Novelty Assessment Report

Paper: TabStruct: Measuring Structural Fidelity of Tabular Data
PDF URL: https://openreview.net/pdf?id=XOPH34Extq
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2025-12-27

Abstract

Evaluating tabular generators remains a challenging problem, as the unique causal structural prior of heterogeneous tabular data does not lend itself to intuitive human inspection. Recent work has introduced structural fidelity as a tabular-specific evaluation dimension to assess whether synthetic data complies with the causal structures of real data. However, existing benchmarks often neglect the interplay between structural fidelity and conventional evaluation dimensions, thus failing to provide a holistic understanding of model performance. Moreover, they are typically limited to toy datasets, as quantifying existing structural fidelity metrics requires access to ground-truth causal structures, which are rarely available for real-world datasets. In this paper, we propose a novel evaluation framework that jointly considers structural fidelity and conventional evaluation dimensions. We introduce a new evaluation metric, global utility, which enables the assessment of structural fidelity even in the absence of ground-truth causal structures. In addition, we present TabStruct, a comprehensive evaluation benchmark offering large-scale quantitative analysis on 13 tabular generators from nine distinct categories, across 29 datasets. Our results demonstrate that global utility provides a task-independent, domain-agnostic lens for tabular generator performance. We release the TabStruct benchmark suite, including all datasets, evaluation pipelines, and raw results.

Core Task Landscape

This paper addresses: **Evaluating Structural Fidelity of Synthetic Tabular Data**

A total of **50 papers** were analyzed and organized into a taxonomy with **16 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Evaluation Frameworks and Methodologies**
- **Structural and Relational Fidelity Assessment**
- **Evaluation Metrics and Measurement Approaches**
- **Privacy-Utility Tradeoff Analysis**
- **Generation Methods and Comparative Analysis**
- **Domain-Specific Applications and Challenges**
- **Data-Centric and Preprocessing Approaches**
- **Survey and Review Studies**

Complete Taxonomy Tree

- Evaluating Structural Fidelity of Synthetic Tabular Data Survey Taxonomy
- Evaluation Frameworks and Methodologies
  - Multi-Dimensional Evaluation Frameworks ★ (5 papers)
  - [0] TabStruct: Measuring Structural Fidelity of Tabular Data (Anon et al., 2026) View paper
  - [18] Critical Challenges and Guidelines in Evaluating Synthetic Tabular Data: A Systematic Review (Esnaola Inaki, 2025) View paper
  - [31] Benchmarking Synthetic Tabular Data: A Multi-Dimensional Evaluation Framework (Sidorenko Andrey, 2025) View paper
  - [38] Synthetic tabular data evaluation in the health domain covering resemblance, utility, and privacy dimensions (Mikel Hernadez, 2023) View paper
  - [39] FEST: A Unified Framework for Evaluating Synthetic Tabular Data (Weijie Niu, 2025) View paper
  - Benchmark Suites and Comparative Studies (5 papers)
  - [2] Benchmarking the fidelity and utility of synthetic relational data (Valter Hudovernik, 2024) View paper
  - [5] Evaluation of synthetic data generators on complex tabular data (Oscar Thees, 2024) View paper
  - [22] A Comparative Study of Open-Source Libraries for Synthetic Tabular Data Generation: SDV vs. SynthCity (Gobbo, 2025) View paper
  - [32] Evaluating Generative Models for Tabular Data: Novel Metrics and Benchmarking (Herurkar, 2025) View paper
  - [50] Evaluating Generative Models for Synthetic Tabular Data: A Comparative Analysis of Fidelity, Diversity, and Generalization (Z Mahovac, 2025) View paper
  - Evaluation Tools and Platforms (5 papers)
  - [9] Syntheval: a framework for detailed utility and privacy evaluation of tabular synthetic data (A. D. Lautrup, 2024) View paper
  - [15] Synthcity: a benchmark framework for diverse use cases of tabular synthetic data (Z Qian, 2023) View paper
  - [27] Synthesizers: A Meta-Framework for Generating and Evaluating High-Fidelity Tabular Synthetic Data (Peter Schneider-Kamp, 2024) View paper
  - [29] A novel and fully automated platform for synthetic tabular data generation and validation (H. Rashidi, 2024) View paper
  - [47] How good is your synthetic data? SynthRO, a dashboard to evaluate and benchmark synthetic tabular data. (Gabriele Santangelo, 2025) View paper

- Structural and Relational Fidelity Assessment
  - Inter-Column Relationship Preservation (4 papers)
  - [1] Evaluating Inter-Column Logical Relationships in Synthetic Tabular Data Generation (Long Yunbo, 2025) View paper
  - [8] LLM-TabLogic: Preserving Inter-Column Logical Relationships in Synthetic Tabular Data via Prompt-Guided Latent Diffusion (Long Yunbo, 2025) View paper
  - [10] How Well Does Your Tabular Generator Learn the Structure of Tabular Data? (Jiang, 2025) View paper
  - [23] Preserving logical and functional dependencies in synthetic tabular data (Schultz, 2024) View paper
  - Structural Complexity and Heterogeneity (4 papers)
  - [11] A Robust Metric for Evaluating the Quality of Synthetic Tabular Data (Mihaela Lechner, 2025) View paper
  - [16] STRUC-BENCH: Are Large Language Models Good at Generating Complex Structured Tabular Data? (Xiangru Tang, 2024) View paper
  - [21] Structured Evaluation of Synthetic Tabular Data (Yang, 2024) View paper
  - [36] TabXEval: Why this is a bad table? an eXhaustive rubric for table evaluation (Vihang Pancholi, 2025) View paper
- Evaluation Metrics and Measurement Approaches
  - Statistical Fidelity and Distributional Similarity (6 papers)
  - [3] On the quality of synthetic generated tabular data (Erica Espinosa, 2023) View paper
  - [28] A universal metric for robust evaluation of synthetic tabular data (Vikram S Chundawat, 2022) View paper
  - [33] TabSynDex: A universal metric for robust evaluation of synthetic tabular data (Chundawat, 2022) View paper
  - [43] Quantitative Comparison of Structural and Distributional Properties of Synthetic Tabular Data in Parkinson's Disease (Shahryar Wasif, 2025) View paper
  - [44] A Quantitative Comparison of Structural and Distributional Properties of Synthetic Tabular Data in Parkinson's Disease (Shahryar Wasif, 2025) View paper
  - [45] Synthetic tabular data validation: A divergence-based approach (Patricia A. ApellÃ¡niz, 2024) View paper
  - Utility-Based and Task-Specific Metrics (3 papers)
  - [17] Evaluating Normalizing Flow Model Variants for Supervised Tabular Data: A Comparative Analysis Using the Data Accuracy Score (Pragati Sachdeva, 2025) View paper
  - [30] Evaluating Fidelity and Machine Learning Utility of Synthetic Tabular Data Generated Using Generative Models (Aaditya Kumar Dhaka, 2025) View paper
  - [42] An Index for Assessing the Fidelity of Synthetic Tabular Data in Classification Tasks: TabDSFidelity (Marcos DÃaz Bastida, 2025) View paper
  - Explainability-Aware and Semantic Metrics (2 papers)
  - [14] What's wrong with your synthetic tabular data? using explainable ai to evaluate generative models (Kapar, 2025) View paper
  - [25] SHAP Distance: An Explainability-Aware Metric for Evaluating the Semantic Fidelity of Synthetic Tabular Data (Ke Yu, 2025) View paper
- Privacy-Utility Tradeoff Analysis (3 papers)
  - [12] On the Fidelity-Privacy Tradeoff of Synthetic Cancer Registry Data (Philipp RÃ¶chner, 2024) View paper
  - [20] Scaling While Privacy Preserving: A Comprehensive Synthetic Tabular Data Generation and Evaluation in Learning Analytics (Qinyi Liu, 2024) View paper
  - [34] A Multi-Faceted Evaluation Framework for Assessing Synthetic Data Generated by Large Language Models (Yuan Ye-feng, 2024) View paper
- Generation Methods and Comparative Analysis
  - Deep Generative Models Evaluation (4 papers)
  - [4] Generative adversarial networks vs large language models: a comparative study on synthetic tabular data generation (Austin A. Barr, 2025) View paper
  - [13] Iterative Application of UMAP-Based Algorithms for Fully Synthetic Healthcare Tabular Data Generation (Carla LÃ¡zaro, 2024) View paper
  - [46] Enhancing Small Tabular Clinical Trial Dataset through Hybrid Data Augmentation: Combining SMOTE and WCGAN-GP (Winston Wang, 2023) View paper
  - [48] Comparative Analysis of Generative AI Techniques for Addressing the Tabular Data Generation Problem in Medical Records (S. S. Aravinth, 2023) View paper
  - LLM-Based Generation Approaches (2 papers)
  - [26] A Text-to-tabular Approach to Generate Synthetic Patient Data using LLMs (Margaux TÃ¶rnqvist, 2024) View paper
  - [37] Struc-bench: Are large language models really good at generating complex structured data? (Tang, 2023) View paper
  - Statistical and Traditional Methods (1 papers)
  - [6] Systematic Review of Generative Modelling Tools and Utility Metrics for Fully Synthetic Tabular Data (Tobias Hyrup, 2024) View paper
- Domain-Specific Applications and Challenges
  - Healthcare and Medical Data (1 papers)
  - [35] Multimedia Appendix 1 : STD Evaluation Supplemental Material of the paper entitled â Synthetic Tabular Data Evaluation in the Health Domain Covering Resemblance , Utility and Privacy Dimensions â (Anon et al., 2022) View paper
  - Financial and Transactional Data (1 papers)
  - [24] Improve Fidelity and Utility of Synthetic Credit Card Transaction Time Series from Data-centric Perspective (Wang, 2024) View paper
- Data-Centric and Preprocessing Approaches (3 papers)
  - [19] Reimagining synthetic tabular data generation through data-centric AI: A comprehensive benchmark (Hansen, 2023) View paper
  - [40] Tab-Shapley: Identifying Top-k Tabular Data Quality Insights (Padala, 2025) View paper
  - [49] Engineering MLOps Pipelines With Data Quality: A Case Study on Tabular Datasets in Kaggle (Matteo Pancini, 2025) View paper
- Survey and Review Studies (2 papers)
  - [7] A comprehensive survey of synthetic tabular data generation (Wang Yili, 2025) View paper
  - [41] High-fidelity synthetic data applications for data augmentation (Zhenchen Wang, 2024) View paper

Narrative

Core task: Evaluating structural fidelity of synthetic tabular data. The field has organized itself around several complementary perspectives. Evaluation Frameworks and Methodologies provide overarching systems for assessing synthetic data quality, often combining multiple dimensions such as statistical resemblance, utility, and privacy. Structural and Relational Fidelity Assessment focuses specifically on whether generated tables preserve inter-column dependencies, logical constraints, and relational integrity—issues that simple marginal or distributional checks may miss. Evaluation Metrics and Measurement Approaches develop concrete scoring functions and distance measures, while Privacy-Utility Tradeoff Analysis examines the tension between data protection and downstream usefulness. Generation Methods and Comparative Analysis benchmarks different synthesizers (GANs, diffusion models, large language models), Domain-Specific Applications tackle challenges in healthcare or finance, Data-Centric and Preprocessing Approaches address data quality before generation, and Survey and Review Studies synthesize the landscape. Together, these branches reflect a maturing discipline that balances theoretical rigor with practical deployment concerns.

Recent work highlights the difficulty of capturing complex structural properties beyond univariate statistics. Multi-Dimensional Evaluation[31] and Critical Evaluation Challenges[18] emphasize that no single metric suffices; evaluators must consider fidelity, diversity, and privacy simultaneously. TabStruct[0] sits squarely within the Multi-Dimensional Evaluation Frameworks branch, proposing a structured approach to assess how well synthetic data preserves intricate dependencies and logical relationships. It shares common ground with Synthetic Tabular Quality[3], which also advocates for holistic quality measures, and with Complex Tabular Evaluation[5], which stresses the need to go beyond simple distributional tests. Meanwhile, works like Inter-Column Logical Relationships[1] and Benchmarking Relational Data[2] drill into specific structural aspects—foreign keys, functional dependencies—that TabStruct[0] aims to incorporate into a unified framework. The central open question remains how to balance computational cost, interpretability, and coverage when evaluating increasingly sophisticated generative models across diverse application domains.

Related Works in Same Category

The following **4 sibling papers** share the same taxonomy leaf node with the original paper:

1. Critical Challenges and Guidelines in Evaluating Synthetic Tabular Data: A Systematic Review

Authors: Esnaola Inaki, Nazia Nafis, MartÃnez-PÃ©rez, Ã□lvaro, IÃ±aki Esnaola, et al. (12 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Generating synthetic tabular data can be challenging, however evaluation of their quality is just as challenging, if not more. This systematic review sheds light on the critical importance of rigorous evaluation of synthetic health data to ensure reliability, relevance, and their appropriate use. Based on screening of 1766 papers and a detailed review of 101 papers we identified key challenges, including lack of consensus on evaluation methods, improper use of evaluation metrics, limited input f...

Relationship Analysis

Both papers belong to the Multi-Dimensional Evaluation Frameworks category, addressing the need to assess synthetic tabular data across multiple dimensions including fidelity, utility, and privacy. They overlap in recognizing the importance of comprehensive evaluation beyond single metrics and both emphasize structural/causal considerations in synthetic data assessment. However, the original paper (TabStruct) introduces a novel global utility metric and provides a large-scale benchmark with 13 generators across 29 datasets, while the candidate paper is a systematic review that surveys existing evaluation practices across 101 papers, identifying critical challenges and proposing reporting guidelines rather than introducing new metrics or benchmarks.

---

2. Benchmarking Synthetic Tabular Data: A Multi-Dimensional Evaluation Framework

Authors: Sidorenko Andrey, Platzer, Michael, Andrey Sidorenko, Scriminaci, et al. (11 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Evaluating the quality of synthetic data remains a key challenge for ensuring privacy and utility in data-driven research. In this work, we present an evaluation framework that quantifies how well synthetic data replicates original distributional properties while ensuring privacy. The proposed approach employs a holdout-based benchmarking strategy that facilitates quantitative assessment through low- and high-dimensional distribution comparisons, embedding-based similarity measures, and nearest-...

Relationship Analysis

Both papers belong to the Multi-Dimensional Evaluation Frameworks category, assessing synthetic tabular data across multiple dimensions including fidelity, utility, and privacy. They overlap in their comprehensive evaluation approach and use of holdout-based benchmarking strategies to assess synthetic data quality. However, TabStruct introduces a novel global utility metric specifically for structural fidelity assessment without requiring ground-truth causal structures and provides a large-scale benchmark across 13 generators and 29 datasets, while the candidate paper focuses on a practical open-source framework (mostlyai-qa) emphasizing embedding-based similarity measures and nearest-neighbor distance metrics for privacy-utility tradeoff evaluation.

---

3. Synthetic tabular data evaluation in the health domain covering resemblance, utility, and privacy dimensions

Authors: Mikel Hernadez, Gorka Epelde, Ane Alberdi, Rodrigo Cilla, Debbie Rankin | **Year/Venue**: 2023 | **URL**: View paper

Abstract

Abstract Background â□□Synthetic tabular data generation is a potentially valuable technology with great promise for data augmentation and privacy preservation. However, prior to adoption, an empirical assessment of generated synthetic tabular data is required across dimensions relevant to the target application to determine its efficacy. A lack of standardized and objective evaluation and benchmarking strategy for synthetic tabular data in the health domain has been found in the literature. Obj...

Relationship Analysis

Both papers belong to the Multi-Dimensional Evaluation Frameworks category, assessing synthetic tabular data across multiple dimensions including fidelity, utility, and privacy. The original paper (TabStruct) introduces structural fidelity as a novel dimension with a focus on causal structure preservation and proposes the global utility metric for evaluation without ground-truth causal graphs, while the candidate paper focuses on orchestrating existing resemblance, utility, and privacy metrics into a standardized evaluation pipeline specifically for health domain applications. The key difference is that TabStruct emphasizes causal structural fidelity as a tabular-specific dimension with new metrics, whereas the candidate paper provides a comprehensive orchestration of conventional evaluation methods without introducing structural/causal considerations.

---

4. FEST: A Unified Framework for Evaluating Synthetic Tabular Data

Authors: Weijie Niu, Alberto Celdran, Karoline Siarsky, Alberto Huertas CeldrÃ¡n, Burkhard Stiller | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Synthetic data generation, leveraging generative machine learning techniques, offers a promising approach to mitigating privacy concerns associated with real-world data usage. Synthetic data closely resembles real-world data while maintaining strong privacy guarantees. However, a comprehensive assessment framework is still missing in the evaluation of synthetic data generation, especially when considering the balance between privacy preservation and data utility in synthetic data. This research ...

Relationship Analysis

Both papers belong to the Multi-Dimensional Evaluation Frameworks category, assessing synthetic tabular data across multiple dimensions including fidelity, utility, and privacy. They overlap in evaluating structural/statistical properties, privacy preservation, and machine learning utility of synthetic data. However, TabStruct introduces a novel global utility metric specifically for structural fidelity assessment without requiring ground-truth causal structures and provides a comprehensive benchmark across 29 datasets and 13 generators, while FEST focuses on implementing a unified Python framework (open-source library) that integrates diverse privacy metrics (attack-based and distance-based) with statistical similarity and ML utility assessments, demonstrated on 3 datasets with 6 generation models.

Contributions Analysis

Overall novelty summary. The paper proposes a multi-dimensional evaluation framework for synthetic tabular data that jointly assesses structural fidelity and conventional quality dimensions, introducing a 'global utility' metric that operates without ground-truth causal structures. It resides in the Multi-Dimensional Evaluation Frameworks leaf, which contains five papers including this one. This leaf sits within the broader Evaluation Frameworks and Methodologies branch, indicating a moderately populated research direction focused on holistic assessment approaches rather than isolated metrics or generation methods.

The taxonomy reveals neighboring work in Benchmark Suites and Comparative Studies (five papers) and Evaluation Tools and Platforms (five papers), both addressing systematic evaluation but with different emphases—standardized comparisons versus software implementation. The Structural and Relational Fidelity Assessment branch (two sub-leaves, eight papers total) focuses specifically on inter-column dependencies and heterogeneity, providing complementary depth to the multi-dimensional perspective. The paper bridges these areas by incorporating structural considerations into a comprehensive framework, distinguishing itself from purely statistical or utility-focused approaches in adjacent branches.

Among thirty candidates examined across three contributions, none yielded clear refutations. The global utility metric examined ten candidates with zero refutable overlaps, suggesting potential novelty in enabling structural assessment without ground-truth causal graphs. The joint evaluation framework and TabStruct benchmark each examined ten candidates with similar results. This limited search scope—thirty papers from semantic retrieval—cannot confirm absolute novelty but indicates that within the examined literature, no prior work explicitly combines these specific elements: causal-structure-agnostic structural metrics, multi-dimensional integration, and large-scale benchmarking across thirteen generators.

Based on top-thirty semantic matches and taxonomy positioning, the work appears to occupy a recognizable but not overcrowded niche. The Multi-Dimensional Evaluation Frameworks leaf contains four siblings, suggesting moderate prior activity in holistic assessment approaches. The absence of refutable candidates within this limited scope suggests the specific combination of contributions may be novel, though exhaustive search across the broader fifty-paper taxonomy and beyond would be necessary to confirm originality conclusively.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Global utility metric for structural fidelity assessment

Description: The authors propose global utility, a novel metric that allows evaluation of how well synthetic tabular data preserves causal structures without requiring access to ground-truth causal graphs, addressing a key limitation of existing structural fidelity metrics that only work on toy datasets with known causal structures.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Dependency-aware synthetic tabular data generation

URL: View paper

Brief Assessment

Dependency-aware Generation[53] focuses on preserving functional and logical dependencies in synthetic tabular data through a hierarchical generation framework, not on evaluating structural fidelity without ground-truth causal graphs. The candidate addresses dependency preservation during generation, while the original addresses post-hoc evaluation metrics.

---

2. Improving the generation and evaluation of synthetic data for downstream medical causal inference

URL: View paper

Brief Assessment

Causal Inference Evaluation[51] focuses on evaluating synthetic data for downstream causal inference tasks (treatment effect estimation), not on assessing structural fidelity of tabular data without ground-truth causal graphs. The candidate addresses treatment-specific evaluation in medical contexts, while the original addresses general tabular data structure preservation across diverse domains.

---

3. A Comparative Study of Open-Source Libraries for Synthetic Tabular Data Generation: SDV vs. SynthCity

URL: View paper

Brief Assessment

SDV vs SynthCity[22] focuses on statistical similarity and predictive utility for synthetic tabular data in low-data regimes, not on structural fidelity or causal structure preservation without ground-truth graphs.

---

4. Preserving logical and functional dependencies in synthetic tabular data

URL: View paper

Brief Assessment

Logical Dependencies[23] focuses on preserving functional and logical dependencies in synthetic tabular data using a Q-function measure, not on evaluating structural fidelity without ground-truth causal graphs. The candidate addresses dependency preservation during generation, while the original addresses post-hoc evaluation of causal structure preservation.

---

5. Evaluation of synthetic data generators on complex tabular data

URL: View paper

Brief Assessment

Complex Tabular Evaluation[5] focuses on evaluating synthetic data generators on complex tabular datasets but does not propose a metric for assessing structural fidelity without ground-truth causal graphs. The candidate's context mentions testing 'fidelity of synthetic data' but provides insufficient detail to determine if this addresses the same novelty claim as the original paper's global utility metric.

---

6. A Quantitative Comparison of Structural and Distributional Properties of Synthetic Tabular Data in Parkinson's Disease

URL: View paper

Brief Assessment

Parkinson's Quantitative Comparison[44] focuses on evaluating synthetic tabular data in Parkinson's disease using correlation stability, PCA, and Jensen-Shannon distance metrics. It does not propose a metric for assessing structural fidelity without ground-truth causal graphs, which is the core novelty of the original paper's global utility metric.

---

7. Structured Evaluation of Synthetic Tabular Data

URL: View paper

Brief Assessment

Structured Evaluation[21] focuses on evaluating synthetic tabular data through a spectrum of distributional structures (marginal, pairwise, leave-one-out, full joint) without requiring ground-truth causal graphs. However, it does not specifically address structural fidelity in terms of causal structure preservation as the original paper does with global utility.

---

8. LLM-TabLogic: Preserving Inter-Column Logical Relationships in Synthetic Tabular Data via Prompt-Guided Latent Diffusion

URL: View paper

Brief Assessment

LLM-TabLogic[8] focuses on preserving inter-column logical relationships in synthetic data generation using LLM reasoning and diffusion models, not on developing evaluation metrics for structural fidelity without ground-truth causal graphs.

---

9. Tabularargn: A flexible and efficient auto-regressive framework for generating high-fidelity synthetic data

URL: View paper

Brief Assessment

Tabularargn[52] focuses on auto-regressive synthetic data generation with emphasis on fidelity, efficiency, and flexibility. It does not propose metrics for evaluating structural fidelity without ground-truth causal graphs, which is the core novelty of the original paper's global utility metric.

---

10. Evaluating Fidelity and Machine Learning Utility of Synthetic Tabular Data Generated Using Generative Models

URL: View paper

Brief Assessment

Fidelity ML Utility[30] focuses on evaluating synthetic tabular data using statistical diagnostics and downstream ML performance, but does not address structural fidelity or causal structure preservation without ground-truth graphs. The candidate's evaluation framework is fundamentally different from the original paper's causal-structure-based approach.

Contribution 2: Evaluation framework jointly considering structural fidelity and conventional dimensions

Description: The authors develop a comprehensive evaluation framework that integrates structural fidelity assessment with traditional evaluation dimensions such as density estimation, ML efficacy, and privacy preservation, providing a more holistic understanding of tabular generator performance than prior work.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Generative Models in Computational Pathology: A Comprehensive Survey on Methods, Applications, and Challenges

URL: View paper

Brief Assessment

Computational Pathology Survey[61] focuses on generative models for pathology image synthesis and clinical applications, not on evaluation frameworks for tabular data generators. The domains are fundamentally different.

---

2. Integration Of Machine Learning and Advanced Computing For Optimizing Retail Customer Analytics

URL: View paper

Brief Assessment

Retail Customer Analytics[64] focuses on retail customer analytics using ML and advanced computing for forecasting, personalization, and operations. It does not address tabular data generation, structural fidelity assessment, or synthetic data evaluation frameworks that combine causal structure preservation with density/privacy metrics.

---

3. Preserving privacy and fidelity via Ehrhart theory

URL: View paper

Brief Assessment

Ehrhart Theory[70] focuses on privacy-preserving data synthesis using differential privacy mechanisms and fidelity measures between original and sanitized databases. It does not address evaluation frameworks for tabular data generators that integrate structural fidelity with density estimation, ML efficacy, and privacy preservation metrics as described in the original paper.

---

4. Differentially Private Graph Data Publishing via Feature-Based Community Detection

URL: View paper

Brief Assessment

Feature-Based Community Detection[66] focuses on privacy-preserving graph data publishing using community detection and quadtree-based density estimation. This addresses graph structure preservation in a privacy context, not the holistic evaluation framework for tabular generators combining structural fidelity with ML efficacy and privacy metrics as proposed in the original paper.

---

5. Differentially private learning of structured discrete distributions

URL: View paper

Brief Assessment

Structured Discrete Distributions[67] focuses on differentially private learning of discrete distributions with privacy constraints, not on evaluation frameworks for tabular data generators that integrate structural fidelity with density estimation and privacy preservation metrics.

---

6. Utilizing synthetic data for privacy-preserving AI modeling in radiomics: a case study *

URL: View paper

Brief Assessment

Radiomics Privacy[69] focuses on privacy-preserving synthetic data generation for medical imaging (radiomics), evaluating fidelity using statistical measures like Jensen-Shannon divergence and Hellinger distance. It does not address structural fidelity in the context of causal structures or tabular data evaluation frameworks.

---

7. Grid-Based Decompositions for Spatial Data under Local Differential Privacy

URL: View paper

Brief Assessment

Grid-Based Decompositions[62] focuses on spatial data privacy under local differential privacy using grid-based methods. It does not address tabular data generation evaluation frameworks or structural fidelity assessment for synthetic data.

---

8. TLPP: Deep-Learning-Based Two-Layer Privacy Preserving Mechanism for Protecting Vehicle Trajectory Data

URL: View paper

Brief Assessment

TLPP[63] focuses on privacy-preserving mechanisms for vehicle trajectory data using GANs and differential privacy, not on evaluation frameworks for tabular data generators that integrate structural fidelity with density estimation and privacy metrics.

---

9. SynthVal: A Framework for Validating Synthetic Medical Images

URL: View paper

Brief Assessment

SynthVal[65] focuses on validating synthetic medical images, not tabular data generators. The candidate addresses image synthesis evaluation, while the original paper develops a framework specifically for tabular data generation with causal structural models.

---

10. OTTER: Optimized Training with Trustworthy Enhanced Replication via Diffusion and Federated VMUNet for Privacy-Aware Medical Segmentation

URL: View paper

Brief Assessment

OTTER[68] focuses on privacy-preserving medical image segmentation using diffusion and federated learning, not on tabular data generation evaluation frameworks combining structural fidelity with density estimation and privacy preservation.

Contribution 3: TabStruct benchmark suite

Description: The authors introduce TabStruct, a large-scale benchmark that evaluates 13 tabular generators across 29 datasets with multiple evaluation dimensions, addressing the limited scope of existing benchmarks and providing datasets, evaluation pipelines, and raw results as open resources.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Tabular and latent space synthetic data generation: a literature review

URL: View paper

Brief Assessment

Latent Space Review[55] focuses on synthetic data generation methods and taxonomies, not on benchmarking tabular generators across datasets with structural fidelity metrics. The review does not present a benchmark suite comparable to TabStruct.

---

2. Comparison of tabular synthetic data generation techniques using propensity and cluster log metric

URL: View paper

Brief Assessment

Propensity Cluster Metric[60] focuses on comparing synthetic data generation techniques using propensity and cluster-log metrics on various dataset types, not on providing a comprehensive benchmark suite with multiple evaluation dimensions and open resources like TabStruct.

---

3. Scaling While Privacy Preserving: A Comprehensive Synthetic Tabular Data Generation and Evaluation in Learning Analytics

URL: View paper

Brief Assessment

Privacy Preserving Learning[20] focuses on privacy-preserving synthetic data generation for learning analytics with three datasets, not on comprehensive benchmarking of tabular generators across structural fidelity dimensions.

---

4. Deep neural networks and tabular data: A survey

URL: View paper

Brief Assessment

Deep Neural Tabular[56] focuses on deep learning methods for tabular data inference and generation tasks, providing an empirical comparison across five datasets. It does not present a benchmark specifically designed to evaluate tabular generators across structural fidelity dimensions or multiple evaluation metrics as TabStruct does.

---

5. A novel and fully automated platform for synthetic tabular data generation and validation

URL: View paper

Brief Assessment

Automated Platform[29] focuses on synthetic tabular data generation and validation for healthcare applications, not on benchmarking tabular generators across multiple datasets and evaluation dimensions as TabStruct does.

---

6. Modeling tabular data using conditional gan

URL: View paper

Brief Assessment

Conditional GAN[59] focuses on generating synthetic tabular data using GANs with mode-specific normalization and conditional generation, not on creating comprehensive benchmarks for evaluating tabular generators across multiple datasets and metrics.

---

7. Systematic assessment of tabular data synthesis

URL: View paper

Brief Assessment

Systematic Assessment[58] evaluates 8 synthesizers on 12 datasets with focus on privacy-preserving synthesis, while the original paper evaluates 13 generators on 29 datasets with emphasis on structural fidelity and causal structure preservation—distinct evaluation scopes and objectives.

---

8. Comprehensive evaluation framework for synthetic tabular data in health: fidelity, utility and privacy analysis of generative models with and without privacy guarantees

URL: View paper

Brief Assessment

Health Comprehensive Framework[54] focuses on evaluating synthetic tabular data in healthcare with emphasis on fidelity, utility, and privacy analysis. While both frameworks evaluate tabular data generators, Health Comprehensive Framework[54] is specifically designed for health domain applications and does not challenge the novelty of TabStruct's large-scale benchmark across 29 datasets with 13 generators from nine categories.

---

9. How Well Does Your Tabular Generator Learn the Structure of Tabular Data?

URL: View paper

Brief Assessment

Structure Learning[10] focuses on evaluating structural fidelity of tabular generators using causal structures, while the original paper's TabStruct benchmark encompasses broader evaluation dimensions (density estimation, privacy preservation, ML efficacy, and structural fidelity) across 29 datasets with 13 generators from 9 categories. The candidate's scope is limited to 7 datasets with expert-validated structures and 9 generators from 8 categories, representing a smaller-scale evaluation framework.

---

10. A comprehensive evaluation framework for synthetic medical tabular data generation

URL: View paper

Brief Assessment

Medical Tabular Framework[57] focuses on evaluating synthetic medical tabular data generators, not general tabular data generators across diverse domains. The scope and domain differ from TabStruct's broader benchmark.

Appendix: Text Similarity Detection

Textual similarity detection checked 34 papers and found 3 similarity segment(s) across 1 paper(s).

The following **1 paper(s)** were detected to have high textual similarity with the original paper. These may represent different versions of the same work, duplicate submissions, or papers with substantial textual overlap. Readers are advised to verify these relationships independently.

1. How Well Does Your Tabular Generator Learn the Structure of Tabular Data?

Detected in: Contribution: contribution_3

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

References

- [0] TabStruct: Measuring Structural Fidelity of Tabular Data View paper
- [1] Evaluating Inter-Column Logical Relationships in Synthetic Tabular Data Generation View paper
- [2] Benchmarking the fidelity and utility of synthetic relational data View paper
- [3] On the quality of synthetic generated tabular data View paper
- [4] Generative adversarial networks vs large language models: a comparative study on synthetic tabular data generation View paper
- [5] Evaluation of synthetic data generators on complex tabular data View paper
- [6] Systematic Review of Generative Modelling Tools and Utility Metrics for Fully Synthetic Tabular Data View paper
- [7] A comprehensive survey of synthetic tabular data generation View paper
- [8] LLM-TabLogic: Preserving Inter-Column Logical Relationships in Synthetic Tabular Data via Prompt-Guided Latent Diffusion View paper
- [9] Syntheval: a framework for detailed utility and privacy evaluation of tabular synthetic data View paper
- [10] How Well Does Your Tabular Generator Learn the Structure of Tabular Data? View paper
- [11] A Robust Metric for Evaluating the Quality of Synthetic Tabular Data View paper
- [12] On the Fidelity-Privacy Tradeoff of Synthetic Cancer Registry Data View paper
- [13] Iterative Application of UMAP-Based Algorithms for Fully Synthetic Healthcare Tabular Data Generation View paper
- [14] What's wrong with your synthetic tabular data? using explainable ai to evaluate generative models View paper
- [15] Synthcity: a benchmark framework for diverse use cases of tabular synthetic data View paper
- [16] STRUC-BENCH: Are Large Language Models Good at Generating Complex Structured Tabular Data? View paper
- [17] Evaluating Normalizing Flow Model Variants for Supervised Tabular Data: A Comparative Analysis Using the Data Accuracy Score View paper
- [18] Critical Challenges and Guidelines in Evaluating Synthetic Tabular Data: A Systematic Review View paper
- [19] Reimagining synthetic tabular data generation through data-centric AI: A comprehensive benchmark View paper
- [20] Scaling While Privacy Preserving: A Comprehensive Synthetic Tabular Data Generation and Evaluation in Learning Analytics View paper

- [21] Structured Evaluation of Synthetic Tabular Data View paper
- [22] A Comparative Study of Open-Source Libraries for Synthetic Tabular Data Generation: SDV vs. SynthCity View paper
- [23] Preserving logical and functional dependencies in synthetic tabular data View paper
- [24] Improve Fidelity and Utility of Synthetic Credit Card Transaction Time Series from Data-centric Perspective View paper
- [25] SHAP Distance: An Explainability-Aware Metric for Evaluating the Semantic Fidelity of Synthetic Tabular Data View paper
- [26] A Text-to-tabular Approach to Generate Synthetic Patient Data using LLMs View paper
- [27] Synthesizers: A Meta-Framework for Generating and Evaluating High-Fidelity Tabular Synthetic Data View paper
- [28] A universal metric for robust evaluation of synthetic tabular data View paper
- [29] A novel and fully automated platform for synthetic tabular data generation and validation View paper
- [30] Evaluating Fidelity and Machine Learning Utility of Synthetic Tabular Data Generated Using Generative Models View paper
- [31] Benchmarking Synthetic Tabular Data: A Multi-Dimensional Evaluation Framework View paper
- [32] Evaluating Generative Models for Tabular Data: Novel Metrics and Benchmarking View paper
- [33] TabSynDex: A universal metric for robust evaluation of synthetic tabular data View paper
- [34] A Multi-Faceted Evaluation Framework for Assessing Synthetic Data Generated by Large Language Models View paper
- [35] Multimedia Appendix 1 : STD Evaluation Supplemental Material of the paper entitled â Synthetic Tabular Data Evaluation in the Health Domain Covering Resemblance , Utility and Privacy Dimensions â View paper
- [36] TabXEval: Why this is a bad table? an eXhaustive rubric for table evaluation View paper
- [37] Struc-bench: Are large language models really good at generating complex structured data? View paper
- [38] Synthetic tabular data evaluation in the health domain covering resemblance, utility, and privacy dimensions View paper
- [39] FEST: A Unified Framework for Evaluating Synthetic Tabular Data View paper
- [40] Tab-Shapley: Identifying Top-k Tabular Data Quality Insights View paper
- [41] High-fidelity synthetic data applications for data augmentation View paper
- [42] An Index for Assessing the Fidelity of Synthetic Tabular Data in Classification Tasks: TabDSFidelity View paper
- [43] Quantitative Comparison of Structural and Distributional Properties of Synthetic Tabular Data in Parkinson's Disease View paper
- [44] A Quantitative Comparison of Structural and Distributional Properties of Synthetic Tabular Data in Parkinson's Disease View paper
- [45] Synthetic tabular data validation: A divergence-based approach View paper
- [46] Enhancing Small Tabular Clinical Trial Dataset through Hybrid Data Augmentation: Combining SMOTE and WCGAN-GP View paper
- [47] How good is your synthetic data? SynthRO, a dashboard to evaluate and benchmark synthetic tabular data. View paper
- [48] Comparative Analysis of Generative AI Techniques for Addressing the Tabular Data Generation Problem in Medical Records View paper
- [49] Engineering MLOps Pipelines With Data Quality: A Case Study on Tabular Datasets in Kaggle View paper
- [50] Evaluating Generative Models for Synthetic Tabular Data: A Comparative Analysis of Fidelity, Diversity, and Generalization View paper
- [51] Improving the generation and evaluation of synthetic data for downstream medical causal inference View paper
- [52] Tabularargn: A flexible and efficient auto-regressive framework for generating high-fidelity synthetic data View paper
- [53] Dependency-aware synthetic tabular data generation View paper
- [54] Comprehensive evaluation framework for synthetic tabular data in health: fidelity, utility and privacy analysis of generative models with and without privacy guarantees View paper
- [55] Tabular and latent space synthetic data generation: a literature review View paper
- [56] Deep neural networks and tabular data: A survey View paper
- [57] A comprehensive evaluation framework for synthetic medical tabular data generation View paper
- [58] Systematic assessment of tabular data synthesis View paper
- [59] Modeling tabular data using conditional gan View paper
- [60] Comparison of tabular synthetic data generation techniques using propensity and cluster log metric View paper
- [61] Generative Models in Computational Pathology: A Comprehensive Survey on Methods, Applications, and Challenges View paper
- [62] Grid-Based Decompositions for Spatial Data under Local Differential Privacy View paper
- [63] TLPP: Deep-Learning-Based Two-Layer Privacy Preserving Mechanism for Protecting Vehicle Trajectory Data View paper
- [64] Integration Of Machine Learning and Advanced Computing For Optimizing Retail Customer Analytics View paper
- [65] SynthVal: A Framework for Validating Synthetic Medical Images View paper
- [66] Differentially Private Graph Data Publishing via Feature-Based Community Detection View paper
- [67] Differentially private learning of structured discrete distributions View paper
- [68] OTTER: Optimized Training with Trustworthy Enhanced Replication via Diffusion and Federated VMUNet for Privacy-Aware Medical Segmentation View paper
- [69] Utilizing synthetic data for privacy-preserving AI modeling in radiomics: a case study * View paper
- [70] Preserving privacy and fidelity via Ehrhart theory View paper