

# Novelty Assessment Report

**Paper:** xRFM: Accurate, scalable, and interpretable feature learning models for tabular data

**PDF URL:** <https://openreview.net/pdf?id=wHuVdpnUFp>

**Venue:** ICLR 2026 Conference Submission

**Year:** 2026

**Report Generated:** 2026-01-01

## Abstract

Inference from tabular data, collections of continuous and categorical variables organized into matrices, is a foundation for modern technology and science. Yet, in contrast to the explosive changes in the rest of AI, the best practice for these predictive tasks has been relatively unchanged and is still primarily based on variations of Gradient Boosted Decision Trees (GBDTs). Very recently, there has been renewed interest in developing state-of-the-art methods for tabular data based on recent developments in neural networks and feature learning methods. In this work, we introduce xRFM, an algorithm that combines feature learning kernel machines with a tree structure to both adapt to the local structure of the data and scale to essentially unlimited amounts of training data. We show that compared to \$31\$ other methods, including recently introduced tabular foundation models (TabPFN-v2) and GBDTs, xRFM achieves best performance across \$100\$ regression datasets and is competitive to the best methods across \$200\$ classification datasets outperforming GBDTs. Additionally, xRFM provides interpretability natively through the Average Gradient Outer Product.

### Disclaimer

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: mingzhang23@m.fudan.edu.cn

## Core Task Landscape

This paper addresses: **Predictive Modeling on Tabular Data**

A total of **50 papers** were analyzed and organized into a taxonomy with **17 categories**.

### Taxonomy Overview

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

- **Deep Learning Architectures for Tabular Data**
- **Foundation Models and Transfer Learning**
- **Generative Modeling and Data Augmentation**
- **Specialized Prediction Tasks and Applications**
- **Benchmarking, Evaluation, and Comparative Studies**
- **Survey and Review Literature**
- **Time Series Forecasting on Tabular Data**
- **Structured Prediction and Theoretical Frameworks**

### Complete Taxonomy Tree

- Predictive Modeling on Tabular Data Survey Taxonomy
- Deep Learning Architectures for Tabular Data
  - Attention-Based and Transformer Architectures (4 papers)
  - [3] Deep neural networks and tabular data: A survey (Borisov, 2022) [View paper](#)
  - [20] A survey on deep tabular learning (Das Subasish, 2024) [View paper](#)
  - [23] Heart disease prediction by tabular modeling with deep learning network and interpretability (Mohammad H Alshayegi, 2025) [View paper](#)
  - [47] Deep learning with tabular data: A self-supervised approach (Tirth Kiranbhai Vyas, 2024) [View paper](#)
  - Specialized Neural Network Designs ★ (4 papers)
  - [0] xRFM: Accurate, scalable, and interpretable feature learning models for tabular data (Anon et al., 2026) [View paper](#)
  - [16] Low-cost and efficient prediction hardware for tabular data using tiny classifier circuits (Konstantinos Iordanou, 2024) [View paper](#)
  - [18] Regularization learning networks: deep learning for tabular datasets (Shavitt, 2018) [View paper](#)
  - [22] An interpretable prototype parts-based neural network for medical tabular data (J Karolczak, 2025) [View paper](#)
  - Graph Neural Networks for Tabular Data (2 papers)
  - [27] Graph neural networks for tabular data learning: A survey with taxonomy and directions (Cheng-Te Li, 2025) [View paper](#)
  - [48] Retrieval & interaction machine for tabular data prediction (Jiarui Qin, 2021) [View paper](#)
- Foundation Models and Transfer Learning
  - Tabular Foundation Models (2 papers)
  - [1] Accurate predictions on small data with a tabular foundation model (Noah Hollmann, 2025) [View paper](#)
  - [13] Anypredict: Foundation model for tabular prediction (Wang Zifeng, 2023) [View paper](#)
  - Cross-Table Transfer Learning (3 papers)
  - [24] TLTD: Transfer learning for tabular data (Maxim Bragilovski, 2023) [View paper](#)
  - [26] Transtab: Learning transferable tabular transformers across tables (Wang Zifeng, 2022) [View paper](#)
  - [50] Unitabe: A universal pretraining protocol for tabular foundation model in data science (Yang Yazheng, 2023) [View paper](#)
  - LLM-Based Approaches for Tabular Prediction (7 papers)
  - [8] Large Language Models (LLMs) on Tabular Data: Prediction, Generation, and Understanding--A Survey (Fang Xi, 2024) [View paper](#)

- [11] Language modeling on tabular data: A survey of foundations, techniques and evolution (Ruan, 2024) [View paper](#)
- [14] Large scale transfer learning for tabular data via language modeling (Josh Gardner, 2024) [View paper](#)
- [17] Unipredict: Large language models are universal tabular predictors (WANG Ruiyu, 2023) [View paper](#)
- [31] Unleashing the potential of large language models for predictive tabular tasks in data science (Yang Yazheng, 2024) [View paper](#)
- [39] Unipredict: Large language models are universal tabular classifiers (WANG Ruiyu, 2023) [View paper](#)
- [44] Quantifying Prediction Consistency Under Fine-tuning Multiplicity in Tabular LLMs (Hamman, 2024) [View paper](#)
- Generative Modeling and Data Augmentation
  - Generative Models for Tabular Synthesis (3 papers)
  - [2] Generative models for tabular data: A review (Dong Keon Kim, 2024) [View paper](#)
  - [4] Goggle: Generative modelling for tabular data by learning relational structure (T Liu, 2023) [View paper](#)
  - [46] Data Augmentation of a Corrosion Dataset for Defect Growth Prediction of Pipelines Using Conditional Tabular Generative Adversarial Networks (Haonan Ma, 2024) [View paper](#)
  - Generative Pre-Training for Prediction (2 papers)
  - [6] Tabbie: Pretrained representations of tabular data (Iida Hiroshi, 2021) [View paper](#)
  - [30] Generative table pre-training empowers models for tabular prediction (Jian Li, 2023) [View paper](#)
- Specialized Prediction Tasks and Applications
  - Missing Value Imputation and Forecasting (1 papers)
  - [9] Diffusion-based time series imputation and forecasting with structured state space models (Alcaraz, 2022) [View paper](#)
  - Multimodal Tabular Prediction (3 papers)
  - [25] Incorporating key features from structured and unstructured data for enhanced carbon trading price forecasting with interpretability analysis (Meiqin Jiang, 2025) [View paper](#)
  - [29] Hyperfusion: A hypernetwork approach to multimodal integration of tabular and medical imaging data for predictive modeling (Daniel Duenias, 2025) [View paper](#)
  - [36] Combining structured and unstructured data for predictive models: a deep learning approach (Dongdong Zhang, 2020) [View paper](#)
  - Domain-Specific Tabular Applications (3 papers)
  - [12] Explaining customer churn prediction in telecom industry using tabular machine learning models (S. S. Poudel, 2024) [View paper](#)
  - [32] A systematic review of machine learning models for predicting outcomes of stroke with structured data (Wenjuan Wang, 2020) [View paper](#)
  - [34] Advances in AI and machine learning for predictive medicine (Alok Sharma, 2024) [View paper](#)
  - Retrieval-Augmented Methods (1 papers)
  - [38] Retrieval-augmented forecasting with tabular time series data (Zichao Li, 2025) [View paper](#)
- Benchmarking, Evaluation, and Comparative Studies
  - Deep Learning vs. Traditional Methods (3 papers)
  - [7] When do neural nets outperform boosted trees on tabular data? (McElfresh, 2023) [View paper](#)
  - [21] Is deep learning on tabular data enough? An assessment (Sheikh Amir Fayaz, 2022) [View paper](#)
  - [35] Tabular data: Deep learning is not all you need (Shwartz-Ziv, 2022) [View paper](#)
  - Benchmark Datasets and Evaluation Protocols (2 papers)
  - [33] Annotatetables: A large tabular dataset with language model annotations (Hu, 2024) [View paper](#)
  - [37] Tabred: Analyzing pitfalls and filling the gaps in tabular deep learning benchmarks (Rubachev, 2024) [View paper](#)
- Survey and Review Literature (2 papers)
  - [5] Representation learning for tabular data: A comprehensive survey (Liu Si-Yang, 2025) [View paper](#)
  - [10] Recent deep learning methods for tabular data (Yejin Hwang, 2023) [View paper](#)
- Time Series Forecasting on Tabular Data (3 papers)
  - [28] Evaluating the fidelity of statistical forecasting and predictive intelligence by utilizing a stochastic dataset (Mohammad Shahin, 2025) [View paper](#)
  - [42] Time series forecasting methods and their applications to particle accelerators (Li, 2023) [View paper](#)
  - [49] Easytime: Time series forecasting made easy (Xiangfei Qiu, 2025) [View paper](#)
- Structured Prediction and Theoretical Frameworks (6 papers)
  - [15] Predicting structured data (D. Touretzky, 2007) [View paper](#)
  - [19] AixelNet: A Pre-trained Model with Table-aware Adaptation for Structured Data Prediction (L Wang, 2025) [View paper](#)
  - [40] Demystifying statistics and machine learning in analysis of structured tabular data (Bardia Khosravi, 2023) [View paper](#)
  - [41] Deep learning with structured data (Mark, 2020) [View paper](#)
  - [43] AlphaML: a clear, legible, explainable, transparent, and elucidative binary classification platform for tabular data (Ahmad Nasimian, 2024) [View paper](#)
  - [45] Predictive analytics with data visualization (Selvaraj Boopathy, 2022) [View paper](#)

## Narrative

Core task: predictive modeling on tabular data. The field has evolved into several major branches that reflect both methodological diversity and application-driven specialization. Deep learning architectures for tabular data explore neural network designs tailored to structured features, often contrasting with traditional gradient-boosted trees. Foundation models and transfer learning investigate how pretrained representations can generalize across datasets, while generative modeling and data augmentation address data scarcity through synthetic sample creation. Specialized prediction tasks span domains from healthcare to finance, and benchmarking studies rigorously compare methods across standardized datasets. Survey and review literature synthesizes these developments, time series forecasting extends tabular methods to temporal data, and structured prediction frameworks provide theoretical grounding. Representative works include Deep Neural Tabular Survey[3] and Tabular Representation Survey[5], which map out architectural trends, and Generative Tabular Review[2], which examines synthetic data strategies.

Within the deep learning architectures branch, specialized neural network designs have emerged to handle tabular data's unique challenges—heterogeneous feature types, missing values, and the need for interpretability. Some approaches like Goggle[4] and Neural Nets Boosted Trees[7] blend neural components with tree-based methods, while others such as Tiny Classifier Circuits[16] and Regularization Learning Networks[18] focus on compact or regularized architectures. The original paper xRFM[0] sits within this specialized design cluster, emphasizing novel network structures that balance expressiveness with tabular data constraints. Compared to neighbors like Prototype Parts Medical[22], which targets interpretable medical predictions, xRFM[0] appears more focused on general-

purpose architectural innovation. This line of work continues to grapple with whether deep learning can consistently outperform classical methods, a debate highlighted by Deep Learning Enough[21] and Deep Learning Not Enough[35], and whether specialized designs can bridge the performance gap observed in many tabular benchmarks.

---

## Related Works in Same Category

The following **3 sibling papers** share the same taxonomy leaf node with the original paper:

### 1. Low-cost and efficient prediction hardware for tabular data using tiny classifier circuits

**Authors:** Konstantinos Jordanou, Timothy Atkinson, Emre Ozer, Jędrzej Kufel, Grace Aligada, et al. (9 authors total) | **Year/Venue:** 2024 • Nature Electronics | **URL:** [View paper](#)

#### Abstract

N/A

#### Relationship Analysis

Both papers belong to the Specialized Neural Network Designs category, focusing on custom architectures for tabular data prediction. While xRFM combines feature learning kernel machines with tree structures for scalability and interpretability through AGOP, the candidate paper focuses on hardware-efficient prediction using tiny classifier circuits, emphasizing low-cost deployment rather than feature learning mechanisms. The key difference lies in xRFM's emphasis on adaptive feature learning and interpretability versus the candidate's focus on hardware efficiency and resource-constrained deployment.

---

### 2. Regularization learning networks: deep learning for tabular datasets

**Authors:** Shavitt, Ira, Segal, Eran | **Year/Venue:** 2018 | **URL:** [View paper](#)

#### Abstract

Despite their impressive performance, Deep Neural Networks (DNNs) typically underperform Gradient Boosting Trees (GBTs) on many tabular-dataset learning tasks. We propose that applying a different regularization coefficient to each weight might boost the performance of DNNs by allowing them to make more use of the more relevant inputs. However, this will lead to an intractable number of hyperparameters. Here, we introduce Regularization Learning Networks (RLNs), which overcome this challenge by ...

#### Relationship Analysis

Both papers belong to the Specialized Neural Network Designs category, developing custom neural architectures for tabular data with domain-specific inductive biases. They overlap in addressing the challenge of making deep learning competitive with tree-based methods on tabular data through specialized regularization and architectural choices. However, the original paper (xRFM) combines feature learning kernel machines with tree-based data partitioning and uses the Average Gradient Outer Product (AGOP) for interpretability, while the candidate paper (RLN) introduces a Counterfactual Loss to learn weight-specific regularization coefficients that produce extremely sparse networks (eliminating 99.8% of edges) for improved performance and interpretability on tabular datasets.

---

### 3. An interpretable prototype parts-based neural network for medical tabular data

**Authors:** J Karolczak, J Stefanowski | **Year/Venue:** 2025 | **URL:** [View paper](#)

#### Abstract

â€¦ a realistic benchmark for evaluating the model's ability to process numerical medical features while addressing class imbalance, a common challenge in clinical predictive modeling [24];

#### Relationship Analysis

Both papers belong to the Specialized Neural Network Designs category, developing custom neural architectures with domain-specific inductive biases for tabular data. They overlap in addressing interpretability for tabular predictive modeling through specialized architectural components: xRFM uses feature learning kernel machines with tree structures and AGOP-based interpretability, while MEDIC employs prototype parts-based networks with learnable discretization and case-based reasoning specifically for medical tabular data. The key difference is that xRFM focuses on scalable feature learning across general tabular domains with local adaptation via tree partitioning, whereas MEDIC is tailored exclusively to medical records with interpretable discretization of diagnostic ranges and prototype-based explanations aligned with clinical reasoning.

---

## Contributions Analysis

**Overall novelty summary.** The paper introduces xRFM, an algorithm combining feature learning kernel machines with an adaptive tree structure for tabular prediction. It resides in the 'Specialized Neural Network Designs' leaf under 'Deep Learning Architectures for Tabular Data', alongside three sibling papers. This leaf represents a moderately populated research direction within a broader taxonomy of 50 papers across approximately 36 topics, suggesting a focused but not overcrowded niche. The work targets general-purpose tabular prediction, contrasting with domain-specific neighbors like medical interpretability methods.

The taxonomy reveals that xRFM's leaf sits within a larger branch exploring neural architectures for tabular data, which includes attention-based transformers and graph neural networks as sibling leaves. Neighboring branches address foundation models, generative pre-training, and retrieval-augmented methods. The 'Specialized Neural Network Designs' scope explicitly excludes general transformers and graph-based methods, positioning xRFM as a custom architecture with domain-specific inductive biases. This placement suggests the work diverges from mainstream transformer adaptations, instead pursuing hybrid kernel-tree designs that balance expressiveness with tabular data constraints.

Among 22 candidates examined across three contributions, the core xRFM algorithm (9 candidates, 0 refutable) and Leaf RFM component (3 candidates, 0 refutable) show no clear prior work overlap within the limited search scope. However, the interpretability contribution via Average Gradient Outer Product (10 candidates, 2 refutable) encounters more substantial prior work. The statistics indicate that the architectural novelty appears stronger than the interpretability mechanism, though the search scale—22 candidates total—means this assessment reflects top-K semantic matches rather than exhaustive coverage. The refutable pairs suggest existing gradient-based interpretability methods may overlap with the proposed approach.

Based on the limited literature search of 22 candidates, the xRFM architecture appears relatively novel within its specialized design niche, while the interpretability component shows more overlap with existing gradient-based methods. The taxonomy context indicates this work occupies a moderately explored research direction, distinct from mainstream transformer or foundation model approaches. A more exhaustive search beyond top-K semantic matches would be needed to fully assess novelty across the broader tabular prediction landscape.

---

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

## Contribution 1: xRFM algorithm combining feature learning kernel machines with adaptive tree structure

**Description:** The authors propose xRFM, a tabular prediction method that integrates Recursive Feature Machines with binary tree-based data partitioning. This enables local feature learning (learning different features for different data subsets) while achieving  $O(n \log n)$  training complexity and  $O(\log n)$  inference complexity.

This contribution was assessed against **9 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

### 1. Navigating the Credit Landscape with Minimal Data: A Transfer Learning and Image-Based Classification Strategy

URL: [View paper](#)

#### Brief Assessment

Credit Transfer Learning[58] focuses on converting tabular credit data to images for CNN-based transfer learning in credit risk prediction. It does not address feature learning kernel machines or adaptive tree structures for general tabular data prediction.

---

### 2. Supervised contrastive representation learning with tree-structured parzen estimator Bayesian optimization for imbalanced tabular data

URL: [View paper](#)

#### Brief Assessment

Supervised Contrastive TPE[56] focuses on contrastive representation learning with Bayesian optimization for imbalanced tabular data, which is a different approach from xRFM's combination of Recursive Feature Machines with tree-based data partitioning for general tabular prediction.

---

### 3. Geodesic Flow Kernels for Semi-Supervised Learning on Mixed-Variable Tabular Dataset

URL: [View paper](#)

#### Brief Assessment

Geodesic Flow Kernels[53] focuses on semi-supervised learning for tabular data using geodesic flow kernels and tree-based embeddings for categorical/continuous variable corruption, not on feature learning kernel machines with adaptive tree structures for scalable tabular prediction.

---

### 4. Instance-based uncertainty estimation for gradient-boosted regression trees

URL: [View paper](#)

#### Brief Assessment

Instance Based Uncertainty[55] focuses on uncertainty estimation for gradient-boosted regression trees using instance-based methods, not on feature learning kernel machines with adaptive tree structures for tabular prediction.

---

### 5. Embracing uncertainty flexibility: harnessing a supervised tree kernel to empower ensemble modelling for 2D echocardiography-based prediction of right ventricular $\hat{\alpha}_1$

URL: [View paper](#)

#### Brief Assessment

Uncertainty Flexibility RV[54] focuses on ensemble regression methods for clinical echocardiography applications using tabular data, not on developing general-purpose feature learning kernel machines with adaptive tree structures for tabular prediction tasks.

---

### 6. Tree-Regularized Tabular Embeddings

URL: [View paper](#)

#### Brief Assessment

Tree Regularized Embeddings[59] uses pretrained tree ensembles to transform raw variables into embeddings for neural networks, not kernel machines. It does not combine feature learning kernel machines with adaptive tree structures as xRFM does.

---

### 7. Autoencoding Random Forests

URL: [View paper](#)

#### Brief Assessment

Autoencoding Random Forests[57] focuses on autoencoding and embedding methods for random forests, not on combining feature learning kernel machines with adaptive tree structures for tabular prediction.

---

### 8. Embracing uncertainty flexibility: harnessing a supervised tree kernel to empower ensemble modelling for 2D echocardiography-based prediction of right ventricular volume

URL: [View paper](#)

#### Brief Assessment

Uncertainty Flexibility Echocardiography[52] focuses on uncertainty quantification for medical imaging (RV volume prediction from echocardiography) using tree-based instance methods, not on general tabular prediction with feature learning kernel machines and adaptive tree structures for scalability.

---

### 9. Leveraging Structural Information in Tree Ensembles for Table Representation Learning

URL: [View paper](#)

#### Brief Assessment

Tree Ensemble Representation[51] focuses on extracting path embeddings from pre-trained tree ensembles (like GBDTs) to improve tabular data representation for deep learning models. This is fundamentally different from xRFM's approach of integrating Recursive Feature Machines with adaptive binary tree-based data partitioning for direct prediction with  $O(n \log n)$  training complexity.

---

## Contribution 2: Leaf RFM: improved kernel-RFM for tabular data

**Description:** The authors develop leaf RFM, an enhanced version of kernel-RFM that uses a more general class of kernels and optionally applies only the diagonal of the AGOP. These modifications introduce axis-aligned bias suitable for tabular data structure and enable better coordinate selection.

This contribution was assessed against **3 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

## 1. Optimally rotated coordinate systems for adaptive least-squares regression on sparse grids

URL: [View paper](#)

### Brief Assessment

Rotated Coordinate Systems[72] focuses on preprocessing for sparse grid regression methods to handle rotated coordinates, not on kernel-RFM enhancements or axis-aligned bias for tabular data prediction.

---

## 2. Topological Activation Maps for Visual Representation Learning from Tabular Data

URL: [View paper](#)

### Brief Assessment

Topological Activation Maps[71] focuses on transforming tabular data to activation maps with hybrid kernels and axis-aligned splits, not on improving kernel-RFM algorithms with diagonal AGOP or general kernel classes for tabular data.

---

## 3. Neural tangent kernels for axis-aligned tree ensembles

URL: [View paper](#)

### Brief Assessment

Neural Tangent Kernels[70] focuses on axis-aligned constraints in soft tree ensembles for theoretical analysis, not on kernel-RFM methods for tabular data prediction. The candidate addresses tree-based models with NTK framework, while the original develops feature learning kernel machines.

---

### Contribution 3: Native interpretability through Average Gradient Outer Product

**Description:** The authors show that xRFM offers built-in interpretability by exposing learned features through AGOP matrices at each leaf. The diagonal entries indicate coordinate relevance while top eigenvectors reveal directions in data most relevant for prediction, enabling understanding of heterogeneous feature importance across data subpopulations.

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. Interpretable QSPR Modeling using Recursive Feature Machines and Multi-scale Fingerprints

URL: [View paper](#)

### Brief Assessment

Interpretable QSPR[63] applies AGOP to molecular property prediction (QSPR modeling), not general tabular data. The candidate focuses on chemical fingerprints and molecular structures, representing a domain-specific application rather than challenging the novelty of AGOP-based interpretability in general machine learning models.

---

## 2. Jacobian Aligned Random Forests

URL: [View paper](#)

### Brief Assessment

Jacobian Aligned Forests[69] uses the expected gradient outer product (EGOP) and Jacobian outer product as a global linear preconditioner for feature transformation in random forests, not for interpretability purposes. The candidate focuses on improving decision boundaries through supervised preconditioning rather than exposing learned features for understanding heterogeneous feature importance across data subpopulations.

---

## 3. Efficient Spike Timing Dependent Plasticity rule for Complex-Valued Neurons

URL: [View paper](#)

### Brief Assessment

The candidate paper focuses on spike-timing dependent plasticity in complex-valued spiking neural networks, not on interpretability through gradient outer products in tabular machine learning models. These are fundamentally different domains and mechanisms.

---

## 4. Mechanism of feature learning in deep fully connected networks and kernel machines that recursively learn features

URL: [View paper](#)

### Prior Art Analysis

Recursive Feature Machines[64] demonstrates that the Average Gradient Outer Product (AGOP) mechanism for interpretability was established prior to the ORIGINAL paper. The candidate paper explicitly describes how AGOP is used to identify features relevant for prediction, with diagonal entries indicating coordinate relevance and top eigenvectors revealing important directions in data. This is the same interpretability mechanism claimed as novel in the ORIGINAL paper's xRFM contribution.

### Evidence

Evidence 1 - **Rationale:** The candidate paper establishes that Recursive Feature Machines use the AGOP mechanism for feature learning, predating the ORIGINAL paper's claim of native interpretability through AGOP. - **Original:** xrfm provides interpretability natively through the average gradient outer product - **Candidate:** the mechanism identified in our work leads to a backpropagation-free method for feature learning with any machine learning model. to demonstrate the effectiveness of this feature learning mechanism, we use it to enable feature learning in classical, non-feature learning models known as kernel machin...

Evidence 2 - **Rationale:** Both papers describe the same interpretability mechanism: using AGOP's diagonal entries and eigenvectors to identify relevant features. The candidate paper establishes this mechanism in the context of Recursive Feature Machines before the ORIGINAL paper's xRFM work. - **Original:** each leaf rfm learns features through a mathematical object known as the average gradient outer product (agop), whose diagonal indicates coordinates relevant for prediction and whose top eigenvectors indicate directions in data most relevant for prediction - **Candidate:** we posit the deep neural feature ansatz, which states that neural feature learning occurs by implementing the average gradient outer product to up-weight features strongly related to model output

---

## 5. Mechanism for feature learning in neural networks and backpropagation-free machine learning models

URL: [View paper](#)

### Prior Art Analysis

Feature Learning Mechanism[60] demonstrates that the Average Gradient Outer Product (AGOP) mechanism for interpretability and feature learning was established prior to the ORIGINAL paper. The candidate paper explicitly states that AGOP 'characterized feature learning in neural networks' and 'enabled feature learning in machine learning models, such as kernel machines, that a priori could not identify task-specific features.' This directly refutes the novelty claim that xRFM offers 'built-in interpretability by exposing learned features through AGOP matrices' as a novel contribution, since the AGOP mechanism for interpretability was already established in the prior work.

## Evidence

Evidence 1 - **Rationale:** This pair shows that Feature Learning Mechanism[60] established AGOP as a mechanism for characterizing and enabling feature learning in machine learning models, including kernel machines, prior to xRFM's claim of native interpretability through AGOP. - **Original:** an additional benefit of xrfm is that it natively provides interpretability by exposing features learned and used for prediction. in particular, each leaf rfm learns features through a mathematical object known as the average gradient outer product (agop), whose diagonal indicates coordinates releva... - **Candidate:** we presented a unifying mathematical mechanism, known as average gradient outer product (agop), that characterized feature learning in neural networks. we provided empirical evidence that agop captured features learned by various neural network architectures, including transformer-based language mod...

Evidence 2 - **Rationale:** This pair demonstrates that the mathematical formulation and interpretability mechanism of AGOP for feature learning was already established in Feature Learning Mechanism[60], undermining the claim that xRFM's use of AGOP for interpretability is a novel contribution. - **Original:** the agop is an estimate of the (uncentered) covariance of the gradients of b and intuitively captures the subspace along which the predictor highly varies (trivedi et al., 2014; kpotufe et al., 2016). the rfm algorithm involves iterating between training a predictive model and using the agop of the ... - **Candidate:** radhakrishnan et al. proposed the deep neural feature ansatz, which states that neural feature learning occurs by up-weighting the features that are most influential on model output, a process that was formulated mathematically in terms of the average gradient outer product and was supported by nume...

---

## 6. Images as weight matrices: Sequential image generation through synaptic learning rules

URL: [View paper](#)

### Brief Assessment

Images Weight Matrices[65] focuses on using outer products for sequential image generation in GANs, not for interpretability of tabular data models through gradient-based feature analysis.

---

## 7. Feature learning as alignment: a structural property of gradient descent in non-linear neural networks

URL: [View paper](#)

### Brief Assessment

Feature Learning Alignment[62] focuses on theoretical analysis of AGOP correlation dynamics in neural networks during training, not on using AGOP for interpretability in tabular models. The contexts are fundamentally different: neural network training dynamics versus interpretability mechanisms in tree-based kernel machines.

---

## 8. Emergence in non-neural models: grokking modular arithmetic via average gradient outer product

URL: [View paper](#)

### Brief Assessment

Grokking Modular Arithmetic[66] focuses on using AGOP for feature learning in modular arithmetic tasks to explain grokking phenomena, not on interpretability of heterogeneous feature importance across data subpopulations in tabular data models.

---

## 9. Grokking Modular Arithmetic Through Group Actions: A Group-Theoretic View of Machine Learning Behavior

URL: [View paper](#)

### Brief Assessment

Grokking Group Actions[68] uses AGOP for analyzing feature learning in modular arithmetic tasks, not for general tabular data interpretability. The candidate focuses on understanding grokking phenomena through group theory rather than providing interpretability for heterogeneous tabular datasets.

---

## 10. Reversed Attention: On The Gradient Descent Of Attention Layers In GPT

URL: [View paper](#)

### Brief Assessment

Reversed Attention[61] focuses on attention mechanisms in GPT models during backpropagation, not on feature learning in tabular data models. The AGOP usage contexts are fundamentally different.

---

## Appendix: Text Similarity Detection

Textual similarity detection checked 25 papers and found 2 similarity segment(s) across 2 paper(s).

The following **2 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. Mechanism of feature learning in deep fully connected networks and kernel machines that recursively learn features

**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.

---

### 2. Emergence in non-neural models: grokking modular arithmetic via average gradient outer product

**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] xRFM: Accurate, scalable, and interpretable feature learning models for tabular data [View paper](#)
- [1] Accurate predictions on small data with a tabular foundation model [View paper](#)
- [2] Generative models for tabular data: A review [View paper](#)
- [3] Deep neural networks and tabular data: A survey [View paper](#)
- [4] Goggle: Generative modelling for tabular data by learning relational structure [View paper](#)
- [5] Representation learning for tabular data: A comprehensive survey [View paper](#)
- [6] Tabbie: Pretrained representations of tabular data [View paper](#)
- [7] When do neural nets outperform boosted trees on tabular data? [View paper](#)
- [8] Large Language Models (LLMs) on Tabular Data: Prediction, Generation, and Understanding--A Survey [View paper](#)
- [9] Diffusion-based time series imputation and forecasting with structured state space models [View paper](#)

- [10] Recent deep learning methods for tabular data [View paper](#)
- [11] Language modeling on tabular data: A survey of foundations, techniques and evolution [View paper](#)
- [12] Explaining customer churn prediction in telecom industry using tabular machine learning models [View paper](#)
- [13] Anypredict: Foundation model for tabular prediction [View paper](#)
- [14] Large scale transfer learning for tabular data via language modeling [View paper](#)
- [15] Predicting structured data [View paper](#)
- [16] Low-cost and efficient prediction hardware for tabular data using tiny classifier circuits [View paper](#)
- [17] Unipredict: Large language models are universal tabular predictors [View paper](#)
- [18] Regularization learning networks: deep learning for tabular datasets [View paper](#)
- [19] AixelNet: A Pre-trained Model with Table-aware Adaptation for Structured Data Prediction [View paper](#)
- [20] A survey on deep tabular learning [View paper](#)
- [21] Is deep learning on tabular data enough? An assessment [View paper](#)
- [22] An interpretable prototype parts-based neural network for medical tabular data [View paper](#)
- [23] Heart disease prediction by tabular modeling with deep learning network and interpretability [View paper](#)
- [24] TLTD: Transfer learning for tabular data [View paper](#)
- [25] Incorporating key features from structured and unstructured data for enhanced carbon trading price forecasting with interpretability analysis [View paper](#)
- [26] Transtab: Learning transferable tabular transformers across tables [View paper](#)
- [27] Graph neural networks for tabular data learning: A survey with taxonomy and directions [View paper](#)
- [28] Evaluating the fidelity of statistical forecasting and predictive intelligence by utilizing a stochastic dataset [View paper](#)
- [29] Hyperfusion: A hypernetwork approach to multimodal integration of tabular and medical imaging data for predictive modeling [View paper](#)
- [30] Generative table pre-training empowers models for tabular prediction [View paper](#)
- [31] Unleashing the potential of large language models for predictive tabular tasks in data science [View paper](#)
- [32] A systematic review of machine learning models for predicting outcomes of stroke with structured data [View paper](#)
- [33] Annotatedtables: A large tabular dataset with language model annotations [View paper](#)
- [34] Advances in AI and machine learning for predictive medicine [View paper](#)
- [35] Tabular data: Deep learning is not all you need [View paper](#)
- [36] Combining structured and unstructured data for predictive models: a deep learning approach [View paper](#)
- [37] Tabred: Analyzing pitfalls and filling the gaps in tabular deep learning benchmarks [View paper](#)
- [38] Retrieval-augmented forecasting with tabular time series data [View paper](#)
- [39] Unipredict: Large language models are universal tabular classifiers [View paper](#)
- [40] Demystifying statistics and machine learning in analysis of structured tabular data [View paper](#)
- [41] Deep learning with structured data [View paper](#)
- [42] Time series forecasting methods and their applications to particle accelerators [View paper](#)
- [43] AlphaML: a clear, legible, explainable, transparent, and elucidative binary classification platform for tabular data [View paper](#)
- [44] Quantifying Prediction Consistency Under Fine-tuning Multiplicity in Tabular LLMs [View paper](#)
- [45] Predictive analytics with data visualization [View paper](#)
- [46] Data Augmentation of a Corrosion Dataset for Defect Growth Prediction of Pipelines Using Conditional Tabular Generative Adversarial Networks [View paper](#)
- [47] Deep learning with tabular data: A self-supervised approach [View paper](#)
- [48] Retrieval & interaction machine for tabular data prediction [View paper](#)
- [49] Easytime: Time series forecasting made easy [View paper](#)
- [50] Unitabe: A universal pretraining protocol for tabular foundation model in data science [View paper](#)
- [51] Leveraging Structural Information in Tree Ensembles for Table Representation Learning [View paper](#)
- [52] Embracing uncertainty flexibility: harnessing a supervised tree kernel to empower ensemble modelling for 2D echocardiography-based prediction of right ventricular volume [View paper](#)
- [53] Geodesic Flow Kernels for Semi-Supervised Learning on Mixed-Variable Tabular Dataset [View paper](#)
- [54] Embracing uncertainty flexibility: harnessing a supervised tree kernel to empower ensemble modelling for 2D echocardiography-based prediction of right ventricular volume [View paper](#)
- [55] Instance-based uncertainty estimation for gradient-boosted regression trees [View paper](#)
- [56] Supervised contrastive representation learning with tree-structured parzen estimator Bayesian optimization for imbalanced tabular data [View paper](#)
- [57] Autoencoding Random Forests [View paper](#)
- [58] Navigating the Credit Landscape with Minimal Data: A Transfer Learning and Image-Based Classification Strategy [View paper](#)
- [59] Tree-Regularized Tabular Embeddings [View paper](#)
- [60] Mechanism for feature learning in neural networks and backpropagation-free machine learning models [View paper](#)
- [61] Reversed Attention: On The Gradient Descent Of Attention Layers In GPT [View paper](#)
- [62] Feature learning as alignment: a structural property of gradient descent in non-linear neural networks [View paper](#)
- [63] Interpretable QSPR Modeling using Recursive Feature Machines and Multi-scale Fingerprints [View paper](#)
- [64] Mechanism of feature learning in deep fully connected networks and kernel machines that recursively learn features [View paper](#)
- [65] Images as weight matrices: Sequential image generation through synaptic learning rules [View paper](#)
- [66] Emergence in non-neural models: grokking modular arithmetic via average gradient outer product [View paper](#)
- [67] Efficient Spike Timing Dependent Plasticity rule for Complex-Valued Neurons [View paper](#)
- [68] Grokking Modular Arithmetic Through Group Actions: A Group-Theoretic View of Machine Learning Behavior [View paper](#)
- [69] Jacobian Aligned Random Forests [View paper](#)
- [70] Neural tangent kernels for axis-aligned tree ensembles [View paper](#)
- [71] Topological Activation Maps for Visual Representation Learning from Tabular Data [View paper](#)
- [72] Optimally rotated coordinate systems for adaptive least-squares regression on sparse grids [View paper](#)