likely intention  $I^* \in \{I_1, \dots, I_n\}$  given the current observed actions, i.e.

$$\begin{aligned} I^* &= \arg \max_{I_i: 1 \leq i \leq n} P(I_i | A_1, \dots, A_m) \\ &= \arg \max_{I_i: 1 \leq i \leq n} \frac{P(I_i) \prod_{j=1}^m P(A_j | I_i, A_1, \dots, A_{j-1})}{P(O)} \end{aligned}$$

The second equation is obtained by applying Bayes' and then Chain rules (Pearl, 1988; Russell and Norvig, 2003). Since the denominator  $P(O)$  is a positive constant, we can ignore it. Then, because of the independency amongst actions, we obtain

$$I^* = \arg \max_{I_i: 1 \leq i \leq n} P(I_i) \prod_{j=1}^m P(A_j | I_i) \quad (5.1)$$

Note that this simplified expression is derived independently of the memory size  $M$ . All the probabilities needed for this computation are to be extracted beforehand using a training plan corpus. There is no update of these probabilities during intention recognizers' life cycle. These constants are arguably obtainable by evolutionary means via lineage (Richardson and Boyd, 2006; Shennan, 2002) – the corpus represents ancestors' received knowledge or acquired experience in interacting with known strategies (Woodward et al., 2009), or the 'like-me' self-experimenting framework (Meltzoff, 2007, 2005).

Also note that if two intentions are assessed with the same probability, then the model predicts the one with higher priority. Priorities of intentions are set depending on the behavioral attitude of the intention recognizer. For example, in Figure 5.2, if  $IR$ 's co-player cooperates in the first round, the co-player's intention can be predicted as either  $AllC$ ,  $WSLS$  or  $TFT$  (since they are assigned the same conditional probability values). On being concerned about  $TFT$ 's and  $WSLS$ 's retaliation after a defection (i.e.  $IR$ 's behavioral attitude),  $WSLS$  and  $TFT$  should have higher priorities than  $AllC$ .

### 5.2.3 Intention-based Decision Making Model

We describe how to acquire a decision making model for an intention recognizer from a training plan corpus. The intention recognizer chooses to play what would provide it with the greatest expected payoff against the recognized strategy (intention).

Namely, from training data we need to extract the function  $\theta(s, I)$ :

$$\theta : \{E, R, T, S, P\}^M \times \{I_1, \dots, I_n\} \rightarrow \{C, D\}$$

deciding what to play (C or D) given the sequence of  $M$  last game states  $s = s_1 \dots s_M$  ( $s_i \in \{E, R, T, S, P\}$  with  $1 \leq i \leq M$ ), and the recognized intention  $I \in \{I_1, \dots, I_n\}$ . It means that the intention recognizer needs to memorize no more than  $M$  last moves of its co-player, besides the fixed set of probabilistic constants in Eq. (5.1). It is done as follows. From the plan sessions in the training corpus for each intention we compute the (per-round) average payoff the intention recognizer would receive with respect to each choice (C or D), for each possible sequence of states  $s$ . The choice giving a greater payoff is chosen. Formally, let  $DS(I)$  be the set of all sequences of actions (plan sessions),  $Sq = A_1 \dots A_k$  ( $A_i \in \Sigma_M$ ,  $1 \leq i \leq k$ ), for intention  $I$  in the corpus and  $\pi(Sq, j)$  the payoff the intention recognizer would get at round  $j$ . In the following, if the sequence (plan session) in which the payoff being computed is clear from the context, we ignore it and simply write  $\pi(j)$ . Thus,

$$\theta(s, I) = \arg \max_{\xi \in \{C, D\}} \Pi_\xi / N_\xi \quad (5.2)$$

where

$$\Pi_\xi = \sum_{A_1 \dots A_k \in DS(I)} \sum_{\substack{1 \leq i \leq k \\ A_i = s\xi}} \sum_{j=i}^k \pi(j)$$

and

$$N_\xi = \sum_{A_1 \dots A_k \in DS(I)} \sum_{\substack{1 \leq i \leq k \\ A_i = s\xi}} \sum_{j=i}^k 1$$

A more general version of this decision making model is provided by considering a discount factor, say  $1/\alpha$  where  $0 < \alpha \leq 1$ , stating by how much distant rounds' payoffs are less important:

$$\theta(s, I) = \arg \max_{\xi \in \{C, D\}} \Pi_{\alpha\xi} / N_{\alpha\xi} \quad (5.3)$$

where

$$\Pi_{\alpha\xi} = \sum_{A_1 \dots A_k \in DS(I)} \sum_{\substack{1 \leq i \leq k \\ A_i = s\xi}} \sum_{j=i}^k \alpha^{j-i} \pi(j)$$

and

$$N_{\alpha\xi} = \sum_{A_1 \dots A_k \in DS(I)} \sum_{\substack{1 \leq i \leq k \\ A_i = s\xi}} \sum_{j=i}^k \alpha^{j-i}$$

The first model is a special case of the second one where  $\alpha = 1$ . This is the most future-optimistic case. On the other hand, if  $\alpha$  is small enough ( $\alpha \approx 0$ ),  $\theta(s, I) = D$ : in a given round of PD, it is always best to defect. This is the future-pessimistic case.

Note that, at the first round, there is no information about the co-player. The intention recognizer cooperates, i.e.  $\theta(E^M, I) = C \quad \forall I \in \{I_1, \dots, I_n\}$ .

## 5.3 Models Acquisition and Evaluation

### 5.3.1 Plan Corpus Generation

With the corpus description provided in Subsection 5.2.1, let us start by generating a plan corpus of four of the most famous strategies within the framework of repeated games of cooperation: *AllC* (always cooperate), *AllD* (always defect), *TFT* and *WSLS* (see above). Not only these strategies constitute the most used corpus of strategies used in this context, as most other strategies can be seen as a high-level composition of the principles enclosed in these strategies. Hence, intention recognizers map their opponent's behaviors to the closest strategy that they know and interact accordingly. When their knowledge is extended to incorporate new strategies, the models can be revised on the fly. However, this issue is beyond the scope of this thesis.

Further, note that here we are not trying to design any optimal strategy for the iterated PD as this is impossible to prove. Indeed, our aim is to show how the ubiquitous cognitive ability of intention recognition in humans (and many other species) (Woodward et al., 2009; Tomasello, 2008; Meltzoff, 2005) was conceivably adopted by natural selection, by using a simple AI intention recognition model. Although we could easily include other more complex strategies, such as those with

	<b>AllC</b>	<b>AllD</b>	<b>TFT</b>	<b>WSLS</b>	<b>Total</b>
<b>Precision</b>	0.859	<b>0.999</b>	0.818	0.579	<b>0.824</b>
<b>Recall</b>	0.859	0.999	0.818	0.575	0.824
<b>Converg.</b>	0.859	0.999	0.719	0.579	0.805

**Table 5.1:** Intention recognition results for each strategy and for the totality.

greater memory size and/or learning capability, and then let the intention recognizers learn to recognize them in a similar manner, by equipping the recognizers with memory capacity concomitant with that of its co-players'. But as noted before, that is not the goal of the present work.


We collect plan sessions of each strategy by playing a random move ( $C$  or  $D$ ) in each round with it. To be more thorough, we can also play all possible combinations for each given number of rounds  $m$ . For example, if  $m = 10$ , there will be  $1024 = 2^{10}$  combinations— $C$  or  $D$  in each round. When noise is present, each combination is played repeatedly several times.

The training corpus to be used here is generated by playing with each strategy all the possible combinations 20 times, for each number of rounds  $m$  from 5 to 10. The testing dataset is generated by playing a random move with each strategy in each round, also for  $m$  from 5 to 10. We continue until obtaining the same number of plan sessions as of the training dataset (corpus). Both datasets are generated in the presence of noise (namely, an intended move is wrongly performed with probability 0.05).

## 5.3.2 Intention Recognition Model

### 5.3.2.1 Evaluation Metrics

For evaluating the intention recognition model, we use three different metrics. *Precision* and *recall* report the number of correct predictions divided by total predictions and total prediction opportunities, respectively. If the intention recognizer always makes a prediction (whenever it has the opportunity), recall is equal to precision. *Convergence* is a metric that indicates how much time the recognizer took to converge on what the current user goal/intention was. Recall the formal definitions of the metrics in Chapter 2.



**Figure 5.1: Decision making model for different values of  $\alpha$ .** If the recognized intention is *Allc* or *AllD*, intention recognizers (IR) always defect, regardless of the current state. If it is *TFT*, IR cooperates when  $\alpha$  is large enough, regardless of the current state. If it is *WSLS*, if the current states are *S* or *T*, IR always defects; otherwise, IR cooperates for large enough  $\alpha$ . This model is acquired for a PD with  $R = 1, S = -1, T = 2, P = 0$ . The model has the same behavior for all PD payoff matrixes used in this chapter.

### 5.3.2.2 Results

The intention recognition model is acquired using the training corpus. Table 5.1 shows the recognition results of the model for the testing dataset, using the three metrics described above. We show the recognition result for each strategy, and for the whole dataset. Given that the training as well as the testing datasets are generated in the presence of noise, the achieved intention recognition performance is quite good. In the next section, we study the performance of players using this intention recognition model together with the intention-based decision making model (called *IR* players) in large scale population settings—particularly to address “What is the role of intention recognition for the emergence of cooperation?”

### 5.3.3 Decision Making Model

The decision making model (in Subsection 5.2.3) is acquired using the training corpus (Figure 5.1). We henceforth use  $\alpha = 1$ , i.e., independently of the PD payoff matrices (used in this chapter), if at the current round the co-player’s recognized intention is an unconditional one (*Allc* or *AllD*), IR always defects, regardless of the current game state; if it is *TFT*, IR always cooperates; and if it is *WSLS*, IR cooperates if and only if the current state is either *R* or *P*<sup>2</sup>. From here it is clear that, similar to

<sup>2</sup>For those who are interested in Logic Programming modeling (other readers can just ignore this footnote), in the following we provide a declarative representation of this IR strategy using

*TFT*, the strategy *IR* in this context bases its decision in the current round solely on the last move of others ( $M = 1$ ).

Figure 5.2 shows how an *IR* using these acquired intention recognition and intention-based decision making model interacts with other strategies, including *AllC*, *AllD*, *TFT*, *WSLS*, and another *IR*, in the absence of noise. Except with *AllD*, *IR* plays *C* in the first two rounds with other strategies: *IR* always plays *C* in the first round, and since others also play *C* (thus, the action is *EC*), they are predicted as a *TFT* (since  $P(EC|ALLC) = P(EC|TFT) = P(EC|WSLS) \gg P(EC|AllD)$ )—therefore, *IR* plays *C* in the second round. Note that here *TFT* is set with a higher priority than *WSLS*, which in turn has a higher priority than *AllC*. In the third round, these strategies are all predicted as *AllC* since they play *C* in the second round (and since  $P(RC|ALLC) > P(RC|WSLS) > P(RC|TFT)$ ). Hence, *IR* plays *D* in this round. The moves of these strategies (the other *IR* plays *D*, others play *C*) classifies *IR* to be *WSLS*, and the other three remain to be *AllC*, since  $P(RD|WSLS) > P(RD|TFT) \gg P(RD|AllC)$ . The two inequalities  $P(RC|WSLS) > P(RC|TFT)$  and  $P(RD|WSLS) > P(RD|TFT)$ , for big enough training corpus, are easily seen to hold: although *TFT* and *WSLS* equally likely play *C* (resp., *D*) after *R*, since *WSLS* corrects mistakes better than *TFT*, mutual cooperations are more frequent in plan sessions for *WSLS* in the training corpus. The reaction in the fourth round classifies *TFT* to be *TFT*, *IR* and *WSLS* to be *WSLS*, and *AllC* to be *AllC*; and like that in the subsequent rounds. From the

---

the intention-based decision making framework in Chapter 3.

1.  $abds([move/1])$ .
2.  $on\_observed(decide) \leftarrow new\_round$ .
3.  $decide \leftarrow move(c)$ .       $decide \leftarrow move(d)$ .  
      $\leftarrow move(c), move(d)$ .
4.  $expect(move(X))$ .
5.  $move(d) \triangleleft move(c) \leftarrow has\_intention(co\_player, allc)$ .  
      $move(d) \triangleleft move(c) \leftarrow has\_intention(co\_player, alld)$ .  
      $move(c) \triangleleft move(d) \leftarrow has\_intention(co\_player, tft)$   
      $move(c) \triangleleft move(d) \leftarrow has\_intention(co\_player, wsls), game\_state(s), (s = 'R'; s = 'P')$ .  
      $move(c) \triangleleft move(d) \leftarrow has\_intention(co\_player, wsls), game\_state(s), (s = 'T'; s = 'S')$ .

Starting a new interaction, an *IR* needs to choose a move, either cooperate (*C*) or defect (*D*) (lines 2-3). Both options are expected, and there are no expectation to the contrary. The *a priori* preferences in line 5 stating which move *IR* prefers to choose given the recognized intention of the co-player (*AllC*, *AllD*, *TFT* or *WSLS*) and the current game state (*'T'*, *'R'*, *'P'* or *'S'*). The built-in predicate *has\_intention/2* in the body of the preferences triggers the intention recognition module to validate if the co-player is most likely to follow a given intention (strategy) specified by the second argument.

fifth round on, *IR* cooperates with *WSLS*, *TFT* and another *IR*. If the number of rounds to be played is very large, up to some big round, these three strategies will be recognized as *AllC* again (since  $P(RC|ALLC) > P(RC|WSLS) > P(RC|TFT)$ ), then the process repeats as from third round. In our corpus, it only happens after more than 100 rounds. In playing with an *AllD*, *IR* cooperates in the first round, and defects in the remaining rounds, since  $P(ED|ALLD) \gg P(ED|I)$  for all  $I \in \{AllC, TFT, WSLS\}$  and furthermore  $P(s|ALLD) \gg P(s|I)$  for all  $I \in \{AllC, TFT, WSLS\}$  and  $s \in \{RD, SD, TD, PD\}$ .

## 5.4 Experiments and Results

In the following we provide analytical results under different evolutionary dynamics as well as using computer-based simulations. We show that the introduction of intention recognition promotes the emergence of cooperation in various settings, even in the presence of noise.

### 5.4.1 Analysis

To begin with, let us consider a population of *AllC*, *AllD*, and *IR* players. They play the iterated PD. Suppose  $m$  ( $m < 100$ ) is the average number of rounds. In the absence of noise, the payoff matrix of *AllC*, *AllD*, and *IR*, in  $m$  rounds, is given by (Figure 5.2,  $\alpha = 1$ )

$$\begin{array}{c}
 \begin{array}{ccc}
 & AllC & AllD & IR \\
 AllC & \left( \begin{array}{c} Rm \\ Tm \\ T(m-2) + 2R \end{array} \right. & \begin{array}{c} Sm \\ Pm \\ P(m-1) + S \end{array} & \begin{array}{c} 2R + S(m-2) \\ T + P(m-1) \\ R(m-1) + P \end{array} \\
 AllD & & & \\
 IR & & & 
 \end{array}
 \end{array}
 \right).$$

In each round *AllC* cooperates. Thus, its co-player would obtain a reward  $R$  if it cooperates and a temptation payoff  $T$  otherwise; hence, in playing with *AllC* (first column of the matrix), another *AllC* obtains  $m$  times of  $R$  since it cooperates in each round; *AllD* obtains  $m$  times of  $T$  since it defects in each round; and *IR* obtains 2 times of  $R$  and  $(m - 2)$  times of  $T$  since it cooperates with *AllC* in the first two rounds and defects in the remaining rounds (Figure 5.2). Other elements of the matrix are computed similarly.

<b>IR:</b> C D D D D ...	<b>IR:</b> C C D D D ...	<b>IR:</b> C C D C C ...
<b>AllD:</b> D D D D D ...	<b>AllC:</b> C C C C C ...	<b>IR:</b> C C D C C ...
<b>IR:</b> C C D D C C C C ...	<b>IR:</b> C C D D C C C C ...	
<b>TFT:</b> C C C D D C C C ...	<b>WSLS:</b> C C C D C C C C ...	

**Figure 5.2:** Interactions of *IR* with *AllC*, *AllD*, *TFT*, *WSLS* and another *IR*, in the absence of noise and  $\alpha = 1$ .

Pairwise comparisons (Hofbauer and Sigmund, 1998; Sigmund, 2010; Nowak, 2006b) of the three strategies lead to the conclusions that *AllC* is dominated by *IR* and that *IR* is an *evolutionary stable strategy* (Hofbauer and Sigmund, 1998) if

$$R(m-1) > T + P(m-2), \quad (5.4)$$

which always holds for  $m \geq 3$  (since  $2R > T + P$  and  $R > P$ ). Evolutionarily stable strategy is a strategy which, if adopted by a population of players, cannot be invaded by any alternative strategy that is initially rare (Hofbauer and Sigmund, 1998). This condition guarantees that once *IR* dominates the population, it becomes stable (for  $m \geq 3$ ).

Furthermore, one can show that

1. *IR* is *risk-dominant* (Hofbauer and Sigmund, 1998; Sigmund, 2010) against *AllD* if  $R(m-1) + S > P(m-1) + T$ , which is equivalent to

$$m > \frac{T + R - S - P}{R - P}. \quad (5.5)$$

For Donation game, it is equivalent to:  $m > 2b/(b-c)$ .

2. *IR* is *advantageous* (Hofbauer and Sigmund, 1998; Sigmund, 2010) against *AllD* if  $R(m-1) + 2S > T + Pm$ , which is equivalent to

$$m > \frac{T + R - 2S}{R - P}. \quad (5.6)$$

For Donation game, it is equivalent to:  $m > (2b+c)/(b-c)$ .

Since *IR* and *AllD* are both evolutionary stable strategies, Eq. (5.5) provides the condition for which *IR* has the greater basin of attraction (Nowak, 2006b; Sigmund,

2010). Eq. (5.6) provides the condition for which natural selection favors an *IR* to replace a population of *AllDs*, i.e. *IR* has a fixation probability greater than the neural one ( $1/N$ ) (Nowak, 2006b; Sigmund, 2010).

### 5.4.2 Evolutionary Simulations

In the presence of noise, it was not easy to provide an exact mathematical analysis. Instead, we will study this case using computer simulations. For convenience and a clear representation of simulation results, we perform our simulations using the Donation game (Sigmund, 2010), i.e.,  $T = b$ ,  $R = b - c$ ,  $P = 0$ ,  $S = -c$ , satisfying that  $b > c > 0$ .

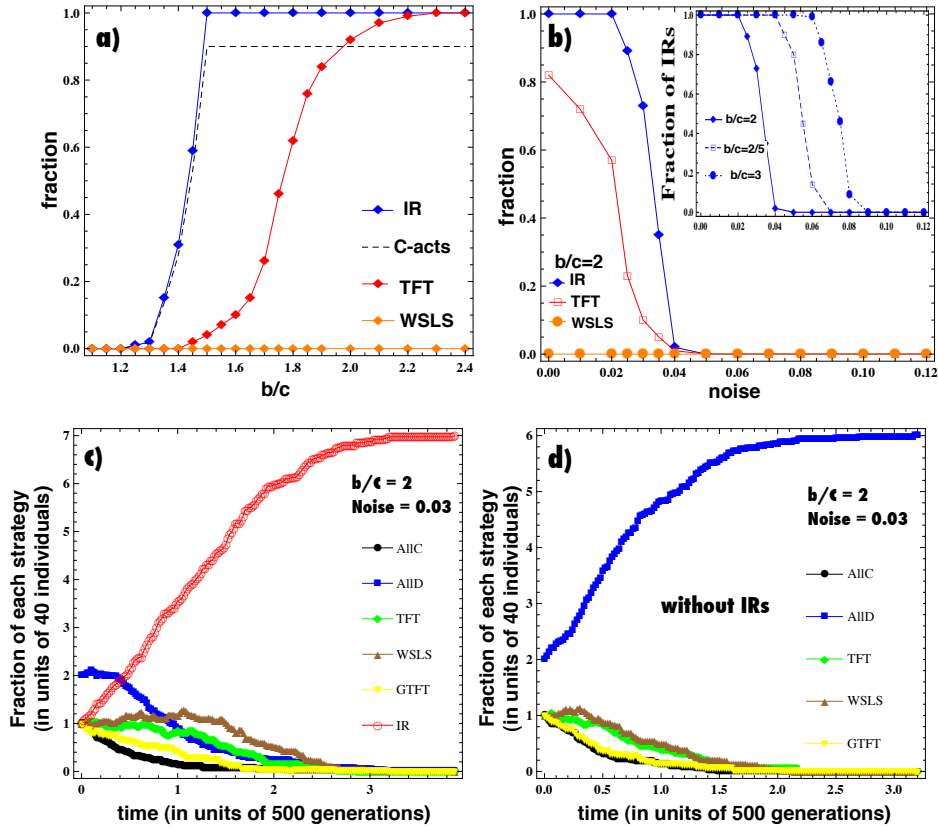
We start with a well-mixed population of size  $N$ , with individuals using different strategies. In each round of a generation, each individual interacts with all others, engaging in a PD game. The payoffs are accumulated over all the rounds. After each generation, an individual is randomly selected from the population, and will adopt the strategy of another randomly selected individual using the pairwise comparison rule (Szabó and Toke, 1998; Traulsen et al., 2006).

The results for some different settings are shown in Figure 5.3. Our results show that *IR* always prevails against other strategies, including *TFT*, *WSLS* and *GTFT*, for different benefit-to-cost ratios  $b/c$ , as well as more robust to noise. Namely, it has a strictly larger range of the benefit-to-cost ratio where cooperation can emerge, and can maintain it under a larger level of noise.

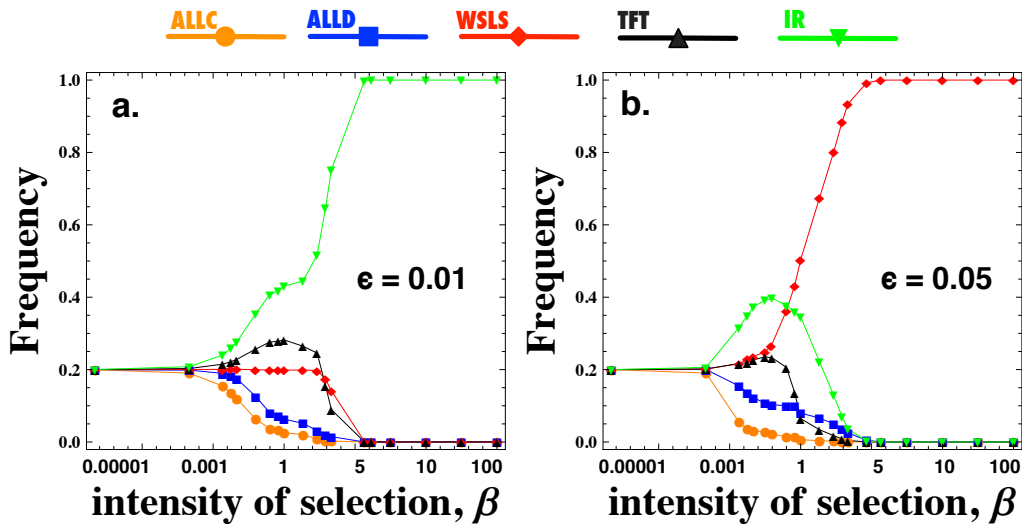
### 5.4.3 Intensities of Selection

We consider a setting where five strategies *AllC*, *AllD*, *TFT*, *WSLS*, and *IR* are present in the population. We compute numerically stationary distributions for variable intensity of selection  $\beta$  (Figure 5.4). The results show that, for small enough noise, the population always spends more time in the homogeneous state of *IR*, especially for strong intensities of selection (Figure 5.4a). When noise is large, *WSLS* wins for strong intensities of selection, but *IR* still wins for the slow ones (Figure 5.4b). Note that in case the intensity of selection  $\beta$  is very small, it is famously known in the literature as *weak selection* (Nowak et al., 2004; Traulsen et al., 2006; Wu et al., 2010). Weak selection describes the situation in which the effects of payoff differences are insignificant, hence the evolutionary dynamics are mainly driven by random fluctuations. This aspect is frequently observed in evolutionary biology

## 5.4. Experiments and Results



**Figure 5.3: Simulation results for Donation game.** In panels (a) and (b), we consider populations of three strategies,  $AUC$ ,  $AUD$  and either  $IR$ ,  $TFT$  or  $WSLS$ —equally distributed at the beginning. We plot the final fraction of  $IR$ ,  $TFT$  and  $WSLS$ . All simulations end up in a homogeneous state (i.e. having only one type of strategy) in less than 5000 generations. Our results show that  $IR$  prevails  $TFT$  and  $WSLS$  for different benefit-to-cost ratios  $b/c$  (panel a) and for different levels of noise (panel b). For a small ratio  $b/c$  (around 1.2),  $IR$  starts having winning opportunity, and from around 1.4 the population always converges to the homogeneous state of  $IR$ . For  $TFT$ , they are 1.4 and 2.1, respectively.  $WSLS$  has no chance to win for  $b/c \leq 2.4$ . The dashed black curve in (a) shows that the fraction of cooperation in the population of  $AUC$ ,  $AUD$  and  $IR$  is monotonic to  $b/c$ . In (b), our result shows that, in the presence of noise,  $IR$  outperforms  $TFT$  and  $WSLS$ . This result is robust to chances on the value of  $b/c$  (the inset of panel (b)). In panels (c) and (d), we consider a more complex setting where the population consists of several types of strategies:  $AUC$ ,  $AUD$ ,  $TFT$ ,  $WSLS$ ,  $GTFT$  (probability of forgiving a defect is 0.5) and  $IR$  (panel (c)) or without  $IR$  (panel (d)). Except for the defective  $AUD$  and  $IR$ , the other strategies are cooperative. Thus, instead of initially being equally distributed, we include a higher fraction of  $AUD$ s in the initial population. Namely, each type has 40 individuals, and  $AUD$  has 80.  $IR$  always wins (panel (c)). However, if  $IR$  individuals are removed,  $AUD$  is the winner (panel (d)), showing how  $IR$ s work as a catalyzer for cooperation. We have tested and obtained similar results for larger population sizes. Finally, in (a) and (b) we show how  $WSLS$  performs badly, as  $WSLS$  needs  $TFT$ s as a catalyst to perform well (Sigmund, 2010)—which can be observed in panels (c) and (d). All results were obtained averaging over 100 runs, for  $m = 10$ ,  $N = 100$  and  $\beta = 0.1$ .



**Figure 5.4: Stationary distribution (in percentage) of each strategy depending on the intensity of selection  $\beta$ .** The population consists of five strategies *ALLC*, *ALLD*, *TFT*, *WSLS*, and *IR*. For small values of  $\beta$ , selection is nearly neutral. Strategy updating is mostly random and frequencies of all strategies are roughly equal. Discrimination between strategies occurs when  $\beta$  increases. (a) When noise is small, IR always wins; (b) When noise is large, IRs wins for small  $\beta$ ; WSLs wins when  $\beta$  is large. All calculations are made with  $b/c = 3$ ,  $m = 10$ ,  $N = 100$ . When noise is present, the average payoffs of each strategy are obtained by averaging  $10^7$  runs.

– a phenotype is often found to be slightly advantageous over another phenotype (Akashi, 1995; Charlesworth and Eyre-Walker, 2007), as well as in human updating strategy in behavioral games (Traulsen et al., 2010).

## 5.5 Discussion – two models of intention recognition for the evolution of cooperation

There are various ways in which the capacity for intention recognition can confer on its holder an evolutionary advantage. In the previous chapter, the intention recognizers assess their co-players' cooperativeness, while in this chapter they recognize the co-players' strategies or patterns of behavior. Both can be found pervasive in real human interactions and communications. When looking for a collaborator/partner, one usually assesses how he/she was cooperative with us in the past, in order to

decide if one should cooperate with him/her in current affairs.

In more complex situations, one might need to assess more complicated patterns of behavior, as was done by the model in this chapter. The patterns could be those of social behavior prototypically represented by the strategies in social dilemmas, but they could be more complicated, e.g. the strategic tactics in military settings (Heinze, 2003; Sunzi and Sawyer, 1994). To enable such more complicated assessments, it is natural that the intention recognizers have more prior experience, especially in interacting with those patterns of behavior (Simon, 1957, 1990; Arthur, 1994; Todd, 2001). The main advantage resulting from this additional complication is that herein the intention recognizer can, to some degree, predict future behavior of the recognized strategist and take advantage from the information. For instance, a pure cooperator *AIC* would always cooperate, hence, the intention recognizer can exploit (i.e. defect with) the *AIC* without having to worry about its future retaliation. On the other hand, knowing that conditional cooperators like *TFT*s would retaliate after a defection, the intention recognizer is cooperative towards them. Differently, the intention recognizers from the previous model cannot distinguish between pure cooperators and conditional cooperators, hence they are cooperative towards them.

Technically, both models have been implemented by supplying Bayesian network models with appropriate characterizations. One might even design a three-layer Bayesian network, similar to the one of the previous chapter, to recognize a co-player's strategy; however, a significant role of the prior acquired plan corpus or experience is to learn a (best-response) intention-based decision making model. Differently, the intention recognizers in the previous chapter resorted instead to a heuristic generalizing the behavior of tit-for-tat, namely, cooperating if the co-player is more likely to cooperate and defecting otherwise.

In both models, the intention recognizers are restricted to exploiting only the direct past interactions with the co-players. However, it is obvious that indirect past observed interactions and other environmental relevant factors (such as noise and population structure) should play an important role in the intention recognition process. We envisage in the future to extend these models to integrate other relevant factors to improve the intention recognition models.

Moreover, when considering intention recognition, or any cognitive capacities in general, one might think of a need to take into account the cost required for

performing the cognition (for instance to gather and memorize more relevant factors necessary for the performance). Our additional analysis shows that the results obtained from both intention recognition models are robust with respect to the cognitive cost, assuming it to be rather small compared to the cost of cooperation (i.e., compared to the parameter  $c$  in the Donation game).

## 5.6 Concluding Remarks

Using the tools of evolutionary game theory, we have addressed explicitly the role played by intention recognition in the evolution of cooperation. We have shown, analytically as well as experimentally, that intention recognition strongly promotes the emergence of cooperation and maintains it well. Given the broad spectrum of problems which are addressed using this cooperative metaphor, our result indicates how intention recognition can be pivotal in social dynamics. Individuals which are equipped with an ability to recognize intention of others, or intention recognizers, can quickly recognize the defectors *AllDs*, thus not being exploited by them as for the *WSLS* strategy. For their own benefit, the intention recognizers can exploit the unconditional cooperators, thus do not suffer random drift as *TFT* and *GTFT*. Furthermore, the intention recognizers are cooperative with the conditional cooperators, including *TFT*, *WSLS*, and players alike themselves.

We have shown that a population with some initial fraction of intention recognizers acting selfishly to achieve greatest benefit can lead to a stable cooperation where intention recognizers come to prevail upon and permeate the population. The intention recognition strategy has a greater range of benefit-to-cost ratios leading to cooperation than the most successful existent strategies, including *TFT*, *WSLS*. We have also shown that it is more robust to noise, as well as does prevail under a various range of intensities of selection.

In addition, our approach of using a plan corpus makes a case for different other AI techniques (Russell and Norvig, 2003) to work with the problem of cooperation. In this work, we studied the role of intention recognition for the emergence of cooperation, but other cognitive abilities are also of great interest and importance, for example pattern recognition algorithms. Classification algorithms (or supervised learning in general) are clearly a good candidate. Indeed, intention recognition can be considered as a classification problem: the sequence of observed actions is clas-

sified into a known strategy. In this regard, intention recognition is similar to the part-of-speech tagging problem (Halteren, van H., 1999) where words need to be probabilistically categorized into one of given tags (e.g. noun, verb, article, etc.): actions correspond to words, while intentions correspond to possible tags. Hence, naturally one can use techniques from this area to implement an intention recognition strategy in the EGT context (see, e.g. my work using a Neural Network model (Han and Marques, 2010)). Moreover, clustering algorithms (or unsupervised learning in general) can be used to categorize the sequences of actions that are not fit with the known strategies. This is a way to learn about unknown strategies, categorize them, revise the model to take them into account (and pass the revised model to the successors).

In short, our study provides new, clear evidence of the evolutionary advantages of having the cognitive skill of intention recognition to deal with problems where cooperation is beneficial, and suggests an approach to exploring other high cognitive abilities.



# Chapter 6

---

## THE EMERGENCE OF COMMITMENTS AND COOPERATION

---

*We may have all come on different ships,  
but we're in the same boat now.*

Martin Luther King Jr. (1929-1968).

**Abstract.** Agents make commitments towards others in order to influence others in a certain way, often by dismissing more profitable options. Most commitments depend on some incentive that is necessary to ensure that the action is in the agent's interest and thus, may be carried out to avoid eventual penalties. The capacity for using commitment strategies effectively is so important that natural selection may have shaped specialized capacities to make this possible. Evolutionary explanations for commitment, particularly its role in the evolution of cooperation, have been actively sought for and discussed in several fields, including Psychology and Philosophy. In this chapter, using the tools of evolutionary game theory, we provide a new model showing that individuals tend to engage in commitments, which leads to the emergence of cooperation

---

<sup>0</sup>This chapter is based on the publication (Han et al., 2012b).

even without assuming repeated interactions. The model is characterized by two key parameters: the punishment cost of failing commitment imposed on either side of a commitment, and the cost of managing the commitment deal. Our analytical results and extensive computer simulations show that cooperation can emerge if the punishment cost is large enough compared to the management cost.

## 6.1 Introduction

Over the last few years, several mechanisms have been pointed out to promote the emergence and maintenance of cooperation. From group and kin relations, memory and reputation based reciprocity mechanisms, to social diversity and context based reactions, grounded or not on incipient levels of cognition, there has been a large improvement on our capacity to understand the roots of animal and human cooperation (Axelrod and Hamilton, 1981; Axelrod, 1984; Hofbauer and Sigmund, 1998; Nowak, 2006b; Santos et al., 2008; Sigmund, 2010; Han et al., 2011a,b). They are certainly hugely important, but are they sufficient? Or, as many have suggested (Boehm, 1999; Hirshleifer, 1999; Humphrey, 1999), might there be other routes to social behavior that have been neglected? Certainly there are. Commitment, which amounts to expressing an intention rather than having it recognized, may stand as another route to cooperation, even in its simplest form, as we purport to show here.

Agents make commitments towards others when they give up options in order to influence others. Most commitments depend on some incentive that is necessary to ensure that the action is in the agent's interest and thus will be carried out (Gintis, 2001), on pain of some heavy penalty. Committers also incur in a small cost when proposing or setting up a commitment so as to make it credible upfront to others, and entice these to accept to commit.

The capacity for using commitment strategies effectively is so important that natural selection may have shaped specialized signaling capacities to make this possible (Skyrms, 2010; Robson, 1990; Santos et al., 2011; Nesse, 2001b; de Vos et al., 2001; Back and Flache, 2008). And it is believed to have an incidence on the emergence of morality (Ruse, 2001). Assuming cooperation to be, at best, just the result of individuals' purely competitive strategies can make it conceptually unstable (Oyama,

2000), most especially in non-iterated or history-free interactions. And it seems possible that the spread of simplistic notions, rooted in science, about the evolutionary origins of social relationships could foster a trend to make these relationships more conflicted, and society more brutal. An antidote is an evolutionary approach to behavior that incorporates a capacity for mutual commitment, shown advantageous for all concerned (Nesse, 2001b), even in non-iterated or memory-free settings.

Hence, our goal in this chapter is to examine, through Evolutionary Game Theory (EGT) (Hofbauer and Sigmund, 1998; Sigmund, 2010), how the most simple of commitment strategies work, and how they can give rise to the emergence of cooperation. We shall do so in the setting of the one-shot Prisoner's Dilemma (PD).

In a nutshell, convincing others of one's credibility in a commitment proposal amounts to submitting to options that change the incentives of the situation. These options, namely commitment cost and penalty for defaulting, can be expressed by the payoffs specified in a game. When opponent players observe meticulously such payoffs, and realize that compliance with a proposed commitment is in the proposing player's best interests, then, given any opponent player's open option to commit, these may change their expectations and behavior accordingly, and adopt as a result a strategy which either accepts commitment proposals or ignores them.

In general, there are four main reasons to believe a commitment will be fulfilled (Nesse, 2001b): i) a commitment can be self-reinforcing if it is secured by incentives intrinsic to the situation; ii) a commitment can be secured by external incentives controlled by third parties; iii) a commitment can be backed by a pledge of reputation; and iv) a commitment can be reinforced by internal emotional motives.

The first two types are secured in much the same way a loan is secured by a collateral. They objectively change the situation so that fulfillment becomes in the individual's best interests. The latter two types do not change the objective contingencies; they are subjective commitments in that they may involve a continued option of renegeing, according to some or other stance extraneous to the game's given payoffs matrix.

In our EGT setting however, we will simply assume that a game's payoff matrix, concerning a set of strategies, summarily ingrains and expresses in its structure the impingement of all such contingencies. For instance, often a capacity for commitment allows individuals to act in ways that reap the benefits of image scoring through maintaining a reputation, or the access of others to a social history of prior

interactions. In this study, for simplicity but also for exhibiting the purity and power of the commitment mechanism, we ignore the effect of repeated interactions (Trivers, 1971), and of any reputation (Nowak and Sigmund, 2005; Pacheco et al., 2006) associated with particular individuals. We aim to show that the simplest of core commitment mechanisms can improve cooperation, and leave any other complications for the future, most promisingly how commitment can be combined with and reinforce other known mechanisms of cooperation. And perhaps surprisingly we can do so. Thus, no credibility of commitment will be taken into account (Castel-franchi and Falcone, 2010) beyond that which is expressed in a game’s payoff matrix. No reputation appraisal of the commitment proposer is made by its co-player, and no historical or social data is even available to do so. Each pairwise interaction is purely based on fixed individual strategies that might involve commitment or the lack thereof. Also, no “cheater or deceit detection” or “intention recognition” is in place (Han et al., 2011a,b). Nevertheless, systematic unconditional bluffing on the part of a player is a possible fixed feature of its strategy, in the sense that, from the start, the player does not intend to fulfill commitments.

It will be seen in our commitment model that players defaulting on their commitments, be they the proposing or the accepting party, will be subject to evolutionary disadvantage for a wide range of parameters.

We show that more elaborate commitment strategies are not strictly necessary for commitment to become evolutionarily advantageous. Neither an aptitude for higher cognition, nor for empathy, nor for mind reading are needed. These aptitudes would only be required for more sophisticated forms of commitment, scaffolded atop the core one. We will explain the evolution, in a population, of the capacity for a simple form of commitment as the result of otherwise being excluded from a group of committed promise abiding cooperators, in the sense that this strategy tends to invade the game playing population under rather general conditions.

The remainder of this chapter is organized as follows. In Section 6.2, we discuss the relevant literature. In Section 6.3, our EGT commitment model and its methods are defined and explained. Forthwith, in Section 6.4, we proffer results obtained with the model, both analytic and via numeric and computer simulations. We conclude the chapter with a discussion section on commitment and its EGT modeling.

## 6.2 Related Work

Evolution of cooperation has been a central research topic of many fields, including Biology, Economics, Artificial Intelligence, Political Science and Psychology (Axelrod, 1984; Hofbauer and Sigmund, 1998; Nowak, 2006b; Santos et al., 2008; Sigmund, 2010; Han et al., 2011b; Hofmann et al., 2011). Several mechanisms responsible for promoting cooperative behavior have been recently identified (see surveys in (Nowak, 2006b; Sigmund, 2010)). In these contexts, several aspects have been shown to play an important role in the emergence of cooperation. Differently, our model does not require any of those aspects, namely it does not assume kinship or in-group relatedness of agents, nor repeated interactions or reputation consideration, nor concrete structures of population distribution. However, we envisage that the mechanism of commitment could reinforce the existing mechanisms of cooperation, e.g., easing the conditions for the emergence of cooperation therein, or leading to high levels of cooperation under otherwise identical conditions. This will be the subject of the future work.

Evolutionary explanations of commitment, particularly its role in the evolution of cooperation, have been actively sought for and discussed in several fields, including Psychology and Philosophy (Nesse, 2001b; Gintis, 2001; Hirshleifer, 2001; Castelfranchi and Falcone, 2010; de Vos et al., 2001; Back and Flache, 2008; Ruse, 2001). But there are only a few computational models that show the evolutionary advantages of commitment in problems where cooperative acts are beneficial (Skyrms, 2010; Robson, 1990; Santos et al., 2011). In addition, often models rely on repeated interactions or long-term relationships (de Vos et al., 2001; Back and Flache, 2008), alike the conditions where Triver’s direct reciprocity (Trivers, 1971) may play a role. Here we provide an analytic model in the framework of evolutionary game theory showing that, with the availability of the mechanism of commitment, cooperation can emerge even without assuming repeated interactions.

Last but not least, it is undoubtedly important to mention the extensive literature of AI and Multi-agent System research on commitment, e.g., (Schelling, 1990; Wooldridge and Jennings, 1999; Hirshleifer, 2001; Castelfranchi and Falcone, 2010; Winikoff, 2007; Harrenstein et al., 2007; Chopra and Singh, 2009). The main concern therein is how to formalize different aspects of commitment and how a commitment mechanism can be implemented in multi-agent interactions to enhance them (e.g. for improved collaborative problem solving (Wooldridge and Jennings, 1999)), espe-

cially in the context of game theory. In contradistinction, our concern is in the nature of an evolutionary explanation of commitment, particularly how it can promote the emergence of cooperation.

## 6.3 Models and Methods

### 6.3.1 Model

Let us consider a commitment variant of the Prisoner's Dilemma game in which a new type of cooperator (denoted by COM\_C) that, before each interaction, asks the co-player whether it commits to cooperate. If the co-player does not so commit, there is no interaction. Both players get 0. Otherwise, if the co-player commits, they then go on to play with each other in the present interaction. If the co-player keeps to its commitment, both players obtain the reward payoff,  $R$ <sup>1</sup>. Otherwise (if the co-player fails its commitment), the proposing or focal player obtains the sucker payoff,  $S$ , and its co-player obtains the temptation payoff,  $T$ . However, the one that fails the commitment will suffer a penalty cost, and its non-defaulting co-player gains a compensation for the potential loss due to its default of fulfilling the commitment. For simplicity, we assume that these two amounts (penalty and compensation) are equal, being denoted by  $\delta$ . The penalty cost can be a real monetary one, e.g., in the form of prior debit (e.g., in the case of accommodation rental) or of a subsequent punishment cost (e.g., commitment was performed in terms of a legal contract, and one who fails commitment must pay a cost to compensate for the other), or an imaginary abstract value, e.g., public spread of good/bad reputation (bad reputation for the one that fails, and sympathy for the other), or even an emotional suffering (Nesse, 2001b; Gintis, 2001; Hirshleifer, 2001; Ruse, 2001). How this cost is set up depends on the types of commitment at work, or the reason for which the commitment is believed to be fulfilled (see Introduction), which topic is beyond the scope of this thesis. However, various techniques can be seen in (Schelling, 1990; Hirshleifer, 2001).

Two players that defect in an interaction obtain the punishment payoff,  $P$ <sup>2</sup>. As

---

<sup>1</sup>Note that here we do not yet take into account any execution noise (see, e.g., (Sigmund, 2010; Pacheco et al., 2006)), that is, under this noise the agents might mis-implement their intended choice, from cooperate to defect or vice versa. Thus, COM\_C will never mis-implement the intended commitment choice. Nonetheless, it is noteworthy that our additional analysis shows that the results obtained in the sequel are robust to a rather high level of execution noise.

<sup>2</sup>For the sake of a clear representation, in our analysis we adopt  $P = 0$  (Santos et al., 2006b;

usual, for the Prisoner's Dilemma, the payoff entries satisfy the ordering,  $T > R > P > S$ , whereas the four possible outcomes can be written down as a payoff matrix

$$\begin{array}{c} C \quad D \\ C \left( \begin{array}{cc} R, R & S, T \\ T, S & P, P \end{array} \right) \\ D \end{array}$$

For setting up a commitment, the proposer must pay a small management cost,  $\epsilon$ . The cost of proposing and setting up the commitment might be high, but it is reasonable to assume that this cost is quite small compared to the mutual benefit of a cooperation strategy guaranteeing commitment,  $\epsilon \ll R$ .

Given the nature of a situation expressed in terms of payoff entries, one can naturally expect that if a proposed punishment cost,  $\delta$ , is high enough compared to the cost of managing the commitment,  $\epsilon$  – to convince and guarantee that cooperation is in the proposer's interest and also drive away potential exploiters – cooperation can emerge, even in the fierce environment of the Prisoner's Dilemma. This penalty and management relation is subject to detailed study below, both analytically and by means of computer simulations.

We consider a finite population of a constant size, consisting of four strategies: COM\_C (as described above), C (always cooperates, without proposing to commit), D (always defects, and does not commit when being asked to), and D\_COM (always defects, though commits when being asked to). Here, we assume that cooperators, including COM\_C and C players, always commit whenever being asked to since they are better off to do so, as cooperation is their default choice, and reasonable commitment deals only are proposed. Hence, for the sake of exposition, the two other (unreasonable) strategies, those of cooperators that refuse to commit and of defectors that propose commitment, are omitted here (they would become eliminated anyway). The former is dominated by the pure cooperator strategy, C, while the latter is by the pure defector strategy, D.

In each round, two random players are chosen from the population for an inter-  
[Sigmund, 2010](#)) (as in the Donation game), even if the more general case can be analyzed in the same manner and portray similar results to the ones presented below.

action. For the row player, the (average) payoff matrix reads

$$\begin{array}{c}
 \begin{array}{c}
 COMC \\
 C \\
 D \\
 DCOM
 \end{array}
 \begin{array}{c}
 COMC \\
 C \\
 D \\
 DCOM
 \end{array}
 \begin{array}{c}
 C \\
 D \\
 C \\
 D
 \end{array}
 \begin{array}{c}
 D \\
 C \\
 D \\
 C
 \end{array}
 \begin{array}{c}
 DCOM \\
 C \\
 D \\
 COMC
 \end{array}
 \end{array}
 \begin{pmatrix}
 R - \epsilon/2 & R - \epsilon & -\epsilon & S + \delta - \epsilon \\
 R & R & S & S \\
 0 & T & P & P \\
 T - \delta & T & P & P
 \end{pmatrix}. \quad (6.1)$$

Note that when a COM\_C interacts with another COM\_C, only one of them pays the cost of having proposed commitment,  $\epsilon$  (e.g., the arbitrary one that proposes). Therefore, the average payoff of a COM\_C in playing with another COM\_C is,  $R - \epsilon/2$ .

### 6.3.2 Methods

Our analysis is based on evolutionary game theory methods for finite populations (Nowak et al., 2004; Imhof et al., 2005) (see [EGT methods](#)).

Now let us recall some important analytic measures which will be used in our analytical study. In a pair-wise comparison of strategy A with strategy B, we say that A is advantageous (against B) if an A mutant has a fixation probability in a population of agents using B greater than that of the neutral selection (which equals the inverse of population size,  $1/N$ ) (Nowak et al., 2004; Nowak, 2006b; Sigmund, 2010). Interestingly, it was shown that this condition holds if


$$(N - 2)\pi_{A,A} + (2N - 1)\pi_{A,B} > (N + 1)\pi_{B,A} + (2N - 4)\pi_{B,B}, \quad (6.2)$$

which, in the limit of large  $N$ , is simplified to

$$\pi_{A,A} + 2\pi_{A,B} > \pi_{B,A} + 2\pi_{B,B}. \quad (6.3)$$

Another important measure to compare the two strategies A and B is which direction the transition is stronger or more probable, an A mutant fixating in a population of agents using B or a B mutant fixating in the population of agents using A. It can be shown that the former is stronger if (Kandori et al., 1993; Sigmund, 2010)

$$(N - 2)\pi_{A,A} + N\pi_{A,B} > (N - 2)\pi_{B,A} + N\pi_{B,B}, \quad (6.4)$$



**Figure 6.1: Stationary distribution and fixation probabilities.** The population spends most of the time in the homogenous state of COM\_C. The black arrows stand for the transitions that are rather stronger than neutral. The strongest transitions are from C to D and D\_COM, and the slowest one is from COM\_C to C. There are rather strong transitions from D and D\_COM to COM\_C. Parameters:  $T = 2$ ,  $R = 1$ ,  $P = 0$ ,  $S = -1$ ;  $\delta = 4$ ;  $\epsilon = 0.05$ ; imitation strength,  $\beta = 1$ ; population size,  $N = 100$ ;  $\mu_N = 1/N$  denotes the neutral fixation probability.

which, in the limit of large  $N$ , is simplified to

$$\pi_{A,A} + \pi_{A,B} > \pi_{B,A} + \pi_{B,B}. \quad (6.5)$$

## 6.4 Results

We compute the fixation probabilities and stationary distribution numerically for small mutation or exploration rates (see [EGT Methods](#)). The population spends most of the time in the homogeneous state where all individuals utilize the commitment strategy (Figure [6.1](#)).

In general, amongst the monomorphic states of the population, the strongest transitions are from C to D and C to D\_COM. The difference of a small cost of proposing commitment,  $\epsilon$ , between COM\_C and C, leads to a near-neutral transition from COM\_C to C. The more intricate transitions are between COM\_C and D or D\_COM, which are the central part of our analysis.

Between D and COM\_C, for  $\epsilon \ll R$ , COM\_C is advantageous. Namely, by a pairwise comparisons of COM\_C and D ([Nowak et al., 2004](#); [Nowak, 2006b](#)) that condition always holds if (in the limit of large  $N$ , see Eq. [\(6.3\)](#))

$$\epsilon < \frac{2R}{5}. \quad (6.6)$$

This inequality also guarantees that, for a population of size  $N > 4$ , the more probable transition is from D to COM\_C, i.e., satisfying that ([Kandori et al., 1993](#); [Sigmund, 2010](#)) (see Eq. [6.5](#))

$$(N - 2)\left(R - \frac{\epsilon}{2}\right) - N\epsilon > 0. \quad (6.7)$$

Similarly, for big enough  $\delta$ , COM\_C is advantageous against D\_COM; namely, if

$$\delta > \frac{T - R - 2S}{3} + \frac{5\epsilon}{6}. \quad (6.8)$$

It guarantees that the transition of D\_COM to COM\_C is more probable than the opposite if

$$(N - 2)\left(R - \frac{\epsilon}{2}\right) + N(S + \delta - \epsilon) > (N - 2)(T - \delta), \quad (6.9)$$

which holds if

$$\delta > \frac{N}{2N-2}(T-R-S + \frac{3\epsilon}{2}). \quad (6.10)$$

For large enough  $N$ , it is simplified to

$$\delta > \frac{T-R-S}{2} + \frac{3\epsilon}{4}. \quad (6.11)$$

Hence, for


$$\delta > \max\left\{\frac{T-R-S}{2} + \frac{3\epsilon}{4}, \frac{T-R-2S}{3} + \frac{5\epsilon}{6}\right\} \quad (6.12)$$

the transition of D\_COM to COM\_C is the more probable one, as well as greater than neutral.


Taking together with the fact that the transition of COM\_C to C is near neutral, one can expect that if the two parameters  $\delta$  and  $\epsilon$  satisfy the inequalities (6.6) and (6.12), COM\_C will prevail – the population will spend most of the time in its homogenous state. This expectation is supported by the numerical results in Figures 6.2 and 6.3. For a given payoff matrix of the PD, for strong enough punishment cost of failing commitment,  $\delta$ , and small enough cost of setting up the commitment,  $\epsilon$ , the population spends most of the time in the homogeneous state of COM\_C (Figure 6.2). In addition, this result is also flexible with respect to the payoff values of the PD (Figure 6.3). For the sake of a clear representation of the result, we use in this numerical experiment the Donation game (Sigmund et al., 2010) – a special case of PD – where  $T = b$ ,  $R = b - c$ ,  $P = 0$ ,  $S = -c$ , satisfying that  $b > c > 0$ ;  $b$  and  $c$  stand for “benefit” and “cost” of cooperation, respectively. It shows that, for given  $\delta$  and  $\epsilon$ , for large enough  $b/c$ , the population spends most of the time in the homogeneous state of COM\_C.

So far, our analytic and numerical results were obtained in the limit of small mutation rates. Next, by extensive computer simulations, we show that this remarkable performance of the commitment strategy COM\_C is flexible with respect to mutation rates (Figure 6.4). Namely, for all the mutation rates up to 0.1, the population always spends most of the time in the homogenous state of COM\_C. It is also noteworthy, that our analytic results for small imitation strengths and under the extremes of low and high mutation or exploration rates — based on the methods described in (Antal et al., 2009) — comply with this simulation results.


All in all, our study exhibits that, in spite of the absence of repeated interactions, reputation effect, network reciprocity, as well as group and kin selection, the strategy



**Figure 6.2:** Contour plot of the frequency of COM\_C as a function of  $\epsilon$  and  $\delta$ . In a population of COM\_C, COM\_D, C, and D individuals, for a wide range of  $\epsilon$  and  $\delta$ , the population spends most of the time in the homogeneous state of COM\_C. The smaller the cost of proposing commitment,  $\epsilon$ , and the greater the punishment cost of failing commitment,  $\delta$ , the greater the frequency of COM\_C. The payoffs being used are,  $T = 2$ ,  $R = 1$ ,  $P = 0$ ,  $S = -1$ ; imitation strength,  $\beta = 1$ ; population size,  $N = 100$ .



**Figure 6.3:** Frequency of each strategy as a function of benefit-to-cost ratio,  $b/c$ , for Donation game ( $T = b$ ,  $R = b - c$ ,  $P = 0$ ,  $S = -c$ , with  $b \geq c$ ). In a population of COM\_C, COM\_D, C, and D individuals, for a large enough benefit-to-cost ratio, the population spends most of the time in the homogeneous state of COM\_C, while D prevails when this ratio is very small. Parameters:  $\delta = 4$ ;  $\epsilon = 0.05$ ; imitation strength,  $\beta = 1$ ; population size,  $N = 100$ .



**Figure 6.4: Frequency as a function of mutation rates.** Symbols indicate results from computer simulations (averaged over  $10^9$  update steps), and solid lines show the numerical approximation results for small mutation rates. The population consists of COM\_C, COM\_D, C, and D individuals. COM\_C always dominates for all the mutation rates up to 0.1. Parameter values:  $T = 2$ ,  $R = 1$ ,  $P = 0$ ,  $S = -1$ ;  $\delta = 4$ ;  $\epsilon = 0.05$ ; imitation strength,  $\beta = 1.0$ ; population size,  $N = 100$ . The simulation results are obtained by averaging 40 runs, and the initial population is equally likely to be in one of the homogenous states.

of commitment proposal may enable the emergence of cooperation. By imposing a high cost for failing a commitment, when compared to the cost of setting up or managing the commitment deal, the commitment cooperative agents COM\_C can get rid of the fake committers (D\_COM) as well as avoid being exploited by the pure defectors (D), while playing approximately equally well against the pure cooperators (C). The results of this study suggest that our specialized capacity for commitment, which might have been shaped by natural selection (Nesse, 2001b) consists in a capacity for managing to impose a high cost of punishment, whether it is monetary or of abstract emotional or reputation value, with a relatively small cost.

We note that there is a significant difference between our commitment model and the works on costly punishment (Ohtsuki et al., 2009; Rand and Nowak, 2011; Dreber et al., 2008; Hauert et al., 2007; Fehr and Gächter, 2002). A commitment deal must be agreed by both sides of it in advance, thereby giving credibility and justification to punish any defaulting player. In addition, the prior agreement gives rise to compensation—the amount of which, in some cases, is agreed explicitly in advance—to the non-defaulting player. This compensation for the non-defaulting player is the significant difference that makes successful those players using the commitment strategy, while those using the costly punishment strategy have only a narrow margin of efficiency (Ohtsuki et al., 2009); does not stand out as a winning strategy (Dreber et al., 2008); nor does it promote cooperation at all when taking into account antisocial punishment (Rand and Nowak, 2011; Herrmann et al., 2008). The compensation might bring benefit to the commitment strategists once an appropriate deal would be arranged.

This suggests that although costly punishment, whether it is social or antisocial, might not promote the evolution of cooperation, what we call ‘justified’ punishment, which is warranted by an appropriate commitment deal, does. This kind of punishment might not be costly at all, and can even bring net benefit to its upholder, hence leading to the emergence of cooperation.

We also note that in our model D\_COM plays worse than D when a strong commitment deal is in place. The commitment proposers COM\_C can even be better off when playing with the fake committers D\_COM than when playing with the pure defectors D, because of the big compensation they can earn from the strong deal. This might lead to a new kind of commitment proposal cheaters, who would make efforts to cheat their co-players into a strong commitment deal with particularly high

punishment cost, and then later deliberately attempt to make them unable to fulfill the commitment deal (perhaps in a secret way to avoid retaliation), hence obtaining the compensation benefit. This kind of behavior might evolve in situations wherein failing a commitment is very costly (for instance destroying opponents' reputation in an election). This new aspect of commitment separates it from costly punishment models in the literature, opening the door to the study of deceit within commitment (Trivers, 2011).

## 6.5 Discussion

Within the general game theory concept of commitment, several distinctions can help separate different subtypes. In particular, some commitments are upfront promises of a next move that can help, while others are upfront threats of a subsequent move that can harm. Commitments can be conditional or unconditional. Threats are usually attempts to influence another person's next move by stating a conditional subsequent move, and that's how we may envisage them. Promises are more likely to be unconditional, and that's how we may conceive of them, though more generally they can be conditional on the other fulfilling a matching promise. Concerning this, we note a difference between a commitment and a convention. A convention is a means for monitoring a commitment: it specifies under what circumstances a commitment can be abandoned and how an agent should behave both locally and towards others when one of these conditions arises (Wooldridge and Jennings, 1999). Commitments can also be just towards oneself, taking into account the evolution of possible futures afforded by actions and events, and the individual's prior and post preferences, in what might be classically seen as a game against nature.

In (Pereira and Han, 2009b,a), three different types of individual commitment – hard, revocable, and momentary – are studied in such an evolution context. Let us recall that commitment, in the context of game theory, is a device or mechanism to decide the outcome with the other party (Schelling, 1990). Schelling distinguishes between commitment pure and simple and commitment that takes the form of a threat. What he calls “ordinary” commitment corresponds, in game theory, to the making of an opening announcement in a sequential play, which we dub preemptive, just before both players make their actual move. To constitute a preemption, a player's announcement action must be irrevocable, that is a promise that is assuredly

kept. Preemptive commitment is not necessarily profitable, because it hinges on the opponent's actual move. Schelling however does not assume the other type of commitment as a "threat", which pertains to the player's move in reaction to the opponent's move. Threats, being conditional, may be of the "if-then-else" form, and can thus combine a threat and a promise, the latter albeit implicit whenever there are just two possible moves. We prefer instead to label "reactive" such so-called threat commitments. In the game context, these occur when the player with the last move irrevocably pledges to respond, in a specified but contingent way, to the opponent's prior choice (Hirshleiffer, 1999).

In a nutshell, some players can be "preemptive" committers – those that always propose and always accept proposed commitments –, others may be "reactive" committers – those that always make a "reactive" statement and comply with the implicit requests in such statements –, while other players, though accepting to commit nevertheless default on their commitment, and even others simply omit and ignore preemptive or reactive commitments in their strategies – they might for instance be persistent defectors or persistent cooperators as we have seen, or, for that matter, follow any other strategy ignorant of commitment. Moreover, in iterated games, commitments can concern future rounds and not just the present one.

We purport to have shown that a simple commitment abiding cooperative strategy can be evolutionarily advantageous even in a non-iterated game setting. But much remains to be explored. In the more general setting and to avoid confusion, it can be helpful to distinguish, even if only conceptually, between "execution moves" and "pre-play moves" (Hirshleiffer, 1999). The terms first move and last move then always refer exclusively to execution moves – the choices that actually generate the payoffs. In contrast, commitments come earlier with respect to execution moves: they are pre-play moves. A preemptive commitment is a pre-play move that allows the player making it to take the first execution move. A reactive commitment, although also a pre-play move, can be made only by the player who has the last execution move. In either case, by giving up on his or her choice through committing, the commitment player leaves the opponent with "the last clear chance to decide the outcome" (Schelling, 1990).

In our present game setting, however, there was no need to make the distinction between the first and the second to play, because each possible player strategy move is exhibited and fixed from the start, as expressed and ingrained in the payoff

matrix. By so introducing the several committed unconditional move strategies – though the payoff is of course conditional on the opponent’s move –, we can emulate what would happen in a round if a move sequence actually existed. Put briefly, our commitment model is of the simplest kind and, moreover, it is brought to bear solely on the very next move fold of a pair of players, with no history available on prior commitments. Nevertheless, it captures core features of commitment, namely the high cost of defaulting to discourage false commitment, and thus make it plausible, and a comparatively small but non-zero cost of commitment proposal to lend it initial credibility. On top of this core model more elaborate models affording commitment can subsequently be rooted, including those involving delayed deceit.

What’s more, commitment (or intention manifestation) and intention recognition, are but two sides of a coin really, and their joint study in the EGT setting is all but unavoidable. We will study the combination of these two mechanisms in the next chapter.

It has become increasingly obvious that maximizing reproductive success often requires keeping promises and fulfilling threats, even when that requires in turn sacrifices regarding individual short-term interests. That natural selection has shaped special mental capacities to make this possible seems likely, including a capacity for commitment (Nesse, 2001b) and for intention recognition (Han et al., 2011a,b). The commitment stance goes yet further, and many aspects of human groups seem shaped by effects of commitments and intention recognition, namely group boundaries, initiation rituals, ideologies, and signals of loyalty to the group (Skyrms, 1996, 2003, 2010). Conversely, many aspects of groups seem to exist largely to facilitate commitment to cooperate and to limit the utility of coercive threats.

The generalized ability for commitment to support cooperative interaction is an important aspect of plasticity in human behavior, and humans support their deal-making in lots of ways. The law is full of instances of people using techniques of commitment to establish the honesty of their intentions, namely through a variety of contracts (Goodenough, 2001). Institutions themselves are supported on committal contracts, and the law of the land proffers methods for constituting and of accountability of social institutions (Searle, 2010).

We believe studies of commitment will benefit greatly from rigorous models that allow for their analytical study and computer simulation, and in particular within the fold of EGT for the better to examine the emergence of complex social behavior.

# Chapter 7

---

## INTENTION RECOGNITION, COMMITMENT AND THE EVOLUTION OF COOPERATION

---

*Commitment is what transforms a promise into reality. It is the words that speak boldly of your intentions. And the actions which speak louder than the words. It is making the time when there is none. Coming through time after time after time, year after year after year. Commitment is the stuff character is made of; the power to change the face of things. It is the daily triumph of integrity over skepticism.*

Abraham Lincoln, 1809-1863

**Abstract.** Individuals make commitments towards others in order to influence others to behave in certain ways. Most commitments may depend on some incentive that is required to ensure that the action is in the agent's best interest and thus, should be carried out to avoid eventual penalties. Similarly, individuals may ground their decision on an accurate assessment of the intentions of others. Hence, both commitments and intention recognition go side by side in behavioral evolution. In this chapter, we analyze the role played by the co-evolution of intention

---

<sup>0</sup>This chapter is based on the publication (Han et al., 2012c).

recognition plus the emergence of commitments, in the framework of the evolution of cooperative behavior. We again resort to the tools of evolutionary game theory, showing how the combination of these two aspects of human behavior can enhance the emergent fraction of cooperative acts under a broad spectrum of configurations.

## 7.1 Introduction

Intention recognition is ubiquitous in many kinds of human interactions and communications, with much documented experimental evidence (Woodward et al., 2009; Meltzoff, 2007; Tomasello, 1999, 2008). It is so critical for human social functioning and the development of key human abilities, such as language and culture, that it might have been shaped by natural selection (Woodward et al., 2009; Tomasello, 2008; Skyrms, 1996; Han et al., 2011a,b) (see again Chapters 4 and 5). Technically, intention recognition can be defined as a process of becoming aware of the intentions or goals of another agent, inferring them through observed actions or their effects on the environment (Kautz and Allen, 1986b; Charniak and Goldman, 1993; Heinze, 2003; Pereira and Han, 2009c, 2011b).

Clearly there might be cases where it is difficult, if not impossible, to recognize the intentions of another agent. It might be your first interaction with someone in your life, and you have no information about him/her which can be used for intention recognition. You also might know someone well, but you still might have very little relevant information in a given situation to predict the intentions with high enough confidence. In such cases, the strategy of proposing a commitment, or manifesting an intention, can help to impose or clarify intentions of others <sup>1</sup>.

Moreover, agents make commitments towards others when they give up options in order to influence others. Most commitments depend on some incentive that is necessary to ensure that an action (or even an intention) is in the agent's interest and thus will be carried out in the future (Gintis, 2001). As previously, the capacity for using commitment strategies effectively is so important that natural selection may have shaped specialized capacities to make this possible (Skyrms, 2010; Robson,

---

<sup>1</sup>Intention is choice with commitment (Cohen and Levesque, 1990; Bratman, 1987; Roy, 2009b). Once an agent intends to do something, it must settle on some state of affairs for which to aim, because of its resource limitation and in order to coordinate its future actions. Deciding what to do established a form of commitment (Cohen and Levesque, 1990; Roy, 2009b). Proposing a commitment deal to another agent consists in asking it to express or clarify its intentions.

1990; Santos et al., 2011; Ruse, 2001; Frank, 2001; Nesse, 2001b; de Vos et al., 2001; Back and Flache, 2008).

One of the commitments we all know is marriage. By giving up the option to leave someone else, spouses gain security and an opportunity for a much deeper relationship that would be impossible otherwise (Nesse, 2001b; Frank, 2001), as it might be risky to assume a partner's intention of staying faithful without the commitment of marriage. A contract is another popular kind of commitment, e.g. for an apartment lease (Frank, 2001). When it is risky to assume another agent's intention of being cooperative, arranging an appropriate contract provides incentives for cooperation. However, for example in accommodation rental, a contract is not necessary when the cooperative intention is of high certainty, e.g. when the business affair is between close friends or relatives. It said, arranging a commitment deal can be useful to encourage cooperation whenever intention recognition is difficult, or cannot be performed with sufficiently high certainty. On the other hand, arranging commitments is not for free, and requires a specific capacity to set it up within a reasonable cost (for the agent to actually benefit from it) (Nesse, 2001b,a) — therefore it should be avoided when opportune to do so. In the case of marriage, partners sometimes choose to stay together without an official commitment when it might be too costly (e.g., it could be against parents' or families' wish, or it may need to be in secret because of their jobs) and/or they strongly trust each other's faithfulness (e.g., because of emotional attachment (Frank, 1988, 2001)). In short, a combination of the two strategies, those of commitment and of intention recognition, seems unavoidable. Nevertheless, intention recognition without actual commitment can be enhanced by costly engagement gifts, in support of sexual selection and attachment (Miller and Todd, 1998; Haselton and Buss, 2001). Furthermore, social emotions can act as ersatz commitment (Frank, 1988).

In this chapter, we start from the model of commitment formation presented in the previous chapter (also see (Han et al., 2012b)), characterized by two key parameters: a punishment cost of failing commitment imposed on either side of a commitment deal, and the cost of managing it. Therein we have shown that, if a strong enough commitment deal can be arranged, that is, with a small enough management cost and a large enough punishment cost, cooperation can emerge in a population of selfish agents.

On top of that model, also using the tools of Evolutionary Game Theory (EGT)

(Maynard-Smith, 1982; Hofbauer and Sigmund, 1998; Nowak, 2006a; Sigmund et al., 2010), we show that combining intention recognition and commitment strategies in a reasonable way can lead to the emergence of improved cooperation, not able to be achieved solely by either strategy. Our study seeks what is a reasonable combination of commitment and intention recognition. We shall do so in the setting of the one-shot Prisoner’s Dilemma (PD) (see [EGT methods](#)).

It will be seen from our model that, in most of the cases, there is a wide range of combination of the intention recognition and commitment strategies, which leads to a strategy that performs better than either strategy solely – in the sense that the population spends more time in the homogeneous state of agents using that strategy (Imhof et al., 2005; Hauert et al., 2007) (see also [EGT methods](#) for finite populations). Our results suggest that, if one can recognize intentions of others with high enough confidence or certainty, one should rely more on it, especially when it is difficult to reach to a conceivably strong commitment deal. It helps to avoid the unnecessary cost of arranging and managing the deal. That is, in a transparent world where people have nothing to hide from each other, contracts are unnecessary. On the other hand, when intention recognition with high precision is difficult (due to, e.g. environment noise, agents have great incentives to hide intentions, or there are not enough observed actions), one should rely more on the commitment strategy, particularly if a reasonable deal can be envisaged.

The remainder of this chapter is structured as follows. In Section 7.2, we introduce the model of commitment, on top of which we integrate the co-evolution of commitment and intention recognition. In Section 7.3, we provide analytical and computer simulations obtained from our model. In Section 7.4, some discussions on the implication of the results are provided. The chapter ends by proffering some concluding remarks.

## 7.2 Models

For ease of following the remainder of this final core chapter of the thesis, we first summarize the commitment variant of the Prisoner’s Dilemma presented in the previous chapter. We then describe our model for the combination of intention recognition and commitment.

### 7.2.1 Commitment variant of the Prisoner's Dilemma

Let us consider a commitment variant of the Prisoner's Dilemma game in which a new type of cooperator (denoted by COM\_C) that, before each interaction, asks the co-player whether it commits to cooperate. If the co-player does not so commit, there is no interaction. Both players get 0. Otherwise, if the co-player commits, they then go on to play with each other in the present interaction. If the co-player keeps to its commitment, both players obtain the reward payoff,  $R$ . Note that here we do not yet take into account execution noise (see, e.g., (Sigmund, 2010)), i.e. the agents might mis-implement their intended choice, from cooperate to defect or vice versa. Thus, COM\_C will never mis-implement the intended commitment choice, all the more so because commitment always entails an initial cost, thus being no point in proposing commitment when not intending to honor it. The payoffs of the commitment PD game, as we shall see, would make such bluffing players inevitably worse off. Otherwise (if the co-player fails its commitment), the proposing or focal player obtains the sucker payoff,  $S$ , and its co-player obtains the temptation payoff,  $T$ . However, the one that fails the commitment, whatever the player, will suffer a penalty cost for its non-defecting co-player to get a compensation. For simplicity, it is assumed that these two amounts are equal, denoted by  $\delta$ . This cost can be a real monetary one, e.g., in the form of prior debit (e.g., in the case of accommodation rental) or of a punishment cost (e.g., commitment was performed in terms of a legal contract, and the one who fails commitment must pay a penalty cost), or an imaginary abstract value, e.g., public spread of good/bad reputation (bad reputation for the one that fails, and sympathy for the other), or even emotional suffering (Nesse, 2001b; Gintis, 2001; Hirshleifer, 2001; Ruse, 2001). How this cost is set up depends on the types of commitment at work, or the reason for which the commitment is believed to become fulfilled (Schelling, 1990; Hirshleifer, 2001; Han et al., 2012b).

Two players that defect in an interaction obtain the punishment payoff,  $P$ . As usual, for the Prisoner's Dilemma, the payoff entries satisfy the ordering,  $T > R > P > S$ , whereas the four possible outcomes can be written down as a payoff matrix

$$\begin{array}{cc} & \begin{array}{cc} C & D \end{array} \\ \begin{array}{c} C \\ D \end{array} & \left( \begin{array}{cc} R, R & S, T \\ T, S & P, P \end{array} \right). \end{array}$$

For setting up a commitment, the proposer must pay a small cost,  $\epsilon$ . The cost of proposing and setting up the commitment might be high, but it is reasonable to assume that this cost is very small compared to the mutual benefit of a cooperation strategy guaranteeing commitment,  $\epsilon \ll R$ .

We consider a finite population of constant size, consisting of four strategies: COM\_C (as described above), C (always cooperates, without proposing a commitment deal), D (always defects, and does not commit when being asked to), and D\_COM (always defects, though commits when being asked to). In the previous chapter, we have shown analytically and by computer simulations that COM\_C dominates the population if the punishment cost  $\delta$  is large enough compared to the management cost  $\epsilon$ , thereby leading to the emergence of cooperation.

In each round, two random players are chosen from the population for an interaction. For the row player, the (average) payoff matrix reads

$$M_1 = \begin{array}{c} \\ COM\_C \\ C \\ D \\ D\_COM \end{array} \begin{array}{c} COM\_C \\ C \\ D \\ D\_COM \end{array} \begin{pmatrix} R - \epsilon/2 & R - \epsilon & -\epsilon & S + \delta - \epsilon \\ R & R & S & S \\ 0 & T & P & P \\ T - \delta & T & P & P \end{pmatrix}. \quad (7.1)$$

### 7.2.2 Combination of intention recognition and commitment

We provide a new strategy, IRCOM, which combines the two strategies, those of intention recognition and commitment. In an interaction, IRCOM recognizes its co-player's intention, which is either to cooperate or to defect, similarly to our model for the evolution of cooperation presented in Chapter 4 (see also (Han et al., 2011a)). A confidence level,  $cl$ , is assigned to the recognition result. It defines the degree of confidence (here in terms of probability) that IRCOM predicts the co-player's intention correctly <sup>2</sup>. In general,  $cl$  follows some probability distribution. As in a

<sup>2</sup>In AI, the problem of intention recognition has been paid attention for several decades, and the main stream is that of probabilistic approaches (Charniak and Goldman, 1993; Heinze, 2003; Bui et al., 2002; Blaylock and Allen, 2004; Armentano and Amandi, 2009). They tackle the problem by assigning probabilities to conceivable intentions (conditional on the current observations), based on which the intentions are ranked. Similarly to (Blaylock and Allen, 2004; Armentano and Amandi, 2009; Han and Pereira, 2011b), in our model, a degree of confidence in terms of a probability measure, is assigned to intentions. Additional discussions on probabilistic approaches to intention recognition are found in Chapter 2.

real intention recognition problem, the distribution should depend on the intention recognition method at work (how efficient it is), the environment IRCOM lives in (is it supportive for gathering relevant information for the recognition process, e.g. observability of co-players' direct and indirect interactions, perception noise, population structure), etc. For example, we can consider different distributions satisfying that the longer IRCOM survives, the more precisely or confidently it performs intention recognition; or, considering the repeated interaction setting in the framework of the iterated PD, the more IRCOM interacts with its co-player, the better it can recognize the co-player's intention (see intention recognition models for the iterated PD in Chapters 4 and 5).

We model  $cl$  by a continuous random variable  $X$  with probability density function  $f(x, U)$ , where  $U$  is a vector characterizing the factors that might influence  $cl$ , including the efficiency of the intention recognition model at work, the environmental factors (e.g., noise, population structure), and the interaction setting (repeated, one-shot, etc.).

If IRCOM is confident enough about the intention recognition process and result, that is  $cl$  is greater than a given, so-called, *confidence threshold*  $\theta \in [0, 1]$ , then in the current interaction IRCOM cooperates if the recognized intention of the co-player is to cooperate, and defects otherwise. The prediction is wrong with probability  $(1-cl)$ . For simplicity, we assume that the prediction is a (continuous) random variable,  $Y$ , uniformly distributed in  $[0, 1]$ . Hence, the probability that IRCOM utilizes intention recognition, but with an incorrect and correct prediction, respectively, can be written as joint probability distributions (Gut, 2009, Chapter 1) (Han et al., 2012e)

$$p_{ic} = P(X > \theta, Y < 1 - X) = \int_{\theta}^{+\infty} \int_{-\infty}^{1-x} f(x, U) dy dx, \quad (7.2)$$

$$p_c = P(X > \theta, Y > 1 - X) = \int_{\theta}^{+\infty} \int_{1-x}^{+\infty} f(x, U) dy dx. \quad (7.3)$$

If  $cl \leq \theta$ , i.e. IRCOM is not confident enough about its intention prediction, it behaves the same as COM\_C (see above). The greater  $\theta$  is, the more cautious IRCOM is about its intention recognition result. Obviously, if  $\theta = 1$ , IRCOM behaves identically to COM\_C; and if  $\theta = 0$ , IRCOM behaves identically to a (pure) intention recognizer (Han et al., 2011a,b) (see Figure 7.1).

We now replace COM\_C with IRCOM, considering a population of four strate-

gies, IRCOM, C, D, and D.COM. For the row player, the (average) payoff matrix reads

$$M = \theta M_1 + M_2, \quad (7.4)$$

where  $M_2$  is the payoff matrix when IRCOM utilizes the intention recognition strategy, i.e. in the case  $cl > \theta$ . To derive  $M_2$ , we consider the case that  $cl$  has a uniform distribution in the interval  $[0, 1]$ , i.e.  $f(x, U) = 1$  for  $x \in [0, 1]$  and 0 otherwise. Note that, on average, this can be considered as the distribution of a very inefficient intention recognition model because the confidence level or precision is a random number uniformly drawn from  $[0, 1]$ . As the prediction of the co-player's intention is only between two options, cooperate and defect, a random choice prediction already has a confidence level of 0.5.

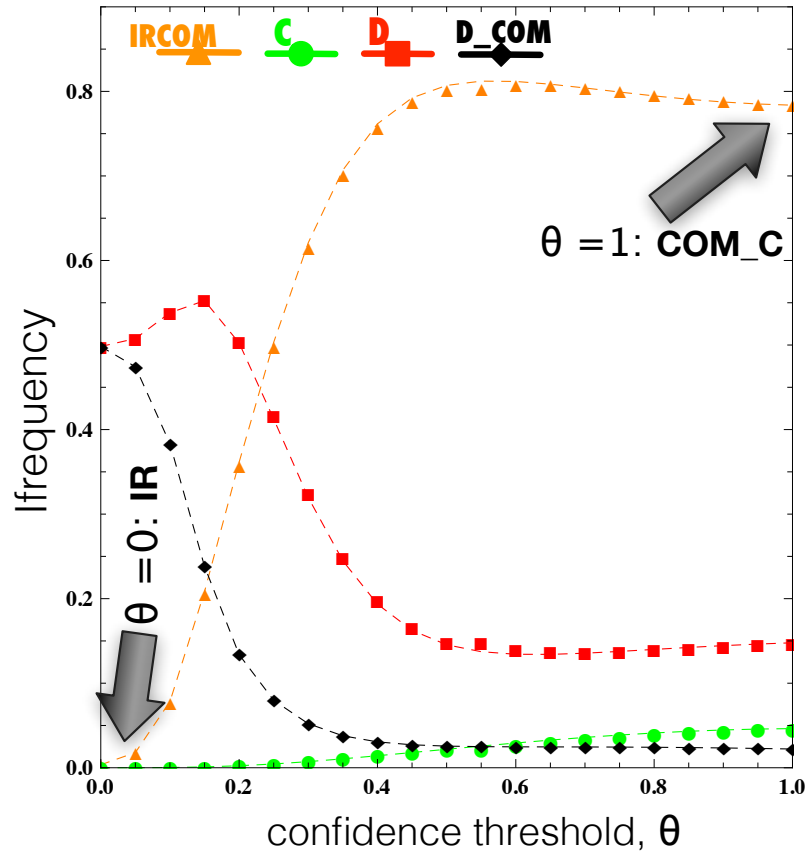
Computing the integrals in Eqs. (7.2) and (7.3), we obtain:  $p_{ic} = (1 - \theta)\frac{1-\theta}{2}$  and  $p_c = (1 - \theta)\frac{1+\theta}{2}$ . Hence,

$$M_2 = (1 - \theta) \begin{pmatrix} m_{11} & m_{12} & m_{13} & m_{14} \\ m_{21} & R & S & S \\ m_{31} & T & P & P \\ m_{41} & T & P & P \end{pmatrix} \quad (7.5)$$

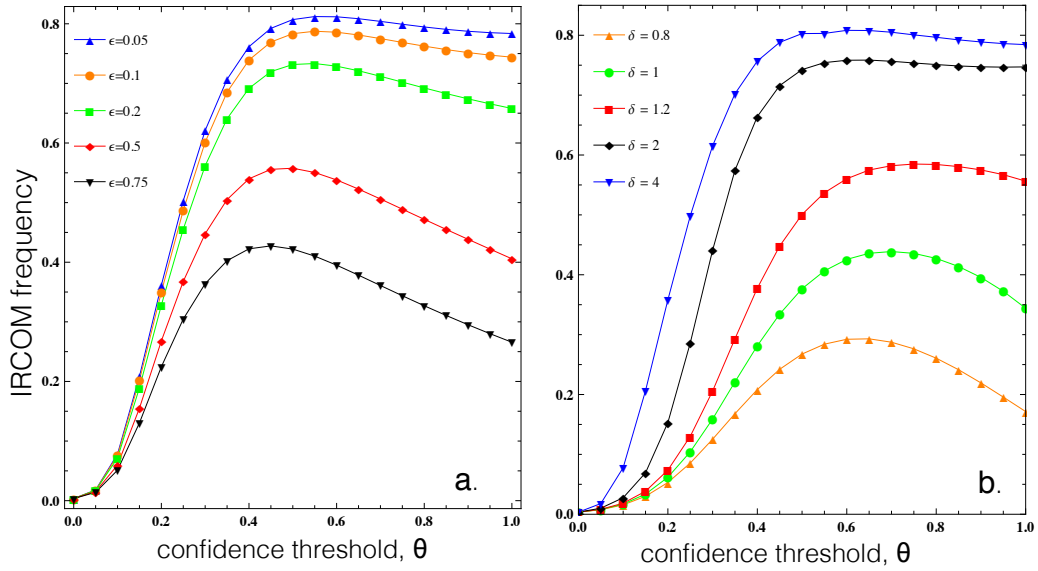
where

$$\begin{aligned} m_{11} &= \frac{1}{4} [P(1 - \theta)^2 + (S + T)(1 - \theta)(1 + \theta) + R(1 + \theta)^2]; \\ m_{12} &= \frac{1}{2} [T(1 - \theta) + R(1 + \theta)]; \quad m_{13} = \frac{1}{2} [S(1 - \theta) + P(1 + \theta)]; \\ m_{14} &= \frac{1}{2} [S(1 - \theta) + P(1 + \theta)]; \quad m_{21} = \frac{1}{2} [S(1 - \theta) + R(1 + \theta)]; \\ m_{31} &= \frac{1}{2} [T(1 - \theta) + P(1 + \theta)]; \quad m_{41} = \frac{1}{2} [T(1 - \theta) + P(1 + \theta)]. \end{aligned}$$

The main subject of the following analysis is to address, given the payoff entries of the PD, and the parameters of the commitment deal IRCOM can manage, how confident about the intention recognition result IRCOM should be to make a decision without relying on commitment proposal. That is, if there is an optimal value of  $\theta$  for an IRCOM to gain greatest net benefit. Our analysis is based on evolutionary game theory methods for finite populations (Nowak et al., 2004; Imhof et al., 2005) (see [EGT methods](#)).



**Figure 7.1:** Frequency of each strategy as a function of confidence threshold,  $\theta$ . Symbols indicate results from computer simulations (averaged over  $10^7$  interactions for each pair of strategies), and dashed curves show the exact numerical results. In a population of IRCOM, COM.D, C, and D individuals, for large enough  $\theta$ , the population spends most of the time in the homogeneous state of IRCOM. The performance of IRCOM decreases when  $\theta$  is too high. It implies that IRCOM should not be too cautious about its intention recognition capacity, i.e. not be too careful to always propose commitment instead of believing in its intention recognition capability and strategy. Parameters:  $\delta = 4$ ;  $\epsilon = 0.05$ ; payoff entries,  $T = 2$ ,  $R = 1$ ,  $P = 0$ ,  $S = -1$ ; population size,  $N = 100$ ; imitation strength,  $\beta = 0.1$ .



**Figure 7.2: Frequency of IRCOM as a function of confidence threshold,  $\theta$ .** (a) We plot for different values of management cost,  $\epsilon$ . In a population of IRCOM, COM\_D, C, and D individuals, for large enough  $\theta$  and small enough  $\epsilon$ , the population spends most of the time in the homogeneous state of IRCOM. The smaller  $\epsilon$ , the better the performance of IRCOM. The performance of IRCOM decreases when  $\theta$  is too high, and the greater  $\epsilon$ , the more significant the decrease is. It implies that the more costly the management of the commitment deal, the more beneficial it is to rely on intention recognition. (b) We plot for different values of punishment cost,  $\delta$ . In a population of IRCOM, COM\_D, C, and D individuals, for large enough  $\theta$  and large enough  $\delta$ , the population spends most of the time in the homogeneous state of IRCOM. The greater  $\delta$ , the better the performance of IRCOM. The performance of IRCOM decreases when  $\theta$  is too high, and the decrease is more significant when  $\delta$  is smaller. It implies that the weaker the commitment deal can be arranged, the more beneficial it is to rely on intention recognition alone more often. Parameters:  $\delta = 4$  in panel (a) and  $\epsilon = 0.05$  in panel (b); payoff entries,  $T = 2$ ,  $R = 1$ ,  $P = 0$ ,  $S = -1$ ; population size,  $N = 100$ ; imitation strength,  $\beta = 0.1$ . The results are computed numerically.

### 7.3 Results

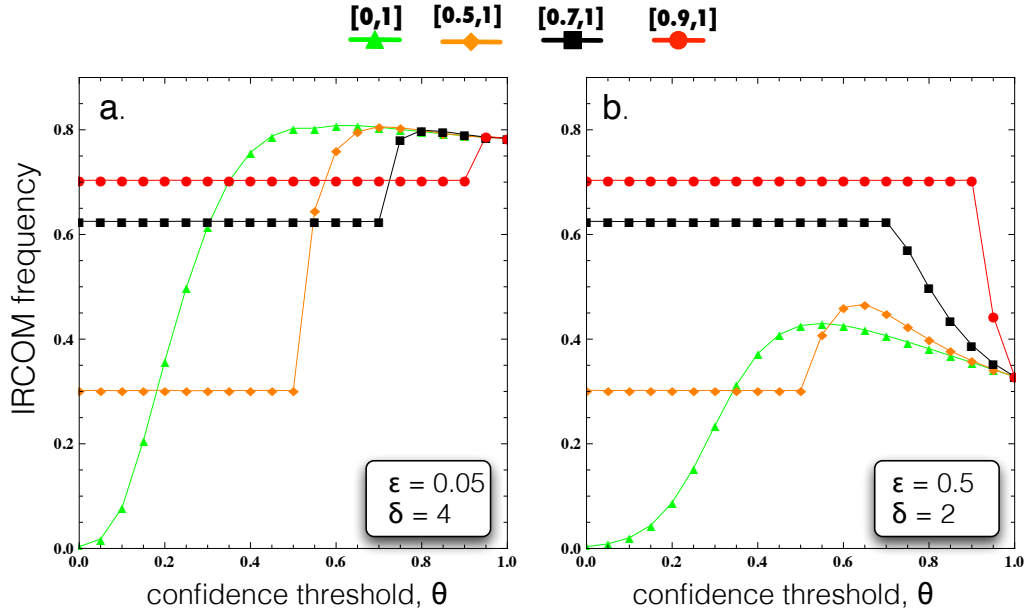
To start with, we compute the stationary distributions analytically and resorting to agent-based simulations (see [EGT methods](#) and [Figure 7.1](#)). The results show that, for a large range of  $\theta$ , IRCOM performs better than COM\_C (i.e. IRCOM with  $\theta = 1$ ), whereas the population spends most of the time in the homogenous state of IRCOM. However, when the confidence threshold is low, defection becomes dominant. This said, when the intention recognition is not of high enough certainty—that is, IRCOM is not confident enough about whether its co-player intends to cooperate or to defect in the current interaction—it would be better off counting on the commitment strategy (this also can be observed in several different configurations in [Figure 7.2](#)).

In [Figure 7.2](#), we analyze the influence of the strength of the commitment deal which can be arranged, on how the intention recognition and commitment strategies can be combined appropriately. Note that the greater the punishment cost,  $\delta$ , and the smaller the management cost,  $\epsilon$ , the stronger the commitment deal. Namely, in [Figure 7.2a](#), fixing  $\delta$ , we plot for different values of management cost,  $\epsilon$ . The performance of IRCOM decreases when  $\theta$  is too high, and the decrease is more dramatic when  $\epsilon$  is greater. It means that the costlier the management of the commitment deal, the more beneficial it is to rely on intention recognition. Next, in [Figure 7.2b](#), fixing  $\epsilon$ , we plot for different values of punishment cost,  $\delta$ . The performance of IRCOM decreases when  $\theta$  is too high, and the decrease is more dramatic when  $\delta$  is smaller. In short, these results imply that the weaker the commitment deal can be arranged, the more beneficial it is to rely on intention recognition.

So far the model has been studied with respect to a very inefficient intention recognition model, where  $cl$  is uniformly distributed in  $[0, 1]$ . It is not surprising that the performance of the intention recognition strategy solely—which corresponds to IRCOM with  $\theta = 0$  (see [Figure 7.1](#))—is very poor ([Figures 7.1](#) and [7.2](#)). In the sequel, let us study the model using more efficient intention recognition models.

We consider different probability distributions of the confidence level  $cl$ , reflecting different levels of efficiency or precision of the intention recognition model at work, given the relevant factors (noise, environment factors, interaction settings, etc.) ([Figure 7.3](#)). Namely, in an increasing order of efficiency,  $cl$  is uniformly drawn from intervals  $[0, 1]$ ,  $[0.5, 1]$ ,  $[0.7, 1]$ , and  $[0.9, 1]$ <sup>3</sup>. The results show that, whenever the

<sup>3</sup>In the context of iterated interactions (e.g. in the framework of the iterated Prisoner's



**Figure 7.3: Frequency of IRCOM as a function of confidence threshold,  $\theta$ , in a population of IRCOM, COM<sub>D</sub>, C, and D individuals.** We consider different probability distributions of the confidence level  $cl$ , reflecting the efficiency or precision of the intention recognition model at work, given the relevant factors (noise, environment, etc.). Namely,  $cl$  is uniformly drawn from  $[0, 1]$ ,  $[0.5, 1]$ ,  $[0.7, 1]$ , and  $[0.9, 1]$ . The results show that if a strong commitment deal can be arranged (panel a), it is better to rely on the commitment strategy using a high enough confidence threshold—even when the intention recognition model is very efficient, while it is more beneficial to rely, even exclusively, on the intention recognition strategy if it is efficient enough, in the case that only weak commitment deals can be arranged (panel b). Parameters:  $\epsilon = 0.05$ ,  $\delta = 4$  (panel a) and  $\epsilon = 0.5$ ,  $\delta = 2$  (panel b); payoff entries,  $T = 2$ ,  $R = 1$ ,  $P = 0$ ,  $S = -1$ ; population size,  $N = 100$ ; imitation strength,  $\beta = 0.1$ ; The payoff matrices in all cases are derived by averaging  $10^7$  interactions of each pair of strategies. The results are computed numerically.

intention recognition model is efficient enough, the intention recognition strategy solely (i.e. IRCOM with  $\theta = 0$ ) performs quite well, complying with the results obtained in (Han et al., 2011a,b), where concrete intention recognition models are deployed.

However, when a quite strong commitment deal can be envisaged (Figure 7.3a), arranging it can still glean some evolutionary advantage. But in case that only weak commitment deals can be arranged (Figure 7.3b), it is then more beneficial to rely, even exclusively, on the intention recognition strategy should it be efficient enough.

## 7.4 Discussion

A general implication of our analysis is that an appropriate combination of the two strategies of commitment and intention recognition often leads to a strategy that performs better than either one solely. It is advantageous to rely on the intention recognition strategy (when reaching sufficiently high confidence about its result) because it helps to avoid the cost of arranging and managing commitment deals, especially when no strong deals can be arranged or envisaged. This result has a similar implication to that obtained in (Johnson and Fowler, 2011), where the authors show that overconfidence might give evolutionary advantage to its holders. In our model, an IRCOM can gain extra net benefit if it is a little overconfident (that is, when using sufficiently small  $\theta$ ), taking the risk to rely on the intention recognition result instead of arranging some commitment deal. Differently, because in our model IRCOM is further guaranteed by an efficient strategy of commitment, being over-overconfident (that is, using a too small  $\theta$ ) and relying exclusively on intention recognition might prevent it from opportunely gaining benefit from the commitment strategy—especially in case the intention recognition model at work is not efficient. This said, the performance of overconfident individuals (Johnson and Fowler, 2011) can be enhanced by relying on the commitment strategy when they need to muster overly high courage (say, in order to decide to claim some resource).

---

Dilemma), these levels of efficiency can be achieved (on average) by considering large enough numbers of interactions between two players (or high enough probabilities of a next interaction or ‘the shadow of future’ (Nowak, 2006b; Sigmund, 2010)), given that the noise is small enough. Normally, the more an intention recognizer interacts with a fixed co-player, the better it predicts its co-player’s intention. For example, this holds for the two intention recognition models described in (Han et al., 2011a,b). Furthermore, in (Frank et al., 1993), the authors present experimental evidence showing that, in a one-shot PD, subjects of only brief acquaintance were able to recognize players with an intention to defect with more than twice chance accuracy.

In the framework where intention recognition is difficult and of high risk, for example, climate change negotiation (Raihani and Aitken, 2011; Milinski et al., 2006; Santos and Pacheco, 2011), military setting (comprising a lot of bluffing) (Schelling, 1990; Leeds, 2003), and international relationships (Lockhart, 1978), our model suggests arranging a strong commitment deal.

## 7.5 Conclusion

Assume simply that we are given an intention recognition method, that affords us a degree of confidence distribution  $cl$  about its predictions, with regard to the intentions of others, and hence their future actions, typically on the basis of their seen actions and surrounding historical and present circumstances. Assume too some commitment model is given us about providing mutual assurances, and involving an initial cost and a penalty for defaulting.

We have shown how to combine together one such general intention recognition method, with a specific commitment model defined for playing the Prisoner's Dilemma (PD), in the setting of Evolutionary Game Theory (EGT), by means of a single payoff matrix extended with a new kind of player, IRCOM, which chooses whether to go by the result of its intention recognition method about a co-player's next move, or to play by the commitment strategy, depending on whether its level of confidence on the intention prediction,  $cl$ , exceeds or not some a given confidence threshold  $\theta$ . Our results indicate that IRCOM is selected by evolution for a broad range of parameters and confidence thresholds.

Then we have studied, for a variety of  $cl$  and  $\theta$ , in the context of PD in EGT, how IRCOM performs in the presence of other well-known non-committing strategies (always cooperate, C, and always defect, D) – plus the strategy that commits when being asked to, but always defects, D\_COM. Analytical and simulation results show under which circumstances, for different  $cl$  and  $\theta$ , and distinct management and punishment costs,  $\epsilon$  and  $\delta$ , the new combined strategy IRCOM proves advantageous and to what degree. And indeed IRCOM proves to be adaptably advantageous over those other strategies, in all circumstances, from a quite small confidence level onwards.

Much remains to be done with respect to further consideration of combining the two strategies of intention recognition and commitment. The two go often together,

and not just in the basic way we have examined. Indeed, they are the two sides of the same coin, one side being an attempt to identify an intention, the other side being a manifestation of an intention. For one, we only considered the case where intention recognition comes first in order to decide on a commitment proposal. But in general, once a commitment is made, intention recognition is a paramount method to follow up on whether the commitment will be honored, on the basis of detecting or not the intermediate actions leading up to commitment fulfillment. Social organizations rely on these mechanisms to structure themselves. Furthermore, the information about commitments can be used to enhance intention recognition. For example, consider the Bayesian Network based intention recognition model in Chapter 4, which operates upon computing mutual trust. One can naturally say that the more often a co-player fulfills commitments, the more the intention recognizer should trust him/her, and that the co-player who is currently in a commitment deal (with the intention recognizer) is more trustworthy (to the intention recognizer) than the one without such a commitment, even he/she has been more cooperative in the past (directly as well as indirectly).

It seems to us that intention recognition, and its use in the scope of commitment, is a foundational cornerstone where we should begin at, naturally followed by the capacity to establish and honor commitments, as a tool towards the successive construction of collective intentions and social organization (Searle, 1995, 2010). Finally, one hopes that understanding these capabilities can be useful in the design of efficient self-organized and distributed engineering applications (Bonabeau et al., 1999), from bio and socio-inspired computational algorithms to swarms of autonomous robotic agents.



# Chapter 8

---

## CONCLUSIONS AND FUTURE WORK

---

*Begin thus from the first act, and proceed;  
and, in conclusion,  
at the ill which thou hast done,  
be troubled, and rejoice for the good.*

Pythagoras, 582 BC - 497 BC

**Abstract.** This final chapter is structured as follows. First, a brief summary of the thesis and the main conclusions are provided. Then some directions for future work are discussed.

### 8.1 Summary and Conclusions

This thesis is split into two parts. In Part I, we study techniques for intention recognition and its applications, which span two chapters.

In Chapter 2 we present a corpus-based method for intention recognition, which is context-dependent and incremental. The method is performed by dynamically constructing a three-layer Bayesian Network (BN) model for intention recognition

(IRBN), from a prior knowledge base consisting of readily maintained and readily constructed fragments of BN. Their simple structures allow easy maintenance by domain experts or their automatic construction from available plan corpora. We have shown that taking into account contextual information appropriately influences the intention recognition process, and can significantly improve its performance, especially when the observed agents might change or abandon their initial goals/intentions. In addition, we have formalized a method for representing relationships amongst intentions in an IRBN, thereby enabling us to more appropriately handle the case where the observed agent might follow multiple intentions at once. These two aspects are apparently indispensable in intentions recognition, but mostly omitted in previous work. Furthermore, the context-dependent aspect of our method is reinforced by a combination of BN and Logic Programming techniques. While the IRBNs represent average general knowledge, Logic Programming techniques can be used for anchoring the IRBNs in specific contexts as they evolve. This way, the size of the IRBNs can significantly be reduced, and sometimes the final solution of the intention recognition process can be reached immediately, without performing BN inference. The method has been evaluated with several benchmarks, including the Linux and IPD plan corpora, with some definite degree of success. The IPD plan corpora themselves also endow new, domain-independent benchmarks, for the evaluation of intention recognition systems.

We admit that the method exhibits several aspects that need further evaluation and improvement, which were pointed out in the chapter's own conclusions (see Section 2.9). Nonetheless, it was developed far enough to be successfully applicable for the modelling work in the evolutionary game theory context in Part II—the more important aim and contribution of this thesis.

In Chapter 3 we present a framework for intention-based decision making, which adopts our previously implemented Evolution Proseption system (Pereira and Han, 2009a,b; Han, 2009) for decision making, and the intention recognition system described in Chapter 2 for recognizing intentions of other agents. Various ways in which recognized intentions of others can be useful for enhancing decision making processes are described and exemplified. The framework serves to bridge intention recognition to those real-world application domains, where it is important and useful to take into account intentions of other agents in the environment, plus

situation specific information and reasoning about it—before making a decision to act. We have demonstrated its applicability in the Elder Care and Ambient Intelligence domains.

In Part II, we study the roles of two cognitive abilities that are ubiquitous in humans, those of intention recognition and commitment, in the evolution of cooperative behavior within a population of agents. It comprises four chapters.

In Chapters 4 and 5, using the tools of evolutionary game theory, we present two different models that explicitly address the roles played by intention recognition in the evolution of cooperation. In both these models we have shown, analytically as well as by computer simulations, that intention recognition strongly promotes the emergence of cooperation and maintains it well, in spite of mutations and noise. We have shown that the strategy of intention recognition outperforms the famous strategies in the framework of the iterated Prisoners' Dilemma, including tit-for-tat (TFT), generous tit-for-tat (GTFT), and win-stay-lose-shift (WSLS).

The model in Chapter 4 is based on a specification of the general Bayesian network structure for intention recognition provided in Chapter 2, now in the framework of social dilemmas. It takes into account mutual trust between intention recognizers and their co-players, accumulated through past direct interactions. The intentions to be recognized are in form of the co-player's motives, that is, to cooperate or to defect in the current interaction, taking into account the player's pairwise interaction history. Differently, the model in Chapter 5 relies on the corpus-based incremental intention method described in Chapter 2, wherein the intentions to be recognized are the co-players' possible known strategies. The method requires a prior plan corpus, which is assumed to stand for (intention recognizers') ancestors' received knowledge or acquired experience in interacting with known strategies (Woodward et al., 2009; Meltzoff, 2007, 2005).

The two different models have shown, to some significant degree, that there are various ways in which the capacity for intention recognition can confer on its holder an evolutionary advantage, specifically in situations where cooperation is profitable. Given the broad spectrum of problems which are addressed using this cooperative metaphor (Axelrod, 1984; Nowak, 2006a; Sigmund, 2010), our results indicate how intention recognition can be pivotal in social dynamics.

In Chapter 6, resorting again to the tools of evolutionary game theory, we provide a computational model showing that even commitment of the simplest form can lead to the emergence of cooperation, albeit in non-iterated settings. The model is characterized by two key parameters: the punishment cost of failing a commitment deal imposed on either side of the deal, and the cost of managing it. Our analytical results and extensive computer simulations exhibit that cooperation can emerge if the punishment cost is large enough compared to the management cost. Our study suggests that humans' specialized capacity for commitment, which might have been shaped by natural selection (Nesse, 2001b), consists in a capacity for managing to impose a high cost of punishment, whether it is monetary or of abstract emotional or reputation value, with a relatively small managing cost. Furthermore, these results have provided important insight into the design of multi-agent systems resorting to commitments to facilitate cooperation (Schelling, 1990; Wooldridge and Jennings, 1999; Hirshleifer, 2001; Castelfranchi and Falcone, 2010; Winikoff, 2007; Harrenstein et al., 2007; Chopra and Singh, 2009).

On top of the commitment model, in Chapter 7, we bring forward a model showing that an appropriate combination of intention recognition and commitment can lead to a strategy performing better than either one by itself, thereby leading to the emergence of improved cooperation.

## 8.2 Future works

### 8.2.1 Part I: Intention Recognition Methods

As mentioned above, there are several aspects of our intention recognition method presented in Chapter 2 that demand further development and assessment. Besides the technical issues pointed out in the chapter's conclusions (Section 2.9), we intend to gather real data for further evaluation, especially regarding the multiple intentions recognition formalism and the combination of Bayesian Network and Logic Programming techniques for context-dependent intention recognition. As another area of future effort, we aim at real deployment of our intention recognition method to further tackle different application domains, e.g. Ambient Intelligence (Han and Pereira, 2010c; Sadri, 2011a), Elder Care (Pereira and Han, 2011a), and several others that are listed in Chapter 3 (see Section 3.5), where intention recognition has

been of increasing importance (Sadri, 2011b,a; Han and Pereira, 2010c; Pereira and Han, 2011b; Geib, 2002; Geib and Goldman, 2001; Haigh et al., 2004; Armentano and Amandi, 2007). The data from these application domains can then be used to furnish further evaluation of the method.

In this thesis we have focused on the problem of *individual* intention recognition, that is, recognizing intentions of a single agent. Many complications arise and are magnified when considering intention recognition of multiple agents, or *collective* intention recognition as it is usually dubbed (Kaminka et al., 2001; Sukthankar, 2007). Collective intention is an important topic actively discussed in philosophical and multi-agent system literature (Bratman, 1999; Malle et al., 2003; Tuomela, 2005; Searle, 2010; Kanno et al., 2003). Most researchers agree that collective intentions are not summative, that is, a collective intention of a group of agents is not merely reducible to a summation of the individual intentions. It requires some “glue” underpinning cognitive attitudes—such as mutual beliefs and mutual awarenesses—to be there amongst the agents to guarantee they hold a collective intention and perform activities cooperatively.

Although there has been much effort made for tackling the problem of recognizing intentions of a group of agents, most, if not all, ignore to address the required “glue” cognitive attitudes, making the assumption that the collective intention is initially assigned to and adopted by the group—like in case of a football team or an army troop (see for instance (Kaminka et al., 2001)). Consequently, they could not fully address the problem of collective intention recognition. They could recognize the collective intention of an arbitrary set of agents assuming beforehand that it existed, but they cannot figure if it was actually so designed. For example, in the Elder Care domain concerning multiple users, there are often cases where the elders’ actions accidentally form a plausible plan—which achieves a conceivable intention—but actually each of them is following his/her own intention. In (Han and Pereira, 2010b), we have done some initial work on this direction, which we intend to develop further in near future.

### 8.2.2 Part II: The Evolution of Cooperation

In this thesis we study the roles of two cognitive abilities which are ubiquitous in humans, those of intention recognition and of arranging commitment, in the evolution of cooperation within a population of agents. There are several potentially

valuable continuations of these inroads.

First, in what regards intention recognition, we have built models focusing only on direct interactions, that is, agents playing the iterated Prisoner's Dilemma might recognize intentions of others (their motives to cooperate or defect, or their strategies) based on past direct interactions. Undoubtedly, the intention recognition process can be improved using observable indirect interactions (Nowak and Sigmund, 2005), the concrete structure of a population (Szabó and Fáth, 2007; Santos et al., 2006b), the information about group membership (García and van den Bergh, 2010) or kinship relations (Hamilton, 1964; West et al., 2002), or similarity (on an arbitrary characteristic) (Riolo et al., 2001) of the recognized agents, etc. Hence, as natural extensions, we will attempt to build models to study how the intention recognition capacity can enable cooperation to emerge in such wider contexts. Another possibility of extension regarding intention recognition is that our three-layer Bayesian network model for intention recognition (see the model in Chapter 4) can provide a general framework for an integrated modeling of several known important factors, such as trust, reputation, commitment, emotions, etc. (Dreber et al., 2008; Ohtsuki et al., 2009; Fehr and Gächter, 2002)—in order to study their combined influence for the evolution of cooperation. Tentatively, such factors can figure in the first layer of the three-layer network, as they might causally affect the way a co-player behaves or what he intends to do. For instance, the co-player's intention to cooperate seems to strongly depend on his emotions towards his co-player, as well as his co-player's reputation and commitments.

Regarding commitment, we have built a computational model showing that a simple form of commitment can promote the evolution of cooperation in a (well-mixed) population of individuals playing the one-shot Prisoner's Dilemma, i.e. without taking into account the structure of the population, reciprocity, kinship, etc. We have also argued that our model can be a good replacement for the models on punishment, which recently have been shown not to promote cooperation well (Rand and Nowak, 2011). In future work, similarly to what has been done with punishment (Ohtsuki et al., 2009; de Weerd and Verbrugge, 2011), we would like to study how the mechanism of commitment works in combination with other known mechanisms of cooperation, including direct and indirect reciprocity, and structured populations.

In addition, in this thesis we only study the evolution of cooperation within the

two-player game interaction context, namely the Prisoner’s Dilemma. But apparently multi-player interactions are unavoidable in real life situations (Hardin, 1968; Hauert et al., 2002; Gokhale and Traulsen, 2010; Han et al., 2012e; Santos et al., 2008; Pacheco et al., 2009), and therein intention recognition might also be useful. In this respect, it also opens the door for studying collective intention recognition modeling and methods (Han and Pereira, 2010b; Sukthankar and Sycara, 2008; Kanno et al., 2003) in a large scale population setting, where there might be group of agents such as alliances and teams attempting collective goals.

Furthermore, in this thesis we have focused on two cognitive abilities, the ability to recognize intentions of others and the ability to arrange appropriate commitments. Obviously, there are several other cognitive abilities that might have important evolutionary roles, which we would like to explore. First choices are those considered part of the “Theory of Mind” theory, including the understanding and attribution of beliefs, desires, preferences, emotions, pretending, deceit, etc. (Premack and Woodruff, 1978; Whiten, 1991; Cheney and Seyfarth, 2007; Trivers, 2011).

---

In summary, we have provided new, clear (theoretical and experimental simulations) evidence revealing the evolutionary advantages of having high cognitive skills, specifically in dealing with problems where cooperation is beneficial. The approach of resorting to AI cognitive modeling methods and the tools of evolutionary game theory that we adopted here can be taken up to further investigate the evolutionary roles of cognition in the evolution of cooperative behavior, on the one hand, and, on the other hand, supply important insight into the global influence of different AI cognitive modeling approaches within a population of agents, thereby facilitating their enhancement.

---



# Bibliography

- Aesop. Fox-crow fable: <http://mythfolklore.net/aesopica/bewick/51.htm>. Last access: November 24, 2011. 5, 73
- Akashi, H. 1995. Inferring weak selection from patterns of polymorphism and divergence at "silent" sites in *Drosophila* DNA. *Genetics*, 139(2):1067–1076. 134
- Albrecht, D. W., Zukerman, I., and Nicholson, A. E. 1998. Bayesian models for keyhole plan recognition in an adventure game. *User Modeling and User-Adapted Interaction*, 8:5–47. 19, 60
- Alferes, J. J., Banti, F., Brogi, A., and Leite, J. A. 2005. The refined extension principle for semantics of dynamic logic programming. *Studia Logica*, 79(2):7–32. 11
- Alferes, J. J., Leite, J. A., Pereira, L. M., Przymusinska, H., and Przymusinski, T. C. 2000. Dynamic updates of non-monotonic knowledge bases. *J. Logic Programming*, 45(1-3):43–70. 11
- Alferes, J. J., Pereira, L. M., and Swift, T. 2004. Abduction in well-founded semantics and generalized stable models via tabled dual programs. *Theory and Practice of Logic Programming*, 4(4):383–428. 11, 63, 64
- Ampatzis, C., Tuci, E., Trianni, V., and Dorigo, M. 2008. Evolution of signaling in a multi-robot system: Categorization and communication. *Adaptive Behavior*, 16(1):5–26. 7
- Antal, T., Traulsen, A., Ohtsuki, H., Tarnita, C. E., and Nowak, M. A. 2009. Mutation-selection equilibrium in games with multiple strategies. *J. Theor. Biol.*, 258:614–622. 149

## BIBLIOGRAPHY

---

- Armentano, M. G. and Amandi, A. 2007. Plan recognition for interface agents. *Artif. Intell. Rev.*, 28:131–162. [4](#), [16](#), [17](#), [19](#), [59](#), [68](#), [85](#), [89](#), [177](#)
- Armentano, M. G. and Amandi, A. 2009. Goal recognition with variable-order markov models. In *Proceedings of the 21st international joint conference on Artificial intelligence*, pages 1635–1640. [4](#), [17](#), [20](#), [21](#), [35](#), [36](#), [37](#), [38](#), [39](#), [120](#), [162](#)
- Arthur, W. B. 1994. Inductive reasoning and bounded rationality. *American Economic Review*, 84(2):406–411. [7](#), [135](#)
- Axelrod, R. 1984. *The Evolution of Cooperation*. Basic Books, ISBN 0-465-02122-2. [8](#), [40](#), [50](#), [90](#), [93](#), [140](#), [143](#), [175](#)
- Axelrod, R. and Hamilton, W. 1981. The evolution of cooperation. *Science*, 211:1390–1396. [93](#), [140](#)
- Back, I. and Flache, A. 2008. The Adaptive Rationality of Interpersonal Commitment. *Rationality and Society*, 20(1):65–83. [9](#), [140](#), [143](#), [159](#)
- Baral, C. 2003. *Knowledge Representation, Reasoning, and Declarative Problem Solving*. Cambridge University Press, New York, NY, USA. [11](#), [51](#), [55](#), [60](#), [64](#)
- Baral, C. and Gelfond, M. 1994. Logic programming and knowledge representation. *Journal of Logic Programming*, 19:73–148. [11](#)
- Baral, C., Gelfond, M., and Rushton, N. 2009. Probabilistic reasoning with answer sets. *Theory and Practice of Logic Programming*, 9(1):57–144. [52](#), [53](#), [63](#)
- Blaylock, N. and Allen, J. 2003. Corpus-based, statistical goal recognition. In *Proceedings of the 18th international joint conference on Artificial intelligence (IJCAI'03)*, pages 1303–1308. [4](#), [17](#), [21](#), [32](#), [120](#)
- Blaylock, N. and Allen, J. 2004. Statistical goal parameter recognition. In Zilberstein, S., Koehler, J., and Koenig, S., editors, *Proceedings of the 14th International Conference on Automated Planning and Scheduling (ICAPS'04)*, pages 297–304. AAAI. [4](#), [17](#), [21](#), [32](#), [35](#), [36](#), [38](#), [39](#), [60](#), [68](#), [120](#), [162](#)
- Blaylock, N. and Allen, J. 2005. Recognizing instantiated goals using statistical methods. In Kaminka, G., editor, *IJCAI Workshop on Modeling Others from Observations (MOO-2005)*, pages 79–86. [35](#)

## BIBLIOGRAPHY

---

- Boehm, C. 1999. The natural selection of altruistic traits. *Human Nature*, 10(3):205–252. [140](#)
- Bonabeau, E., Dorigo, M., and Theraulaz, G. 1999. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, USA. [7](#), [171](#)
- Börgers, T., Sarin, R., and Street, G. 1997. Learning through reinforcement and replicator dynamics. *Journal of Economic Theory*, 77:1–14. [91](#)
- Bratman, M. E. 1987. *Intention, Plans, and Practical Reason*. The David Hume Series, CSLI. [7](#), [9](#), [16](#), [17](#), [20](#), [28](#), [32](#), [40](#), [55](#), [62](#), [63](#), [84](#), [89](#), [99](#), [158](#)
- Bratman, M. E. 1992. Planning and the stability of intention. *Minds and Machines*, 2(1):1–16. [16](#), [17](#), [20](#), [55](#)
- Bratman, M. E. 1999. *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge University Press. [7](#), [62](#), [89](#), [99](#), [177](#)
- Brewka, G. and Eiter, T. 1999. Preferred answer sets for extended logic programs. *Artif. Intell.*, 109:297–356. [11](#)
- Brown, S. M. 1998. *A decision theoretic approach for interface agent development*. PhD thesis, Faculty of the Graduate School of Engineering of the Air Force Institute of Technology Air University, USA. [20](#)
- Bui, H. 2003. A general model for online probabilistic plan recognition. In Gottlob, G. and Walsh, T., editors, *Proceedings of the 8th International Joint Conference on Artificial Intelligence (IJCAI'03)*, pages 1309–1318. [21](#), [89](#)
- Bui, H., Venkatesh, S., and West, G. 2002. Policy recognition in the abstract hidden markov model. *Journal of Artificial Intelligence Research*, 17:451–499. [5](#), [19](#), [98](#), [162](#)
- Burmeister, B., Arnold, M., Copaciu, F., and Rimassa, G. 2008. BDI-agents for agile goal-oriented business processes. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems: industrial track, AAMAS '08*, pages 37–44. [62](#)
- Castelfranchi, C. and Falcone, R. 2010. *Trust Theory: A Socio-Cognitive and Computational Model (Wiley Series in Agent Technology)*. Wiley. [142](#), [143](#), [176](#)

## BIBLIOGRAPHY

---

- Cesta, A. and Pecora, F. 2004. The robocare project: Intelligent systems for elder care. In *AAAI Fall Symposium on Caring Machines: AI in Elder Care*. 77
- Charlesworth, J. and Eyre-Walker, A. 2007. The other side of the nearly neutral theory, evidence of slightly advantageous back-mutations. *Proceedings of the National Academy of Sciences of the United States of America*, 104(43):16992–16997. 134
- Charniak, E. and Goldman, R. P. 1990. Plan recognition in stories and in life. In *Proceedings of the Fifth Annual Conference on Uncertainty in Artificial Intelligence*, UAI '89, pages 343–352, Amsterdam, The Netherlands, The Netherlands. North-Holland Publishing Co. 4, 85
- Charniak, E. and Goldman, R. P. 1991. A probabilistic model of plan recognition. In *Proceedings of the ninth National conference on Artificial intelligence - Volume 1 (AAAI'1991)*, AAAI'91, pages 160–165. AAAI Press. 20
- Charniak, E. and Goldman, R. P. 1993. A Bayesian model of plan recognition. *Artificial Intelligence*, 64(1):53–79. 4, 7, 16, 17, 19, 20, 89, 98, 120, 158, 162
- Cheney, D. L. and Seyfarth, R. M. 2007. *Baboon Metaphysics: The Evolution of a Social Mind*. University Of Chicago Press. 6, 90, 120, 179
- Chopra, A. K. and Singh, M. P. 2009. Multiagent commitment alignment. In *Proceedings of the 8th International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, pages 937–944. 143, 176
- Cohen, P. R. and Levesque, H. J. 1990. Intention is Choice with Commitment. *Artificial Intelligence*, 42(2-3):213–261. 9, 17, 40, 62, 99, 158
- Cohen, P. R., Perrault, C. R., and Allen, J. F. 1981. Beyond question answering. In *Strategies for Natural Language Processing*, pages 245–274. Lawrence Erlbaum Associates. 4, 5
- Conati, C., Gertner, A. S., Vanlehn, K., and Druzdzel, M. J. 1997. On-Line Student Modeling for Coached Problem Solving Using Bayesian Networks. In *Proceedings of the Sixth International Conference on User Modeling*, pages 231–242. 19

## BIBLIOGRAPHY

---

- Cook, D. J., Augusto, J. C., and Jakkula, V. R. 2009. Ambient intelligence: Technologies, applications, and opportunities. *Pervasive and Mobile Computing*, 5(4):277–298. [76](#), [77](#)
- Cozman, F. G. 2004. Axiomatizing noisy-or. In de Mantaras, R. L. and Saitta, L., editors, *In Proceedings of the 16th European Conference on Artificial Intelligence (ECAI-04)*, pages 979–980. IOS Press. [27](#), [28](#)
- Darwin, C. 1859. *On the origin of species by means of natural selection*. Cambridge-London. Reprinted in Harvard University Press (1964). [30](#). [8](#)
- de Vos, Smaniotto, R., and Elsas, D. 2001. Reciprocal altruism under conditions of partner selection. *Rationality and Society*, 13(2):139–183. [9](#), [140](#), [143](#), [159](#)
- de Weerd, H. and Verbrugge, R. 2011. Evolution of altruistic punishment in heterogeneous populations. *Journal of Theoretical Biology*, 290(0):88 – 103. [178](#)
- Dean, T. and Kanazawa, K. 1992. Probabilistic Temporal reasoning. In *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI-88)*, pages 524–528. AAAI. [60](#)
- Diez, F. J. 1993. Parameter adjustment in Bayes networks. the generalized noisy or-gate. In *Proceedings of Conference on Uncertainty in Artificial Intelligence (UAI'93)*, pages 99–105. Morgan Kaufmann. [18](#), [28](#)
- Dreber, A., Rand, D. G., Fudenberg, D., and Nowak, M. A. 2008. Winners don't punish. *Nature*, 452(7185):348–351. [91](#), [153](#), [178](#)
- Duong, T., Phung, D., Bui, H., and Venkatesh, S. 2009. Efficient duration and hierarchical modeling for human activity recognition. *Artif. Intell.*, 173:830–856. [17](#)
- Falk, A., Fehr, E., and Fischbacher, U. 2008. Testing theories of fairness—intentions matter. *Games and Economic Behavior*, 62(1):287 – 303. [50](#)
- Fehr, E. and Gächter, S. 2002. Altruistic punishment in humans. *Nature*, 415:137–140. [91](#), [153](#), [178](#)
- Forbes, J., Huang, T., Kanazawa, K., and Russell, S. 1995. The batmobile: towards a Bayesian automated taxi. In *Proceedings of the 14th international joint conference on Artificial intelligence - Volume 2*, pages 1878–1885. [19](#), [60](#)

## BIBLIOGRAPHY

---

- Frank, R. H. 1988. *Passions Within Reason: The Strategic Role of the Emotions*. W. W. Norton and Company. 10, 159
- Frank, R. H. 2001. Cooperation through Emotional Commitment. In Nesse, R. M., editor, *Evolution and the capacity for commitment*, pages 55–76. New York: Russell Sage. 10, 159
- Frank, R. H., Gilovich, T., and Regan, D. T. 1993. The evolution of one-shot cooperation: An experiment. *Ethology and Sociobiology*, 14(4):247 – 256. 169
- Friedewald, M., Costa, O. D., Punie, Y., Alahuhta, P., and Heinonen, S. 2005. Perspectives of ambient intelligence in the home environment. *Telemat. Inf.*, 22:221–238. 80
- Friedewald, M., Vildjiounaite, E., Punie, Y., and Wright, D. 2007. Privacy, identity and security in ambient intelligence: A scenario analysis. *Telematics and Informatics*, 24(1):15 – 29. 80
- Fudenberg, D. and Imhof, L. A. 2005. Imitation processes with small mutations. *Journal of Economic Theory*, 131:251–262. 95
- Fudenberg, D. and Levine, D. K. 1998. *The Theory of Learning in Games*, volume 1 of *MIT Press Books*. The MIT Press. 121
- García, J. and van den Bergh, J. C. 2010. Evolution of parochial altruism by multilevel selection. *Evolution and Human Behavior*, 32(4):277–287. 178
- Gazzaniga, M. 2011. *Who’s in Charge?: Free Will and the Science of the Brain*. HarperCollins. 120
- Geib, C. W. 2002. Problems with intent recognition for elder care. In *Proceedings of AAAI Workshop Automation as Caregiver*. 4, 77, 85, 177
- Geib, C. W. 2004. Assessing the complexity of plan recognition. In *Proceedings of the Conference of the American Association of Artificial Intelligence (AAAI’2004)*, pages 507–512. 20
- Geib, C. W. and Goldman, R. P. 2001. Plan recognition in intrusion detection systems. In *Proceedings of the DARPA Information Survivability Conference and Exposition (DISCEX)*. 5, 6, 77, 177

- Geib, C. W. and Goldman, R. P. 2003. Recognizing plan/goal abandonment. In *Proceedings of the 18th international joint conference on Artificial intelligence (IJCAI'03)*, pages 1515–1517. [5](#), [21](#), [59](#)
- Geib, C. W. and Goldman, R. P. 2009. A probabilistic plan recognition algorithm based on plan tree grammars. *Artificial Intelligence*, 173(2009):1101–1132. [4](#), [5](#), [7](#), [16](#), [17](#), [19](#), [20](#), [21](#), [59](#), [89](#), [120](#)
- Gintis, H. 2001. Beyond selfishness in modeling human behavior. In Nesse, R. M., editor, *Evolution and the capacity for commitment*. New York: Russell Sage. [9](#), [140](#), [143](#), [144](#), [158](#), [161](#)
- Glymour, C. 2001. *The Mind's Arrows: Bayes Nets and Graphical Causal Models in Psychology*. MIT Press. [98](#)
- Gogate, V. and Dechter, R. 2011. SampleSearch: Importance sampling in presence of determinism. *Artificial Intelligence*, 175(2):694–729. [18](#)
- Gokhale, C. S. and Traulsen, A. 2010. Evolutionary games in the multiverse. *Proceedings of the National Academy of Sciences of the United States of America*, 107(12):5500–5504. [179](#)
- Goldman, R. P., Geib, C. W., and Miller, C. 1999. A new model of plan recognition. In *Proceedings of Conference on Uncertainty in Artificial Intelligence (UAI'99)*, pages 245–254. [20](#)
- Goodenough, O. R. 2001. Law and the biology of commitment. In Nesse, R. M., editor, *Evolution and the capacity for commitment*, pages 262–291. New York: Russell Sage. [156](#)
- Gut, A. 2009. *An Intermediate Course in Probability*. Springer Publishing Company, Incorporated, 2nd edition. [163](#)
- Gutierrez, A., Campo, A., Santos, F. C., Monasterio-Huelin, and Dorigo, M. 2009. Social odometry: Imitation based odometry in collective robotics. *International Journal of Advanced Robotics Systems*, 2(6):129–136. [7](#)
- Haigh, K., Kiff, L., Myers, J., Guralnik, V., Geib, C., Phelps, J., and Wagner, T. 2004. The Independent LifeStyle Assistant (I.L.S.A.): AI Lessons Learned. In

- Proceedings of Conference on Innovative Applications of Artificial Intelligence*, pages 852–857. [4](#), [77](#), [85](#), [177](#)
- Halteren, van H. 1999. *Syntactic Wordclass Tagging*. Kluwer Academic Publishers, Dordrecht, The Netherlands. [137](#)
- Hamilton, W. 1964. The genetical evolution of social behaviour. i. *Journal of Theoretical Biology*, 7(1):1 – 16. [8](#), [90](#), [178](#)
- Han, T. A. 2009. Evolution prospection with intention recognition via computational logic. Master’s thesis, Technical University of Dresden, Germany and Universidade Nova de Lisbon, Portugal. [12](#), [62](#), [63](#), [85](#), [174](#)
- Han, T. A. and Marques, N. 2010. Towards encoding background knowledge with temporal extent into neural networks. In Bi, Y. and Williams, M.-A., editors, *Knowledge Science, Engineering and Management*, volume 6291 of *Lecture Notes in Computer Science*, pages 63–75. Springer Berlin / Heidelberg. [137](#)
- Han, T. A. and Pereira, L. M. 2010a. Anytime intention recognition via incremental Bayesian network reconstruction. In *AAAI 2010 Fall Symposium on Proactive Assistant Agents (PAA 2010)*, pages 20–25. AAAI. [11](#), [12](#), [15](#), [16](#)
- Han, T. A. and Pereira, L. M. 2010b. Collective intention recognition and elder care. In *AAAI 2010 Fall Symposium on Proactive Assistant Agents (PAA 2010)*, pages 26–31. AAAI. [51](#), [85](#), [177](#), [179](#)
- Han, T. A. and Pereira, L. M. 2010c. Proactive intention recognition for home ambient intelligence. In *Proceedings of 5th Workshop on Artificial Intelligence Techniques for Ambient Intelligence (AITAmI’10), Ambient Intelligence and Smart Environments - vol.8*, pages 91–100. IOS Press. [4](#), [12](#), [13](#), [61](#), [68](#), [76](#), [85](#), [176](#), [177](#)
- Han, T. A. and Pereira, L. M. 2011a. Context-dependent and incremental intention recognition via Bayesian network model construction. (*Submitted for journal publication*). [11](#), [12](#), [13](#), [15](#)
- Han, T. A. and Pereira, L. M. 2011b. Context-dependent incremental intention recognition through Bayesian network model construction. In Nicholson, A., editor, *Proceedings of the Eighth UAI Bayesian Modeling Applications Workshop*

- (*UAI-AW 2011*), volume 818, pages 50–58. CEUR Workshop Proceedings. [4](#), [5](#), [11](#), [12](#), [15](#), [68](#), [120](#), [162](#)
- Han, T. A. and Pereira, L. M. 2011c. Intention-based decision making with evolution prospection. In Antunes, L. and Pinto, H., editors, *Proceedings of 15th Portuguese Conference on Artificial Intelligence (EPIA '2011)*, pages 254–267. LNAI 7026, Springer LNAI. [6](#), [12](#), [13](#), [56](#), [61](#), [90](#)
- Han, T. A., Pereira, L. M., and Santos, F. C. 2011a. Intention recognition promotes the emergence of cooperation. *Adaptive Behavior*, 19(3):264–279. [12](#), [13](#), [20](#), [59](#), [63](#), [97](#), [140](#), [142](#), [156](#), [158](#), [162](#), [163](#), [169](#)
- Han, T. A., Pereira, L. M., and Santos, F. C. 2011b. The role of intention recognition in the evolution of cooperative behavior. In Walsh, T., editor, *Proceedings of the 22nd international joint conference on Artificial intelligence (IJCAI'2011)*, pages 1684–1689. AAAI. [12](#), [13](#), [40](#), [48](#), [59](#), [63](#), [119](#), [140](#), [142](#), [143](#), [156](#), [158](#), [163](#), [169](#)
- Han, T. A., Pereira, L. M., and Santos, F. C. 2012a. Corpus-based intention recognition leads to the emergence of cooperative behavior. In (*submitted to journal of Artificial Life*). [119](#)
- Han, T. A., Pereira, L. M., and Santos, F. C. 2012b. The emergence of commitments and cooperation. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'2012)*. ACM. (Forthcoming). [12](#), [13](#), [139](#), [159](#), [161](#)
- Han, T. A., Pereira, L. M., and Santos, F. C. 2012c. Intention Recognition, Commitment, and The Evolution of Cooperation. In *Proceedings of the 2012 IEEE Congress on Evolutionary Computation*. IEEE Press. (Accepted). [12](#), [13](#), [157](#)
- Han, T. A., Ramli, C. K., and Damásio, C. V. 2008. An implementation of extended P-log using XASP. In *Proceedings of International Conference on Logic Programming (ICLP'08)*, pages 739–743. Springer LNCS 5366. [52](#), [53](#), [63](#)
- Han, T. A., Ramli, C. K., and Damásio, C. V. 2009. Tabling for P-log probabilistic query evaluation. In *New Trends in Artificial Intelligence, Proceedings of 14th Portuguese Conference on Artificial Intelligence (EPIA 2009)*, pages 115–126. [52](#), [63](#)

- Han, T. A., Saptawijaya, A., and Pereira, L. M. 2012d. Moral reasoning under uncertainty. In *Proceedings of the 18th International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR-18)*. Springer LNAI. (Forthcoming). [17](#), [63](#), [85](#)
- Han, T. A., Traulsen, A., and Gokhale, C. S. 2012e. On equilibrium properties of evolutionary multiplayer games with random payoff matrices. *Theoretical Population Biology*. (Accepted). [163](#), [179](#)
- Hardin, G. 1968. The tragedy of the commons. *Science*, 162:1243–1248. [8](#), [90](#), [179](#)
- Harrenstein, P., Brandt, F., and Fischer, F. 2007. Commitment and extortion. In *Proceedings of the 6th international joint conference on Autonomous agents and MultiAgent systems, AAMAS '07*, New York, NY, USA. ACM. [143](#), [176](#)
- Haselton, M. G. and Buss, D. M. 2001. Error management theory: A new perspective on biases in cross-sex mind reading. *Journal of Personality and Social Psychology*, 78(1):81–91. [159](#)
- Hauert, C., Monte, S. D., Hofbauer, J., and Sigmund, K. 2002. Volunteering as red queen mechanism for cooperation in public goods games. *Science*, 296(5570):1129. [DOI: 10.1126/science.1070582]. [179](#)
- Hauert, C. and Schuster, H. 1997. Effects of increasing the number of players and memory size in the iterated prisoner’s dilemma: a numerical approach. *Royal Society B: Biological Sciences*, 264(1381):513–519. [121](#)
- Hauert, C., Traulsen, A., Brandt, H., Nowak, M. A., and Sigmund, K. 2007. Via freedom to coercion: The emergence of costly punishment. *Science*, 316:1905–1907. [91](#), [95](#), [153](#), [160](#)
- Hayes, P. 1973. Acomputation and deduction. In *Proceedings of the 2nd MFCS Symposium, Czechoslovak Academy of Sciences*, pages 105–118. [11](#)
- Heinze, C. 2003. *Modeling Intention Recognition for Intelligent Agent Systems*. PhD thesis, The University of Melbourne, Australia. [4](#), [7](#), [16](#), [18](#), [20](#), [22](#), [52](#), [58](#), [84](#), [85](#), [89](#), [90](#), [99](#), [135](#), [158](#), [162](#)
- Herrmann, B., Thöni, C., and Gächter, S. 2008. Antisocial Punishment Across Societies. *Science*, 319(5868):1362–1367. [153](#)

## BIBLIOGRAPHY

---

- Hirshleifer, J. 2001. Game-theoretic interpretations of commitment. In Nesse, R. M., editor, *Evolution and the capacity for commitment*, pages 77–93. New York: Russell Sage. [143](#), [144](#), [161](#), [176](#)
- Hirshleiffer, J. 1999. There are many evolutionary pathways to cooperation. *Journal of Bioeconomics*, (1):73–93. [140](#), [155](#)
- Hofbauer, J. and Sigmund, K. 1998. *Evolutionary Games and Population Dynamics*. Cambridge University Press. [8](#), [9](#), [43](#), [44](#), [45](#), [49](#), [59](#), [91](#), [93](#), [131](#), [140](#), [141](#), [143](#), [160](#)
- Hofmann, L.-M., Chakraborty, N., and Sycara, K. 2011. The evolution of cooperation in self-interested agent societies: a critical study. In *The 10th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '11*, pages 685–692. [143](#)
- Hong, J. 2001. Goal recognition through goal graph analysis. *Journal of Artificial Intelligence Research*, 15:1–30. [4](#), [19](#), [85](#)
- Horvitz, E., Breese, J., Heckerman, D., Hovel, D., and Rommelse, K. 1998. The lumière project: Bayesian user modeling for inferring the goals and needs of software users. In Cooper, G. F., , and Moral, S., editors, *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence*, pages 256–265. Morgan Kaufmann. [17](#), [59](#)
- Huber, M. and Simpson, R. 2004. Recognizing the plans of screen reader users. In *In Workshop on Modeling Other Agents from Observations (MOO2004)*. [19](#)
- Humphrey, N. K. 1999. Varieties of altruism. *Social Research*, (64):199–209. [140](#)
- Iacoboni, M., Molnar-Szakacs, I., Gallese, V., Buccino, G., Mazziotta, J. C., and Rizzolatti, G. 2005. *PLoS Biol*, 3(3):e79. [4](#), [120](#)
- Imhof, L. A., Fudenberg, D., and Nowak, M. A. 2005. Evolutionary cycles of cooperation and defection. *Proceedings of the National Academy of Sciences of the United States of America*, 102:10797–10800. [8](#), [91](#), [95](#), [98](#), [110](#), [146](#), [160](#), [164](#)
- Imhof, L. A., Fudenberg, D., and Nowak, M. A. 2007. Tit-for-tat or win-stay, lose-shift? *Journal of Theoretical Biology*, 247(3):574 – 580. [115](#)

## BIBLIOGRAPHY

---

- Johnson, D. D. P. and Fowler, J. H. 2011. The evolution of overconfidence. *Nature*, 477(7364):317–320. [169](#)
- Johnson-Laird, P. N. 2010. Mental models and human reasoning. *Proceedings of the National Academy of Sciences of the United States of America*, 107(43):18243–18250. [7](#)
- Kakas, A., Kowalski, R., and Toni, F. 1993. Abductive logic programming. *Journal of Logic and Computation*, 2(6):719–770. [11](#), [64](#)
- Kaminka, G. A., Tambe, D. V. P. M., Pynadath, D. V., and Tambe, M. 2002. Monitoring teams by overhearing: A multi-agent plan-recognition approach. *Journal of Artificial Intelligence Research*, 17:2002. [63](#), [69](#)
- Kaminka, G. A., Wendler, J., and Ronen, G. 2001. New challenges in multi-agent intention recognition: Extended abstract. In *Proceedings of AAAI Fall symposium on Intent Inference for Collaborative Tasks*. AAAI Press. [177](#)
- Kandori, M., Mailath, G. J., and Rob, R. 1993. Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56. [146](#), [148](#)
- Kanno, T., Nakata, K., and Furuta, K. 2003. A method for team intention inference. *International Journal of Human–Computer Studies*, 33(4):393–413. [177](#), [179](#)
- Kareev, Y. 1995. Through a narrow window: working memory capacity and the detection of covariation. *Cognition*, 56(3):263 – 269. [7](#)
- Karlin, S. and Taylor, H. E. 1975. *A First Course in Stochastic Processes*. Academic Press, New York. [95](#)
- Kautz, H. and Allen, J. 1986a. Generalized plan recognition. In *Proceedings of the Conference of the American Association of Artificial Intelligence (AAAI'1986)*, pages 32–38. AAAI. [7](#), [19](#), [89](#)
- Kautz, H. A. and Allen, J. F. 1986b. Generalized plan recognition. In *Proceedings of 1986 Conference of the American Association for Artificial Intelligence (AAAI'1986)*, pages 32–37. AAAI. [158](#)
- Keng, N. and Yun, D. Y. Y. 1989. A planning / scheduling methodology for the constrained resources problem. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 999–1003. [32](#)

## BIBLIOGRAPHY

---

- Kim, K.-J. and Lipson, H. 2009. Towards a "theory of mind" in simulated robots. In *Proceedings of the 11th Annual Conference Companion on Genetic and Evolutionary Computation Conference: Late Breaking Papers*, GECCO '09, pages 2071–2076, New York, NY, USA. ACM. 121
- Kording, K. P. and Wolpert, D. M. 2004. Bayesian integration in sensorimotor learning. *Nature*, 427(6971):244–247. 98
- Kowalski, R. 2011. *Computational Logic and Human Thinking: How to be Artificially Intelligent*. Cambridge University Press. 11
- Kraus, S. 1997. Negotiation and cooperation in multi-agent environments. *Artificial Intelligence*, 94(1-2):79–98. 7, 90
- Laskey, K. 2008. MEBN: A language for first-order Bayesian knowledge bases. *Artificial Intelligence*, 172(2-3):140 – 178. 18, 29, 51
- Laskey, K. and Mahoney, S. 1997. Network fragments: Representing knowledge for constructing probabilistic models. In *Proceedings of Conference on Uncertainty in Artificial Intelligence (UAI'97)*. Morgan Kaufmann. 18
- Leeds, B. A. 2003. Alliance Reliability in Times of War: Explaining State Decisions to Violate Treaties. *International Organization*, 57(04):801–827. 170
- Lesh, N. 1998. *Scalable and Adaptive Goal Recognition*. PhD thesis, University of Washington. 4, 32, 35, 60, 85
- Lindgren, K. and Nordahl, M. G. 1994. Evolutionary dynamics of spatial games. *Physica D: Nonlinear Phenomena*, 75(1-3):292 – 309. 9, 91
- Linux-Plan-Corpus. <http://www.cs.rochester.edu/research/cisd/resources/linux-plan/>. Last access: November 21 2010. 35, 36
- Lloyd, J. W. 1987. *Foundations of Logic Programming*. Springer-Verlag. 11
- Lockhart, C. 1978. Flexibility and commitment in international conflicts. *International Studies Quarterly*, 22(4):545–568. 170
- Macy, M. W. and Flache, A. 2002. Learning dynamics in social dilemmas. *Proceedings of the National Academy of Sciences of the United States of America*, 99:7229–7236. 92, 121

## BIBLIOGRAPHY

---

- Madani, O., Bui, H., and Yeh, E. 2009. Efficient online learning and prediction of users' desktop actions. In *Proceedings of the 21st international joint conference on Artificial intelligence (IJCAI'2009)*, pages 1457–1462. [17](#), [59](#)
- Mahoney, S. M. and Laskey, K. B. 1996. Network engineering for complex belief networks. In *Proceedings of Conference on Uncertainty in Artificial Intelligence (UAI'96)*, pages 389–396. Morgan Kaufmann. [18](#)
- Malle, B., Moses, L., and Baldwin, D. 2003. *Intentions and intentionality: foundations of social cognition*. Bradford Books. MIT Press. [177](#)
- Mao, W. and Gratch, J. 2004. A Utility-Based Approach to Intention Recognition. In *AAMAS 2004 Workshop on Agent Tracking: Modeling Other Agents from Observations*. [4](#), [85](#)
- Masuda, N. and Ohtsuki, H. 2009. *Bulletin of Mathematical Biology*, 71:1818–1850. [98](#)
- Maynard-Smith, J. 1982. *Evolution and the Theory of Games*. Cambridge University Press, Cambridge. [9](#), [91](#), [160](#)
- Meltzoff, A. N. 2005. Imitation and other minds: the “like me” hypothesis. In *Perspectives on imitation: From neuroscience to social science. Imitation, human development, and culture*, pages 55–77. Cambridge, MA: MIT Press. [120](#), [124](#), [126](#), [175](#)
- Meltzoff, A. N. 2007. The framework for recognizing and becoming an intentional agent. *Acta Psychol Amst*, 124(1):26–43. [6](#), [90](#), [120](#), [121](#), [124](#), [158](#), [175](#)
- Milinski, M., Semmann, D., Krambeck, H. J., and Marotzke, J. 2006. Stabilizing the Earth's climate is not a losing game: Supporting evidence from public goods experiments. *Proceedings of the National Academy of Sciences of the United States of America*, 103:3994–3998. [50](#), [170](#)
- Miller, G. F. and Todd, P. M. 1998. Mate choice turns cognitive. *Trends in Cognitive Sciences*, 2(5):190 – 198. [159](#)
- Morgan, T. J. H., Rendell, L. E., Ehn, M., Hoppitt, W., and Laland, K. N. 2011. The evolutionary basis of human social learning. *Proceedings of the Royal Society B: Biological Sciences*. [50](#)

## BIBLIOGRAPHY

---

- Nakahara, K. and Miyashita, Y. 2005. Understanding intentions: Through the looking glass. *Science*, 308(5722):644–645. [4](#), [120](#)
- Narendra, K. and Thathachar, M. 1989. *Learning Automata: An introduction*. Prentice-Hall International, Boston, USA. [91](#)
- Natarajan, S., Tadepalli, P., Dietterich, T. G., and Fern, A. 2008. Learning first-order probabilistic models with combining rules. *Annals of Mathematics and Artificial Intelligence*, 54:223–256. [18](#)
- Nesse, R. M. 2001a. *Evolution and the capacity for commitment*. Russell Sage Foundation series on trust. Russell Sage. [10](#), [159](#)
- Nesse, R. M. 2001b. Natural selection and the capacity for subjective commitment. In Nesse, R. M., editor, *Evolution and the capacity for commitment*, pages 1–44. New York: Russell Sage. [9](#), [10](#), [89](#), [140](#), [141](#), [143](#), [144](#), [153](#), [156](#), [159](#), [161](#), [176](#)
- Nowak, M. A. 2006a. *Evolutionary Dynamics: Exploring the Equations of Life*. Harvard University Press, Cambridge, MA. [8](#), [9](#), [90](#), [91](#), [98](#), [160](#), [175](#)
- Nowak, M. A. 2006b. Five rules for the evolution of cooperation. *Science*, 314(5805):1560. DOI: 10.1126/science.1133755. [8](#), [90](#), [110](#), [131](#), [132](#), [140](#), [143](#), [146](#), [148](#), [169](#)
- Nowak, M. A., Sasaki, A., Taylor, C., and Fudenberg, D. 2004. Emergence of cooperation and evolutionary stability in finite populations. *Nature*, 428:646–650. [98](#), [132](#), [146](#), [148](#), [164](#)
- Nowak, M. A. and Sigmund, K. 1992. Tit for tat in heterogeneous populations. *Nature*, 355:250–253. [8](#), [41](#), [91](#), [93](#)
- Nowak, M. A. and Sigmund, K. 1993. A strategy of win-stay, lose-shift that outperforms tit-for-tat in prisoner’s dilemma. *Nature*, 364:56–58. [41](#), [93](#)
- Nowak, M. A. and Sigmund, K. 2005. Evolution of indirect reciprocity. *Nature*, 437(1291-1298). [8](#), [91](#), [111](#), [142](#), [178](#)
- Ohtsuki, H. and Iwasa, Y. 2006. The leading eight: Social norms that can maintain cooperation by indirect reciprocity. *Journal of Theoretical Biology*, 239(4):435 – 444. [98](#), [111](#)

## BIBLIOGRAPHY

---

- Ohtsuki, H., Iwasa, Y., and Nowak, M. A. 2009. Indirect reciprocity provides only a narrow margin of efficiency for costly punishment. *Nature*, 457(7601):79–82. [91](#), [153](#), [178](#)
- Oyama, S. 2000. *Evolution's Eye: A Systems View of the Biology-Culture Divide*. Durham, N.C.: Duke University Press. [140](#)
- Pacheco, J. M., Santos, F. C., and Chalub, F. A. C. C. 2006. Stern-judging: A simple, successful norm which promotes cooperation under indirect reciprocity. *PLoS Comput. Biol.*, 2(12):e178. [8](#), [91](#), [111](#), [142](#), [144](#)
- Pacheco, J. M., Santos, F. C., Souza, M. O., and Skyrms, B. 2009. Evolutionary dynamics of collective action in n-person stag hunt dilemmas. *Proc. R. Soc. B*, 276:315–321. [179](#)
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann. [11](#), [18](#), [21](#), [23](#), [27](#), [28](#), [29](#), [98](#), [124](#)
- Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge University Press. [11](#), [18](#), [21](#), [23](#), [98](#)
- Pereira, L. M., Dell'Acqua, P., and Lopes, G. 2011a. Inspecting and preferring abductive models. In *Handbook on Reasoning-based Intelligent Systems*. World Scientific Publishers. (in press). [11](#), [63](#), [65](#)
- Pereira, L. M., Dell'Acqua, P., and Lopes, G. 2011b. On preferring and inspecting abductive models. In *Proceedings of 11th International Symposium on Practical Aspects of Declarative Languages*, pages 1–15. Springer LNCS 5418. [65](#)
- Pereira, L. M. and Han, T. A. 2009a. Evolution prospection. In *Proceedings of International Symposium on Intelligent Decision Technologies (KES-IDT'09)*, pages 51–63. Springer Studies in Computational Intelligence 199. [12](#), [62](#), [63](#), [79](#), [80](#), [154](#), [174](#)
- Pereira, L. M. and Han, T. A. 2009b. Evolution prospection in decision making. *Intelligent Decision Technologies*, 3(3):157–171. [11](#), [12](#), [56](#), [62](#), [63](#), [70](#), [71](#), [75](#), [79](#), [80](#), [82](#), [85](#), [154](#), [174](#)
- Pereira, L. M. and Han, T. A. 2009c. Intention recognition via causal Bayes networks plus plan generation. In *Progress in Artificial Intelligence, Proceedings of 14th*

## BIBLIOGRAPHY

---

- Portuguese International Conference on Artificial Intelligence (EPIA'09)*, pages 138–149. Springer LNAI 5816. [5](#), [7](#), [18](#), [20](#), [22](#), [51](#), [52](#), [69](#), [73](#), [99](#), [158](#)
- Pereira, L. M. and Han, T. A. 2011a. Elder care via intention recognition and evolution propection. In *select extended papers from the 18th International Conference on Applications of Declarative Programming and Knowledge Management (INAP'09)*, pages 170–187. Springer, LNAI 6547. [4](#), [12](#), [13](#), [25](#), [61](#), [68](#), [76](#), [85](#), [176](#)
- Pereira, L. M. and Han, T. A. 2011b. Intention recognition with evolution propection and causal Bayesian networks. In *Computational Intelligence for Engineering Systems 3: Emergent Applications*, pages 1–33. Springer. [4](#), [7](#), [12](#), [13](#), [15](#), [18](#), [20](#), [22](#), [25](#), [51](#), [52](#), [56](#), [61](#), [68](#), [73](#), [90](#), [98](#), [99](#), [158](#), [177](#)
- Pereira, L. M. and Lopes, G. 2009. Prospective logic agents. *International Journal of Reasoning-based Intelligent Systems (IJRIS)*, 1(3/4):200–208. [63](#)
- Pfeffer, A., Koller, D., Milch, B., and Takusagawa, K. T. 1999. Pook: A system for probabilistic object-oriented knowledge representation. In *Proceedings of Conference on Uncertainty in Artificial Intelligence (UAI'93)*. [18](#)
- Pinker, S., Nowak, M. A., and Lee, J. J. 2008. The logic of indirect speech. *Proceedings of the National Academy of Sciences of the United States of America*, 105(3):833–838. [5](#), [7](#), [84](#), [90](#)
- Premack, D. and Woodruff, G. 1978. Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, 1(4):515–526. [120](#), [179](#)
- Pynadath, D. V. and Wellman, M. P. 1995. Accounting for context in plan recognition, with application to traffic monitoring. In *Proceedings of Conference on Uncertainty in Artificial Intelligence (UAI'95)*. [4](#), [19](#), [20](#), [28](#), [32](#), [49](#), [59](#), [60](#), [85](#)
- Raihani, N. and Aitken, D. 2011. Uncertainty, rationality and cooperation in the context of climate change. *Climatic Change*, 108(1):47–55. [170](#)
- Ramírez, M. and Geffner, H. 2010. Probabilistic plan recognition using off-the-shelf classical planners. In *Proceedings of the Conference of the American Association of Artificial Intelligence (AAAI'2010)*. [17](#)
- Rand, D. G. and Nowak, M. A. 2011. The evolution of antisocial punishment in optional public goods games. *Nature Communications*, 2:434. [91](#), [153](#), [178](#)

## BIBLIOGRAPHY

---

- Rao, A. S. and Georgeff, M. P. 1995. BDI Agents: from theory to practice. In *Proceeding of First International Conference on Multiagent Systems*, pages 312–319. [62](#)
- Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M. W., Fogarty, L., Ghirlanda, S., Lillicrap, T., and Laland, K. N. 2010. Why copy others? insights from the social learning strategies tournament. *Science*, 328(5975):208–213. [44](#), [45](#), [49](#), [50](#), [94](#)
- Richerdson, P. J. and Boyd, R. 2006. *Not By Genes Alone: How Culture Transforms Human Evolution*. The University of Chicago Press. [44](#), [120](#), [124](#)
- Riolo, R. L., Cohen, M. D., and Robert, A. 2001. Evolution of cooperation without reciprocity. *Nature*, pages 441–443. [178](#)
- Rizzolatti, G. and Craighero, L. 2004. The mirror-neuron system. *Annual Review of Neuroscience*, 27:169–192. [4](#), [120](#)
- Robson, A. 1990. Efficiency in evolutionary games: Darwin, Nash, and the secret handshake. *Journal of Theoretical Biology*, 144(3):379–396. [7](#), [9](#), [90](#), [140](#), [143](#), [158](#)
- Roy, O. 2009a. Intentions and interactive transformations of decision problems. *Synthese*, 169(2):335–349. [6](#), [90](#)
- Roy, O. 2009b. *Thinking before Acting: Intentions, Logic, Rational Choice*. PhD thesis, ILLC Dissertation Series DS-2008-03, Amsterdam. [9](#), [16](#), [17](#), [32](#), [55](#), [62](#), [63](#), [84](#), [158](#)
- Roy, P., Bouchard, B., Bouzouane, A., and Giroux, S. 2007. A hybrid plan recognition model for alzheimer’s patients: interleaved-erroneous dilemma. In *Proceedings of IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, pages 131–137. [4](#), [6](#), [76](#), [77](#), [83](#), [85](#)
- Ruse, M. 2001. Morality and commitment. In Nesse, R. M., editor, *Evolution and the capacity for commitment*, pages 221–236. New York: Russell Sage. [9](#), [140](#), [143](#), [144](#), [159](#), [161](#)
- Russell, S. J. and Norvig, P. 2003. *Artificial Intelligence: A Modern Approach*. Pearson Education. [62](#), [124](#), [136](#)

## BIBLIOGRAPHY

---

- Sadri, F. 2010. Intention recognition with event calculus graphs and weight of evidence. In *In Proceedings 4th International Workshop on Human Aspects in Ambient Intelligence*. (accepted). [5](#)
- Sadri, F. 2011a. Ambient intelligence: A survey. *ACM Computing Survey*, 43:36:1–36:66. [76](#), [77](#), [85](#), [176](#), [177](#)
- Sadri, F. 2011b. Logic-based approaches to intention recognition. In *Handbook of Research on Ambient Intelligence: Trends and Perspectives*, pages 346–375. [4](#), [5](#), [16](#), [19](#), [85](#), [177](#)
- Sandholm, T. W. and Crites, R. H. 1995. Multiagent reinforcement learning in the iterated prisoner’s dilemma. *Biosystems*, 37:147–166. [121](#)
- Santos, F. C. and Pacheco, J. M. 2005. Scale-free networks provide a unifying framework for the emergence of cooperation. *Phys. Rev. Lett.*, 95:098104. [8](#), [91](#)
- Santos, F. C. and Pacheco, J. M. 2011. Risk of collective failure provides an escape from the tragedy of the commons. *Proceedings of the National Academy of Sciences of the United States of America*, 108(26):10421–10425. [50](#), [170](#)
- Santos, F. C., Pacheco, J. M., and Lenaerts, T. 2006a. Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proceedings of the National Academy of Sciences of the United States of America*, 103:3490–3494. [8](#), [91](#), [92](#)
- Santos, F. C., Pacheco, J. M., and Lenaerts, T. 2006b. Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proceedings of the National Academy of Sciences of the United States of America*, 103:3490–3494. [144](#), [178](#)
- Santos, F. C., Pacheco, J. M., and Skyrms, B. 2011. Co-evolution of pre-play signaling and cooperation. *Journal of Theoretical Biology*, 274(1):30–35. [7](#), [9](#), [90](#), [95](#), [140](#), [143](#), [159](#)
- Santos, F. C., Santos, M. D., and Pacheco, J. M. 2008. Social diversity promotes the emergence of cooperation in public goods games. *Nature*, 454:214–216. [9](#), [91](#), [140](#), [143](#), [179](#)

## BIBLIOGRAPHY

---

- Sayama, H., Farrell, D. L., and Dionne, S. D. 2011. The effects of mental model formation on group decision making: An agent-based simulation. *Complexity*, 16(3):49–57. [121](#)
- Schelling, T. C. 1990. *The strategy of conflict*. London: Oxford University Press. [143](#), [144](#), [154](#), [155](#), [161](#), [170](#), [176](#)
- Schmidt, C., Sridharan, N., and Goodson, J. 1978. The plan recognition problem: An intersection of psychology and artificial intelligence. *Artificial Intelligence*, 11(1978):45–83. [19](#)
- Schrempf, O. C., Albrecht, D., and Hanebeck, U. D. 2007. Tractable probabilistic models for intention recognition based on expert knowledge. In *Proceedings of International Conference on Intelligent Robots and Systems*, pages 1429–1434. IEEE. [20](#)
- Searle, J. R. 1995. *The Construction of Social Reality*. New York: The Free Press. [63](#), [171](#)
- Searle, J. R. 2010. *Making the Social World: The Structure of Human Civilization*. Oxford University Press. [7](#), [62](#), [63](#), [156](#), [171](#), [177](#)
- Selten, R. 2001. What is bounded rationality? In *Bounded rationality: The adaptive toolbox*, pages 13–36. Cambridge, MA: MIT Press. [7](#)
- Shennan, S. 2002. *Genes, Memes and Human History: Darwinian Archaeology and Cultural Evolution*. London: Thames and Hudson Ltd. [120](#), [124](#)
- Sigmund, K. 2010. *The Calculus of Selfishness*. Princeton University Press. [8](#), [40](#), [41](#), [44](#), [49](#), [50](#), [59](#), [90](#), [92](#), [93](#), [94](#), [98](#), [107](#), [110](#), [113](#), [114](#), [131](#), [132](#), [133](#), [140](#), [141](#), [143](#), [144](#), [145](#), [146](#), [148](#), [161](#), [169](#), [175](#)
- Sigmund, K., Silva, H. D., Traulsen, A., and Hauert, C. 2010. Social learning promotes institutions for governing the commons. *Nature*, 466:7308. [9](#), [43](#), [44](#), [45](#), [50](#), [91](#), [149](#), [160](#)
- Simon, H. A. 1957. A behavioral model of rational choice. In *Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting*. New York: Wiley. [7](#), [135](#)

## BIBLIOGRAPHY

---

- Simon, H. A. 1990. Invariants of human behavior. *Annual review of psychology*, 41:1–19. [7](#), [135](#)
- Singh, M. P. 1991. Intentions, commitments and rationality. In *In 13th Annual Conference of the Cognitive Science Society*, pages 493–498. [16](#), [17](#), [62](#), [84](#), [99](#)
- Singla, P. and Mooney, R. J. 2011. Abductive markov logic for plan recognition. In *Proceedings of the Conference of the American Association of Artificial Intelligence (AAAI'2011)*, pages 1069–1075. [19](#), [35](#), [39](#)
- Skyrms, B. 1996. *Evolution of the Social Contract*. Cambridge University Press. [156](#), [158](#)
- Skyrms, B. 2003. *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press. [156](#)
- Skyrms, B. 2010. *Signals: Evolution, Learning, and Information*. Oxford University Press. [7](#), [9](#), [90](#), [140](#), [143](#), [156](#), [158](#)
- Spirtes, P., Glymour, C., and Scheines, R. 2000. *Causation, Prediction, and Search (Second Edition)*. MIT Press. [98](#)
- Srinivas, S. 1993. A generalization of the noisy-or model. In *Proceedings of Conference on Uncertainty in Artificial Intelligence (UAI'93)*. Morgan Kaufmann. [18](#), [27](#), [28](#)
- Sukthankar, G. and Sycara, K. 2008. Robust and efficient plan recognition for dynamic multi-agent teams. In *Proceedings of International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1383–1388. [85](#), [179](#)
- Sukthankar, G. R. 2007. *Activity Recognition for Agent Teams*. PhD thesis, Robotics Institute School of Computer Science, Carnegie Mellon University. [85](#), [177](#)
- Sunzi, R. D. and Sawyer, M.-C. 1994. *The art of war = [Sun-tzu ping fa]*. History and warfare. Westview Press, Boulder. [135](#)
- Szabó, G. and Fáth, G. 2007. Evolutionary games on graphs. *Phys Rep*, 97-216(4-6). [9](#), [91](#), [178](#)
- Szabó, G. and Toke, C. 1998. *Phys Rev E*, 58:69–73. [94](#), [132](#)

## BIBLIOGRAPHY

---

- Tahboub, K. A. 2006. Intelligent human-machine interaction based on dynamic Bayesian networks probabilistic intention recognition. *J. Intell. Robotics Syst.*, 45:31–52. [4](#), [16](#), [18](#), [19](#), [20](#), [22](#), [52](#), [58](#), [85](#), [98](#), [99](#)
- Takano, M. and Arita, T. 2006. Asymmetry between Even and Odd Levels of Recursion in a Theory of Mind. pages 405–411. Proceedings of ALife X. [121](#)
- Tenenbaum, J. B. and Griffiths, T. L. 2001. Generalization, similarity, and bayesian inference. *Behavioral and Brain Sciences*. [98](#)
- Todd, P. 2001. Fast and frugal heuristics for environmentally bounded minds. In *Bounded rationality: The adaptive toolbox*, pages 51–70. Cambridge, MA: MIT Press. [7](#), [135](#)
- Tomasello, M. 1999. *The Cultural Origins of Human Cognition*. Harvard University Press. [6](#), [7](#), [90](#), [158](#)
- Tomasello, M. 2008. *Origins of Human Communication*. MIT Press. [6](#), [7](#), [84](#), [90](#), [121](#), [126](#), [158](#)
- Traulsen, A., Hauert, C., De Silva, H., Nowak, M. A., and Sigmund, K. 2009. Exploration dynamics in evolutionary games. *Proc. Natl. Acad. Sci. USA*, 106(3):709–712. [95](#)
- Traulsen, A. and Nowak, M. A. 2006. Evolution of cooperation by multilevel selection. *Proceedings of the National Academy of Sciences of the United States of America*, 103(29):10952. [8](#), [90](#)
- Traulsen, A., Nowak, M. A., and Pacheco, J. M. 2006. Stochastic dynamics of invasion and fixation. *Phys. Rev. E*, 74:11909. [44](#), [94](#), [95](#), [132](#)
- Traulsen, A., Pacheco, J. M., and Nowak, M. A. 2007. Pairwise comparison and selection temperature in evolutionary game dynamics. *J. Theor. Biol.*, 246:522–529. [94](#)
- Traulsen, A., Semmann, D., Sommerfeld, R. D., Krambeck, H.-J., and Milinski, M. 2010. Human strategy updating in evolutionary games. *Proceedings of the National Academy of Sciences of the United States of America*, 107:2962–2966. [50](#), [134](#)

## BIBLIOGRAPHY

---

- Trivers, R. 2011. *Deceit and Self-Deception: Fooling Yourself the Better to Fool Others*. Penguin Books, Limited. [7](#), [90](#), [154](#), [179](#)
- Trivers, R. L. 1971. The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46:35–57. [8](#), [91](#), [98](#), [142](#), [143](#)
- Tu, P. H., Son, T. C., and Baral, C. 2007. Reasoning and planning with sensing actions, incomplete information, and static causal laws using answer set programming. *Theory Pract. Log. Program.*, 7:377–450. [69](#)
- Tu, P. H., Son, T. C., Gelfond, M., and Morales, A. R. 2011. Approximation of action theories and its application to conformant planning. *Artificial Intelligence*, 175(1):79 – 119. [69](#)
- Tuomela, R. 2005. We-intentions revisited. *Philosophical Studies*, 125(3):327–369. [177](#)
- Turaga, P., Chellappa, R., Subrahmanian, V. S., and Udrea, O. 2008. Machine Recognition of Human Activities: A Survey. *Circuits and Systems for Video Technology, IEEE Transactions on*, 18(11):1473–1488. [17](#)
- Tuyls, K. and Parsons, S. 2007. What evolutionary game theory tells us about multiagent learning. *Artificial Intelligence*, 171:406–416. [43](#), [50](#)
- van Hees, M. and Roy, O. 2008. Intentions and plans in decision and game theory. In *Reasons and intentions*, pages 207–226. Ashgate Publishers. [6](#), [62](#), [90](#)
- Van Segbroeck, S., de Jong, S., Nowé, A., Santos, F. C., and Lenaerts, T. 2010. Learning to coordinate in complex networks. *Adaptive Behavior*, 18:416–427. [7](#), [90](#), [91](#), [121](#)
- Vukov, J., Santos, F. C., and Pacheco, J. M. 2011. Incipient Cognition Solves the Spatial Reciprocity Conundrum of Cooperation. *PLoS ONE*, 6(3):e17939. [98](#)
- Wang, S., Szalay, M. S., Zhang, C., and Csermely, P. 2008. Learning and innovative elements of strategy adoption rules expand cooperative network topologies. *PLoS ONE*, 3(4):e1917. [98](#)
- West, S., Griffin, A., and Gardner, A. 2007. Evolutionary explanations for cooperation. *Current Biology*, 17:R661–R672. [90](#)

## BIBLIOGRAPHY

---

- West, S. A., Pen, I., and Griffin, A. S. 2002. Cooperation and Competition Between Relatives. *Science*, 296(5565):72–75. [8](#), [90](#), [178](#)
- Whiten, A. 1991. *Natural theories of mind: evolution, development, and simulation of everyday mindreading*. B. Blackwell. [120](#), [179](#)
- Wikipedia. Burglar alarm: <http://en.wikipedia.org/wiki/burglaralarm>. Last access: November 17, 2011. [80](#)
- Wilensky, R. 1983. *Planning and Understanding*. Addison-Wesley. [19](#)
- Winikoff, M. 2007. Implementing commitment-based interactions. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, AAMAS '07, pages 868–875, New York, NY, USA. ACM. [143](#), [176](#)
- Wolpert, D. and Ghahramani, Z. 2004. *Bayes rule in perception, action and cognition*. Oxford University Press. [98](#)
- Woodward, A. L., Sommerville, J. A., Gerson, S., Henderson, A. M., and Buresh, J. 2009. The emergence of intention attribution in infancy. In Ross, B. H., editor, *The Psychology of Learning and Motivation*, volume 51 of *Psychology of Learning and Motivation*, pages 187 – 222. Academic Press. [6](#), [89](#), [90](#), [120](#), [121](#), [124](#), [126](#), [158](#), [175](#)
- Wooldridge, M. 2000. *Reasoning About Rational Agents*. MIT Press. [62](#)
- Wooldridge, M. 2002. Reasoning about rational agents. *The Journal of Artificial Societies and Social Simulation*, 5(1). [62](#)
- Wooldridge, M. and Jennings, N. R. 1999. The cooperative problem-solving process. In *Journal of Logic and Computation*, pages 403–417. [143](#), [154](#), [176](#)
- Wu, B., Altrock, P. M., Wang, L., and Traulsen, A. 2010. Universality of weak selection. *Phys. Rev. E*, 82(4):046106. [132](#)
- Xiang, Y. P. and Poh, K. L. 2005. Constructing Bayesian network in a changing world. In *Proceedings of AAAI Spring Symposium on Challenges to Decision Support in a Changing World*, pages 108–113. AAAI. [18](#), [51](#)
- XSB 2009. The XSB system version 3.2 vol. 2: Libraries, interfaces and packages. [52](#), [54](#), [56](#), [63](#)

## BIBLIOGRAPHY

---

Zanlungo, F. 2007. *Microscopic dynamics of artificial life systems*. PhD thesis. [121](#)