Foundation Models for Structured Business Data

Jure Leskovec
Co-Founder and Chief Scientist at, Kumo
Professor at Stanford University

Kumo

The Two Halves of the Enterprise Brain

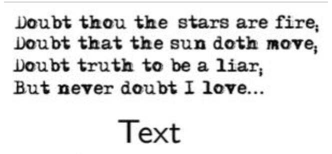

Unstructured Data

(Language & Perception)

- Text, documents, images
- Reasoning, summarizing, and generating content
- Answer the "what" and "why" from human knowledge

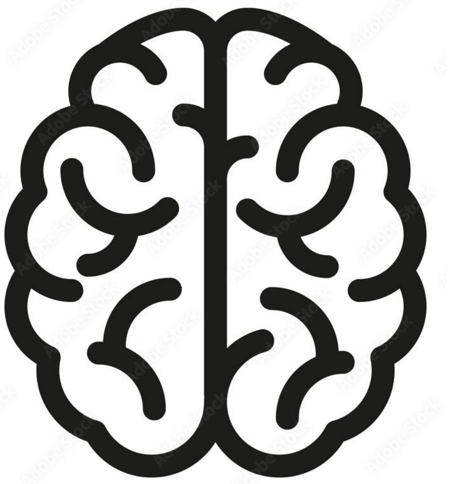

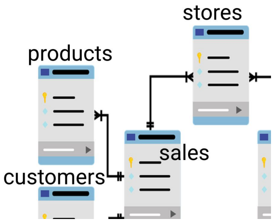

Structured, Relation Data

(Business Operations)

- Databases of customers, products, transactions, and supply chains
- This data holds the patterns that predict what will happen. It's the ground truth of the business.

Kumo

The AI revolution is Incomplete

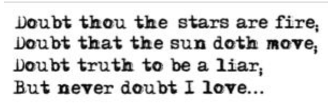
Doubt thou the stars are fire;
Doubt that the sun doth move;
Doubt truth to be a liar;
But never doubt I love...

Text

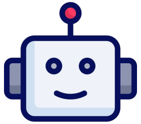

LLMs

LLM - Reasoning

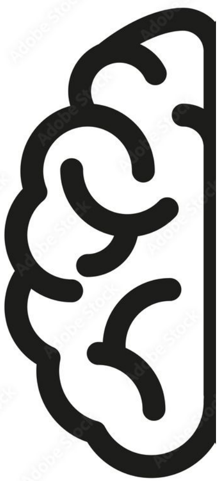

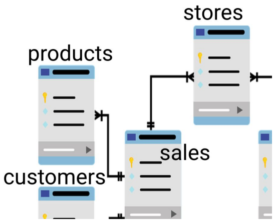

Reasoning Brain:
Understands "*what is in this contract?*"

Predictive/Analytical Brain:
Answers "*which customer will churn next?*"

Kumo

The AI revolution is Incomplete

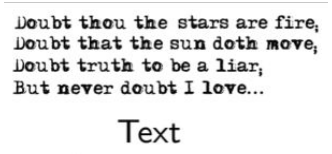
Doubt thou the stars are fire;
Doubt that the sun doth move;
Doubt truth to be a liar;
But never doubt I love...

Text

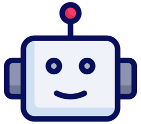

LLMs

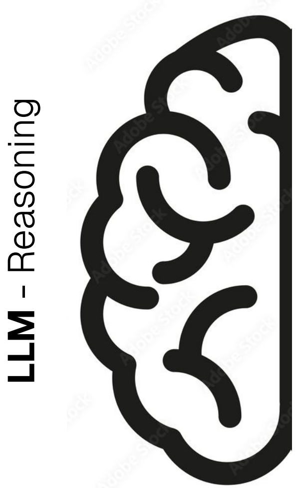

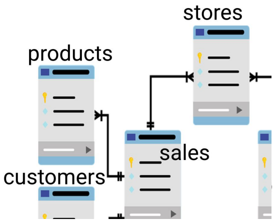

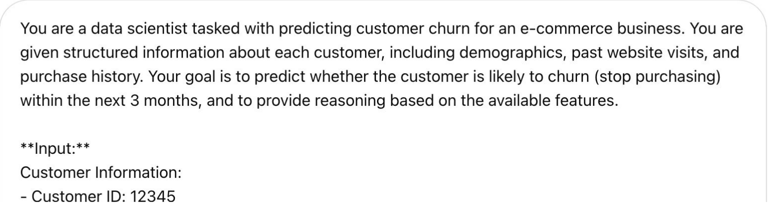
You are a data scientist tasked with predicting customer churn for an e-commerce business. You are given structured information about each customer, including demographics, past website visits, and purchase history. Your goal is to predict whether the customer is likely to churn (stop purchasing) within the next 3 months, and to provide reasoning based on the available features.

Input:
Customer Information:
- Customer ID: 12345

LLMs cannot reason
over structured data effectively

Reasoning Brain:
Understands "*what is in this contract?*"

Predictive/Analytical Brain:
Answers "*which customer will churn next?*"

Kumo

The AI revolution is Incomplete

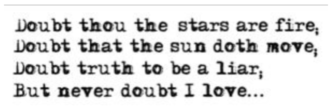
Doubt thou the stars are fire;
Doubt that the sun doth move;
Doubt truth to be a liar;
But never doubt I love...

Text

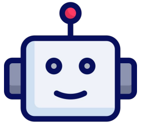

LLMs

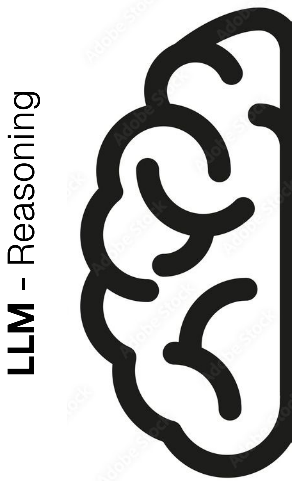

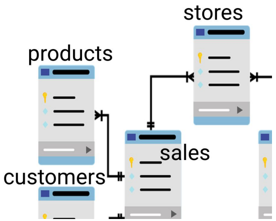

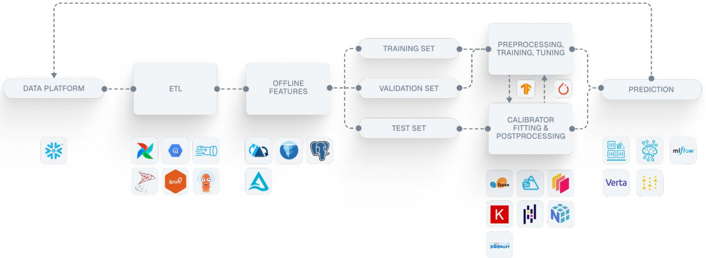

Manually build an ML model per task

Reasoning Brain:
Understands "*what is in this contract?*"

Predictive/Analytical Brain:
Answers "*which customer will churn next?*"

Kumo

Core Problem & Our Insight

We need both halves of the brain to work together!

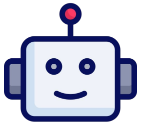

LLMs

LLM - Reasoning

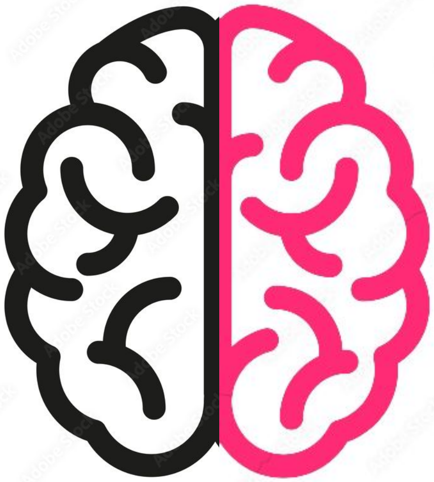

RFM - Prediction

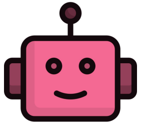

Relational Foundation Models

Reasoning Brain (LLM):
Understands "*what is in this contract?*"

Predictive/Analytical Brain:
Answers "*which customer will churn next?*"

Relational Foundation Model

Relational Foundation Model (RFM) is designed specifically for structured business data to make powerful forecasts & predictions.

You can simply point RFM at your data and ask predictive questions:

- *'Which leads are most likely to convert in the next 30 days?'*
- *'What products are at the highest risk of a stockout next quarter?'*
- *'Which transactions show the highest probability of fraud?'*

Kumo

RFM: The missing piece in the AI puzzle

RFM is powered by Graph Transformer Models

Simply connect your database

One universal model for business data.

Ask predictive questions

Kumo

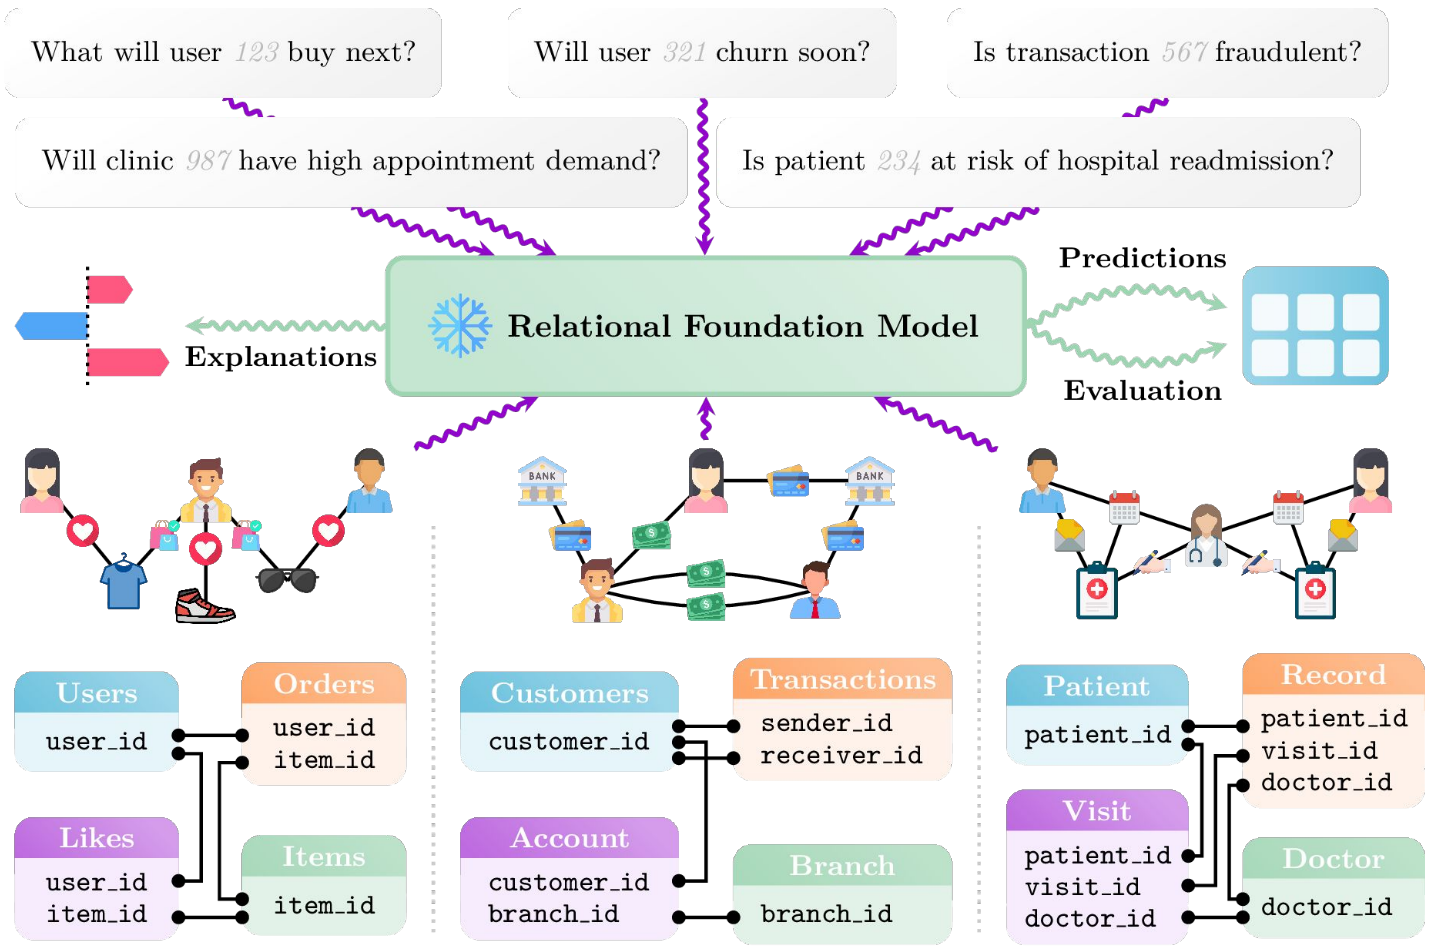

(a) E-Commerce (b) Finance (c) Health-Care

Kumo

Relational Graph Transformer

Key Observation: Kumo turn the web of connections around a data point into a sequence, like a sentence, and feed it into an attention network (like ChatGPT).

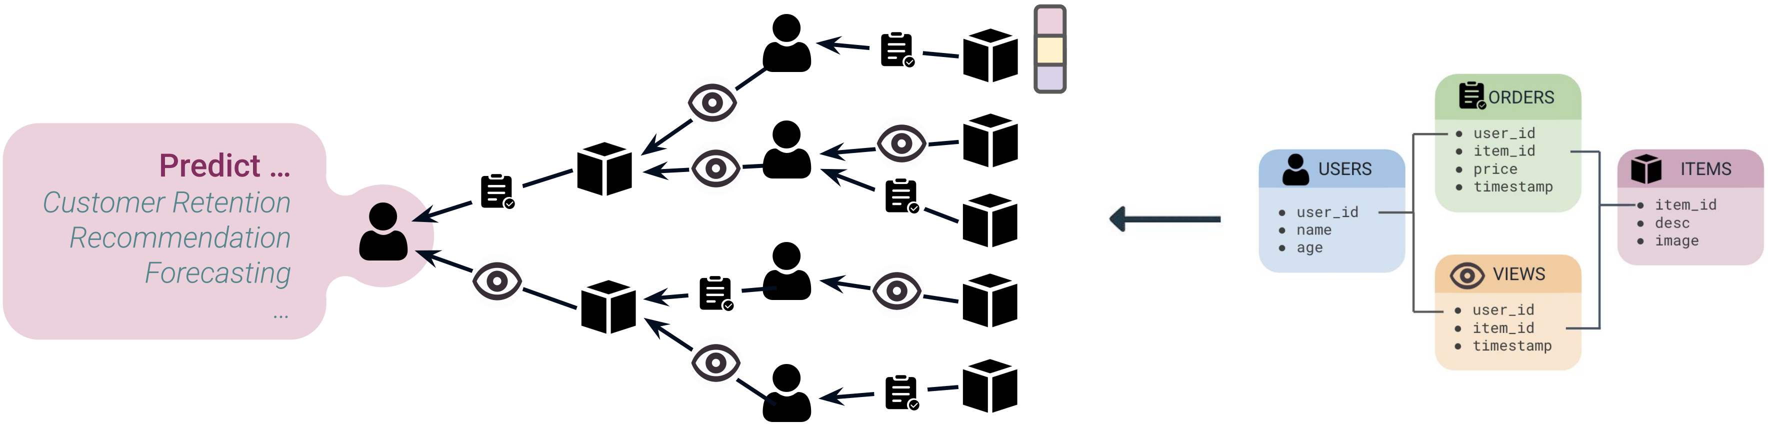

Kumo

Relational Graph Transformer

Key Observation: Kumo turn the web of connections around a data point into a sequence, like a sentence, and feed it into an attention network (like ChatGPT).

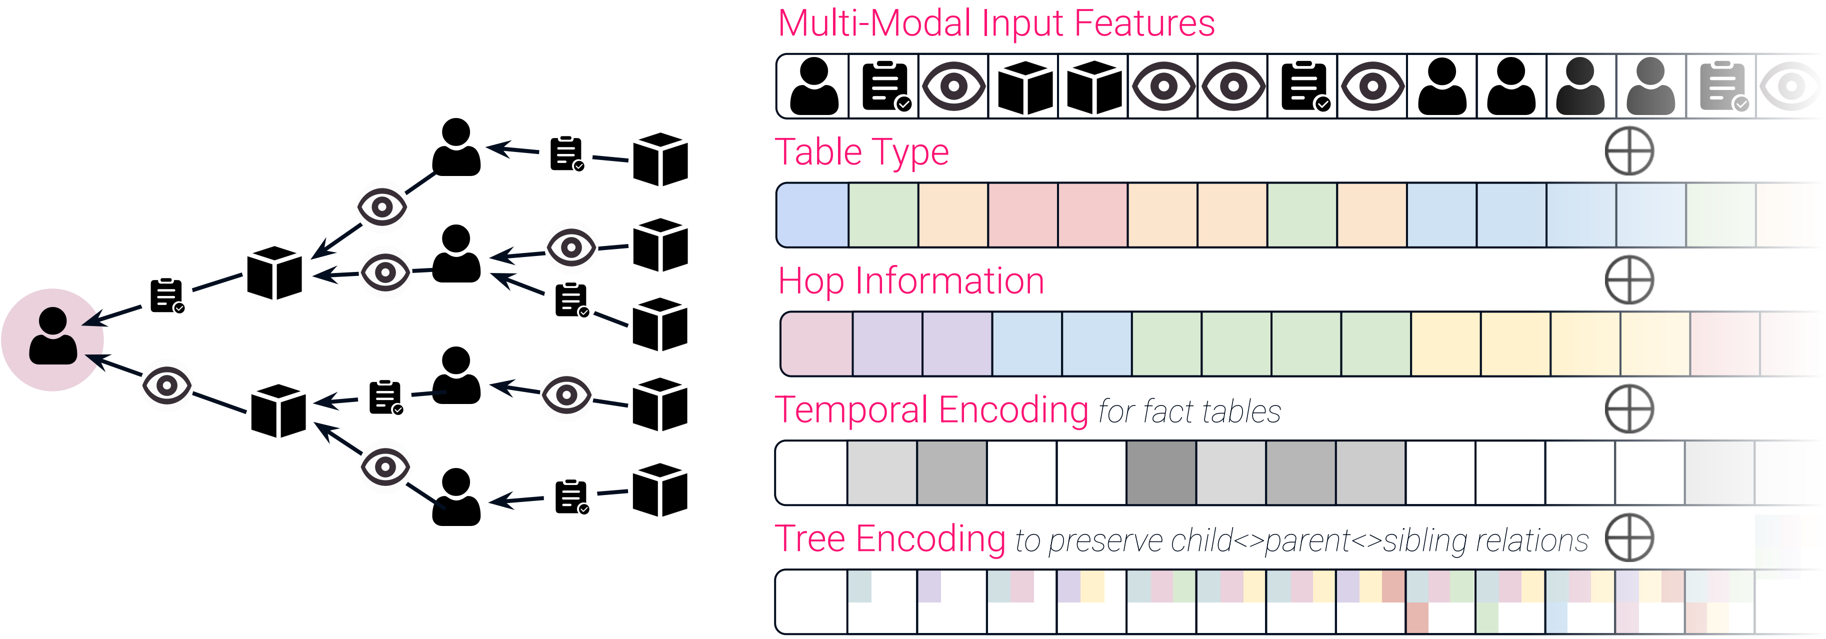

Kumo

Foundational Model for Relational Data

Chat

box

Google Drive

Doubt thou the stars are fire;
Doubt that the sun doth move;
Doubt truth to be a liar;
But never doubt I love...

Text

Natural Language Prompt
Context: Unstructured Docs

Common-sense in-context reasoning

Text output

Kumo

Foundational Model for Relational Data

Chat

box

Google Drive

Doubt thou the stars are fire;
Doubt that the sun doth move;
Doubt truth to be a liar;
But never doubt I love...

Text

Natural Language Prompt
Context: Unstructured Docs

Common-sense in-context reasoning

Text output

Relational Foundation Model:

snowflake

databricks

Google BigQuery

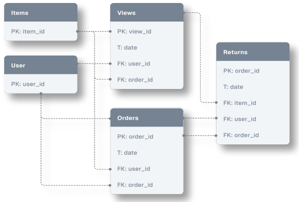

Relational Prompt
Context: Relational Data
Predict at-risk customers

Relational in-context reasoning

AI Predictions
Embeddings output

Kumo

We can talk to KumoRFM with Natural Language when combined with LLMs via MCP

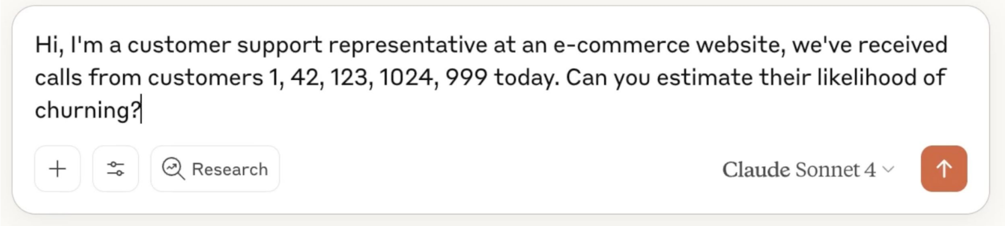

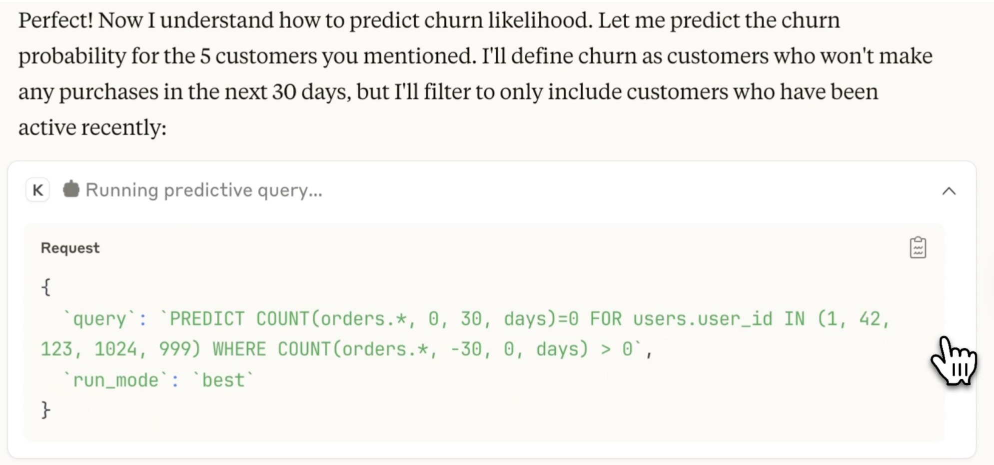

Kumo

We can talk to KumoRFM

first categorical value

```
PREDICT FIRST(orders.type, 0, 7)
FOR EACH users.user_id IN (0, 1, 2)
```

logical value

```
PREDICT COUNT(orders.*, 0, 7) > 0
FOR EACH users.user_id IN (0, 1, 2)
```

sum of numerical values

```
PREDICT SUM(orders.value, 0, 7)
FOR EACH users.user_id IN (0, 1, 2)
```

set of foreign keys

```
PREDICT LIST_DISTINCT(orders.item_id, 0, 7)
FOR EACH users.user_id IN (0, 1, 2)
```

- Interact via Predictive Query Language, capable of a broad set of task types
- It has a label definition (PREDICT clause) and entity definition (FOR clause)
- Additional filters and aggregations can be applied both to label and entity clause

Kumo

Kumo Relational Foundation Model

Get predictions from data in real-time with a few lines of code. No ML expertise required.

Personalized Recommendations	Predict customer's purchases next week	Predict now Evaluate

Check out the quick tutorial and try more use cases yourself!


```
PREDICT COUNT(orders.*, 0, 30, days) > 0 FOR users.user_id=1001
```

Kumo

Benefits of RFMs

Better performance:

+10% accuracy improvement through relational context

Faster time to value:

95% reduction in data preparation effort, handles any tasks on the fly

Scalability:

Handles complex, relational data across billions of records

RelBench Datasets Benchmarks

(12 tasks: Amazon, H&M, StackExchange, Clinical, Avito, Hangtime)

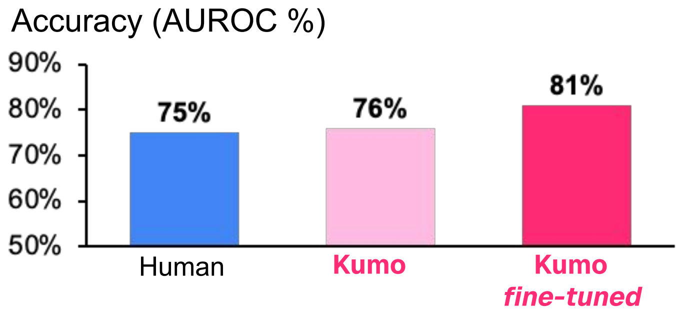

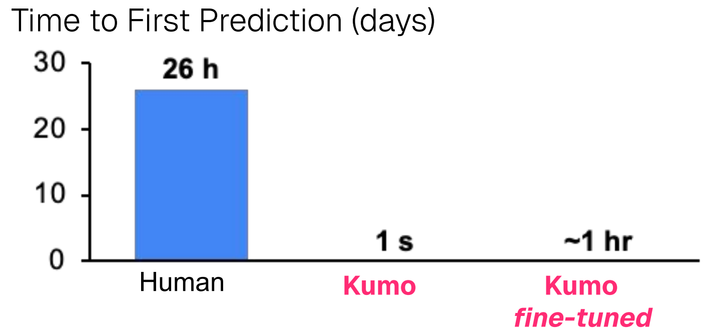

RELBENCH RELATIONAL DEEP LEARNING BENCHMARK by Stanford

Our Unified Architecture

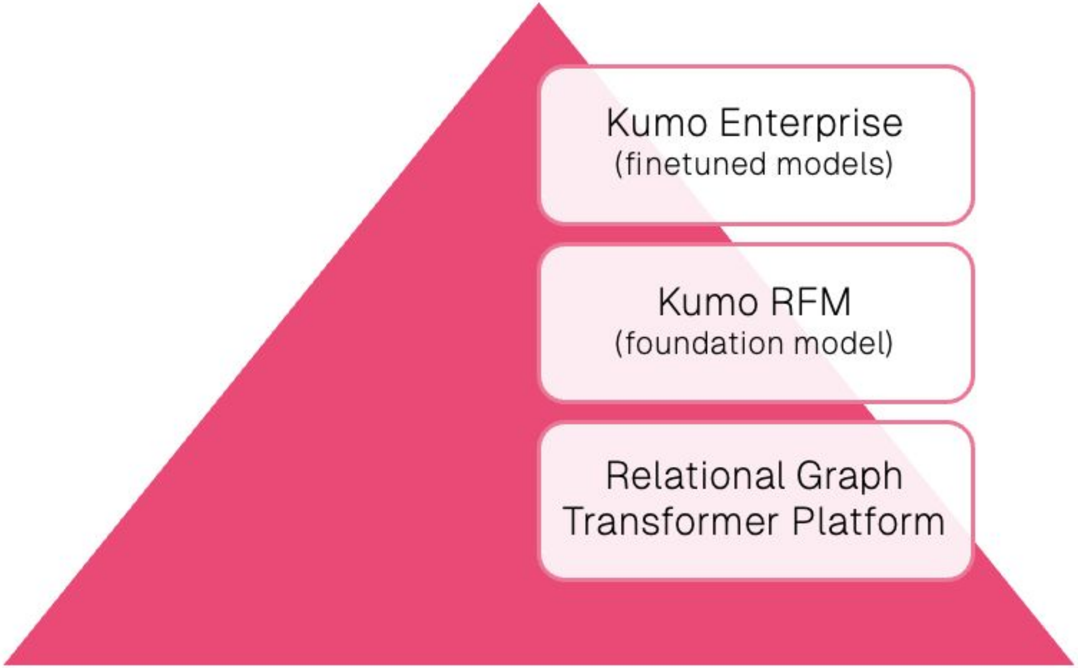

Last year 6.5k jobs, this year on track to 25k jobs.

In a week in July '25: 262 models tuned, 405 inference jobs, 767 trillion nodes, 2079 trillion edges (82PB data)

The Kumo Ecosystem

We are the chosen predictive engine for the data ecosystem

Platform Partners

- Amazon Bedrock
- SAP
- snowflake
- SIEMENS
- salesforce
- databricks

Customers

- P&G
- coinbase
- chime
- DOORDASH
- reddit
- tubi
- FAIRE
- Allstate
- databricks
- Discord
- snowflake

...

Kumo

RFMs address an LLM capability gap

- Not predicting text from text but outcomes from business data
- Just like LLMs replaced time-consuming, difficult, and expensive NLP work, RFMs replace time-consuming, difficult, and expensive predictive modeling work
- Since RFMs fills a gap, it's not replacing or competing with LLMs.
- RFMs are complementary to LLMs and most businesses will end up using both

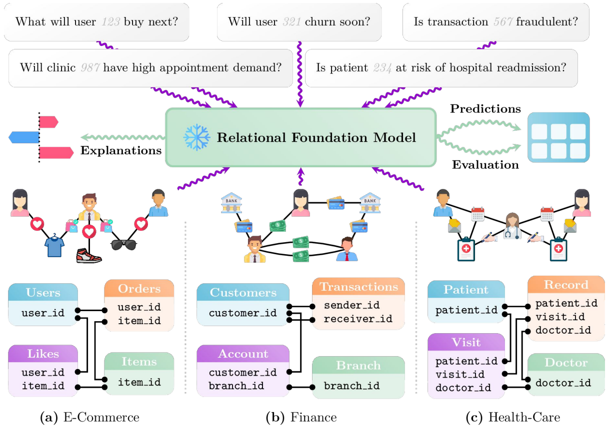


```python
from kumoai.rfm import Graph, KumoRFM

df_dict = {
  'users': pd.DataFrame(...),
  'orders': pd.DataFrame(...),
  'items': pd.DataFrame(...),
}

graph = Graph.from_data(df_dict)
rfm = KumoRFM(graph)

query = "PREDICT COUNT(orders.*, 0, 30)>0 FOR users.user_id=1"
rfm.query(query)
```

Building with Kumo

Kumo

Agents need predictive AI tools

First generation agents:

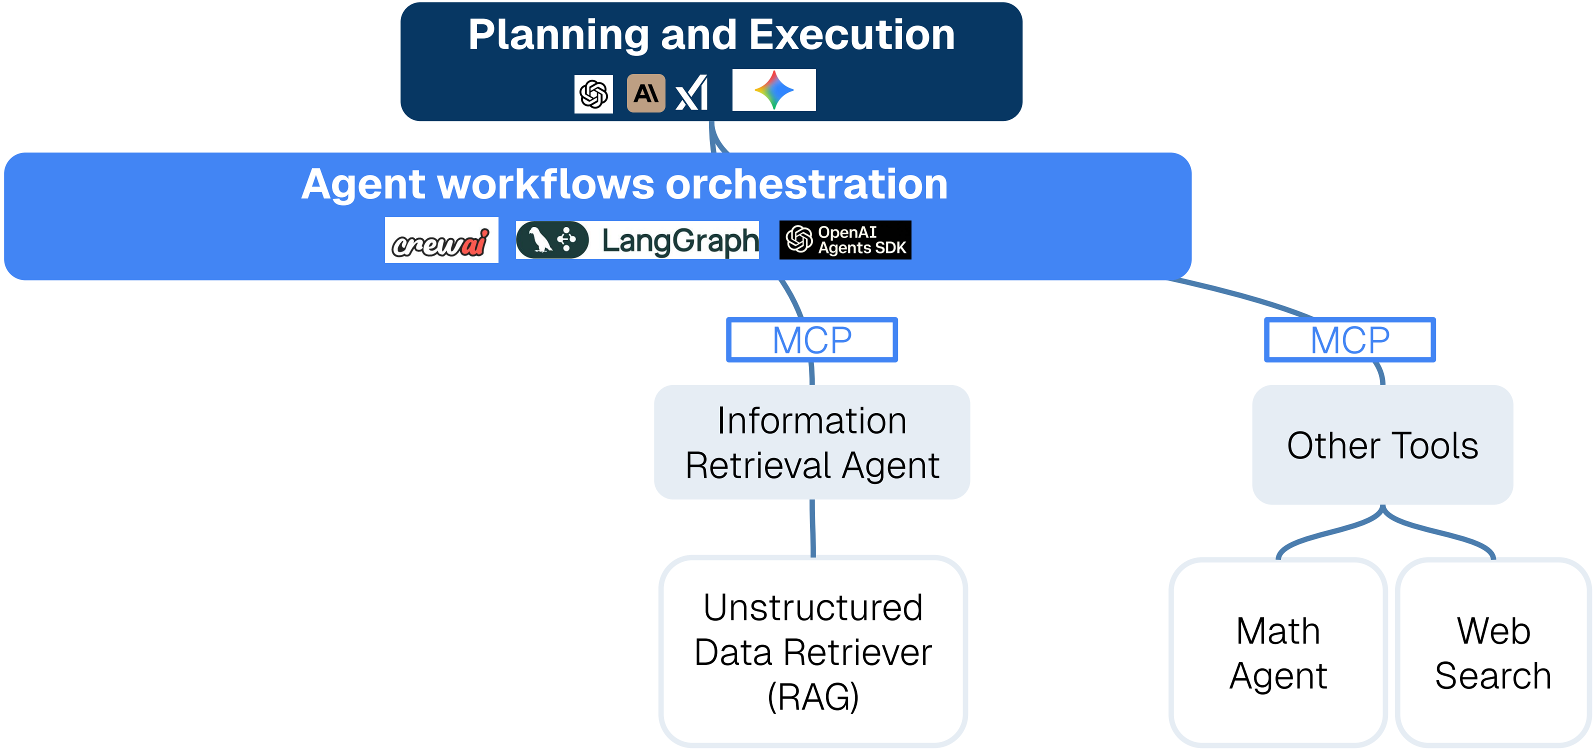

Agents need predictive AI tools

Second generation agents:

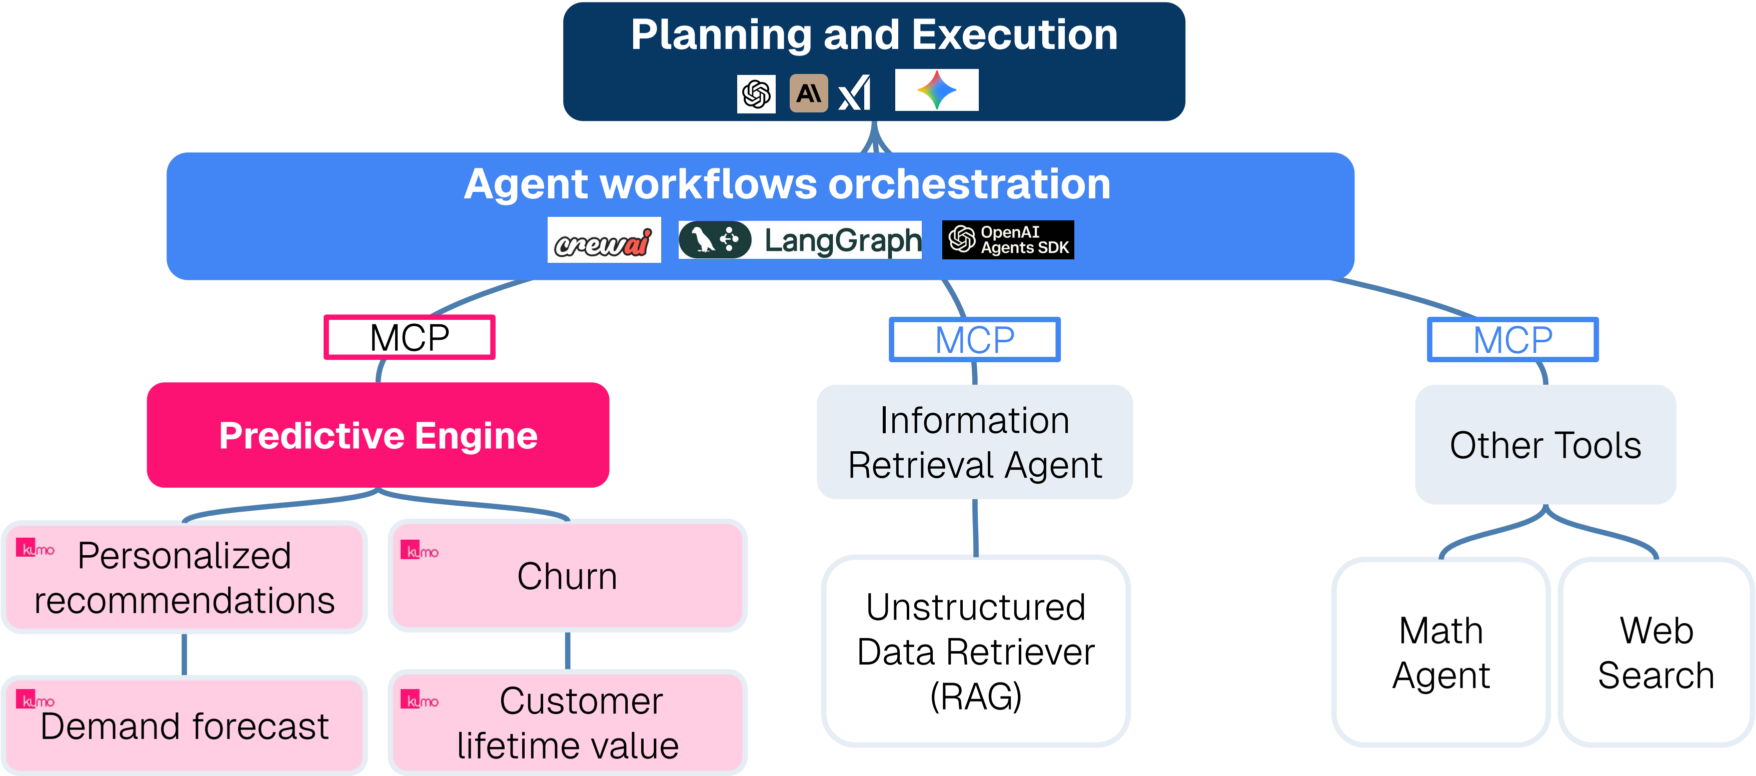

Example: Insurance Renewals Agent

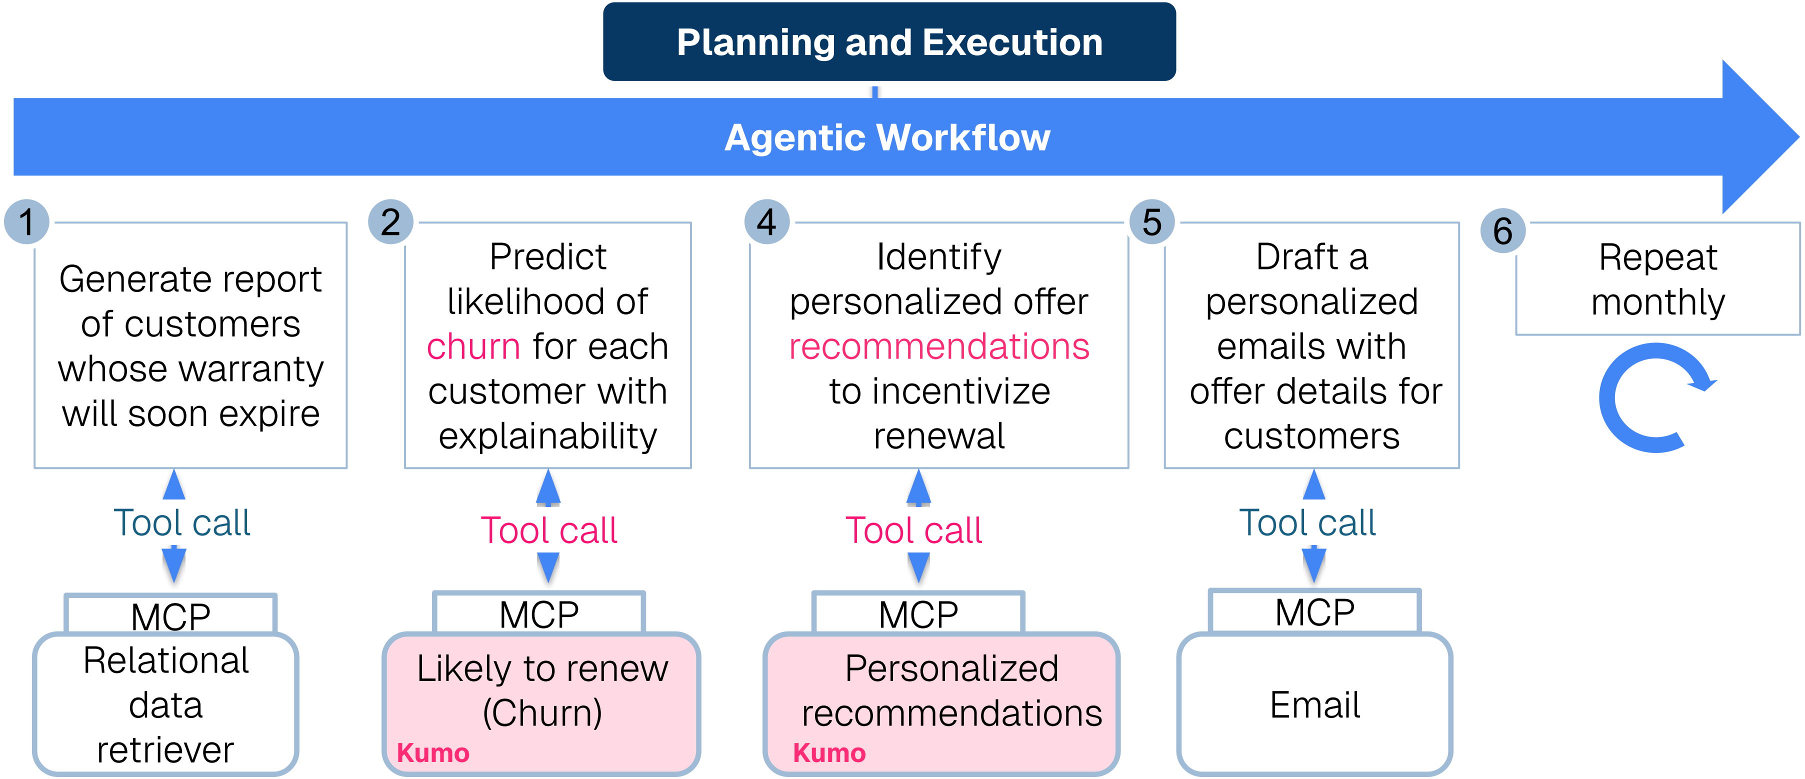

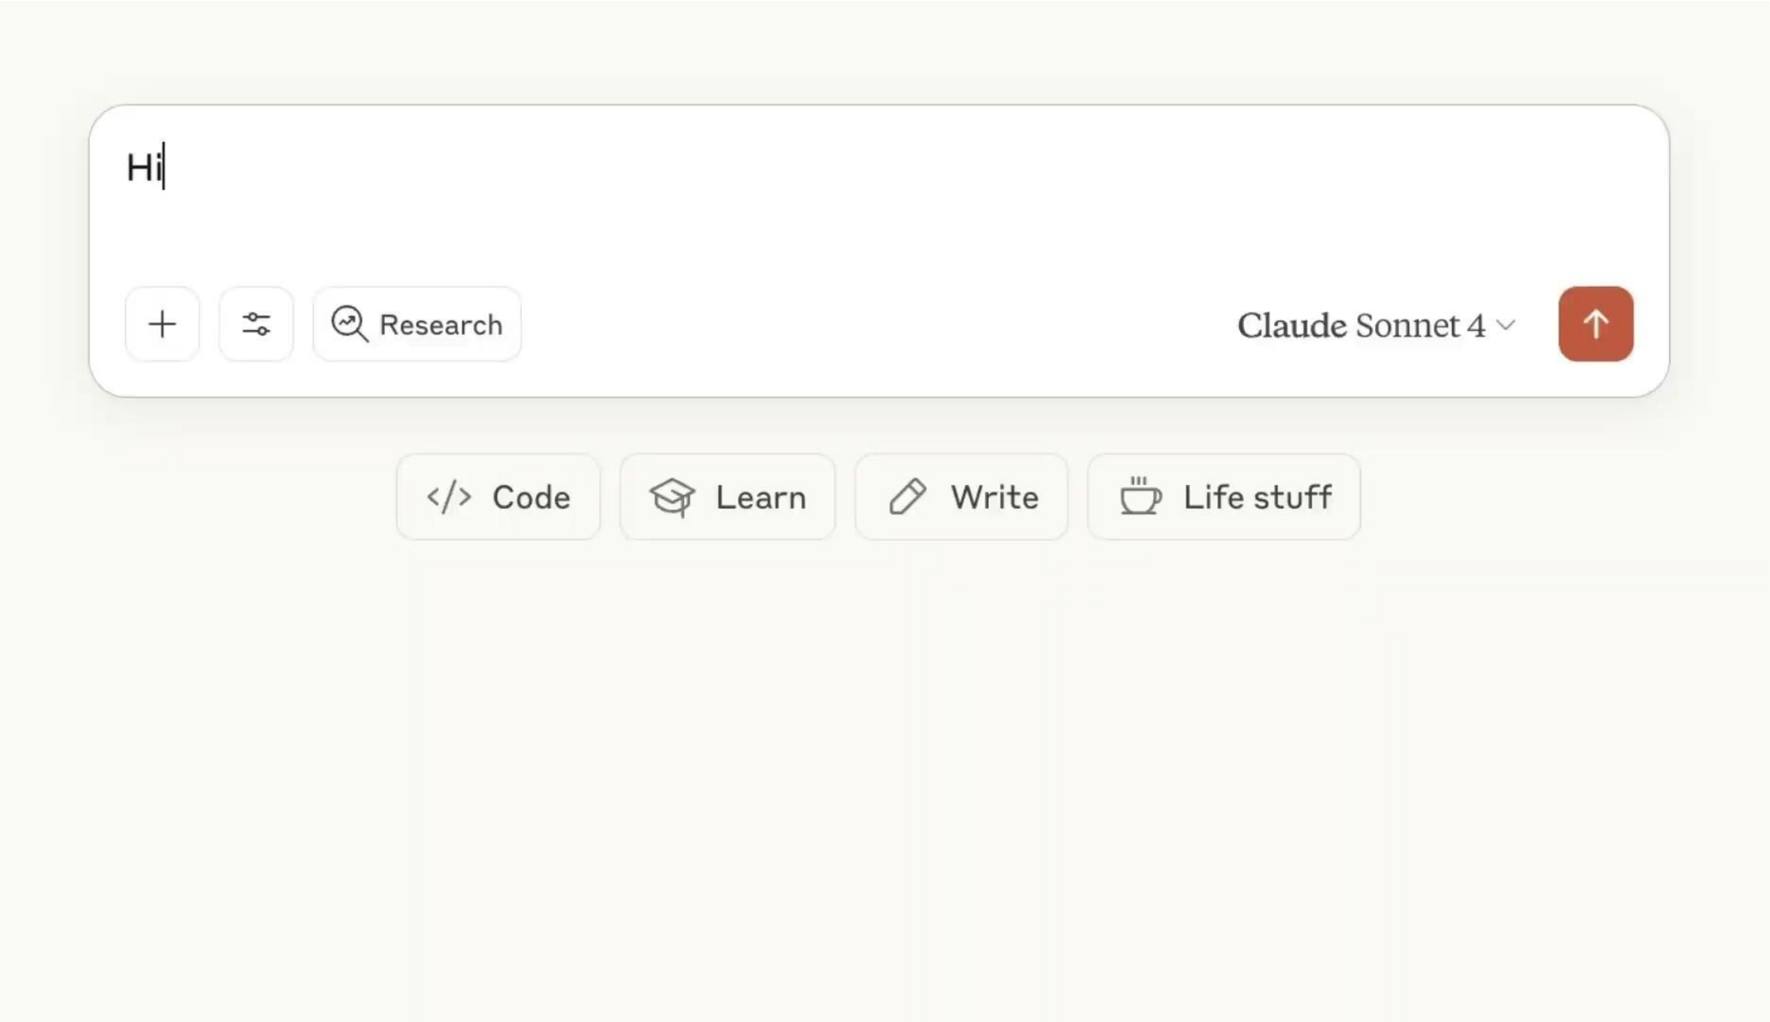
Hi
Research
Claude Sonnet 4
Code
Learn
Write
Life stuff

Kumo

Agents for Sales and Marketing

Customer Behavior Predictions

Churn Risk

Predict which customers are likely to leave based on usage patterns, support tickets, and purchases.

Upsell / Cross-sell Opportunities

Identify which customers are most likely to buy additional or complementary products.

Sales Forecasting

Pipeline forecasting

Predict the probability of deals closing, improving revenue forecasting accuracy.

Lead Scoring

Rank leads by conversion likelihood, so sales teams focus on the most promising prospects.

Marketing Optimization

Campaign Response Predictions

Forecast which customers are most likely to engage with specific campaigns.

Customer Satisfaction Pred.

Estimate Net Promoter Score (NPS) before feedback to proactively engage detractors

Sales Operational Efficiency

Sales Rep Performance Pred.

Forecast which reps may need coaching/training based on deal activity data.

Support Resource Planning

Predict workload in customer sales support to allocate staff efficiently.

Kumo

Sales and Marketing … myriad more use-cases

Customer Behavior Predictions		Sales Forecasting	Marketing Optimization	Service & Support	Sales Operational Efficiency
Churn Risk	Next Best Action (NBA) Prediction	Pipeline Forecasting	Campaign Response Predictions	Ticket/Escalation Prediction	Sales Rep Performance Predictions
Engagement Risk	Feature Adoption Prediction	Deal Closing Probability & Timeline	Personalized Recommendations	Customer Satisfaction Prediction	Sales Rep Training Recommendation
Upsell / Cross-sell Opportunities	Meeting / Demo Attendance Likelihood	Lead Scoring	Optimal Contact Timing	Self-service Deflection	Resource Planning
Customer Lifetime Value (CLV)	Customer Advocacy / Reference	Quota Attainment	Optimal Pricing or Discount Prediction	Customer Case Routing Prediction	Revenue Leakage Prevention
Early Fraud / Abuse Detection	Predictive Relationship Modelling	Cross-Team Collaboration Prediction	Click Through Prediction	Support Volume Forecasting	Data Quality Degradation Prediction

Kumo

Agents for the Enterprise

Every category of enterprise software will be enhanced with agents and require forward-looking predictions

Sales and Marketing Agents	Workforce & HR	Finance & Compliance
Operations & Supply Chain Expansion	Industry-Specific Agents (health, energy, ...)	

Kumo

A traditional ML approach does not scale

Model permutations become intractable

Per customer features

Custom fields or schemas (e.g., SAP Z fields, Salesforce custom fields)

Per customer models

Legal constraints prohibiting cross-training across customer data

Per region/ site models

Differing schemas across BUs / subsidiaries

Disparate systems

e.g., a result of M&A

What kinds of projects are people building with KumoRFM?

Kumo

Health coverage matching

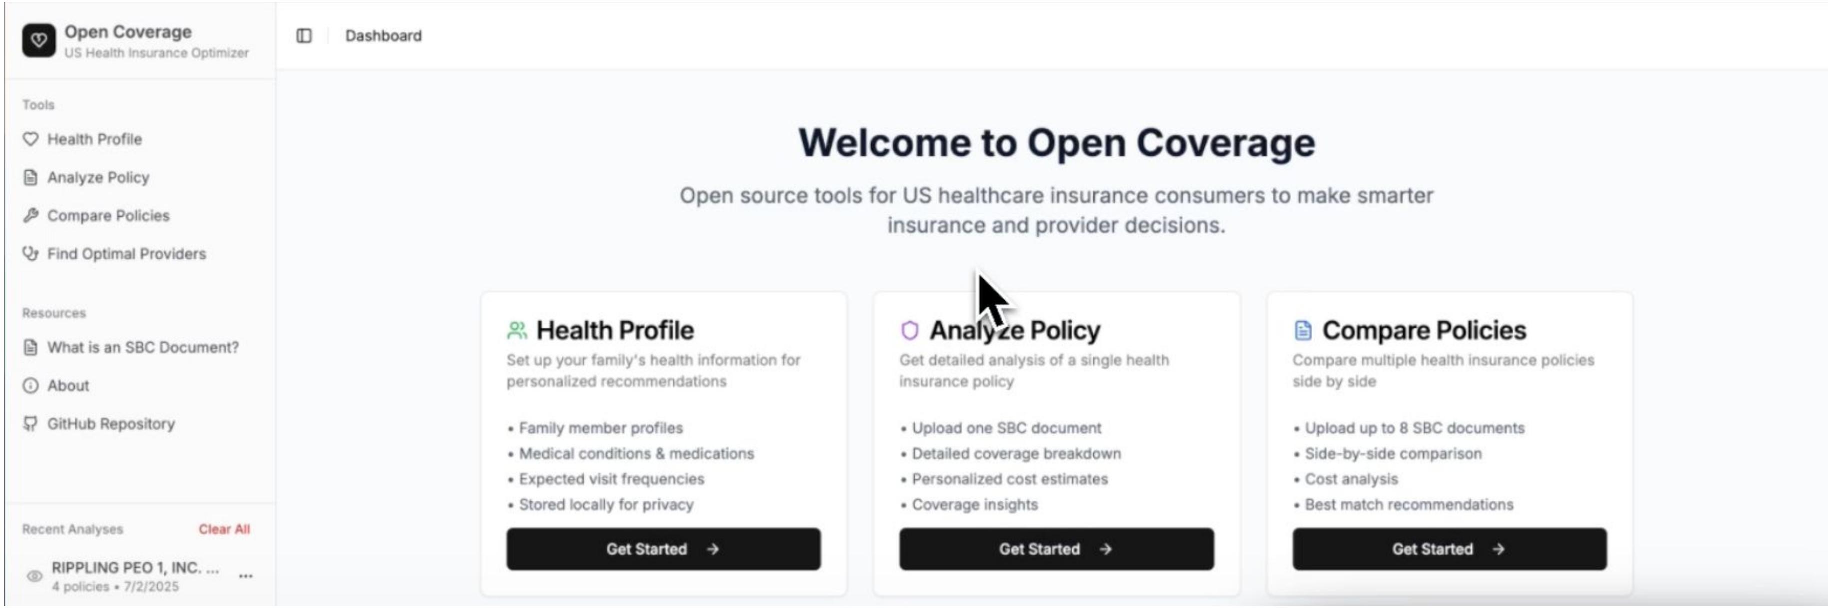

An application that empowers U.S. consumers to make smarter insurance decisions and reduce out-of-pocket spending through personalized healthcare utilization forecasting and preference matching.

Kumo

Energy grid optimization

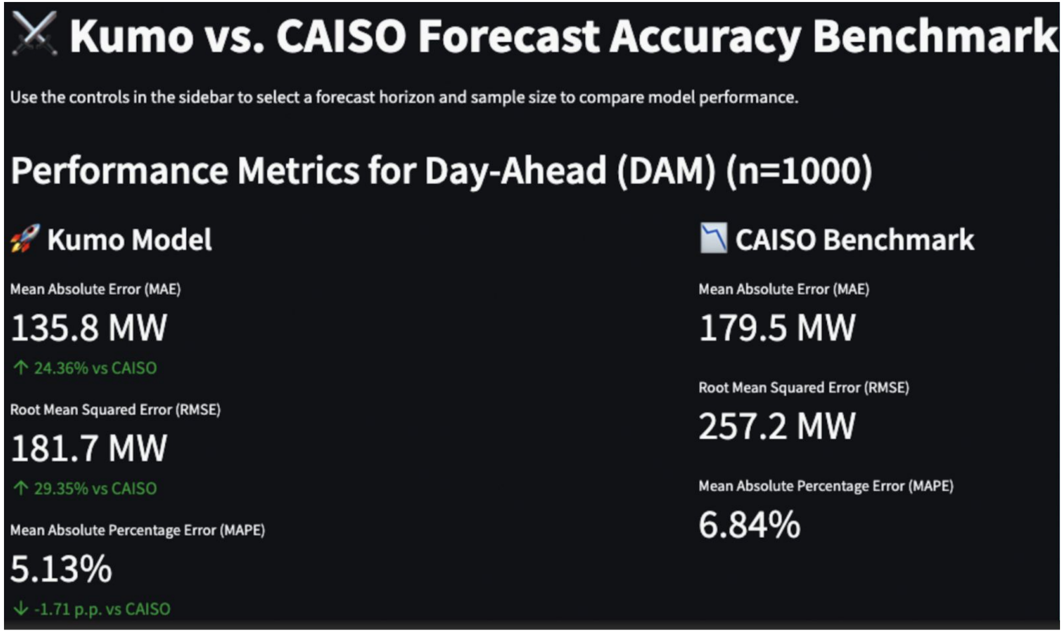

A localized energy prediction platform that outperforms institutional forecasts by 30%, with applicability to reduce electricity over purchasing from wholesalers by Load Serving Entities.

Kumo

Wildfire risk assessment predictions

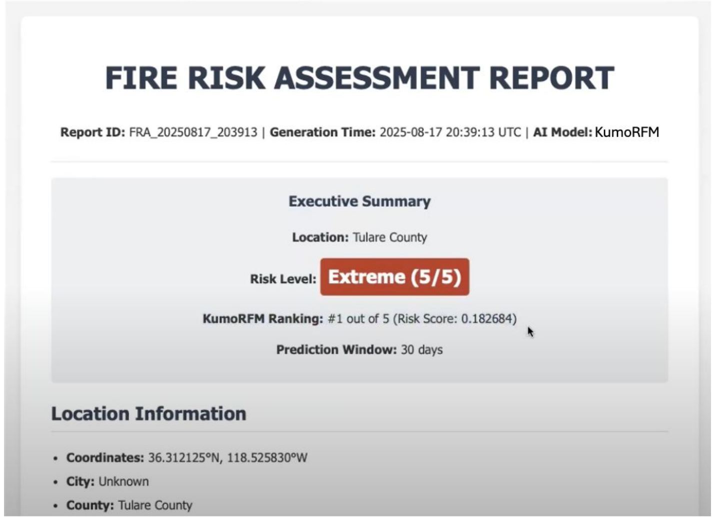
FIRE RISK ASSESSMENT REPORT

Report ID: FRA_20250817_203913 | **Generation Time:** 2025-08-17 20:39:13 UTC | **AI Model:** KumoRFM

Executive Summary

Location: Tulare County

Risk Level: Extreme (5/5)

KumoRFM Ranking: #1 out of 5 (Risk Score: 0.182684)

Prediction Window: 30 days

Location Information

- **Coordinates:** 36.312125°N, 118.525830°W
- **City:** Unknown
- **County:** Tulare County

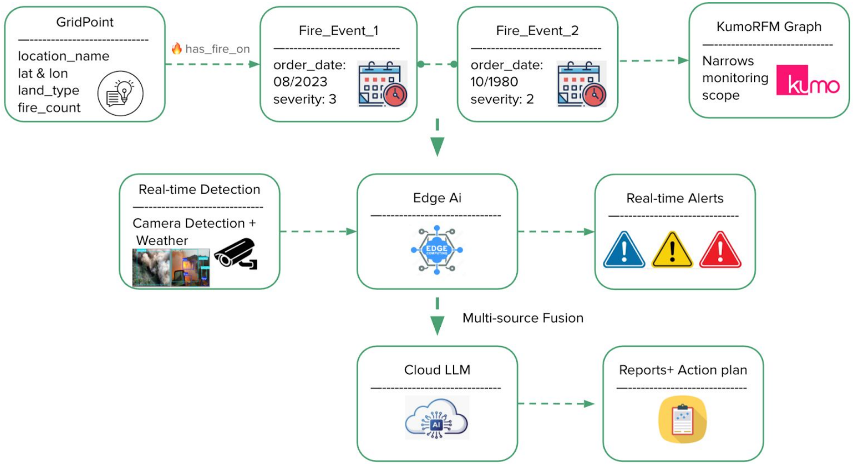

This three-layer AI architecture combines predictive modeling, real-time detection, and intelligent decision-making for comprehensive wildfire risk management.

A wildfire risk prediction system, combining KumoRFM with real-time edge sensor data and large language models. KumoRFM serves as the prediction engine to forecast wildfire risk at a granular, geo-grid-based level, aiding wildfire proactive measure prioritization and response planning.

Kumo

Hospital Capacity Planning

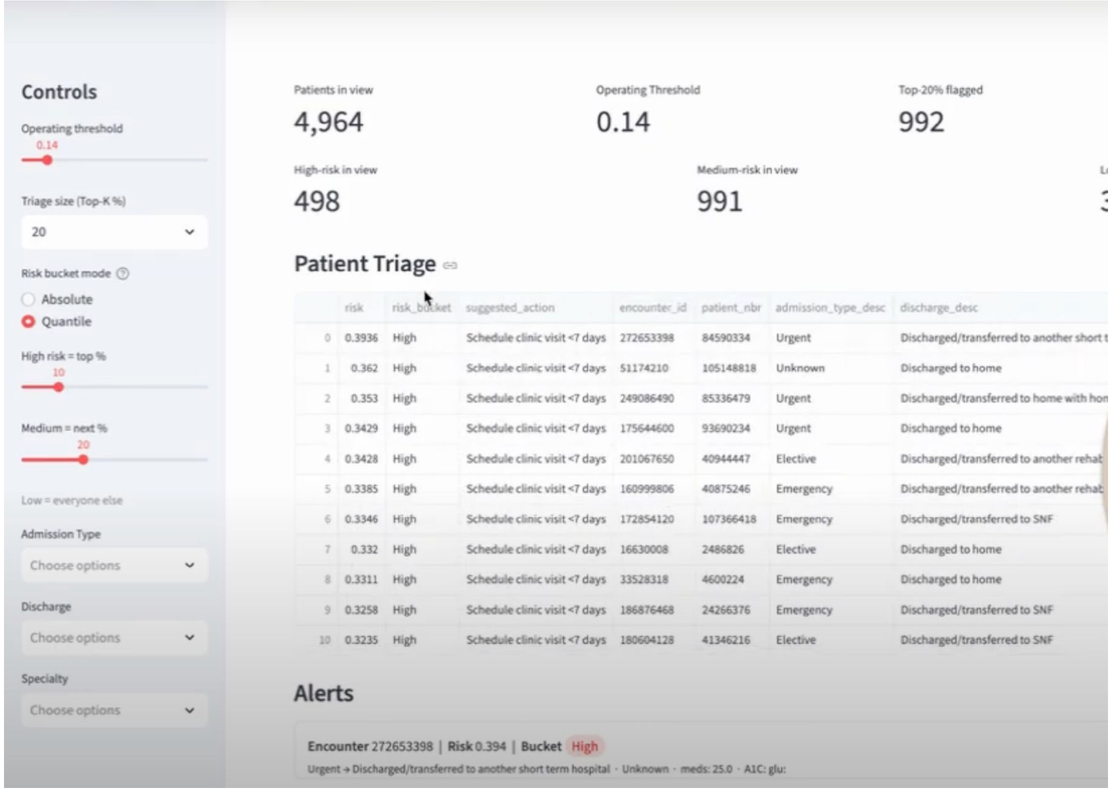

- Patient readmission prediction
- Emergency room volume forecasting based on seasonal patterns, holiday events, partner hospital utilization, and readmission rate prediction
- Workforce requirements forecasting based on volume forecasting

Kumo

Medicare provider fraud detection

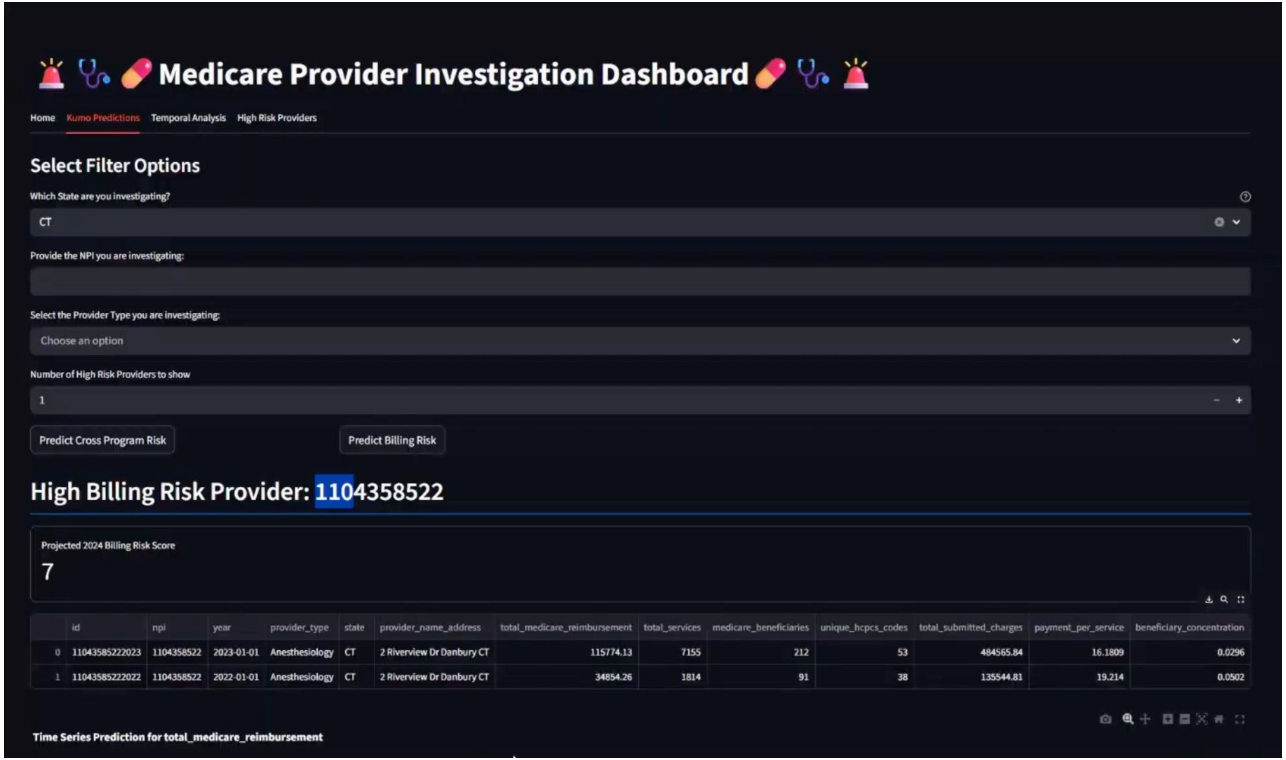

- Billing risk
- Cross Program Risk
- Anomalous or Overprescription Risk

Kumo

Hyper-personalized e-commerce

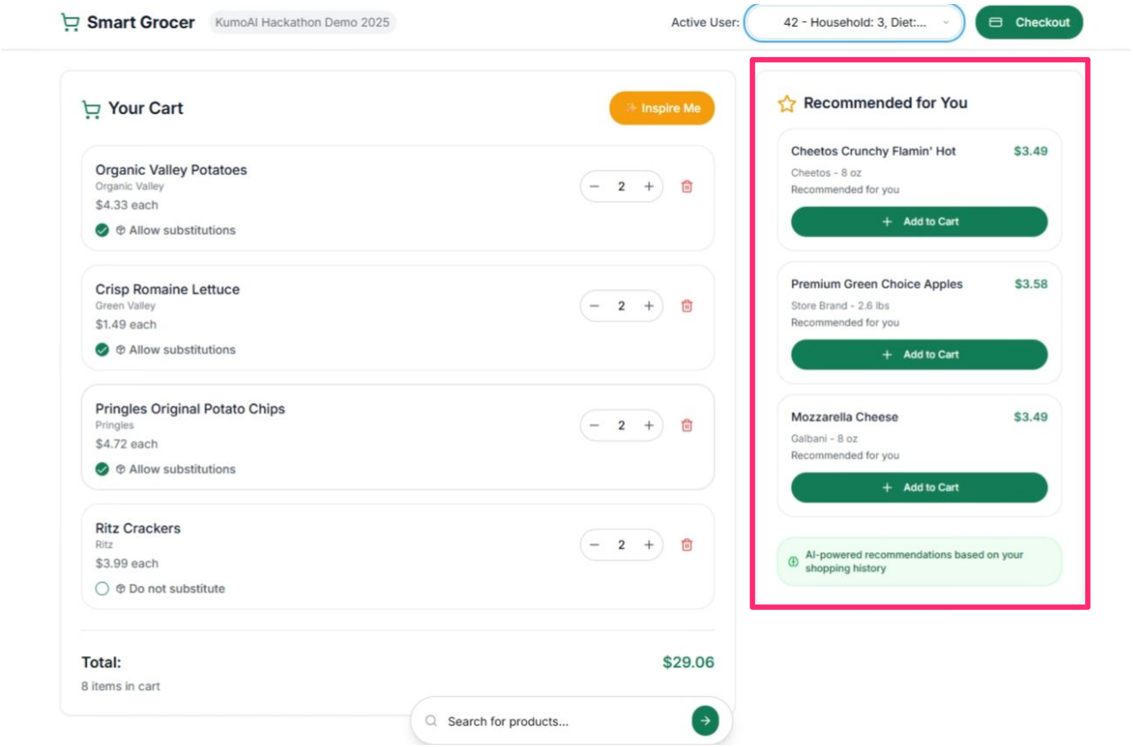

- Shopping agents and recommendations
- Recommendations based on shopping cart additions and item interactions
- Cart abandonment detection and recommendations to re-engage

KumoRFM: Summary

What is it? **A foundation model for business data!**

What can it do? **Make Zero-shot predictions on relational data!**

Can we make it even better? **Fine-tune on a particular task!**

Why does this matter?

- Predictions become commoditized
- Predictive models are democratized
- Prediction-driven applications become possible
- Data becomes even more important
- Fine-tuning empowers high-performance use-cases

Thank You & Next Steps

Share what you've learned on Linkedin/X!
Tag @kumo_ai_team

Get a KumoRFM API key and start building with your data!
kumorfm.ai

Quickstart notebook

https://tinyurl.com/hellokumo

Thinking about using Kumo?

Email hello@kumo.ai to learn more about enterprise options

Foundation Models for Structured Business Data

Jure Leskovec
Co-Founder and Chief Scientist at, Kumo
Professor at Stanford University

Kumo

The Two Halves of the Enterprise Brain

Unstructured Data
(Language & Perception)

- Text, documents, images
- Reasoning, summarizing, and generating content
- Answer the "what" and "why" from human knowledge

Structured, Relation Data
(Business Operations)

- Databases of customers, products, transactions, and supply chains
- This data holds the patterns that predict what will happen. It's the ground truth of the business.

Kumo

The AI revolution is Incomplete

Two halves of the enterprise brain:

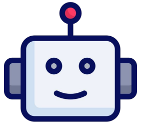

LLMs

LLM - Reasoning

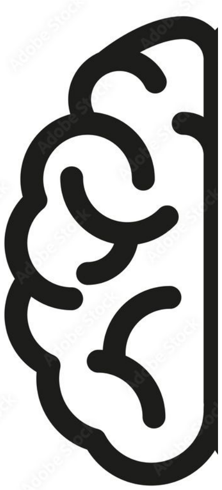

LLMs cannot reason over structured data effectively

Reasoning Brain (LLM):
Understands "*what is in this contract?*"

Predictive/Analytical Brain:
Answers "*which customer will churn next?*"

Kumo

The AI Revolution is Incomplete

For structured data one has to manually build a model for each separate task

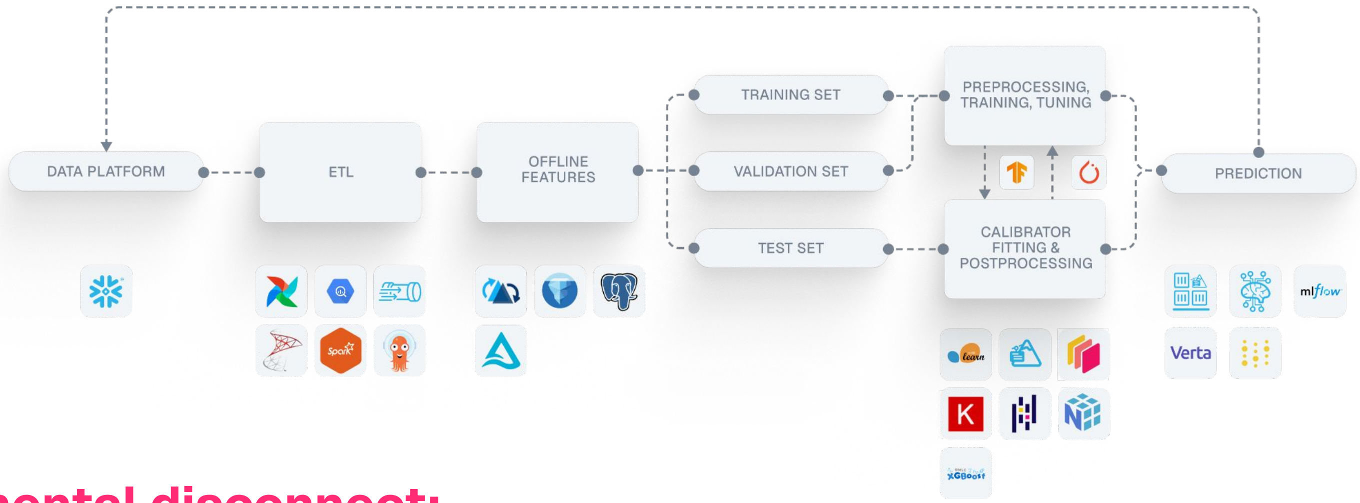

Fundamental disconnect:

- Unstructured data: Ask LLM any question
- Structured data: +6-months to build a model for a **single question**

Kumo

Core Problem & Our Insight

We need both halves of the brain to work together!

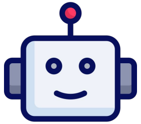

LLMs

LLM - Reasoning

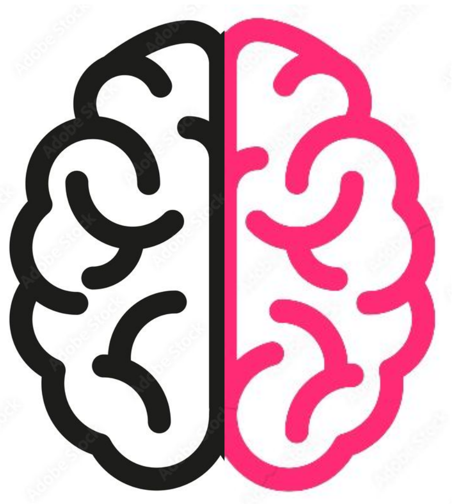

RFM - Prediction

Reasoning Brain (LLM):
Understands "*what is in this contract?*"

Predictive/Analytical Brain (RFM):
Answers "*which customer will churn next?*"

Relational Foundation Model

Kumo has built the world's first Relational Foundation Model (RFM).
Designed specifically business data to make powerful forecasts & predictions.

You can simply point RFM at your database and ask predictive questions:

- *'Which leads are most likely to convert in the next 30 days?'*
- *'What products are at the highest risk of a stockout next quarter?'*
- *'Which transactions show the highest probability of fraud?'*

Need forward-looking forecasts that quantify the impact of actions on your KPIs.

RFM: The missing piece in the AI puzzle

We are building the Predictive Brain for the Enterprise.

We turn operational data into your most powerful predictive asset.

Simply connect your database

One universal model for business data.

Ask predictive questions

Relational Foundation Models (RFMs)

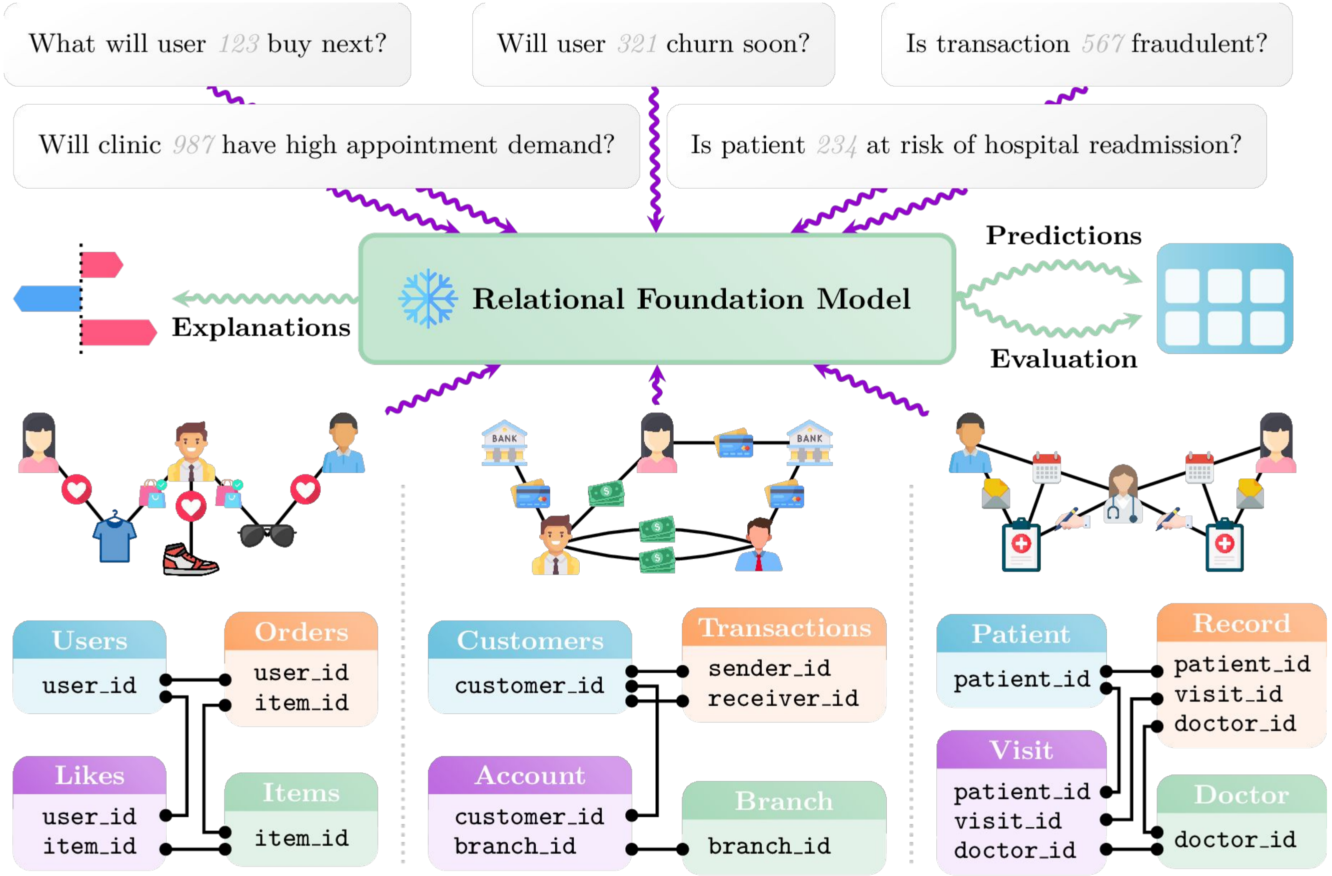

(a) E-Commerce (b) Finance (c) Health-Care

Foundational Model for Relational Data

Large Language Model (LLM):

Doubt thou the stars are fire;
Doubt that the sun doth move;
Doubt truth to be a liar;
But never doubt I love...

Text

Natural Language Prompt
Context: Unstructured Docs

Common-sense in-context reasoning

Text output

Relational Foundation Model:

snowflake
databricks
Google Cloud
amazon S3

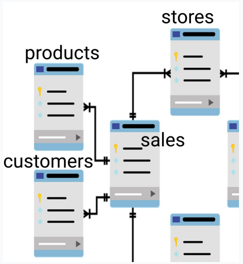

Relational Prompt
Context: Relational Data

Predict churn for active customers:

```
PREDICT COUNT(Sales.*,0,30,days)=0
FOR EACH Customers.ID
WHERE COUNT(Sales.*,-90,0,days)>0
```

Relational reasoning

Predictions
Embeddings

Prompting KumoRFM

We can *talk* to KumoRFM through the **Predictive Query Language Interface**

- It has a label definition (PREDICT clause) and entity definition (FOR clause)
- Additional filters can be applied both to label and entity clause
- Supports aggregations, binary operations and logical operations
- Capable of handling a broad set of task types:

first categorical value

```
PREDICT FIRST(orders.type, 0, 7)
FOR EACH users.user_id IN (0, 1, 2)
```

(a) Node multi-class/label classification

logical value

```
PREDICT COUNT(orders.*, 0, 7) > 0
FOR EACH users.user_id IN (0, 1, 2)
```

(b) Node binary classification

sum of numerical values

```
PREDICT SUM(orders.value, 0, 7)
FOR EACH users.user_id IN (0, 1, 2)
```

(c) Node regression

set of foreign keys

```
PREDICT LIST_DISTINCT(orders.item_id, 0, 7)
FOR EACH users.user_id IN (0, 1, 2)
```

(d) Link prediction

Explainability

- RFM is **fully-differentiable** and enables **gradient-based explanation techniques**
- Importance scores are computed on the **cell level** rather than on the feature level
- Conversion to **textual summary**

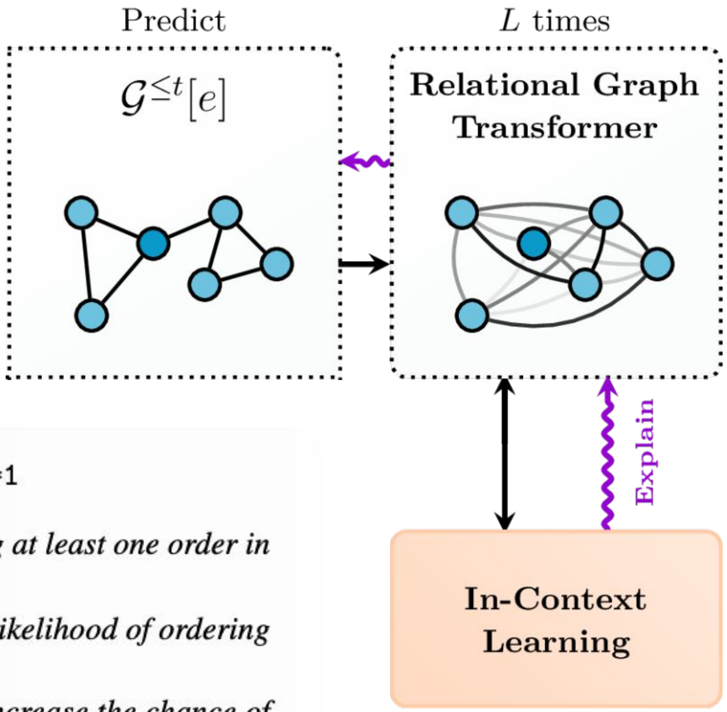


```
PREDICT COUNT(orders.*, 0, 30) > 0 FOR users.user_id=1
```

The model predicts that the user has a moderate likelihood of placing at least one order in the upcoming month. Key factors influencing this prediction include:

- ***Order Count:*** *Users with only a few past orders have a very low likelihood of ordering soon, while those with more orders show increased probabilities.*
- ***Order Date Recency:*** *Recent orders (6-12 months ago) greatly increase the chance of placing new orders soon.*
- ***Fashion News Frequency and Club Membership:*** *Users who regularly receive fashion news or have active club membership status show higher probabilities of ordering.*

Kumo

Benefits of RFMs

Better performance:
+10% accuracy improvement through relational context

Faster time to value:
95% reduction in data preparation effort, handles any taks on the fly

Scalability:
Handles complex, relational data across billions of records

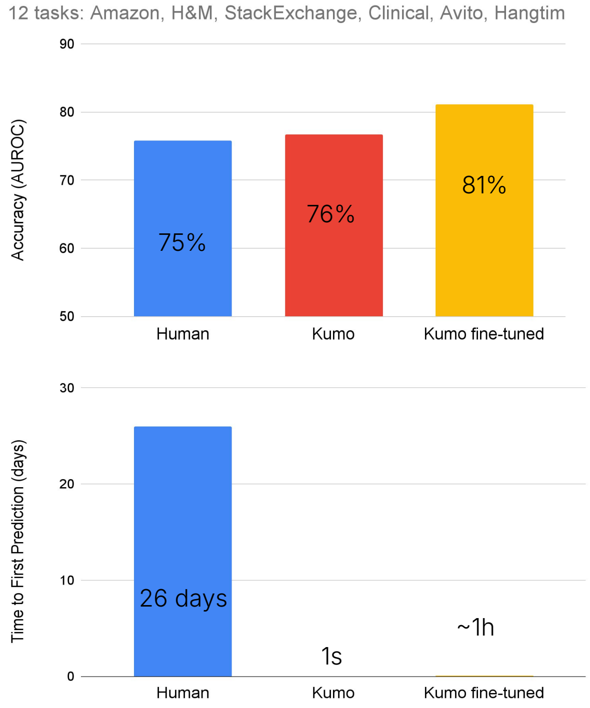

Stanford RelBench

RELBENCH
RELATIONAL DEEP LEARNING BENCHMARK

LLMs and RFMs

RFMs are filling the hole left by LLMs:

- Not predicting text from text but outcomes from business data
- Just like LLMs replaced time-consuming, difficult, and expensive NLP work, RFMs replace time-consuming, difficult, and expensive predictive modeling work
- Since RFMs fills a hole, it's not replacing or competing with LLMs.
- RFMs are complementary to LLMs and most businesses will end up using both

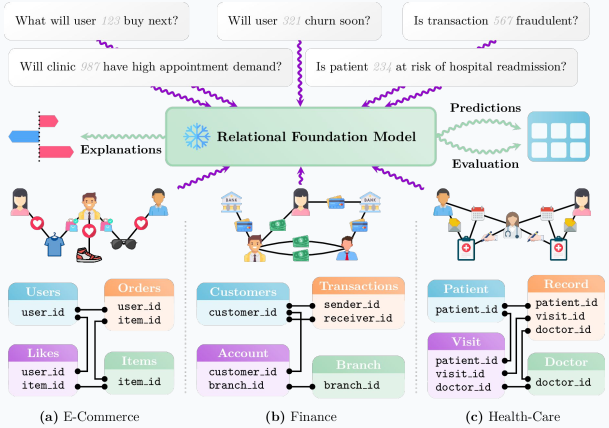


```
from kumoai.rfm import Graph, KumoRFM

df_dict = {
  'users': pd.DataFrame(...),
  'orders': pd.DataFrame(...),
  'items': pd.DataFrame(...),
}

graph = Graph.from_data(df_dict)
rfm = KumoRFM(graph)

query = "PREDICT COUNT(orders.*, 0, 30)>0 FOR users.user_id=1"
rfm.query(query)
```

Bulding with Kumo

What types of predictive tasks and tools may benefit Claudius, a shop management agent?

Demand Forecasting: Predic … at times (weekly,monthly) to reduces o…

Customer Purchase Predicti… …ers are likely to buy next and tailor promotions or r…

Product-Level Margin Foreca… …em category using historic sales, costs, discounts…

Dynamic Pricing: Adjust prici… …s.

Campaign Effectiveness: Pr… …onversions.

Supplier Lead Times: Anticipate delivery/restocking delays (OTIF - On Time in Full predictions) and adjust orders proactively.

Scenario Simulation: Predict the impact on margins of different strategies, such as adjusting prices, introducing a loyalty discount, or swapping suppliers.

Make a case for agents and applications needing forward looking predictive/scoring ability.

What types of predictive tasks and tools may benefit Claudius, a shop management agent?

Demand Forecasting: Predict which products will sell, in what quantity, and at what times (weekly,monthly) to reduces overstock and stockouts

Customer Purchase Predictions & Personalized Offers: Anticipate what customers are likely to buy next and tailor promotions or recommendations.

Product-Level Margin Forecasting: Forecast profit margin by product/SKU and item category using historic sales, costs, discounts, and seasonality to optimize profit margin.

Dynamic Pricing: Adjust pricing based on demand, competition, or inventory levels.

Campaign Effectiveness: Predict which marketing strategies will lead to higher conversions.

Supplier Lead Times: Anticipate delivery/restocking delays (OTIF - On Time in Full predictions) and adjust orders proactively.

Scenario Simulation: Predict the impact on margins of different strategies, such as adjusting prices, introducing a loyalty discount, or swapping suppliers.

Agents need predictive AI tools

Majority of First-Generation Agents, Including Claudius

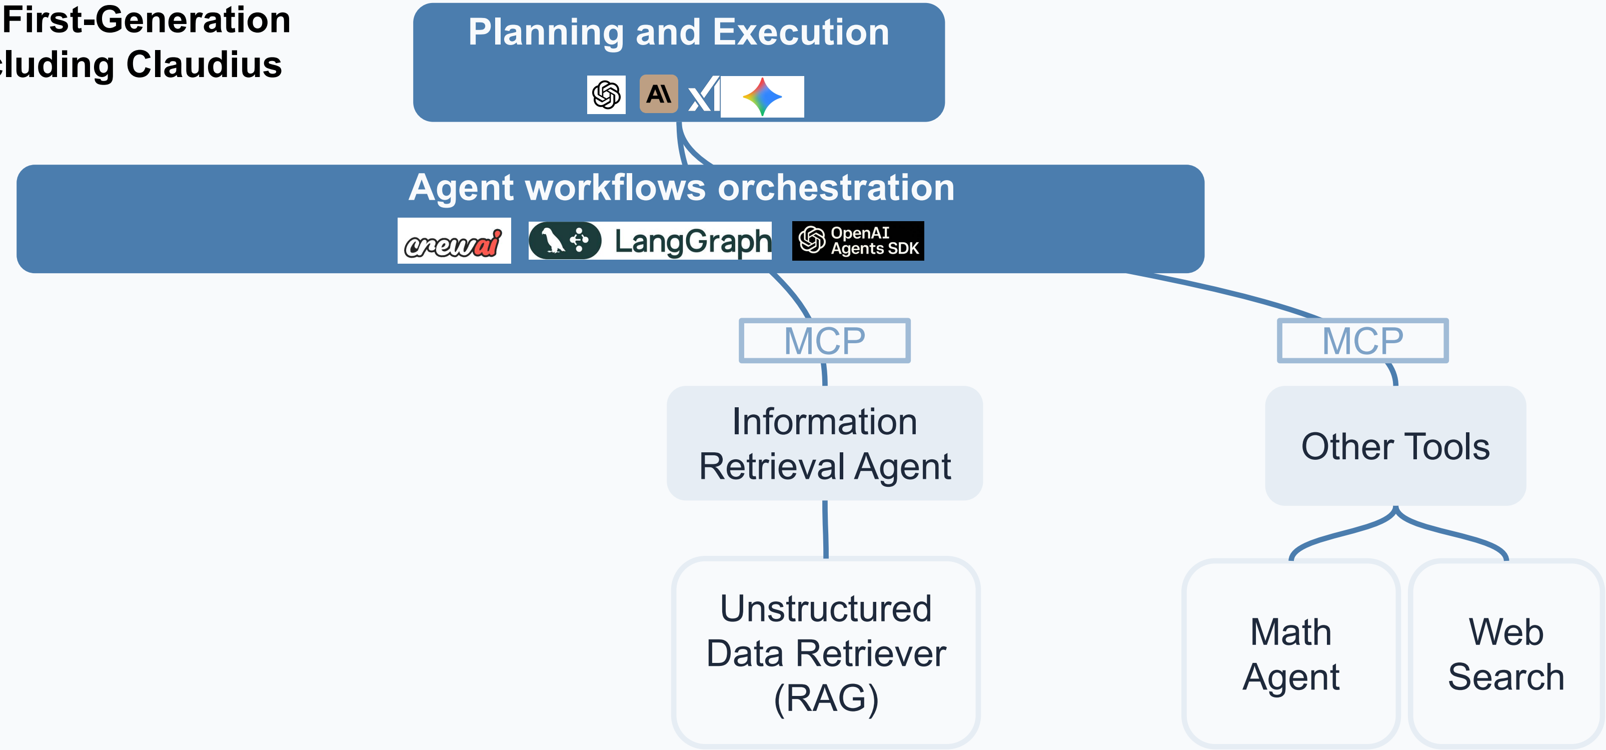

Agents need predictive AI tools

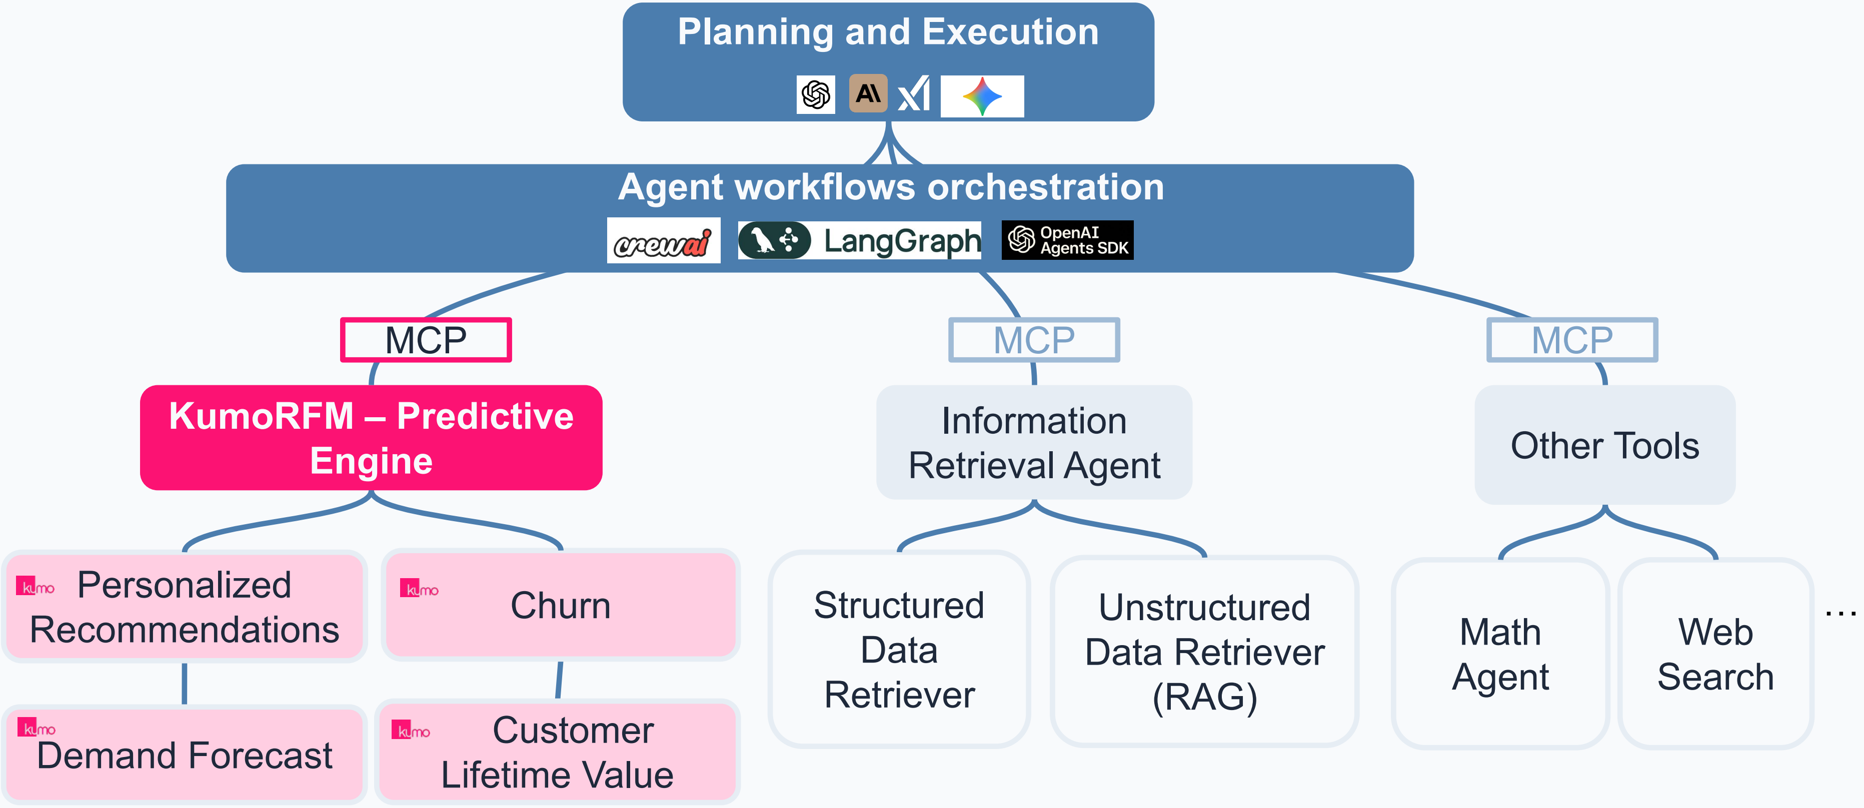

.. And many other Predictive AI tasks

Without the correct tools, LLMs hallucinate over structured data

To make accurate predictions, you need **data** + **algorithms** to extract patterns from the data

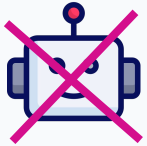

LLMs cannot reason over this data effectively

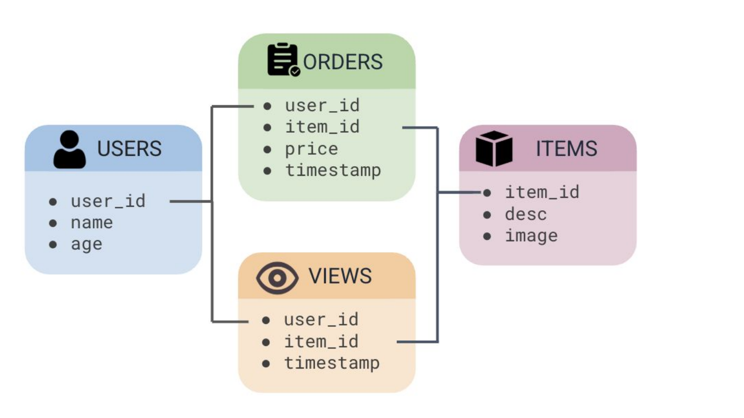

For instance, to predict User (001)'s item preference, you'll need information about

- User 001
- Other users
- Other information across the database (e.g. items, views, orders)

And an algorithm to optimize for the prediction

LLMs are not built for this purpose

Model Context Protocol (MCP) acts as a bridge between agent and service / tool

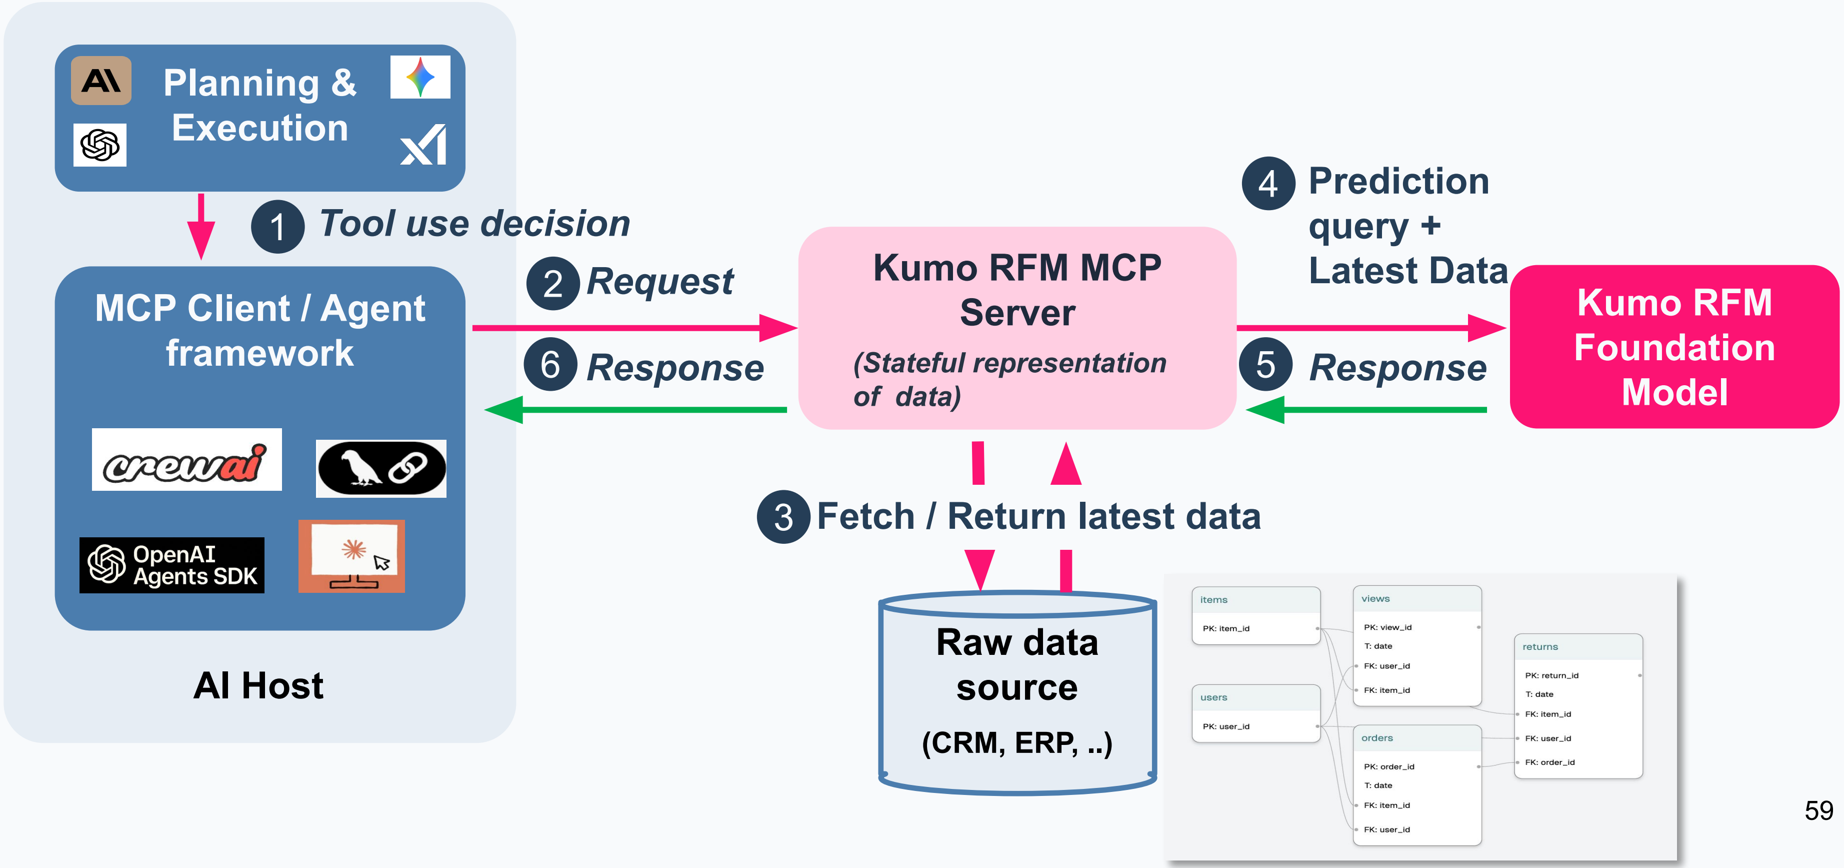

The KumoRFM MCP server exposes KumoRFM's capabilities through carefully designed tools

🤔Knowledge tools

`get_docs`: Retrieves documentation for graph setup and Predictive Query Language (PQL) syntax

📊 I/O Operations

`find_table_files`: Discovers CSV/Parquet files in local directories or S3

`inspect_table_files`: Examines schemas and sample data to understand structure

🤖 Prediction Execution

`predict`: Executes Kumo predictive queries to generate predictions

`evaluate`: Assesses prediction quality on holdout data / historical ground truth

🕸 Graph Management

`inspect_graph_metadata`: Views current graph structure, relationships, and semantic types

`update_graph_metadata`: Configures table relationships, primary keys, and time columns

`get_mermaid`: Generates visual entity-relationship diagrams

`lookup_table_rows`: Retrieves specific rows from tables

`materialize_graph`: Builds the graph structure for predictions

`Lookup_table_rows` : Lookup rows in the raw data frame of a table for a list of primary keys

Extended Warranty Renewals Agent

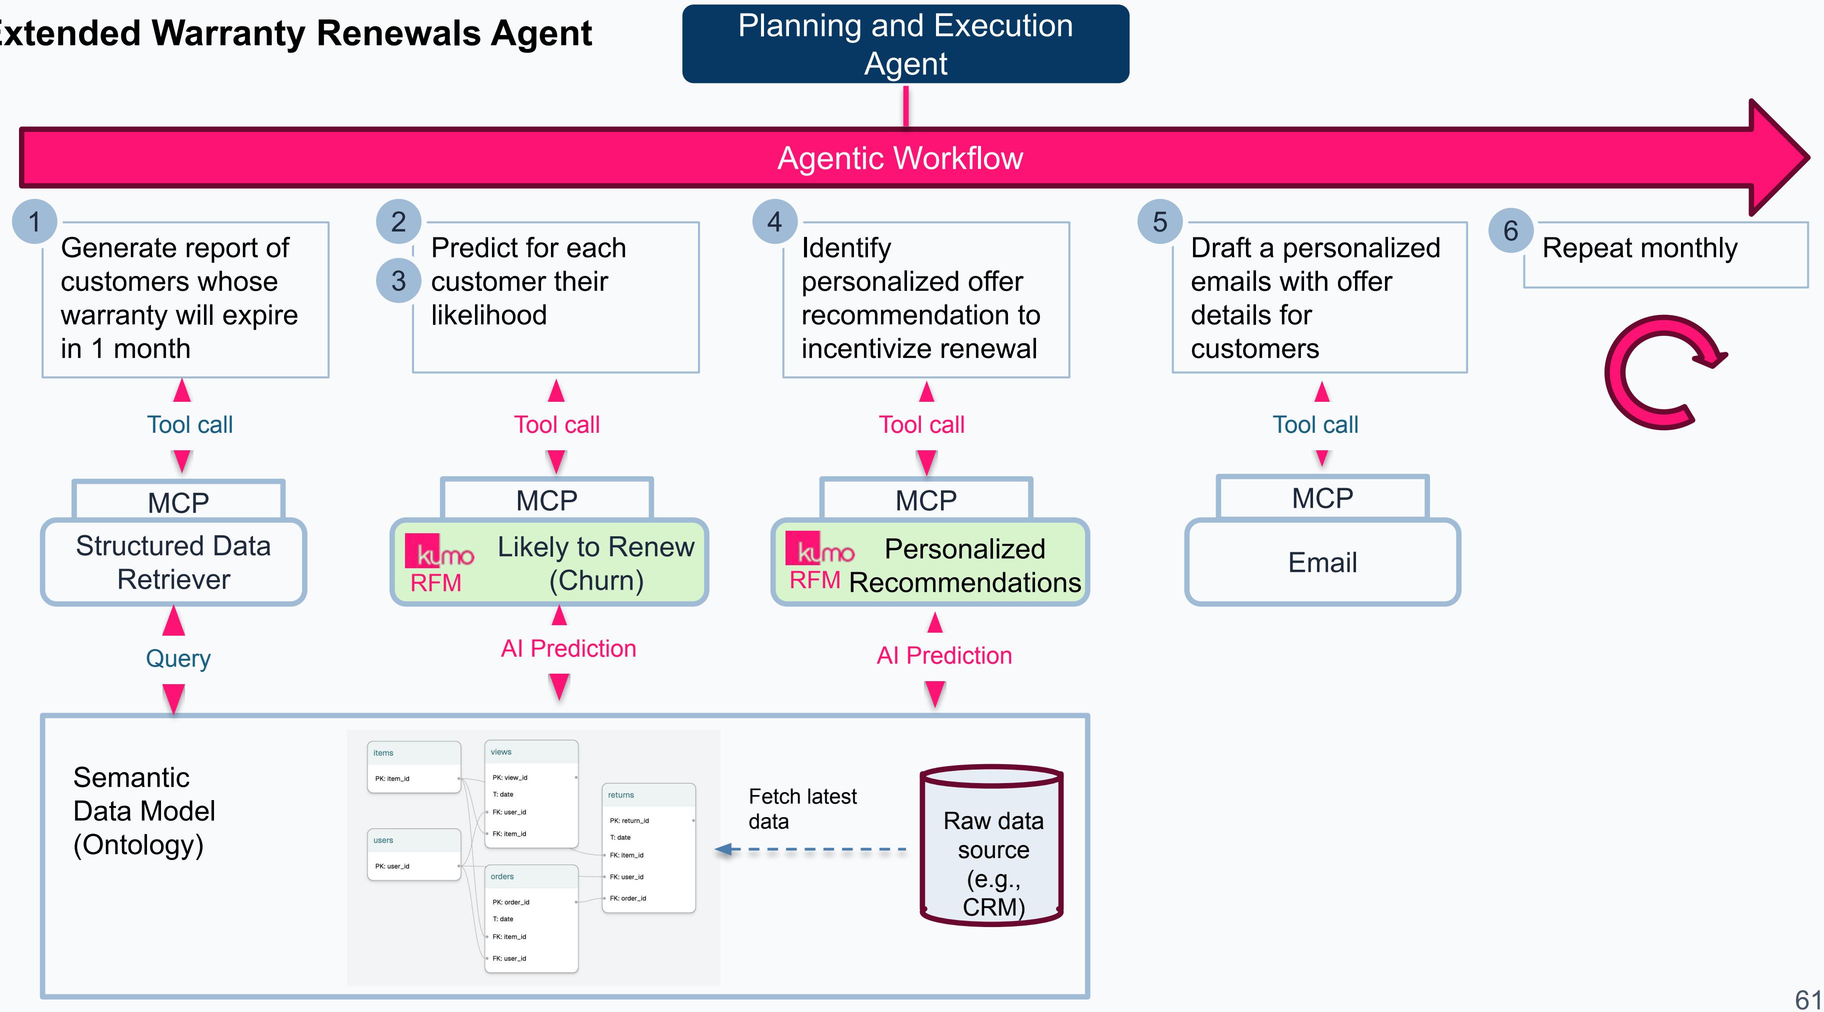

What kinds of projects are people building with KumoRFM?

Health coverage matching

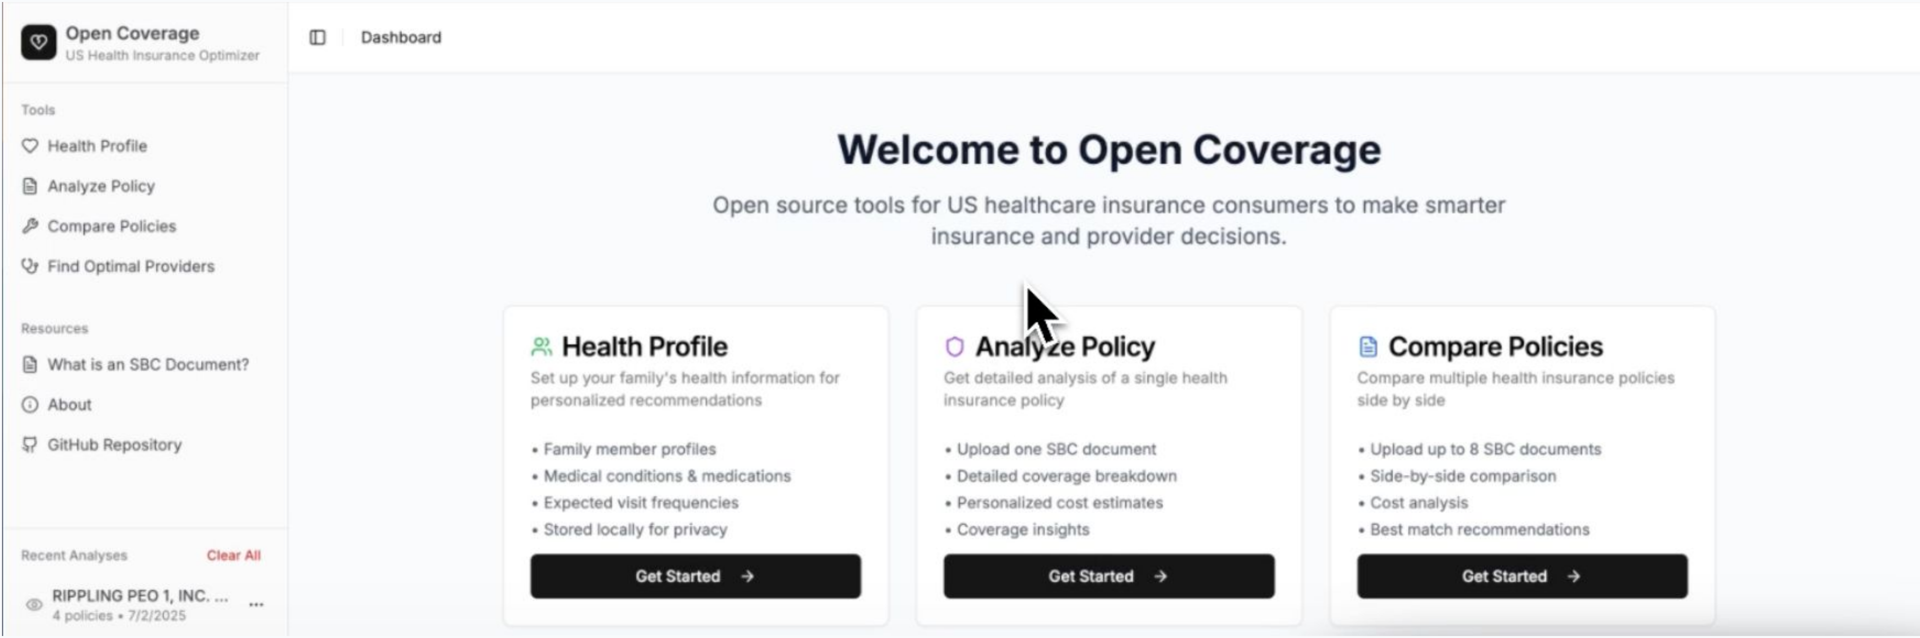

An application that empowers U.S. consumers to make **smarter insurance decisions** and reduce out-of-pocket spending through **personalized healthcare utilization forecasting** and **preference matching**.

Energy grid optimization

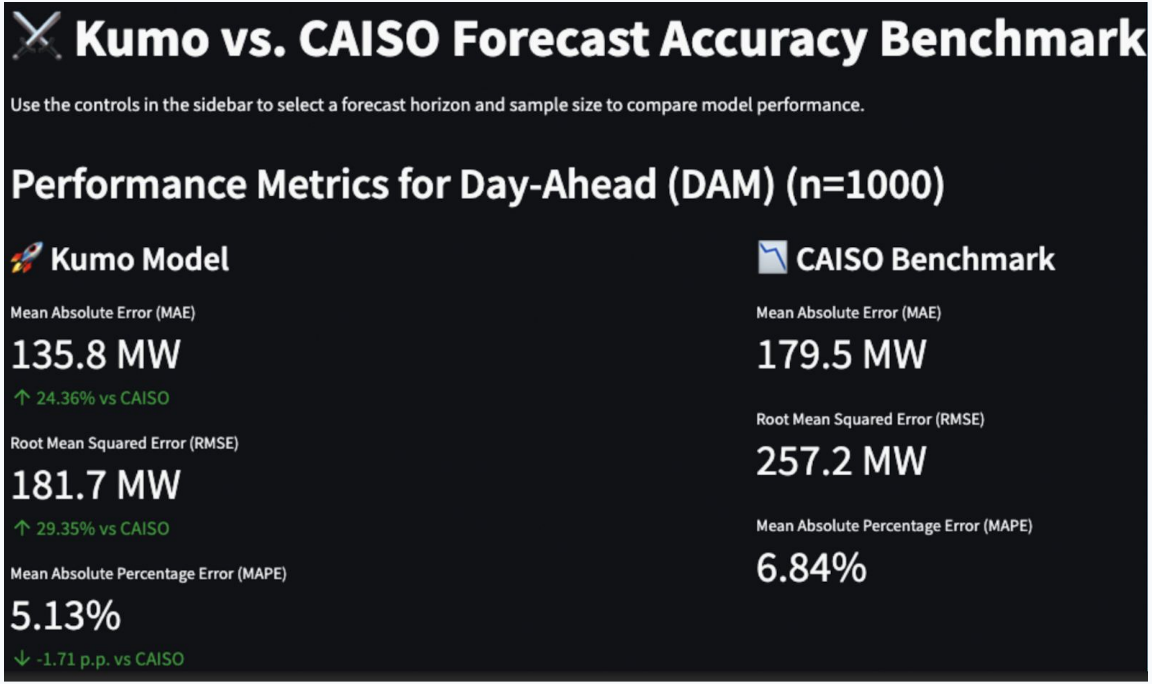

A **localized energy prediction platform** that outperforms institutional forecasts by 30%, with applicability to **reduce electricity over purchasing from wholesalers** by Load Serving Entities. KumoRFM is used to generate local electricity forecasts using data including weather, distributed solar timing, local electric vehicle data, and community demand cycles.

Wildfire risk assessment predictions

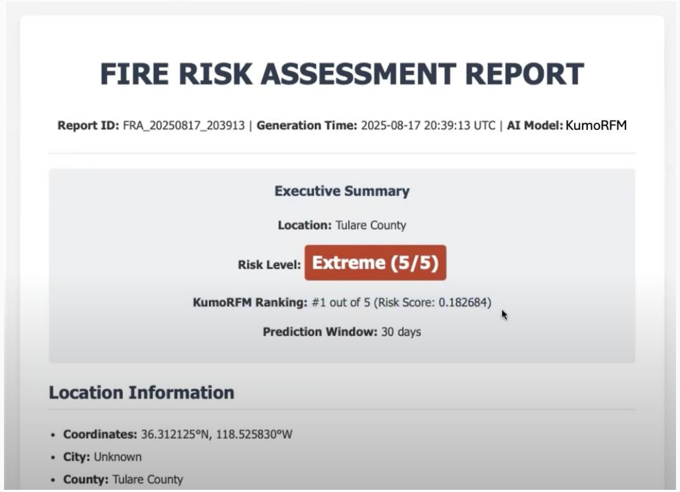

FIRE RISK ASSESSMENT REPORT

Report ID: FRA_20250817_203913 | **Generation Time:** 2025-08-17 20:39:13 UTC | **AI Model:** KumoRFM

Executive Summary

Location: Tulare County

Risk Level: Extreme (5/5)

KumoRFM Ranking: #1 out of 5 (Risk Score: 0.182684)

Prediction Window: 30 days

Location Information

- **Coordinates:** 36.312125°N, 118.525830°W
- **City:** Unknown
- **County:** Tulare County

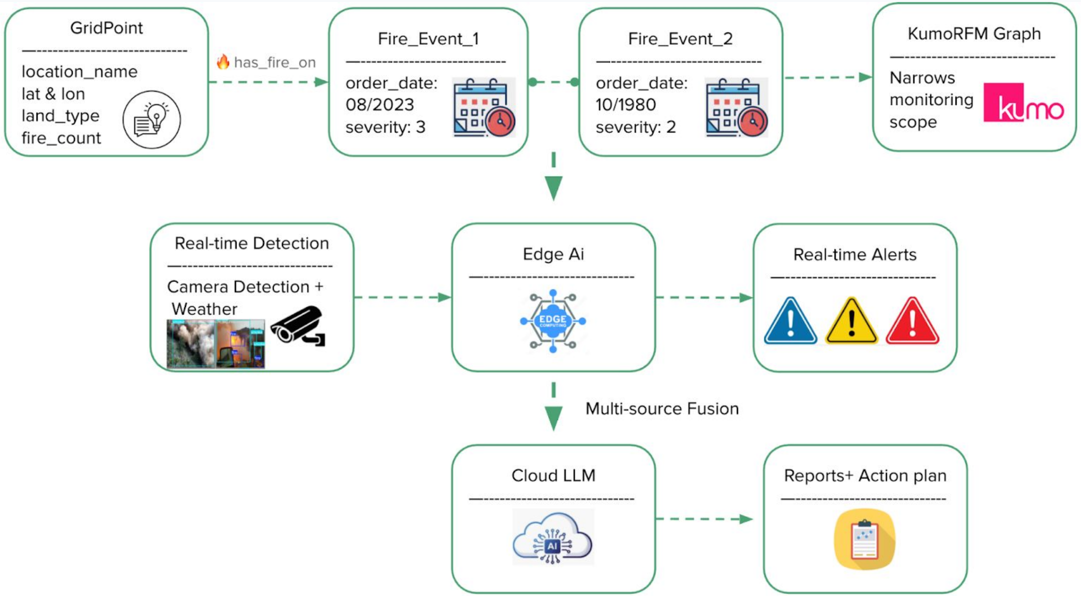

This three-layer AI architecture combines predictive modeling, real-time detection, and intelligent decision-making for comprehensive wildfire risk management.

A **wildfire risk prediction system**, combining KumoRFM with real-time edge sensor data and large language models. KumoRFM serves as the prediction engine to forecast wildfire risk at a granular, geo-grid-based level, aiding wildfire proactive measure prioritization and response planning.

Patient readmission and hospital utilization predictions & capacity planning

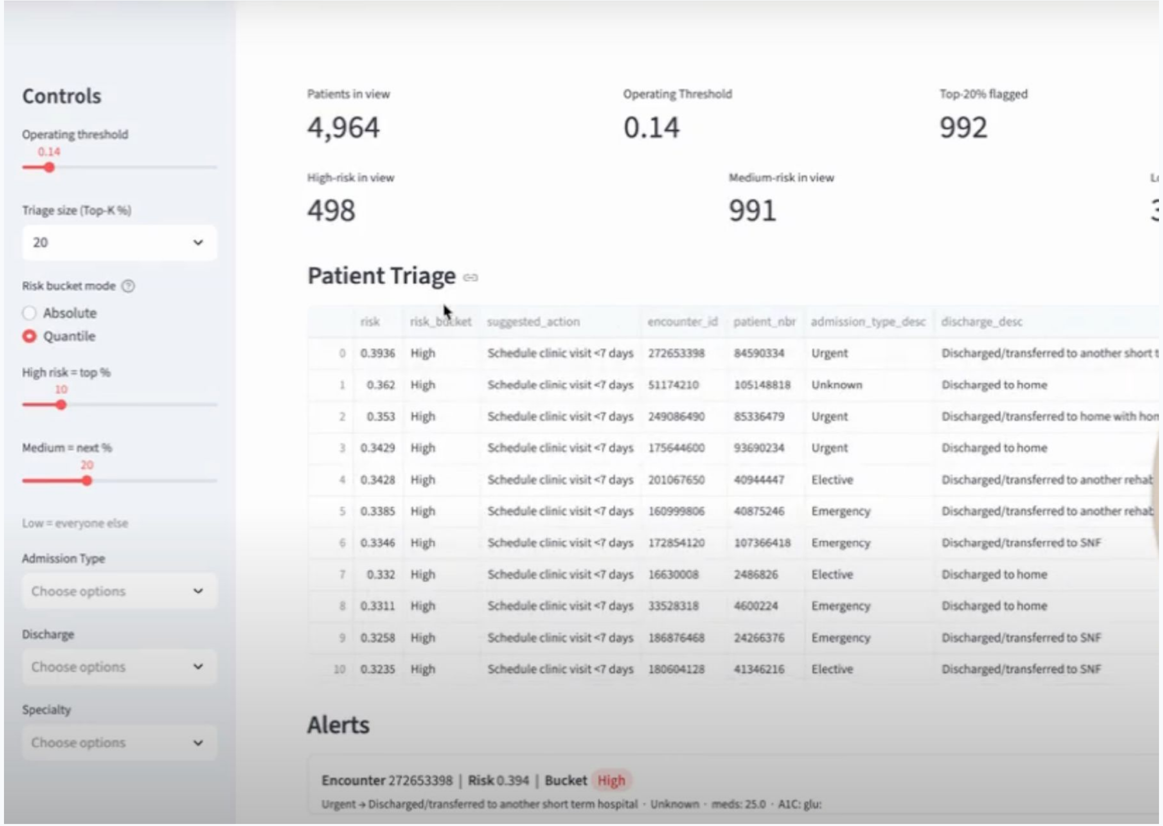

	risk	risk_bucket	suggested_action	encounter_id	patient_nbr	admission_type_desc	discharge_desc
0	0.3936	High	Schedule clinic visit <7 days	272653398	84590334	Urgent	Discharged/transferred to another short t
1	0.362	High	Schedule clinic visit <7 days	51174210	105148818	Unknown	Discharged to home
2	0.353	High	Schedule clinic visit <7 days	249086490	85336479	Urgent	Discharged/transferred to home with hon
3	0.3429	High	Schedule clinic visit <7 days	175644600	93690234	Urgent	Discharged to home
4	0.3428	High	Schedule clinic visit <7 days	201067650	40944447	Elective	Discharged/transferred to another rehab
5	0.3385	High	Schedule clinic visit <7 days	160999806	40875246	Emergency	Discharged/transferred to another rehab
6	0.3346	High	Schedule clinic visit <7 days	172854120	107366418	Emergency	Discharged/transferred to SNF
7	0.332	High	Schedule clinic visit <7 days	16630008	2486826	Elective	Discharged to home
8	0.3311	High	Schedule clinic visit <7 days	33528318	4600224	Emergency	Discharged to home
9	0.3258	High	Schedule clinic visit <7 days	186876468	24266376	Emergency	Discharged/transferred to SNF
10	0.3235	High	Schedule clinic visit <7 days	180604128	41346216	Elective	Discharged/transferred to SNF

Medicare provider fraud detection

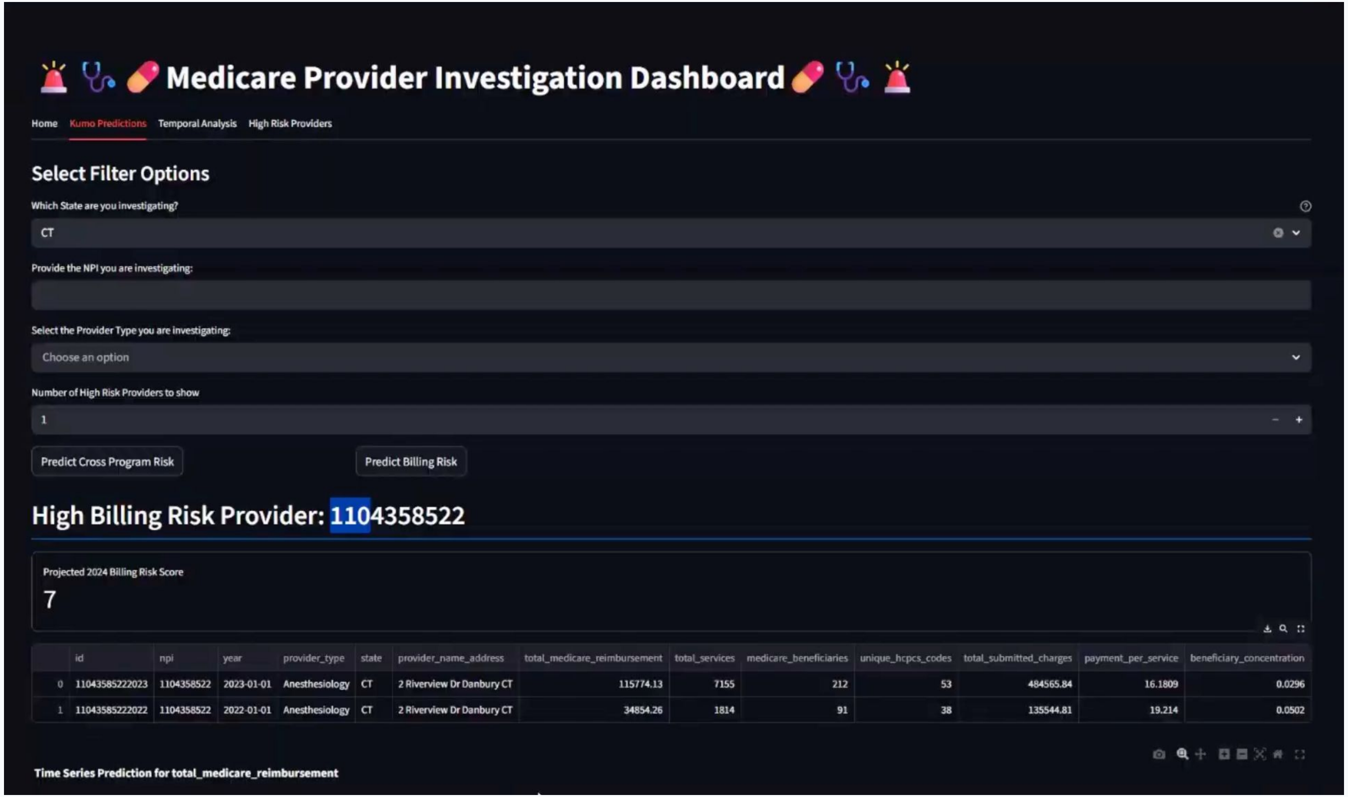

	id	npi	year	provider_type	state	provider_name_address	total_medicare_reimbursement	total_services	medicare_beneficiaries	unique_hcpcs_codes	total_submitted_charges	payment_per_service	beneficiary_concentration
0	11043585222023	1104358522	2023-01-01	Anesthesiology	CT	2 Riverview Dr Danbury CT	115774.13	7155	212	53	484565.84	16.1809	0.0296
1	11043585222022	1104358522	2022-01-01	Anesthesiology	CT	2 Riverview Dr Danbury CT	34854.26	1814	91	38	135544.81	19.214	0.0502

Hyper-personalized e-commerce and shopping agents and recommendations

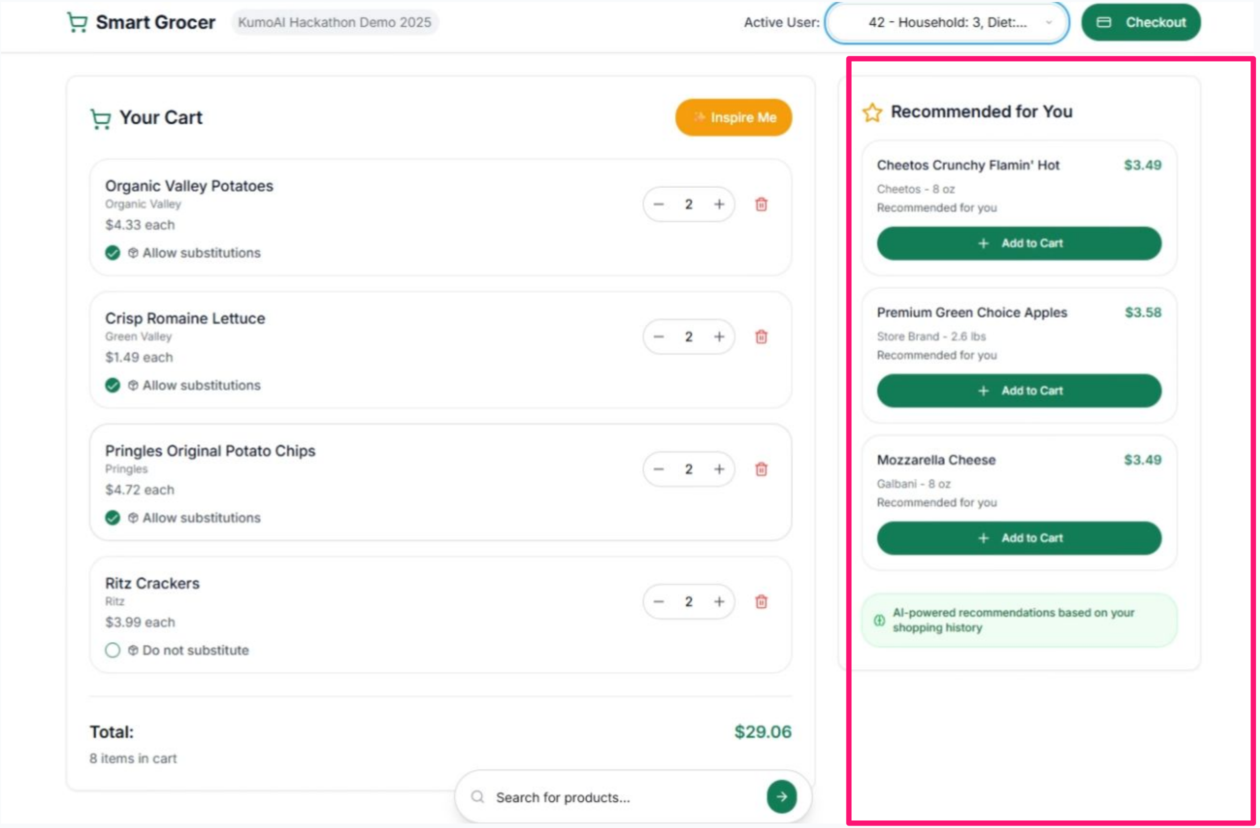

KumoRFM: Summary

What is it? **A foundation model for business data!**

What can it do? **Make Zero-shot predictions on relational data!**

Can we make it even better? **Fine-tune on a particular task!**

Why does this matter?

- Predictions become commoditized
- Predictive models are democratized
- Prediction-driven applications become possible
- Data becomes even more important
- Fine-tuning empowers high performance use-cases

Thank You & Next Steps

Share what you've learned on Linkedin/X!
Tag @kumo_ai_team

Get a KumoRFM API key and start building with your data!
kumorfm.ai

Quickstart notebook

https://tinyurl.com/hellokumo

Thinking about using Kumo?

Email hello@kumo.ai to learn more about enterprise options

The END

Why is building a RFM hard?

Relational Foundation Models present unique and substantial challenges:

Learn across and adapt to diverse database schemas:
- Arbitrary number of tables and columns
- Different types of relationships (*e.g.*, one-to-many, many-to-many)

Heterogeneity of column types:
- Numericals, categoricals, free text, *etc*
- Divergent semantic meaning of columns
- Proprietary or opaque information (*e.g.*, custom upstream embeddings, hashed identifiers)

Complex Task Type Definition:
- Going beyond missing cell imputation
- Temporal forecasting, *e.g.*, predicting inventory demand by utilizing past sales, supplier reliability, seasonal trends, and macroeconomics effects

- How can one interact with such a model?
- How does the neural network look like?
- How should such a model be trained?
- How can one efficiently apply it in real-time?

Supervised *vs.* In-Context Learning

Traditional modeling

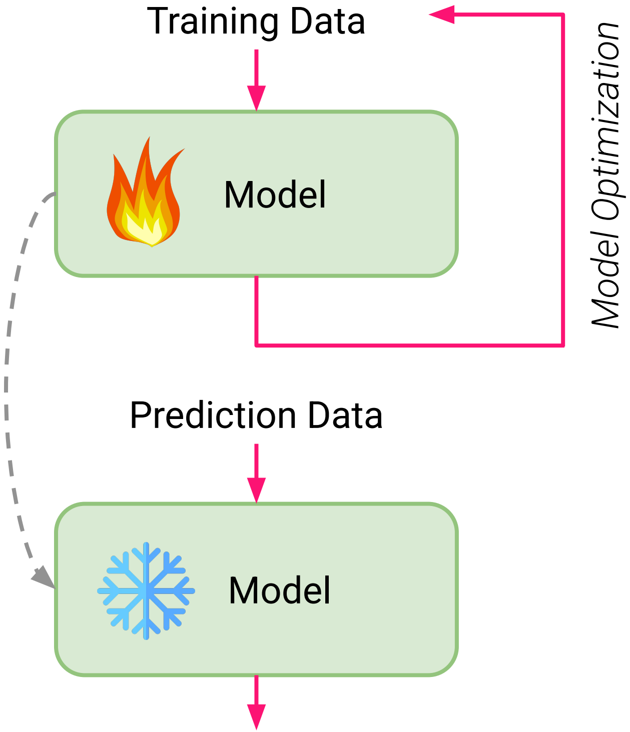

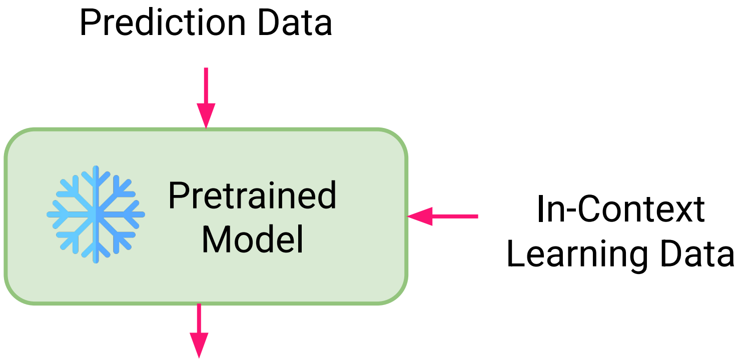

Pre-trained the model performs "optimization" within a **single forward pass**

How it works? Relational Deep Learning

https://arxiv.org/abs/2312.04615

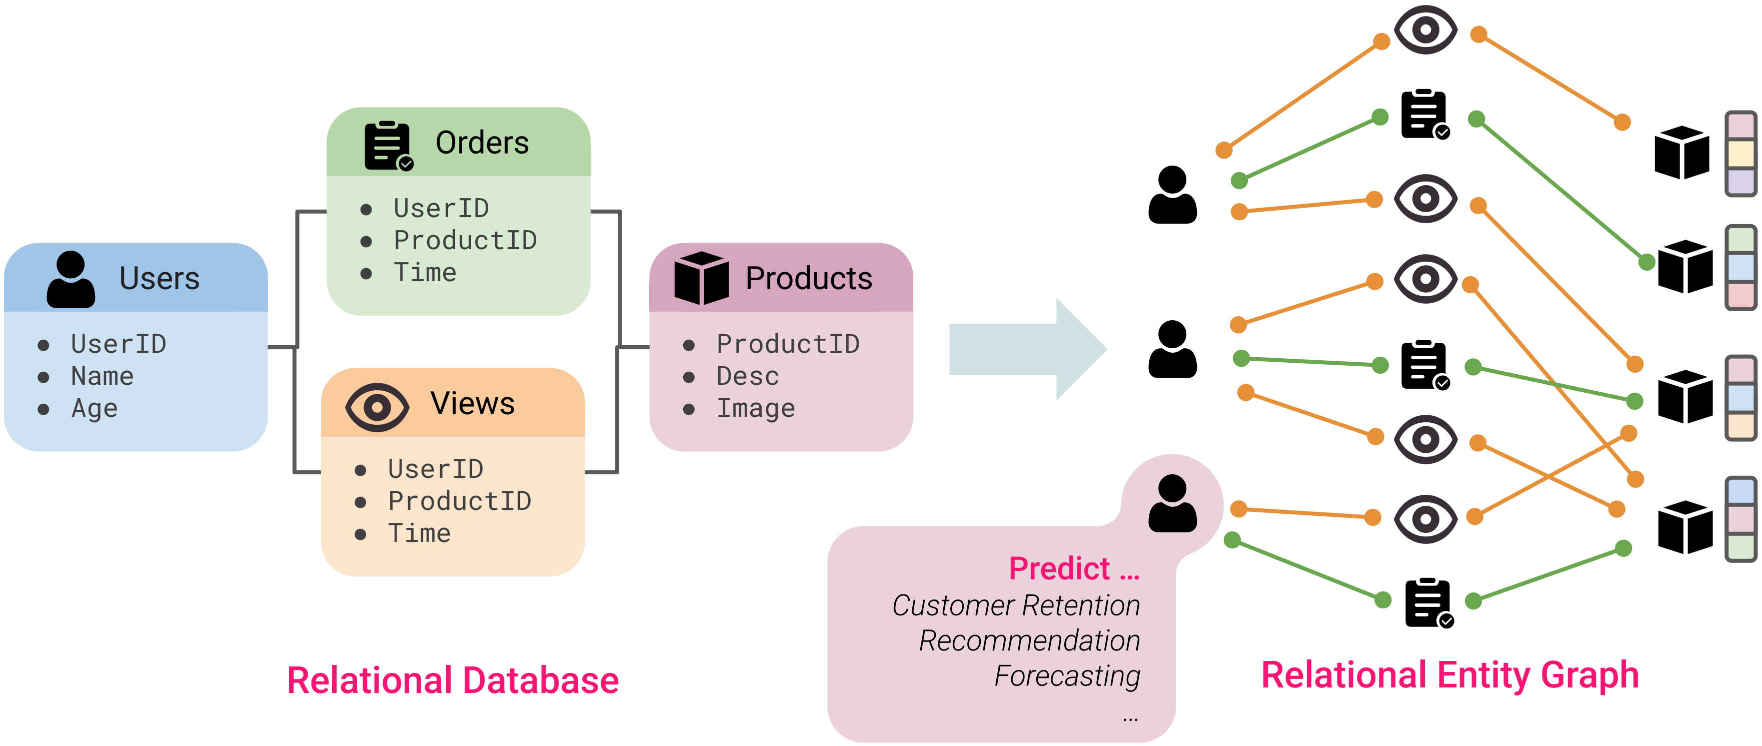

with a **Pretrained Relational Graph Transformer**

How it works? Relational Deep Learning

https://arxiv.org/abs/2312.04615

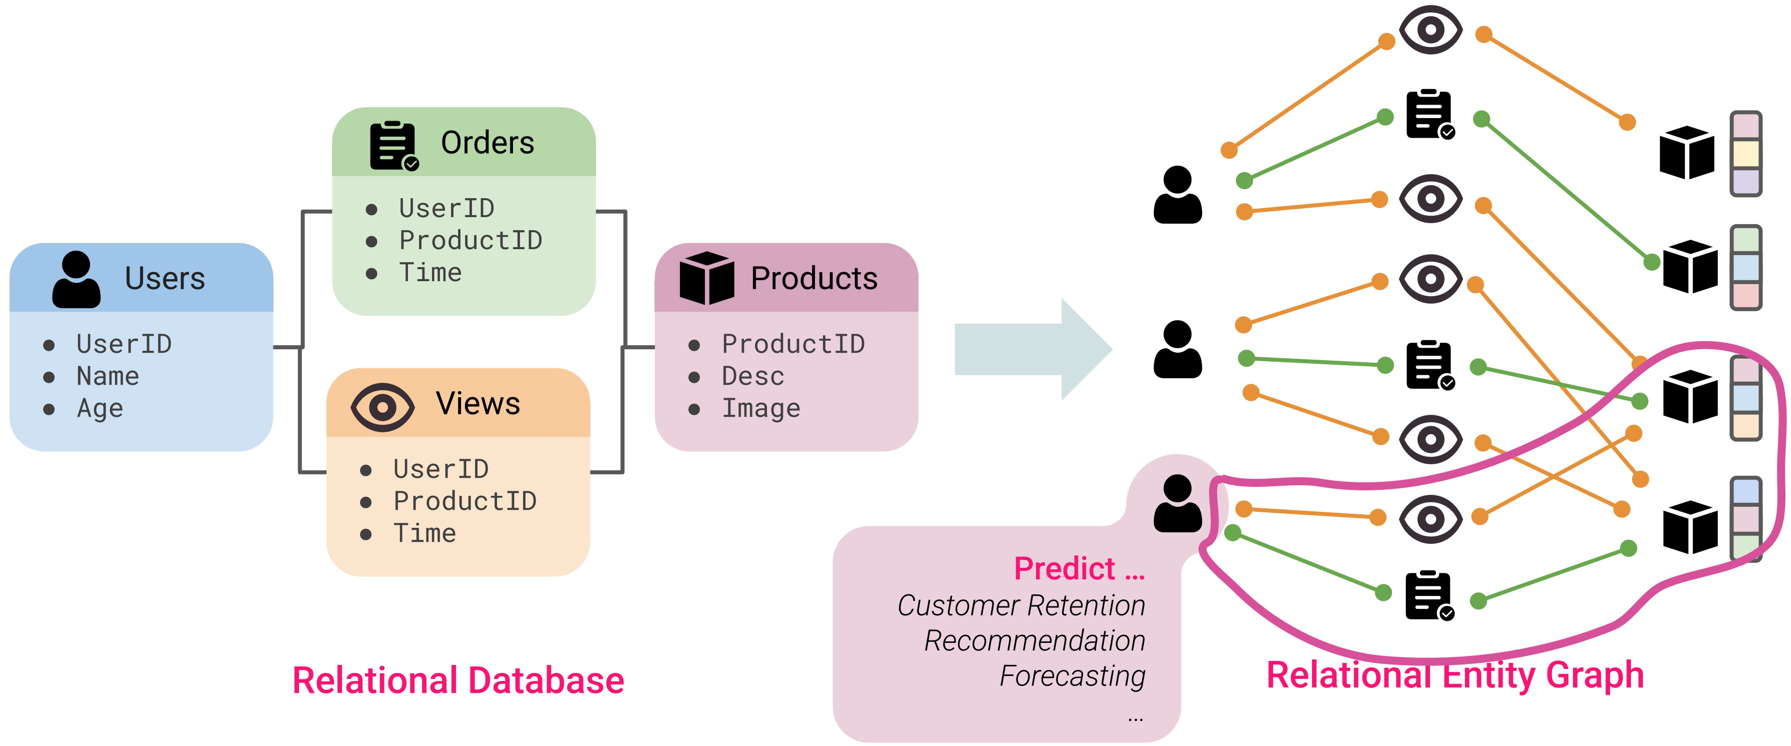

with a **Pretrained Relational Graph Transformer**

How it works? Relational Deep Learning

https://arxiv.org/abs/2312.04615

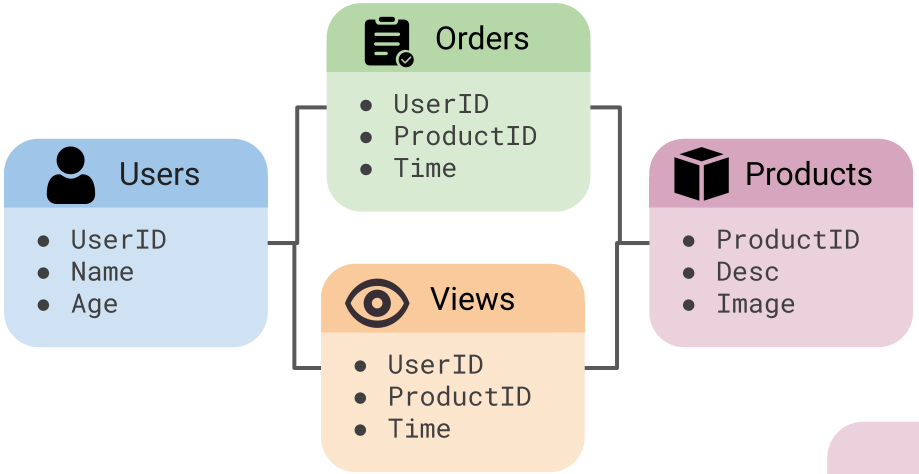

Relational Database

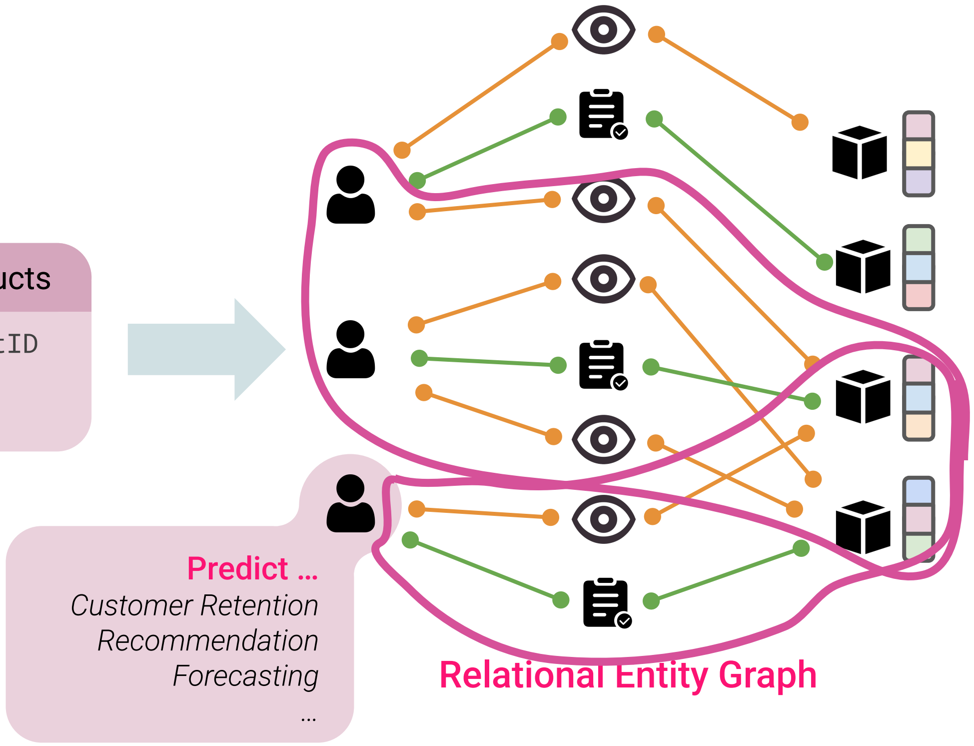

Relational Entity Graph

with a **Pretrained Relational Graph Transformer**

Architecture: Relational Graph Transformer

Tokenizes the subgraph, attaches graph-specific **positional encodings**, and applies **multi-head attention**

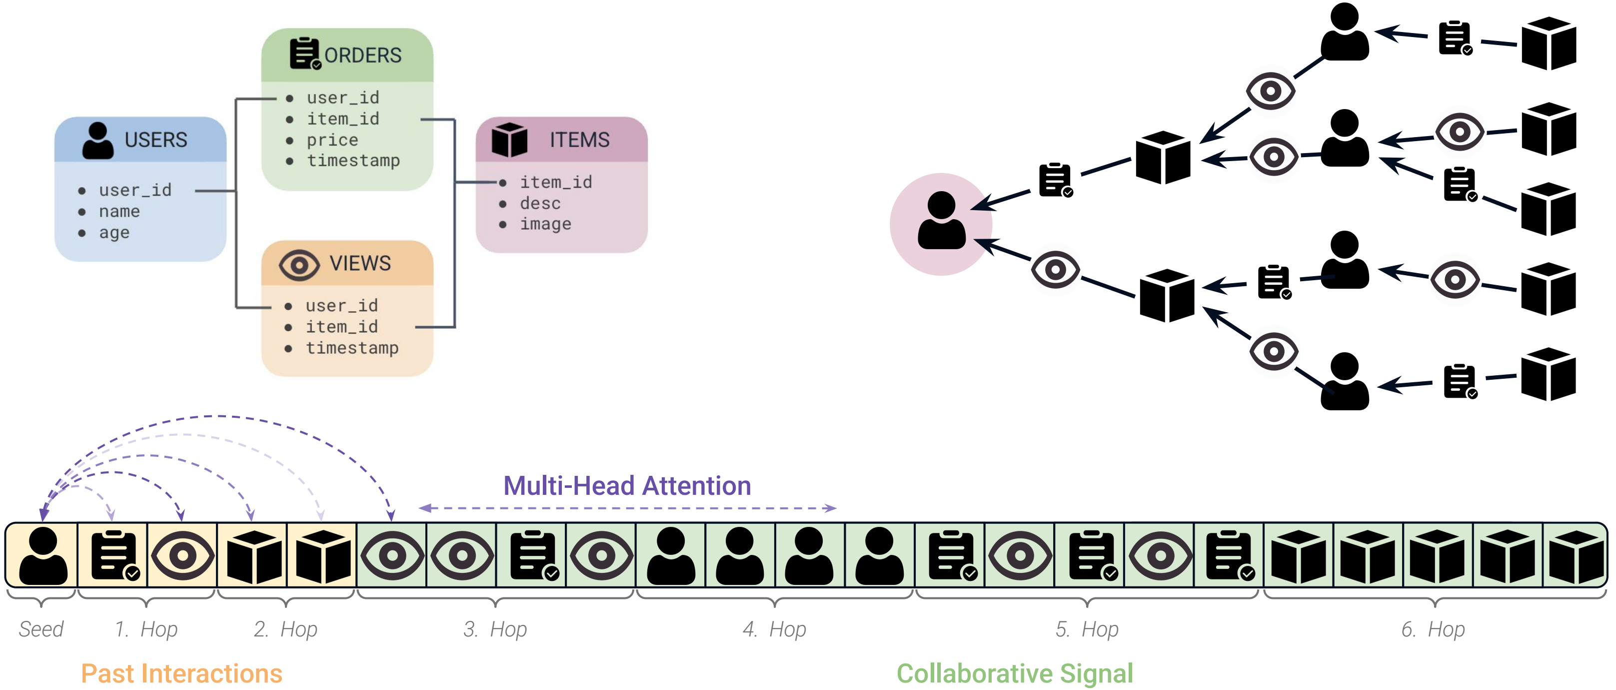

https://kumo.ai/research/relational-graph-transformers

Relational Graph Transformer

Graph Transformers can **attend** across
multiple **columns**, *multiple* **tables** and *multiple* **hops**

They can learn ...

1 Filters

- Last fact
- Last *k* facts
- Facts in last week
- Upcoming holiday
- ...

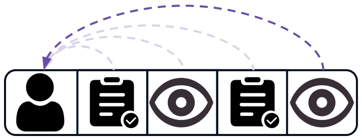

2 Correlations

- Bought together
- Repeated/regular patterns
- Trends
- Time between
- ...

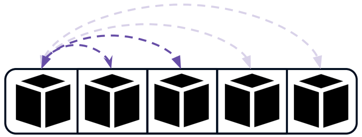

3 Aggregations

- Weighted average
- Summation *(degree count)*
- Standard deviation
- ...

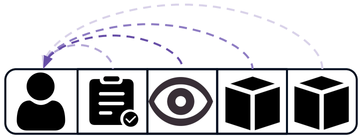

Relational Foundation Model

Predictive Query → “Prompt”

```
PREDICT COUNT(TRANSACTIONS,0,30,days)=0
FOR EACH CUSTOMERS.CUSTOMER_ID
```

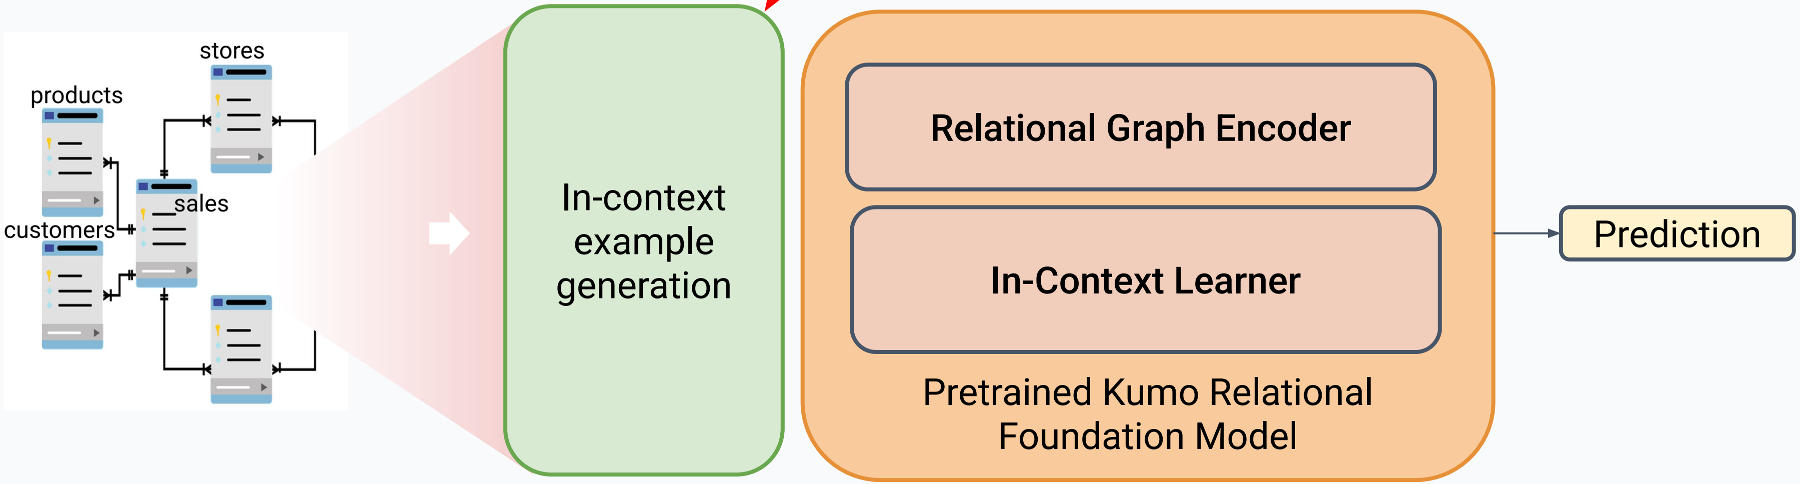

Foundation Model: In-Context Learning

Predictive Query → "Prompt"

```
PREDICT COUNT(TRANSACTIONS,0,30,days)=0
FOR EACH CUSTOMERS.CUSTOMER_ID
```

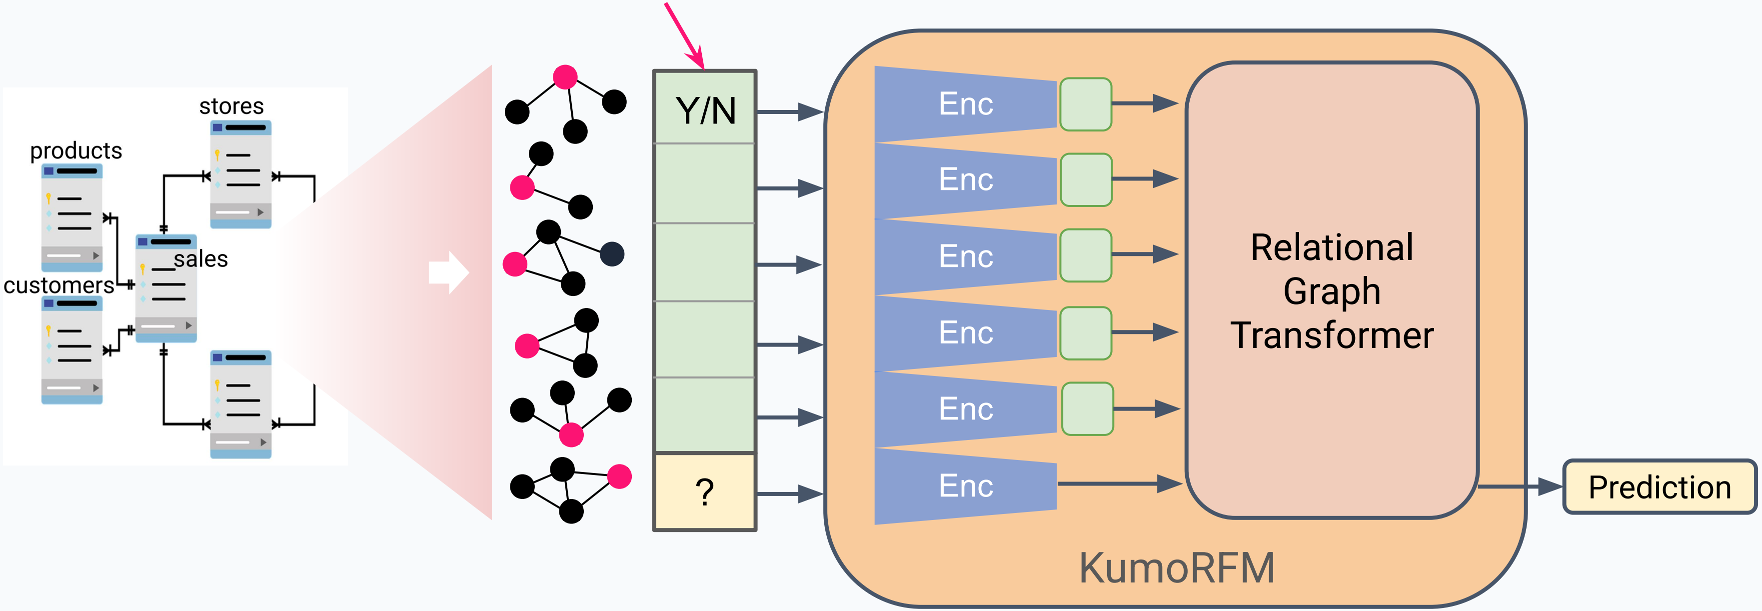

KumoRFM

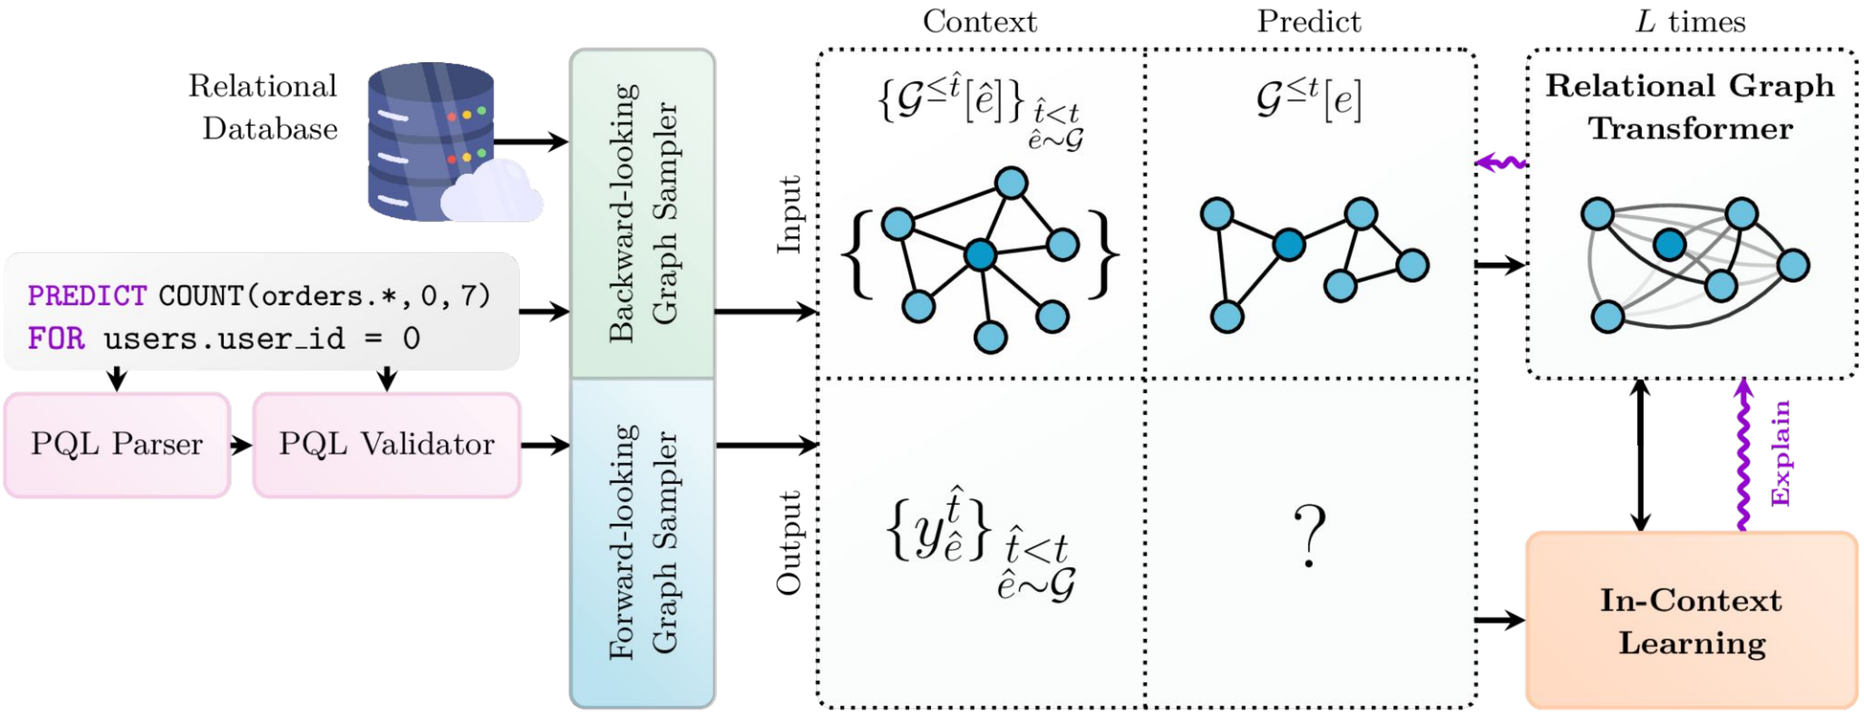
Relational Database
PREDICT COUNT(orders.*, 0, 7)
FOR users.user_id = 0
PQL Parser
PQL Validator
Backward-looking Graph Sampler
Forward-looking Graph Sampler
Input
Output
Context
Predict
$\{\mathcal{G}^{\leq \hat{t}}[\hat{e}]\}_{\substack{\hat{t}<t \\ \hat{e}\sim\mathcal{G}}}$
$\mathcal{G}^{\leq t}[e]$
$\{y_{\hat{e}}^{\hat{t}}\}_{\substack{\hat{t}<t \\ \hat{e}\sim\mathcal{G}}}$
?
L times
Relational Graph Transformer
Explain
In-Context Learning

KumoRFM

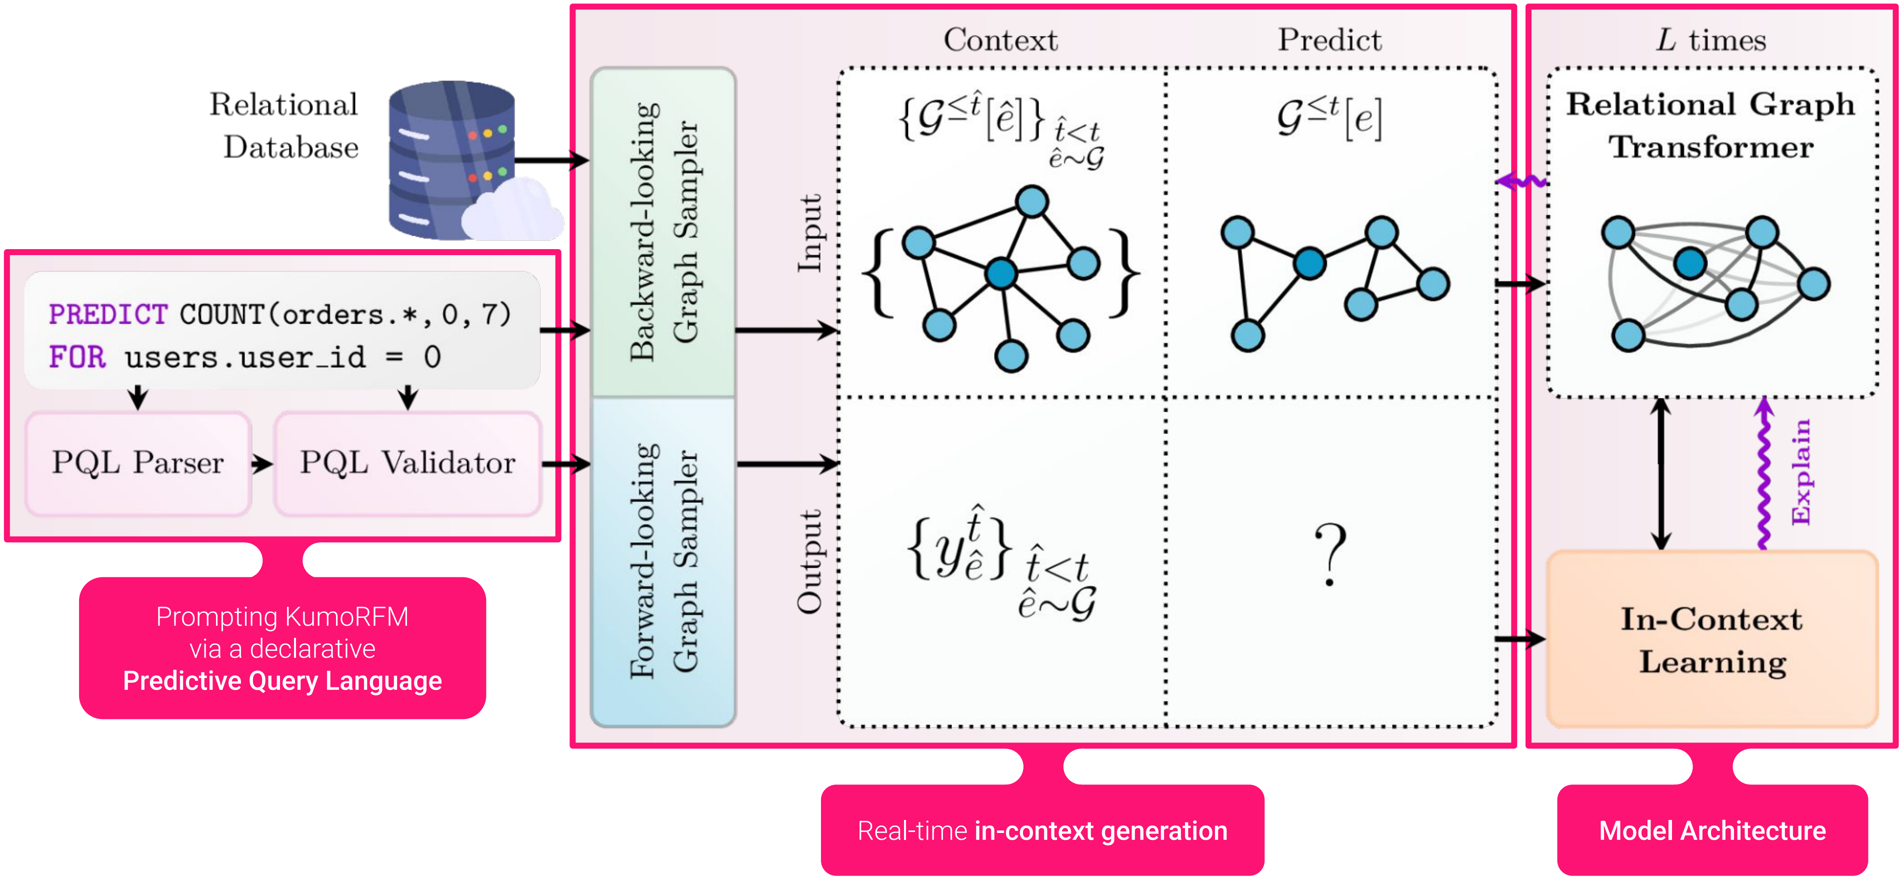

Relational Database
PREDICT COUNT(orders.*, 0, 7)
FOR users.user_id = 0
PQL Parser
PQL Validator
Prompting KumoRFM via a declarative Predictive Query Language
Backward-looking Graph Sampler
Forward-looking Graph Sampler
Context
Predict
Input
Output
$\{\mathcal{G}^{\leq \hat{t}}[\hat{e}]\}_{\hat{e} \sim \mathcal{G}}^{\hat{t}<t}$
$\mathcal{G}^{\leq t}[e]$
$\{y_{\hat{e}}^{\hat{t}}\}_{\hat{e} \sim \mathcal{G}}^{\hat{t}<t}$
?
Real-time in-context generation
L times
Relational Graph Transformer
Explain
In-Context Learning
Model Architecture

Prompting KumoRFM

We can *talk* to KumoRFM through the **Predictive Query Language Interface**

- It has a label definition (`PREDICT` clause) and entity definition (`FOR` clause)
- Additional filters can be applied both to label and entity clause
- Supports aggregations, binary operations and logical operations
- Capable of handling a broad set of task types:

first categorical value

```
PREDICT FIRST(orders.type, 0, 7)
FOR EACH users.user_id IN (0, 1, 2)
```

(a) Node multi-class/label classification

logical value

```
PREDICT COUNT(orders.*, 0, 7) > 0
FOR EACH users.user_id IN (0, 1, 2)
```

(b) Node binary classification

sum of numerical values

```
PREDICT SUM(orders.value, 0, 7)
FOR EACH users.user_id IN (0, 1, 2)
```

(c) Node regression

set of foreign keys

```
PREDICT LIST_DISTINCT(orders.item_id, 0, 7)
FOR EACH users.user_id IN (0, 1, 2)
```

(d) Link prediction

Explainability

- RFM is **fully-differentiable** and enables **gradient-based explanation techniques**
- Importance scores are computed on the **cell level** rather than on the feature level
- Conversion to **textual summary**

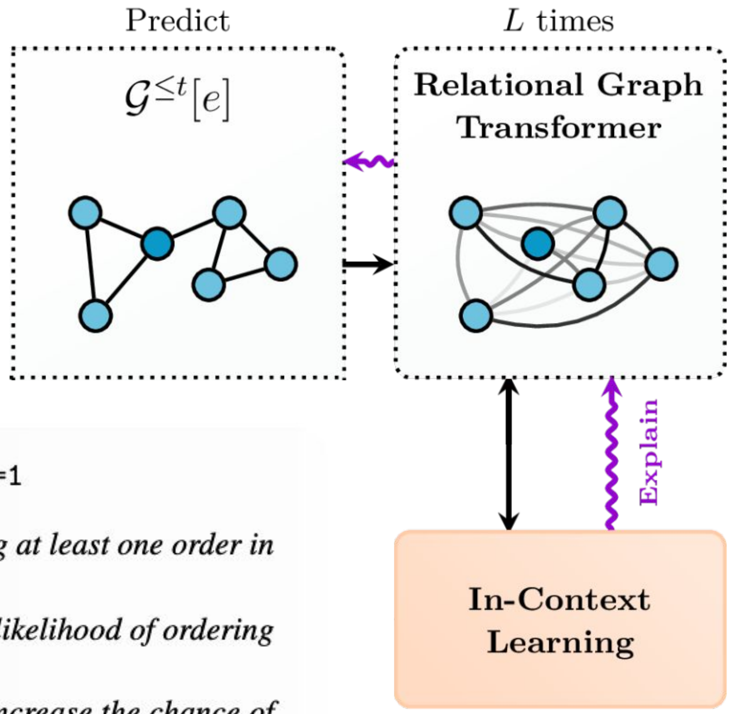


```
PREDICT COUNT(orders.*, 0, 30) > 0 FOR users.user_id=1
```

The model predicts that the user has a moderate likelihood of placing at least one order in the upcoming month. Key factors influencing this prediction include:

- ***Order Count:*** *Users with only a few past orders have a very low likelihood of ordering soon, while those with more orders show increased probabilities.*
- ***Order Date Recency:*** *Recent orders (6-12 months ago) greatly increase the chance of placing new orders soon.*
- ***Fashion News Frequency and Club Membership:*** *Users who regularly receive fashion news or have active club membership status show higher probabilities of ordering.*

Quantitative Accuracy Estimation

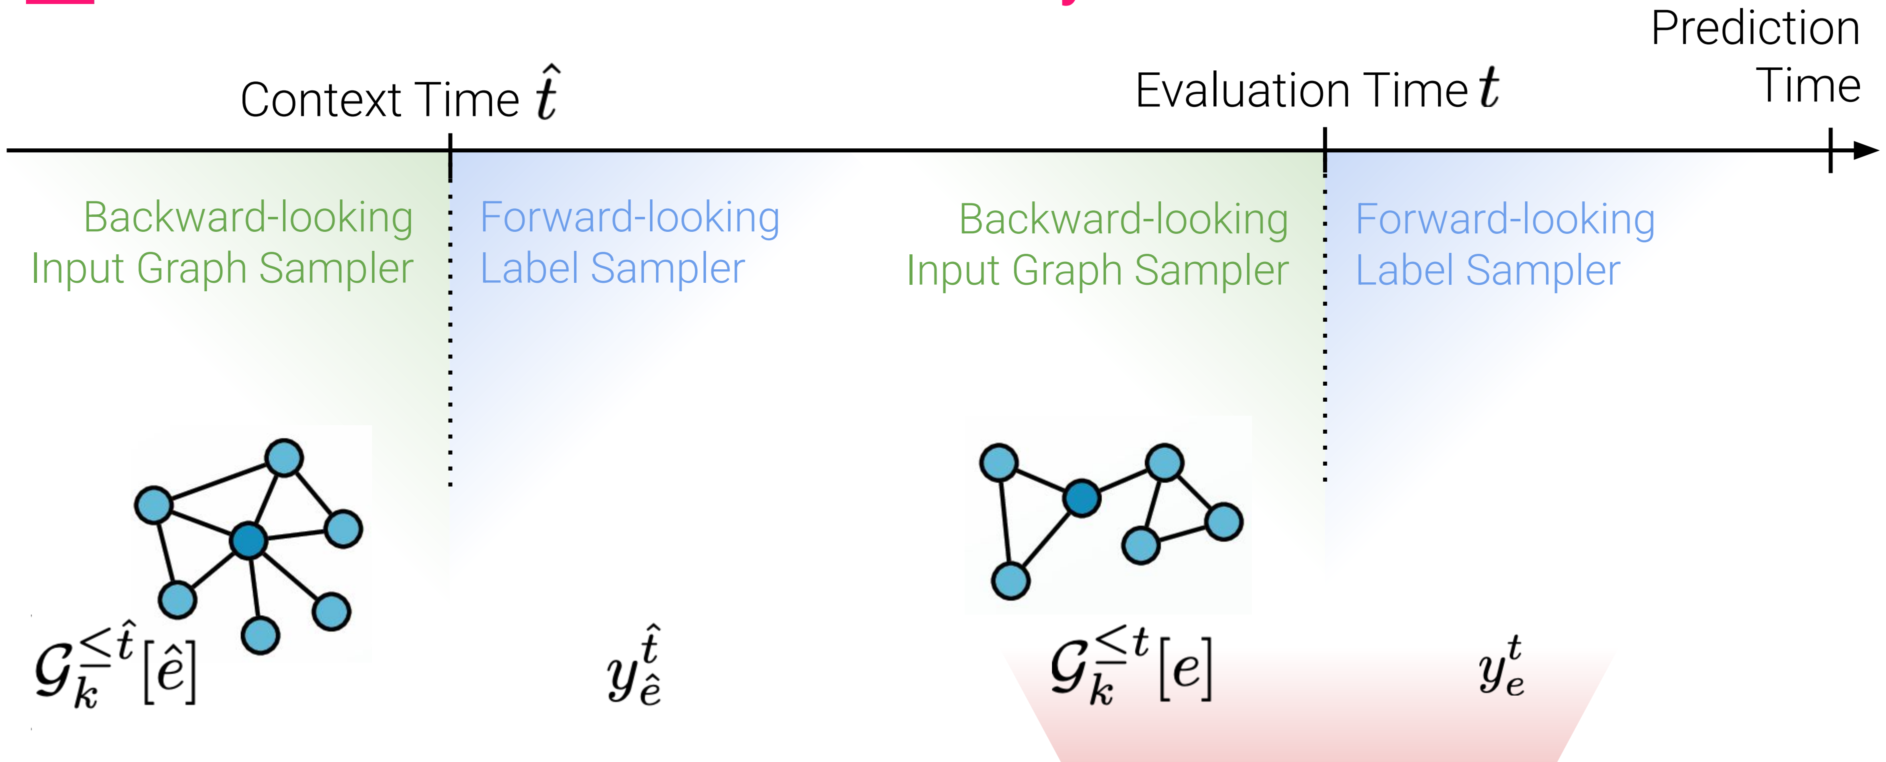

Build **trust into model predictions** by verifying its **correctness over historical data**

auroc	ap
0.91	0.72

Benchmarks

Results

We trained KumoRFM on a mixture of
publicly available relational databases and **synthetic data**
For label generation, we utilize a **random Predictive Query generator**

We evaluated KumoRFM on

RELBENCH
RELATIONAL DEEP LEARNING BENCHMARK

7 relational databases
30 predictive tasks (classification, regression recommendation)

Dataset	Domain	#Tasks
rel-amazon	E-commerce	7
rel-avito	E-commerce	4
rel-event	Social	3
rel-f1	Sports	3
rel-hm	E-commerce	3
rel-hm	Social	5
rel-trial	Medical	5
Total		30

KumoRFM has not seen any RelBench datasets during its pre-training phase, which guarantees no leakage!

Baselines

- **LightGBM:** Supervised ensemble of decision trees
- **Data Scientist:** An expert data scientist that solves each task by manual feature engineering
- **RDL:** End-to-end supervised GNN
- **LLM:** A `Llama 3.2` 3B model that is asked to do in-context predictions

Entity Classification

Dataset	Task	LightGBM (Supervised)	Data Scientist (Supervised)	RDL (Supervised)	LLM (Foundational)	KumoRFM (in-context) (Foundational)	KumoRFM (fine-tuned) (Foundational)
rel-amazon	user-churn	52.22	67.60	**70.42**	62.55	67.29	**70.47**
	item-churn	62.54	81.80	**82.81**	73.41	79.93	**82.83**
rel-avito	user-visits	53.05	—	66.20	53.36	64.85	**78.30**
	user-clicks	53.60	—	65.90	54.07	64.11	**66.83**
rel-event	user-repeat	53.05	—	76.89	53.36	76.08	**80.64**
	user-ignore	79.93	—	81.62	68.65	**89.20**	**89.43**
rel-f1	driver-dnf	68.86	69.80	72.62	80.03	**82.41**	**82.63**
	driver-top3	73.93	82.40	75.54	87.11	**91.07**	**99.62**
rel-hm	user-churn	55.21	69.00	69.88	63.81	67.71	**71.23**
rel-stack	user-engagement	63.39	90.30	90.59	81.23	87.09	**90.70**
	user-badge	63.43	86.20	88.86	79.99	80.00	**89.86**
rel-trial	study-outcome	70.09	**72.00**	68.60	59.17	70.79	71.16
Average ↑		62.44	—	75.83	68.06	**76.71**	**81.14**

KumoRFM is on par with best baselines without tuning, and outperforms them when fine-tuned.

https://kumo.ai/research/kumo_relational_foundation_model.pdf

Time-to-First-Prediction

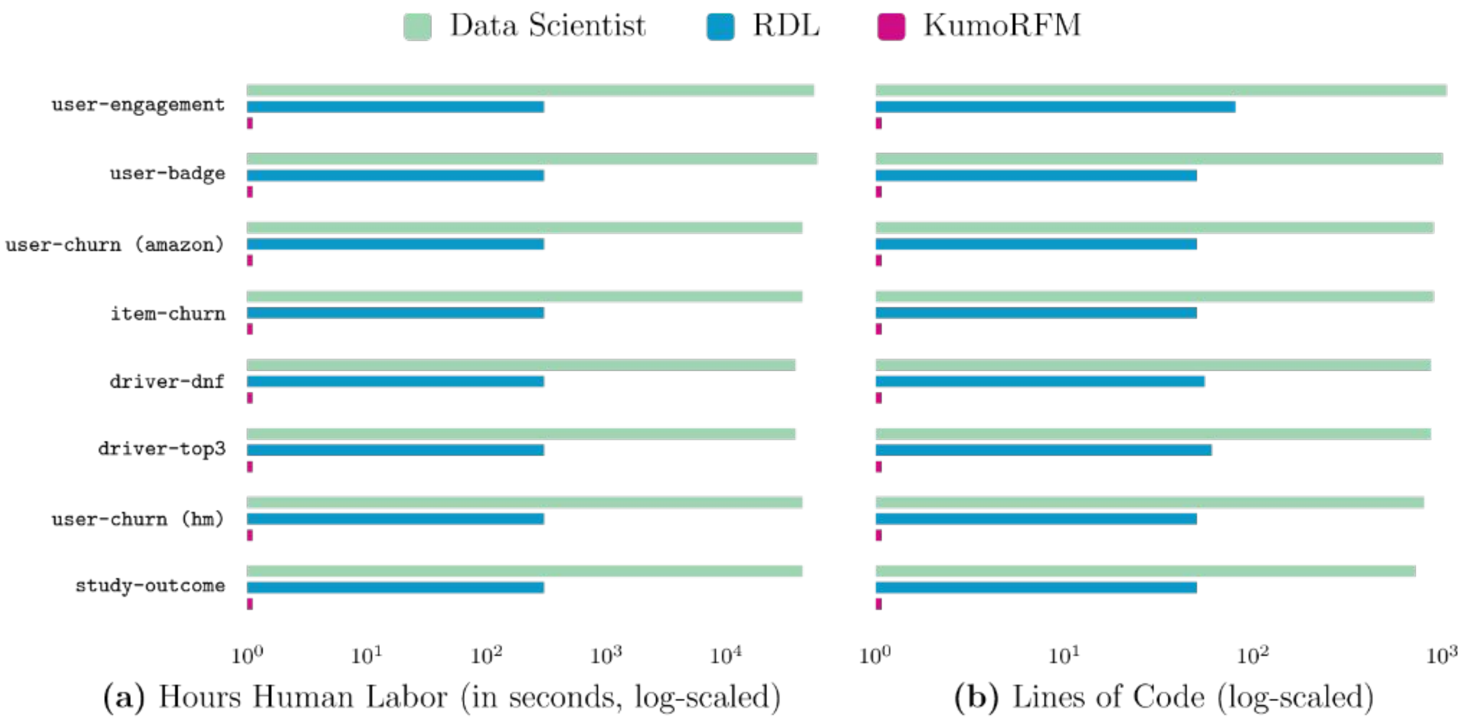

KumoRFM is 1,000 - 10,000 faster than alternatives!

KumoRFM: Summary

What is it? **A foundation model for business data!**

What can it do? **Make Zero-shot predictions on relational data!**

Can we make it even better? **Fine-tune on a particular task!**

Why does this matter?

- Predictions become commoditized
- Predictive models are democratized
- Prediction-driven applications become possible
- Data becomes even more important
- Fine-tuning empowers high performance use-cases

Q&A, Thank You & Next Steps

Share what you've learned on Linkedin/X! Tag @kumo_ai_team

Get a KumoRFM API key and start building with your data! kumorfm.ai

Quickstart notebook

https://tinyurl.com/hellokumo

Thinking about using KumoRFM?
Email hello@kumo.ai to learn more about enterprise options