

TabAttackBench: A benchmark for adversarial attacks on tabular data

Zhipeng He ^{id a,b,*}, Chun Ouyang ^{id a,b}, Lijie Wen ^{id c}, Cong Liu ^{id d}, Catarina Moreira ^{id e,b,f}

^a School of Information Systems, Queensland University of Technology, Brisbane, Australia

^b Center for Data Science, Queensland University of Technology, Brisbane, Australia

^c School of Software, Tsinghua University, Beijing, China

^d NOVA Information Management School, NOVA University of Lisbon, Lisboa, Portugal

^e Data Science Institute, University of Technology Sydney, Sydney, Australia

^f INESC-ID/Instituto Superior Técnico, University of Lisboa, Lisboa, Portugal

ARTICLE INFO

Keywords:

Adversarial attack
Tabular data
Benchmark
Machine learning
Robustness

ABSTRACT

Adversarial attacks pose a significant threat to machine learning models by inducing incorrect predictions through imperceptible perturbations to input data. While these attacks are well studied in unstructured domains such as images, their behaviour on tabular data remains underexplored due to mixed feature types and complex inter-feature dependencies. This study introduces a comprehensive benchmark that evaluates adversarial attacks on tabular datasets with respect to both effectiveness and imperceptibility. We assess five white-box attack algorithms (FGSM, BIM, PGD, DeepFool, and C&W) across four representative models (LR, MLP, Tab-Transformer and FT-Transformer) using eleven datasets spanning finance, energy, and healthcare domains. The benchmark employs four quantitative imperceptibility metrics (proximity, sparsity, deviation, and sensitivity) to characterise perturbation realism. The analysis quantifies the trade-off between these two aspects and reveals consistent differences between attack types, with ℓ_∞ -based attacks achieving higher success but lower subtlety, and ℓ_2 -based attacks offering more realistic perturbations. The benchmark findings offer actionable insights for designing more imperceptible adversarial attacks, advancing the understanding of adversarial vulnerability in tabular machine learning.

1. Introduction

In recent years, the field of machine learning has seen substantial advancements, leading to the deployment of models across a wide range of applications. However, with these advancements comes increasing concern about the robustness and security of models, particularly in the context of adversarial attacks. Adversarial attacks involve the intentional manipulation of input data to deceive machine learning models, causing incorrect or misleading outputs (Szegedy et al., 2014). This area of research has drawn significant attention as researchers strive to understand and mitigate the vulnerabilities in various types of data and models. Adversarial perturbations to images involve pixel intensity modifications (Weng et al., 2024), spatial transformations (Aydin & Temizel, 2023), texture perturbations (Geirhos et al., 2018), and localised patches (Wang et al., 2025) that cause dramatic misclassifications while remaining visually imperceptible in Computer Vision (CV). Similarly, in Natural Language Processing (Zhang et al., 2020), attacks typically involve word substitutions (Yang et al., 2023), character-level modifications (Rocamora et al., 2024), or syntactic transformations (Asl et al.,

2024) that preserve semantic meaning while fooling text classifiers (Gao et al., 2024). Adversarial vulnerabilities have also been demonstrated in audio processing (Nouredine et al., 2023) through amplitude modifications (Ko et al., 2023), frequency perturbations (Abdullah et al., 2019), and psychoacoustic masking (Qin et al., 2019) that cause speech recognition systems to misinterpret commands. By addressing the vulnerabilities in these types of data, researchers aim to develop more robust and secure machine learning systems across various domains.

1.1. Challenges in adversarial attacks on tabular data

Tabular data, structured yet rich in semantics, heterogeneity, and interdependencies, is prevalent in domains such as finance, healthcare, and e-commerce (Borisov et al., 2022). These datasets often contain vital information used for decision-making processes such as credit scoring (Mushava & Murray, 2024), medical diagnosis (Khosravi et al., 2023), and fraud detection (Yi et al., 2023). Despite their significance, machine learning models trained on tabular data remain comparatively underexplored regarding their vulnerabilities to adversarial manipulation. Even

* Corresponding author.

E-mail addresses: zhipeng.he@hdr.qut.edu.au (Z. He), c.ouyang@qut.edu.au (C. Ouyang), wenj@tsinghua.edu.cn (L. Wen), cliu@novaims.unl.pt (C. Liu), catarina.pintomoreira@uts.edu.au (C. Moreira).

<https://doi.org/10.1016/j.eswa.2025.130491>

Received 1 June 2025; Received in revised form 12 October 2025; Accepted 16 November 2025

Available online 19 November 2025

0957-4174/© 2025 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Table 1

Illustrative example of adversarial perturbations in the *German Credit* dataset. Minor yet realistic feature changes can alter the model prediction while preserving record plausibility.

Feature	Original Record	Perturbed Record
Duration (months)	24	31.3 (↑)
Credit amount (DM)	2384	4,606 (↑)
Age (years)	64	58.6 (↓)
Account check status	< 0 DM	< 0 DM
Credit history	existing credits paid back duly	same
Purpose	domestic appliances	same
Property	real estate	real estate
Housing	rent	rent
Model prediction	No Default (✓)	Default (×)

though small, carefully crafted perturbations can significantly alter predictions (Ballet et al., 2019), as illustrated in Table 1.

Adversarial perturbations in tabular data arise through subtle yet realistic modifications of feature values — such as adjusting a numerical attribute by a small increment, or flipping a categorical feature within its valid set — to shift a model’s decision boundary while maintaining a realistic record. For instance, in electronic health records, marginal changes to clinical indicators might lead to misclassification of disease risk (An et al., 2019); in fraud detection, slightly altering transaction frequency or amount could conceal illicit behaviour (Lunghi et al., 2023); and in loan approval systems, minor modifications to income or credit utilisation ratios could influence acceptance outcomes (Gunasekaran et al., 2023).

These examples illustrate the broader issue: adversarial vulnerabilities in tabular data are not confined to a specific domain but are inherent to the data’s structural and semantic complexity. Unlike image or text modalities, tabular datasets combine heterogeneous feature types (categorical and numerical) often with varying scales, missing values, and intricate interdependencies (Borisov et al., 2022). These complexities make applying adversarial attacks to tabular data more intricate compared to image or text data.

1.2. Benchmarking adversarial attacks on tabular data

To systematically study above challenges for tabular data, it is essential to establish *benchmarks*—standardised frameworks that evaluate how different attack algorithms perform under controlled and comparable conditions on tabular data. While considerable progress has been made in understanding adversarial attacks on different modalities such as images, graphs, and time series (see Table 2), there remains a relatively underexplored area: adversarial attacks on tabular data. Existing benchmarks have primarily focused on image data, employing metrics such as attack success rate, robust accuracy, and norm-based metrics (e.g., ℓ_∞ , ℓ_2) to quantify the strength of the adversarial perturbations. While these metrics are well-suited for image data, where imperceptibility is measured by slight pixel changes that remain “indistinguishable to the human eye” (Goodfellow et al., 2015), they do not translate directly to tabular data.

A recent effort on tabular data, Simonetto et al. (2024b), provides a valuable benchmark focusing on dataset and model-level evaluation under a single attack setup, emphasising data augmentation and robust accuracy. In contrast, our benchmark is designed to systematically evaluate and compare *multiple attack algorithms* across various tabular datasets. Crucially, it integrates the concept of imperceptibility for tabular data by focusing on four key quantitative properties: **Proximity**, **Sparsity**, **Deviation**, and **Sensitivity** (He et al., 2025). These properties ensure that adversarial examples closely resemble original data, minimise feature alterations, and respect the statistical distribution of the data.

By integrating measures of *attack effectiveness* and *imperceptibility*, we propose a benchmark that jointly evaluates how successful and how realistic adversarial attacks are on tabular data.

1.3. Contribution

This study seeks to address three open research questions:

1. *How effective are existing adversarial attack algorithms on tabular data?*
2. *How imperceptible are existing adversarial attacks when evaluated against the imperceptibility properties?*
3. *To what extent can existing adversarial attacks achieve effectiveness while remaining imperceptible?*

To answer these questions, this paper makes the following key contributions:

- **Comprehensive benchmark for tabular data.** We introduce a benchmark that systematically evaluates representative *white-box* adversarial attacks on diverse tabular datasets. The benchmark jointly assesses *attack effectiveness* (e.g., Attack Success Rate) and *imperceptibility*, providing the first unified evaluation framework for structured data.
- **Imperceptibility metrics grounded in prior theory.** Building upon our earlier formal framework (He et al., 2025), we operationalise four quantitative imperceptibility metrics, including *Proximity*, *Sparsity*, *Deviation*, and *Sensitivity*. These metrics capture how closely adversarial examples align with original data distributions, extending the notion of “indistinguishability” beyond image norms to heterogeneous tabular features.
- **Empirical analysis of trade-offs.** Through extensive experiments across multiple datasets and models, we reveal quantitative trade-offs between attack success and imperceptibility. Our findings highlight which attack families tend to favour realism versus disruption, offering practical guidance for future algorithm design.

The remainder of this paper unfolds as follows. Section 2 reviews existing research and taxonomies that define the scope of this benchmark. Section 3 then introduces the proposed benchmarking methodology, detailing how both effectiveness and imperceptibility are systematically assessed across representative attacks and datasets. Section 4 reports the experimental results and analyses the quantitative trade-offs between attack success and imperceptibility. Finally, Section 5 reflects on the key insights, practical implications, and future research directions arising from this work.

2. Background and related work

The study of adversarial machine learning originates from the recognition that predictive models, despite demonstrating strong performance under standard evaluation, can be highly fragile when faced with carefully crafted perturbations. The seminal work by Szegedy et al. (2014) first revealed that small, often imperceptible perturbations applied to image data could cause state-of-the-art classifiers to produce incorrect predictions. Subsequent work by Goodfellow et al. (2015) provided both a theoretical explanation and efficient algorithms for generating such adversarial examples, sparking a broad and active line of research. At its core, adversarial machine learning investigates the interaction between an adversary, who deliberately perturbs inputs within allowable modifications, and a predictive model, whose robustness is evaluated against these perturbations.

Robustness analysis requires us to extend the standard learning framework by considering how models behave when their inputs are deliberately modified. Not all modifications are meaningful: some perturbations would lead to invalid or semantically inconsistent inputs. To address this, ML introduces the notion of *perturbation sets*, which formalise the allowable neighbourhoods around a given instance.

Table 2

Overview of existing benchmarks on adversarial attacks across different data types, attack scopes, and evaluation metrics. Our work uniquely focuses on benchmarking multiple *attack algorithms* and their *imperceptibility* on tabular data.

Benchmark	Data Type	Attack Type	Evaluation Metric
Jin et al. (2024)	Image	Transferable Attacks	Attack Transferability Score
Dong et al. (2020)	Image	White-box, Black-box Attacks	ℓ_∞ Norm, Robust Accuracy
Zheng et al. (2025)	Image	Black-box Attacks	Attack Success Rate, Query Count
Croce et al. (2021)	Image	ℓ_∞, ℓ_2 Norm-based Attacks	ℓ_∞, ℓ_2 Norm, Robust Accuracy
Hingun et al. (2023)	Image	Patch-based Attacks	Attack Success Rate, Realism Score
Cinà et al. (2025)	Image	Gradient-based Attacks	Attack Success Rate
Zheng et al. (2021)	Graph	Graph Modification	Robust Accuracy
Li et al. (2021)	VQA	Visual-Textual Attacks	Robust Accuracy
Siddiqui et al. (2020)	Time-series	Temporal Attacks	Attack Success Rate
Simonetto et al. (2024b)	Tabular	Single Tabular Attack	Robust Accuracy
<i>Our paper</i>	<i>Tabular</i>	<i>Multiple Tabular Attacks</i>	<i>Attack Success Rate, Imperceptibility Metrics</i>

Definition 1 (Perturbation set). For an instance $x \in \mathcal{X}$, a *perturbation set* $\Delta(x) \subseteq \mathcal{X}$ is the collection of all instances x' that are considered valid modifications of x . Formally, $\Delta(x)$ defines the allowable neighbourhood of x under which adversarial examples may be sought.

In the context of *tabular data*, the construction of $\Delta(x)$ must respect the heterogeneous nature of features. For continuous features, perturbations are typically restricted to bounded changes (e.g., $\|x' - x\|_\infty \leq \epsilon$) or minimal perturbation (e.g., $\|x' - x\|_2$), while for categorical features, perturbations are constrained to valid substitutions within the feature's domain. This ensures that all perturbed instances $x' \in \Delta(x)$ remain semantically meaningful and valid.

An *adversarial example* is an instance obtained by deliberately perturbing a valid input within its perturbation set, with the goal of challenging the model's prediction.

Definition 2 (Adversarial example). Let $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ be an ML classifier, and $(x, y) \in \mathcal{X} \times \mathcal{Y}$ be a correctly classified instance, i.e., $f_\theta(x) = y$. An *adversarial example* for x is a perturbed instance $\tilde{x} \in \Delta(x)$, often expressed as $\tilde{x} = x + \delta$ with $\delta \in \Delta(x)$, such that $f_\theta(\tilde{x}) \neq y$. If the prediction is instead forced into a specified target label $y_t \neq y$, then \tilde{x} is referred to as a *targeted adversarial example*, satisfying $f_\theta(\tilde{x}) = y_t$.

Formally, adversarial examples are generated by an *adversarial attack*, which searches for $\tilde{x} \in \Delta(x)$ that maximises some misclassification objective. Whether the resulting adversarial example actually causes a prediction change depends on the model under attack.

Definition 3 (Adversarial attack). Let $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ be a classifier, $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$ a loss function, and $\Delta(x) \subseteq \mathcal{X}$ a perturbation set. An *adversarial attack* is an optimisation procedure that, given a correctly classified instance (x, y) , searches over candidates $x' \in \Delta(x)$ and outputs an adversarial example $\tilde{x} \in \Delta(x)$ that satisfies a misclassification objective.

2.1. The taxonomy of adversarial attacks

Given the wide variety of adversarial attack strategies proposed in the literature, researchers have established several taxonomies to characterise and compare them (Assion et al., 2019; Biggio & Roli, 2018; Papernot et al., 2016; Pelekis et al., 2025; Sadeghi et al., 2020). These taxonomies are not mutually exclusive, but together they provide a systematic framework for understanding the assumptions, objectives, and mechanisms of different attacks. Taken collectively, these dimensions define the *threat model* of an adversary, which specifies the conditions under which an attack is mounted. Four of the most widely adopted dimensions for classification in Fig. 1, including: (1) the attacking stage in the model lifecycle (*Adversary's Influence*), (2) the adversary's knowledge of the model (*Adversary's Knowledge*), (3) the constraints imposed on perturbations (*Adversary's Capabilities*), and (4) the goal of the attack in terms of prediction outcomes (*Adversary's Goals*).

2.1.1. Adversary's influence

A complementary dimension for classifying adversarial attacks concerns the *influence* or attack surface available to the adversary. This axis describes when and where the adversary acts relative to the model lifecycle and data pipeline. Three primary influence categories are commonly recognised:

- **Evasion (test-time) attacks** (Biggio et al., 2013; Biggio & Roli, 2018): the adversary perturbs inputs at inference time so that the deployed model produces incorrect outputs without altering the training data or model parameters.
- **Poisoning (training-time) attacks** (Biggio & Roli, 2018): the adversary injects or modifies examples in the training set (or manipulates the training process) to compromise model behaviour at deployment, e.g., by raising the error on a target subgroup or enabling backdoor triggers.
- **Privacy (information) attacks** (Shokri et al., 2017): the adversary seeks to extract sensitive information about training data, model internals, or users (e.g., via membership inference, model inversion or parameter extraction), thereby compromising confidentiality rather than directly causing misclassification.

In this benchmark, we concentrate on *evasion attacks* at test time, as they directly instantiate the adversarial risk framework introduced earlier and provide the most immediate measure of the model robustness when deployed in real-world decision-making settings.

2.1.2. Adversary's knowledge

Attacks are commonly distinguished by the amount of information available to the adversary. They are categorised into three primary types of attacks based on the adversary's knowledge of the target system (Papernot et al., 2016):

- **White-box settings:** the adversary has full access to the model architecture, parameters, and gradients, enabling highly effective gradient-based optimisation methods.
- **Grey-box settings:** partial information may be available (e.g., the architecture is known but weights are hidden), requiring approximate or transfer-based approaches.
- **Black-box settings:** the adversary has no access to the internal details of the model and can only query the model to observe outputs, motivating query-efficient or transferability-based attacks.

Among these settings, this paper focuses on the *white-box* setting, as it represents the worst-case scenario for model vulnerability, offering a rigorous basis for analysing robustness before extending consideration to more restrictive adversarial assumptions.

2.1.3. Adversary's constraints

A further dimension of adversarial taxonomies concerns the *constraints* imposed on perturbations. These constraints determine the

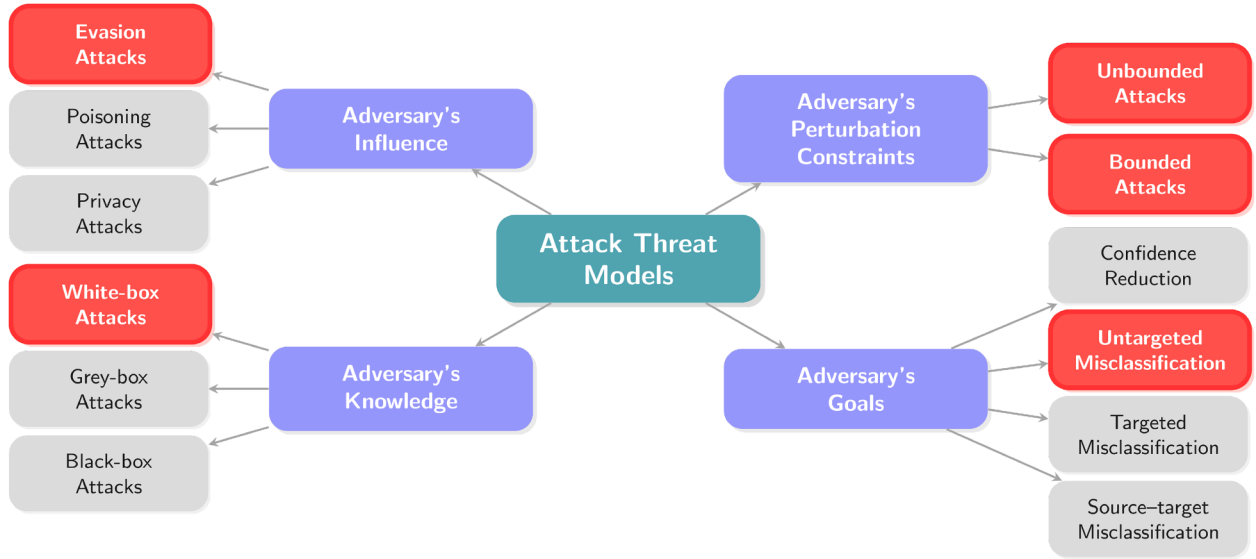


Fig. 1. A taxonomy of adversarial attack threat models. The taxonomy spans four primary dimensions: *Adversary's Influence*, *Adversary's Knowledge*, *Adversary's Perturbation Constraints* and *Adversary's Goals*. Highlighted categories indicate the specific the research scope in this benchmark.

extent to which inputs may be altered while still being considered valid, thereby balancing the adversary's power against the requirement that perturbed instances remain realistic and meaningful. Two broad categories are typically recognised:

Unbounded attacks (Carlini & Wagner, 2017; Moosavi-Dezfooli et al., 2016). the adversary faces no explicit restriction on the magnitude of perturbations. The objective is to find the smallest possible modification δ that changes the classifier's prediction.

Definition 4 (Unbounded adversarial attack). Let $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ be an ML classifier, and $(x, y) \in \mathcal{X} \times \mathcal{Y}$ a correctly classified instance. An *unbounded adversarial attack* seeks the adversarial example $\tilde{x} = x + \delta$ with the smallest possible perturbation δ that changes the classifier's decision, i.e.,

$$\min_{\delta} \|\delta\| \quad \text{subject to } f_\theta(x + \delta) \neq y.$$

2.1.3.1. Bounded attacks (Goodfellow et al., 2015; Madry et al., 2018). the adversary is restricted by a perturbation budget, typically expressed as $\|\delta\| \leq \epsilon$. Within this budget, the adversary seeks to maximise prediction error while keeping perturbed inputs close to the original.

Definition 5 (Bounded adversarial attack). Let $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ be an ML classifier, and $\Delta(x) = \{x' \in \mathcal{X} : \|x' - x\|_p \leq \epsilon\}$ a norm-bounded perturbation set. A *bounded adversarial attack* restricts the adversary to perturbations within the budget ϵ . Formally, it searches for an adversarial example $\tilde{x} \in \Delta(x)$ such that:

$$f_\theta(\tilde{x}) \neq y,$$

with the perturbation bounded by $\|\tilde{x} - x\|_p \leq \epsilon$.

2.1.4. Adversary's goals

The final dimension of the taxonomy concerns the *goals* of the adversary (Assion et al., 2019; Papernot et al., 2016), which specify the intended effect of perturbations on model predictions. These goals determine the success criterion of an attack and influence both the choice of optimisation objective and the evaluation metrics. Four canonical categories are widely recognised:

- **Confidence reduction:** the adversary reduces the model's prediction confidence without necessarily changing the predicted class.

- **Untargeted misclassification:** the adversary seeks to induce any incorrect prediction, without control over which incorrect class is produced.
- **Targeted misclassification:** the adversary forces the model to predict a specific incorrect class.
- **Source-target misclassification:** the adversary specifies both the source class and the desired incorrect target class.

Among these objectives, this benchmark concentrates on *untargeted misclassification*. This choice reflects the practical relevance of arbitrary misclassification in tabular decision systems, where any incorrect prediction can undermine reliability, and the computational efficiency of untargeted attacks, which avoid the additional complexity of enforcing specific outcomes.

Taken together, these four dimensions provide a systematic framework for characterising adversarial attacks in Fig. 1. In this benchmark, we adopt the *evasion* setting under a *white-box* threat model, with perturbations considered in both *bounded* and *unbounded* forms to capture practical as well as theoretical perspectives, and adversarial goals defined in terms of *untargeted misclassification*. This combination represents the most rigorous and practically relevant formulation of adversarial robustness in tabular data.

2.2. White-box adversarial attacks

Most adversarial attack algorithms were originally developed and studied in the context of Computer Vision (CV), where inputs are represented as high-dimensional images and perturbations can be visualised directly. This domain has served as the primary testbed for formulating adversarial objectives, designing optimisation procedures, and benchmarking robustness. Although the focus of this benchmark is on tabular data, understanding these canonical attack methods is essential, as many tabular attack strategies are adapted or inspired from techniques first introduced in the CV literature.

2.2.1. Fast gradient sign method

The fast gradient sign method (FGSM) attack (Goodfellow et al., 2015) works by calculating the gradient of the neural network with respect to the input data and using the sign of this gradient to determine the direction of the perturbation.

Definition 6 (FGSM). Given a classifier $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ with loss function ℓ , the FGSM adversarial example for input \mathbf{x} with label y is

$$\tilde{\mathbf{x}} = \mathbf{x} + \epsilon \operatorname{sign}(\nabla_{\mathbf{x}} \ell(f_\theta(\mathbf{x}), y)),$$

where $\epsilon > 0$ controls the perturbation magnitude and $\operatorname{sign}(\cdot)$ is applied elementwise.

It is equivalent to taking a single ℓ_∞ -bounded step of size ϵ in the steepest ascent direction. Its single-step nature makes it weaker than iterative attacks and prone to failure against robust models. Moreover, the uniform ℓ_∞ perturbation assumption does not align well with the heterogeneous constraints of tabular data.

2.2.2. Iterative method

The *Basic Iterative Method (BIM)* (Kurakin et al., 2017) extends FGSM by applying it multiple times with a smaller step size, iteratively updating the adversarial example. Projected Gradient Descent (Madry et al., 2018) generalises this idea by introducing random starts and projecting perturbed inputs back into the allowed perturbation set at each iteration, making it one of the strongest first-order adversaries.

Definition 7 (BIM and PGD). Given a classifier $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$, loss function ℓ , and perturbation budget ϵ , iterative adversarial examples are generated as

$$\mathbf{x}^{(t+1)} = \operatorname{Proj}_{\mathbf{x}, \epsilon}(\mathbf{x}^{(t)} + \alpha \operatorname{sign}(\nabla_{\mathbf{x}} \ell(f_\theta(\mathbf{x}^{(t)}), y))),$$

where α is the step size and $\operatorname{Proj}_{\mathbf{x}, \epsilon}(\cdot)$ ensures that $\mathbf{x}^{(t)}$ remains within the ℓ_∞ ball of radius ϵ around \mathbf{x} . The final adversarial example is denoted $\tilde{\mathbf{x}} = \mathbf{x}^{(T)}$.

- **BIM:** uses deterministic initialisation $\mathbf{x}^{(0)} = \mathbf{x}$, with projection implemented as clipping.
- **PGD:** generalises BIM by starting from a random perturbation $\mathbf{x}^{(0)} \in B_\infty(\mathbf{x}, \epsilon)$ and projecting back at every step.

2.2.3. Carlini and Wagner attack

The Carlini and Wagner attack (Carlini & Wagner, 2017), also known as the C&W attack, is a family of optimisation-based methods that aims to find the minimum perturbation that can cause a misclassification in a targeted deep neural network. Unlike gradient-sign methods, the C&W attack formulates the problem as a constrained optimisation under different ℓ_p norms, yielding adversarial examples with minimal perturbation.

Definition 8 (C&W attack). Given an input \mathbf{x} with label y and classifier f_θ , the C&W attack seeks

$$\min_{\delta} \|\delta\|_p + c \cdot g(\mathbf{x} + \delta),$$

where δ is the perturbation, $c > 0$ balances perturbation and misclassification, and $g(\cdot)$ is a surrogate loss term that enforces adversarial misclassification. A common choice is the logit-based hinge:

$$g(\mathbf{x}') = \max(f_\theta(\mathbf{x}')_y - \max_{i \neq y} f_\theta(\mathbf{x}')_i, -\kappa),$$

where $f_\theta(\mathbf{x}')_i$ is the logit score for class i and $\kappa \geq 0$ is a confidence parameter controlling attack strength. The adversarial example is $\tilde{\mathbf{x}} = \mathbf{x} + \delta$.

2.2.4. DeepFool

DeepFool (Moosavi-Dezfooli et al., 2016) formulates adversarial perturbations as the minimal change required to cross the classifier's decision boundary. Unlike FGSM or PGD, which use fixed-magnitude gradient steps, DeepFool approximates the boundary locally as linear and iteratively projects the input toward the nearest boundary until misclassification occurs.

Definition 9 (DeepFool). Formulated under the ℓ_2 norm, DeepFool assumes a locally linear decision boundary $\mathcal{F} = \{\mathbf{x} : \mathbf{w}^\top \mathbf{x} + b = 0\}$ for some weight vector \mathbf{w} and bias b . It computes the minimal perturbation

$$\mathbf{r}_*(\mathbf{x}) = \arg \min_{\delta} \|\delta\|_2 \quad \text{s.t.} \quad f(\mathbf{x} + \delta) \neq f(\mathbf{x}),$$

which corresponds to the orthogonal projection of \mathbf{x} onto the separating hyperplane \mathcal{F} . The adversarial example is then updated iteratively as

$$\mathbf{x}^{(t+1)} = \mathbf{x}^{(t)} + \delta^{(t)},$$

where $\delta^{(t)}$ is the projection step, until misclassification is achieved. The final adversarial example is denoted $\tilde{\mathbf{x}} = \mathbf{x}^{(T)}$.

In the linear binary case, this reduces to the closed-form solution

$$\mathbf{r}_*(\mathbf{x}) = -\frac{f(\mathbf{x})}{\|\mathbf{w}\|_2^2} \mathbf{w}, \quad \mathbf{x}^{adv} = \mathbf{x} + \mathbf{r}_*(\mathbf{x}).$$

2.3. Adversarial attacks on tabular data

Adversarial attacks on tabular data aim to manipulate inputs such that ML models produce incorrect predictions while keeping modifications imperceptible (He et al., 2025). Within this context, adversarial machine learning on tabular data faces two central challenges, including *effectiveness*, i.e., the ability to reliably degrade model performance through computationally efficient attacks; and *imperceptibility*, i.e., ensuring that adversarial examples respect domain constraints, preserve feature interdependencies, and avoid unrealistic or easily detectable modifications.

White-box gradient-based effective attacks. A canonical baseline is the PGD attack (Madry et al., 2018), which iteratively updates adversarial examples by ascending the gradient of the loss while constraining perturbations within a bounded ℓ_p ball. In the context of tabular data, PGD has been widely adopted as a benchmark due to its simplicity and strong empirical performance. Building on this foundation, several variants introduce domain-specific constraints to account for the heterogeneous and semantically structured nature of tabular features. Early work by Ballet et al. (2019) demonstrate the adaptation of canonical methods such as FGSM and PGD to tabular data, showing that even simple gradient-based attacks can mislead models when applied to structured inputs. More recently, Constrained Adaptive PGD (CAPGD) (Simonetto et al., 2022, 2024a,b) refines this approach by adapting step sizes and directions to respect feasibility conditions such as valid categorical values and bounded numerical ranges. These methods illustrate how canonical gradient-based attacks can be adapted for the particular challenges of tabular domains.

Imperceptibility and constraint-aware attacks. A key challenge in tabular attack is imperceptibility: adversarial examples should not only cause misclassification but also remain realistic and semantically valid. Low-ProFool (Ballet et al., 2019) introduces feature-importance weighting, biasing perturbations toward low-salience features to improve imperceptibility. Similarly, The Feature Importance Guided Attack (FIGA) (Gressel et al., 2021) perturbs the most influential features identified by importance scores, maximising impact with minimal changes. Subsequent research refined constraint handling: Mathov et al. (2022) enforce consistency across feature interdependencies, while Chernikova and Oprea (2022) ensures feasibility in security-critical datasets. Cost-aware frameworks (Kireev et al., 2023) further restrict perturbations according to domain-specific budgetary limits, such as financial cost in fraud detection. More recently, CAPGD (Simonetto et al., 2024a) unified gradient-based and search-based optimisation to efficiently generate valid adversarial examples, accompanied by a benchmarking framework across diverse tabular datasets and models (Simonetto et al., 2024b).

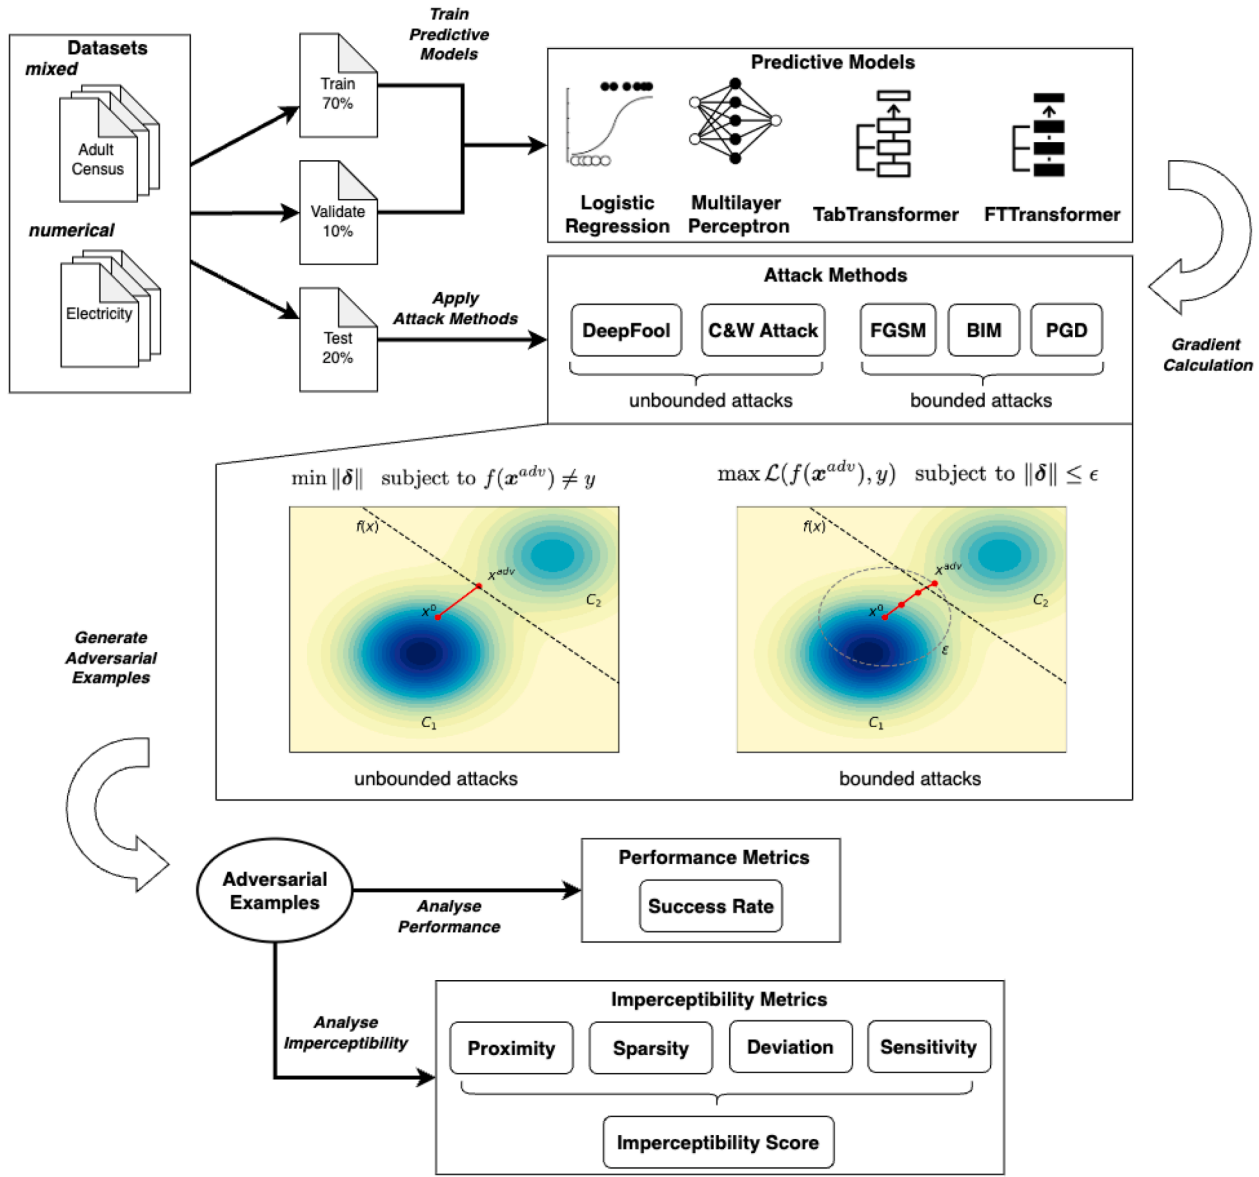


Fig. 2. Roadmap of the proposed evaluation framework for benchmarking adversarial attacks on tabular data. The pipeline extends standard machine learning workflows by adding an adversarial attack stage, linking dataset preparation, model training, attack generation, and evaluation to provide a comprehensive measure of effectiveness and imperceptibility.

3. Methodology

Integrating adversarial attacks as an additional phase in the standard machine learning pipeline enables a systematic assessment of model robustness against adversarial perturbations. Inspired by the established machine learning benchmark guidelines (Thiyagalingam et al., 2022), we design our benchmark to evaluate adversarial attacks on tabular data. It comprises three sequential stages: (i) data preparation and model training, where clean tabular datasets are pre-processed and predictive models are optimised to establish baseline performance; (ii) adversarial example generation, in which attack algorithms craft perturbed inputs under a white-box assumption with full access to model parameters and gradients; and (iii) robustness evaluation, where the perturbed samples are assessed using both effectiveness (attack success rate) and imperceptibility metrics. This design enables consistent, reproducible comparison across models and attacks while isolating the influence of perturbations from data preprocessing or training variability. Fig. 2 summarises the overall benchmark pipeline and its component interactions.

3.1. Datasets

The selection of datasets for the benchmark follows several criteria to ensure both suitability and effectiveness.

1. Each dataset must support classification tasks and contain at least two classes, enabling the evaluation of binary classification performance.
2. To maintain computational feasibility and avoid excessive dimensionality, the total number of features, including those generated by one-hot encoding of categorical attributes, is restricted to no more than 200.
3. Datasets should be representative of real-world scenarios, with sufficient diversity of instances to meaningfully assess model robustness and generalisation.

In addition to our selection criteria, we also refer to datasets that are commonly adopted in the deep learning benchmark for tabular data (Gorishniy et al., 2021), ensuring consistency with prior work and

Table 3

Data profiles of the 11 datasets used in the benchmark, including the total number of instances (n_{total}), instances for training (n_{train}), validation ($n_{validate}$), and adversarial perturbation (n_{test}), as well as the number of numerical features (d_{num}), categorical features (d_{cat}), one-hot encoded (categorical) features ($d_{encoded}$), and the total number of features (d_{total}).

Dataset	Domain	n_{total}	n_{train}	n_{val}	n_{test}	d_{num}	d_{cat}	$d_{encoded}$	d
Adult	Finance	32,561	22,792	3256	6513	6	8	99	105
Electricity	Energy	45,312	31,717	4532	9063	7	1	7	14
COMPAS	Justice	16,644	11,650	1665	3329	8	8	50	58
Higgs	Physics	1,000,000	700,000	100,000	200,000	28	0	0	28
house_16H	Finance	22,784	15,948	2279	4557	16	0	0	16
jm1	Software	10,885	7619	1089	2177	21	0	0	21
BreastCancer	Healthcare	569	398	57	114	30	0	0	30
WineQuality-White	Chemistry	4898	3428	490	980	11	0	0	11
WineQuality-Red	Chemistry	1599	1119	160	320	11	0	0	11
phoneme	Speech	5404	3782	541	1081	5	0	0	5
MiniBooNE	Physics	130,064	91,044	13,007	26,013	50	0	0	50

Table 4

Profiles of the five adversarial attack methods and the Gaussian noise baseline. Profiles of the five adversarial attack methods and the Gaussian noise baseline. Type abbreviations: G = Gradient-based, O = Optimisation-based, R = Random baseline; B = Bounded, U = Unbounded.

Method	Type	Optimisation Strategy	Constraint
FGSM	G, B	One-step gradient sign	ℓ_∞ ball, ϵ budget
BIM	G, B	Iterative gradient sign	ℓ_∞ ball, ϵ budget
PGD	G, B	Iterative with projection	ℓ_∞ ball, ϵ budget
DeepFool	O, U	Iterative boundary projection	Closest decision boundary (ℓ_2)
C&W	O, U	Gradient descent with custom loss	Minimal ℓ_2 perturbation
Gaussian Noise	R, B	Random sampling from $\mathcal{N}(0, \sigma^2)$	ℓ_∞ norm, ϵ budget

comparability across studies. Table 3 summarises the profiles of the 11 datasets¹ used, including their size, feature composition, and data splits for training, validation, and adversarial evaluation.

3.2. Adversarial attacks

Following the scope established in Section 2.1, all adversarial examples in this benchmark are generated using white-box, untargeted attacks with full gradient access to the predictive models. Although adversarial methods for tabular data have begun to emerge, no systematic benchmark currently exists. Consequently, our selection process draws on established benchmarks in computer vision, with careful screening to ensure that the chosen methods can be extended to tabular settings. We prioritise attack algorithms that:

1. The selected attack methods should be applicable to tabular data.
2. The selected attack methods should be designed for white-box attack.

Based on these considerations, five widely used attack methods were selected (Table 4). These span different levels of complexity and optimisation strategy, and all have demonstrated strong effectiveness in prior domains:

- **Bounded gradient-based attacks:** FGSM (Goodfellow et al., 2015), BIM (Kurakin et al., 2017) and PGD (Madry et al., 2018).
- **Unbounded optimisation-based attacks:** DeepFool (Moosavi-Dezfooli et al., 2016) and C&W (Carlini & Wagner, 2017).

In addition, Gaussian noise is included as a baseline. This serves as a reference to compare deliberate adversarial perturbations against random noise injection.

¹ WineQuality (White) and WineQuality (Red) originate from the same dataset but are treated as two separate datasets in this benchmark.

3.3. Predictive models

The choice of predictive models is guided by three criteria: *diversity* (covering classical and modern architectures), *interpretability* (balancing transparency with complexity), and *performance* (ensuring competitive accuracy on tabular tasks). Based on these principles, we evaluate four representative models:

1. **Logistic regression (LR):** A simple, interpretable linear baseline that establishes performance lower bounds and vulnerability benchmarks.
2. **Multilayer perceptron (MLP):** A foundational neural network adept at capturing nonlinear patterns, offering a mid-complexity comparison point.
3. **TabTransformer (Huang et al., 2020):** An attention-based model that processes tabular features via transformer layers, leveraging contextual relationships among features.
4. **FTTransformer (Gorishniy et al., 2021):** A transformer-based architecture that tokenizes numerical and categorical features, enabling unified processing through self-attention mechanisms.

3.4. Evaluation metrics

These metrics are grouped into two complementary categories: *effectiveness*, which measures whether adversarial examples succeed in misleading the predictive model, and *imperceptibility*, which evaluates how closely the perturbed instances adhere to the properties of imperceptible tabular data.

3.4.1. Effectiveness

The primary metric of attack effectiveness is the attack success rate (ASR), which is defined as the proportion of adversarial examples that cause misclassification relative to the total number of perturbed instances. This metric provides a direct indication of how vulnerable the

predictive model is to each attack.

$$\text{ASR}(f_\theta) = \frac{1}{n} \sum_{i=1}^n \mathbb{1} \left\{ \exists \mathbf{x}' \in \Delta(\mathbf{x}_i) : f_\theta(\mathbf{x}') \neq y_i \right\}$$

3.4.2. Imperceptibility

Following our prior work (He et al., 2025), we employ four quantitative imperceptibility metrics including proximity, sparsity, sensitivity and deviation.

Definition 10 (Proximity). For an input $\mathbf{x} \in \mathcal{X}$ and its adversarial example $\tilde{\mathbf{x}} \in \Delta(\mathbf{x})$, the perturbation size can be quantified using ℓ_p norms:

$$D_p(\tilde{\mathbf{x}}, \mathbf{x}) = \|\tilde{\mathbf{x}} - \mathbf{x}\|_p = \begin{cases} \left(\sum_{j=1}^d |\tilde{x}_j - x_j|^p \right)^{1/p}, & p \in \{1, 2\}, \\ \max_{j \in \{1, \dots, d\}} |\tilde{x}_j - x_j|, & p \rightarrow \infty. \end{cases}$$

Here, d denotes the feature dimension. Over a dataset $D = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$, the average ℓ_p proximity is

$$\overline{D}_p(D) = \frac{1}{n} \sum_{i=1}^n D_p(\tilde{\mathbf{x}}_i, \mathbf{x}_i),$$

which measures the mean perturbation size across all adversarial examples. In this benchmark, we utilise the ℓ_2 norm (Euclidean distance) for evaluating proximity.

Definition 11 (Sparsity). For an input $\mathbf{x} \in \mathcal{X}$ and its adversarial example $\tilde{\mathbf{x}} \in \Delta(\mathbf{x})$, the *sparsity* of the perturbation is the number of features whose values differ between \mathbf{x} and $\tilde{\mathbf{x}}$. It is measured by the ℓ_0 pseudo-norm:

$$\text{Spa}(\tilde{\mathbf{x}}, \mathbf{x}) = \|\tilde{\mathbf{x}} - \mathbf{x}\|_0 = \sum_{j=1}^d \mathbb{1}\{\tilde{x}_j \neq x_j\},$$

where d denotes the feature dimension. The corresponding *fractional sparsity* measures the fraction of features changed per instance:

$$\text{SpaFrac}(\tilde{\mathbf{x}}, \mathbf{x}) = \frac{1}{d} \text{Spa}(\tilde{\mathbf{x}}, \mathbf{x}).$$

Over a dataset $D = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$, the *sparsity rate* is defined as the average fractional sparsity of adversarial examples:

$$\text{SpaR}(D) = \frac{1}{n} \sum_{i=1}^n \text{SpaFrac}(\tilde{\mathbf{x}}_i, \mathbf{x}_i) = \frac{1}{nd} \sum_{i=1}^n \sum_{j=1}^d \mathbb{1}\{\tilde{x}_{i,j} \neq x_{i,j}\}.$$

Definition 12 (Deviation). For an input $\mathbf{x} \in \mathcal{X}$, its adversarial example $\tilde{\mathbf{x}} \in \Delta(\mathbf{x})$, and the covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$ of the data distribution \mathcal{P} with mean $\boldsymbol{\mu}$, the *deviation* of $\tilde{\mathbf{x}}$ from \mathcal{P} is measured by the *Mahalanobis Distance (MD)*:

$$\text{MD}(\tilde{\mathbf{x}}, \boldsymbol{\mu}) = \sqrt{(\tilde{\mathbf{x}} - \boldsymbol{\mu})^\top \Sigma^{-1} (\tilde{\mathbf{x}} - \boldsymbol{\mu})}.$$

An adversarial example is considered an *outlier* if

$$\text{MD}^2(\tilde{\mathbf{x}}, \boldsymbol{\mu}) > \chi_{\alpha, d}^2,$$

where $\chi_{\alpha, d}^2$ denotes the critical value of the chi-squared distribution with d degrees of freedom at significance level α . Over a dataset $D = \{(\mathbf{x}_i, y_i)\}_{i=1}^n \sim \mathcal{P}$, the *outlier rate* is defined as

$$\text{OR}(D) = \frac{1}{n} \sum_{i=1}^n \mathbb{1} \left(\text{MD}^2(\tilde{\mathbf{x}}_i, \mathbf{x}_i) > \chi_{\alpha, d}^2 \right),$$

which measures the proportion of adversarial examples flagged as outliers.

Definition 13 (Sensitivity). Over a dataset $D = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$, where each feature vector $\mathbf{x}_i = (x_{i,1}, \dots, x_{i,d})$ has d features. Among them, let $d_{\text{num}} \leq d$ denote the number of numerical features. For the j -th numerical feature, let σ_j denote its standard deviation across the dataset D .

Given an input $\mathbf{x} \in \mathcal{X}$ and its adversarial example $\tilde{\mathbf{x}} \in \Delta(\mathbf{x})$, the *sensitivity* of the perturbation is defined as the variance-normalised ℓ_1 distance:

$$\text{SEN}(\tilde{\mathbf{x}}, \mathbf{x}) = \sum_{j=1}^{d_{\text{num}}} \frac{|\tilde{x}_j - x_j|}{\sigma_j}.$$

The *average sensitivity* over the dataset D is defined as

$$\overline{\text{SEN}}(D) = \frac{1}{n} \sum_{i=1}^n \text{SEN}(\tilde{\mathbf{x}}_i, \mathbf{x}_i) = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^{d_{\text{num}}} \frac{|\tilde{x}_{i,j} - x_{i,j}|}{\sigma_j},$$

which measures the mean variance-normalised perturbation size across all adversarial examples.

Definition 14 (Imperceptibility Score). The overall *imperceptibility* of an adversarial attack is quantified by a composite metric that integrates four complementary measures: average ℓ_2 proximity (\overline{D}_2), sparsity rate (SpaR), outlier rate (OR), and average sensitivity (SEN). Each captures a distinct aspect of how perturbations alter data realism and is assigned equal importance ($w_i = 1$). To ensure comparability, all metrics are normalised to the range $[0, 1]$: sparsity and outlier rates are inherently bounded, while \overline{D}_2 and SEN are rescaled via min-max normalisation. The resulting *Imperceptibility Score (IS)* is defined as the weighted harmonic mean:

$$\text{IS} = \frac{4}{\frac{1}{\overline{D}_2} + \frac{1}{\text{SpaR}} + \frac{1}{\text{OR}} + \frac{1}{\text{SEN}}}.$$

3.5. Implementation details and reproducibility

To ensure transparency and reproducibility, all experiments in this benchmark were implemented in Python 3.12 using PyTorch 2.4 and executed on a workstation equipped with an NVIDIA A100 GPU (80 GB VRAM), AMD EPYC 7543 CPU, and 256 GB RAM. The random seed was fixed to 42 across all stages, including dataset partitioning, model training, and adversarial generation, to guarantee deterministic behaviour. All code, pretrained models, and configuration files are publicly available in the companion repository², ensuring full reproducibility of the reported results.

Data partitioning and preprocessing. Each dataset was partitioned using stratified sampling to preserve class distributions, allocating 70% of the data for training, 10% for validation, and 20% for testing and adversarial evaluation. Constant or duplicate features were removed, missing values were imputed using the median for numerical features and the mode for categorical features, and categorical variables were one-hot encoded. Numerical features were normalised to the $[0, 1]$ range using min-max scaling. These consistent procedures maintain comparability across datasets while preserving their intrinsic statistical characteristics.

Training protocol. All predictive models were trained for 20 epochs with a batch size of 512, using the Adam optimiser with a learning rate of 1×10^{-3} and cross-entropy loss. Dropout ($p = .2$) was applied to the MLP and transformer-based models for regularisation. Each transformer used 6 layers and 8 attention heads, with hidden-layer dimensionality fixed at 64 units to maintain comparable model capacity across architectures. The same hyperparameters were used for every dataset to facilitate consistent benchmarking rather than per-dataset optimisation. Table 5 summarises the main configurations.

Attack configuration. All adversarial examples were generated under the white-box assumption with full gradient access. Five attack methods were implemented: FGSM, BIM, PGD, DeepFool, and C&W, along with Gaussian noise as a random baseline. For FGSM, BIM, and PGD, perturbation budgets $\epsilon \in \{0.01, 0.03, 0.05, 0.1, 0.3, 0.5, 1.0\}$ were applied, with

² <https://github.com/ZipengHe/TabAttackBench>

Table 5

Key hyperparameter configurations for the four predictive models. The ‘-’ denotes that the corresponding hyperparameter does not apply in this model.

Parameter	LR	MLP	TabTransformer	FTTransformer
Training epochs			20	
Batch size			512	
Learning rate			1×10^{-3}	
Optimizer			Adam	
Loss function			Cross-entropy	
Hidden layers	-	64-32	-	-
Activation	-	ReLU	ReLU	ReLU
Dropout rate	-	0.2	0.2	0.2
Transformer depth	-	-	6	6
Attention heads	-	-	8	8
Embedding dimension	-	-	64	64

Table 6

Accuracy of four predictive models across 11 datasets. Deep learning models generally outperform LR and all models exceed 63% accuracy.

Dataset	LR	MLP	TabTransformer	FTTransformer
<i>Adult</i>	0.834	0.8337	0.8328	0.799
<i>BreastCancer</i>	0.9386	0.9737	0.9035	0.9737
<i>Compas</i>	0.6654	0.6738	0.7053	0.6858
<i>Electricity</i>	0.6607	0.7635	0.762	0.7712
<i>Higgs</i>	0.6366	0.7234	0.6951	0.7296
<i>MiniBooNE</i>	0.7724	0.8372	0.8497	0.8402
<i>WineQuality-Red</i>	0.7219	0.7344	0.7281	0.7344
<i>WineQuality-White</i>	0.6745	0.7469	0.7316	0.752
<i>house_16H</i>	0.7029	0.8578	0.8251	0.8466
<i>jm1</i>	0.8075	0.8098	0.8066	0.8107
<i>phoneme</i>	0.7095	0.7872	0.7882	0.8002

BIM and PGD using $T=10$ iterations and a relative step size $\alpha = \epsilon/T$. DeepFool employed a maximum of 50 iterations and an overshoot factor of 0.02, while C&W used 10 binary search steps, an initial constant $c_0=0.001$, and confidence $\kappa=0$. Each attack was run on all trained models, and evaluation metrics were averaged over five independent runs.

4. Evaluation

Building on the benchmark design and implementation described in Section 3, this section presents the empirical evaluation of the proposed framework.

4.1. Task 1: effectiveness evaluation

We begin by evaluating model accuracy to establish a baseline prior to adversarial testing. If accuracy falls below 60%, a model may be trivially deceived, rendering robustness evaluation less meaningful. Table 6 reports the accuracy of the four predictive models across the eleven datasets. The three deep learning models generally outperform LR, although LR achieves comparable results on *Adult*, *jm1*, and *COMPAS*. Overall, All models exceed 63% accuracy, providing a reliable basis for subsequent adversarial evaluation.

4.1.1. Effectiveness: mixed datasets

Fig. 3 shows that the *Electricity* dataset demonstrates unique characteristics compared to other mixed datasets. Across all model architectures, every attack achieves consistently high success rates, all approaching 100% success on LR, MLP, and TabTransformer. This indicates that the feature distribution or decision boundaries in *Electricity* are especially susceptible to adversarial perturbations. Unlike *Adult*

and *COMPAS*, where attack types diverge in performance, *Electricity* exhibits uniform vulnerability.

For *Adult* and *COMPAS*, clearer differences emerge. The ℓ_∞ -based bounded attacks (FGSM, BIM, PGD) consistently outperform ℓ_2 -based unbounded attacks (C&W attack and DeepFool) on these datasets. This suggests that bounded perturbations under the ℓ_∞ ball are particularly effective for mixed-type tabular data, as they can alter critical features directly without being constrained by overall perturbation magnitude.

Architecturally, Transformer-based models exhibit greater robustness than LR and MLP. Both TabTransformer and FTTransformer require larger ϵ values before achieving comparable attack success rate (ASR), especially on *Adult* dataset. This points to attention mechanisms and deeper architectures providing some inherent robustness to adversarial manipulation in mixed tabular data.

4.1.2. Effectiveness: numerical datasets

Figs. 4-6 summarise results on the eight numerical datasets, which exhibit more consistent patterns than mixed datasets, while still showing dataset-specific behaviours.

On *Higgs* and *house_16H* (Fig. 4), all three ℓ_∞ -based attacks (FGSM, BIM, PGD) produce nearly identical success curves across models. The *BreastCancer* dataset (Fig. 5) provides the most diverse response to different attack methods. Here, PGD consistently outperforms other attacks against transformer-based models. FTTransformer under BIM shows a non-monotonic trend, with attack success occasionally decreasing as ϵ increases. The *WineQuality* datasets (*Red* and *White*; Figs. 5 and 6) show that ℓ_∞ -based attacks require substantially lower epsilon values to achieve high success rates compared to ℓ_2 -based approaches. DeepFool eventually achieves comparable success but requires much larger perturbation budgets, limiting its efficiency. The *phoneme* and *MiniBooNE* datasets (Fig. 6) further confirm the superiority of ℓ_∞ -based attacks, with all three methods (FGSM, PGD, BIM) demonstrating nearly identical performance trajectories and consistently outperform ℓ_2 -based approaches across models.

Across the numerical datasets, LR models are consistently the most vulnerable, often showing sharp threshold effects where ASR increases abruptly at particular ϵ values. This reflects the ease with which linear decision boundaries can be exploited by adversarial perturbations. In contrast, MLP and transformer models generally resist small perturbations and only degrade at larger ϵ values, though the exact patterns vary by dataset. The C&W attack exhibits the least consistent behaviour, sometimes barely outperforming random noise yet occasionally achieving competitive success on specific dataset-model combinations. Its strong dependence on data and architecture characteristics limits its reliability as a general-purpose attack for tabular domains.

4.2. Task 2: imperceptibility assessment

To evaluate imperceptibility in a fair and consistent manner, we first establish a systematic strategy for selecting attack budgets. For each model-dataset-attack combination, we identify the ϵ value at which the ASR curve first reaches a plateau, beyond which further increases yield little or no improvement. The specific attack budgets selected through this methodology are detailed in Appendix A. However, direct comparisons across different models and datasets would be misleading, since the optimal ϵ values vary significantly between these contexts. To address this, we identify the most frequently occurring ϵ value for each attack method across all experiments. These representative settings, listed in Table 7, are then adopted as our standardised benchmark parameters for subsequent analyses in Task 2 and Task 3. Full imperceptibility results can be found in Appendix B.

4.2.1. Task 2.1: sparsity

Our analysis of sparsity patterns reveals distinct behavioural characteristics among adversarial attack methods while highlighting the influ-

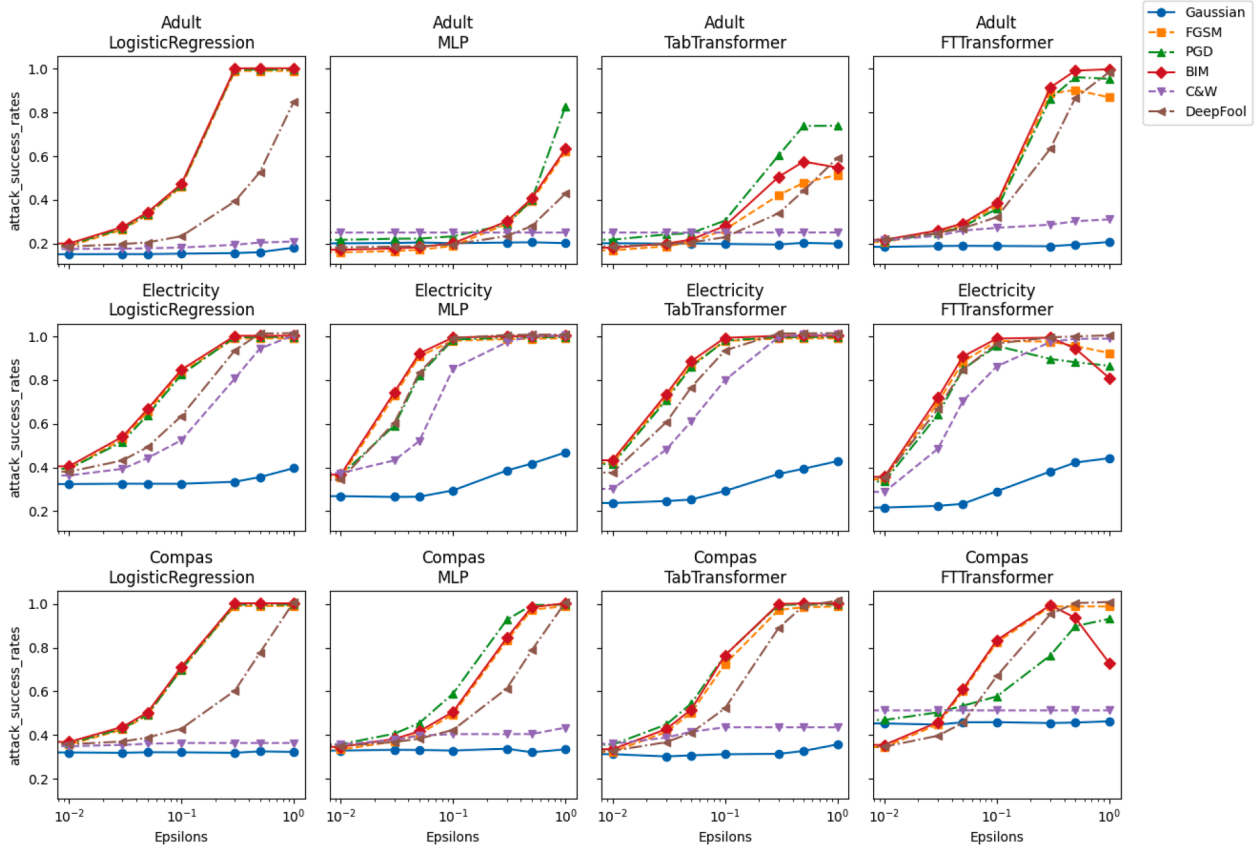


Fig. 3. Attack success rate (ASR) of evaluated attack methods on all three mixed datasets (*Adult*, *COMPAS*, *Electricity*) and four models. *Electricity* shows uniformly high vulnerability across all attacks and models, while *Adult* and *COMPAS* reveal clear differences between ℓ_∞ -based (FGSM, BIM, PGD) and ℓ_2 -based (DeepFool, C&W attack) methods. Transformer models generally require larger perturbations to reach comparable success rates, indicating greater robustness than LR and MLP.

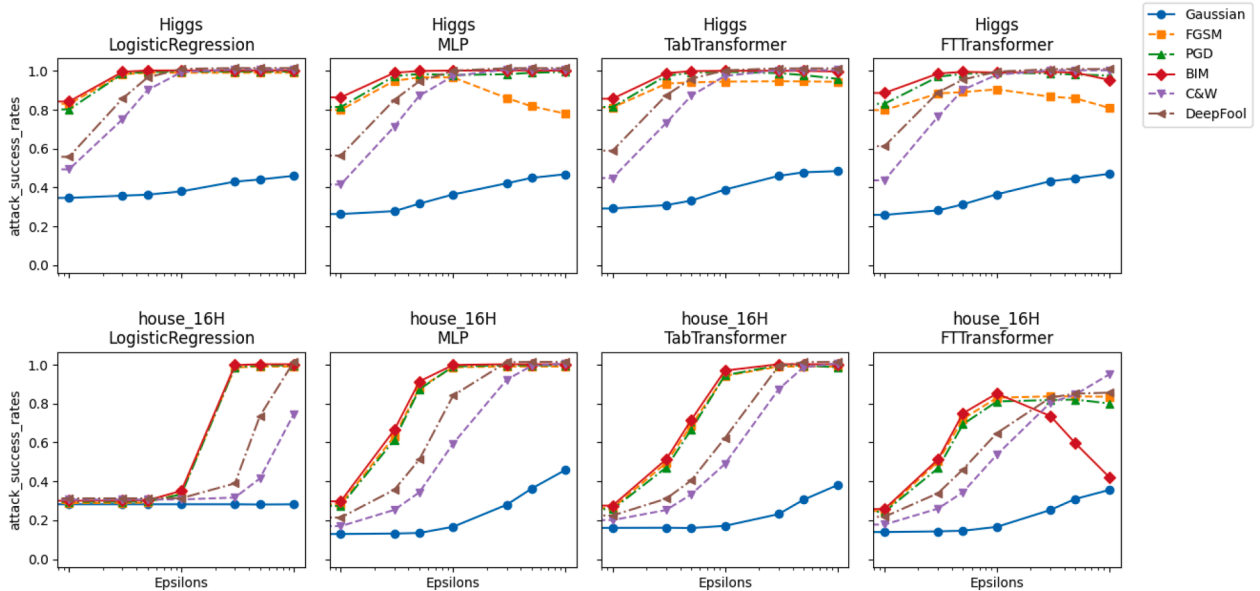


Fig. 4. Attack success rate (ASR) of evaluated attack methods on two (out of eight) numerical datasets and four ML models.

ence of dataset dimensionality and model architecture. The *sparsity rate* measure quantifies the proportion of features modified in adversarial examples, with higher values indicating more features being perturbed.

Sparsity: mixed datasets. Results for the mixed datasets (Figs. 7–9) reveal systematic differences between categorical and numerical features.

A consistent pattern emerges across most attacks: FGSM, BIM, C&W, and DeepFool strongly prioritise numerical features when targeting neural architectures (MLP, TabTransformer, FT-Transformer). In categorical-dominant datasets such as *Adult* (105 features, 99 categorical; Fig. 7) and *COMPAS* (58 features, 50 categorical; Fig. 7), these methods perturb numerical features heavily (72–99% sparsity) while

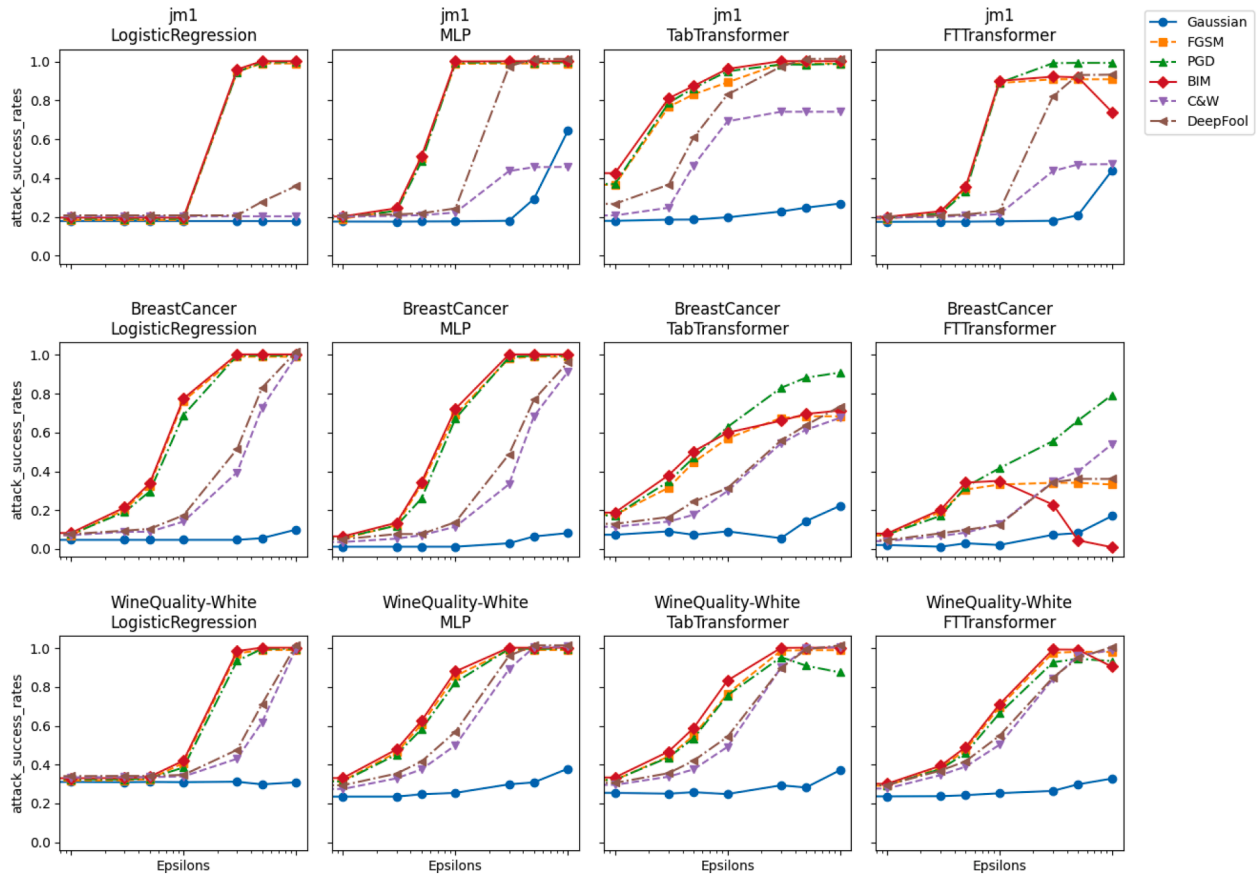


Fig. 5. (Cont.) Attack success rate (ASR) of evaluated attack methods on another three (out of eight) *numerical* datasets and four ML models.

Table 7

Representative ϵ values for each attack method, identified as the most frequent across all model-dataset combinations. These standardised settings balance effectiveness with comparability, and are used in the imperceptibility (Task 2) and trade-off (Task 3) analyses.

Attacks	Gaussian	FGSM	BIM	PGD	C&W	DeepFool
ϵ	1	0.3	0.3	0.3	1	1

leaving categorical ones nearly untouched (0–1%). This systematic selectivity indicates that these attacks inherently favour continuous variables, largely ignoring one-hot encoded categorical dimensions despite their prevalence.

Among the evaluated attacks, PGD stands out as the only method that consistently modifies categorical features, achieving around 50% sparsity across all models and datasets. Its projection mechanism appears uniquely suited to navigating the discrete one-hot space, distinguishing it from other ℓ_∞ -based bounded attacks.

Turning to model-specific effects, the behaviour of LR differs markedly from neural models. Under FGSM, PGD, and BIM, LR exhibits moderate categorical sparsity (40–52%) across all mixed datasets. This suggests that LR encodes information in a way that makes categorical features more vulnerable, an effect particularly pronounced in *COMPAS*, where categorical sparsity under FGSM and BIM reaches 52%, compared to 0% for neural networks.

The *Electricity* dataset (14 features: 7 categorical, 7 numerical; Fig. 9) illustrates attack behaviour in a balanced feature space. Even here, ℓ_∞ -based attacks retain their numerical preference: FGSM and BIM achieve 99% sparsity on numerical features while leaving categorical features

unaltered (0%) when attacking neural models. This persistence confirms the algorithmic bias of these methods toward continuous features. ℓ_2 -based attacks show mixed outcomes. C&W generally produces very low overall sparsity (0–2%) across mixed datasets but achieves moderate sparsity on numerical features in *Electricity* (64–69%). DeepFool shows an even stronger numerical preference, altering 40–74% of numerical features while leaving categorical features virtually untouched (0–0.4%). These patterns underscore the difficulty of ℓ_2 -based methods in handling one-hot encoded categorical variables.

These findings collectively demonstrate that with the exception of PGD, current adversarial attacks on tabular data overwhelmingly target numerical features while almost entirely neglecting categorical dimensions, regardless of their prevalence.

Sparsity: numerical datasets. For numerical datasets, as illustrated in Fig. 10, we observe distinct patterns of feature perturbation across different attack methods.

The three ℓ_∞ -based attacks (FGSM, PGD, and BIM) consistently modify nearly all available features, producing sparsity rates between 80–100% across models and datasets. This tendency holds regardless of feature dimensionality, from the low-dimensional *phoneme* dataset (5 features, Fig. 10g) to the high-dimensional *MiniBooNE* dataset (50 features, Fig. 10g). Among them, PGD is particularly aggressive, frequently driving sparsity close to 99–100%, especially on transformer-based models.

In contrast, the two ℓ_2 -based optimisation attacks show far more selective behaviour. C&W is the most variable: in some cases (e.g., LR on *jm1*, Fig. 10c) it modifies almost no features at all (0%), while in others it perturbs around half (50–74%). DeepFool generally occupies a middle ground, with sparsity rates between 17–80%, and its most selective behaviour is observed on the *Higgs* dataset (17–24%, Fig. 10a). Overall,

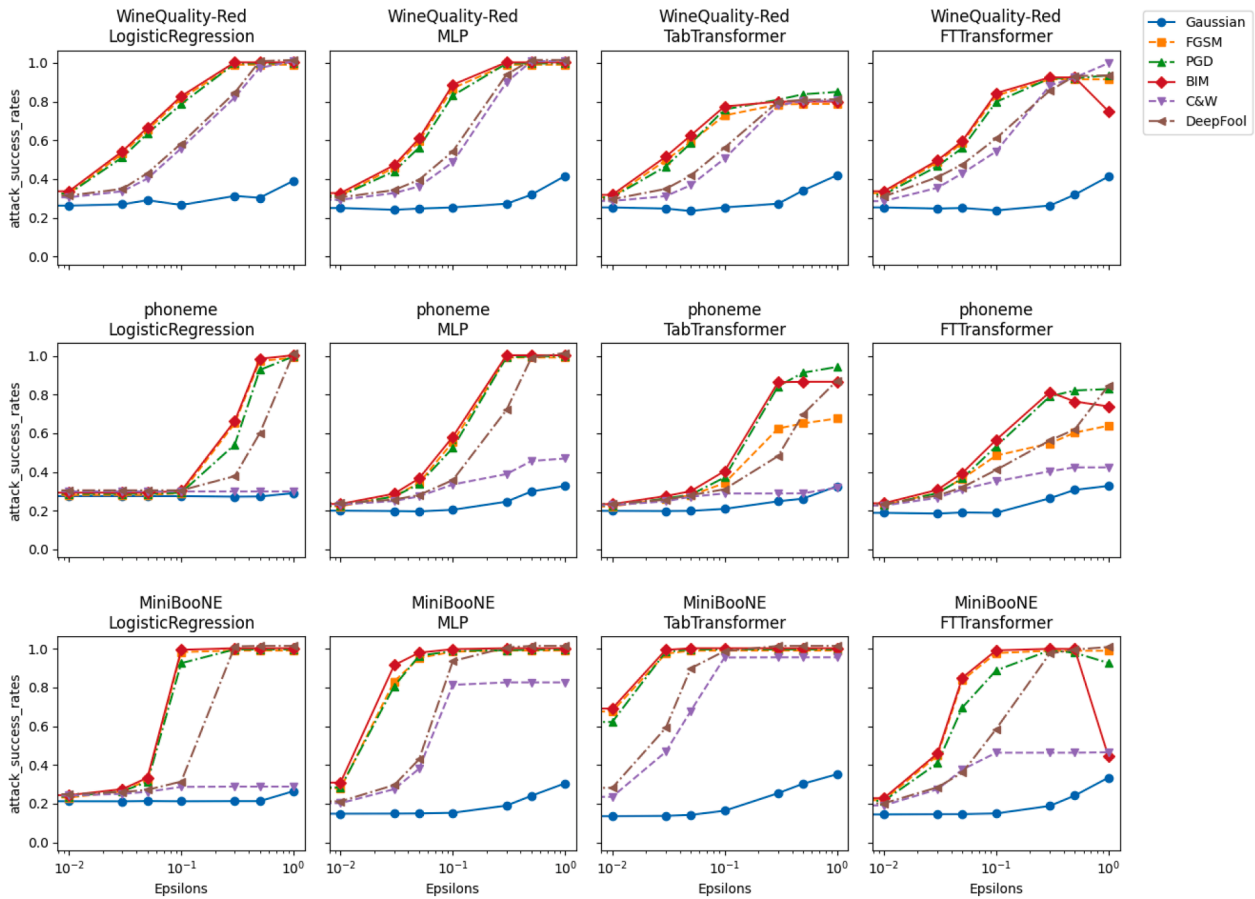


Fig. 6. (Cont.) Attack success rate (ASR) of evaluated attack methods on the remaining three (out of eight) numerical datasets and four ML models.

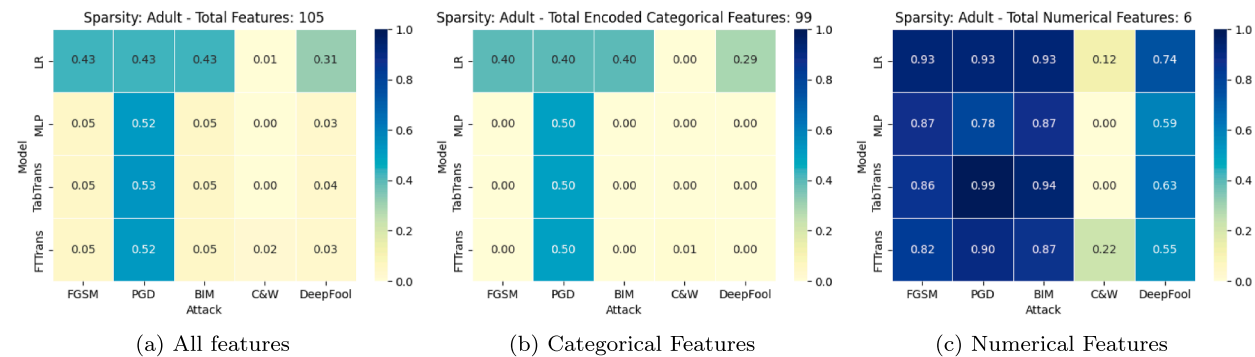


Fig. 7. Sparsity results (*sparsity rate*) of five evaluated attack methods and four models on the *Adult* dataset. Despite categorical dominance (99 of 105 features), most attacks primarily perturb numerical features, leaving categorical ones virtually untouched.

C&W appears highly sensitive to dataset–model combinations, aggressively pruning its perturbations to a small subset of influential features, while DeepFool is more stable across architectures but still adapts its sparsity according to dataset characteristics.

Feature dimensionality exerts a clear influence on ℓ_2 attacks. Both C&W and DeepFool reduce their sparsity on larger datasets, concentrating on fewer features as dimensionality grows. For example, on the 50-feature *MiniBooNE* dataset, C&W modifies only 5–28% of features, compared with 50–74% on small datasets such as *phoneme* (5 features). An important counterexample is the *BreastCancer* dataset (Fig. 10d), where all methods exhibit very high sparsity rates (> 80%), suggesting that every feature contributes strongly to classification and thus becomes a viable target.

Model architecture further shapes sparsity outcomes. LR exhibits the greatest variability, particularly under C&W and DeepFool, with sparsity rates ranging from negligible to substantial depending on the dataset. TabTransformer also shows notable fluctuations across attack types, whereas MLP yields more consistent sparsity patterns.

4.2.2. Task 2.2: proximity

Our proximity analysis measures how close adversarial examples remain to their original samples in the feature space using ℓ_2 distance metrics. The heatmaps in Figs. 11 and 12 reveal distinct patterns across attack types, model architectures, and datasets that provide important insights into the imperceptibility of different adversarial approaches.

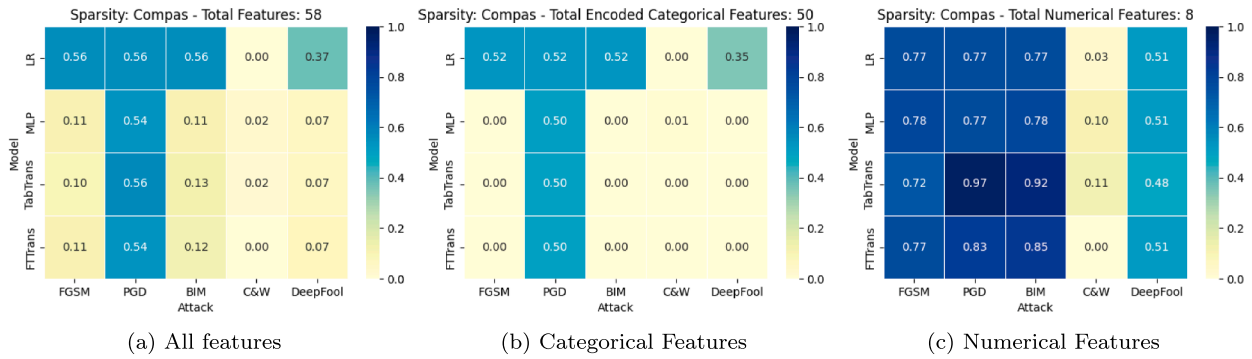


Fig. 8. Sparsity results (*sparsity rate*) of evaluated attack methods and four models on the COMPAS dataset. As with *Adult*, attacks largely ignore categorical features, though LR shows moderate categorical sparsity, highlighting model-specific vulnerability.

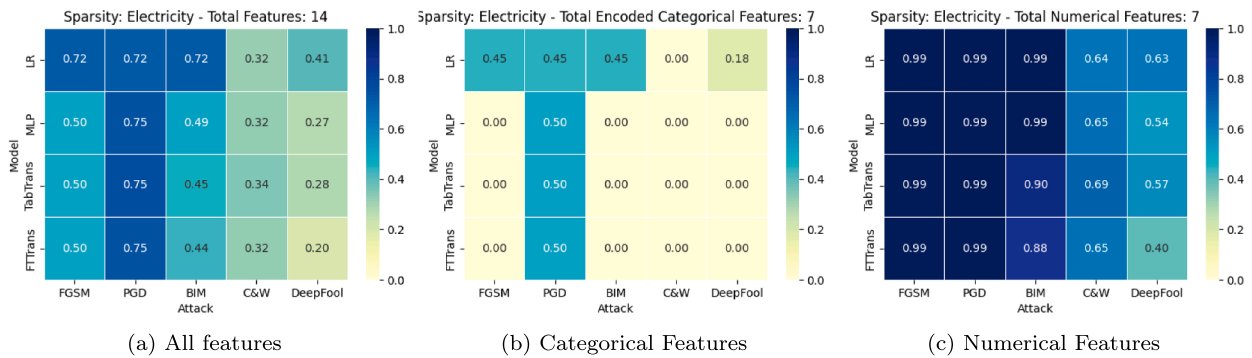


Fig. 9. Sparsity results (*sparsity rate*) of evaluated attack methods and four models on the Electricity dataset. Even with a balanced feature space (7 categorical, 7 numerical), attacks strongly prioritise numerical features, with categorical perturbations near zero.

Proximity: mixed dataset. The mixed datasets (*Adult*, *Electricity*, and *COMPAS*) reveal a consistent divide between attack types: ℓ_2 -based methods (C&W, DeepFool) generate adversarial examples that remain much closer to the originals, while ℓ_∞ -based methods (FGSM, BIM, PGD) produce substantially larger perturbations. Within this overall pattern, each dataset shows distinctive behaviour.

For the *Adult* dataset (Fig. 11a), C&W achieves the smallest distances (0.00–0.13), followed by DeepFool (0.30–0.41), whereas ℓ_∞ -based methods range much higher (0.64–1.96). PGD stands out for producing the most distant adversarial examples, reaching 4.42 in LR and 2.25 in FTTransformer. This suggests that PGD’s iterative optimisation, while effective, does so at the cost of proximity in this dataset.

The *Electricity* dataset (Fig. 11b) exhibits smaller distances overall compared to *Adult*, yet the ℓ_2 vs ℓ_∞ gap persists. The difference is particularly pronounced in LR, where ℓ_∞ attacks average around 0.89, while C&W and DeepFool remain far closer (0.17 and 0.10, respectively). Neural models show minimal variation within each attack type, indicating that proximity here depends more on dataset structure than model architecture.

The *COMPAS* dataset (Fig. 11c) continues the trend of ℓ_2 -based methods achieving the closest adversarial examples (0.00–0.07 for C&W). However, PGD again produces extreme outcomes, particularly in FTTransformer, where distances reach the highest (3.14) across all mixed datasets. This highlights that certain dataset-model-attack combinations can yield disproportionately poor proximity, underscoring the variability of imperceptibility in mixed-feature tabular data.

Proximity: numerical dataset. Numerical datasets generally exhibit smaller proximity distances than mixed datasets, with adversarial examples often remaining closer to the original inputs. The gap between

ℓ_2 - and ℓ_∞ -based methods is narrower here, though ℓ_2 attacks (C&W, DeepFool) still tend to preserve proximity better overall.

The *Higgs* dataset (Fig. 12a) provides the clearest case of uniformly close adversarial examples. Across all models and attacks, ℓ_2 distances remain very small (0.01–0.25), with ℓ_2 -based methods slightly closer (0.01–0.07) than ℓ_∞ -based methods (0.19–0.25). This suggests that the feature space in *Higgs* makes it easy to find nearby adversarial examples regardless of attack strategy.

In contrast, the *house_16H* dataset (Fig. 12b) shows pronounced variability across models. For LR and TabTransformer, ℓ_∞ -based attacks yield much larger distances (0.91–0.95), while for MLP and FTTransformer, the same attacks produce considerably closer adversarial examples (0.29–0.36). Similarly, the *jm1* dataset (Fig. 12c) highlights how optimisation choices within ℓ_∞ methods affect proximity. While distances are consistent for LR (1.13), TabTransformer shows wide variation: BIM finds relatively close examples (0.35) whereas FGSM produces much more distant ones (0.96).

Some datasets exhibit particularly extreme proximity behaviour. In *BreastCancer* (Fig. 12d), ℓ_∞ -based attacks generate some of the highest distances observed across all numerical datasets. For instance, PGD reaches 2.11 against FTTransformer, more than six times the distance of VIM (0.34) on the same model. The *WineQuality* datasets (Fig. 12e and f) show moderate distances overall, but with dataset-specific patterns: in *WineQuality-White*, PGD produces consistently higher values (up to 1.33 in LR) compared to FGSM and BIM, while this gap is less evident in *WineQuality-Red*. The *phoneme* dataset (Fig. 12g) reveals strong model dependence: PGD yields the most distant examples for LR (1.27), while FGSM does so for TabTransformer (1.05). In contrast, C&W achieves exceptionally close adversarial examples across models, reaching zero for LR and just 0.10 for transformer models. Finally, the *MiniBooNE* dataset (Fig. 12h) shows similar extremes, with PGD generating distances of

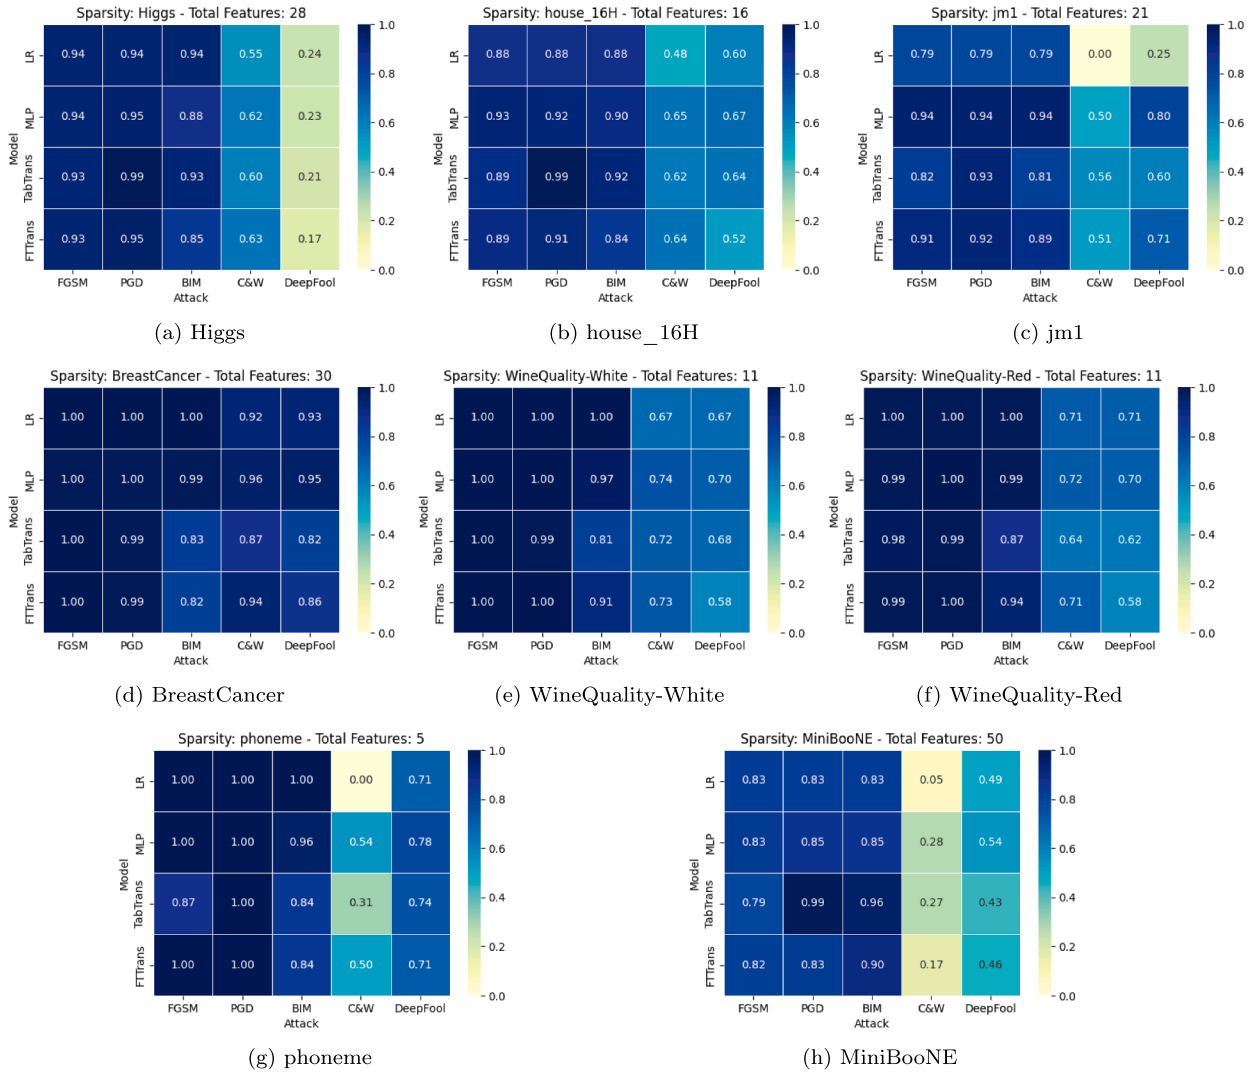


Fig. 10. Sparsity results of evaluated attack methods and four models on eight numerical datasets. ℓ_∞ attacks (FGSM, BIM, PGD) consistently modify nearly all features, while ℓ_2 attacks (C&W, DeepFool) are more selective, with sparsity strongly influenced by dataset dimensionality.




Fig. 11. Proximity results (ℓ_2 distance) of evaluated attack methods and four models on all three mixed datasets. ℓ_2 -based unbounded attacks (C&W, DeepFool) consistently produce adversarial examples much closer to the originals, while ℓ_∞ -based methods (FGSM, BIM, PGD) generate larger perturbations.

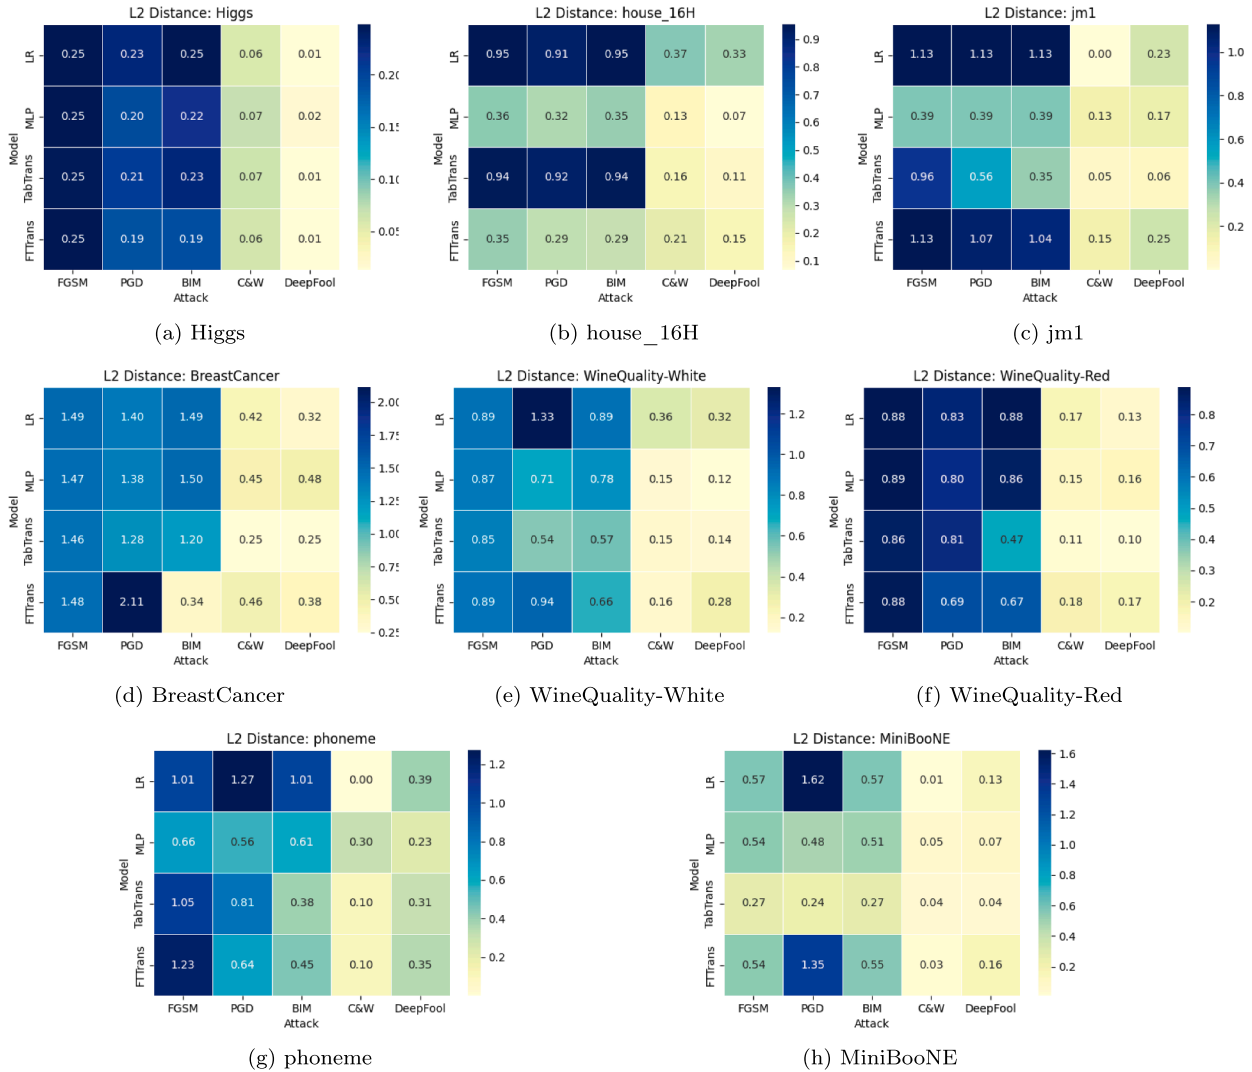


Fig. 12. Proximity results (ℓ_2 distance) of evaluated attack methods and four models on eight numerical datasets. Overall, ℓ_2 -based attacks (C&W, DeepFool) yield much closer adversarial examples, while ℓ_∞ -based methods (FGSM, BIM, PGD) often generate more distant perturbations. Architectural effects are evident, with LR generally requiring larger perturbations than MLP and transformer models.

1.62 (LR) and 1.35 (FTTransformer), far above those produced by other methods.

From an architectural perspective, LR generally produces the largest ℓ_2 distances, particularly under ℓ_∞ attacks, reflecting the need for larger perturbations to cross its linear decision boundaries. Transformer-based models exhibit more variable proximity patterns: in some cases, such as *jm1*, they yield closer adversarial examples, while in others, such as *BreastCancer*, they require substantially larger perturbations. MLP models typically fall between these two extremes, showing more stable values across datasets.

4.2.3. Task 2.3: deviation

Our deviation analysis examines how significantly the adversarial examples differ from the original data distribution. The heatmaps presented in Figs. 13 and 14 reveal clear patterns in the outlier rates produced by different attack algorithms across model architectures and datasets.

Deviation: mixed dataset. Across three mixed datasets, a consistent divide emerges between attack types. ℓ_∞ -based methods (FGSM, BIM, PGD) typically drive adversarial examples far outside the original data distribution, yielding very high outlier rates, whereas ℓ_2 -based methods (C&W, DeepFool) preserve distributional alignment far more effectively.

The *Adult* dataset (Fig. 13a) exemplifies this pattern: ℓ_∞ -based attacks achieve near-perfect outlier rates (close to 1.00) across all models, while ℓ_2 -based methods remain much lower, with C&W at 0.14–0.34 and DeepFool at 0.20–0.46. In *Electricity* (Fig. 13b), the overall magnitude of outlier rates decreases, yet the divide persists. ℓ_∞ -based attacks still dominate (0.60–0.88), with TabTransformer particularly vulnerable, while ℓ_2 -based methods again stay closer to the original distribution (0.02–0.39). The *COMPAS* dataset (Fig. 13b) shows the starkest contrast: all ℓ_∞ -based attacks reach a perfect outlier rate of 1.00 across models, while C&W achieves consistently minimal rates (0.06–0.09) and DeepFool produces moderate ones (0.22–0.51).

Deviation: numerical dataset. Across numerical datasets, a clear pattern emerges: ℓ_∞ attacks (FGSM, BIM, PGD) generally drive outlier rates close to 1.0, while ℓ_2 methods (C&W, DeepFool) tend to preserve distributional alignment more effectively, though with important dataset-specific variations.

The *Higgs* dataset (Fig. 14a) provides the clearest baseline: ℓ_∞ attacks yield near-perfect outlier rates, while ℓ_2 attacks remain low (0.12–0.25). By contrast, both *house_16H* and *jm1* (Fig. 14b and c) show elevated outlier rates even for ℓ_2 methods, reaching up to 0.65 for C&W and above 0.80 for DeepFool. A similar trend is seen in *MiniBooNE* (Fig. 14h),




Fig. 13. Deviation results (*outlier rates*) of evaluated attack methods and four models on all three *mixed* dataset. ℓ_∞ -based attacks (FGSM, BIM, PGD) consistently push adversarial examples far outside the data distribution (outlier rates near 1.0), while ℓ_2 -based methods (C&W, DeepFool) generally preserve distributional alignment.

where adversarial examples from all methods frequently fall outside the data distribution, indicating that these feature spaces are particularly sensitive to perturbations regardless of norm constraints.

The *BreastCancer* dataset (Fig. 14d) is the strongest anomaly. Here, even small perturbations push samples out of distribution, with all methods, including C&W and DeepFool, producing high outlier rates (0.67–0.93). This suggests that the dataset’s compact, tightly clustered structure leaves little room for imperceptible modification.

The *WineQuality* datasets (Fig. 14e and f) show contrasting behaviours. White wine data align with the overall trend: almost all examples by ℓ_∞ attacks are out-of-distribution (0.98–1.00), while ℓ_2 methods yield lower but still non-trivial deviation (0.30–0.73). In contrast, the red wine dataset is less extreme and outlier rates are lower and more variable, with ℓ_∞ attacks occasionally dropping below 1.0 (e.g., BIM at 0.82 on TabTransformer). Finally, the *phoneme* dataset (Fig. 14g) illustrates the role of model architecture. While LR collapses under ℓ_∞ attacks (outlier rates of 1.0), TabTransformer shows remarkable resilience, with BIM reducing deviation to just 0.11. This highlights how complex architectures can sometimes accommodate perturbations without generating deviated adversarial examples.

4.2.4. Task 2.4: sensitivity

Our sensitivity analysis examines how adversarial attacks handle narrow-guard feature perturbation, particularly for features with narrow distributions in tabular data. The heatmaps in Figs. 15 and 16 reveal complex patterns that vary significantly across datasets, attack algorithms, and model architectures. Rather than showing consistent behaviours, the sensitivity metrics highlight the contextual nature of how perturbations interact with narrowly distributed features.

Sensitivity: mixed dataset. Across mixed datasets, a clear divide emerges between attack types. C&W consistently yields the lowest sensitivity scores (0.00–0.06), reflecting its ℓ_2 objective that discourages large deviations in any single feature. By contrast, ℓ_∞ -based attacks generally induce higher sensitivity, with FGSM and BIM often producing the most extreme perturbations to narrow-distribution features, while PGD shows a somewhat more controlled pattern across models and datasets.

Model architecture further shapes these outcomes. LR models typically exhibit relatively low sensitivity, even under ℓ_∞ attacks (e.g., 0.13 in *Adult*, 0.41 in *Electricity*), suggesting that their linear boundaries limit perturbations to narrow-guard features. In contrast, MLP models can amplify sensitivity substantially, as seen in *COMPAS* where FGSM and BIM reach 0.99. Transformer-based models usually fall between these extremes, but can occasionally produce very high spikes, such as the TabTransformer in *Adult* where FGSM reaches 3.37 — the highest sensitivity score observed across all mixed datasets.

Dataset-specific outcomes illustrate these broader trends. In *Adult* (Fig. 15a), the gap between LR (0.13) and TabTransformer (3.37) under

FGSM shows how architectural complexity can magnify perturbations on narrow features. The *Electricity* dataset (Fig. 15b) displays uniformity across ℓ_∞ methods within each model: all three attacks yield identical scores (e.g., 0.41 in LR, 0.13–0.19 in neural models), indicating that the dataset’s distribution constrains perturbation behaviour. In *COMPAS* (Fig. 15c), MLP records much larger scores (0.99 for FGSM and BIM) than transformer models (0.47–0.61), while PGD consistently produces smaller values than the other ℓ_∞ methods, reflecting differences in optimisation strategy.

Sensitivity: numerical dataset. Numerical datasets reveal diverse patterns of sensitivity, with some showing very limited perturbation of narrow-distribution features and others displaying pronounced vulnerability. The *Higgs* (Fig. 16a) and *MiniBooNE* (Fig. 16h) datasets stand out for their uniformly low sensitivity scores across all attack-model combinations (typically 0.00–0.06). This suggests that either narrow-distribution features are absent or that adversarial examples can be generated without significantly altering them.

In contrast, the *jm1* dataset (Fig. 16c) produces substantially higher scores, particularly for LR under ℓ_∞ attacks (up to 1.23), indicating that narrowly distributed features strongly influence predictions in this dataset. More complex models, such as TabTransformer and FTTransformer, reduce these values (0.29–0.61), suggesting a lesser reliance on such features. The *phoneme* dataset (Fig. 16g) also shows high sensitivity, especially in transformer models targeted by FGSM and PGD, confirming that some datasets contain narrow features that adversaries can exploit.

The *WineQuality* datasets (Fig. 16e and f) occupy a middle ground. In *White*, PGD produces notably higher scores (0.31–0.54), particularly in LR and FTTransformer, whereas *Red* exhibits more uniform behaviour, with ℓ_∞ attacks clustering together and ℓ_2 methods consistently lower.

From an architectural perspective, LR amplifies dataset effects: it shows either very high sensitivity (as in *jm1*) or very low (as in *Higgs*). Transformer-based models display greater variability, sometimes producing elevated scores (as in *phoneme*) and sometimes low values (as in *MiniBooNE*). MLP models generally lie between these extremes.

Across attack methods, C&W consistently yields the lowest sensitivity values, reflecting its ℓ_2 -based optimisation that avoids strong perturbation of individual features. PGD, by contrast, shows the widest variation, producing high scores in some datasets (e.g., *WineQuality-White*, *phoneme*) and more moderate values elsewhere, suggesting that its iterative procedure adapts strongly to dataset-specific structures.

Overall, sensitivity to narrow-distribution feature perturbation is highly contextual, driven by the interplay of dataset characteristics, model reliance on specific features, and the perturbation strategy of the attack. These results caution against generalising attack behaviours and highlight the need for dataset-specific evaluation when assessing imperceptibility from a sensitivity perspective.

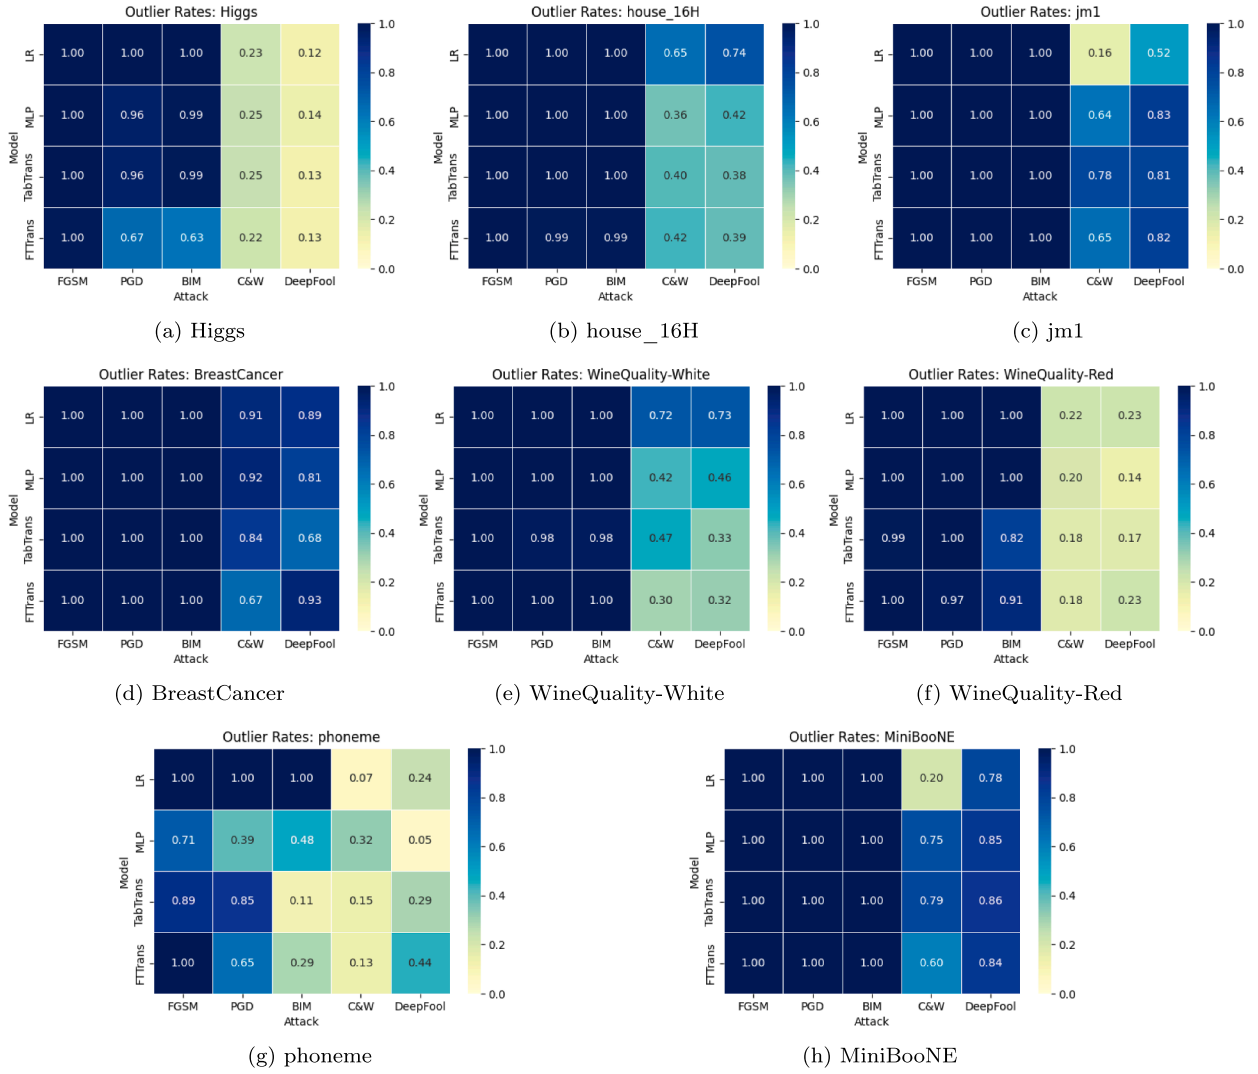


Fig. 14. Deviation results (*outlier rate*) of evaluated attack methods and four models on eight *numerical* datasets. ℓ_∞ attacks (FGSM, BIM, PGD) generally push nearly all samples out-of-distribution, while ℓ_2 methods (C&W, DeepFool) preserve distributions more effectively but can still produce high deviation on compact datasets such as *BreastCancer*.

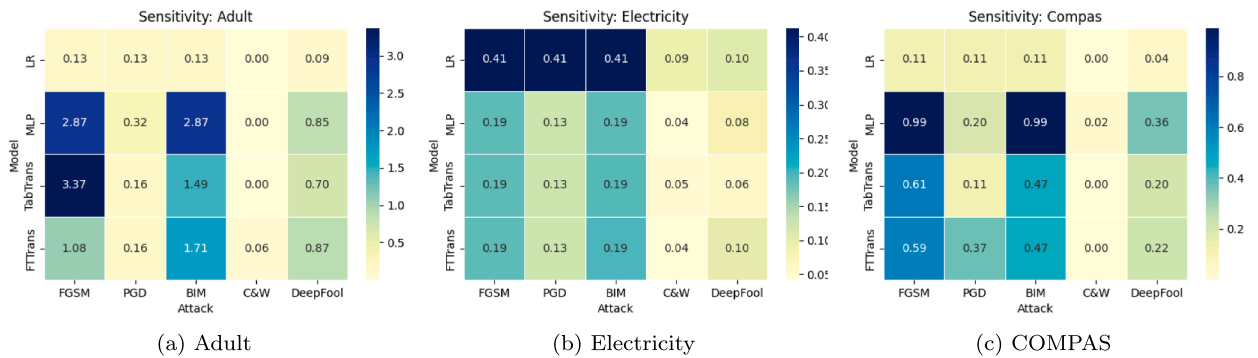


Fig. 15. Sensitivity results (*average sensitivity*) of evaluated attack methods and four models on all three *mixed* dataset. ℓ_2 -based attacks (C&W, DeepFool) consistently yield the lowest sensitivity, preserving narrow-distribution features, while ℓ_∞ -based methods (FGSM, BIM, PGD) show higher variation across architectures and datasets.

4.3. Task 3: trade-off analysis

Analysing the relationship between attack success rate (ASR) and imperceptibility score (IS) provides critical insights into the relationship between effectiveness and imperceptibility of adversarial attacks on tab-

ular data. By visualising this relationship through a 2D density plot in Fig. 17, we can discern patterns that illuminate the interplay between these two crucial factors. We take the Gaussian noise method as baseline and select the maximum ASR value (0.659) and the minimum IS value (0.181) from all adversarial examples generated by Gaussian noise. The

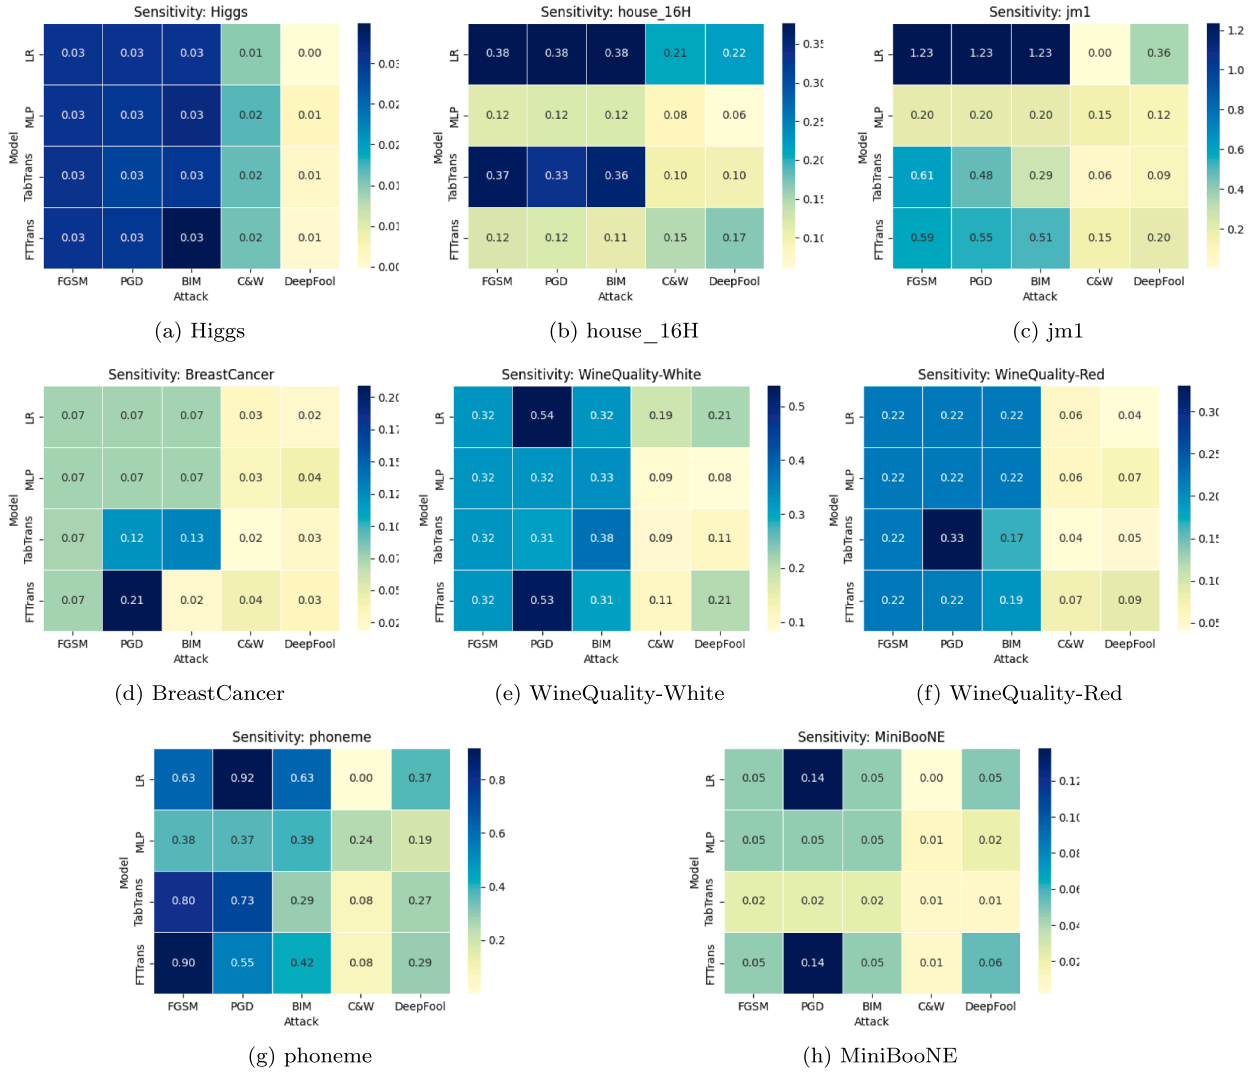


Fig. 16. Sensitivity results (*average sensitivity*) of evaluated attack methods and four models on eight *numerical* datasets. Overall, ℓ_2 -based attacks (C&W, DeepFool) again preserve narrow-distribution features most effectively, while ℓ_∞ -based attacks (FGSM, BIM, PGD) show strong dataset- and model-dependent variability, with *jm1* and *phoneme* producing the highest sensitivities and *Higgs* and *MiniBooNE* the lowest.

graphs were divided into four distinct sections based on two thresholds, enabling us to categorise different scenarios and gain a clearer understanding of their impact.

Effective and imperceptible (High ASR, Low IS). The most desirable outcome for adversarial attacks occurs when examples successfully fool models while remaining nearly indistinguishable from original data. The density plot reveals that DeepFool consistently achieves co-optimal performance, with its highest density region falling in this quadrant. DeepFool’s iterative approach of finding minimal perturbations to cross decision boundaries clearly excels at preserving tabular data characteristics while maintaining high effectiveness. C&W also demonstrates strong performance in this quadrant for a portion of its examples, though it shows a bimodal distribution across both imperceptible regions. This suggests that C&W can achieve the ideal results in many cases but may sometimes sacrifice effectiveness to maintain imperceptibility.

Effective but perceptible (High ASR, High IS). This quadrant contains attacks that successfully mislead models but make noticeable modifications to the data. The density plots show that FGSM, PGD, and BIM consistently fall into this category, achieving high attack success rates at the cost of more significant data alterations. These ℓ_∞ -based

attacks effectively fool models but often modify features in ways that could compromise data integrity or be detected in quality control processes.

Ineffective but imperceptible (Low ASR, Low IS). Attacks in this quadrant make subtle changes that preserve data characteristics but fail to successfully mislead models. C&W shows a significant density in this region, indicating that it sometimes generates examples that maintain excellent imperceptibility but cannot effectively fool the model. This highlights C&W’s explicit optimisation for minimal perturbations, which can sometimes come at the expense of attack effectiveness.

Ineffective and perceptible (Low ASR, High IS). The least desirable outcome occurs when attacks make noticeable changes yet fail to mislead the model. Gaussian noise predominantly falls in this category, confirming its poor performance as a baseline comparison. Its high-density region centres on moderate ASR values with high imperceptibility scores, demonstrating why random noise is considered both ineffective and easily perceptible.

Correlation analysis. To further investigate the trade-off between attack effectiveness and imperceptibility, we conducted a correlation




Fig. 17. The 2D density plot shows the attack success rate (ASR) and imperceptibility score (IS) across Gaussian noise and five different attack methods. The plot is divided into four sectors based on the maximum ASR value (0.659) and the minimum IS value (0.181) observed among all adversarial examples generated by Gaussian noise. Gaussian noise is considered an ineffective and perceptible method for generating adversarial examples for tabular data. FGSM, PGD, and BIM are categorised as effective but perceptible methods. C&W attack has two high-density regions: one that is Effective and Imperceptible, and another that is Ineffective but Imperceptible. Most of DeepFool attack’s high-density regions fall into the Effective and Imperceptible sector.

Table 8

Correlation between Attack Success Rate and Imperceptibility Metrics by Attack Type. Note: Bold values (positive correlation) indicate trade-off; Italic values (negative correlation) indicate efficiency. Abbreviations: SpaR = Sparsity Rate, Sens = Sensitivity, OR = Outlier Rate, IS = Imperceptibility Score.

Attack	ℓ_2	SpaR	Sens	OR	IS	Avg
Gaussian	-0.059	0.077	0.375	0.359	0.494	0.249
FGSM	0.220	0.229	-0.283	0.278	0.410	0.171
BIM	0.344	0.296	-0.269	0.218	0.472	0.212
PGD	0.081	0.252	0.212	0.428	0.390	0.273
C&W	0.316	0.326	0.185	0.221	0.297	0.269
DeepFool	<i>green-0.536</i>	0.206	<i>green-0.427</i>	-0.212	-0.206	<i>green-0.235</i>

analysis between Attack Success Rate (ASR) and five imperceptibility metrics across six attack methods. Table 8 presents the Pearson correlation coefficients for each attack-metric pair. PGD, BIM, FGSM and C&W demonstrate positive correlations with an average r ranging from 0.171 to 0.273, indicating a clear trade-off where higher perturbation magnitudes are associated with higher success rates. In contrast, DeepFool exhibits negative correlations across most metrics (avg $r = -0.235$), demonstrating efficient behaviour where it achieves high success rates with lower perturbation magnitudes (ℓ_2 : $r = -0.536$, Sensitivity: $r = -0.427$). These findings suggest that the trade-off of effectiveness and imperceptibility is attack-dependent rather than universal (See more analysis in Appendix C).

Overall performance comparison. The density plots and correlation analysis provide clear evidence for ranking the overall performance of different attack methods:

1. DeepFool emerges as the most favourable approach, consistently generating examples that are both highly effective and imperceptible. Its iterative linearisation of decision boundaries enables precise identification of minimal perturbations needed to cross classification boundaries, resulting in subtle modifications that maintain data integrity while achieving high ASR (Moosavi-Dezfooli et al., 2016).
2. C&W shows mixed results with two distinct behaviour patterns — one group achieving the simultaneous strength and another maintaining imperceptibility at the cost of effectiveness.
3. The ℓ_∞ -based attacks (FGSM, PGD, and BIM) prioritise effectiveness over imperceptibility, making them suitable for scenarios where attack success is more important than maintaining data characteristics.
4. Gaussian noise serves as an appropriate baseline, demonstrating poor performance in both dimensions as expected.

This analysis provides valuable guidance for selecting appropriate attack methods based on specific requirements for tabular data scenarios, highlighting the fundamental trade-off between effectiveness and imperceptibility in adversarial machine learning.

5. Discussion

5.1. Investigating the inverse relationship between BIM attack budget and success rate

As presented in Section 4.1, our Task 1 evaluation results across both mixed and numerical datasets reveal an intriguing and counterintuitive phenomenon. While increasing epsilon ϵ values generally leads to improved success rates for most attack methods, the BIM attack on the FTTransformer model shows a notable decline in success rates at higher

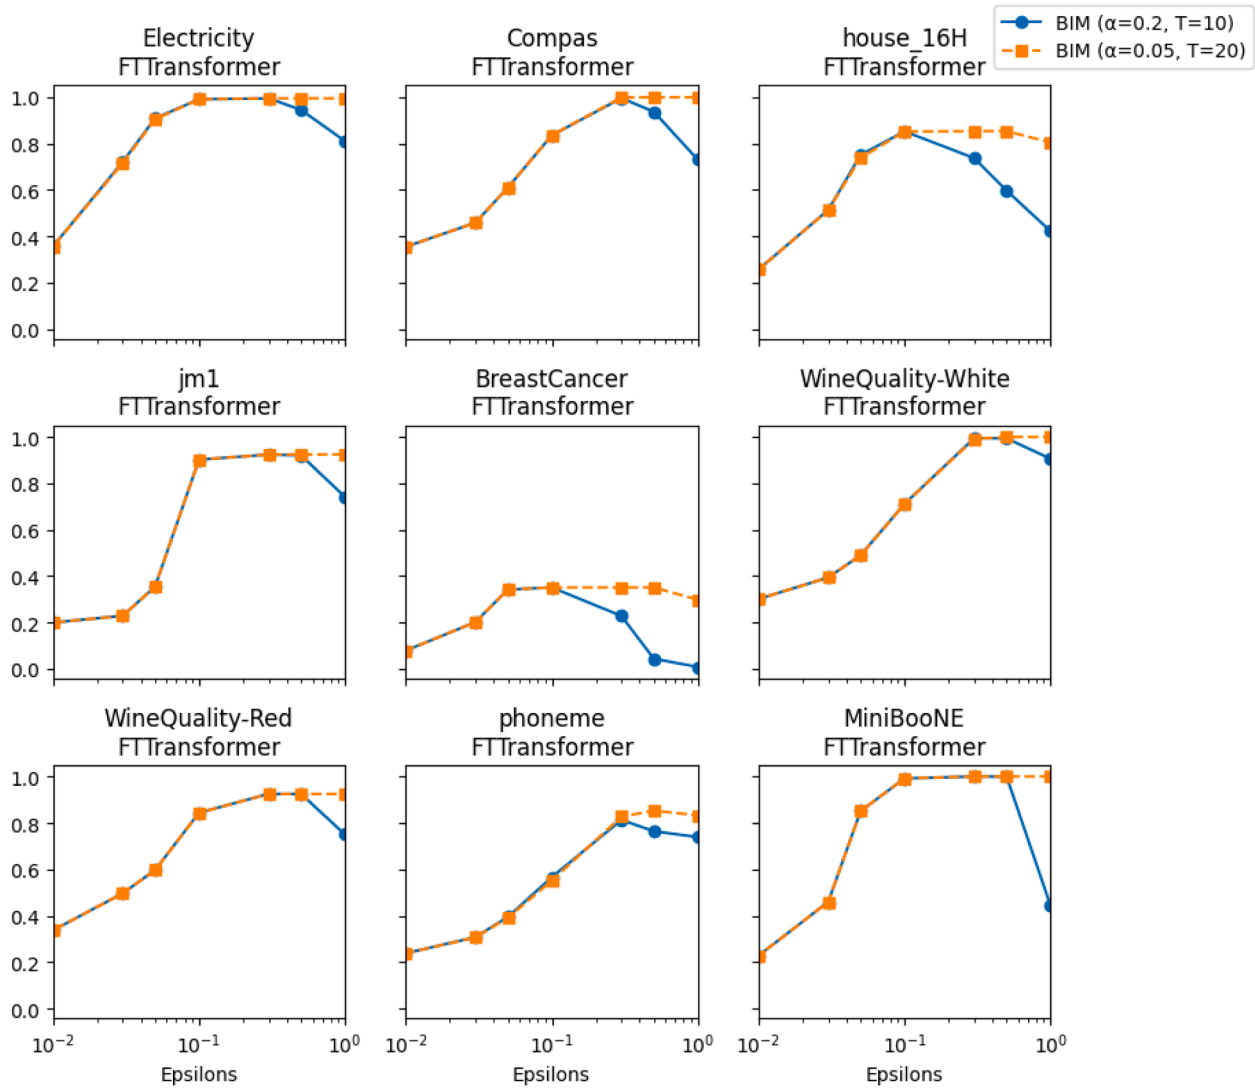


Fig. 18. BIM attacks on FTTransformer stop dropping in attack success rate after adjusting step size (α) hyperparameters. The orange lines, representing the adjusted BIM implementation, consistently maintain high attack success rates across all epsilon values, eliminating the dramatic drops observed with the default parameters.

perturbation budgets. This inverse relationship between attack budget and effectiveness contradicts conventional adversarial attack theory, where larger perturbation budgets typically enable more successful attacks.

The plots in Figs. 3 to 6 clearly demonstrate this unexpected pattern across multiple datasets, including *Electricity*, *COMPAS*, *house_16H*, *BreastCancer*, and *MiniBooNE*. In these cases, the attack success rate of BIM initially increase with epsilon values but then significantly decline at higher epsilon values, sometimes dropping dramatically. For example, on the *BreastCancer* dataset, the success rate drops from approximately 35% to nearly 0% at the highest epsilon value, while on *MiniBooNE*, it plummets from 100% to about 40%.

The primary factor may explain this counterintuitive behaviour, which is gradient saturation effects. In BIM's iterative approach, the step size (α) and number of iterations (T) play critical roles. When using default parameters ($\alpha=0.2$, $T=10$), the relatively large step size may cause overshooting at higher epsilon values. This occurs because BIM computes gradients with respect to the input and takes steps in that direction. As epsilon increases, these steps can become too large, causing the attack to miss optimal adversarial regions and produce less effective perturbations.

Our follow-up experiment demonstrates that adjusting BIM's hyperparameters can mitigate this issue. By reducing the step size (α) from 0.2 to 0.05 and increasing iterations (T) from 10 to 20, in Fig. 18, we observe that the modified BIM attack maintains high success rates even at larger epsilon values across all datasets. This confirms that the original decline was primarily due to optimisation challenges rather than fundamental limitations of the attack method.

This finding has important implications for adversarial attack research on tabular data:

- FTTransformer models possess unique adversarial robustness characteristics that differ from other model architectures.
- Attack hyperparameters require careful tuning based on both the model architecture and dataset characteristics.

5.2. Exploring design strategies for effective and imperceptible adversarial attacks on tabular data

In light of the results from analysing the relationship between attack success rate (ASR) and imperceptibility score (IS), achieving an optimal balance between effectiveness and imperceptibility is crucial in

Table 9
The encoding methods employed in recent papers on adversarial attacks targeting tabular data.

Paper	Year	Encoding method	Distance metric
Ballet et al. (2019)	2019	Drop all categorical features	ℓ_p -norm & Weighted ℓ_p -norm
Mathov et al. (2022)	2021	Label encoding	ℓ_2 -norm
Chernikova and Oprea (2022)	2022	One-hot encoding	ℓ_2 -norm
Cartella et al. (2021)	2021	One-hot encoding	ℓ_2 -norm
Kireev et al. (2023)	2022	Discrete continuous features	Cost function
Zhou et al. (2022)	2022	Discrete continuous features	ℓ_1 -norm

designing adversarial attack algorithms for tabular data. One notable observation is that ℓ_∞ attacks tend to generate highly effective adversarial examples, whereas ℓ_2 attacks are more adept at producing imperceptible examples. The key challenge lies in finding the equilibrium between these two aspects.

To design effective and imperceptible adversarial attack algorithms for tabular data, several strategies can be explored:

- **Optimisation techniques:** Employing advanced optimisation techniques can enhance the efficiency of adversarial attack algorithms. Techniques such as evolutionary algorithms, genetic algorithms, or gradient-based optimisation methods can be tailored to optimise both ASR and IS simultaneously, thereby facilitating the creation of more effective and imperceptible adversarial examples.
- **Feature engineering:** Leveraging domain-specific knowledge and feature engineering techniques can enhance the robustness and imperceptibility of adversarial attacks. By identifying and manipulating key features within the tabular data that are most susceptible to manipulation, attackers can craft adversarial examples that achieve their objectives while minimising perceptible changes to the data.

5.3. Evaluating the suitability of one-hot encoding for adversarial attacks on tabular data

Our sparsity evaluation in Section 4 revealed that categorical features are rarely perturbed under one-hot encoding. Since each category is represented by multiple binary indicators, altering a categorical attribute requires flipping at least two dimensions, which increases the perturbation cost. Consequently, optimisation algorithms tend to modify numerical features instead, resulting in dense perturbations within continuous attributes and limited diversity across feature types.

Several existing studies have explored different strategies for handling categorical features for adversarial attacks on tabular data in Table 9. Ballet et al. (2019) proposed dropping all categorical features and using the ℓ_p -norm and Weighted ℓ_p -norm as distance metrics. Mathov et al. (2022) suggested using label encoding for categorical data, though they did not specify the distance metric employed. Chernikova and Oprea (2022) and Cartella et al. (2021) both used one-hot encoding and applied the ℓ_2 -norm as their distance metric. On the other hand, Kireev et al. (2023) and Zhou et al. (2022) opted for discretising continuous features and used a cost function as a distance measure, without specifying a particular norm.

Moreover, exploring alternative distance metrics presents a promising direction for future research. Traditional metrics like the ℓ_p -norm may not be well-suited for the mixed data types often found in tabular datasets. Metrics such as Gower's distance (Gower, 1971), which can handle mixed types of data (continuous, ordinal, and categorical), could provide a more accurate measure of similarity for tabular data. Additionally, other categorical feature similarity measures, such as those proposed by Cost and Salzberg (1993) and Le and Ho (2005), offer potential improvements by considering the unique characteristics of categorical data. By integrating these distance metrics into the design of adversarial attack algorithms, researchers can develop more effective and nuanced methods that are better tailored to the complexities of tabular data.

6. Conclusion

In this paper, we conducted a comprehensive benchmark analysis of adversarial attacks on tabular data, focusing on both their *effectiveness* and *imperceptibility*. Using a diverse set of 11 datasets, encompassing both mixed and numerical data types, we evaluated the performance of five different adversarial attacks across four predictive models. The results reveal substantial variation in attack performance depending on model type and dataset characteristics, underscoring the unique challenges of adversarial robustness in structured data.

Our benchmark demonstrates that highly effective attacks often compromise imperceptibility, while perturbations that remain realistic tend to achieve lower success rates. These findings quantify the trade-off between attack potency and subtlety, establishing a reproducible reference point for future work on tabular adversarial robustness. By providing unified evaluation metrics, standardised preprocessing, and open-source implementations, this study offers a consistent foundation for analysing vulnerabilities and improving defences in tabular machine learning systems.

Future work. While this benchmark establishes a strong foundation for evaluating adversarial robustness in tabular data, several directions remain open for future exploration:

- **Incorporating defence mechanisms:** Extending the benchmark to include state-of-the-art adversarial defences (e.g., adversarial training, feature denoising, or certified robustness) will allow a more holistic assessment of both attack and defence strategies in practical deployments.
- **Expanding domain coverage:** Future iterations can include additional real-world domains such as cybersecurity, fraud detection, and medical diagnostics to further validate generalisability and cross-domain consistency.
- **Benchmarking black-box attacks:** Beyond white-box settings, future work will extend the benchmark to include black-box and adaptive attack scenarios (Brendel et al., 2017; Cartella et al., 2021; Chen et al., 2017). Such evaluations will enhance the practical relevance of the benchmark by reflecting real-world conditions where model internals are not fully observable.
- **Non-uniform feature importance:** This study assumes uniform feature contribution, yet real-world tabular data often contain features with heterogeneous influence. Investigating non-uniform or adaptive adversarial perturbations will provide deeper insight into model vulnerability under realistic conditions (Erdemir et al., 2021; Nandy et al., 2023).

CRedit authorship contribution statement

Zhipeng He: Conceptualisation, Methodology, Software, Investigation, Writing – original draft, Visualisation; **Chun Ouyang:** Conceptualisation, Methodology, Investigation, Writing – original draft; **Lijie Wen:** Methodology, Investigation; **Cong Liu:** Supervision; **Catarina Moreira:** Methodology, Investigation, Supervision.

Data availability

The implementation code, including data processing scripts and experimental pipelines, is openly available at <https://github.com/ZhipengHe/TabAttackBench>.

Acknowledgements

The reported research forms part of part of a Ph.D. project supported by QUT Postgraduate Research Award (QUTPRA) at Queensland University of Technology (QUT), Australia. This work was supported by national funds through FCT (Fundação para a Ciência e a Tecnologia), under the project UID/04152/2025 — Centro de Investigação em Gestão de Informação (MagIC)/NOVA IMS and UID/PRR/04152/2025. The work reported in this article was also partially supported under the auspices of the UNESCO Chair on AI & VR.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Selected attack budgets (ϵ) by ASR

This appendix summarises the selected attack budgets (ϵ) used in the main evaluation. For each dataset-model combination, [Table A.10](#) reports the value of ϵ that achieved the highest attack success rate (ASR) for each attack algorithm.

Appendix B. Full experimental results

This appendix reports the full experimental results for all adversarial attack configurations considered in the paper. [Table B.11](#) lists ASR and all imperceptibility metrics for each dataset, model, attack type, and attack budget ϵ , where ℓ_2^* and SEN^* denote normalised values.

Table A.10

Best attack budget (ϵ) settings for four models on different datasets from the evaluation.

Datasets	Model	Gaussian	FGSM	PGD	BIM	C&W	DeepFool
Adult	LR	1	0.3	0.3	0.3	0.5	1
Adult	MLP	0.5	1	1	1	0.01	1
Adult	TabTrans	0.5	1	0.5	0.5	0.01	1
Adult	FTTrans	1	0.3	0.5	0.5	0.5	1
Electricity	LR	1	0.3	0.3	0.3	1	0.5
Electricity	MLP	1	0.1	0.1	0.1	0.3	0.3
Electricity	TabTrans	1	0.1	0.1	0.1	0.3	0.3
Electricity	FTTrans	1	0.1	0.1	0.1	0.3	0.3
COMPAS	LR	0.5	0.3	0.3	0.3	0.1	1
COMPAS	MLP	0.3	0.5	0.5	0.5	1	1
COMPAS	TabTrans	1	0.3	0.3	0.3	0.1	0.5
COMPAS	FTTrans	0.07	0.3	1	0.3	0.01	0.5
Higgs	LR	1	0.07	0.07	0.07	0.3	0.1
Higgs	MLP	1	0.07	0.07	0.07	0.3	0.1
Higgs	TabTrans	1	0.07	0.07	0.07	0.3	0.1
Higgs	FTTrans	1	0.07	0.07	0.07	0.3	0.1
house_16H	LR	0.01	0.3	0.3	0.3	1	1
house_16H	MLP	1	0.1	0.1	0.1	0.5	0.3
house_16H	TabTrans	1	0.3	0.3	0.3	0.5	0.5
house_16H	FTTrans	1	0.1	0.1	0.1	1	0.5
jm1	LR	0.01	0.5	0.5	0.5	0.01	1
jm1	MLP	1	0.1	0.1	0.1	0.5	0.5
jm1	TabTrans	1	0.3	0.3	0.3	0.3	0.5
jm1	FTTrans	1	0.3	0.3	0.3	0.5	0.5
BreastCancer	LR	1	0.3	0.3	0.3	1	1
BreastCancer	MLP	1	0.3	0.3	0.3	1	1
BreastCancer	TabTrans	1	0.3	0.5	0.5	1	1
BreastCancer	FTTrans	1	0.3	1	0.1	1	0.5
Wine-White	LR	1	0.3	0.5	0.3	1	1
Wine-White	MLP	1	0.3	0.3	0.3	0.5	0.5
Wine-White	TabTrans	1	0.3	0.3	0.3	0.5	0.5
Wine-White	FTTrans	1	0.3	0.5	0.3	1	1
Wine-Red	LR	1	0.3	0.3	0.3	1	0.5
Wine-Red	MLP	1	0.3	0.3	0.3	0.5	0.5
Wine-Red	TabTrans	1	0.3	0.5	0.3	0.5	0.5
Wine-Red	FTTrans	1	0.3	0.3	0.3	1	0.5
phoneme	LR	1	0.5	1	0.5	0.01	1
phoneme	MLP	1	0.3	0.3	0.3	1	1
phoneme	TabTrans	1	1	1	0.3	1	1
phoneme	FTTrans	1	1	0.5	0.3	0.5	1
MiniBooNE	LR	1	0.1	0.3	0.1	0.1	0.3
MiniBooNE	MLP	1	0.1	0.1	0.1	0.1	0.3
MiniBooNE	TabTrans	1	0.07	0.07	0.07	0.1	0.3
MiniBooNE	FTTrans	1	0.1	0.3	0.1	0.1	0.5

Table B.11Full experimental results from all adversarial attack configurations. ℓ_2^* and SEN* refer to normalised value.

Dataset	Model	Attack	ϵ	ASR	SpaR	ℓ_2	ℓ_2^*	SEN	SEN*	OR	IS
Adult	LR	Gaussian	1	0.1962	0.5216	0.7160	0.3660	0.0796	0.0740	0.2011	0.0436
Adult	LR	FGSM	0.3	0.9997	0.4289	1.9565	1.0000	0.1324	0.1230	0.9986	0.0807
Adult	LR	PGD	0.3	0.9997	0.4289	1.9460	0.9946	0.1324	0.1230	0.9980	0.0807
Adult	LR	BIM	0.3	0.9997	0.4289	1.9565	1.0000	0.1324	0.1230	0.9986	0.0807
Adult	LR	C&W	1	0.1999	0.9976	0.0510	0.0261	0.0058	0.0054	0.1598	0.0050
Adult	LR	DeepFool	1	0.8334	0.3272	0.4113	0.2102	0.0894	0.0830	0.5236	0.0463
Adult	MLP	Gaussian	1	0.2168	0.5229	0.7205	0.3683	0.0800	0.0743	0.2302	0.0450
Adult	MLP	FGSM	0.3	0.2985	0.0496	0.6323	0.3232	1.0415	0.9674	0.7269	0.0396
Adult	MLP	PGD	0.3	0.2906	0.5177	1.3321	0.6809	0.0973	0.0904	0.3941	0.0593
Adult	MLP	BIM	0.3	0.2985	0.0496	0.6323	0.3232	1.0415	0.9674	0.7269	0.0396
Adult	MLP	C&W	1	0.2412	0.9975	0.0000	0.0000	0.0000	0.0000	0.1598	0.0005
Adult	MLP	DeepFool	1	0.4141	0.0339	0.3624	0.1852	0.8500	0.7895	0.5651	0.0270
Adult	TabTrans	Gaussian	1	0.2143	0.5228	0.7169	0.3664	0.0796	0.0739	0.2335	0.0450
Adult	TabTrans	FGSM	0.3	0.4307	0.0491	0.6303	0.3222	1.0413	0.9672	0.9964	0.0398
Adult	TabTrans	PGD	0.3	0.6069	0.5298	1.3461	0.6880	0.0951	0.0883	0.9992	0.0644
Adult	TabTrans	BIM	0.3	0.5025	0.0544	0.5985	0.3059	0.9274	0.8614	0.9994	0.0426
Adult	TabTrans	C&W	1	0.2417	0.9975	0.0009	0.0005	0.0001	0.0001	0.1614	0.0006
Adult	TabTrans	DeepFool	1	0.5765	0.0364	0.3054	0.1561	0.6964	0.6469	0.4621	0.0272
Adult	FTTrans	Gaussian	1	0.2220	0.5210	0.7190	0.3675	0.0804	0.0746	0.2040	0.0440
Adult	FTTrans	FGSM	0.3	0.8967	0.0469	0.6353	0.3247	1.0766	1.0000	0.9663	0.0385
Adult	FTTrans	PGD	0.3	0.8687	0.5234	1.3617	0.6960	0.0963	0.0895	0.9408	0.0647
Adult	FTTrans	BIM	0.3	0.9154	0.0495	0.6310	0.3225	1.0201	0.9475	0.9764	0.0400
Adult	FTTrans	C&W	1	0.3057	0.8866	0.1345	0.0687	0.0588	0.0546	0.1386	0.0246
Adult	FTTrans	DeepFool	1	0.9754	0.0370	0.3029	0.1548	0.8710	0.8091	0.1986	0.0256
Electricity	LR	Gaussian	1	0.4116	0.7505	0.7795	0.3984	0.6045	0.5615	0.9332	0.1496
Electricity	LR	FGSM	0.3	0.9997	0.7224	0.8874	0.4536	0.4125	0.3831	0.6398	0.1291
Electricity	LR	PGD	0.3	0.9996	0.7224	0.8524	0.4357	0.4125	0.3832	0.6454	0.1278
Electricity	LR	BIM	0.3	0.9997	0.7224	0.8874	0.4536	0.4125	0.3831	0.6398	0.1291
Electricity	LR	C&W	1	0.9961	0.9911	0.1673	0.0855	0.0925	0.0859	0.0506	0.0230
Electricity	LR	DeepFool	1	1.0000	0.4342	0.0986	0.0504	0.1054	0.0979	0.3918	0.0291
Electricity	MLP	Gaussian	1	0.4835	0.7465	0.7695	0.3933	0.6042	0.5612	0.9350	0.1488
Electricity	MLP	FGSM	0.3	0.9960	0.4985	0.7107	0.3632	0.5645	0.5244	0.7767	0.1260
Electricity	MLP	PGD	0.3	0.9972	0.7488	0.7146	0.3652	0.3812	0.3540	0.8317	0.1237
Electricity	MLP	BIM	0.3	0.9977	0.4968	0.6977	0.3566	0.5669	0.5265	0.8937	0.1279
Electricity	MLP	C&W	1	0.9947	0.9916	0.0732	0.0374	0.0506	0.0470	0.0407	0.0139
Electricity	MLP	DeepFool	1	0.9938	0.3805	0.0487	0.0249	0.0838	0.0779	0.2522	0.0173
Electricity	TabTrans	Gaussian	1	0.4433	0.7485	0.7713	0.3942	0.6023	0.5594	0.9266	0.1487
Electricity	TabTrans	FGSM	0.3	0.9992	0.4984	0.7017	0.3587	0.5647	0.5245	1.0000	0.1301
Electricity	TabTrans	PGD	0.3	0.9993	0.7520	0.6161	0.3149	0.3795	0.3525	1.0000	0.1202
Electricity	TabTrans	BIM	0.3	0.9996	0.4632	0.5672	0.2899	0.6127	0.5691	1.0000	0.1198
Electricity	TabTrans	C&W	1	0.9989	0.9939	0.0796	0.0407	0.0488	0.0454	0.0990	0.0177
Electricity	TabTrans	DeepFool	1	0.9999	0.3804	0.0356	0.0182	0.0582	0.0541	0.1222	0.0123
Electricity	FTTrans	Gaussian	1	0.4571	0.7491	0.7745	0.3958	0.5995	0.5569	0.9399	0.1491
Electricity	FTTrans	FGSM	0.3	0.9861	0.4983	0.6976	0.3565	0.5649	0.5247	0.9549	0.1291
Electricity	FTTrans	PGD	0.3	0.9015	0.7490	0.6441	0.3292	0.3800	0.3530	0.9547	0.1215
Electricity	FTTrans	BIM	0.3	0.9943	0.4490	0.5273	0.2695	0.6107	0.5672	0.9649	0.1148
Electricity	FTTrans	C&W	1	0.9849	0.9941	0.0674	0.0345	0.0455	0.0422	0.0166	0.0092
Electricity	FTTrans	DeepFool	1	0.9949	0.3839	0.0885	0.0452	0.1510	0.1403	0.1350	0.0259
COMPAS	LR	Gaussian	1	0.3373	0.5432	0.7204	0.3682	0.1202	0.1117	0.7890	0.0681
COMPAS	LR	FGSM	0.3	1.0000	0.5645	1.6790	0.8582	0.1146	0.1064	1.0000	0.0755
COMPAS	LR	PGD	0.3	1.0000	0.5645	1.6686	0.8529	0.1146	0.1064	1.0000	0.0755
COMPAS	LR	BIM	0.3	1.0000	0.5645	1.6790	0.8582	0.1146	0.1064	1.0000	0.0755
COMPAS	LR	C&W	1	0.3545	0.9986	0.0017	0.0009	0.0002	0.0002	0.0822	0.0007
COMPAS	LR	DeepFool	1	0.9940	0.3953	0.2639	0.1349	0.0408	0.0379	0.4530	0.0264
COMPAS	MLP	Gaussian	1	0.3491	0.5431	0.7180	0.3670	0.1194	0.1109	0.7866	0.0677
COMPAS	MLP	FGSM	0.3	0.8414	0.1146	0.6900	0.3527	0.5851	0.5434	1.0000	0.0699
COMPAS	MLP	PGD	0.3	0.9303	0.5463	1.0876	0.5559	0.1173	0.1090	1.0000	0.0729
COMPAS	MLP	BIM	0.3	0.8414	0.1146	0.6900	0.3527	0.5851	0.5434	1.0000	0.0699
COMPAS	MLP	C&W	1	0.4245	0.9444	0.0733	0.0375	0.0236	0.0219	0.0920	0.0123
COMPAS	MLP	DeepFool	1	0.9940	0.0745	0.2722	0.1391	0.3579	0.3324	0.5056	0.0394
COMPAS	TabTrans	Gaussian	1	0.3716	0.5415	0.7232	0.3696	0.1209	0.1123	0.7801	0.0683
COMPAS	TabTrans	FGSM	0.3	0.9820	0.1060	0.6623	0.3385	0.6096	0.5662	0.9997	0.0664
COMPAS	TabTrans	PGD	0.3	0.9988	0.5672	1.0631	0.5434	0.1121	0.1041	1.0000	0.0709
COMPAS	TabTrans	BIM	0.3	0.9970	0.1316	0.6307	0.3223	0.4689	0.4356	1.0000	0.0718
COMPAS	TabTrans	C&W	1	0.4263	0.9987	0.0074	0.0038	0.0010	0.0009	0.0909	0.0014
COMPAS	TabTrans	DeepFool	1	0.9994	0.0726	0.1340	0.0685	0.2171	0.2017	0.2381	0.0270
COMPAS	FTTrans	Gaussian	1	0.4776	0.5394	0.7173	0.3667	0.1198	0.1113	0.7736	0.0677
COMPAS	FTTrans	FGSM	0.3	0.9985	0.1126	0.6855	0.3504	0.5950	0.5526	1.0000	0.0692
COMPAS	FTTrans	PGD	0.3	0.7696	0.5499	1.0582	0.5409	0.1165	0.1082	0.9949	0.0723
COMPAS	FTTrans	BIM	0.3	0.9949	0.1232	0.5849	0.2989	0.4662	0.4330	1.0000	0.0681
COMPAS	FTTrans	C&W	1	0.5074	0.9982	0.0000	0.0000	0.0000	0.0000	0.0628	0.0005
COMPAS	FTTrans	DeepFool	1	0.9991	0.0777	0.1651	0.0844	0.2515	0.2336	0.2913	0.0311

Table B.11
Continued.

Dataset	Model	Attack	ϵ	ASR	SpaR	ℓ_2	ℓ_2^*	SEN	SEN*	OR	IS
Higgs	LR	Gaussian	1	0.4752	0.9336	0.8177	0.4180	0.2425	0.2252	0.9998	0.1127
Higgs	LR	FGSM	0.3	1.0000	0.9362	1.2939	0.6613	0.1836	0.1706	1.0000	0.1063
Higgs	LR	PGD	0.3	1.0000	0.9362	1.2392	0.6334	0.1836	0.1706	1.0000	0.1056
Higgs	LR	BIM	0.3	1.0000	0.9362	1.2939	0.6613	0.1836	0.1706	1.0000	0.1063
Higgs	LR	C&W	1	1.0000	0.9857	0.0585	0.0299	0.0147	0.0137	0.2339	0.0095
Higgs	LR	DeepFool	1	1.0000	0.5986	0.0144	0.0073	0.0050	0.0046	0.1228	0.0032
Higgs	MLP	Gaussian	1	0.4834	0.9349	0.8176	0.4179	0.2414	0.2242	1.0000	0.1124
Higgs	MLP	FGSM	0.3	0.8697	0.9373	1.2987	0.6638	0.1833	0.1703	1.0000	0.1063
Higgs	MLP	PGD	0.3	0.9864	0.9522	1.0700	0.5469	0.1803	0.1675	1.0000	0.1020
Higgs	MLP	BIM	0.3	0.9999	0.8958	1.1412	0.5833	0.1928	0.1790	1.0000	0.1066
Higgs	MLP	C&W	1	0.9998	0.9890	0.0703	0.0359	0.0171	0.0159	0.2470	0.0110
Higgs	MLP	DeepFool	1	0.9996	0.6796	0.0188	0.0096	0.0077	0.0071	0.1499	0.0044
Higgs	TabTrans	Gaussian	1	0.4995	0.9343	0.8193	0.4188	0.2426	0.2253	0.9998	0.1128
Higgs	TabTrans	FGSM	0.3	0.9563	0.9278	1.2783	0.6534	0.1854	0.1723	1.0000	0.1067
Higgs	TabTrans	PGD	0.3	0.9918	0.9976	1.1295	0.5773	0.1719	0.1597	1.0000	0.1005
Higgs	TabTrans	BIM	0.3	0.9982	0.9363	1.2146	0.6208	0.1833	0.1702	1.0000	0.1051
Higgs	TabTrans	C&W	1	0.9956	0.9880	0.0669	0.0342	0.0163	0.0151	0.2489	0.0105
Higgs	TabTrans	DeepFool	1	0.9982	0.6394	0.0164	0.0084	0.0071	0.0066	0.1411	0.0040
Higgs	FTTrans	Gaussian	1	0.4855	0.9340	0.8188	0.4185	0.2422	0.2250	0.9998	0.1126
Higgs	FTTrans	FGSM	0.3	0.8775	0.9344	1.2902	0.6595	0.1840	0.1709	1.0000	0.1064
Higgs	FTTrans	PGD	0.3	0.9919	0.9600	0.9245	0.4726	0.1788	0.1661	0.9995	0.0987
Higgs	FTTrans	BIM	0.3	0.9955	0.8740	0.9650	0.4933	0.1973	0.1833	1.0000	0.1043
Higgs	FTTrans	C&W	1	0.9988	0.9892	0.0625	0.0319	0.0160	0.0148	0.2232	0.0101
Higgs	FTTrans	DeepFool	1	1.0000	0.6784	0.0157	0.0080	0.0065	0.0060	0.1384	0.0038
house_16H	LR	Gaussian	1	0.2967	0.9626	0.7776	0.3975	0.5257	0.4883	1.0000	0.1517
house_16H	LR	FGSM	0.3	0.9952	0.8966	0.9539	0.4876	0.3766	0.3498	1.0000	0.1427
house_16H	LR	PGD	0.3	0.9903	0.8968	0.9142	0.4673	0.3766	0.3498	1.0000	0.1409
house_16H	LR	BIM	0.3	0.9952	0.8966	0.9539	0.4876	0.3766	0.3498	1.0000	0.1427
house_16H	LR	C&W	1	0.7369	0.9929	0.3681	0.1881	0.2080	0.1932	0.6498	0.0770
house_16H	LR	DeepFool	1	0.9993	0.6018	0.3277	0.1675	0.2207	0.2050	0.7418	0.0725
house_16H	MLP	Gaussian	1	0.4742	0.9220	0.7923	0.4050	0.5526	0.5133	1.0000	0.1541
house_16H	MLP	FGSM	0.3	1.0000	0.9472	1.0094	0.5159	0.3517	0.3266	1.0000	0.1420
house_16H	MLP	PGD	0.3	1.0000	0.9299	0.8950	0.4574	0.3584	0.3329	1.0000	0.1379
house_16H	MLP	BIM	0.3	1.0000	0.9163	0.9432	0.4821	0.3641	0.3382	1.0000	0.1407
house_16H	MLP	C&W	1	0.9978	0.9964	0.1374	0.0702	0.0869	0.0807	0.3717	0.0334
house_16H	MLP	DeepFool	1	1.0000	0.8143	0.0675	0.0345	0.0607	0.0564	0.4237	0.0203
house_16H	TabTrans	Gaussian	1	0.3950	0.9279	0.7909	0.4043	0.5520	0.5127	1.0000	0.1541
house_16H	TabTrans	FGSM	0.3	0.9998	0.9068	0.9405	0.4807	0.3699	0.3435	1.0000	0.1413
house_16H	TabTrans	PGD	0.3	1.0000	0.9928	0.9157	0.4680	0.3337	0.3100	1.0000	0.1360
house_16H	TabTrans	BIM	0.3	1.0000	0.9345	0.9447	0.4828	0.3560	0.3307	1.0000	0.1399
house_16H	TabTrans	C&W	1	0.9976	0.9963	0.1642	0.0839	0.1022	0.0950	0.4158	0.0391
house_16H	TabTrans	DeepFool	1	1.0000	0.7428	0.1126	0.0575	0.0999	0.0928	0.3799	0.0315
house_16H	FTTrans	Gaussian	1	0.3704	0.9351	0.7878	0.4027	0.5445	0.5058	1.0000	0.1534
house_16H	FTTrans	FGSM	0.3	0.8479	0.9102	0.9585	0.4899	0.3699	0.3435	1.0000	0.1421
house_16H	FTTrans	PGD	0.3	0.8238	0.9304	0.7085	0.3621	0.3526	0.3275	1.0000	0.1271
house_16H	FTTrans	BIM	0.3	0.7364	0.8775	0.6852	0.3502	0.3466	0.3219	0.9997	0.1237
house_16H	FTTrans	C&W	1	0.9451	0.9812	0.2130	0.1089	0.1473	0.1368	0.4233	0.0507
house_16H	FTTrans	DeepFool	1	0.8466	0.7388	0.2312	0.1182	0.2510	0.2331	0.3901	0.0603
jm1	LR	Gaussian	1	0.1925	0.8952	0.7096	0.3627	0.8589	0.7978	1.0000	0.1635
jm1	LR	FGSM	0.3	0.9573	0.8155	0.6913	0.3533	0.7366	0.6842	1.0000	0.1537
jm1	LR	PGD	0.3	0.9499	0.8165	0.6906	0.3530	0.7362	0.6838	1.0000	0.1537
jm1	LR	BIM	0.3	0.9573	0.8155	0.6913	0.3533	0.7366	0.6842	1.0000	0.1537
jm1	LR	C&W	1	0.1925	0.9993	0.0000	0.0000	0.0000	0.0000	0.1551	0.0005
jm1	LR	DeepFool	1	0.3436	0.3140	0.2277	0.1164	0.3551	0.3299	0.5160	0.0601
jm1	MLP	Gaussian	1	0.6587	0.8882	0.7300	0.3731	0.8799	0.8173	1.0000	0.1661
jm1	MLP	FGSM	0.3	0.9995	0.9534	1.1506	0.5881	0.5868	0.5450	1.0000	0.1793
jm1	MLP	PGD	0.3	0.9995	0.9534	1.1487	0.5871	0.5868	0.5450	1.0000	0.1793
jm1	MLP	BIM	0.3	0.9995	0.9534	1.1505	0.5881	0.5868	0.5450	1.0000	0.1793
jm1	MLP	C&W	1	0.4483	0.9999	0.1335	0.0682	0.1504	0.1397	0.6373	0.0415
jm1	MLP	DeepFool	1	0.9995	0.7997	0.1679	0.0858	0.1195	0.1110	0.8336	0.0437
jm1	TabTrans	Gaussian	1	0.2834	0.8876	0.6921	0.3538	0.8408	0.7810	1.0000	0.1607
jm1	TabTrans	FGSM	0.3	1.0000	0.9004	0.9625	0.4920	0.6091	0.5658	1.0000	0.1695
jm1	TabTrans	PGD	0.3	0.9890	0.9723	0.5614	0.2870	0.4785	0.4445	1.0000	0.1291
jm1	TabTrans	BIM	0.3	1.0000	0.9096	0.3550	0.1814	0.2877	0.2673	1.0000	0.0884
jm1	TabTrans	C&W	1	0.7327	0.9999	0.0497	0.0254	0.0600	0.0558	0.7781	0.0173
jm1	TabTrans	DeepFool	1	1.0000	0.7323	0.0648	0.0331	0.0943	0.0876	0.8071	0.0231
jm1	FTTrans	Gaussian	1	0.4534	0.8971	0.7328	0.3745	0.8681	0.8063	1.0000	0.1663
jm1	FTTrans	FGSM	0.3	0.9196	0.9467	1.1269	0.5760	0.5855	0.5439	1.0000	0.1779
jm1	FTTrans	PGD	0.3	0.9977	0.9607	1.0661	0.5449	0.5543	0.5149	1.0000	0.1721
jm1	FTTrans	BIM	0.3	0.9233	0.9492	1.0364	0.5297	0.5149	0.4782	1.0000	0.1660
jm1	FTTrans	C&W	1	0.4667	0.9999	0.1490	0.0761	0.1553	0.1442	0.6516	0.0447
jm1	FTTrans	DeepFool	1	0.9233	0.7761	0.3426	0.1751	0.2691	0.2500	0.8199	0.0822

Table B.11
Continued.

Dataset	Model	Attack	ϵ	ASR	SpaR	ℓ_2	ℓ_2^*	SEN	SEN*	OR	IS
BreastCancer	LR	Gaussian	1	0.1140	1.0000	0.8718	0.4456	0.0988	0.0918	1.0000	0.0666
BreastCancer	LR	FGSM	0.3	1.0000	0.9977	1.4918	0.7625	0.0743	0.0690	1.0000	0.0568
BreastCancer	LR	PGD	0.3	1.0000	0.9977	1.4039	0.7176	0.0743	0.0690	1.0000	0.0566
BreastCancer	LR	BIM	0.3	1.0000	0.9977	1.4918	0.7625	0.0743	0.0690	1.0000	0.0568
BreastCancer	LR	C&W	1	0.9825	0.9991	0.4157	0.2125	0.0314	0.0292	0.9107	0.0251
BreastCancer	LR	DeepFool	1	1.0000	0.9363	0.3159	0.1615	0.0233	0.0216	0.8860	0.0190
BreastCancer	MLP	Gaussian	1	0.0965	1.0000	0.8610	0.4401	0.0965	0.0896	1.0000	0.0654
BreastCancer	MLP	FGSM	0.3	0.9912	0.9956	1.4714	0.7521	0.0744	0.0691	1.0000	0.0569
BreastCancer	MLP	PGD	0.3	0.9912	0.9976	1.3779	0.7043	0.0743	0.0690	1.0000	0.0565
BreastCancer	MLP	BIM	0.3	1.0000	0.9921	1.4967	0.7650	0.0747	0.0694	1.0000	0.0571
BreastCancer	MLP	C&W	1	0.9035	0.9997	0.4474	0.2287	0.0344	0.0319	0.9223	0.0272
BreastCancer	MLP	DeepFool	1	0.9474	0.9676	0.4797	0.2452	0.0433	0.0402	0.8148	0.0327
BreastCancer	TabTrans	Gaussian	1	0.2368	1.0000	0.8826	0.4511	0.0920	0.0855	1.0000	0.0634
BreastCancer	TabTrans	FGSM	0.3	0.6842	0.9987	1.4568	0.7446	0.0742	0.0689	1.0000	0.0567
BreastCancer	TabTrans	PGD	0.3	0.8333	0.9993	0.9238	0.4722	0.0737	0.0685	1.0000	0.0540
BreastCancer	TabTrans	BIM	0.3	0.6579	0.8644	0.8821	0.4508	0.0807	0.0749	1.0000	0.0570
BreastCancer	TabTrans	C&W	1	0.6667	0.9991	0.2530	0.1293	0.0185	0.0172	0.8421	0.0154
BreastCancer	TabTrans	DeepFool	1	0.7193	0.8890	0.2479	0.1267	0.0282	0.0262	0.6829	0.0212
BreastCancer	FTTrans	Gaussian	1	0.1842	1.0000	0.9038	0.4620	0.0959	0.0891	1.0000	0.0655
BreastCancer	FTTrans	FGSM	0.3	0.3509	0.9975	1.4770	0.7549	0.0743	0.0690	1.0000	0.0568
BreastCancer	FTTrans	PGD	0.3	0.5614	0.9979	0.8884	0.4541	0.0731	0.0679	1.0000	0.0534
BreastCancer	FTTrans	BIM	0.3	0.2281	0.8359	0.5681	0.2904	0.0588	0.0546	0.9231	0.0422
BreastCancer	FTTrans	C&W	1	0.5351	0.9989	0.4630	0.2366	0.0400	0.0371	0.6721	0.0304
BreastCancer	FTTrans	DeepFool	1	0.3509	0.9225	0.7440	0.3803	0.0608	0.0564	0.9250	0.0451
Wine-White	LR	Gaussian	1	0.3235	0.9994	0.8094	0.4137	0.5372	0.4989	1.0000	0.1560
Wine-White	LR	FGSM	0.3	0.9827	0.9997	0.8930	0.4564	0.3227	0.2997	1.0000	0.1332
Wine-White	LR	PGD	0.3	0.9378	0.9997	0.8370	0.4278	0.3227	0.2997	1.0000	0.1306
Wine-White	LR	BIM	0.3	0.9827	0.9997	0.8930	0.4564	0.3227	0.2997	1.0000	0.1332
Wine-White	LR	C&W	1	0.9878	0.9936	0.3592	0.1836	0.1859	0.1727	0.7231	0.0737
Wine-White	LR	DeepFool	1	1.0000	0.6745	0.3166	0.1618	0.2130	0.1978	0.7265	0.0713
Wine-White	MLP	Gaussian	1	0.3939	0.9993	0.8239	0.4211	0.5428	0.5042	1.0000	0.1575
Wine-White	MLP	FGSM	0.3	1.0000	0.9997	0.8665	0.4429	0.3227	0.2997	1.0000	0.1320
Wine-White	MLP	PGD	0.3	1.0000	0.9997	0.7106	0.3632	0.3223	0.2994	1.0000	0.1239
Wine-White	MLP	BIM	0.3	1.0000	0.9761	0.7791	0.3982	0.3312	0.3076	1.0000	0.1287
Wine-White	MLP	C&W	1	1.0000	0.9952	0.1535	0.0784	0.0890	0.0827	0.4265	0.0359
Wine-White	MLP	DeepFool	1	1.0000	0.7468	0.1251	0.0639	0.0841	0.0781	0.4643	0.0317
Wine-White	TabTrans	Gaussian	1	0.3857	0.9993	0.8174	0.4178	0.5450	0.5062	1.0000	0.1573
Wine-White	TabTrans	FGSM	0.3	0.9969	0.9997	0.8463	0.4326	0.3227	0.2997	1.0000	0.1311
Wine-White	TabTrans	PGD	0.3	0.9571	1.0000	0.5375	0.2747	0.3122	0.2900	0.9776	0.1101
Wine-White	TabTrans	BIM	0.3	0.9990	0.8524	0.5702	0.2914	0.3755	0.3488	0.9775	0.1180
Wine-White	TabTrans	C&W	1	0.9980	0.9950	0.1559	0.0797	0.0938	0.0871	0.4755	0.0372
Wine-White	TabTrans	DeepFool	1	1.0000	0.7315	0.1615	0.0825	0.1215	0.1129	0.3449	0.0400
Wine-White	FTTrans	Gaussian	1	0.3429	0.9997	0.8063	0.4121	0.5364	0.4983	1.0000	0.1557
Wine-White	FTTrans	FGSM	0.3	0.9867	0.9995	0.8852	0.4524	0.3227	0.2998	1.0000	0.1328
Wine-White	FTTrans	PGD	0.3	0.9327	0.9998	0.6832	0.3492	0.3215	0.2986	1.0000	0.1221
Wine-White	FTTrans	BIM	0.3	0.9939	0.9417	0.6582	0.3364	0.3071	0.2853	0.9990	0.1174
Wine-White	FTTrans	C&W	1	0.9857	0.9954	0.1579	0.0807	0.1113	0.1034	0.3002	0.0383
Wine-White	FTTrans	DeepFool	1	0.9969	0.7512	0.2762	0.1412	0.2088	0.1940	0.3153	0.0600
Wine-Red	LR	Gaussian	1	0.4063	0.9965	0.8256	0.4220	0.3677	0.3416	0.9923	0.1371
Wine-Red	LR	FGSM	0.3	1.0000	0.9955	0.8850	0.4523	0.2197	0.2040	1.0000	0.1101
Wine-Red	LR	PGD	0.3	1.0000	0.9955	0.8296	0.4240	0.2197	0.2040	1.0000	0.1083
Wine-Red	LR	BIM	0.3	1.0000	0.9955	0.8850	0.4523	0.2197	0.2040	1.0000	0.1101
Wine-Red	LR	C&W	1	1.0000	0.9835	0.1703	0.0871	0.0557	0.0517	0.2156	0.0278
Wine-Red	LR	DeepFool	1	1.0000	0.7188	0.1288	0.0658	0.0418	0.0388	0.2313	0.0218
Wine-Red	MLP	Gaussian	1	0.4313	0.9960	0.8200	0.4191	0.3668	0.3407	0.9855	0.1366
Wine-Red	MLP	FGSM	0.3	1.0000	0.9940	0.8872	0.4535	0.2200	0.2044	1.0000	0.1102
Wine-Red	MLP	PGD	0.3	1.0000	0.9991	0.7989	0.4083	0.2188	0.2032	1.0000	0.1071
Wine-Red	MLP	BIM	0.3	1.0000	0.9946	0.8605	0.4398	0.2199	0.2042	1.0000	0.1094
Wine-Red	MLP	C&W	1	1.0000	0.9844	0.1546	0.0790	0.0570	0.0529	0.1969	0.0270
Wine-Red	MLP	DeepFool	1	1.0000	0.7287	0.1618	0.0827	0.0679	0.0631	0.1563	0.0283
Wine-Red	TabTrans	Gaussian	1	0.4344	0.9974	0.8297	0.4241	0.3701	0.3438	1.0000	0.1379
Wine-Red	TabTrans	FGSM	0.3	0.7938	0.9753	0.8615	0.4404	0.2173	0.2018	0.9882	0.1083
Wine-Red	TabTrans	PGD	0.3	0.8125	0.9983	0.5685	0.2906	0.2108	0.1958	0.9231	0.0944
Wine-Red	TabTrans	BIM	0.3	0.7969	0.8841	0.4695	0.2399	0.1669	0.1550	0.8157	0.0774
Wine-Red	TabTrans	C&W	1	0.7969	0.9811	0.1109	0.0567	0.0412	0.0383	0.1843	0.0203
Wine-Red	TabTrans	DeepFool	1	0.7969	0.6510	0.1385	0.0708	0.0652	0.0606	0.1765	0.0268
Wine-Red	FTTrans	Gaussian	1	0.4281	0.9960	0.8259	0.4222	0.3699	0.3435	1.0000	0.1376
Wine-Red	FTTrans	FGSM	0.3	0.9250	0.9926	0.8800	0.4498	0.2203	0.2047	1.0000	0.1101
Wine-Red	FTTrans	PGD	0.3	0.9250	0.9982	0.6888	0.3521	0.2156	0.2003	0.9730	0.1018
Wine-Red	FTTrans	BIM	0.3	0.9250	0.9509	0.6710	0.3430	0.1935	0.1798	0.9122	0.0945
Wine-Red	FTTrans	C&W	1	0.9938	0.9854	0.1810	0.0925	0.0736	0.0684	0.1792	0.0316
Wine-Red	FTTrans	DeepFool	1	0.9250	0.7067	0.2618	0.1338	0.1308	0.1215	0.3007	0.0492

Table B.11
Continued.

Dataset	Model	Attack	ϵ	ASR	SpaR	ℓ_2	ℓ_2^*	SEN	SEN*	OR	IS
phoneme	LR	Gaussian	1	0.3062	1.0000	0.7827	0.4001	0.6939	0.6445	0.9879	0.1652
phoneme	LR	FGSM	0.3	0.6605	1.0000	0.6507	0.3326	0.3756	0.3489	0.7213	0.1214
phoneme	LR	PGD	0.3	0.5449	1.0000	0.5967	0.3050	0.3746	0.3480	0.6078	0.1140
phoneme	LR	BIM	0.3	0.6605	1.0000	0.6507	0.3326	0.3756	0.3489	0.7213	0.1214
phoneme	LR	C&W	1	0.2905	0.9178	0.0000	0.0000	0.0000	0.0000	0.0669	0.0005
phoneme	LR	DeepFool	1	0.9991	0.7093	0.3903	0.1995	0.3698	0.3435	0.2398	0.0743
phoneme	MLP	Gaussian	1	0.3432	1.0000	0.8085	0.4132	0.7091	0.6586	0.9865	0.1683
phoneme	MLP	FGSM	0.3	1.0000	1.0000	0.6562	0.3354	0.3756	0.3489	0.7123	0.1215
phoneme	MLP	PGD	0.3	0.9954	1.0000	0.5568	0.2846	0.3746	0.3479	0.3931	0.1010
phoneme	MLP	BIM	0.3	1.0000	0.9630	0.6144	0.3141	0.3931	0.3651	0.4838	0.1110
phoneme	MLP	C&W	1	0.4607	0.9627	0.3044	0.1556	0.2441	0.2268	0.3233	0.0671
phoneme	MLP	DeepFool	1	1.0000	0.7872	0.2305	0.1178	0.1866	0.1733	0.0463	0.0273
phoneme	TabTrans	Gaussian	1	0.3414	1.0000	0.7911	0.4044	0.6967	0.6471	0.9837	0.1660
phoneme	TabTrans	FGSM	0.3	0.6337	0.8701	0.5664	0.2895	0.3268	0.3036	0.8248	0.1101
phoneme	TabTrans	PGD	0.3	0.8455	1.0000	0.4353	0.2225	0.3511	0.3261	0.2352	0.0783
phoneme	TabTrans	BIM	0.3	0.8603	0.8417	0.3839	0.1962	0.2920	0.2712	0.1129	0.0534
phoneme	TabTrans	C&W	1	0.3099	0.9475	0.0971	0.0496	0.0760	0.0706	0.1493	0.0241
phoneme	TabTrans	DeepFool	1	0.8557	0.7524	0.3103	0.1586	0.2714	0.2521	0.2908	0.0668
phoneme	FTTrans	Gaussian	1	0.3432	1.0000	0.8006	0.4092	0.6977	0.6480	0.9811	0.1668
phoneme	FTTrans	FGSM	0.3	0.5578	1.0000	0.6559	0.3353	0.3756	0.3489	0.9287	0.1265
phoneme	FTTrans	PGD	0.3	0.7956	1.0000	0.4553	0.2327	0.3584	0.3329	0.2698	0.0836
phoneme	FTTrans	BIM	0.3	0.8131	0.8480	0.4499	0.2299	0.4178	0.3881	0.2856	0.0864
phoneme	FTTrans	C&W	1	0.4191	0.9656	0.1861	0.0951	0.1459	0.1355	0.2075	0.0424
phoneme	FTTrans	DeepFool	1	0.8363	0.7611	0.3481	0.1779	0.2914	0.2707	0.4425	0.0779
MiniBooNE	LR	Gaussian	1	0.2798	0.9978	0.7557	0.3863	0.1493	0.1387	1.0000	0.0852
MiniBooNE	LR	FGSM	0.3	1.0000	0.9977	1.6490	0.8429	0.1380	0.1282	1.0000	0.0915
MiniBooNE	LR	PGD	0.3	1.0000	0.9977	1.6217	0.8289	0.1380	0.1282	1.0000	0.0914
MiniBooNE	LR	BIM	0.3	1.0000	0.9977	1.6490	0.8429	0.1380	0.1282	1.0000	0.0915
MiniBooNE	LR	C&W	1	0.2794	0.9827	0.0100	0.0051	0.0022	0.0020	0.2083	0.0020
MiniBooNE	LR	DeepFool	1	1.0000	0.7701	0.1316	0.0673	0.0481	0.0447	0.7855	0.0256
MiniBooNE	MLP	Gaussian	1	0.3204	0.9998	0.7529	0.3848	0.1477	0.1372	1.0000	0.0846
MiniBooNE	MLP	FGSM	0.3	1.0000	0.9983	1.5766	0.8059	0.1376	0.1278	1.0000	0.0909
MiniBooNE	MLP	PGD	0.3	0.9961	0.9990	1.4559	0.7441	0.1374	0.1276	1.0000	0.0899
MiniBooNE	MLP	BIM	0.3	1.0000	0.9953	1.5173	0.7755	0.1379	0.1281	1.0000	0.0906
MiniBooNE	MLP	C&W	1	0.8172	0.9961	0.0526	0.0269	0.0110	0.0102	0.7517	0.0078
MiniBooNE	MLP	DeepFool	1	1.0000	0.8357	0.0889	0.0455	0.0254	0.0236	0.8508	0.0155
MiniBooNE	TabTrans	Gaussian	1	0.3671	0.9987	0.7553	0.3860	0.1482	0.1376	1.0000	0.0848
MiniBooNE	TabTrans	FGSM	0.3	1.0000	0.9983	1.5165	0.7751	0.1376	0.1278	1.0000	0.0905
MiniBooNE	TabTrans	PGD	0.3	1.0000	0.9999	1.5001	0.7667	0.1371	0.1273	1.0000	0.0901
MiniBooNE	TabTrans	BIM	0.3	1.0000	0.9863	1.4737	0.7532	0.1391	0.1292	1.0000	0.0908
MiniBooNE	TabTrans	C&W	1	0.9470	0.9970	0.0362	0.0185	0.0077	0.0072	0.7906	0.0057
MiniBooNE	TabTrans	DeepFool	1	1.0000	0.8480	0.0464	0.0237	0.0116	0.0108	0.8604	0.0078
MiniBooNE	FTTrans	Gaussian	1	0.3499	0.9998	0.7550	0.3859	0.1478	0.1373	1.0000	0.0847
MiniBooNE	FTTrans	FGSM	0.3	1.0000	0.9981	1.5529	0.7937	0.1377	0.1279	1.0000	0.0908
MiniBooNE	FTTrans	PGD	0.3	0.9970	0.9987	1.3504	0.6902	0.1375	0.1277	1.0000	0.0892
MiniBooNE	FTTrans	BIM	0.3	1.0000	0.9710	0.9056	0.4629	0.1140	0.1058	1.0000	0.0738
MiniBooNE	FTTrans	C&W	1	0.4610	0.9933	0.0305	0.0156	0.0082	0.0076	0.6035	0.0056
MiniBooNE	FTTrans	DeepFool	1	0.9997	0.8382	0.1968	0.1006	0.0668	0.0620	0.8454	0.0356

Appendix C. Trade-off analysis

Figs. C.19–C.21 present correlation summaries, heatmaps, and scatter plots that further illustrate how different attacks balance effectiveness and imperceptibility across datasets and models.




Fig. C.19. Average correlation between Attack Success Rate (ASR) and imperceptibility metrics by attack type. Positive correlations (red bars) indicate a trade-off where higher perturbation magnitudes lead to higher success rates. Negative correlations (green bars) indicate efficient attacks that achieve high success with lower perturbations. Gray bars indicate no clear relationship. Dashed lines mark the weak correlation threshold ($|r| = 0.2$).

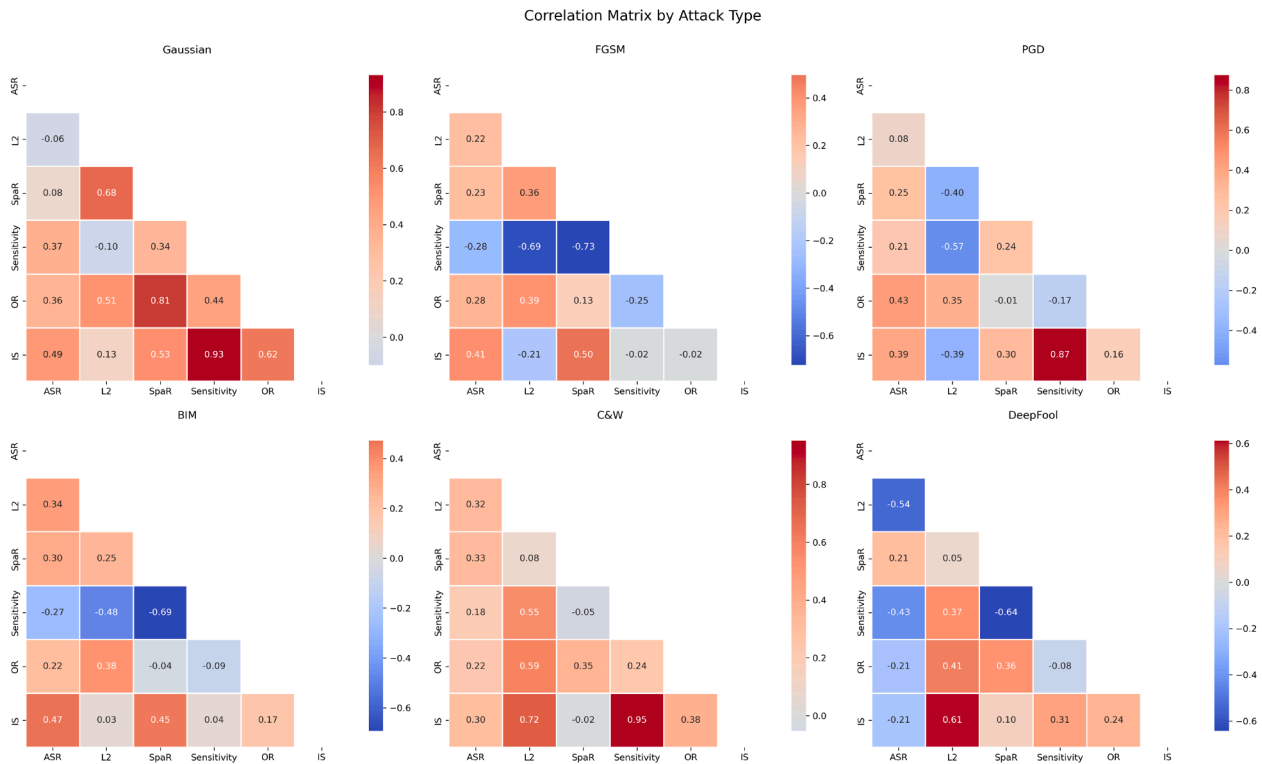


Fig. C.20. Correlation heatmap showing the relationship between Attack Success Rate and five imperceptibility metrics for each attack type. Red cells indicate positive correlations (trade-off), blue cells indicate negative correlations (efficient), and white cells indicate no correlation.

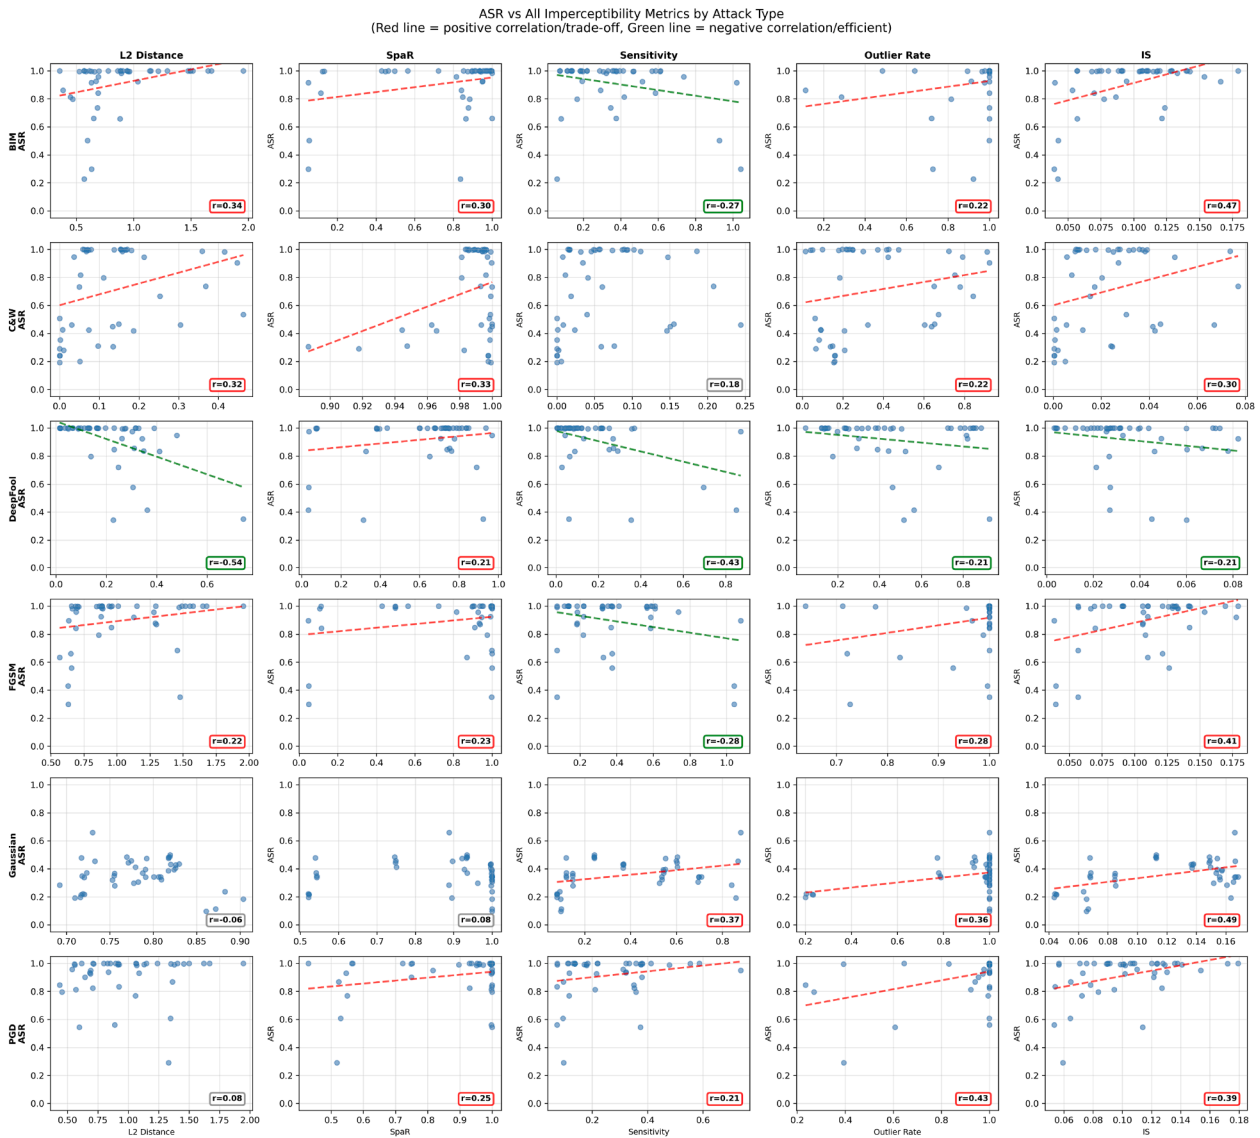


Fig. C.21. Scatter plots of ASR versus imperceptibility metrics by attack type. Trend lines show correlation direction: red (trade-off) or green (efficient). Each attack shows distinct patterns across different imperceptibility dimensions.

References

Abdullah, H., Garcia, W., Peeters, C., Traynor, P., Butler, K. R. B., & Wilson, J. (2019). Practical hidden voice attacks against speech and speaker recognition systems. In *26th annual network and distributed system security symposium, NDSS 2019, San Diego, California, USA, February 24–27, 2019*. <https://doi.org/10.14722/ndss.2019.23362>

An, S., Xiao, C., Stewart, W. F., & Sun, J. (2019). Longitudinal adversarial attack on electronic health records data. In *The world wide web conference, WWW 2019, San Francisco, CA, USA, May 13–17, 2019* (pp. 2558–2564). ACM. <https://doi.org/10.1145/3308558.3313528>

Asl, J. R., Rafiei, M. H., Alohaly, M., & Takabi, D. (2024). A semantic, syntactic, and context-aware natural language adversarial example generator. *IEEE Transactions on Dependable and Secure Computing*, 21(5), 4754–4769. <https://doi.org/10.1109/TDSC.2024.3359817>

Assion, F., Schlicht, P., Greßner, F., Günther, W., Hüger, F., Schmidt, N. M., & Rasheed, U. (2019). The attack generator: A systematic approach towards constructing adversarial attacks. In *IEEE conference on computer vision and pattern recognition workshops, CVPR workshops 2019, Long Beach, CA, USA, June 16–20, 2019* (pp. 1370–1379). Computer Vision Foundation / IEEE. <https://doi.org/10.1109/CVPRW.2019.00177>

Aydin, A., & Temizel, A. (2023). Adversarial image generation by spatial transformation in perceptual colorspace. *Pattern Recognition Letters*, 174, 92–98. <https://doi.org/10.1016/j.patrec.2023.09.003>

Ballet, V., Renard, X., Aigrain, J., Laugel, T., Frossard, P., & Detryniecki, M. (2019). Imperceptible adversarial attacks on tabular data. *arXiv preprint arXiv:1911.03274*.

Biggio, B., Corona, I., Maiorca, D., Nelson, B., Šrđić, N., Laskov, P., Giacinto, G., & Roli, F. (2013). Evasion attacks against machine learning at test time. In *Machine learning and knowledge discovery in databases – european conference, ECML PKDD 2013, Prague, Czech Republic, September 23–27, 2013* (pp. 387–402). Springer. https://doi.org/10.1007/978-3-642-40994-3_25

Biggio, B., & Roli, F. (2018). Wild patterns: Ten years after the rise of adversarial machine learning. *Pattern Recognition*, 84, 317–331. <https://doi.org/10.1016/J.PATCOG.2018.07.023>

Borisov, V., Leemann, T., Seßler, K., Haug, J., Pawelczyk, M., & Kasneci, G. (2022). Deep neural networks and tabular data: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 35(6), 7499–7519. <https://doi.org/10.1109/TNNLS.2022.3229161>

Brendel, W., Rauber, J., & Bethge, M. (2017). Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. *arXiv preprint arXiv:1712.04248*.

Carlini, N., & Wagner, D. A. (2017). Towards evaluating the robustness of neural networks. In *2017 IEEE symposium on security and privacy, SP 2017, San Jose, CA, USA, May 22–26, 2017* (pp. 39–57). IEEE Computer Society. <https://doi.org/10.1109/SP.2017.49>

Cartella, F., Anunção, O., Funabiki, Y., Yamaguchi, D., Akishita, T., & Elshocht, O. (2021). Adversarial attacks for tabular data: Application to fraud detection and imbalanced data. In *Proceedings of the workshop on artificial intelligence safety 2021 (SafeAI 2021) co-located with the thirty-fifth AAAI Conference on Artificial Intelligence (AAAI 2021), virtual, February 8, 2021*. https://ceur-ws.org/Vol-2808/Paper_4.pdf

Chen, P., Zhang, H., Sharma, Y., Yi, J., & Hsieh, C. (2017). ZOO: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *Proceedings of the 10th ACM workshop on artificial intelligence and security, AISec@CCS 2017* (pp. 15–26). ACM <https://doi.org/10.1145/3128572.3140448>

- Chernikova, A., & Oprea, A. (2022). FENCE: Feasible evasion attacks on neural networks in constrained environments. *ACM Transactions on Privacy and Security*, 25(4), 34:1–34:34. <https://doi.org/10.1145/3544746>
- Cinà, A. E., Rony, J., Pintor, M., Demetrio, L., Demontis, A., Biggio, B., Ayed, I. B., & Roli, F. (2025). AttackBench: Evaluating gradient-based attacks for adversarial examples. In *Proceedings of the AAAI conference on artificial intelligence, AAAI-25, February 25 – March 4, 2025, Philadelphia, PA, USA* (pp. 2600–2608). AAAI. <https://doi.org/10.1609/aaai.v39i3.32263>
- Cost, S., & Salzberg, S. (1993). A weighted nearest neighbor algorithm for learning with symbolic features. *Machine Learning*, 10, 57–78. <https://doi.org/10.1023/A:1022664626993>
- Croce, F., Andriushchenko, M., Sehwan, V., Debenedetti, E., Flammarion, N., Chiang, M., Mittal, P., & Hein, M. (2021). RobustBench: A standardized adversarial robustness benchmark. In *Proceedings of the neural information processing systems track on datasets and benchmarks 1, NeurIPS datasets and benchmarks 2021, December 2021, virtual*.
- Dong, Y., Fu, Q., Yang, X., Pang, T., Su, H., Xiao, Z., & Zhu, J. (2020). Benchmarking adversarial robustness on image classification. In *2020 IEEE/CVF Conference on computer vision and pattern recognition, CVPR 2020, Seattle, WA, USA, June 13–19, 2020* (pp. 318–328). <https://doi.org/10.1109/CVPR42600.2020.00040>
- Erdemir, E., Bickford, J., Melis, L., & Aydoore, S. (2021). Adversarial robustness with non-uniform perturbations. In *Advances in neural information processing systems 34: Annual conference on neural information processing systems 2021, NeurIPS 2021, December 6–14, 2021, virtual* (pp. 19147–19159).
- Gao, C., Gu, K., Vosoughi, S., & Mehnaz, S. (2024). Semantic-preserving adversarial example attack against BERT. In *Proceedings of the 4th workshop on trustworthy natural language processing (TrustNLP 2024)* (pp. 202–207). ACL. <https://doi.org/10.18653/v1/2024.trustnlp-1.17>
- Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F. A., & Brendel, W. (2018). Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *7th international conference on learning representations, ICLR 2019, New Orleans, LA, USA, May 6–9, 2019*.
- Goodfellow, I. J., Shlens, J., & Szegedy, C. (2015). Explaining and harnessing adversarial examples. In *3rd international conference on learning representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015*.
- Gorishniy, Y., Rubachev, I., Khrulkov, V., & Babenko, A. (2021). Revisiting deep learning models for tabular data. In *Advances in neural information processing systems 34: Annual conference on neural information processing systems 2021, NeurIPS 2021, December 6–14, 2021, virtual* (pp. 18932–18943).
- Gower, J. C. (1971). A general coefficient of similarity and some of its properties. *Biometrics*, (pp. 857–871). <https://doi.org/10.2307/2528823>
- Gressel, G., Hegde, N., Sreekumar, A., Radhakrishnan, R., Harikumar, K., Achuthan, K. et al. (2021). Feature importance guided attack: A model agnostic adversarial attack. *arXiv preprint arXiv:2106.14815*
- Gunasekaran, M. et al. (2023). Generating and defending against adversarial examples for loan eligibility prediction. In *2023 International conference on system, computation, automation and networking (ICSCAN)* (pp. 1–6). IEEE. <https://doi.org/10.1109/ICSCAN58655.2023.10395585>
- He, Z., Ouyang, C., Alzubaidi, L., Barros, A., & Moreira, C. (2025). Investigating imperceptibility of adversarial attacks on tabular data: An empirical analysis. *Intelligent Systems with Applications*, 25, 200461. <https://doi.org/10.1016/J.ISWA.2024.200461>
- Hingun, N., Sitawarin, C., Li, J., & Wagner, D. (2023). REAP: A large-scale realistic adversarial patch benchmark. In *IEEE/CVF International conference on computer vision, ICCV 2023, Paris, France, October 1–6, 2023* (pp. 4617–4628). IEEE. <https://doi.org/10.1109/ICCV51070.2023.00428>
- Huang, X., Khetan, A., Cvitkovic, M., & Karnin, Z. (2020). Tabtransformer: Tabular data modeling using contextual embeddings. *arXiv preprint arXiv:2012.06678*.
- Jin, Z., Zhang, J., Zhu, Z., & Chen, H. (2024). Benchmarking transferable adversarial attacks. In *Workshop on AI systems with confidential computing (AISCC) 2024, 26 February 2024, San Diego, CA, USA*. <https://doi.org/10.14722/aiscc.2024.23017>
- Khosravi, B., Weston, A. D., Nugen, F., Mickley, J. P., Kremers, H. M., Wyles, C. C., Carter, R. E., & Taunton, M. J. (2023). Demystifying statistics and machine learning in analysis of structured tabular data. *The Journal of Arthroplasty*, 38(10), 1943–1947. <https://doi.org/10.1016/j.arth.2023.08.045>
- Kireev, K., Kulynych, B., & Troncoso, C. (2023). Adversarial robustness for tabular data through cost and utility awareness. In *30th annual network and distributed system security symposium, NDSS 2023, San Diego, California, USA, February 27 – March 3, 2023*. <https://doi.org/10.14722/ndss.2023.24924>
- Ko, K., Kim, S., & Kwon, H. (2023). Multi-targeted audio adversarial example for use against speech recognition systems. *Computers & Security*, 128, 103168. <https://doi.org/10.1016/J.COSE.2023.103168>
- Kurakin, A., Goodfellow, I. J., & Bengio, S. (2017). Adversarial examples in the physical world. In *5th international conference on learning representations, ICLR 2017, Toulon, France, April 24–26, 2017*.
- Le, S. Q., & Ho, T. B. (2005). An association-based dissimilarity measure for categorical data. *Pattern Recognition Letters*, 26(16), 2549–2557. <https://doi.org/10.1016/J.PATREC.2005.06.002>
- Li, L., Lei, J., Gan, Z., & Liu, J. (2021). Adversarial VQA: A new benchmark for evaluating the robustness of VQA models. In *2021 IEEE/CVF International conference on computer vision, ICCV 2021, Montreal, QC, Canada, October 10–17, 2021* (pp. 2022–2031). IEEE. <https://doi.org/10.1109/ICCV48922.2021.00205>
- Lunghi, D., Simitsis, A., Caelen, O., & Bontempi, G. (2023). Adversarial learning in real-world fraud detection: Challenges and perspectives. In *Proceedings of the second ACM data economy workshop, DEC 2023, Seattle, WA, USA, 18 June 2023* (pp. 27–33). ACM. <https://doi.org/10.1145/3600046.3600051>
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., & Vladu, A. (2018). Towards deep learning models resistant to adversarial attacks. In *6th international conference on learning representations, ICLR 2018, Vancouver, BC, Canada, April 30 – May 3, 2018*.
- Mathov, Y., Levy, E., Katzir, Z., Shabtai, A., & Elovici, Y. (2022). Not all datasets are born equal: On heterogeneous tabular data and adversarial examples. *Knowledge-Based Systems*, 242, 108377. <https://doi.org/10.1016/J.KNSYS.2022.108377>
- Moosavi-Dezfooli, S., Fawzi, A., & Frossard, P. (2016). Deepfool: A simple and accurate method to fool deep neural networks. In *2016 IEEE Conference on computer vision and pattern recognition, CVPR 2016, Las Vegas, NV, USA, June 27–30, 2016* (pp. 2574–2582). IEEE Computer Society. <https://doi.org/10.1109/CVPR.2016.282>
- Mushava, J., & Murray, M. (2024). Flexible loss functions for binary classification in gradient-boosted decision trees: An application to credit scoring. *Expert Systems with Applications*, 238, 121876. <https://doi.org/10.1016/J.ESWA.2023.121876>
- Nandy, J., Chauhan, J., Saket, R., & Raghuvver, A. (2023). Non-uniform adversarial perturbations for discrete tabular datasets. In *Proceedings of the 32nd ACM international conference on information and knowledge management* (pp. 1887–1896). ACM. <https://doi.org/10.1145/3583780.3614992>
- Noureddine, K., Kheddar, H., & Maazouz, M. (2023). Adversarial example detection techniques in speech recognition systems: A review. In *2023 2nd international conference on electronics, energy and measurement (IC2EM)*. IEEE. <https://doi.org/10.1109/IC2EM59347.2023.10419688>
- Papernot, N., McDaniel, P. D., Jha, S., Fredrikson, M., Celik, Z. B., & Swami, A. (2016). The limitations of deep learning in adversarial settings. In *IEEE European symposium on security and privacy, EuroS&P 2016, Saarbrücken, Germany, March 21–24, 2016* (pp. 372–387). IEEE. <https://doi.org/10.1109/EUROSP.2016.36>
- Pelekis, S., Koutroubas, T., Blika, A., Berdelis, A., Karakolis, E., Ntanos, C., Spiliotis, E., & Askounis, D. (2025). Adversarial machine learning: A review of methods, tools, and critical industry sectors. *Artificial Intelligence Review*, 58(8), 226. <https://doi.org/10.1007/S10462-025-11147-4>
- Qin, Y., Carlini, N., Cottrell, G. W., Goodfellow, I. J., & Raffel, C. (2019). Imperceptible, robust, and targeted adversarial examples for automatic speech recognition. In *Proceedings of the 36th international conference on machine learning, ICML 2019, 9–15 June 2019, Long Beach, California, USA* (pp. 5231–5240).
- Rocamora, E. A., Wu, Y., Liu, F., Chrysos, G. G., & Cevher, V. (2024). Revisiting character-level adversarial attacks for language models. In *Forty-first international conference on machine learning, ICML 2024, Vienna, Austria, July 21–27, 2024*.
- Sadeghi, K., Banerjee, A., & Gupta, S. K. S. (2020). A system-driven taxonomy of attacks and defenses in adversarial machine learning. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 4(4), 450–467. <https://doi.org/10.1109/TETCI.2020.2968933>
- Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP), San Jose, CA, USA, May 22–26, 2017* (pp. 3–18). IEEE. <https://doi.org/10.1109/SP.2017.41>
- Siddiqui, S. A., Dengel, A., & Ahmed, S. (2020). Benchmarking adversarial attacks and defenses for time-series data. In *Neural information processing – 27th international conference, ICONIP 2020, Bangkok, Thailand, November 23–27, 2020* (pp. 544–554). Springer. https://doi.org/10.1007/978-3-030-63836-8_45
- Simonetto, T., Dyrnishi, S., Ghamizi, S., Cordy, M., & Traon, Y. L. (2022). A unified framework for adversarial attack and defense in constrained feature space. In *Proceedings of the thirty-first international joint conference on artificial intelligence, IJCAI 2022, Vienna, Austria, 23–29 July 2022*. <https://doi.org/10.24963/IJCAI.2022/183>
- Simonetto, T., Ghamizi, S., & Cordy, M. (2024a). Constrained adaptive attack: Effective adversarial attack against deep neural networks for tabular data. In *Advances in neural information processing systems 38: Annual conference on neural information processing systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10–15, 2024*.
- Simonetto, T., Ghamizi, S., & Cordy, M. (2024b). TabularBench: Benchmarking adversarial robustness for tabular deep learning in real-world use-cases. In *Advances in neural information processing systems 38: Annual conference on neural information processing systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10–15, 2024*.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I. J., & Fergus, R. (2014). Intriguing properties of neural networks. In *2nd international conference on learning representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014*.
- Thiyagalangam, J., Shankar, M., Fox, G., & Hey, T. (2022). Scientific machine learning benchmarks. *Nature Reviews Physics*, 4(6), 413–420. <https://doi.org/10.1038/s42254-022-00441-7>
- Wang, J., Li, F., & He, L. (2025). A unified framework for adversarial patch attacks against visual 3d object detection in autonomous driving. *IEEE Transactions on Circuits and Systems for Video Technology*, 35(5), 4949–4962. <https://doi.org/10.1109/TCSVT.2025.3525725>
- Weng, J., Luo, Z., Lin, D., & Li, S. (2024). Comparative evaluation of recent universal adversarial perturbations in image classification. *Computers & Security*, 136, 103576. <https://doi.org/10.1016/J.COSE.2023.103576>
- Yang, Y., Huang, P., Cao, J., Ma, F., Zhang, J., & Li, J. (2023). Quantifying robustness to adversarial word substitutions. In *Machine learning and knowledge discovery in databases: Research track – european conference, ECML PKDD 2023, Turin, Italy, September 18–22, 2023* (pp. 95–112). Springer. https://doi.org/10.1007/978-3-031-43412-9_6
- Yi, Z., Cao, X., Pu, X., Wu, Y., Chen, Z., Khan, A. T., Francis, A., & Li, S. (2023). Fraud detection in capital markets: A novel machine learning approach. *Expert Systems with Applications*, 231, 120760. <https://doi.org/10.1016/J.ESWA.2023.120760>
- Zhang, W. E., Sheng, Q. Z., Alhazmi, A., & Li, C. (2020). Adversarial attacks on deep-learning models in natural language processing: A survey. *ACM Transactions on Intelligent Systems and Technology*, 11(3), 24:1–24:41. <https://doi.org/10.1145/3374217>
- Zheng, M., Yan, X., Zhu, Z., Chen, H., & Wu, B. (2025). BlackboxBench: A comprehensive benchmark of black-box adversarial attacks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(9), 7867–7885. <https://doi.org/10.1109/TPAMI.2025.3574432>

Zheng, Q., Zou, X., Dong, Y., Cen, Y., Yin, D., Xu, J., Yang, Y., & Tang, J. (2021). Graph robustness benchmark: Benchmarking the adversarial robustness of graph machine learning. In *Proceedings of the neural information processing systems track on datasets and benchmarks 1, NeurIPS datasets and benchmarks 2021, December 2021, virtual*.

Zhou, J., Zaidi, N. A., Zhang, Y., & Li, G. (2022). Discretization inspired defence algorithm against adversarial attacks on tabular data. In *Advances in knowledge discovery and data mining - 26th pacific-asia conference, PAKDD 2022, Chengdu, China, May 16-19, 2022* (pp. 367–379). Springer https://doi.org/10.1007/978-3-031-05936-0_29