GenAI for Predictive Analytics in Insurance

Frank Hutter

frank@priorlabs.ai

The GenAI Revolution

GenAI has transformed text & images

Text

OpenAI DeepMind deepseek

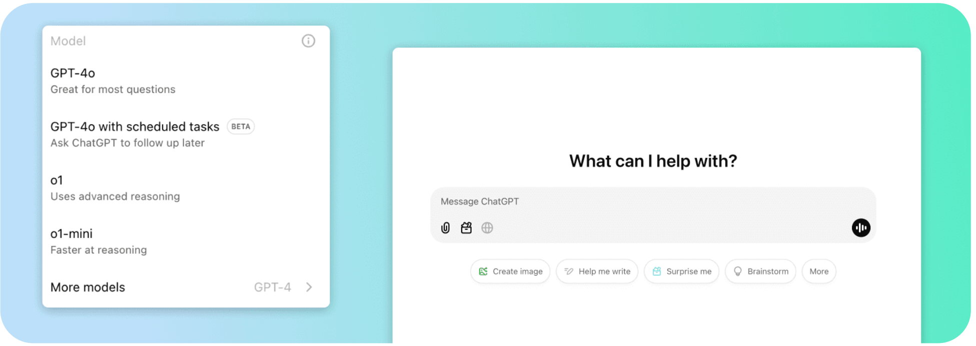

Image

Midjourney black forest labs.

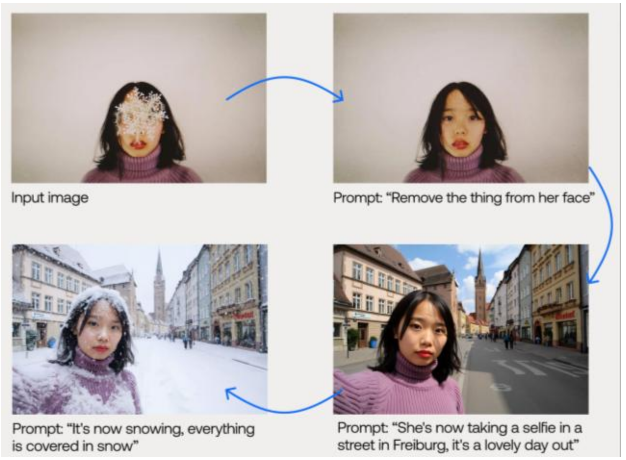

How is GenAI relevant for insurance?

Super helpful in many ways facing customers

- Call centers: virtual assistants
- Emails: customer engagement & claims communication
- Marketing: content creation
- Claims processing: summarizing documents

Not helpful for the core business

- Numbers
- Statistics
- Estimating risk
- Predictive analytics

What are Predictive Analytics Tasks for Insurers?

- What will the **cost of claims** be for this policyholder?
- What **premium** should we set for this customer?
- What is the expected **customer lifetime value**?
- Is this claim likely to be **fraudulent**?
- What **risk category** does this customer fall into?
- Will this customer **churn** or purchase **additional coverage**?
- How will **loss ratios** and **reserves** develop over time?
- How are **fraud patterns** or **climate-driven risks** evolving?
- ...

Most data driving daily enterprise decisions is structured in tables

Customer Datasets

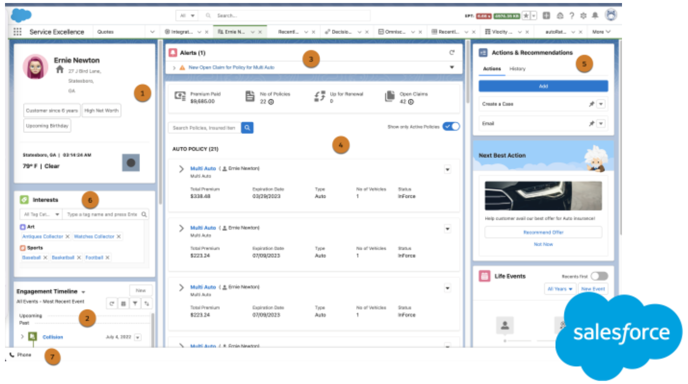

Risk & Underwriting Tables

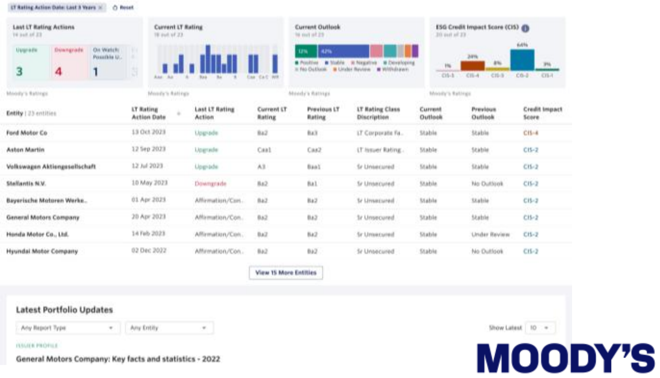

Product Catalogs

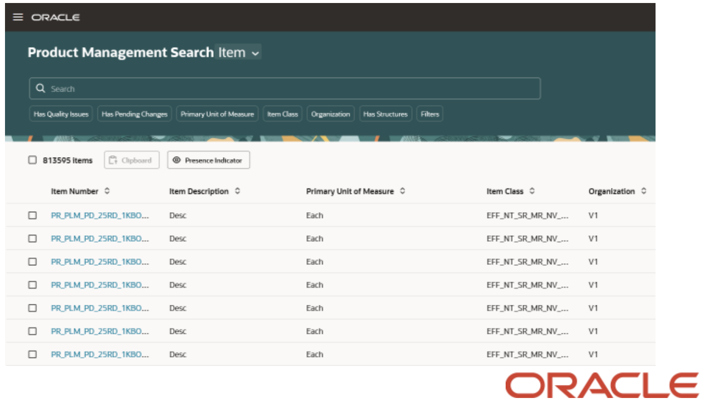

Transaction Logs

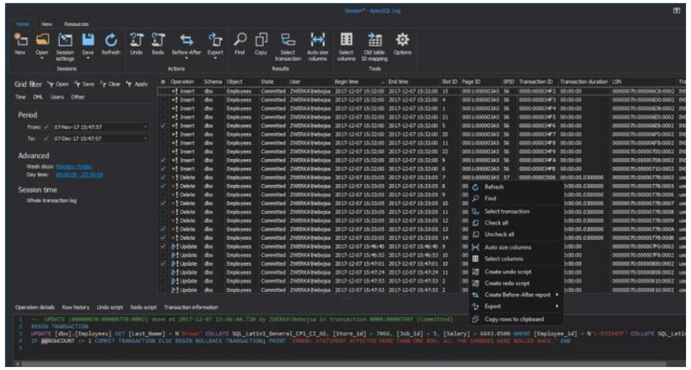

GenAI has transformed text & images

But our most valuable data is organized in tables

Text

OpenAI

DeepMind

ANTHROP\C

deepseek

ALEPH ALPHA

Tables, Time Series & Databases

?

Images

OpenAI

DeepMind

deepseek

Midjourney

black forest labs.

Outline

- Motivation
- TabPFN: The GenAI Revolution for Tabular Data
- Prior Labs: the Startup around TabPFN
- TabPFN for Insurance

Concrete Example of Tabular Prediction (Insurance Lead Prediction)

City Code	Region Code	Accomodation Insurance	Reco Type Insurance	Upper Age	Lower Age	Is Spouse	Health Indicator	Holding Policy Duration	Holding Policy Type	Reco Policy Category	Reco Policy Premium	Response
C3	3213	Rented	Individual	36	36	No	X1	14+	3.0	22	11628.0	0
C8	2190	Rented	Individual	44	44	No	X2	3.0	1.0	16	10404.0	0
C9	1785	Rented	Individual	52	52	No	X2	5.0	1.0	22	15264.0	1
C1	3175	Owned	Joint	75	73	Yes	X4	9.0	4.0	17	29344.0	1
C15	3497	Owned	Joint	52	43	No	X1	14.0	3.0	1	27283.2	0
C9	855	Rented	Individual	21	21	No	X5			16	3744.0	1
C1	520	Rented	Individual	24	24	No	X5			19	7800.0	0
C3	927	Rented	Joint	49	47	Yes	X4	3.0	3.0	18	22568.0	1
C2	1036	Owned	Individual	64	64	No		14+	3.0	3	19830.0	0
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
C2	1016	Owned	Individual	48	48	No		14+	3.0	21	15180.0	1
C3	1471	Rented	Individual	31	31	No	X1	5.0	1.0	16	10926.0	? 0.24
C2	2659	Owned	Individual	68	68	No	X2	1.0	4.0	11	15456.0	? 0.91
C2	2779	Rented	Individual	33	33	No		3.0	4.0	22	6800.0	? 0.55

Tabular prediction task (probabilistic predictions)

GenAI for Predictive Analytics?

- You could just put your table into GPT-5 and aks for predictions
  - But that’s a bad idea: LLMs are very poor at statistical reasoning
  - Decade-old tabular ML methods work much better

CatBoost	0.743
RandomForest	0.747
XGBoost	0.6915
ChatGPT-5	0.2775

But GenAI is much easier to use

Traditional machine learning approach

Data science project (hours to months)

Missing value imputation | Feature selection

Model selection & training | Hyperparameter tuning

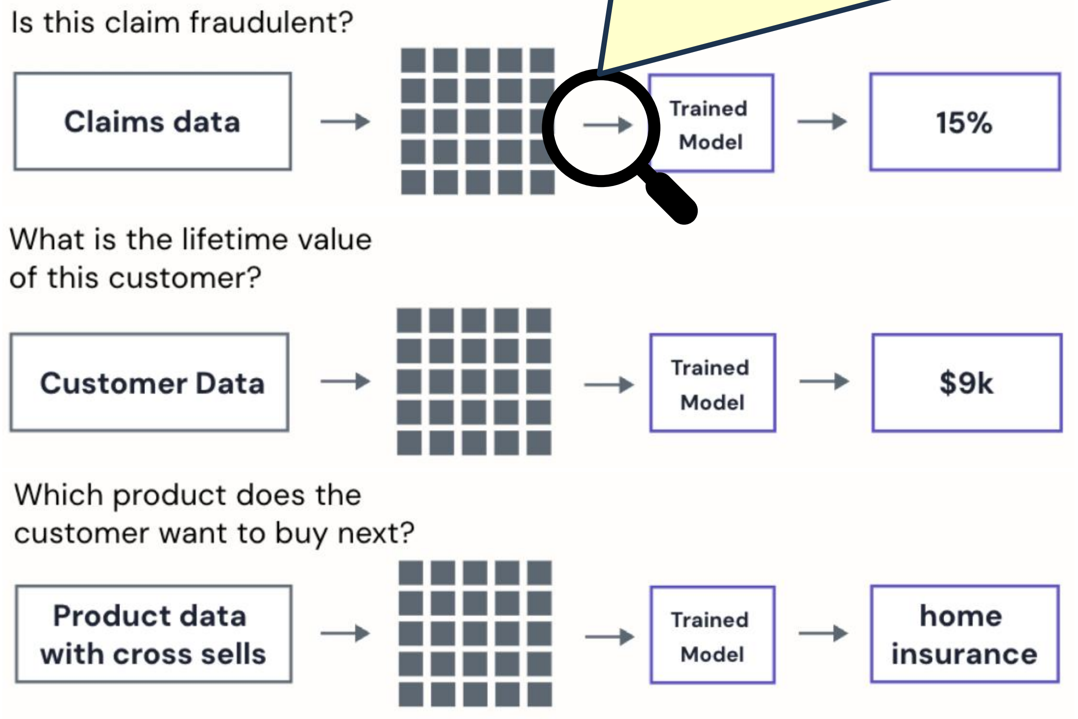

- Lengthy task-specific model selection & training
- Individual siloed models

GenAI

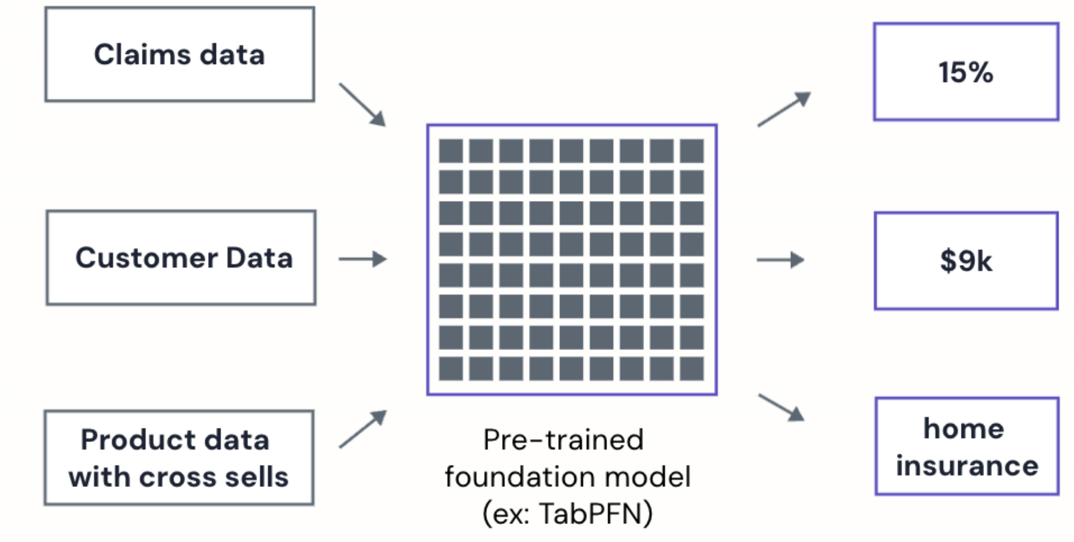

- Off-the-shelf use without retraining
- Can quickly be customized to new use cases

TabPFN: new GenAI method revolutionizing predictive analytics

TabPFN v2
Publication in *Nature*

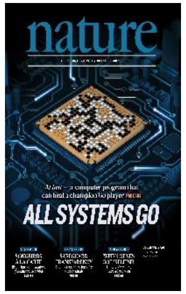

More accurate than previous ML on

>96% of use-cases

More data efficient than previous ML

Only 50% of the data needed for same accuracy

More than
1.7 Million Downloads

Open-source Software with 4K stars on Github

OPEN SOURCE USED AT

Inkitt · entrostat · sky · ETH · EAC · KTH · KBase · PIC NIC

The standard tabular data prediction problem

City Code	Region Code	Accomodation Type	Upper Age	Health Indicator	Reco Policy Category	Response
C8	2190	Rented	44	X2	16	0
C9	1785	Rented	52	X2	22	1
C3	679	Owned	28		17	0
C1	3175	Owned	75	X4	17	1
C15	3497	Owned	52	X1	1	0
C1	530	Owned	59		18	1
C28	600	Owned	21	X2	21	1
C27	1097	Owned	59	X3	13	0
⋮	⋮	⋮	⋮	⋮	⋮	⋮
C3	1484	Rented	20	X3	2	0
C24	2748	Owned	67	X1	17	?
C27	811	Owned	75	X1	20	?
C2	256	Owned	42	X1	20	?

X_{train} y_{train}

X_{test} y_{test}

TabPFN High-Level Overview

TabPFN is trained on synthetic data to take entire datasets as inputs and predict in a forward pass

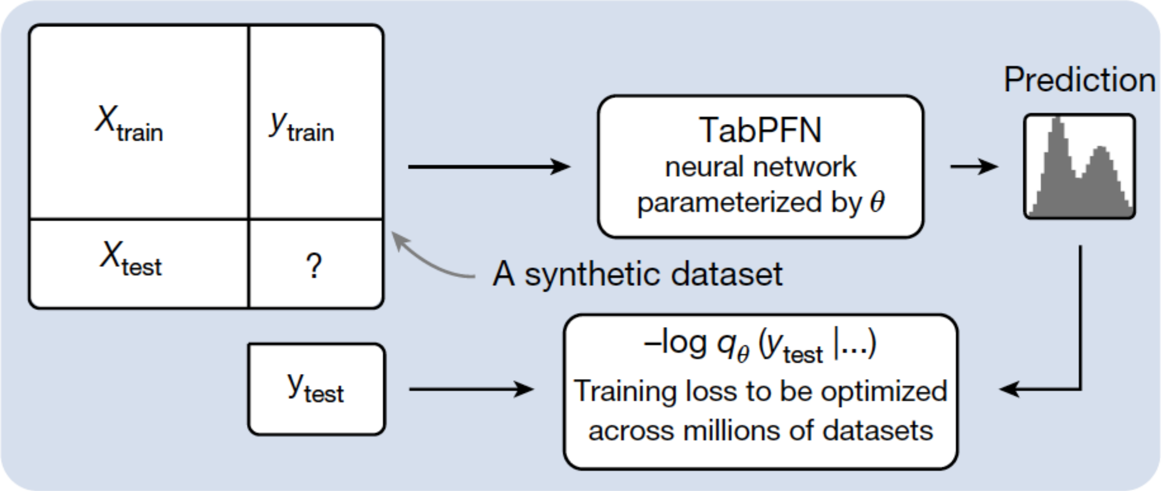

TabPFN can now be applied to arbitrary unseen real-world datasets

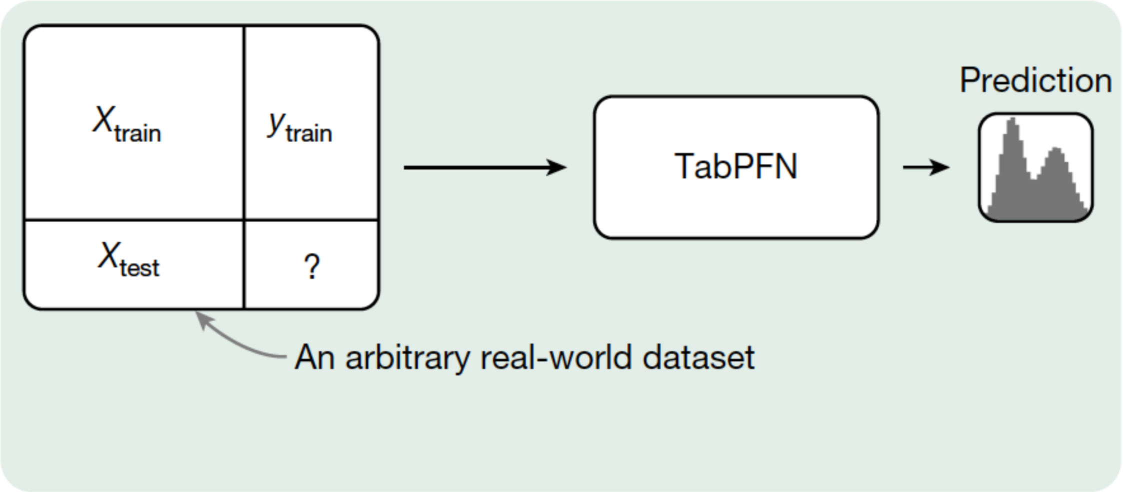

TabPFN: a learned tabular prediction algorithm

- Executed in a forward pass of a transformer neural network
- In contrast to previous handcoded algorithms, TabPFN has been *learned* to do well

Learned across which datasets?

- 130 million synthetic data sets sampled entirely from scratch
- Based on highly parameterized structural **causal models** $\rightarrow$ generalizes to arbitrary domains

Clear New State-of-the-Art Results Across Many Datasets

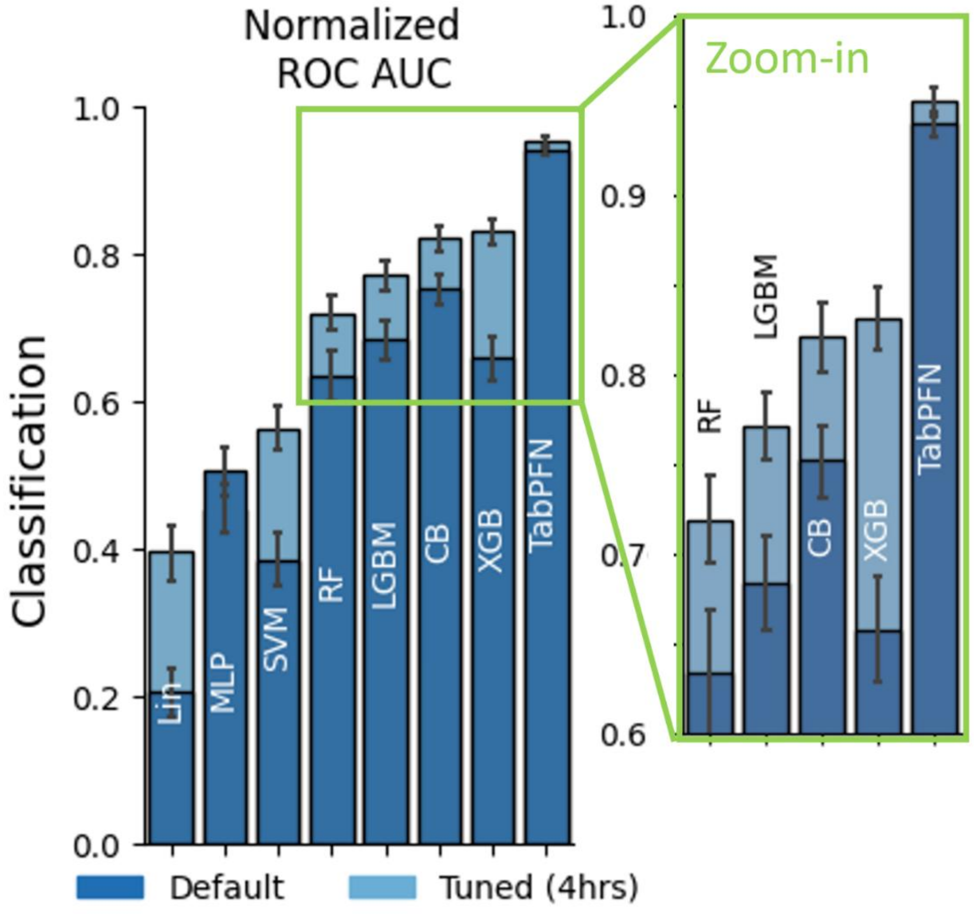

Result across 29 datasets:
better in 5s than other methods in 4h

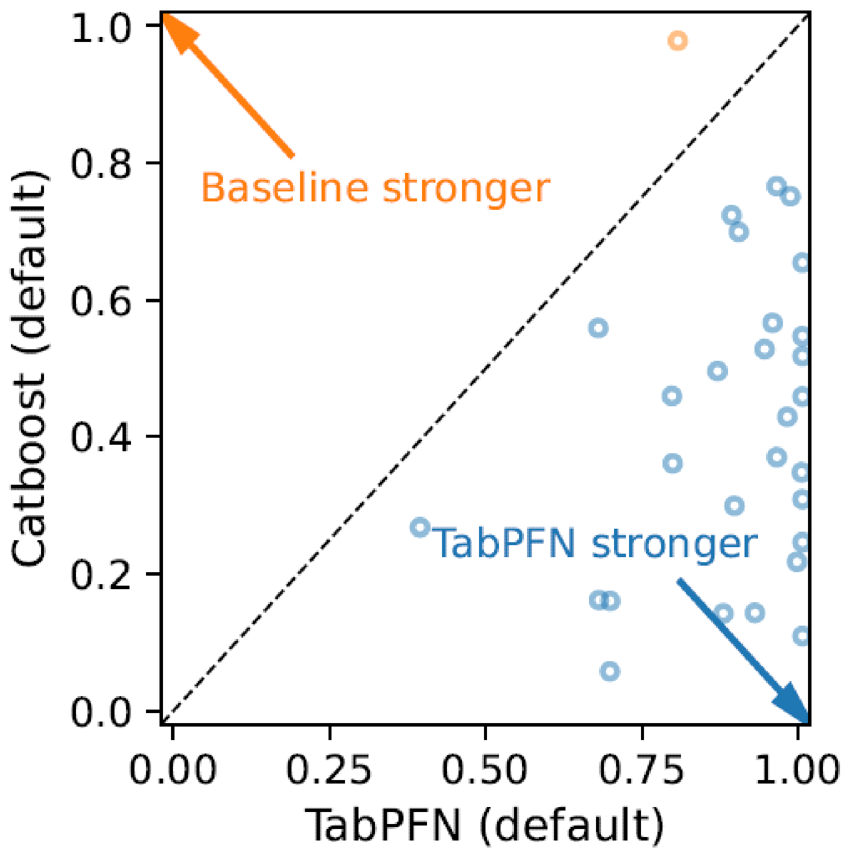

Time Series Forecasting

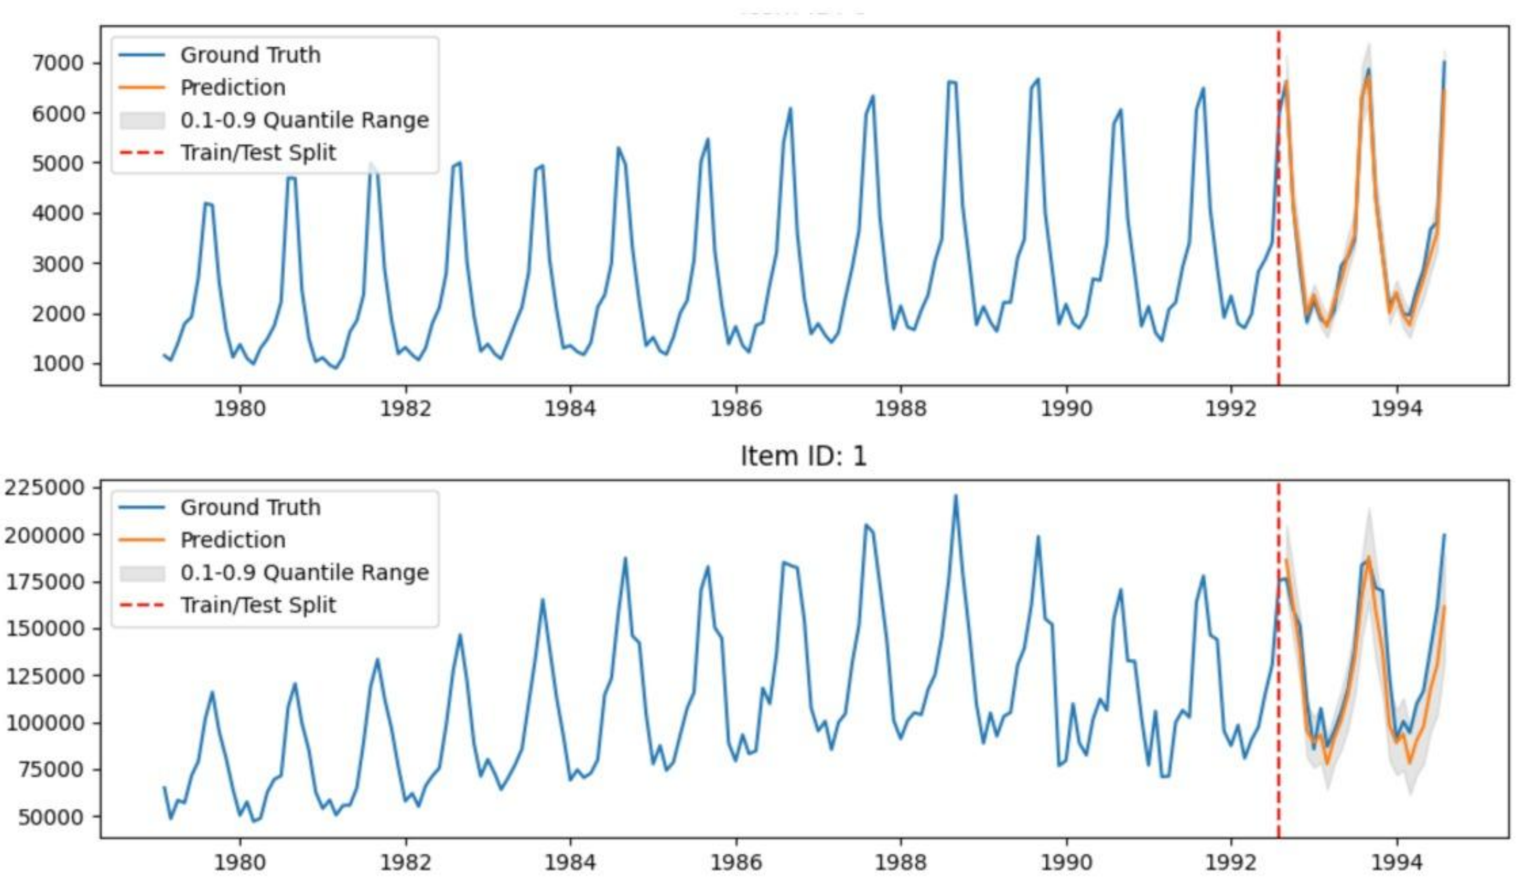

How will XYZ develop over time?

- Claim frequency
- Claim severity
- Liabilities from past underwriting
- Capital & solvency
- Demand
- ...

Time series forecasting as tabular prediction

- **TabPFN-TS**: state-of-the-art forecasting model, **outperforming Google and Amazon**
- Even though it **only ever saw synthetic tabular datasets** (no time series, no real data)

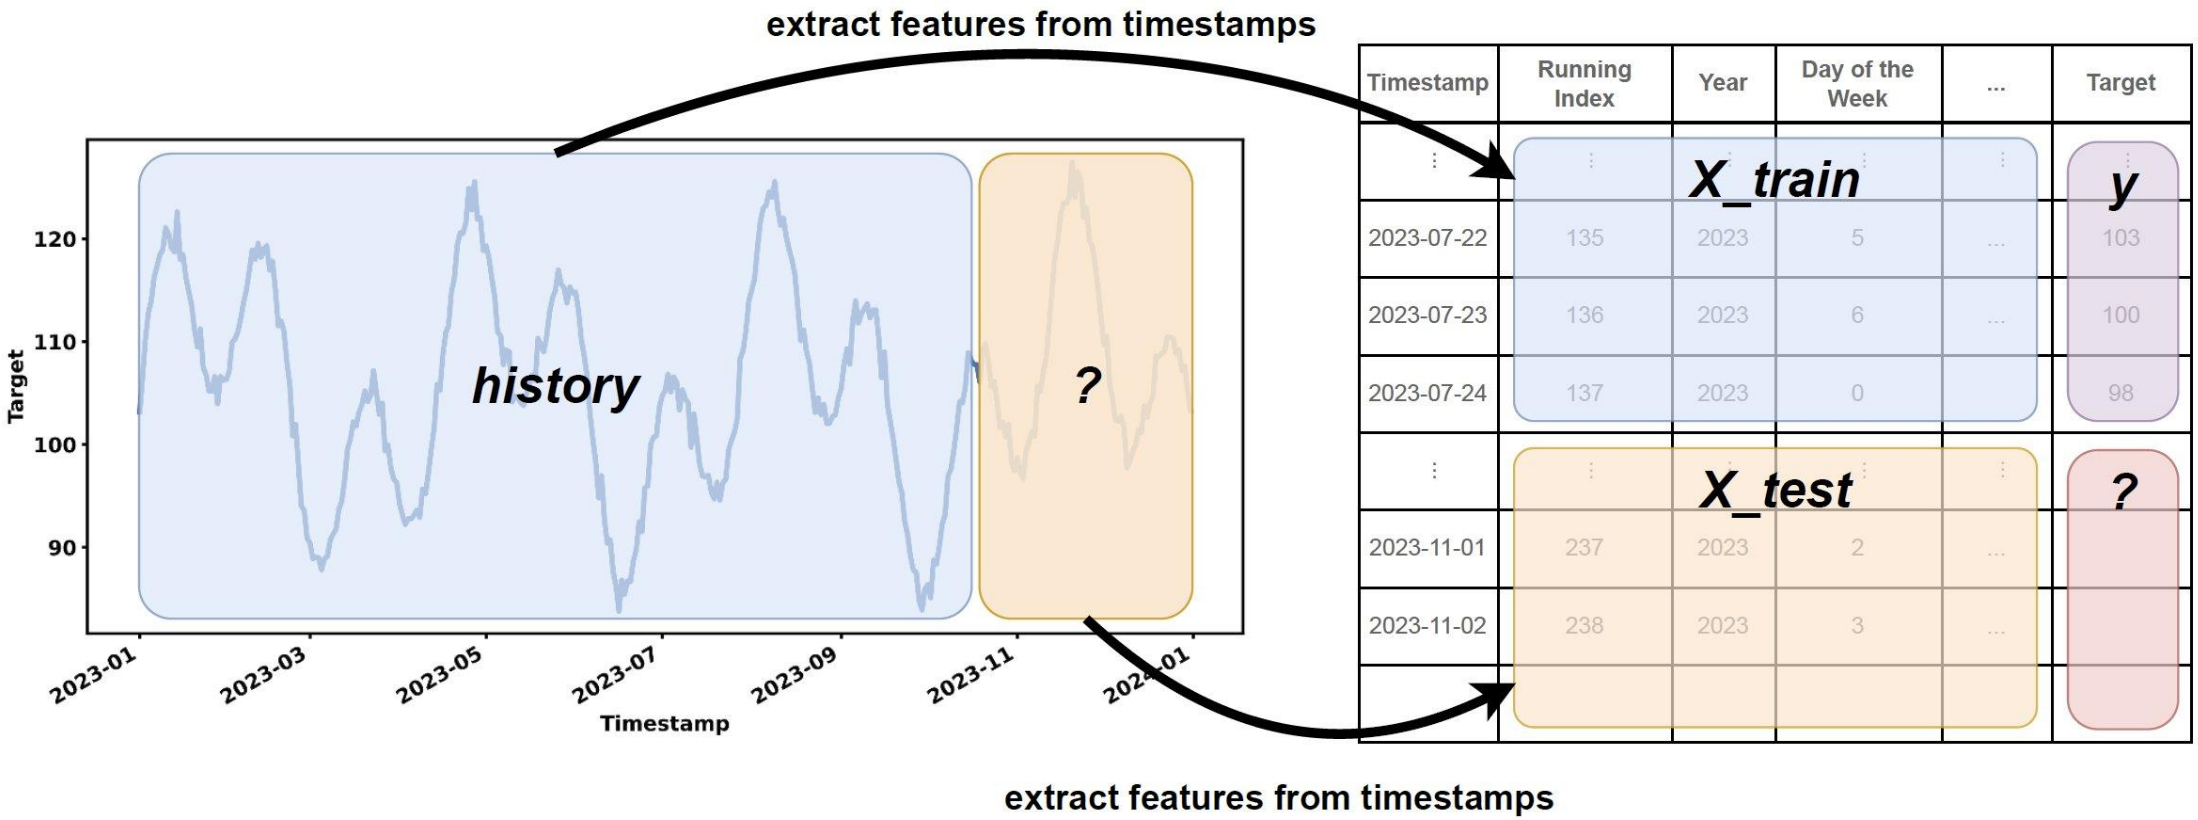

Interpretability: Smooth Predictions & Reliable Uncertainty Estimates

Input data with 2 classes

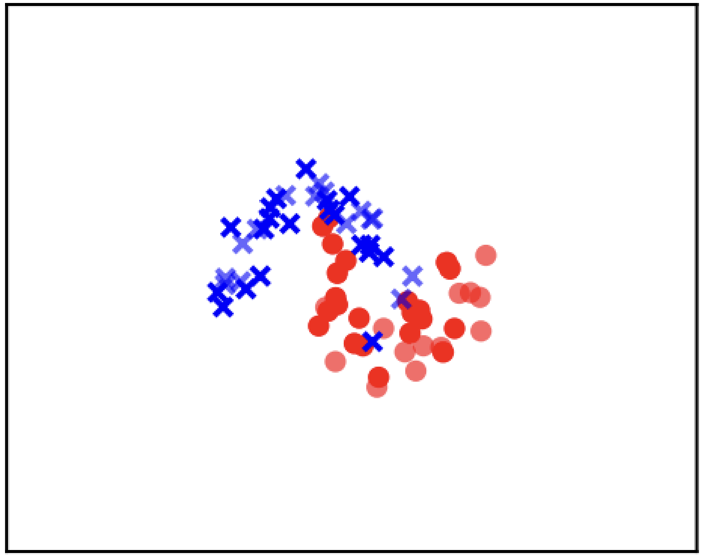

Predictions with CatBoost

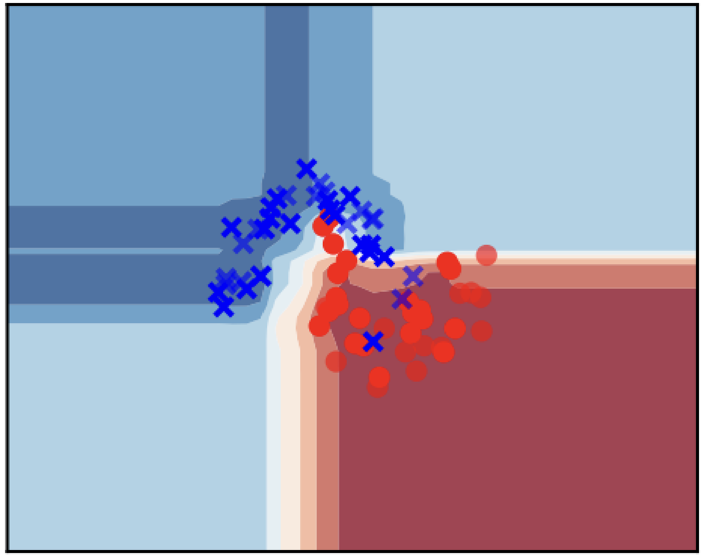

Predictions with TabPFN

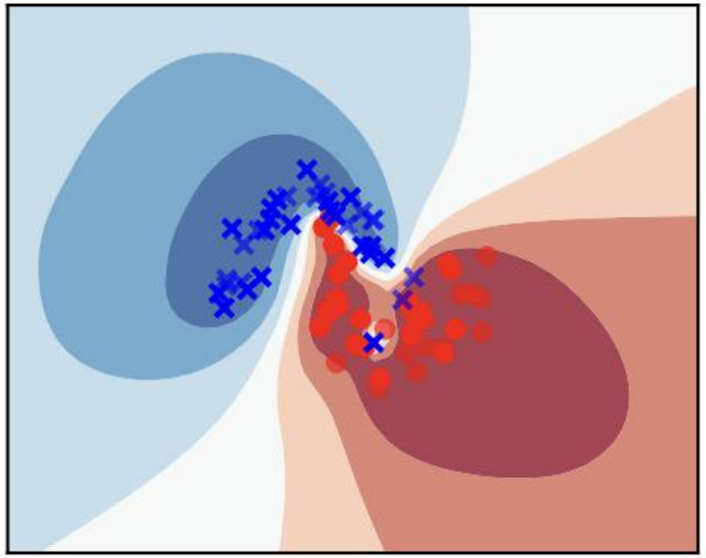

Dark blue: sure that blue class
White: completely unsure
Dark red: sure that red class

TabPFN makes very smooth predictions

Interpretability: effect of each feature clearer for TabPFN

SHAP analysis yields
more reliable results for TabPFN

- Reason: much smoother predictions than boosted trees (see previous slide)

- **Interpretation:**
  - Younger customers are less credit-worthy
  - If they already have an account with funds that helps

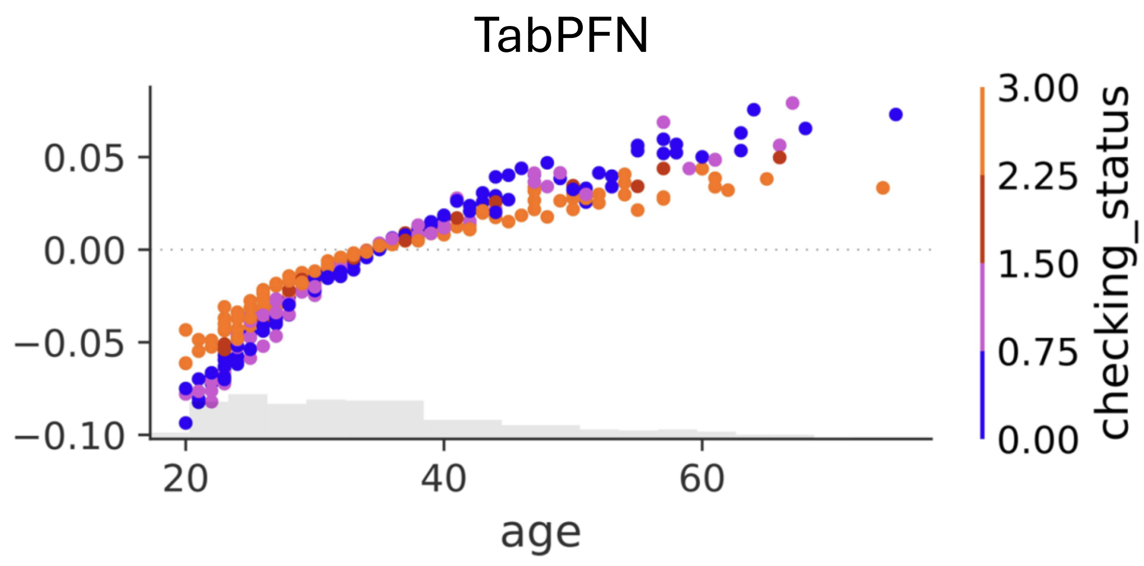

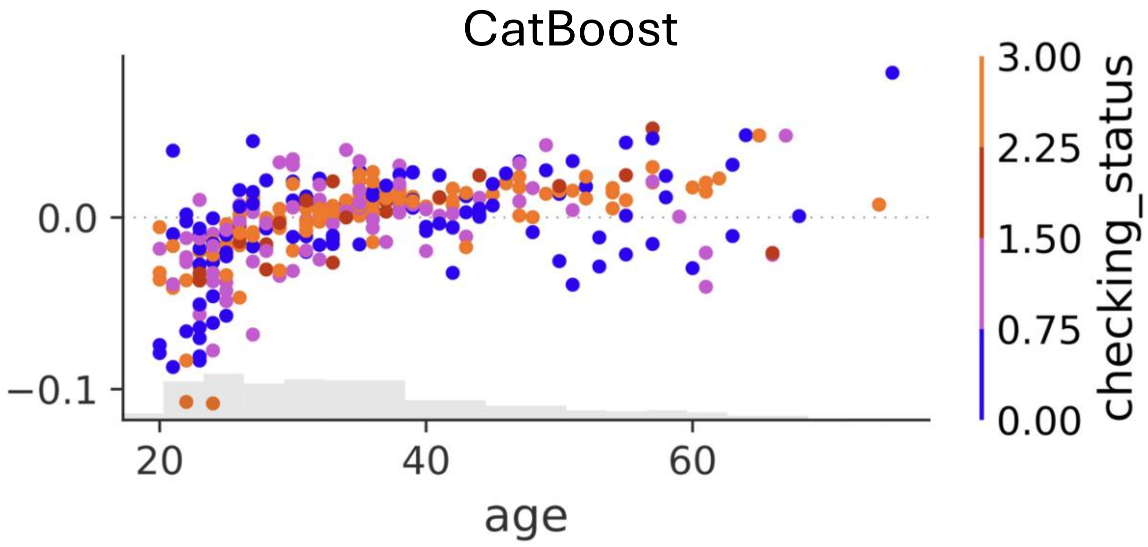

Novel Capabilities: a Single Model Across Tasks

- **Removes need for manual selection between many models & hyperparameters**
  - Faster
  - More reproducible

- **Simplified operations**
  - Audit & approve model once, then use everywhere (like ChatGPT)
  - Reduced maintenance: simple pipeline to serve the model

Novel capabilities: Finetuning TabPFN to your specific data

- You can **finetune ChatGPT** (or Llama) to your own types of text
  - E.g., lots of claims texts
  - Can dramatically improve results

- Likewise, **you can finetune TabPFN to your own tabular datasets**
  - Can dramatically improve results
  - In contrast to traditional ML, which always learns from scratch

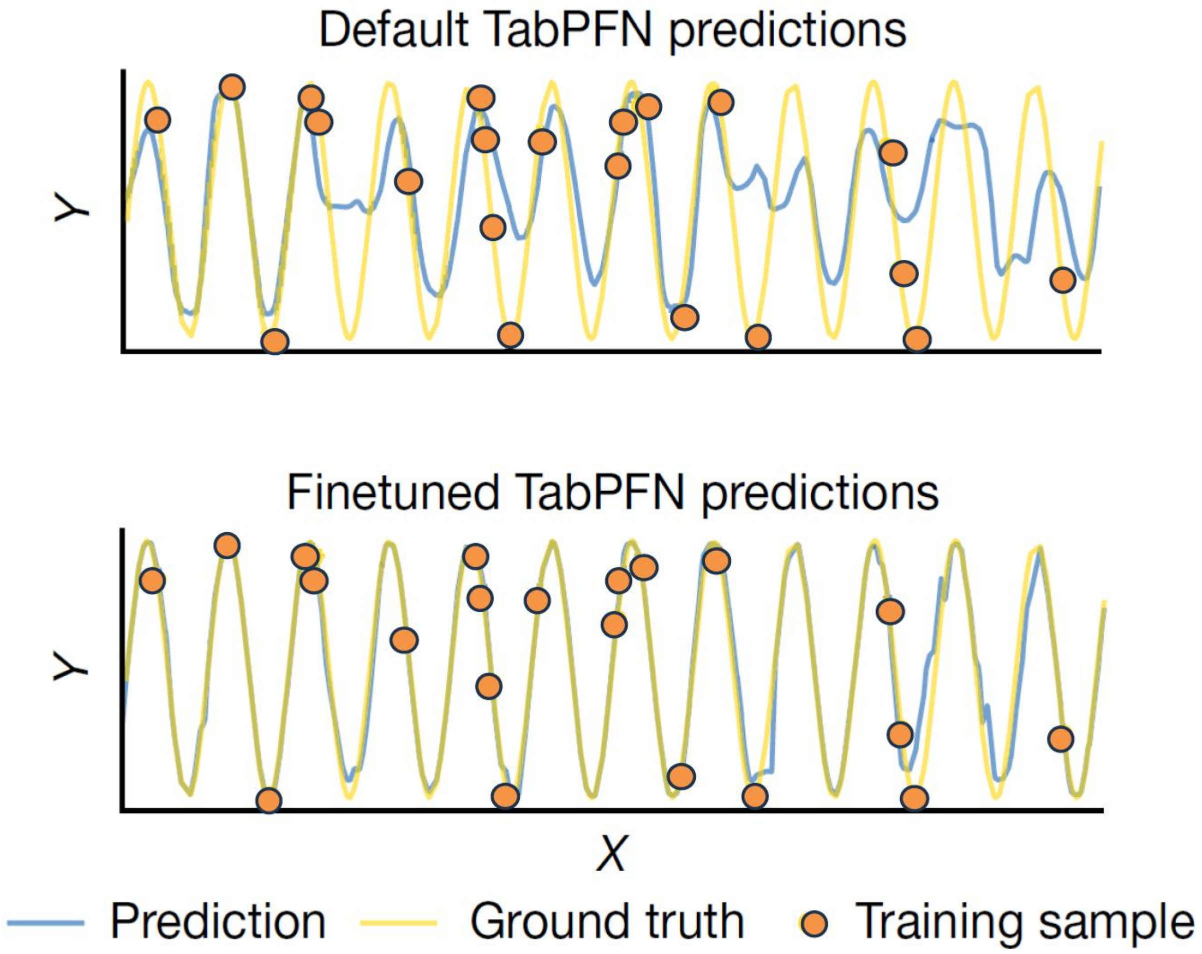

Novel Capabilities: Proper What-If Analyses

City Code	Region Code	Accomodation Type	Age	Health Indicator	Reco Policy Category	Response
C8	2190	Rented	44	X2	16	0
C3	679	Owned	28		17	0
C1	3175	Owned	75	X4	17	1
⋮	⋮	⋮	⋮	⋮	⋮	⋮
C3	1484	Rented	20	X3	2	0
C24	2748	Owned	67	X1	17	?

- E.g. "What would happen if I recommend a different policy category?"
  - Standard ML models: what will be the response for this individual?
  - Interventional questions: what will be the response for this individual **if I set A to a**?
  - **Standard ML models are systematically biased for what-if questions**, but this is possible with TabPFN extension Do-PFN

Summary of TabPFN: Our GenAI Model for Tabular Data

TabPFN is a transformer-based foundation model that yields high-accuracy predictions without the need for task-specific training.

- **Speed:** single model, predictions in seconds
- **Works with less data:** 50% of the data suffices for same performance
- **Versatility:** handles missing values, outliers, uninformative features, categorical & numerical data
- **Fine-tunability to your tabular data sets:** just like finetuning LLMs to your text data

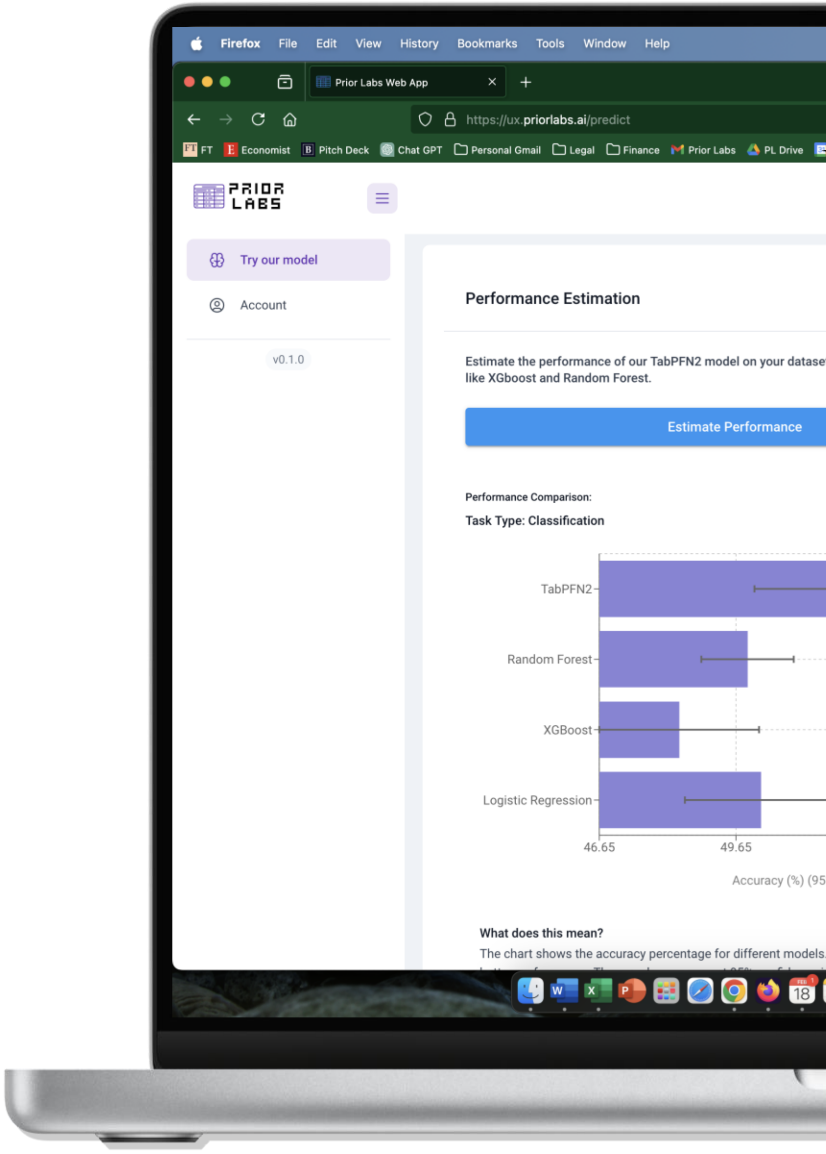

Outline

- Motivation
- TabPFN: The GenAI Revolution for Tabular Data
- ➡ Prior Labs: the Startup around TabPFN
- TabPFN for Insurance

GenAI has transformed text & images

But our most valuable data is organized in tables

Text

OpenAI
DeepMind
ANTHROP\C
deepseek
ALEPH ALPHA

Tables, Time Series & Databases

PRIOR LABS

Images

OpenAI
DeepMind
deepseek
Midjourney
black forest labs.

Prior Labs: the First Startup on Tabular Foundation Models

- **VC-backed startup**, founded November 2024

- Foundation model company – ***the OpenAI of tabular prediction***

- **€9M pre-seed funding** in February 2025
  - Europe's largest VC
  - Strategic VC: best-performing hedge fund
  - Top angels from HuggingFace, DeepMind, etc

- **Hired top talent** from
  - MIT, Oxford, Cambridge, ETH, …
  - Google, Meta, G-Research, QuantCo, …

Prior Labs products & capabilities

CURRENT MODELS

Available & under development

TabPFN v2

TabPFN-R

TabPFN-TS

TabPFN v3

CAPABILITIES / PREDICTION

Tasks that span across all models

Classification

Regression

Forecasting

Anomaly Detection

Data Generation

Multi-Table Reasoning

Fine-Tuning

Text Integration

Distribution Shifts

Interpretability

INDUSTRIES

Our models are used across industries

INSURANCE

FINANCE

HEALTHCARE

ENERGY

INDUSTRIALS

AND MANY OTHERS

APPLICATION AREAS

Exemplary use case: predict whether to intubate before it's too late

New AI Tool to support decisions around patient intubation

by University of Warwick

edited by Gaby Clark, reviewed by Andrew Zinin — Editors' notes

"TabPFN's accuracy in predicting the outcome of NIV has been extremely impressive. ***We are very enthusiastic about its potential to improve patient outcomes*** *and hope that further development will enable it to be rolled out across the NHS."*

Surgeon Commander Tim Scott

NHS
University Hospitals of North Midlands
NHS Trust

TabPFN is used in NHS University Hospitals today, set to save lives ❤️

Exemplary use case: immunoprofiling cancer treatments

BostonGene

- Machine-learning-powered blood test and analysis platform
- Identifies immune profiles and enables better treatment prediction

53%
error reduction

Differentiating blood samples of cancer patient vs healthy individuals (ROC 0.91 vs 0.81)

90%
time saved

TabPFN reduced model development times by 90% as no hyperparameter tuning was required even with new data.

Enhanced efficiency, allowing scientists to concentrate on meaningful clinical analysis rather than machine learning

Exemplary use case: Helping fraud check & credit risk teams

Taktile

- TabPFN is available on the **Taktile risk decision automation platform** to **enhance fraud checks and credit risk anomaly detection** for its clients

Simplification of ops with a foundation model

- No need to retrain when there is new data
- Significant savings on validation and infrastructure overhead

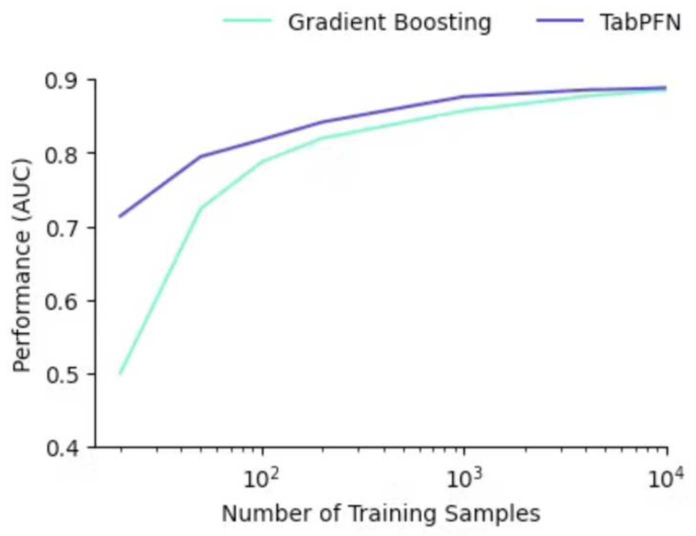

Fraud Checks

- Only few cases require manual intervention
  → sparse, small fraud datasets for which TabPFN works best

Credit Risk & Anomaly Detection

- Limited & irregular data for credit risks and anomalies

Large interest from algorithmic trading firms

Sadly, I can’t say more about that …

Ease of use: Google sheets integration

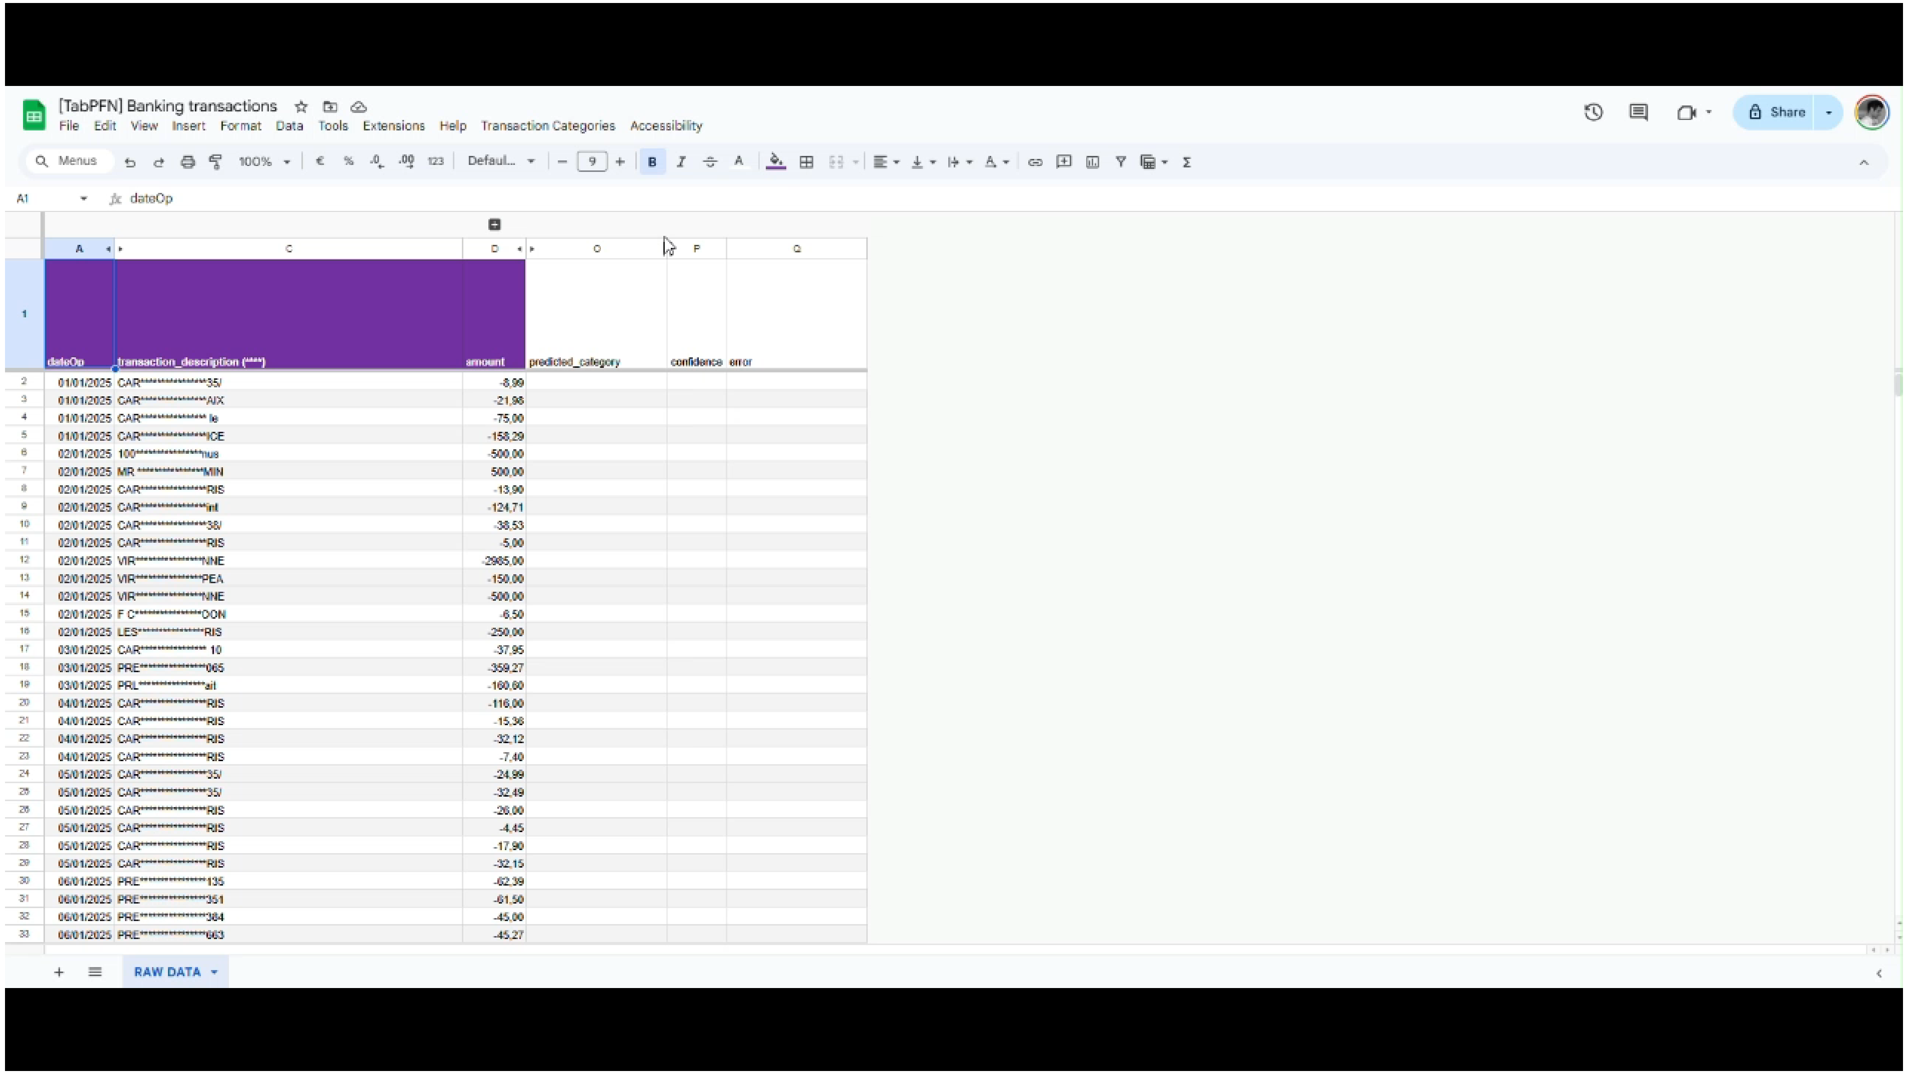

Contribution by the open source community – Benjamin Lalanne ☺

Outline

- Motivation
- TabPFN: The GenAI Revolution for Tabular Data
- Prior Labs: the Startup around TabPFN
- ➡ TabPFN for Insurance

A New Benchmark of 10 Public Insurance Datasets

- **Car Insurance Claim Prediction**
  - Based on demographics, technical details of vehicle
  - Based on demographic + driving info
  - Based on demographics, vehicle details, details on incidents
- **Cross selling**
  - Car Insurance Cold Calls (demographics, financial info, campaign details)
  - Caravan Insurance (aggregate data: other policy holders in zip-code area)
  - Health Insurance (demographics, insurance type, policy details)
  - Travel Insurance (demographics, behavioural features)
- **Travel insurance**
  - Speed of settling corporate claims (type & magnitude of incidents, outstanding reserves)
  - Lapse (demographics, health conditions, travel history)
- **Prudential Life Insurance**
  - Risk category (demographics, employment & insurance history, medical data)

Results on this Insurance Benchmark

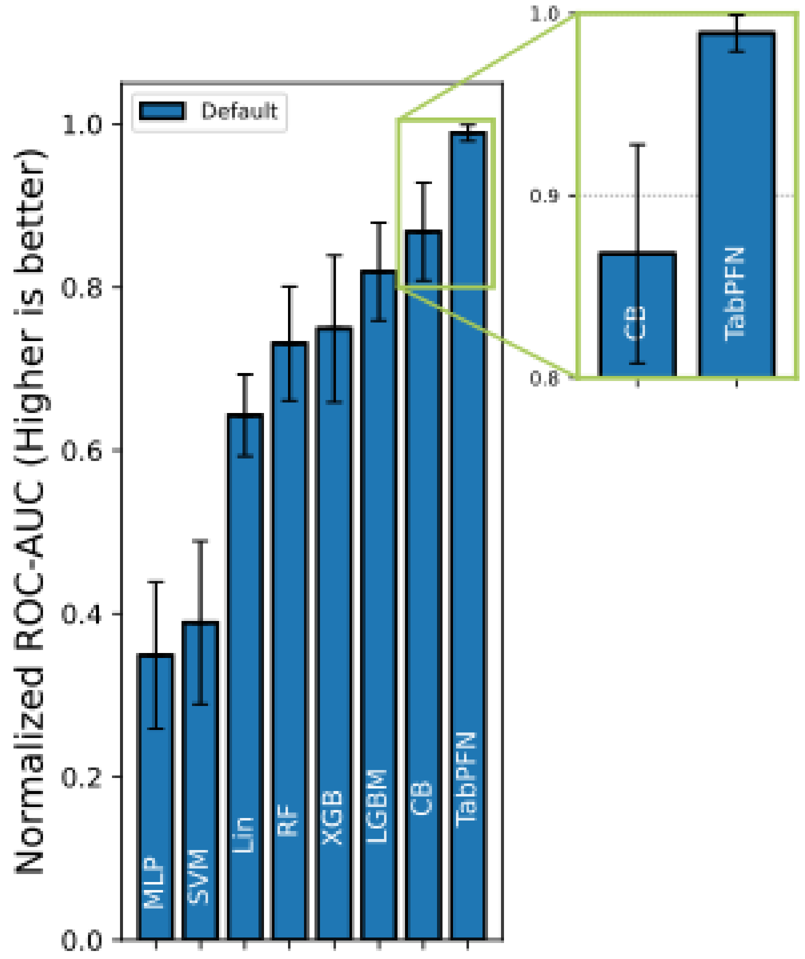

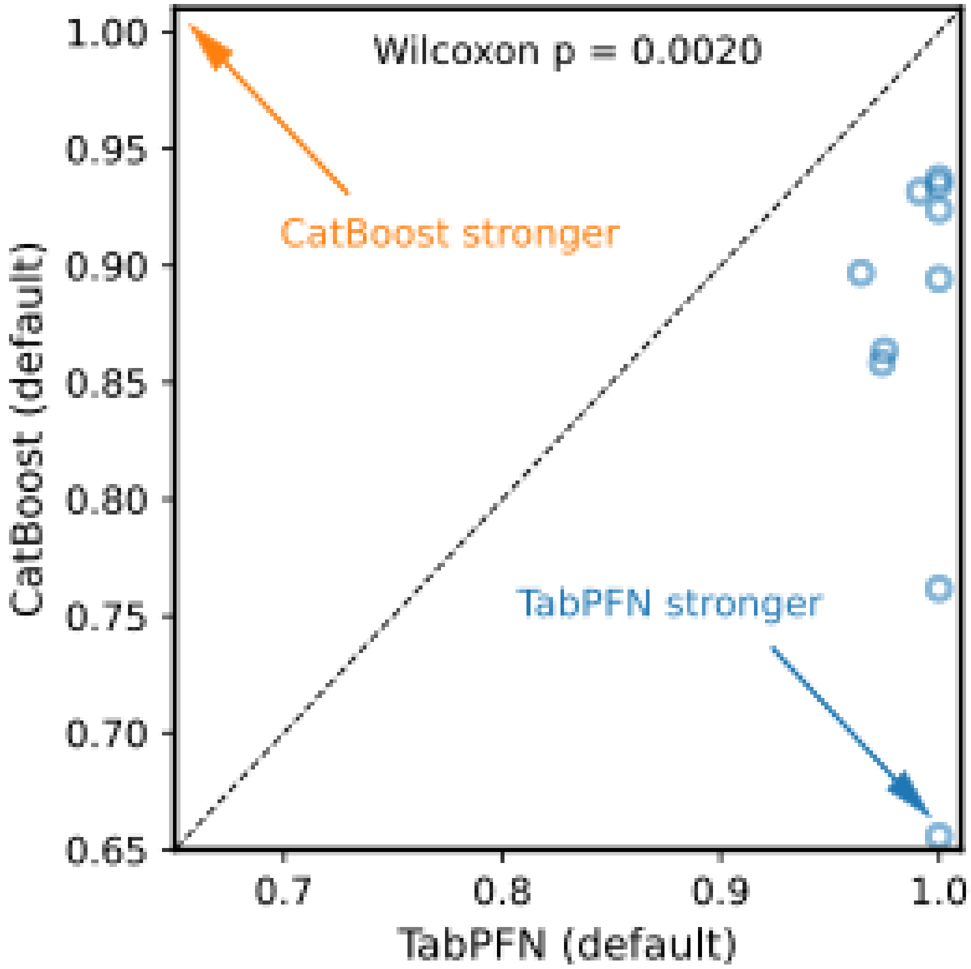

TabPFN substantially outperforms the baselines in all 10 cases

Example: prudential life insurance use case

- Goal: replace lengthy (30-day) process to obtain life insurance with instant quotes
- Accurate risk classification is crucial
- Fine-tuning TabPFN improves yet further

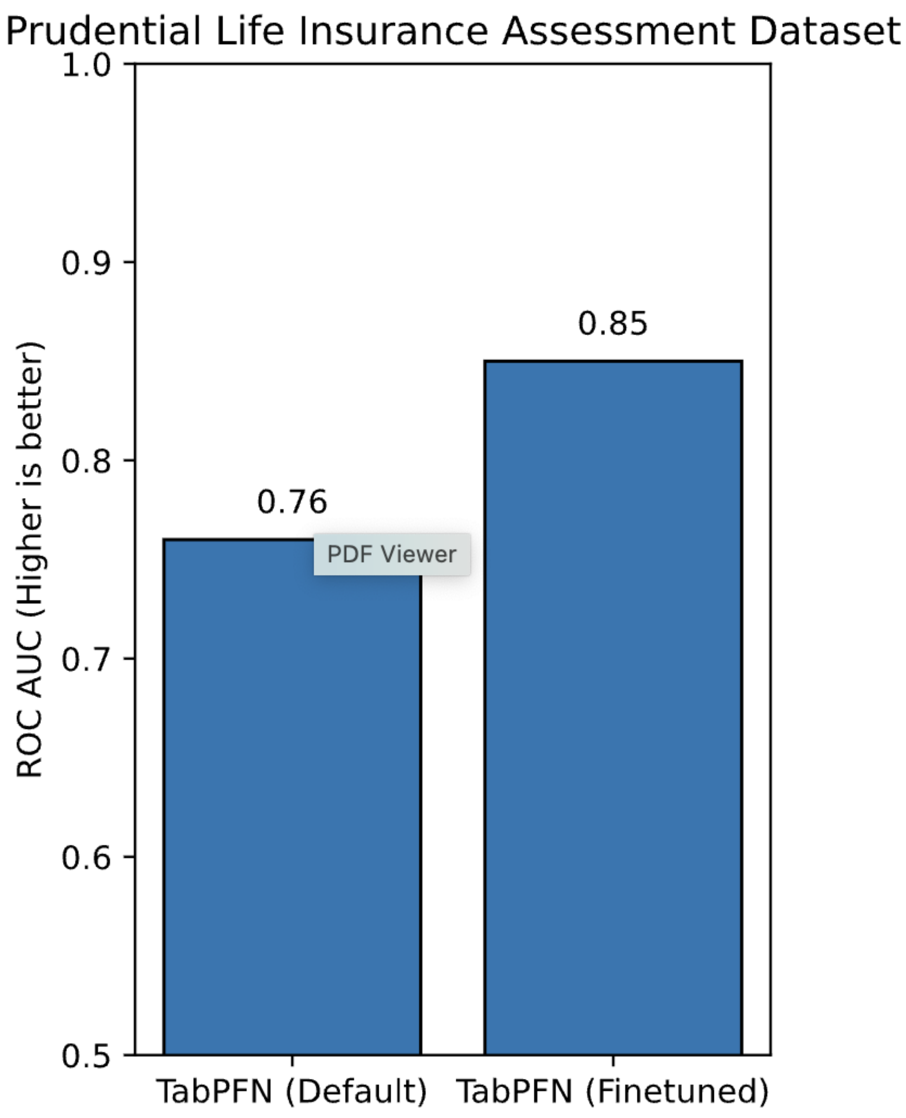

How to Get Started with TabPFN

[Web Client / API]

[Local / On-Prem]

Private VPC

Collaboration & Commercial Support

- Custom licensing
- Access to TabPFN v2.5+ with instant upgrades
- Fine-tuning on proprietary data
- Dedicated support with implementation help
- Joint case study & co-branded success story
- Educational content:
  webinars, internal research talks

The Road Ahead

- **Better models**
  - We're continuously improving our models
- **Finetuning on company-specific insurance data**
  - Creating proprietary models for our partners
- **Partnership to build a tabular foundation model specific to insurance**
  - Integrate all available data to build a model from scratch
- **Agentic system**
  - Support in data gathering, preparation & cleaning
  - Focussed on data analysts & data-savvy managers
  - Chat with your data

Take aways

- **Tabular data is ubiquitous** in insurance applications
- **GenAI revolution** of predictive analytics on tabular data
- **Open source insurance benchmark**
  - Substantial improvements across 10 public datasets
  - Very easy to try out – open source code
- **Entirely new possibilities**
  - Single model
  - Fine-tuning to your company's specific data

priorlabs.ai/insurance