Tabular Data: The Foundation for Enterprise AI

Research note · July 2026

Foundation models reshaped machine learning for language and images: one model, pretrained once, applied across many tasks without task-specific training. The data that dominates enterprise systems, however, is neither text nor images — it is **tabular**: rows and columns of transactions, master data, and time series. This note reviews the recent research that establishes tabular data as a foundation-model modality in its own right — what these models are, how they perform, and where they still fall short.

Tabular data's footprint, and the research gap

Tabular data is "*among the most prevalent data types in machine learning classification and regression applications*" (Representation Learning for Tabular Data, 2025). Despite that ubiquity, it long received comparatively little foundation-model research. A 2024 ICML position paper stated the imbalance directly: tabular data is "*the dominant modality in many fields, yet it is given hardly any research attention and significantly lags behind in terms of scale and power*" (van Breugel & van der Schaar, 2024). The authors argue for *large tabular models* that contextualise a dataset against related ones, rather than treating each table in isolation.

What a tabular foundation model is

The current exemplar is **TabPFN** (Hollmann et al., *Nature*, 2025), a *prior-data fitted network*. Rather than training on the user's data, the network is pretrained once on a large collection of synthetically generated tabular datasets — produced from causal models with a bias toward simpler structures — to learn a general-purpose learning algorithm. At use time it performs **in-context learning**: the labelled training rows and the query rows pass through the network in a single forward pass, with attention alternating across rows and columns, and predictions emerge with no gradient training on the new data. There are no hyperparameters to tune and no model to fit per dataset.

Empirical standing

On small tabular problems the approach is strong. In the *Nature* evaluation, TabPFN matched or exceeded a tuned ensemble of the best baseline methods while running in about **2.8 seconds**, against roughly **four hours** of tuning for those baselines (Hollmann et al., 2025). The *Nature* model targets datasets up to roughly 10,000 samples; later versions extend the range substantially — TabPFN-3 (2026) to on the order of a million rows.

Independent benchmarking both tempers and sharpens the picture. **TabArena**, a continuously maintained ("living") benchmark introduced at NeurIPS 2025, reports that gradient-boosted trees remain strong on realistic tabular datasets, that deep-learning methods catch up given larger time budgets and ensembling, that **foundation models are strongest on smaller datasets**, and that **ensembling across model families advances the state of the art** (TabArena, 2025). No single method dominates everywhere — a finding with direct consequences for how these models should be deployed.

Limitations and open problems

Tabular foundation models are not a universal solution. Reported performance degrades on high-dimensional and large-scale data — the regime where gradient-boosted trees and tuned deep networks remain competitive. More fundamentally, a 2025 position paper from SAP researchers argues that models trained on single tables in isolation miss the *operational context* — the rules, procedures, and domain logic that govern how enterprise data is created and used — and proposes grounding models in "*semantically linked tables*" instead (Klein & Hoffart, 2025). Treating a table as a self-contained artifact is, on this view, an incomplete abstraction for enterprise use.

Enterprise trajectory

The commercial signal matches the research momentum. In 2026 **SAP acquired Prior Labs**, the group behind TabPFN, and **committed to invest more than €1 billion over four years** to develop it as an independent frontier lab for — in SAP's words — "*the structured data that runs the world's businesses*" (terms of the acquisition were not disclosed) (SAP, 2026). SAP states that the TabPFN series, published in *Nature*, "*set the state-of-the-art on tabular benchmarks across hundreds of independent academic studies.*"

Relevance to demand planning in the ERP

Demand forecasting is a tabular problem — history, drivers, and horizons arranged as rows and columns — and the research above bears on it directly.

Two findings are especially load-bearing. First, a pretrained model requires **no per-item fitting and no hyperparameter search**, which is what makes it practical to forecast an entire product range on a routine cycle rather than a curated subset; the *Nature* result of seconds against hours of tuning is the difference between forecasting everything and forecasting what there was time for. Second, TabArena's finding that **foundation models are strongest on smaller datasets** is well matched to demand planning, where an individual product's history is short by machine-learning standards — often a few dozen periods, and for a newly launched item, none at all.

The limitations matter just as much. TabArena's conclusion that no single method dominates everywhere argues against treating any model's output as the final word. The SAP position paper on systemic context goes further: a table read in isolation misses the operational rules and domain logic surrounding it. Both point to the same design conclusion — the model should produce a **baseline**, and the planner must retain the ability to override it with knowledge the table does not contain: a promotion, a customer commitment, a known shift in the market.

Demand Planning Add-In for SAP S/4HANA Cloud ERP is built on that basis. A tabular foundation model generates the baseline forecast from sales history already held in the ERP. The planner corrects history where actuals mislead and overrides any forecast period where judgement is required, with model output and manual adjustment kept as separate figures so each contribution stays visible. Products with no history of their own are phased in from a reference product rather than skipped. The agreed result is released as planned independent requirements, driving MRP and supply planning inside the same system.

References

1. Hollmann, N., Müller, S., Purucker, L., Krishnakumar, A., Körfer, M., Hoo, S. B., Schirrmeister, R. T., & Hutter, F. "Accurate predictions on small data with a tabular foundation model." *Nature* **637**, 319–326 (2025). https://www.nature.com/articles/s41586-024-08328-6

2. van Breugel, B., & van der Schaar, M. "Position: Why Tabular Foundation Models Should Be a Research Priority." *ICML 2024*, PMLR 235. https://proceedings.mlr.press/v235/van-breugel24a.html
3. "TabArena: A Living Benchmark for Machine Learning on Tabular Data." arXiv:2506.16791 (2025); NeurIPS 2025 Datasets & Benchmarks Track. https://arxiv.org/abs/2506.16791
4. Klein, T., & Hoffart, J. (SAP SE). "Position: Foundation Models for Tabular Data within Systemic Contexts Need Grounding." arXiv:2505.19825 (2025). https://arxiv.org/abs/2505.19825
5. "Representation Learning for Tabular Data: A Comprehensive Survey." arXiv:2504.16109 (2025). https://arxiv.org/abs/2504.16109
6. SAP SE. "SAP to Acquire Prior Labs to Establish a Globally Leading Frontier AI Lab in Europe." SAP News, May 2026. https://news.sap.com/2026/05/sap-to-acquire-prior-labs-establish-frontier-ai-lab-europe/

Disclaimer. This note reflects published research and publicly available information as of July 2026. Tabular foundation models, benchmark results, and vendor offerings in this area are evolving rapidly; all figures and quotations are attributed to the cited sources and may be superseded by later work.

Integrated Business Planning Solutions LLC · solveibp.com