A Work Project, presented as part of the requirements for the Award of a master's degree in business Analytics from the Nova School of Business and Economics.

TITLE OF WORK PROJECT: A Roadmap to Ai for Ngos: A Cáritas Case Study - An Adapted Data Governance Framework for Ai Readiness in Federated Nonprofit Organizations

STUDENT FULL NAME: Ryan Myers

Work project carried out under the supervision of:

Bernardo Forbes Godinho Costa

11/12/2025

Abstract

INDIVIDUAL CONTRIBUTION

Federated nonprofits face unique AI implementation challenges due to data fragmentation across autonomous entities. Existing governance frameworks address corporate hierarchies rather than decentralized nonprofits requiring voluntary coordination. This research adapts and synthesizes established data governance frameworks (Alation, 2024; DAMA International, 2017; Khatri & Brown, 2010) specifically for federated nonprofit organizations preparing for AI implementation. The adapted framework comprises three components: a three-tier governance architecture defining decision rights across central, federated, and autonomous levels; a five-level maturity model establishing AI readiness thresholds; and implementation principles balancing standardization with organizational autonomy.

GROUP PROJECT APPLICATION

The framework was validated through a field lab engagement with Cáritas Portugal, a federation of 20 autonomous dioceses. The team conducted organizational assessment, developed technical solutions for data aggregation and quality monitoring, and created implementation recommendations. Assessment revealed Level 1 maturity despite sophisticated technical infrastructure, confirming that technical platforms alone do not ensure governance readiness. The team developed a 24-month roadmap addressing governance structure, data standardization, and early AI use cases.

Keywords: Data governance, federated nonprofits, AI readiness, organizational autonomy, maturity model

This master's thesis was developed under the Social Sector Database project, led by the Nova SBE Data, Operations & Technology Knowledge Center. The project is part of the Social Equity Initiative - a partnership between the "la Caixa" Foundation, BPI Bank, and Nova School of Business and Economics.

GROUP PROJECT

Introduction

Cáritas Portuguesa maintains a unique position of influence throughout Portugal's social intervention and humanitarian assistance sector. The organization serves as the Catholic Church's national body to coordinate charitable work between ecclesiastical missions and social policy frameworks and civil society organizations. The analysis of Cáritas Portuguesa becomes essential because of its historical significance and its current role in addressing evolving poverty patterns and social weaknesses and community strength. The thesis requires an accurate introduction to Cáritas Portuguesa because it establishes the institutional context needed for future analysis.

Social requirements in Portugal have developed into more complex issues throughout the previous several decades. The combination of economic downturns with increasing expenses and population changes and migration patterns and environmental threats requires social protection systems to achieve both close relationships and flexible responses and unified coordination. The nationwide institution Cáritas Portuguesa leads local network activation through its national framework which supports international partnerships. The organization maintains its effectiveness through its extensive reach and decentralized system and established reputation which enables it to deliver transparent interventions that fulfill social objectives.

Historical Emergence and Institutional Identity

Cáritas Portuguesa established itself in 1956 during a period when Portugal lacked sufficient welfare programs, and its citizens faced widespread poverty in both rural and urban areas while lacking access to public social services. The institution began as a national charity initiative grounded in Catholic Church traditions to support marginalized groups. It was created to organize charitable work, establish a system for diocesan coordination, and strengthen the Church's capacity

to respond to emergencies and persistent social challenges.

The entity bases its identity on Catholic teachings about caritas, which require practical expression through solidarity and service to promote human dignity. It stands apart from secular service organizations because its foundation is rooted in Catholic principles. Its work is guided by religious teachings that define charity as an essential aspect of Christian life and a catalyst for social transformation.

Cáritas emphasizes in its official documents that it supports the Church's mission by helping build caring communities that protect vulnerable people, advocate for justice, and promote the dignity of every person. Cáritas Portuguesa functions as a social sector institution operating under both ecclesiastical and civic frameworks.

Over time, the entity has strengthened its operational capabilities through professionalization and expanded its service range to address emerging patterns of poverty and social fragility. It upholds its ethical foundations through its dual role as a religious institution and a civic actor, allowing it to collaborate with public bodies, private organizations, and international partners.

Organizational Structure and Governance Model

Cáritas Portuguesa operates through several organizational levels that combine national coordination with local autonomy. The governance model follows the principle of subsidiarity because it supports decision making at the level that is closest to those who require assistance.

The national office acts as the central institution that brings together the twenty diocesan Cáritas structures. It prepares national activity plans, maintains relations with the Conferência Episcopal Portuguesa, represents Portugal within Caritas Internationalis and Caritas Europa, and supports communication across dioceses. Although the national office promotes coherence in administrative matters such as training and general project management procedures, its responsibilities do not include defining technical data concepts or creating formal data rules. These areas remain uncoordinated across dioceses, which is consistent with findings that Cáritas has an established organizational governance system but does not have a formal structure dedicated to data governance.

The governance model relies on four statutory bodies: the National Board, the General Council, the Permanent Commission, and the Fiscal Council. These bodies carry out tasks related to institutional planning, oversight, representation, and compliance. Their scope concerns overall organizational administration rather than data management. This means that authority over data processes is not formally assigned at the national level.

Each of the twenty Cáritas Diocesanas operates with legal and canonical autonomy. This autonomy allows each diocese to adapt national orientations to its own context and develop programs that match local circumstances. However, it also contributes to variation in data practices because dioceses manage AidHound data independently. There is no national system that harmonizes these procedures.

At the local level, Cáritas Paroquiais units and parish groups maintain community presence. They identify needs, provide immediate forms of assistance, and refer more complex cases to diocesan services. Their proximity to communities strengthens the local dimension of Cáritas but also results in diverse data routines, since information is recorded according to the practices of each parish and diocese.

This governance model supports national coordination while allowing dioceses and parishes to address local needs. At the same time, the absence of formal data governance explains why Cáritas has well-defined organizational bodies without a unified approach to data definitions or data management.

The IV Strategic Framework

The IV Strategic Framework for 2024–2030 serves as the current strategic direction of Cáritas Portuguesa. The institution uses this framework to achieve its mission through a long-term strategy which adapts to Portugal's social and economic changes and international developments. The framework establishes a systematic structure for planning activities, resource distribution, and performance assessment. Cáritas bases its first strategic objective on developing a unified method to achieve Integral Human Development through multiple approaches. The entity works to improve its ability to detect and handle various poverty forms while building ecological sustainability and developing emergency readiness systems. The institution dedicates itself to addressing both urgent requirements and the fundamental elements which create vulnerability.

Cáritas dedicates its second strategic objective to building up its network structure. The entity works to strengthen diocesan partnerships while developing shared spiritual values, deepening organizational unity, and promoting synodality which enables Church members to listen and participate in collective decision-making. The network development process receives support from training programs and successful practice-sharing initiatives.

The institution focuses on sustainability and transformation as its third strategic objective. It works to develop effective leadership through younger generations while seeking diverse funding sources and implementing governance standards in its operations. Digital transformation has become its fundamental element for organizational growth. The framework recognizes that social organizations need to master technological change to stay effective and transparent while maintaining coherence in today's fast-paced environment. The strategic framework functions as a directional tool which leads Cáritas Portuguesa operations while preserving its mission throughout social changes.

Main Areas of Intervention

Cáritas Portuguesa delivers its services through multiple intervention areas which address the various social requirements of Portugal and delivers preventive and remedial and transformative services to its clients. Social intervention stands as the main focus of its operational activities. Cáritas organizations in diocesan and parish areas deliver assistance to families dealing with economic problems and food shortages and housing problems and multiple sources of stress. The organization provides individualized social support through trained professionals who create personalized plans to help clients access public and community resources. Cáritas Portuguesa supports elderly people who experience social isolation and financial challenges and dependency while running programs to enhance social connections in both urban and rural areas.

Furthermore, it provides humanitarian assistance through its domestic social intervention programs. The entity activates its disaster response system to help families in need while joining international emergency relief efforts through Caritas Internationalis. This demonstrates its dedication to help people in need through its international humanitarian work.

Cáritas dedicates resources to educational programs and institutional growth initiatives. The organization runs training sessions for staff members and volunteers while sharing methodological resources and establishing thematic practice communities for employability and senior support and migration programs. The training programs at the organization teach staff members both practical skills and ethical and spiritual service principles to maintain network unity.

The organization maintains international cooperation as one of its essential activities and works with multiple countries to execute development projects which focus on food security and psychosocial care and community development. The institution supports these international partnerships through national and international funding which helps build stronger civil society organizations worldwide.

The Observatório Cáritas conducts social research to generate reports which track poverty and exclusion patterns and provide data for internal planning and public discussions. The research activities of the organization help it better understand new social problems while enabling its participation in European and national policy discussions.

As Cáritas Portuguesa continues to expand its interventions and manage its social programs, the organization faces several challenges in handling vast amounts of sensitive data while maintaining accountability and transparency. Ensuring trust among beneficiaries, donors, and partner institutions requires ethical management of information as well as the adoption of technologies that respect privacy of individuals and support informed decision-making. In this

context, integrating human-centered AI systems can be critical in enhancing operational efficiency, providing insights, and supporting staff and volunteers through intelligent tools. Addressing issues such as AI literacy, trust, transparency, and data privacy is therefore essential for Cáritas to

8. GDPR in a Social Context

Data has become a central part of today’s world, leaving behind digital traces with nearly every action, interaction, or transaction. These data trails are stored, analyzed, and often used for commercial purposes, from profiling behavior to predicting actions, rather than improving society in more meaningful ways (Saqr, 2024). In order to address the growing privacy and ethical concerns, the European Union introduced the General Data Protection Regulation (GDPR), adopted in 2016 and put in action in May 2018. The regulation aims to protect individual rights, uphold ethical data use, and harmonize privacy standards across public and private sectors, including finance, healthcare, research, and digital services (Saqr, 2024). GDPR’s influence has extended globally, inspiring countries such as Japan, South Korea, Kenya, Chile, and Argentina to adopt similar privacy frameworks.

Our client, Cáritas routinely processes highly sensitive personal data related to health, housing, migration status, financial hardship, and family conditions in order to design and deliver targeted social support. As the organization increasingly explores digital tools to improve its operations, the risks associated with data misuse, profiling, and automated decision-making become important. In this context, GDPR functions as both a protective shield for beneficiaries and a guiding framework for Cáritas to ensure that innovation remains aligned with the value needed to protect individuals.

The GDPR is built on strong principles such as lawfulness, transparency, purpose limitation, data minimisation, fairness, and accountability (Voigt & Von dem Bussche, 2017). It also grants people concrete rights: the right to access their data, correct it, erase it, restrict processing, and object to automated decision-making (Kuner, 2020)**.** This legal framework actually came into life before the modern AI innovation, but it plays a central role in regulating how AI systems use and process personal data. Many discussions on AI governance in Europe start from the assumption that GDPR is still the backbone of data protection in an AI-driven environment (Sartor, 2020).

Today, several apps rely heavily on personal data. On one hand, personal information is often used to train machine-learning models. On the other hand, those same models are then applied to new personal data to make predictions or decisions about individuals. AI uses almost any kind of personal information, whether it is behavioural, biometric, social, or even just inferred to predict patterns or produce insights. This turns personal data, and the insights derived from it, into valuable assets (Sartor, 2020).

AI also automated decision-making much easier in areas that traditionally required human judgment, combining many variables and criteria. In several domains, automated decisions can be more consistent, cheaper, or even more accurate than human decisions because machines do not suffer from cognitive biases. However, it would be incorrect to say that these systems do not make mistakes. AI systems can sometimes rely on existing social inequalities for judgement, or, amplify biased training data, or create new forms of discrimination. This leads us to concerns about surveillance, manipulation, and a broader shift toward what is now called “surveillance capitalism” or even a “surveillance state” (Zuboff, 2019).

Zuboff (2019) defines surveillance capitalism as an economic system that turns human experience into data that can be monitored, analysed, and even sold to third parties for profit. Instead of relying on physical labour to generate money, this new model, adopted by big tech companies, extracts value directly from people’s personal lives, namely tracking every single action that they do in order to monetize it. The model keeps recording what people click on, where they go, who they talk to, and even what they might think or feel. Companies collect behavioural traces, often without clear consent, and feed them into predictive models designed to forecast and influence future behaviour. In this sense, surveillance capitalism is not just a business strategy but rather a form of power that grows by knowing individuals better than they know themselves and producing accurate information about them. Zuboff (2019) also argues that this creates a massive imbalance: companies know more about us than we know about them, and this imbalance allows them to shape choices and dictate actions.

This perspective points directly into GDPR values, because this regulation was created to address exactly these risks. The GDPR's core principles, which are transparency, purpose limitation, minimisation, and accountability, push back against the idea that companies can collect unlimited amounts of data simply because it is profitable or technologically feasible. Under GDPR, organisations must justify every purpose for which they process personal data, show that individuals understand what is happening, and limit data collection to what is strictly necessary. These rules that companies must follow under GDPR are not consistent with the approach that they are now taking, which is to collect as much information as possible for potential later use, or even for monetization purposes (Sartor, 2020; Zuboff, 2019).

In addition to that, GDPR gives control to individuals on how their data is used. It gives them the exclusive right to access, delete, or challenge the processing of their data, and it places strict limits on automated decision-making. In truth, the GDPR regulations cannot overthrow surveillance capitalism but they act as a counterforce by setting legal and ethical boundaries in a world that naturally pushes toward ever more digital control over our lives.

9. AI Under the GDPR

Although the GDPR does not explicitly treat AI or mention it as a potential threat to how personal data is used,, many of its rules implicitly apply to AI systems. There is a natural tension between GDPR principles, like purpose limitation, data minimisation, and restrictions on automated decisions, and the realities of AI that we see on a daily basis, which thrive on large, flexible datasets used to train models and improve performance (Sartor, 2020).

However, we can manage these tensions intelligently. For example, purpose limitation, which states that the data can only be used for the original purpose for which it was collected, can be applied flexibly: data can sometimes be reused if the new purpose is compatible with the original one. Also, data minimisation does not always mean using fewer data points. It can also mean reducing identifiability through data anonymization techniques. However, information obligations still apply because individuals have to understand what the system does and why their data is processed despite the complexity of these AI systems.

Furthermore, profiling and automated decision-making are allowed when certain safeguards are present, even though there is still debate about how much explanation organisations must provide to the user. However, privacy is by default a must-have feature in AI systems, otherwise, personal data is at risk and companies might be in violation of GDPR.

10. The Problem of Data Inference

Data inference can be defined as a data security threat in which valuable information about individuals can be inferred from data acquired by the system, even if this data was not directly collected (Wang et al., 2025). For example, retail stores can infer information about customers' age, income level, or interests based on their shopping history.

Over time, this kind of threat has grown to a high degree, particularly in the age of artificial intelligence. One example is the Cambridge Analytica scandal in 2014, when several information such as personality traits and political preferences of millions of Facebook users were drawn from psychological tests, allowing for targeted advertising during the US presidential elections in 2016. More recently in 2022, DiDi Global, a giant ride-hailing company in China, was fined over 1.2 billion USD by Chinese authorities for the illegal collection and inference of user travel patterns, highlighting that such threats are not just theoretical but have concrete consequences today (Wang et al., 2025).

Data inference can be grouped into three types. State inference, which uses real-time and dynamic data to predict the current states of targets, such as vehicle movements for location, or social behaviors for recommendations. Parameter inference relies on long-term data to deduce deeper information like personality traits Finally, joint inference combines both approaches, enabling simultaneous insights into dynamic behaviors and static attributes. Wang argues that although there is extensive research that focuses on specific theoretical problems around data inference, this narrow focus tends to underestimate the broader risks and complexity of inference threats in real-world situations, especially today in the artificial intelligence era. Moreover, since traditional data security frameworks often evaluate only the instant value of the data itself, like guessing a password or decrypting a bank account credentials, the risks of inference attacks, which is data that can be collected beyond what we have in our hands, are often overlooked (Wang et al., 2025).

There are three potential factors limiting the effectiveness of traditional inference methods, according to Wang (2025). Firstly, data heterogeneity presents a challenge because information collected from different sources can be in different formats and quality, making it difficult to combine and to detect hidden relationships. Secondly, the ambiguity of data relationships, and the fact that two pieces of data may seem distinct, makes it hard to draw links between the data and consequently make inferences. Finally, the complexity of inference increases proportionally as the number of data categories increases, making it challenging to explore all possible links between the data points. These technical obstacles have always presented real challenges, shaping data-protection frameworks that assume conventional inference methods cannot fully overcome these obstacles. However, AI systems are actually starting to challenge these obstacles, since they are now able to infer better and more accurate insights with improved efficiency.

11. The Modern Challenges of GDPR and Artificial Intelligence

There are several challenges that have been discussed in research papers about the implementation of GDPR principles, and how they can possibly collide with AI systems.

Data minimization is a principle under GDPR that requires organizations to collect, process, and store only the personal data that is strictly necessary for a specific purpose (Yang et al., 2025). The aim is to reduce privacy risks such as unlawful use of data, unauthorized access, and data breaches by limiting exposure in the first place (Ganesh et al., 2025). Although this principle is widely followed by global regulations along with GDPR such as CPRA, and LGPD, its practical implementation remains challenging because the principle is still lacking a clear picture about where data minimization ends and how we can establish a framework for AI systems to take into consideration this principle (Ganesh et al., 2025).

Garnish (2025) argues that current debate about data minimization ignores two key points: first of all, data minimization is a principle applied at an individual level. This means that a data point that is deemed irrelevant for one person could be crucial for another, and secondly, minimization is closely linked to privacy risks, especially risks associated with re-identification attacks. This means that the more data we have about an individual, the easier it becomes to identify them (Ganesh et al., 2025).

This research about data minimization also reviews regulatory foundations and identifies two pillars: purpose limitation (data must be collected only for a clear, specific purpose and used only for that purpose) and data relevance (we can only keep the data that is truly necessary for the purpose). In machine learning, this translates to selecting only the data that helps the performance of the model while ignoring the rest (Ganesh, 2025).

The authors highlight three major implications. First, data minimization is by default individualized because different individuals require different types or amounts of information. For example, loan applications in banks usually rely on income, job history, job security, and possibly other information which may or may not be relevant based on each person’s application. Secondly, minimization depends on the data, especially in situations where companies select relevant attributes from already collected large datasets. Thirdly, while regulations such as GDPR implicitly assume that minimization improves privacy, this assumption is not always realistic, simply because datasets often contain correlations between attributes. Therefore, removing a feature does not always prevent data inference through reconstruction techniques (Ganesh et al., 2025).

Diana Camões (2024), in her paper entitled “The Challenges of the GDPR in the Era of Artificial Intelligence: What Can We Expect From the Future?”, believes that the GDPR remains the fundamental legal framework for protecting people’s data even after the introduction of the European Union’s AI Act in 2024. Camões (2024) argues that although the AI Act adds a new modern layer to the regulatory landscape by specifically addressing risks posed by AI systems, it does not replace the original framework or the baseline, that is GDPR. Instead, the two legal bases operate together, sometimes harmoniously and sometimes in tension, depending on the type of the

AI system and the data in question.

Camões (2024) also believes that one of the key challenges today is determining when each law applies. If an AI system processes personal data, for example in credit scoring, hiring, or algorithmic profiling, then both the GDPR and the AI Act must be respected simultaneously. If the system uses no personal data at all, then only the AI act is in place. In the same manner, some research or academic literature may fall outside the AI Act but remains fully covered by the GDPR in the case of studies conducted on people. This shows that the modern regulatory ecosystem is complex and also proves the difficulty of ensuring consistent protection across different use cases.

In her paper, Camões (2024) focuses deeply on Article 22 of the GDPR, which regulates decisions made solely through automated processing that have legal or significant effects on individuals. According to the article, people have the right to not be subject to decisions made exclusively by automated processing if such processing will have a significant impact in their lives (Barbosa & Félix, 2022).

There has been extensive debate on whether Article 22 should be interpreted as a right by which individuals can object to automated processing, or a prohibition that restricts companies by default unless specific conditions are met. Camões (2024) suggests that Article 22 is fundamentally a general prohibition against fully automated decision-making. This interpretation significantly enhances user protection because it places the burden on companies that are required to justify the use of automated decisions, and ensures safeguards like human intervention whenever necessary.

In 2023, the Court of Justice of the European Union ruled that automated decisions with significant effects fall under Article 22 and are considered lawful only if one of the three exceptions applies: contractual necessity, explicit consent, or a specific legal authorization (Camões, 2024). The author considers this judgment a major milestone because it reinforces a strict interpretation of Article 22 at a time when AI systems are becoming more complex. The ruling emphasises that individuals must not be subject to invisible or algorithmic evaluations that are difficult to challenge, and that organisations must design AI processes with transparency and accountability in mind.

Despite its broad reach and ambitious goals, GDPR has proven challenging to interpret and implement. Organizations must comply with detailed requirements, including maintaining

records of processing activities, notifying breaches, appointing data protection officers, and adhering to ambiguous concepts such as “undue delay” or “disproportionate effort” (Saqr, 2024). These clauses often allow for varying interpretations, as seen when Facebook delayed breach notifications for nearly two months while claiming compliance (Saqr, 2024).

High-profile enforcement actions highlight GDPR’s authority, but fines often fail to deter large corporations. Amazon was fined 877 million Euros in 2021 for cookie policy violations, while Meta and Google faced hundreds of millions in penalties for breaches involving WhatsApp, YouTube, and other services (Saqr, 2024). Compounding the issue, enforcement is slowed by limited resources, legal complexities, and the increasing volume of complaints, leaving many unresolved since GDPR’s inception.

Saqr (2024) believes that these challenges do not mean that GDPR itself is failing, but they show how complex it is to apply privacy rules in data-heavy environments. For an organization like Cáritas, this highlights the need for tools that can handle data responsibly while helping them work more efficiently and ethically.

Enabling AI adoption in federated nonprofit organizations requires addressing multiple interdependent challenges simultaneously. Recognizing this complexity, our research team collaborated to develop a comprehensive AI enablement roadmap for Caritas Portugal, a federation of 20 legally autonomous dioceses.

Through engagement with Caritas leadership via exploratory meetings and collaborative assessment, we evaluated the organization's current state across these dimensions while developing integrated recommendations.

This section presents our team's validation of the data governance framework and its application to Caritas Portugal. While other team members addressed technical architecture, ethical considerations, and implementation planning, this component examines the governance foundations that enable AI adoption, assessing Caritas's current governance maturity and validating whether the adapted framework accurately diagnoses readiness in federated nonprofit contexts.

1. Preliminary Framework Application

1.1 Validation Approach

This section presents a preliminary application of the framework to Caritas Portugal rather than comprehensive empirical validation. Given time constraints inherent in the Field Lab format, the assessment demonstrates the framework's diagnostic utility while acknowledging methodological limitations that future research must address.

The Caritas assessment draws on two semi-structured interviews with national coordination leadership conducted in September and October 2025, supplemented by review of organizational documentation including the AidHound platform structure and internal reports. This data collection strategy enabled evaluation of federation-level governance structures, which represent the primary diagnostic criteria for maturity assessment. The presence or absence of Tier 1 coordination mechanisms, documented federation-wide standards, and systematic quality monitoring processes can be definitively established through central leadership interviews, as these structures would necessarily be visible to national coordination. However, this national-level assessment cannot speak to operational realities within individual dioceses, informal coordination mechanisms at regional levels, or actual data collection practices across the 20 autonomous entities. Comprehensive validation would require diocese-level interviews, observation of data management processes, and quantitative quality assessment across the federation.

Within these constraints, the assessment establishes that Caritas lacks documented federation-wide standards, operates without a formal governance council, and has no systematic quality monitoring processes. These absences constitute Level 1 maturity regardless of diocese-level variations. This application serves as proof-of-concept that the framework provides diagnostic utility while leaving implementation feasibility, timeline accuracy, and generalizability for future empirical testing.

1.2 Framework Positioning

The adapted framework builds on established data governance principles while adapting them for nonprofit contexts where hierarchical enforcement is unavailable. The three-tier architecture adapts corporate federated governance concepts (Alation, 2025; Khatri & Brown, 2010) for voluntary coordination rather than mandated compliance, a distinction validated through the Caritas assessment where dioceses operate as legally independent entities with no hierarchical relationship to national coordination. The maturity model addresses technical requirements for successful AI implementation identified by Miller et al. (2025) and Saeed et al. (2025), who emphasize that federated systems require standardized schemas while avoiding data heterogeneity that degrades model performance. The framework's progression from Level 1 heterogeneity to Level 3 standardization systematically addresses these technical prerequisites.

The framework differs from corporate models that assume hierarchical authority and financial incentives, relying instead on value-based participation incentives appropriate for mission-driven organizations. Unlike centralized nonprofit governance guidance (NetHope, 2025) that assumes unified organizational structures, this framework explicitly addresses legal independence of entities and designs coordination mechanisms appropriate for autonomous organizations. The Caritas case confirmed that generic guidance proved inadequate for federated structures, requiring the specificity that the three-tier architecture provides while maintaining necessary flexibility. The framework also addresses nonprofit-specific constraints identified by TechSoup and Tapp Network (2025): 76% of nonprofits lack AI strategy and only 6% have internal expertise. These constraints informed the framework's phased 24-month roadmap and emphasis on basic standardization before technical integration.

1.3 Application to Caritas Portugal

Project Context and Assessment Methodology:

This framework application was conducted as part of a field lab engagement developing a data and AI enablement strategy for Caritas Portugal, a federated nonprofit comprising 20 legally autonomous dioceses operating under shared mission. Assessment occurred through structured interviews with national Caritas leadership utilizing the evaluation methodology outlined in

Section 3.3's self-assessment framework, focusing on the four key dimensions: infrastructure, standards, quality management, and governance structures.

Current State Assessment:

Caritas Portugal presents a critical disconnect between technical infrastructure and governance maturity. The organization has implemented a third-party data management platform (AidHound) with standardized forms deployed across all 20 dioceses, suggesting technical infrastructure capabilities associated with Level 2-3 maturity. However, detailed operational assessment reveals fundamental governance deficiencies that conclusively place Caritas at Level 1 maturity.

Infrastructure Assessment:

While technical integration exists through AidHound, the data aggregation process exhibits fundamental quality issues. Data flows from individual dioceses to the central office where it undergoes a second aggregation process, creating compounding risks of quality degradation at each transformation stage. More critically, while data cleaning occurs within this pipeline, national coordination cannot identify who performs this function or what procedures govern the process. This represents complete absence of governance oversight for mission-critical data quality activities. When evaluated against the self-assessment framework's infrastructure dimension, Caritas demonstrates "some integration" (Level 2 indicator) but cannot reliably aggregate data due to unknown cleaning procedures and double-aggregation inconsistencies.

Standards Assessment:

Assessment revealed complete absence of standardized data definitions across the federation. No documented definitions exist for core concepts such as "beneficiary," "emergency assistance," or "service type." Different dioceses interpret these fundamental terms inconsistently, making cross-diocese data aggregation unreliable despite technical integration. The centralized dashboard represents the "common denominator" data elements all dioceses report, but no federation-wide standards mandate what data must be collected or how it should be defined. This finding directly validates the framework's emphasis on semantic consistency as prerequisite for federated AI applications. Caritas's lack of standardization creates exactly the "data heterogeneity" that "degrades model performance" (Saeed et al., 2025). Against the standards dimension, Caritas falls clearly into Level 1: no documentation of core definitions.

Quality Management Assessment:

No systematic data quality monitoring exists within the federation. Quality issues are discovered reactively when they cause operational problems rather than through proactive monitoring processes. No defined metrics assess accuracy, completeness, or consistency. The unknown data cleaning procedures compound this challenge, quality management cannot occur without visibility into quality processes. Additionally, no impact analysis utilizes the collected data, representing fundamental failure to leverage data for organizational learning or evidence-based decision-making. Against the quality dimension, Caritas clearly demonstrates Level 1 characteristics: reactive-only quality management with no proactive monitoring.

Governance Structure Assessment:

Assessment revealed absence of formal governance structures coordinating data management across dioceses. No federation council establishes standards (Tier 1 coordination absent), no regional working groups facilitate peer learning (Tier 2 coordination absent), and while individual dioceses maintain local control (Tier 3 autonomy exists by default), this operates without coordination mechanisms or minimum standards. Decision rights remain undefined: no clear authority determines what data should be collected, how terms should be defined, or what quality standards apply.

Maturity Level Determination:

Despite technical infrastructure suggesting Level 2-3 capabilities, operational assessment conclusively places Caritas Portugal at Level 1 maturity. The organization exhibits all defining characteristics: data managed without formal coordination processes, no standardization of definitions across autonomous entities, no quality monitoring beyond reactive error correction, and absence of governance structures or defined decision rights.

This assessment validates a critical framework insight: technical infrastructure does not equal governance maturity. Organizations can possess sophisticated platforms while remaining at Level 1 if they lack the governance structures, standardized processes, and quality management that enable effective data utilization. Caritas exemplifies this disconnect, AidHound provides technical capability, but absence of governance prevents its effective deployment for AI applications or even basic federation-wide analytics.

Discussion

2.1 Theoretical Contribution

This research makes three theoretical contributions to data governance and nonprofit management literature.

First, it extends federated data governance theory from corporate to nonprofit contexts, demonstrating how governance frameworks must adapt when hierarchical enforcement and financial incentives are unavailable. While Alation (2025) and others have established federated governance principles, they assume organizational contexts fundamentally different from nonprofits. This research shows that federated governance in nonprofits requires explicit design for voluntary coordination, consensus-based decision-making, and mission-driven rather than profit-driven incentives. The three-tier architecture operationalizes these adaptations, providing a model for future research on governance in decentralized mission-driven organizations.

Second, the research addresses a critical gap by connecting data governance frameworks to AI readiness requirements in federated contexts. Existing AI readiness literature (Hiniduma et al., 2024; Saeed et al., 2025; Miller et al., 2025) identifies technical data requirements but does not address how federated organizations can achieve these requirements across autonomous entities. This framework bridges these literatures, demonstrating how governance structures can systematically address technical requirements like semantic consistency, data quality thresholds, and sufficient data availability in federated contexts.

Third, the five-level maturity model provides a practical assessment tool for federated nonprofits, filling a methodological gap in the literature. While maturity models exist for centralized organizations (DAMA International, 2017), few address federated structures, and none specifically address AI readiness in nonprofit contexts. This model enables both assessment and planning, making abstract concepts like "data quality" and "governance maturity" concrete and measurable.

2.2 Practical Implications

For nonprofit practitioners, this adapted framework provides actionable guidance for AI preparation that has been largely absent from the field.

The framework resolves a critical challenge facing federated nonprofits: how to achieve sufficient data standardization for AI while respecting entity autonomy. The three-tier architecture shows that this is not an either or choice but a matter of clearly delineating decision rights. Central coordination can establish minimum standards enabling AI applications without controlling local operations.

Regional coordination can facilitate voluntary learning without imposing uniformity. Local entities can maintain operational control while benefiting from federation-wide AI capabilities.

The maturity model provides realistic expectations about AI readiness timelines. The TechSoup and Tapp Network (2025) finding that only 7.4% of nonprofits have successfully adopted AI reflects, in part, organizations attempting AI without adequate data foundations. By clearly defining Level 3 as the minimum threshold and providing a 24-month roadmap to reach it, the framework helps organizations avoid premature implementation that wastes limited resources.

The framework also addresses resource constraints realistically. Phase 1 requires no technical infrastructure, only governance structure establishment and definition development—activities feasible even for small nonprofits with limited budgets. Phase 2 introduces quality monitoring and templates, requiring some investment but not sophisticated technology. Only Phase 3 requires substantial technical investment, and by this point organizations have 12 months of governance maturity demonstrating commitment and feasibility.

2.3 Limitations

Several significant limitations must be acknowledged. Most fundamentally, the framework presents theoretically-grounded guidance with preliminary application rather than rigorous empirical validation. The Caritas assessment relied on two interviews with national leadership and organizational documentation review, providing visibility into federation-level governance structures but limited insight into diocese-level operations, informal coordination mechanisms, or practitioners' lived experience. The Level 1 finding reflects observable absence of formal mechanisms at the national level but cannot comprehensively characterize the federation's full governance landscape.

Comprehensive validation would require longitudinal case studies following multiple federations through implementation to test the 24-month timeline and maturity progression, cross-organizational comparison to establish generalizability, and diocese-level data collection to validate whether federation-level assessments reflect entity realities. Such work represents a multi-year research agenda beyond master's thesis scope but essential for establishing empirical soundness.

The framework assumes baseline organizational capacity that resource-constrained federations may lack and focuses narrowly on governance structures without deeply engaging change management challenges, organizational culture barriers, or political dynamics of voluntary coordination among autonomous entities. While this framework addresses governance prerequisites for AI implementation, it does not substantively engage AI ethics, algorithmic fairness, or responsible AI considerations for organizations serving vulnerable populations, dimensions addressed by other team members in the collaborative project. The framework's value lies in providing theoretically-sound structural guidance for an under-addressed problem while acknowledging substantial empirical work remains necessary to validate practical utility and refine implementation guidance.

2.4 Future Research Directions

This research opens several avenues for future investigation. Empirical validation studies implementing this framework across multiple federated nonprofits would provide crucial insights into actual implementation challenges, realistic timelines, and necessary adaptations. Comparative research examining how different organizational characteristics affect governance implementation success would enable more targeted adaptation guidance. Integration research connecting this data governance framework to AI ethics frameworks would provide comprehensive guidance for responsible AI implementation. Tool development research creating practical implementation resources would lower barriers to adoption. Cross-sector research examining framework applicability to other federated structures would test broader applicability.

Conclusion on Data Governance

Federated nonprofit organizations face critical AI implementation challenges due to data fragmentation and inconsistent governance. This research adapts established data governance frameworks (Alation, 2024; DAMA International, 2017; Khatri & Brown, 2010; Miller et al., 2025; Saeed et al., 2025) to design a comprehensive framework for voluntary, consensus-based coordination in federated contexts. The framework's three-tier architecture and five-level maturity model provide a practical implementation path.

Application to Caritas Portugal validated the framework's core insight: technical infrastructure does not equal governance maturity, and this disconnect is the primary barrier to AI readiness. The significance extends beyond operational considerations, federated nonprofits serve vulnerable populations and deploying AI on inadequate data foundations risks harm through biased predictions or inequitable resource allocation. By establishing governance prerequisites through framework adaptation, this research enables responsible AI deployment that upholds the dignity and rights of those served.

Overview of Cáritas Portuguesa’s Data and AI Challenges

Cáritas Portuguesa is at an intriguing turning point. On paper, it already has a contemporary data platform (AidHound) distributed across all 20 dioceses, indicating some degree of digital maturity. In practice, however, after some interviews we had with their data sector team they continue to operate with severely fragmented data practices: there are no standard definitions for essential concepts, no institutional governing body for data, and no regular data quality monitoring. As a result, while possessing infrastructure that seems to be in Levels 2-3, the company is classified as Level 1 - Ad Hoc according to our maturity model.

This mismatch between infrastructure and governance is not limited to Cáritas. Data governance research emphasizes that technology alone does not provide data value. Organizations must also establish who has decision-making authority over data, how standards are developed, and who is accountable for quality (Khatri & Brown, 2010). When these components are lacking, data is viewed as a byproduct of activities rather than a controlled asset. That is exactly the pattern found at Cáritas: data flows exist, but the organization has yet to formalize who owns definitions, who may update forms, and how quality concerns are handled.

1.1 AI, Data, and Nonprofits: The Broader Picture

Cáritas's challenges are part of a larger trend in the charity sector. According to TechSoup and Tapp Network's 2025 "State of AI in Nonprofits" benchmark assessment, over 80% of nonprofits are using or exploring AI tools, but only about a quarter have a formal AI strategy (DiGiovanni, 2025). Many are experimenting with AI but lack clear governance, ethical guidelines, and internal expertise, creating an “AI governance gap” between experimentation and responsible use. Initiatives like NOVA SBE’s Social Sector Database and Leadership for Impact Knowledge Center demonstrate how data is being positioned as a foundation for social impact (NOVA SBE Leadership for Impact, 2017; Nova School of Business and Economics, 2020). Cáritas operates in this evolving landscape but, from an internal governance standpoint, still functions at a foundational level. The following analysis details the specific data and AI challenges identified within this context.

Our team's diagnostic assessment (Section 4.3) confirmed that Cáritas, while possessing a modern data platform, lacks the governance structures needed to translate data into strategic value. This insight directly informed our objective: to develop a Proof-of-Concept application that addresses these gaps by enhancing data utility and embedding responsible AI principles into the user workflow.

1.2 Why These Challenges Matter for AI Implementation

AI and these governance gaps are fundamentally linked. Evidence from federated learning shows that data heterogeneity and inconsistent quality, such as that found at Cáritas, compromise model performance and risk biased outcomes. Simultaneously, the sector-wide "AI governance gap" increases the likelihood of deploying tools that are misaligned with organizational values or that expose vulnerable beneficiaries to unfair decisions (DiGiovanni, 2025). For an organization like Cáritas operating in welfare support, this is not only an operational risk but a legal one: the EU AI Act classifies such systems as 'high-risk,' necessitating robust governance to ensure compliance (Lomba & Roldão, 2025). Beyond compliance, deploying AI without guardrails risks undermining the values of justice, dignity, and fairness that international principles for trustworthy AI are designed to protect (OECD, 2019). This underscores a core lesson from data governance literature: without explicit decision rights, standards, and accountability frameworks, AI initiatives are prone to ethical and operational failure (Khatri & Brown, 2010; Bento, 2021; Van Grembergen et al., 2004). The remainder of this section outlines targeted strategies to address these gaps and build a responsible foundation for AI adoption.

2. Improving Communication and Governance

The lack of a defined governance structure that is applicable to all dioceses is a problem for Cáritas Portuguesa's present data practices. Although the organization has technological tools in place, there isn't a structured entity in charge of setting data standards, organizing data quality procedures, or making sure dioceses utilize and interpret data consistently. This kind of gap tends to cause long-term fragmentation in federated organizations since each independent unit develops its own procedures, terminology, and priorities. Over time, these variances complicate coordination and skew aggregated analysis. According to governance studies, coherence in decentralized systems requires regular communication channels, shared decision-making powers, and predictable structures (Khatri & Brown, 2010). These ideas clearly apply to Cáritas, where autonomy is crucial yet coordination is becoming more and more important for data-driven work and the use of AI in the future.

2.1 Federation Data Governance Council (Level 1)

Establishing a Federation Data Governance Council (FDGC) is a useful first step for Cáritas. It is not necessary for this entity to be big or bureaucratic. It only has to be acknowledged as the body in charge of making decisions that have an impact on every diocese.

Composition:

A rotating group of diocesan members and representatives from national coordination should serve on the council. This guarantees that all members see government as something that is co-created rather than something that is imposed from above, and that smaller dioceses are not left out.

Mandate and responsibilities

The FDGC would have the authority to make decisions regarding:

- establishing and upholding fundamental data definitions utilized throughout the federation.

- establishing minimal requirements for data quality in important domains.

- verifying any modifications made to data gathering templates.

- setting guidelines for sharing and accessing data.

- supervising the creation of Cáritas's acceptable-use policies and AI strategy.

Functioning as a coordinating body, the FDGC ensures definitional and quality consistency, preventing drift in line with federated governance principles (Van Grembergen et al., 2004).

2.2 Tier 3 Local Data Stewards

Accountability for data is frequently still implicit at the local level. Formalizing each diocese's local data steward position offers a clear point of reference without introducing undue bureaucracy.

Definition of a role

The data steward is not an expert in information technology. Coordination is more important in this role:

- ensuring that employees comprehend the significance of indicators and fields.

- examining data for glaring mistakes before submission.

- taking part in regional gatherings.

-recording local codes or exceptions.

This is consistent with the larger body of research on data governance, which shows that better data quality and more successful application of central standards are linked to operational role clarity (Khatri & Brown, 2010).

2.3 Communication Protocols and Shared Tools

Governance frameworks function only through consistent communication, which can be enabled by simple, practical tools. A key example is establishing a **shared knowledge database**: a central repository for housing essential materials like formal data definitions, process documentation, policy documents, and training resources. This creates a single, accessible source of truth that supports consistent practice across the organization. Cáritas can keep this central repository that includes, the data dictionary, documentation for dashboards, guidelines for quality checks and short explanatory videos, for example. This does not need complicated software. An organized Google Drive or SharePoint setup could be sufficient.

Any update to the dashboards or forms should be accompanied by a date, description, and action required. This will prevent misinterpretation and ensure that all dioceses have the same information. Good governance requires standard communication that can be operationalized through basic, practical tools-just like a **change-log** implementation.

Finally, a regular **quarterly data forum**. A 90-minute online meeting with representatives from all dioceses-allows the organization to discuss reoccurring difficulties, share improvements from various dioceses, and present new projects, including AI pilots.

3. Building Data Literacy and AI Awareness

Cáritas Portuguesa cannot go forward with more advanced data practices or appropriate AI usage until data literacy is improved across the federation. Our online meeting revealed that many employees are confused about the significance of important fields, lack confidence in reading dashboards, and frequently regard data entry as an administrative chore rather than a vital component of service delivery. This is not uncommon in the charity sector: research suggests that organizations frequently underinvest in staff competencies compared to their objectives for data-driven work (Van Grembergen et al., 2004). Even well-designed governance mechanisms fail to affect behavior unless they are accompanied by focused training and explicit explanations of how data supports the objective.

The idea isn't to turn social workers into analysts. Instead, Cáritas requires a foundational level of common understanding so that employees can enter data accurately, analyze information consistently, and grasp the ethical implications of AI tools. Building this foundation is critical for attaining Level 3 of the maturity model, when standardization and governance become operational rather than theoretical.

3.1 A Three-Level Data Literacy Programme

A planned yet lightweight data literacy training can bring clarity and alleviate anxiety about data activities. The emphasis should be on practical learning that is directly relevant to Cáritas' environment.

Level 1: Foundations (for all staff)

Staff must grasp what data is and why it matters. Training at this level should include:

- the goal of data collecting in social assistance services

- fundamental notions including fields, records and time periods.

- common errors and how to avoid them.

- the relationship between reliable data and improved services.

Short sessions with basic examples work well. Many NGOs have seen considerable advantages by implementing quick, repeated learning rather than lengthy technical seminars.

Level 2: Operational Data Use (for Data Stewards and Coordinators).

For individuals who often work with data, the next step is to establish operational competence.

- reading and Interpreting AidHound Dashboards.

- understanding how indicators are calculated.

- identifying outliers or contradictory values.

- documenting local customs and deviations.

This level assures that each diocese has at least one person capable of connecting everyday operations with governance requirements.

Level 3: Data for Decision-Making (Management and National Coordination)

The leadership level is focused on using data for planning and oversight.

- interpreting trends with confidence.

- understanding the limits of available data.

- using fundamental analytical reasoning.

- raising knowledge of AI principles like prediction, categorization, and model limits.

Senior leadership is important to the long-term success of data culture. When leaders utilize data to make decisions and speak openly about constraints, staff attitudes frequently change more rapidly (Bryson et al., 2015).

4. Strengthening Data Foundations: Definitions, Documentation and Quality

4.1 Clarifying Definitions and Indicators

Cáritas faces challenges in advancing beyond basic reporting due to a lack of clarity around key concepts. Interviews revealed that different dioceses have varying interpretations of terms like "beneficiary," "service," and "emergency assistance." In a federated structure, these differences lead to inconsistencies that hinder effective analysis across the organization. Research on data governance indicates that such variability can adversely affect data quality and compromise the reliability of future AI applications (Khatri & Brown, 2010). A great way to tackle this would be by developing a focused federation data dictionary (FDD) for core operational concepts, defining each with its scope, valid values, and an owner. Also, we could complement this with a dashboard indicator glossary that explains calculations and limitations on site, reducing confusion and building trust.

4.2 Documentation Standards for New Fields

To foster a consistent and bright future, we encourage Cáritas to implement a straightforward standard for documenting any new data fields added to AidHound or local forms. Each field should define its purpose, include a clear definition, and outline the possible values it can hold. By ensuring that only fields that meet this documentation standard receive approval from the governance council, we can embrace clarity and reduce local variations. This proactive approach will empower the federation, paving the way for a more unified and effective system.

4.2 Local and Federation-level Quality Assurance

At the local level, dioceses should implement simple review routines, such as a second-person check before submission, to catch obvious errors and maintain data integrity. Complementing this, a quarterly federation-wide quality dashboard would enable each diocese to monitor performance indicators (completeness, error rates, timeliness), fostering transparency, peer awareness, and targeted improvement across the organization (Bryson et al., 2015).

5. Impact Assessment, AI Risk and Early Use-Cases

5.1 Strengthening Impact Measurement and AI Risk Awareness

Cáritas currently collects extensive operational data, but the lack of universal outcome metrics prevents accurate federation-wide comparisons. Implementing a simple, realistic set of indicators, such as improvements in basic needs, access to long-term support, or duration of aid, would enable the organization to consistently measure success and understand beneficiary progress over time.

When discussing impact today, it is impossible to ignore AI. Any organization handling sensitive social data has to consider the risks that automated tools may create. Even basic systems might accidentally promote prejudices or affect judgments that should be made entirely by humans. To avoid this, Cáritas requires an established AI Impact and Risk Assessment methodology that applies to any new AI tools, including informal experiments. This doesn't have to be complex. A simple survey asking what data the system uses, if it affects vulnerable populations, how human control is assured, and what its potential risks are would already put Cáritas ahead of many other NGOs. This form of review is compatible with international guidelines emphasizing openness, justice, and accountability in AI systems (OECD, 2019).

5.2 Early, Safe AI Use-Cases and a Responsible Path Forward

With these precautions in place, Cáritas may begin testing early, low-risk AI use cases. The simplest and safest apps are those that support rather than replace employees. For example, AI-generated summaries of case notes can decrease administrative labor while leaving full professional judgement to humans. Another useful application is anticipating service demand using aggregated, anonymized data. This helps the federation prepare resources without affecting decisions for individual beneficiaries. AI can also assist with quality checks, flagging unusual patterns or clear inconsistencies before reports are completed. These early use-cases allow Cáritas to learn how AI behaves in practice, using controlled scenarios with minimal risk.

By linking impact assessment with AI governance, Cáritas builds a responsible path forward. It avoids the common mistake of rushing into AI without understanding what technology is meant to achieve or what risks it may introduce.

6. Technical Readiness, Partnerships and a Practical Roadmap

6.1 Stabilizing Infrastructure and Leveraging Partnerships

Technical improvements are necessary, but they should be realistic and aligned with Cáritas's actual capacity. The most urgent demand is to consolidate data flows so that information flows smoothly from dioceses to the federation, avoiding needless transformations. Today, some data is processed twice before reaching the national level, which increases the chance of errors and makes the process difficult to understand. Mapping the entire data pipeline and defining one “source of truth” for federation-wide reporting would resolve many of these issues. Ensuring that each record contains a consistent beneficiary identifier also supports more reliable analysis over time.

Partnerships are another essential part of technical readiness. Cáritas does not need to develop everything internally. Universities, particularly NOVA SBE, can support training, impact evaluation and early AI development in a controlled environment. Technology partners can help improve dashboards, validation rules and secure data storage without taking over the organization’s autonomy. What matters is establishing clear agreements about data protection and decision-making so that Cáritas remains responsible for the direction of its digital transformation.

6.2 A Sustainable Growth Blueprint

With these foundations in place, Cáritas can implement a straightforward, phased roadmap that is appropriate for its size and resources. In the short term, the top objective is to create governance structures, publish the data dictionary, establish minimal documentation requirements, and implement basic quality checks. The following step focuses on integration: stabilizing the pipeline, enhancing dashboards, and testing one or two AI tools with close human supervision. The last step entails growing successful processes, recording lessons gained, and ensuring that governance procedures become part of daily operations rather than one-time projects. This grounded approach avoids overpromising, reduces risk, and enables steady progress toward a more coherent, evidence-based, and AI-ready environment (Bryson et al., 2019).

Ultimately, meaningful transformation for Cáritas begins not with advanced technology, but with stronger coordination, clearer standards, and greater confidence around data. By building this foundation of governance and trust, the organization can responsibly adopt AI, starting with low-risk use cases, ensuring that innovation aligns with its mission of serving vulnerable communities with dignity and fairness.

Cáritas on the news

It was reported by Jornal de Notícias, the Portuguese newspaper, in November 2025 that Cáritas is about to kick off a house-building campaign for victims of violence. The project received around €700,000 in funding, as is expected to be launched during the Christmas season. During this campaign, Cáritas will be selling candles, with the hopes of allocating 65% of total funds to finance the construction of these shelters (Costa, 2025).

The journal also reported at an earlier date that Cáritas assisted more than 1,500 people last year, of which 874 were new requests for assistance. Cáritas was able to leverage its facilities as well as supporting entities host migrants in search of shelter, children, young people and families from ethnic minorities. The same institution reported a delivery of 363 food baskets, and provided financial support amounting to 32,000 euros to support housing, food, health, education, mobility, and other needs (Costa, 2025).

By looking at the above information provided by the journal, we notice that Cáritas is keeping track of its fund-raisings, expenses and future plans with clear figures, as reported by the journal.

Cáritas, just like social organizations that operate nationally or globally, pays close attention to the way it handles and tracks its projects, which is important for keeping things transparent and accountable. Setting clear targets like this also helps build trust with donors, who can see exactly where their money is going.

However, it’s not just about big projects. Cáritas also keeps tabs on its day-to-day assistance programs. Last year, it helped over 1,500 people, including migrants, children, and minority families, and it tracked the number of food baskets distributed and the €32,000 in financial aid given for things like housing, health, and education (Costa, 2025). Keeping track of all these numbers allows Cáritas to figure out early needs, and also which social groups are in critical conditions and might need urgent support.

According to Expresso (2025), Cáritas Portugal recently stressed the urgent need for renewed efforts to combat poverty and social exclusion, as recent progress has been limited and vulnerable to reversal due to global geopolitical risks. The organization also ran an assessment of national conditions using indicators from the National Institute of Statistics (INE) and Eurostat. While the population at risk of poverty or social exclusion dropped by a low percentage, from 21.1% in 2019 to 19.7% in 2024, Cáritas warns that this improvement is fragile and largely driven by favorable economic conditions, such as small employment growth, rising wages, pensions, and social transfers.

The evaluation that Cáritas performed also highlights some weaknesses in Portugal's social support system, namely low contributions from social transfers, particularly for children, and calls for increased investments and more effective public policies dedicated to the social context. Cáritas also points to a rising number of requests from immigrants, whose vulnerabilities are sometimes underestimated by official statistics or completely ignored, which is preventing these marginalized groups from housing services, labor market integration, and access to social support.

On top of that, housing deprivation remains a major issue, which is part of the general housing crisis in the country. The number of homeless people has risen sharply, from 6,000 in 2018 to 13,100 in 2023, and access to adequate housing has become increasingly difficult due to rising rents and overcrowding. Children and low-income families are also affected, with long-term consequences for education, health, and future opportunities. Cáritas emphasizes that in order to overcome these challenges, the efforts must be multidimensional, structural, and targeted to hopefully prioritize the most excluded groups (Expresso, 2025).

Taking into consideration these on-going projects at Cáritas, an AI-supported application could significantly strengthen the organization's ability to plan, prioritize, and monitor its social intervention projects, especially given the importance and urgency of the challenges it is currently addressing in relation to housing, poverty, and fund-raising. One of the most immediate benefits would be the platform's ability to analyze the organization's existing data, such as food basket deliveries, financial support amounts, housing requests, and migrant assistance cases, and generate structured recommendations based on the data. Instead of manually reviewing spreadsheets, staff could receive automated insights about emerging patterns, such as neighborhoods where requests are increasing, demographic groups showing higher vulnerability, or services that may soon exceed capacity. This would allow decision-makers to shift resources more efficiently, anticipate demand, and justify funding requests using clear evidence.

The app could also help Cáritas improve the way it manages its financial support systems. Since the organization allocates funds for housing, healthcare, food, mobility, and education, an our dashboards could break down spending by category, identify repeated assistance patterns, and flag cases where deeper intervention might be needed. This level of visibility helps prevent resource misallocation and ensures that support aligns with the organization's mission and strategic priorities. In particular, during large campaigns, such as the recent housing initiative, the system could track how funds are being used in real time, monitor progress toward construction milestones,

11.1 Examples of KPIs

Sheet Type	Example KPIs	Description
Pedido Apoio Pontual	Number of requests; Total amount requested	Measures demand for emergency financial assistance
Pedido Prioridade às Crianças	Age distribution; Types of expenses	Captures patterns in child-focused support
Pedido Vales	Count of 5€ and 10€ vouchers	Tracks material assistance volumes
Pedido Cartões	Total top-up amount; average per card	Shows cash-equivalent support workload
Vales	Voucher inventory	Measures stock and flow of vouchers
Atribuição Vales	Allocation by Cáritas branch; demographic distribution	Monitors who receives vouchers and why

Group Part (continued)

12. Advanced KPI Framework

The KPI framework implemented in the *Cáritas Data Explorer* functions as the analytical backbone of the application. Its purpose is to transform raw data records into actionable indicators that can be easily interpreted by operational staff, programme managers and strategic decision-makers. Each KPI is designed to reflect the specific characteristics and operational goals associated with the different datasets used by Cáritas. While the underlying computations may appear relatively straightforward from a programming perspective, they form a crucial interpretative layer that converts numerical data into meaningful organisational insights.

12.1 Principles Behind KPI Design:

The KPI structure used in the dashboard follows four guiding principles that ensure their usefulness across the organization:

Relevance

KPI selection was determined by how relevant an indicator would be for the organization's operational context. Metrics such as the total value of "Verba Solicitada" (requested financial support) or the motive of the request provide Cáritas with a greater practical value than more abstract statistical measures. This ensures that every indicator directly contributes to programme monitoring or resource allocation discussions.

Interpretability

KPIs are expressed in clear and unambiguous units (e.g., number of requests, euros allocated). This allows stakeholders without statistical or technical expertise to easily interpret results with accuracy.

Comparability

Where possible, KPIs are structured so they remain consistent across different sheet types. For instance, demographic distributions appear in multiple dashboards, enabling cross-programme comparisons and facilitating a unified understanding of beneficiary profiles. Such alignment improves the overall coherence of the analytical environment and supports multi-programme evaluation.

Efficiency

The KPI engine is designed to operate efficiently by performing calculations only on the filtered subset of data selected by the user. This approach ensures that the dashboard remains responsive, even when larger files are uploaded.

Together, these principles shape both the numerical summaries and the visual components of the dashboard. As a result, the KPI system extends beyond visual reporting and becomes a functional decision-support tool that strengthens Cáritas' capacity to analyse, interpret, and act upon its data.

13. KPI Logic by Sheet Type

The Cáritas Data Explorer processes six dataset categories, each reflecting a distinct operational workflow. Accordingly, the KPI logic was designed to capture the essential characteristics and decision-relevant patterns of each dataset. The following sections summarise the key indicators implemented in each of the datasets.

Pedido Apoio Pontual (Emergency Financial Support Requests)

This sheet represents urgent assistance cases, typically involving financial instability or inability to meet essential expenses. KPIs focus on quantifying demand and identifying core drivers of vulnerability.

a) Volume of Requests: Count of records, used to assess workload and detect seasonal fluctuations.

b) Total Value Requested: Sum of Verba Solicitada, indicating financial pressure on emergency support budgets.

c) Gender Distribution: Proportion of male/female beneficiaries, highlighting demographic patterns such as prevalence of single-parent households.

d) Motive Distribution: Categorisation of reasons for support (e.g., Baixos rendimentos, Saúde, Desemprego), revealing dominant socioeconomic stressors.

Pedido Prioridade às Crianças (Child-Focused Support Programme)

This sheet concerns a national programme focusing on minors. KPIs quantify programme usage and the types of needs presented.

a) Total Number of Cases: Measures the programme's reach.

b) Total Value of Requests: Indicates the overall intensity of financial demand.

c) Requested Value from the Programme: Isolates the amount expected to be funded through this specific programme.

d) Disability Status Prevalence: Identifies cases involving children with additional vulnerabilities (e.g., disabilities).

e) Types of Expenses Requested: Summary of categories such as healthcare, school materials, or emergency items, usually visualised in bar charts.

f) Age Distribution: Histogram-based view of age groups most represented among requests.

Pedido Vales (Voucher Request Form)

This sheet records requests for vouchers (typically 5€ or 10€).

a) Total number of 5€ and 10€ vouchers requested

b) Total monetary value represented by these vouchers

c) Proportional distribution between voucher types

Together, these indicators reflect material needs, food insecurity patterns and household budget constraints.

Pedido Cartões (Rechargeable Financial Cards)

This sheet documents the use of prepaid or rechargeable support cards.

a) Number of cards loaded

b) Total amount loaded onto these cards

c) Average load per card

d) Distribution of motives for support

These KPIs help identify recurring needs, financial strain among households, and operational trends in card-based assistance.

Vales (Voucher Inventory Dataset)

This sheet functions as the voucher stock inventory.

a) Number of vouchers

b) Total value of stored vouchers

c) Average voucher value

d) Distribution of voucher values across denominations

These metrics support supply management and internal financial accountability.

Atribuição Vales (Voucher Allocation Records)

This sheet captures the actual distribution of vouchers to beneficiaries.

a) Number of allocations

b) Total value allocated

c) Value distributions by Cáritas diocesan units

d) Demographic patterns (age group, gender)

e) Beneficiary categories (“Novo”, “Recorrente”, etc.)

f) Motive for allocation

These KPIs provide insights into how material support is delivered and which populations benefit most from it.

14. Visualisation Strategy

The visualisation strategy in the application focuses on presenting KPIs in a clear and straightforward manner, ensuring that users can easily interpret the information derived from each dataset. The chart types used are simple and usual, such as bar charts, pie charts, and histograms, chosen according to the characteristics of data. Categorical variables, such as motives for support or voucher types, are displayed using bar charts to highlight differences in frequency, while demographic proportions are typically shown through pie charts to provide an immediate sense of distribution. Histograms, on the other hand, are used to present age-related information or voucher value distributions. By matching each KPI with the most suitable visual form, the dashboard ensures that patterns and trends are quickly recognisable and that users can extract the intended insights without difficulty.

15. User Interface and Interaction Design

The dashboard employs a layered user interface designed to maximise transparency, simplicity, and accessibility for users. When a file is uploaded, each detected sheet type is presented as a separate tab, allowing users to navigate between datasets without losing context. Within each tab, the interface opens with a preview of the first 200 rows, masked when PII protection is enabled, helping users verify if the correct file has been uploaded and acknowledge the available fields before examining the KPIs. Filters, implemented through Streamlit's st.expander, facilitate data exploration, allowing users to work the dataset by selecting date ranges, demographic categories such as gender or nationality, or specific parishes and diocesan units. The computed KPIs are displayed using Streamlit's metric component, providing a simple way to interpret the most important indicators. Finally, an insight generation panel is placed at the bottom of each tab, ensuring that users first engage with the information displayed before reading AI-generated interpretations, reducing the risk of cognitive bias, in which users solely rely on narrative explanations without contextual understanding.

16. Synthetic Data Rubric (Main Section)

Because Cáritas' operational data could not be used for development due to privacy restrictions, the project relied entirely on synthetic datasets that imitate the structure and statistical patterns of the real data without containing identifiable information. The purpose of synthetic data was to enable realistic testing, replicate standard operational scenarios, and validate both KPI logic and filtering behaviour while maintaining full anonymity. The data therefore needed to balance realism with anonymity, functioning as a structurally accurate but non-sensitive substitute for real records.

To achieve this, each sheet type was generated according to its required structure, including the consistent positioning of mandatory fields, realistic value ranges such as ages between 0 and 100, authentic categorical labels like “Feminino” and “Masculino” and typical values for monetary fields. Categorical variables, such as age groups, motives, gender and nationality, were produced through weighted sampling to reflect realistic distributions. Numerical variables, including requested amounts, voucher quantities, and card top-up values, were generated using truncated normal, Poisson, or uniform distributions depending on the variable type. Textual descriptions were constructed using templates representing common situations such as employment instability, health-related expenses or rental pressure.

Although artificial, the datasets preserved internal logic: totals matched quantities and unit values, chronological sequences were maintained, and records mirrored typical operational workflows. This ensured that the synthetic data behaved like real data from a functional perspective, allowing the dashboard to be tested and validated effectively without compromising privacy.

17. Analytical Insight Module

The analytical insight module serves as the interpretive component of the Cáritas Data Explorer, transforming the descriptive outputs of KPIs and charts into higher-level observations. It analyses only the currently filtered dataset and builds a structured, context-aware prompt that ensures the resulting narrative remains directly relevant to Cáritas’ operational needs. To ensure this, the module generates a compact data profile containing row counts, column names, top categories for qualitative variables, numerical aggregates and frequency distributions. This profile allows the system to identify patterns such as demographic concentrations, unusual financial amounts, or distinctive motive distributions without exposing raw data.

The profile is constructed by transforming the filtered DataFrame into a compact JSON object that includes the sheet name, number of rows, column structure, and summaries of categorical and numerical variables. This approach protects confidentiality while ensuring that insights always correspond to the user's active filters. The prompt incorporates this JSON profile and follows a defined structure that encourages the analysis to produce multi-paragraph insights, contextual interpretation and operational recommendations, with explicit constraints, such as specific number of insights and recommendations or academic quality writing, to guarantee depth and consistency across sheet types.

In the Streamlit interface, the module is positioned at the end of each sheet tab within an expandable panel, reinforcing a workflow in which users first explore the data preview, then the KPIs and charts, and only afterwards consult the interpretive insights. This ordering mitigates the risk of over-reliance on the AI output and fosters comparison between narrative interpretation and the previously observed evidence. Overall, the insight module enhances organizational decision-making by highlighting trends, detecting patterns in demographic and financial data, and identifying potential operational issues.

18. System Limitations

Despite its advantages, the Cáritas Data Explorer has several structural and operational limitations that must be acknowledged. Starting with data-related constraints, the system relies on consistent uploaded file structures, therefore, if diocesan submissions significantly modify templates, automated detection could fail. Additionally, although the application can detect them, it cannot correct substantive data-entry errors, and the synthetic data used during development, while structurally accurate, cannot fully capture the totality of nuances that real operations may face.

Secondly, technical limitations arise from the application's need to run on local machines or secure servers, making performance dependent on user hardware and potentially slower with large datasets. Furthermore, Streamlit reruns the script at each interaction, which may feel unfamiliar to some users, and the lack of persistent storage prevents long-term data retention or multi-user dashboards without further development.

Moreover, the performance of the system may be harmed if there are variations in the data governance framework between dioceses, for instance, if dioceses collect or record data with different methods it weakens the comparability power of the app. Likewise, incomplete data submissions within dioceses limit national-level insights. Finally, the insight module cannot integrate external contextual factors, such as sudden regional crises, therefore, observed trends must not be mistaken for causal relationships, and PII masking, although essential, limits granular analyses involving individual-level timelines or repeated cases.

19. Overall Assessment

The Streamlit-based architecture successfully delivers a functional, user-friendly and analytically robust dashboard that addresses directly to Cáritas' core challenges, including fragmented data submissions, the need for rapid reporting, limited technical infrastructure, confidentiality requirements and the demand for strategic insights. While the system has recognized limitations, the overall approach remains technically sound, ethically appropriate, and feasible within the organization's operational context. Its modular design further ensures long-term adaptability, allowing new datasets, KPIs, or analytical components to be added with minimal disruption.

What Now?

From Insight to Action at Cáritas Portuguesa

We’ve diagnosed the problem, defined the frameworks, and built a prototype. The question shifts

from what could be done to what should happen next.

1. Starting small, but starting now.

Governance doesn’t require a big budget, it requires commitment.

Next week: Our straightforward suggestion starts by scheduling the first Federation Data Governance Council meeting. Invite one representative from five dioceses. Use the meeting to agree on single definitions, for example of the term “beneficiary.”

That’s it. One clear definition is the first brick in a trustworthy data foundation.

2. Pilot with purpose, not perfection.

The team already has AidHound. The Streamlit dashboard we built is a functional prototype. **Next month:** Running a three-diocese pilot. Use the dashboard to summarize Cárita's July aid requests. Not aiming for flawless integration, but rather for learning.

What works? What confuses people? The goal isn’t a perfect report, it’s a shared understanding of what AI can and cannot do for Cáritas Portuguesa teams.

3. Building literacy by doing, not just training.

Data literacy grows when people use tools that make their jobs easier.

Next quarter: Host a "data lunch" where staff from one diocese share how they used the dashboard to spot a trend in voucher requests. Learning happens peer-to-peer. Their local stewards will become their best trainers.

4. Treating ethics as a practice, not a policy.

An AI Acceptable Use Policy is essential, but policies gather dust unless they're lived.
Next step: Before deploying any AI tool, there are three questions to address:

1. Who could this affect unfairly?
2. Where is the human in the loop?
3. What will we do if it's wrong?

This is crucial. Ethics is how you act when no one is watching.

5. Leveraging partnerships with clear boundaries.

NOVA SBE, tech providers, and other Caritas networks are ready to help, but Cáritas must steer.
Next conversation: Implementation of a one-page collaboration charter for any external partner:

- What data stays inside Cáritas servers
- Who owns the outputs
- How decisions get made

6. Measuring what matters: impact, not activity.

It's easy to count reports generated or hours saved. Harder to track dignity preserved or trust strengthened.

Next evaluation: Pick one pilot and measure:

- Did staff feel more confident in decisions?
- Were beneficiaries heard more clearly?
- Did resources reach people faster?

The real risk isn’t moving too slowly. It’s standing still.

Cáritas Portuguesa has the mission, the network, and now the roadmap. What they do next will determine whether AI remains a buzzword or becomes a genuine force for good in the communities they serve.

Our suggestion is simple:

Start where you are. Use what you have. Do what you can. The rest will follow.

Important Sidenote

A distinctive characteristic of the current document is its descriptive orientation. This is consistent with the objectives of a field lab, where the emphasis lies in understanding an organisational problem, designing a solution, and demonstrating its practical relevance. The report follows this logic by presenting the operational context of Cáritas Portuguesa, detailing the system’s architecture, and explaining the rationale behind its technical components. Such an approach is appropriate for a project of applied nature as the primary goal is to articulate how the developed tool responds to the institution’s data management needs. Nonetheless, it is important to acknowledge that a more evaluative dimension could strengthen the analytical depth of the work. While the descriptive sections provide clarity on how each subsystem functions, they do not extensively assess the implications of those choices or compare them to alternative pathways. A more evaluative approach would include reflection on the trade-offs made throughout the development process, particularly regarding performance, usability, scalability, and organisational fit. This kind of examination not only demonstrates technical awareness but also signals an appreciation for the broader strategic factors influencing system deployment within a humanitarian organisation such as Cáritas. For instance, with more time a further discussion could address whether different architectural decisions could have reduced long-term maintenance costs and simplified onboarding for regional teams. It could also consider whether more advanced data validation procedures might have reduced manual cleaning efforts, or also how a more integrated design could have supported closer alignment with AidHound or other internal systems for the future. Please note that these reflections would not be criticizing the system as inadequate but rather situate the final design within a wider set of potential avenues and constraints.

Introduction Individual

Successful implementation of AI requires foundational data governance infrastructure that many organizations lack. Despite widespread AI adoption efforts, 95 percent of AI initiatives fail (Challapally and MIT NANDA 2025). These challenges extend beyond technical limitations to encompass workflow compatibility, organizational dynamics, and system characteristics. Challenges are particularly acute within federated nonprofit structures—organizations composed of multiple autonomous entities operating under a shared mission, such as the American Red Cross (769 chapters), United Way (1,800 local organizations), and faith-based networks like Caritas Portugal—where data fragmentation and inconsistent governance create additional implementation barriers.

1.1 The Gap in Literature

While federated structures are common in the nonprofit sector, there is limited research addressing data governance frameworks specifically designed for these decentralized organizations preparing for AI implementation. Existing literature on federated data governance focuses primarily on for-profit enterprises (Alation, 2024), where hierarchical authority and financial incentives enable compliance enforcement mechanisms unavailable in nonprofit federations. Humanitarian organizations have developed comprehensive data protection frameworks addressing privacy and ethical considerations in AI deployment (Beduschi et al. 2024), but these focus on compliance safeguards rather than the operational governance structures required to standardize data quality and schemas across independent entities.

This gap is particularly critical as federated nonprofits increasingly seek to implement AI and advanced analytics, which require consistent data quality and standardized schemas across all autonomous entities. Without appropriate governance frameworks that respect organizational autonomy while enabling coordination, these organizations risk implementing AI prematurely with inadequate data foundations.

1.2 Research Objective and Contribution

This research addresses this gap by adapting established data governance frameworks specifically for federated nonprofit structures preparing for AI implementation. The adapted framework synthesizes corporate federated governance principles (Alation, 2024; Khatri & Brown, 2010), data maturity progression models (DAMA International, 2017), and recent AI readiness requirements (Miller et al., 2025; Saeed et al., 2025) to establish operational standards and coordination mechanisms necessary for AI readiness while preserving entity independence. This work forms part of a larger team effort examining AI implementation at Caritas Portugal, where this governance framework provides the foundational data standards required for the AI applications proposed in the team's broader AI strategy roadmap.

Research Question: What governance framework can enable AI readiness in federated nonprofit organizations while respecting the autonomy of independent entities?

Section 2: Background

2.1: Data Governance Basics

2.1.1: Defining Data Governance

This research defines data governance as the system of decision rights, accountabilities, and processes that ensures appropriate management of data assets to meet organizational objectives (adapted from DAMA International, 2017; Khatri & Brown, 2010). In federated nonprofit contexts, this definition emphasizes the critical question of decision rights—determining which standards and processes are decided centrally versus locally—while maintaining accountability across autonomous entities.

2.1.2: Why Data Governance Matters in AI preparation

Data governance is the essential bridge between an organization's raw data and its capacity to successfully implement AI. The primary purpose of data governance in an AI context is to ensure a fundamental prerequisite: data quality. AI models are critically dependent on the quality of the data they are trained on; as the adage states, "garbage in, garbage out."

A comprehensive survey of data readiness for AI by Hiniduma et al. (2024) categorizes data quality as the foremost pillar of AI readiness. The authors state, "Data quality ensures that the data used to train AI models is accurate, complete, and consistent. High-quality data minimizes the risk of errors in AI outputs, leading to more reliable and trustworthy models" (Hiniduma et al., 2024). When this quality is compromised, the direct result is "inaccurate and unstable models."

In a federated nonprofit structure, where data is fragmented across autonomous entities, data governance is the only mechanism to address this challenge. It provides the "decision rights and accountabilities" (as defined in 2.1.1) necessary to create the consistent, standardized, and high-quality data that AI models require to function effectively. Without this governance framework, any AI initiative will fail because it lacks a usable data foundation.

2.1.3 Ideal Data Quality Standards

AI itself is just machine learning, and the key ways to make sure that we do not end up with the "garbage in, garbage out" failure of AI and machine learning models is by establishing the most "ideal" data standards put into place for our federated organizations. The "Gold Standard" of data standards across multiple organizations is "well represented across all frameworks" as "accuracy", "completeness", "consistency", and "timeliness". To better understand a clearly defined definition of each of these standards, the following are how we will define them:

- **Accuracy:** Data must be correct, precise, and valid. This includes both the values themselves and the labels.
- **Completeness:** Data must be free of missing values.
- **Consistency:** Data must be uniform and free of contradictions across different systems or entries.
- **Timeliness (or Currentness):** Data must be recent enough to be relevant to the model's task.

2.1.4 Adding Data Quality Standards for AI

While the foundational dimensions of accuracy and completeness are essential, modern AI models—and federated systems in particular—introduce stricter and more complex requirements. Simply having clean data at each node is not enough. The "gold standard" for a federated context must also include:

Homogeneity and Representativeness. This is perhaps the most critical requirement. The ideal state is a dataset that is homogenous and representative, meaning the data is similarly distributed across all nodes and accurately reflects the real-world problem (Saeed et al., 2025). This ideal stands in direct contrast to the key challenges of "data heterogeneity," "statistical heterogeneity," "class imbalances," and "data skew," which are all known to degrade model performance and lead to unreliable predictions (Saeed et al., 2025).

Semantic Consistency and Interoperability. For data to be aggregated from different autonomous sources, as in a federated nonprofit, it must have a shared meaning. This requires "semantics" and "portability" (Miller et al., 2025). In practical terms, data schemas, definitions, and categories must be standardized so that a "client" in one chapter's database means the exact same thing as a "client" in another's. This semantic link is what allows models to "understand and process" the combined data (Miller et al., 2025).

Sufficient Quantity and Availability. A small, siloed dataset, even if high-quality, is often insufficient. Modern AI models "require large quantities of available and accessible data to learn from" (Miller et al., 2025). This need for large-scale data makes the aggregation of data from all federated entities not just beneficial, but necessary.

2.2 Challenges for non-profits

While nonprofits show a growing interest in AI, they face significant hurdles in adoption. Many organizations are still in the early stages, and successful, comprehensive implementation remains limited. A 2025 benchmark report from TechSoup and Tapp Network provides a clear statistical picture of these internal challenges. The primary findings show a widespread lack of strategy, with 76% of nonprofit organizations having no AI strategy and 80% having no AI-acceptable use policy; furthermore, 10% of respondents are unsure if their organization even has a strategy. This is compounded by limited in-house expertise, as AI exploration is often driven by just a few individuals; only 5.9% of respondents identify as AI experts, and in 42% of organizations, only one or two people are trying to learn how to use AI. Decision-making is also siloed and overburdened, with one staff member responsible for these decisions in 43% of organizations. Consequently, successful adoption rates are low, as only 7.4% of nonprofits report they have successfully adopted AI to address core challenges, while 26.2% are not currently using AI at all. Finally, beyond these technical barriers, respondents expressed ethical and qualitative concerns: over half (53.4%) find it "too early" to determine if AI will prioritize ethical principles, citing specific fears of "privacy breach, data leakage," the "untrusted" accuracy of AI, and "losing the personal touch" critical to their mission. (TechSoup & Tapp Network, 2025)

Section 3: Adapted Framework - Three-Tier Governance Model for Federated Nonprofits

This section presents a three-tier data governance framework adapted specifically for federated nonprofit organizations preparing for AI implementation. The framework synthesizes established corporate governance principles with nonprofit operational realities to address the unique challenges of establishing data standards across legally autonomous entities while respecting organizational independence. It consists of three components: (1) a three-tier governance architecture defining decision rights at central, federated, and autonomous levels; (2) a five-level data quality maturity model for assessing AI readiness; and (3) an implementation roadmap guiding organizations through the governance development process.

3.1 Framework Development Approach and Foundational Models

This governance framework builds upon and adapts several established models rather than developing entirely novel concepts. The three-tier architecture draws heavily from Alation's (2025) corporate federated governance principles, adapting their decision rights framework for nonprofit contexts where hierarchical enforcement is unavailable. The maturity model structure follows DAMA International's (2017) established progression methodology, while incorporating AI-specific requirements identified by recent federated learning research (Miller et al., 2025; Saeed et al., 2025). Khatri and Brown's (2010) seminal work on governance decision rights provides the theoretical foundation for delineating central versus local authorities.

The innovation lies not in creating new governance concepts, but in synthesizing these established approaches to address the specific intersection of federated nonprofit structures and AI readiness, a context not adequately addressed by existing frameworks designed for either corporate environments or centralized nonprofits.

3.2 Three Tier Governance Architecture

Data governance models for federated structures typically follow three organizational approaches: centralized (corporate model), decentralized (affiliate model), or hybrid (franchise model) (Blackbaud, 2023). While these models address IT purchasing and organizational coordination, they do not specifically address data governance for AI readiness. Similarly, federated data governance frameworks exist in corporate contexts (Alation, 2024), but these assume hierarchical enforcement mechanisms and financial incentives unavailable in nonprofit federations. This framework adapts a three-tier DATA GOVERNANCE architecture specifically designed for federated nonprofits preparing for AI implementation, emphasizing decision rights and accountability structures appropriate for autonomous mission-driven organizations where coordination is voluntary.

Tier 1, the Central Coordination Layer, exists to establish minimum standards that enable cross-entity AI applications. Its decision rights include: standardizing core data definitions (e.g., "beneficiary," "service") to ensure the semantic consistency required for AI models to process combined data (Miller et al., 2025); setting minimum quality thresholds for accuracy (≥95%), completeness, and consistency, as poor quality yields "inaccurate and unstable models" (Hiniduma et al., 2024); establishing a baseline for privacy compliance and consent protocols (IASC, 2023); and managing dispute resolution for standards. This tier is governed by a federation council with entity representation, operating through consensus rather than mandate. Standards are voluntary, but participation in federated AI requires compliance. For example, Caritas national coordination could define "emergency assistance" consistently across all 20 dioceses, enabling cross-diocese demand prediction while each diocese controls its own local service delivery.

Tier 2, Federated Governance, is designed to enable peer coordination and shared learning among entities. Its key rights involve developing shared data collection templates adapted to regional contexts, establishing peer-to-peer working groups for common challenges (e.g., homeless services data), negotiating data sharing agreements between entities, and coordinating the adoption of shared platforms. This tier operates through voluntary structures like regional working groups and cross-entity data steward networks, not hierarchical authority. For instance, Northern Portugal dioceses could coordinate on homeless services data standards, sharing insights while maintaining independent operations, which addresses the challenge that federated systems work best when data is "similarly distributed" (Saeed et al., 2025).

Tier 3, the Autonomous Entity Layer, preserves local decision-making and operational flexibility. Entities retain all key rights beyond the Tier 1 minimums, including: defining all local program-specific data fields, maintaining direct control over beneficiary relationships and data collection consent, choosing their own operational data systems and collection methods (provided they meet or exceed Tier 1 standards), and implementing additional or higher quality standards based on local needs. Governance at this level is managed by local data stewards and entity-specific policies. For example, individual Caritas dioceses would continue to control their own volunteer management systems and local demographic data collection, independent of the national minimums.

This three-tier architecture balances the competing demands of federated nonprofit governance. Tier 1 provides sufficient standardization for AI applications, addressing the technical requirement for "semantics and portability" across autonomous sources (Miller et al., 2025). Tier 2 facilitates voluntary coordination and learning. Tier 3 preserves the entity autonomy essential to mission-driven organizations. By clearly delineating decision rights at each level, the framework avoids the ambiguity that causes governance failures while respecting the voluntary coordination model inherent to nonprofit federations.

3.3: Five-Level Data Quality Maturity Model

To assess AI readiness and guide governance implementation, this framework includes a five-level data quality maturity model adapted from established data management maturity frameworks (DAMA International, 2017). Organizations progress through these levels as they develop capabilities in data infrastructure, standardization, quality management, and governance structures. Level 3 represents the minimum threshold for AI implementation, addressing the requirement that AI models "require large quantities of available and accessible data" with sufficient quality and consistency (Miller et al., 2025). Organizations at Levels 1 or 2 must implement the three-tier governance architecture to progress toward AI readiness.

Level 1: Ad Hoc (Not AI-Ready)

Organizations at Level 1 manage data without formal processes or coordination. Each autonomous entity maintains separate spreadsheets or basic databases with no standardization across the federation. Data definitions vary significantly—one chapter's "beneficiary" may mean something different than another's—making aggregation impossible. There is no quality monitoring beyond reactive correction when errors cause operational problems. No governance structure exists; data management decisions are made individually by whoever handles the data at each entity.

AI Feasibility: None. The data heterogeneity and lack of standardization at Level 1 creates exactly the conditions that "degrade model performance and lead to unreliable predictions" (Saeed et al., 2025). Without consistent definitions and quality standards, any attempt at federated AI will fail.

Typical State: Most small federated nonprofits begin at Level 1, especially those that grew organically without central coordination.

Level 2: Aware (Not AI-Ready)

Organizations recognize the need for better data management but lack implementation. Some entities have migrated from spreadsheets to databases, but systems remain disconnected with minimal integration. Informal standards may exist—perhaps documented in an email or shared document—but are inconsistently applied across entities. Quality checks are manual and irregular, often triggered by specific reporting needs rather than systematic monitoring. Governance structures are discussed in meetings, and a working group may have been formed, but formal decision rights and accountability remain undefined.

AI Feasibility: Very limited. While awareness is growing, the lack of implementation means data still cannot be reliably aggregated. The "inconsistent labeling and quality variance" that persists at Level 2 significantly degrades potential AI model performance (Saeed et al., 2025).

Typical State: Many medium-sized federations operate at Level 2, having recognized data challenges but struggling with resource constraints and competing priorities that prevent full implementation.

Level 3: Defined (Minimum AI-Ready Threshold)

This level represents the minimum threshold for AI implementation. Organizations have documented data standards covering core definitions, and a majority of entities (typically 70-80%) have adopted these standards in their operational systems. Integrated databases enable cross-entity data aggregation for at least mission-critical data elements. Regular quality monitoring with defined metrics occurs quarterly or monthly, identifying issues proactively rather than reactively. The three-tier governance structure described in Section 3.2 is operational: a central coordination council meets regularly, regional working groups facilitate peer learning, and local data stewards are appointed and active.

AI Feasibility: Basic AI applications are possible. The standardization and quality management at Level 3 addresses the fundamental requirement that federated systems need "completeness, consistency, accuracy, and representativeness" (Saeed et al., 2025). However, AI applications remain limited to well-defined use cases with tolerance for some data quality issues. Advanced applications requiring real-time data or complex aggregations are not yet feasible.

Typical State: Mature federations that have invested 18-24 months in governance implementation typically reach Level 3, representing a significant organizational achievement.

Level 4: Managed (Fully AI-Ready)

Organizations at Level 4 have mature data governance fully implemented across all entities. A unified data platform enables seamless federation-wide aggregation and analysis. Standards are consistently enforced with metrics tracking compliance, and entities failing to meet standards receive support. Automated quality processes catch and correct issues in real-time or near-real-time. Clear accountability structures with defined roles ensure sustained governance—data stewards have authority and resources, governance councils have executive support, and the system operates smoothly with minimal intervention.

AI Feasibility: Full. Organizations at Level 4 can successfully implement advanced AI applications including predictive analytics, resource optimization, and beneficiary matching. The automated quality management ensures the "proper data preparation, including cleaning, normalization, and augmentation" essential for AI success (Saeed et al., 2025). The unified platform provides the "large quantities of available and accessible data" that modern AI models require (Miller et al., 2025).

Typical State: Only the most sophisticated federated nonprofits currently operate at Level 4, typically those with dedicated data teams and executive commitment to data-driven decision-making.

Level 5: Optimizing (AI-Native)

The highest maturity level represents organizations where data governance is strategic, proactive, and continuously improving. Real-time federated learning infrastructure enables advanced AI applications across the organization. Standards themselves evolve based on data-driven insights—the organization learns from its AI implementations and refines governance accordingly. AI drives quality improvement through automated monitoring, anomaly detection, and even automated correction of certain error types. Governance is anticipatory, identifying future data needs before they become critical. The organization views data as a strategic asset and governance as a competitive advantage.

AI Feasibility: Advanced and innovative. Organizations at Level 5 can implement cutting-edge AI applications and contribute to the field's advancement. They serve as models for other federated nonprofits.

Typical State: Very few organizations currently operate at Level 5. This represents an aspirational state for most federations, achievable only after years of sustained investment in data governance and AI capabilities.

Most federated nonprofits currently operate at Level 1 or 2. Reaching Level 3 typically requires 18-24 months of focused governance implementation following the roadmap in Section 3.4. Organizations should not attempt AI below Level 3, as inadequate data foundations lead to 95% failure rates (Challapally and MIT NANDA 2025).

3.4 DATA GOVERNANCE ROADMAP

Implementing this adapted governance framework requires a phased approach. The following 24-month roadmap guides federated nonprofits from their current state to Level 3 AI readiness (minimum threshold), organized into four sequential phases.

Phase 1: Assessment & Foundation (Months 1-6) Organizations conduct maturity assessment across all entities, map the current data landscape, establish the Tier 1 governance council, define 10-15 core data definitions for mission-critical concepts, and identify 2-3 pilot entities. This phase delivers a maturity assessment report, governance charter, and core definitions document, with success requiring full entity participation, operational governance council, and consensus approval of definitions.

Phase 2: Standardization (Months 7-12) Organizations roll out core definitions to all entities with training, establish data quality monitoring with defined metrics (accuracy ≥95%, completeness ≥95%), create Tier 2 regional coordination groups in 2-3 regions, deploy standardized data collection templates, and appoint local data stewards. Success requires 75% entity adoption of core definitions, operational quality monitoring, and active data stewards at all entities.

Phase 3: Integration (Months 13-18) Organizations implement a unified data platform, integrate data from at least 80% of entities, establish data sharing agreements, conduct the first comprehensive cross-entity quality audit, and pilot a basic AI application (e.g., demand forecasting) with volunteer entities. Success requires 80% entity integration, cross-entity data quality score ≥70%, and AI pilot demonstrating technical feasibility

Phase 4: Optimization (Months 19-24) Organizations expand the AI pilot to production deployment, implement automated quality monitoring for critical fields, establish continuous improvement processes, document lessons learned, and create a 3-year governance sustainability plan. Success is confirmed when Level 3 maturity is achieved, the AI application delivers measurable operational value, and governance becomes self-sustaining.

Roadmap Adaptation Guidance:

This roadmap assumes a medium-sized federation (10-30 entities) with limited technical resources. Organizations should adapt the timeline based on:

Size: Smaller federations (<10 entities) may compress to 15-18 months; larger federations (>50 entities) may require 30-36 months

Resources:

Organizations with dedicated data staff can accelerate; volunteer-dependent organizations should extend timelines to avoid burnout

Starting Point: Level 1 organizations need the full 24 months; Level 2 organizations already discussing governance can begin at Phase 2

Critical Success Factors:

Executive sponsorship, adequate resourcing (even if minimal), entity buy-in through inclusive governance, and realistic expectations about timeline

Common Pitfalls to Avoid:

Organizations should avoid four critical mistakes: (1) imposing standards without assessment undermines buy-in from excluded entities; (2) overly complex standards fail without adequate resources; (3) attempting AI before Level 3 wastes resources and damages confidence; and (4) treating governance as a one-time project rather than allocating resources for ongoing maintenance ensures degradation to lower maturity levels.

Individual Work Conclusion

Federated nonprofits face unique AI implementation challenges where the 95% failure rate is amplified by data fragmentation and entity autonomy. Existing corporate governance models cannot address these voluntary coordination structures.

This research adapts and synthesizes established data governance frameworks specifically for federated nonprofits. The Three-Tier Architecture balances minimum standardization with entity autonomy, while the Five-Level Maturity Model defines clear readiness thresholds and a 24-month implementation roadmap.

By establishing governance before technology, this adapted framework provides federated nonprofits a practical blueprint to collectively build the data infrastructure required for successful AI implementation while fulfilling their mission.

References

AidHound. (n.d.). AidHound platform documentation [Internal operational descriptions provided for institutional use].

Alation. (2024). Federated Data Governance Model: Critical Path for Organization Data Maturity Transformation Journey. White Paper. June.

Amazon Web Services. (2024). Registry of Open Data on AWS. Retrieved May 21, 2024, from https://registry.opendata.aws/

Aning, Tabiri E., & Nishant Agnihotri. (2024, March 13). Data Security: A Study of Data Anonymization and Several Techniques. https://ssrn.com/abstract=4757543

Araujo, Theo, Natali Helberger, Sophie Kruikemeier, & Claes H. de Vreese. (2020). In AI We Trust? Perceptions about Automated Decision-Making by Artificial Intelligence. AI & Society, 35, 611–623. https://doi.org/10.1007/s00146-019-00931-w

Aristidou, A., Dunckelman, A., & Fankuchen, S. (2025). How AI can deepen nonprofit relationships. Stanford Social Innovation Review, Fall 2025. https://doi.org/10.48558/YQ8C-9H56

Arroyabe, M. F., Arranz, C. F. A., Fernández De Arroyabe, I., & Fernández de Arroyabe, J. C. (2024). Analyzing AI adoption in European SMEs: A study of digital capabilities, innovation, and external environment. Technology in Society, 79, Article 102733.

Arrillaga, Elisha Smith, Seara Grundhoefer, & Christina Im. (2025, September 28). AI with Purpose: How Foundations and Nonprofits Are Thinking About and Using Artificial Intelligence. https://cep.org/report-backpacks/ai-with-purpose-how-foundations-and-nonprofits-are-thinking-about-and-using-artificial-intelligence/?section=intro

Banks, N., Hulme, D., & Edwards, M. (2015). NGOs, States, and Donors Revisited: Still Too Close for Comfort? World Development, 66, 707–718.

Barbosa, S., & Félix, S. (2022). Algorithms and the GDPR: An Analysis of Article 22. Universidade Nova de Lisboa. https://protecaodedadosue.cedis.fd.unl.pt/wp-content/uploads/2022/10/3.-Sandra-Barbosa.pdf

Beduschi, Ana, Massimo Marelli, and Aaron Martin, eds. (2024). Data Protection in Humanitarian Action: Responding to Crises in a Data-Driven World. London: Routledge.

Bento, P. I. P. R. (2021). A sustainable approach for data governance in organizations [Master's thesis, NOVA School of Business and Economics]. https://run.unl.pt/handle/10362/134831

Blackbaud. (2023, April 19). 3 Federated Models for Larger Nonprofits and How Each Model Impacts Decision Making. The ENGAGE Blog by Blackbaud.

Bryson, J. M., Crosby, B. C., & Stone, M. M. (2015). Designing and implementing cross-sector collaborations: Needed and challenging. Public Management Review, 17(5), 673–689. https://doi.org/10.1080/14719037.2015.1025851

Camões, D. (2024). The Challenges of the GDPR in the Era of Artificial Intelligence: What Can We Expect from the Future? e-Publica, 11(3). https://e-publica.pt/api/v1/articles/127705-the-challenges-of-the-gdpr-in-the-era-of-artificial-intelligence-what-can-we-expect-from-the-future.pdf

Caritas Internationalis. (n.d.). Mission, identity, and organizational norms. Caritas Internationalis.

Caritas Portuguesa. (2024). Plano Estratégico Cáritas em Portugal 2024–2030. Cáritas Portuguesa.

Caritas Portuguesa. (2025). Plano de Atividades e Orçamento 2025. Cáritas Portuguesa.

Caritas Portuguesa. (n.d.). Website institucional. https://caritas.pt

Caritas Portuguesa. (n.d.). Cáritas Build-Out to Help with Thesis. Internal document provided by the organization.

Cáritas Portuguesa. (n.d.). Formulários internos de apoio social [Excel file].

Castelo, N., Bos, M. W., & Lehmann, D. R. (2019). Task-Dependent Algorithm Aversion. Journal of Marketing Research, 56(5), 809–825. https://doi.org/10.1177/002224371986

Challapally, Aditya, and MIT NANDA. (2025). The GenAI Divide: State of AI in Business 2025. Cambridge, MA: MIT NANDA Initiative.

Charity Navigator. (2024). Annual Report FY2024. https://www.charitynavigator.org/content/dam/cn/cn/pdf/CN-AnnualReport-FY24.pdf

Costa, R. R. (2025, November 17). Cáritas Braga Launches Campaign to Build Houses for Victims of Violence. JN. https://www.jn.pt/pais/artigo/Cáritas-de-braga-com-campanha-para-construir-casa-para-vitimas-de-violencia/18020831

DAMA International. (2017). DAMA-DMBOK: Data Management Body of Knowledge (2nd ed.). Technics Publications.

DiGiovanni, J. (2025, January 24). Game-changing 2025 AI benchmark report: How artificial intelligence is revolutionizing the nonprofit sector. Tapp Network Blog. https://blog.tappnetwork.com/tapp-and-techsoup-release-2025-ai-benchmark-report

Dorsch, J., & Deroy, O. (2024). Quasi-Metacognitive Machines: Why We Don't Need Morally Trustworthy AI and Communicating Reliability Is Enough. Philosophy & Technology, 37(62). https://doi.org/10.1007/s13347-024-00752-w

EDPS. (2023). TechDispatch #2/2023: Explainable Artificial Intelligence (XAI). https://www.edps.europa.eu/system/files/2023-11/23-11-16_techdispatch_xai_en.pdf

Efthymiou, Ioannis P. (2023). "The Role of Artificial Intelligence in Revolutionizing NGOs' Work: Opportunities and Challenges." Journal of Politics and Ethics in New Technologies and AI, 1(1), 1–20. https://www.researchgate.net/publication/372824365_The_Role_of_Artificial_Intelligence_in_Revolutionizing_NGOs'_Work

Elamin, M. O. I. (2024). Modernizing the Charitable Sector through Artificial Intelligence: Enhancing Efficiency and Impact. Journal of Ecohumanism, 3(4), 3426–3443. https://doi.org/10.62754/joe.v3i4.3856

European Commission. (2019). Ethics Guidelines for Trustworthy AI. https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai

European Commission. (2025, August 1). AI Act. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

European Union. (2024, August 1). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

Expresso. (2025, March 18). Cáritas Deixa Alertas sobre Imigrantes, Sem-Abrigo e Acesso à Habitação e Pede Impulso Urgente contra a Pobreza e Exclusão.

https://expresso.pt/sociedade/2025-03-18-Cáritas-deixa-alertas-sobre-imigrantes-sem-abrigo-e-acesso-a-habitacao-e-pede-impulso-urgente-contra-a-pobreza-e-exclusao-2b16ecb5

Ganesh, P., Tran, C., Shokri, R., & Fioretto, F. (2025). The Data Minimization Principle in Machine Learning. https://doi.org/10.48550/arXiv.2405.19471

Glikson, E., & Woolley, A. W. (2020). Human Trust in Artificial Intelligence: Review of Empirical Research. Academy of Management Annals, 14(2), 627–660. https://doi.org/10.5465/annals.2018.0057

Goldkind, L., Jian, M., & Fink, A. (2025). AI in the Nonprofit Human Services: Distinguishing Between Hype, Harm, and Hope. Human Service Organizations, 49(3), 225–236. https://doi.org/10.1080/23303131.2024.2427459

Hiniduma, Kaveen, Suren Byna, and Jean Luca Bez. (2025). Data Readiness for AI: A 360-Degree Survey. ACM Computing Surveys, 57(9), 219.

Inter-Agency Standing Committee (IASC). (2023). Operational Guidance: Data Responsibility in Humanitarian Action. Geneva: IASC Data Responsibility Working Group, April.

Khatri, Vijay, and Carol V. Brown. (2010). Designing Data Governance. Communications of the ACM, 53(1), 148-152. https://doi.org/10.1145/1629175.1629210

Kim, M., Charbonneau, É., & Sowa, J. E. (2024). The Nonprofit Starvation Cycle: The Extent of Overhead Ratios' Manipulation, Distrust, and Ramifications. Nonprofit and Voluntary Sector Quarterly, 54.

Klingbeil, A., Grützner, C., & Schreck, P. (2024). Trust and Reliance on AI—An Experimental Study on the Extent and Costs of Overreliance on AI. Computers in Human Behavior, 160, 108352. https://doi.org/10.1016/j.chb.2024.108352

Krause, D. (2025, January 4). AI Agents and Automation in Small Non-Profit Organizations' Accounting Functions. SSRN. https://www.researchgate.net/publication/387723705_AI_Agents_and_Automation_in_Small_Non-Profit_Organizations'_Accounting_Functions

Kuner, C. (2020). Transborder Data Flows and Data Privacy Law (2nd ed.). Oxford University Press.

Larsson, S., & Heintz, F. (2020). Transparency in Artificial Intelligence. Internet Policy Review, 9(2), 1–16. https://doi.org/10.14763/2020.2.1469

Leitoniene, S., Waniak-Michalak, H., & Perica, I. (2025). The Digitalization of Non-Governmental Organizations: Factors and Organizational Characteristics. International Review on Public and Nonprofit Marketing, 22(3), 549–569.

Lomba, P., & Roldão, M. S. (2025, June 20). AI Act: Commission guidelines on high-risk systems. PLMJ. https://www.plmj.com/en/knowledge/informative-notes/AI-Act-Commission-guidelines-on-high-risk-systems/33988

Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to Medical Artificial Intelligence. Journal of Consumer Research, 46(4), 629–650. https://doi.org/10.1093/jcr/ucz013

Miller, Russell, Sai Hin Matthew Chan, Harvey Whelan, and João Gregório. (2025). A Comparison of Data Quality Frameworks: A Review. Big Data and Cognitive Computing, 9(4), 93.

Mitchell, T. (2025). Trust and Transparency in Artificial Intelligence. Philosophy & Technology, 38, 87. https://doi.org/10.1007/s13347-025-00916-2

Muehlemann, S. (2025). Artificial Intelligence Adoption and Workplace Training. https://www.sciencedirect.com/science/article/pii/S0167268125003257

NetHope. (2020, October 6). AI in the Humanitarian Sector. https://nethope.org/articles/ai-in-the-humanitarian-sector/

NetHope. (2025). Data Governance Toolkit: A Guide to Implementing Data Governance in Nonprofits. Falls Church, VA: NetHope.

Nova School of Business and Economics. (2017). Leadership for Impact Knowledge Center. https://www.novasbe.unl.pt/en/leadership-for-impact

Nova School of Business and Economics. (2020). Social sector database. https://www.novasbe.unl.pt/en/about-us/projects-for-a-better-future/social-equity-initiative/social-database

OECD. (2019). Recommendation of the Council on Artificial Intelligence. https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449

OECD. (2024, December). Digital public infrastructure for digital governments. https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/12/digital-public-infrastructure-for-digital-governments_11fe17d9/ff525dc8-en.pdf

Organisation for Economic Co-operation and Development. (2019). OECD AI Principles overview. https://oecd.ai/en/ai-principles

Peck, D. (2025, January 3). Employee Training Statistics, Trends, and Data in 2025. https://www.devlinpeck.com/content/employee-training-statistics

Pizzi, Michael, Mila Romanoff, and Tim Engelhardt. (2021). AI for Humanitarian Action: Human Rights and Ethics. International Review of the Red Cross. https://doi.org/10.1017/S1816383121000011

Plaisance, G. (2025). Artificial Intelligence (AI) in the Context of Nonprofits and Philanthropy: Suspicion and Hope for Researchers and Organizations. https://doi.org/10.1002/nvsm.70022

Riedl, M. O. (2020). Human-Centered Artificial Intelligence and Machine Learning. https://arxiv.org/abs/1901.11184

Rizun, N., Edelmann, N., Janowski, T., & Revina, A. (2025). AI-enabled co-creation for evidence-based policymaking. https://doi.org/10.59490/dgo.2025.991

Rosenberger, J., Kuhlemann, S., Tiefenbeck, V., Kraus, M., & Zschech, P. (2025). The Impact of Transparency in AI Systems on Users' Data-Sharing Intentions: A Scenario-Based Experiment. https://arxiv.org/abs/2502.20243

Saeed, Nawraz, Mohamed Ashour, and Maggie Mashaly. (2025). Comprehensive Review of Federated Learning Challenges: A Data Preparation Viewpoint. Journal of Big Data, 12, 153.

Saqr, M. (2022). Is GDPR Failing? A Tale of the Many Challenges in Interpretations, Applications, and Enforcement. International Journal of Health Sciences (Qassim), 16(5), 1–2.

Sartor, G. (2020). [Full citation needed - AI governance and GDPR as backbone]

Saxena, S., & D'Ignazio, C. (2024). Operationalizing AI ethics through multi-stakeholder co-creation: Learning, power, and governance in practice. AI & Society. https://doi.org/10.1007/s00146-024-01937-9

Scott, K. T. (2025). AI For the Common Good: Exploring Global Nonprofit Framing of Artificial Intelligence. https://www.proquest.com/openview/cc52a75e1e7c59c7adda03e42f448024/1?pq-origsite=gscholar&cbl=18750&diss=y

Shneiderman, B. (2020, February 23). Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy. https://arxiv.org/abs/2002.04087v1

TechSoup. (2024, April 25). Case Study: Nonprofits Leveraging Microsoft 365 Copilot for Impact. https://blog.techsoup.org/posts/case-study-nonprofits-leveraging-microsoft-365-copilot-for-impac

TechSoup and Tapp Network. (2025). The State of AI in Nonprofits: 2025 Benchmark Report on Adoption, Impact, and Trends. San Francisco: TechSoup. https://page.techsoup.org/ai-benchmark-report-2025

Tapp Network & TechSoup. (2024). The State of AI in Nonprofits: 2025 Benchmark Report on Adoption, Impact, and Trends.

Uddin, S. F., Khan, I., Aleem, S. A., & Farheen, M. (2025). AI in Employee Training and Development: A Literature Review of Opportunities and Challenges. International Journal of Science and Management Studies, 8(5), 373–377. https://doi.org/10.51386/25815946/ijsms-v8i5p123

UNICEF Evaluation Office. (2024). An Operational Framework for Machine Learning in Evaluation. New York: UNICEF. https://www.unicef.org/evaluation/media/5606/file/ML%20Operational%20Framework%20FINAL.pdf.pdf

Van Grembergen, W., De Haes, S., & Guldentops, E. (2004). Structures, processes and relational mechanisms for IT governance. https://doi.org/10.4018/978-1-59140-140-7.ch001

Voigt, P., & von dem Bussche, A. (2017). The EU General Data Protection Regulation (GDPR): A Practical Guide. Springer.

Wang, Z., Liu, T., Liu, Y., Zio, E., & Guan, X. (2025). Data Inference: Data Security Threats in the AI Era. Engineering, 52, 29–33. https://doi.org/10.1016/j.eng.2025.08.007

Wong, P. H., & Santos, R. (2023). The demand for algorithmic transparency in public governance: An experimental study. https://www.novasbe.unl.pt/en/research/knowledge/working-paper-series/working-paper/the-demand-for-algorithmic-transparency-in-public-governance-an-experimental-study

Yang, Q., Lepore, C., Eynard, J., & Laborde, R. (2025). From Theory to Practice: Data Minimisation and Technical Review of Verifiable Credentials under the GDPR. Computer Law & Security Review, 57, 106138. https://hal.science/hal-05250812

Appendix

Roadmap Implications

Given Level 1 assessment, Caritas requires the full 24-month governance implementation roadmap. The organization must progress sequentially through all four phases:

Phase 1 (Months 1-6): Caritas must establish the Tier 1 governance council, conduct comprehensive maturity assessment across all dioceses, and define core data definitions for fundamental concepts. The current "common denominator" dashboard approach must evolve into deliberate standardization with documented definitions and mandated data collection standards.

Phase 2 (Months 7-12): Address the data cleaning opacity by appointing local data stewards at each diocese, documenting quality procedures, and establishing proactive quality monitoring with defined metrics (accuracy ≥95%, completeness ≥95%). The double-aggregation process requires redesign to ensure single-point aggregation with transparent quality controls.

Phase 3 (Months 13-18): Leverage existing AidHound infrastructure but implement proper integration protocols, data sharing agreements between dioceses, and comprehensive quality audit processes. The technical platform exists; governance processes do not.

Phase 4 (Months 19-24): Establish continuous improvement processes and demonstrate AI value through pilot applications, but only after governance foundations become operational.

Any attempt to implement AI applications before achieving Level 3 maturity would fail due to inadequate data foundations—precisely the scenario this framework prevents. Caritas's current state validates the framework's core thesis: federated nonprofits require governance structures before pursuing AI, regardless of technical infrastructure sophistication.

Maturity Model Overview Table

Level	Data Infrastructure	Standards and Definitions	Quality Management	AI Ready?
1	Siloed spreadsheets	No documented standards	Reactive error correction	No
2	Minimal Integration	Informal, inconsistent standards	Manual Spot checks	No
3	Majority of information integrated	Core standards adopted	Regular monitoring	Yes (Minimum Threshold)
4	Complete unified Platform	Standards enforced	Automated processes	Yes

Level	Data Infrastructure	Standards and Definitions	Quality Management	AI Ready?
5	Real-Time Federated data	Standards enforced but consistently refined	AI-driven quality	Yes (Advanced AI capabilities)

3 Tier Governance Structure

Decision Type	Tier 1 Central	Tier 2 Regional	Tier 3 Local
Establishes definitions	Defines for everyone	Uses definition	Uses definitions
Minimum data Quality	Sets the Floor (ex 95% accuracy)	Can exceed if desired	Must meet floor, can exceed
Privacy & consent rules	Creates baseline policy	May add regional requirements	Obtains actual consent, follows policy
Local Program Details	No authority	Can coordinate with peers	Full control
Sharing Data between entities	Approval power on frameworks	Negotiate specific agreements	Decides whether to participate

Self-Assessment Framework

Organizations can assess their current maturity level by answering four key questions:

1. Infrastructure: Can we currently aggregate data from all our entities into a unified view for analysis?

• No integration = Level 1

• Some integration = Level 2

• Majority integrated = Level 3

• All integrated = Level 4+

2. Standards: Are our core data definitions documented and consistently applied across entities?

• No documentation = Level 1

• Documented but inconsistent = Level 2

• Documented and mostly adopted = Level 3

• Fully standardized and enforced = Level 4+

3. Quality: Do we have regular, proactive monitoring of data quality with defined metrics?

• Reactive only = Level 1

• Manual spot-checks = Level 2

• Regular monitoring = Level 3

• Automated monitoring = Level 4+

4. Governance: Are roles, responsibilities, and decision rights clearly defined and operational?

• No governance = Level 1

• Discussed but not implemented = Level 2

• Basic structure operational = Level 3

• Mature with clear accountability = Level 4+