Ciencia y Tecnología
Secretaría de Ciencia, Humanidades, Tecnología e Innovación

GrowGraph

Building a Research Digital Twin: From Reactive Management to Strategic Foresight

Secretary of Science, Humanities, Technology and Innovation
Mexico
21 November 2025

Alexander Belikov, GrowGraph
alexander-belikov.github.io

Talk Outline

1. The AI Revolution & Why Research Can't Wait
2. The Digital Twin Solution: KGs + Agentic AI
3. Proven Pilots: Real Impact in Research Workflows
4. Agentic Graph Optimization Platform

Gen AI technical revolution

Generative AI reached 40% adoption among U.S. adults (ages 18-64) in ~ 2 years.
For comparison it took Internet ~ 5 years and the Personal Computer (PC) ~ 12 years

Early adopters are already seeing tangible returns, with reports suggesting that for every $1 invested in GenAI, the average ROI is ~3.7x

Figure 1
Technology Adoption Rates, by Number of Years Since Introduction

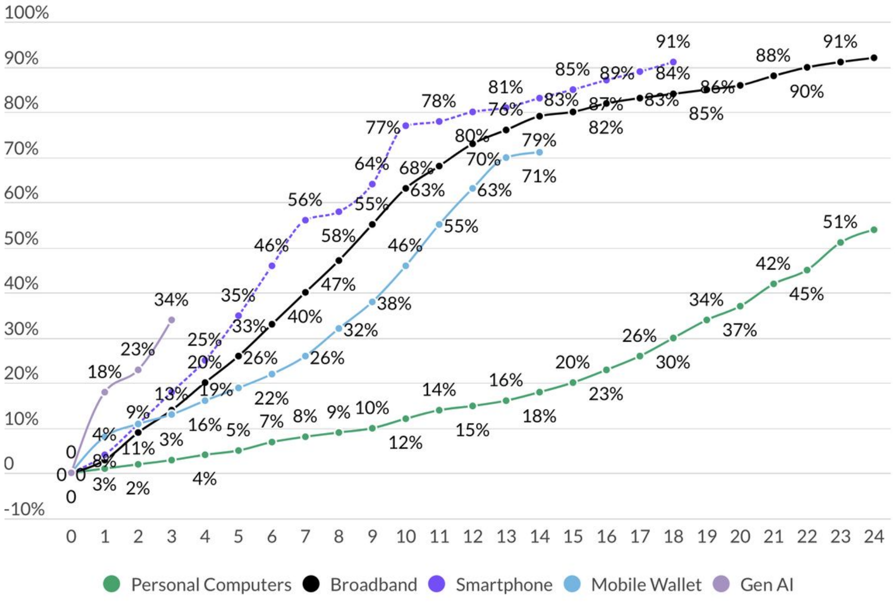

Source: PYMNTS Intelligence Analysis from Varied Sources

Gen AI Impact

While we haven't had time to react…

The design and art economy has already changed

- Creative production has shifted from hours to seconds, radically reducing the cost of ideation and iteration.
- Individuals now have studio-level capabilities, lowering barriers and enabling solo creators to compete with professional design teams.
- Demand has moved from “finished assets” to “infinite variations,” reshaping client expectations and pricing models.
- The value has shifted from manual execution to creative direction, making prompts, taste, and curation more important than technical skills.
- Agencies and brands are rethinking workflows around AI-native pipelines, where human designers supervise, refine, and orchestrate large volumes of generated content.

Area	Gen AI Impact
Software Development	Resource Shift: Automating 40% of coding tasks means shifting developer resources from maintenance/bug fixing to strategic innovation (R&D). Re-evaluate team structure and hiring profiles.
Education	With 80% reportedly using LLMs: puts at risk are Foundational Research practices, such as Information Retrieval, Critical Thinking and Problem-Solving.
Search	Financial Allocation: Marketing budgets are shifting rapidly from traditional search engine optimization (SEO) to AI-driven content generation and personalization.

The Exponential Conceptual Leap

The world is changing so fast, you have to run just to stay in place.

The acceleration of AI is not defined by faster software updates, but by the continuous, conceptual breakthroughs that redefine what is possible in a single cycle. New paradigms emerge monthly, not annually.

- Generative AI (Chatbots): Human-level text synthesis and creation.
- Agentic AI: Systems that plan, execute, and adapt complex, multi-step tasks.
- Multimodal Systems: Seamlessly unifying text, image, code, and sensory data.
- Knowledge-Augmented AI (Deep Research): Fusing the creativity of LLMs with the rigor of structured facts and knowledge graphs (e.g., RAG, Neural-Symbolic methods).

Generative AI: the Multi-Tool Interface

The most striking feature of Generative AI (GenAI) is not its capability, but its universal interface.

Unlike traditional software, where humans had to adapt to a new application for every new task, GenAI uses a single, familiar interface - the natural language chat prompt - to execute a vast and ever-expanding range of functions.

GenAI is the ultimate multi-tool, unifying cognitive tasks into a single dialogue:

Traditional Tool/Process	GenAI as the Multi-Tool
Reading & Digesting	Summarization: Condensing a 100-page document into a few key bullet points.
Data Entry & Manual Review	Information Extraction: Pulling structured data (names, dates, figures) from unstructured text.
Spreadsheets & Formulas	Calculation & Accounting: Performing complex analysis and financial modeling via simple prompts.
Brainstorming & Drafting	Content Generation: Drafting emails, reports, presentations, or creative copy in seconds.
Research & Consultation	Inquiry & Counseling: Providing first-line context, acting as a sounding board, or surfacing 'deep research' connections.
Programming & IT	Code Generation: Writing, debugging, translating, and documenting software on demand.

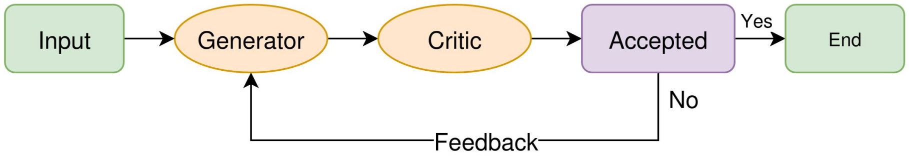

Agentic AI

Agents are LLMs with initiative:
It's a Large Language Model that doesn't just answer questions - it sets goals, makes plans, and takes action.

The Go-Getter Flow: Operates on an iterative loop: Plan → Execute → Critique → Refine.

Knowledge-Grounded Reasoning: Uses Retrieval-Augmented Generation (RAG) to find, verify, and incorporate real-time data from knowledge bases.

Self-Correction is Key: Includes an internal "Critique" mechanism to evaluate its own output and fix mistakes autonomously.

Tool Masters: Connects the LLM to specialized tools (like code execution, APIs, or database queries) to interact with the world.

From **Answerer** to **Investigator**: Shifts the AI role from a reactive chat partner to a proactive, independent researcher.

Dynamic Goal Pursuers: Capable of breaking down a complex, high-level goal into a sequence of smaller, manageable tasks.

Organizational Implications of Generative AI Adoption

The transition of GenAI from individual productivity booster to a fundamental, company-wide capability is a complex strategic and operational challenge. The gains that are obvious at the individual level are much more difficult to achieve at the institutional level.

Early attempts by large companies to realize massive, quick workforce reductions based on AI often failed because they underestimated the complexity of this deployment. The potential gains are undeniable, but they require a structured, systemic approach.

The Critical Adoption Challenges

- Finding the Internal Customer (The 'Where to Start' Problem):
    - Focus: Identifying high-value, repetitive, and well-defined tasks (e.g., data review, legal summarization, tier-one customer support) where GenAI can deliver demonstrable, measurable gains.
    - Failure Point: Trying to apply a general-purpose AI solution to vague, ill-defined, or highly unique problems.
- Technical Deployment (The 'Last Mile' Problem):
    - Challenge: Moving from pilot to production requires robust, scalable, and secure AI Platforms that can handle enterprise data security, regulatory compliance, and massive user loads.
    - Solution: Developing a central, modular architecture that supports incremental, composable pilots and manages the entire lifecycle of models and agents.
- Making it Comfortable to Use (The 'Ecosystem' Problem):
    - Key: AI tools must be seamlessly embedded in the existing ecosystem (e.g., within Microsoft Office, Slack, Salesforce, or proprietary company software) rather than creating new interfaces.
    - Goal: Minimizing context switching and integrating the AI into the natural flow of work to ensure high user adoption.

The Path Forward: Incremental, Composable Integration

AI adoption is inevitable, but failures are expected. The most successful strategy avoids large, all-or-nothing transformations. Instead, institutions should focus on:

- Composability: Building small, self-contained AI services that can be mixed and matched.
- Incremental Pilots: Starting with small, focused projects that demonstrate clear ROI before scaling.
- Platform Integration: Ensuring every pilot integrates back into a unified, secure AI platform to prevent "agent sprawl" and siloed efforts. This approach allows the organization to learn, adapt, and build institutional knowledge while mitigating large-scale risk.

Foundation Layers for Enterprise AI Adoption

Moving from a pilot to enterprise-scale AI means confronting the physics and economics of compute. The decision to build versus buy involves a massive commitment.

GPU Cluster Investment (The Hard Constraint):

- Minimum Entry Size: A production-ready cluster for top-tier players starts at 8,000 GPUs.
- Power Demand: Requires significant dedicated infrastructure, typically 10–15 MW.
- Lead Time: Building on-premises implies a substantial time commitment:
  - 1–2 years if full infrastructure (transformer, cooling) needs to be established.
  - 6 months if no hardware is yet secured.
  - 2-3 months if hardware is available.
- Support: Requires significant in-house expertise, with a support team of approximately 10–20 specialized DevOps/SRE personnel.
- Cost/Pricing: Public cloud prices are anticipated to drop, but an on-premise foundation remains a huge capital and operational expenditure

The Model & Platform Layers

Model Layer (Vendor vs. Open Source)

- Vendor Solutions (LLM APIs): Outsourced, hosted APIs (e.g., GPT, Gemini, Claude) offer the fastest time-to-market with the lowest operational overhead.
- Open Source Models: Running self-hosted solutions (e.g., Llama) offers maximum control, better security, and cheaper inference at massive scale, but requires the massive infrastructure described above. They may be perceived by some, particularly AI engineers, as "not up-to-date" for cutting-edge use cases.

Application & Orchestration Layer (The Productivity Engine):

Bridges the model with enterprise processes, transforming LLM calls into intelligent, measurable workflows.

- Agentic Frameworks: Tools like Langgraph and CrewAI enable complex multi-step reasoning, memory, and self-correction.
- Knowledge Retrieval & Structuring: Libraries like LlamaIndex facilitate Retrieval-Augmented Generation (RAG), grounding LLMs in proprietary, internal data.
- Enterprise Management: This requires robust DevOps solutions and Monitoring (like Google ADK) for managing deployment, tracking costs, ensuring security, and guaranteeing performance (latency, accuracy).

What is a Knowledge Graph?

A **Knowledge Graph (KG)** is a **graph-based data structure** where entities (nodes) and their relationships (edges) are explicitly represented.
It integrates **structured and unstructured data** into a unified, queryable format.

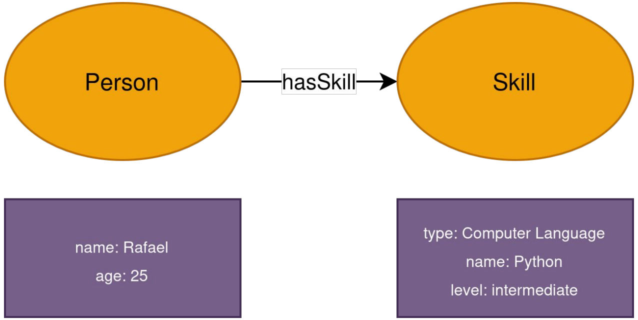

Core Components

- **Entities**: e.g., companies, people, publications, events
- **Relations**: e.g., *employs*, *cites*, *reports_to*
- **Attributes**: properties of nodes/edges
- **Semantics**: often enhanced with ontologies or schemas for interpretability

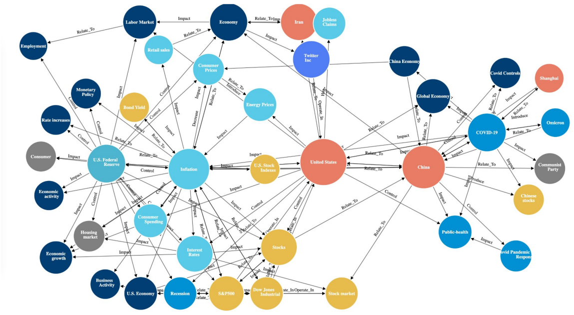

Why It Matters

- Handles **heterogeneous, relational data** natively (e.g., linking funding to publication metrics)
- Enables **non-tabular** (e.g., Graph Neural Networks, community detection)
- Captures **contextual and temporal dependencies** (which might be causal)

The Synergy of LLMs and KG

Data → LLM → KG : LLM can generate Knowledge Graphs from unstructured data.

LLM → KG → LLM → User : KGs help to eliminate LLM hallucination via RAG: by adding relevant, verified parts of the Knowledge Graph to the LLM's prompt in a Retrieval-Augmented Generation (RAG) approach, we solve the hallucination problem by grounding the LLM in factual relationships. The KG serves as the memory and source of truth.

KG provides content, LLM provides shape. But KG goes beyond this simple interaction!
⇒ **Statistical Rigor**: They act as the foundation for modern machine learning, enabling Graph Neural Networks (GNNs) to detect anomalies, cluster entities, and predict unknown relationships across the entire dataset, optimize for metrics of interest.

The Epistemological Challenge of Organizational Measurement

The Pitfalls of One-Dimensional Metrics

The Illusion of Simplicity: We grew up in a world driven by single, monolithic metrics:

- Financial Returns for businesses (ROI).
- GDP for national performance.
- Hirsch Index (H-Index) for researchers.

The Epistemological Question: Why are we sure that the overall "goodness" or strategic success of a complex organization can be reflected by just one or two metrics?

The Maximization Paradox: Maximizing one metric often leads to the minimization or catastrophic failure of others.

- **Historical Example**: Soviet-era maximization of steel output led to massive ecological damage and economic inefficiency.
- **Organizational Example**: Maximizing short-term productivity (e.g., working 16 hours/day) leads to long-term efficiency decay and burnout.

The Problem of **Diminishing Returns**

- Convexity of Utility: Microeconomics teaches us about the convexity of the utility function. Our sensitivity to the first unit of return is vastly higher than to the thousandth.
    - The utility gained from a researcher's first citation is much higher than that gained from their 1001st.
- Strategic Blindness: Focusing on continuous maximization of a high-value metric with limited resources simply means harvesting one area to the point of future catastrophe in another.

The Mandate for Multi-Dimensional Optimization

From Scalar ROI to Portfolio Optimization

The function of any organization is inherently complex, defined by a portfolio of inputs and outputs. To optimize performance, we must define it in multi-dimensional terms.

Defining the Performance Space

- **Input Portfolio**: Investments, R&D spend, Human Capital (HC).
- **Output Portfolio** (KPIs): Financial returns, academic performance, societal impact, ecological compliance, talent retention, etc.
- **The Trade-Off Challenge**: Selecting metrics is not enough; we must understand the trade-offs. Maximizing one metric (e.g., Speed) may require minimizing another (e.g., Safety).
- **The Decision-Making Apex**: At the highest level, a decision should evaluate: What does 1 invested dollar generate?
    - Condition A: Given the uncertain scenario (e.g., recession, technological boom).
    - Condition B: Given the organization's preference vector (the desired weighting of long-term vs. short-term returns).

The Hierarchy of Performance

- **Organizational performance** is ultimately defined by the performance of its **constituents** and their interactions.
- We must track a **hierarchy of metrics** that starts at the highest level (global impact) and **cascades** down to the single employee.
- **Role**: Each component of the system functions best when it is optimized according to its specific vocation and contribution to the overall strategic goals.

Knowledge Graphs and Strategic Foresight

Enabling Multi-Dimensional KPIs

Knowledge Graphs (KGs) are the foundational technology for managing and optimizing this multi-dimensional complexity.

Beyond Silos: Transversal Metric

- The KG Advantage: KGs are uniquely suited to connect data across traditional organizational silos (HR, Finance, R&D, Operations).
- Transversal Metrics: KGs enable the evaluation of implicit and transversal metrics - connections that cannot be easily measured in tabular or siloed data.
    - Example: Tracing a specific R&D investment ($) to a resulting Societal Impact metric via Talent Mobility (HR) and Scientific Output (R&D).
- Agentic Evaluation: This structured data allows Agentic AI to run sophisticated simulations:
    - Evaluate the certainty of any metric prediction, translating outputs into traceable scenarios.
    - Serve as the single source of truth for all supporting evidence, ensuring traceable decision-making.

The Strategic Foresight Outcome

The convergence of multi-dimensional KPIs tracked by KGs, and optimized by Agentic AI, allows the organization to move from reactive management to Strategic Foresight - maximizing the utility function across time horizons and trade-offs to ensure long-term, sustainable performance.

The Digital Twin: Modeling Reality for Optimization

The Digital Twin is the operational culmination of structured data (KG) and multi-dimensional performance tracking (Metrics).

DT moves the organization from measurement to emulation and prediction, by incorporating real-time observations of interconnected entities at different scales.

It is engineered to:

- **Emulate Dynamics**: Simulate the system's behavior over time under various conditions.
- **Uncover Causal Relations**: Move beyond correlation to identify why certain outcomes occur.
- **Produce Statistical Predictions**: Forecast future states and resource needs with verifiable confidence.
- **Adaptability**: Remain flexible with respect to scope widening and ontology extension as the organization evolves.

The Foundational Architecture of the Digital Twin

The Digital Twin is a sophisticated, continuously updated decision-making engine:

- **Knowledge Graph**: Provides the structured, semantic "source of truth." It models the interconnected entities and their relationships across silos.
- **Models (Simulation & AI)**: Executes the predictive and dynamic emulation. This includes physics-based models, machine learning models, and complex Agentic AI frameworks.
- **Policy Validation (Optimization Engine)**: The engine for strategic decision-making. It tests hypothetical changes (policies) against the Twin to evaluate trade-offs against the organization's multi-dimensional metrics before deploying changes to the real world.

Potential Use Cases

Reviewer Discovery & Conflict-of-Interest Checking

Principal client: Grant evaluation offices / National Funding Agency directorates

Value: Automates reviewer selection, reduces bias, ensures conflict-of-interest compliance.

Current pain points:

Reviewer selection currently takes up to 8 weeks and has 30% conflict-of-interest rate

Quantified benefits:

Reduces reviewer discovery time by 80%

Funding Portfolio Optimization

Principal client: Budget and Program Planning Divisions

Value: Optimizes allocation across programs, detects redundancy, improves ROI of funding.

Potential Use Cases

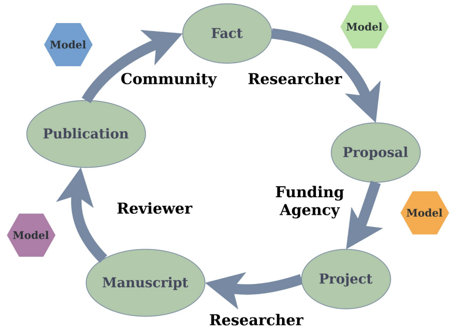

Talent Mobility & Career Path Analytics

Principal client: Human Resources and Talent Development Offices
Value: Tracks national talent flows, detects shortages, supports career policy.

Monitoring Scientific Impact & Accountability

Principal clients: Researchers, Accountability/audit divisions; Government oversight bodies (Cour des Comptes), Parliament reporting units
Value: Track grants to outputs, automate reporting, increase transparency and trust.

Funding Metrics: Funding the Frontier

A system that tracks Funding Impact!

An interconnected data collection of 7M research grants, 140M scientific publications, 160M patents, 10.9M policy documents, 800K clinical trials, and 5.8M newsfeeds, with 1.8B citation linkages among these entities.

The effectiveness and usability of the system is evaluated using case studies and expert interviews.

Wang et al., [arXiv:2509.16323]

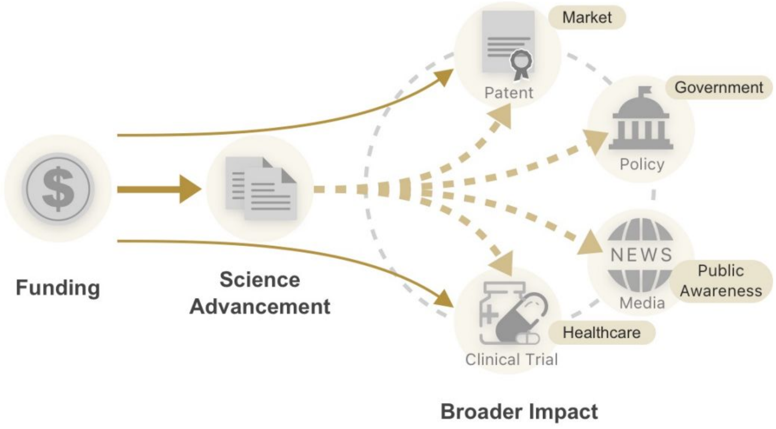

XSI: Knowledge Graph Generation

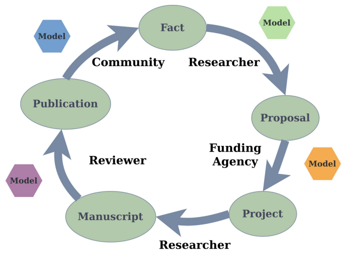

Input Data: **arXiv**, **biorXiv**, **medrXiv**

Relation Extraction: **Syntactic Features**

Entity Linking: **wiki, CHEBI, NCBI, omim**

Graph Representations:

Raw KG -> Redux KG -> Reference KG

Reference KG: KG of entities with weighted edges

arXiv:2502.13912

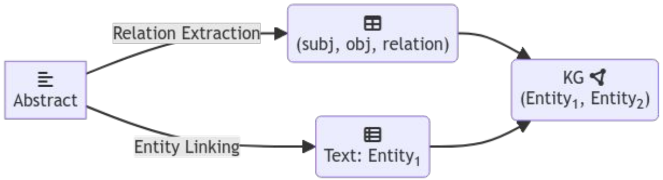

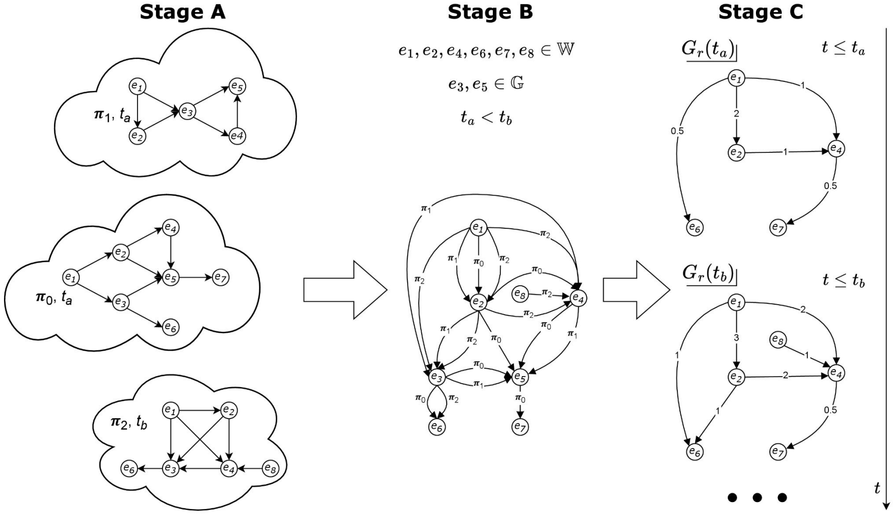

What is Semantic Impact?

Total entity graph: *Γ(t)* at time *t*.
Publication subgraph: *γ*
Semantic impact (SI):

$$J_\gamma = \frac{1}{|\gamma|} \sum_{e \in \gamma} \left(\frac{w_e(\gamma \bigcup \Gamma(t+\Delta))}{w_e(\gamma \bigcup \Gamma(t))} - 1 \right)$$

Low impact case: publication subgraph repeats the edges already present in the reference graph at time *t* and these edges do not have higher weight at time *t+Δ.*
High impact case: publication subgraph introduces new edges in the reference graph at time *t* and these edges acquire high weight by time *t+Δ.*

💪Less biased than citation counts, correlates with disruptive research

Models, Features

Features: network and diffusion features for publication subgraph γ (info flow through γ)

Develop models for prediction of $r = \log J_\gamma$ and **prediction error** $\varepsilon = |r - r^*|$, with a horizon of 3 years, where r^* is the ground truth value.

Out-of-sample model performance (R-squared) is >0.7 for **SI** and 0.2 for **SI error.**

NB: Covid trough.

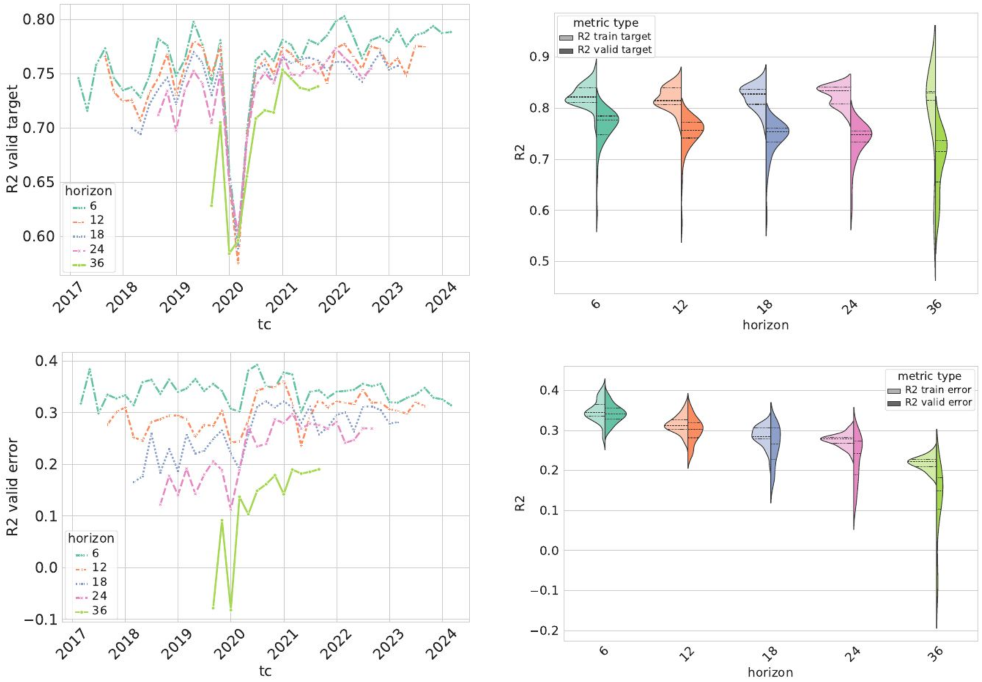

Efficient Frontier?

Financial mathematics analogy
impact : return, *error* : volatility

No time series: each publication is a single observation.

impact: $r = \log J_\gamma$
prediction error $\varepsilon = |r - r^*|$, with a prediction horizon of $\boldsymbol{t}$ years, where r^* is the ground truth value

NB: not at all Mean-Variance problem but an interesting analogy.

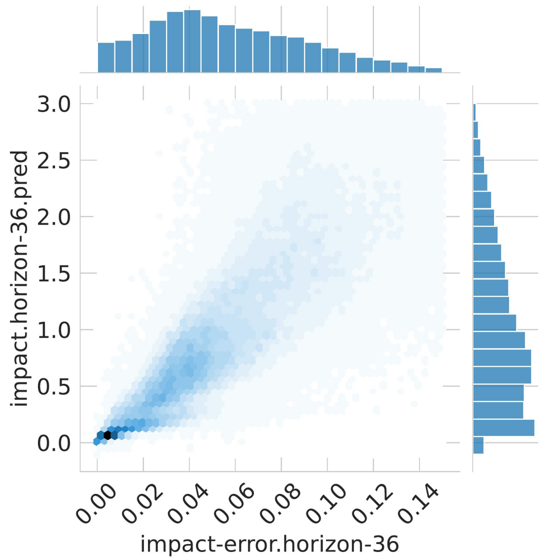

Portfolio Optimization For Academic Publications

For each period we take available publications, fix the fraction α of publications to choose, predict out of sample (future)

Integer Linear programming (0-1 knapsack)

$$\max \sum_{i \in I} (r_i - \epsilon_i) x_i$$

$$s.t. \frac{1}{|I|} \sum_{i \in I} x_i \leq \alpha$$

$$x_i \in \{0, 1\}, \quad \forall i \in I$$

Prediction of at *t+3* years

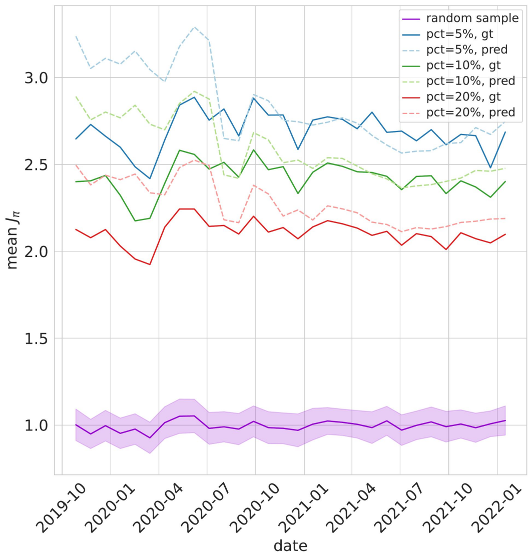

Biomedical Preprint Predicted Semantic Scores Online

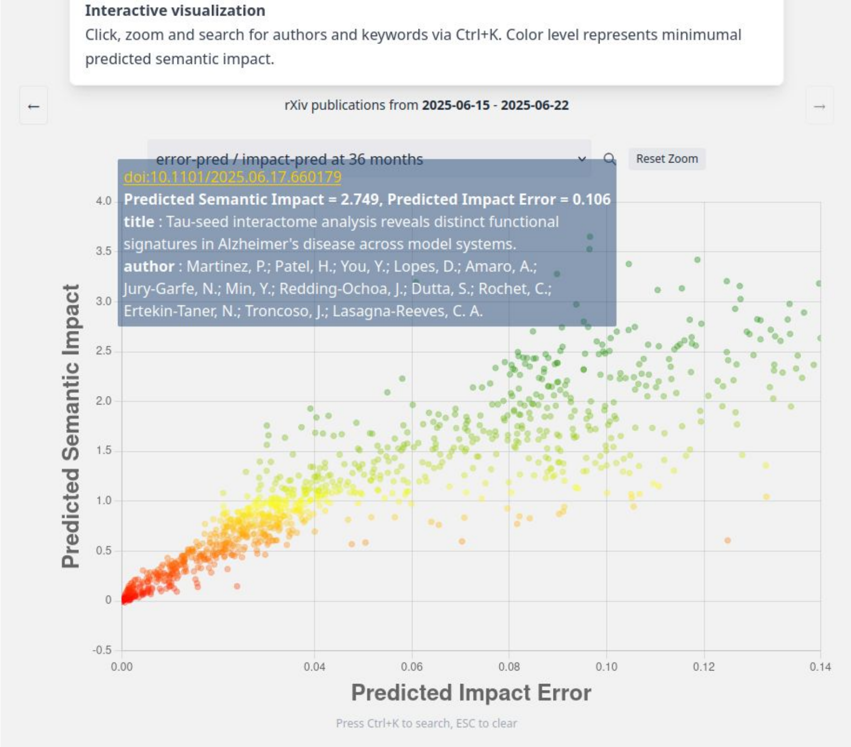

Results

- SI has low but statistically significant correlation with citation counts (~ 0.2)

Shortcomings

- Quality of research execution is not evaluated

Future Steps

- Extend the model beyond biomedical domain
- Add the quality of research execution

https://growgraph.dev/app-public/xsi

Prediction of robust scientific facts from literature

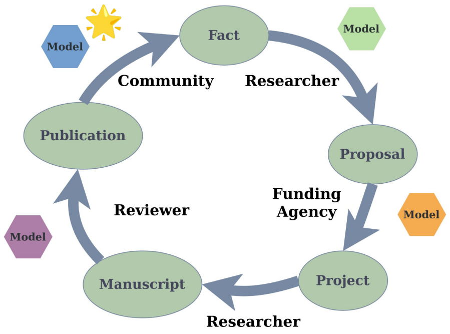

Literature

Experiment

Rzhetsky, A. et al, 2004

Poon,H. et al, 2014

Library of Integrated Network-Based Cellular Signatures
Subramanian, A. et al, et al, 2017

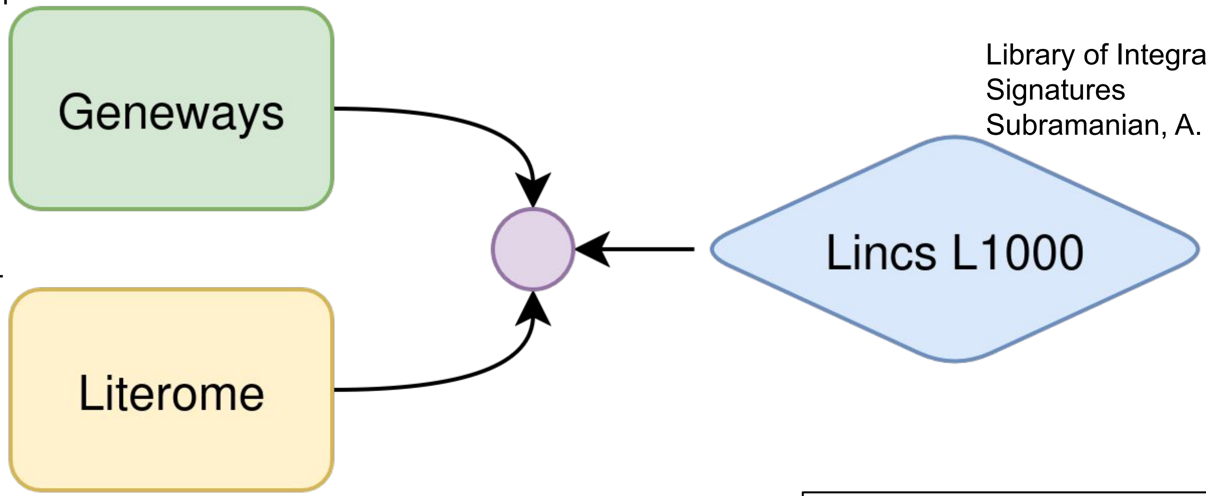

Belikov et al., *Nature Machine Intelligence* **4**, 445 (2022)

The goal of study: To use large-scale experimental data (LINCS L1000) to rigorously validate and predict the correctness of gene-interaction relationships extracted from the noisy biomedical literature."

Integral Modeling Results

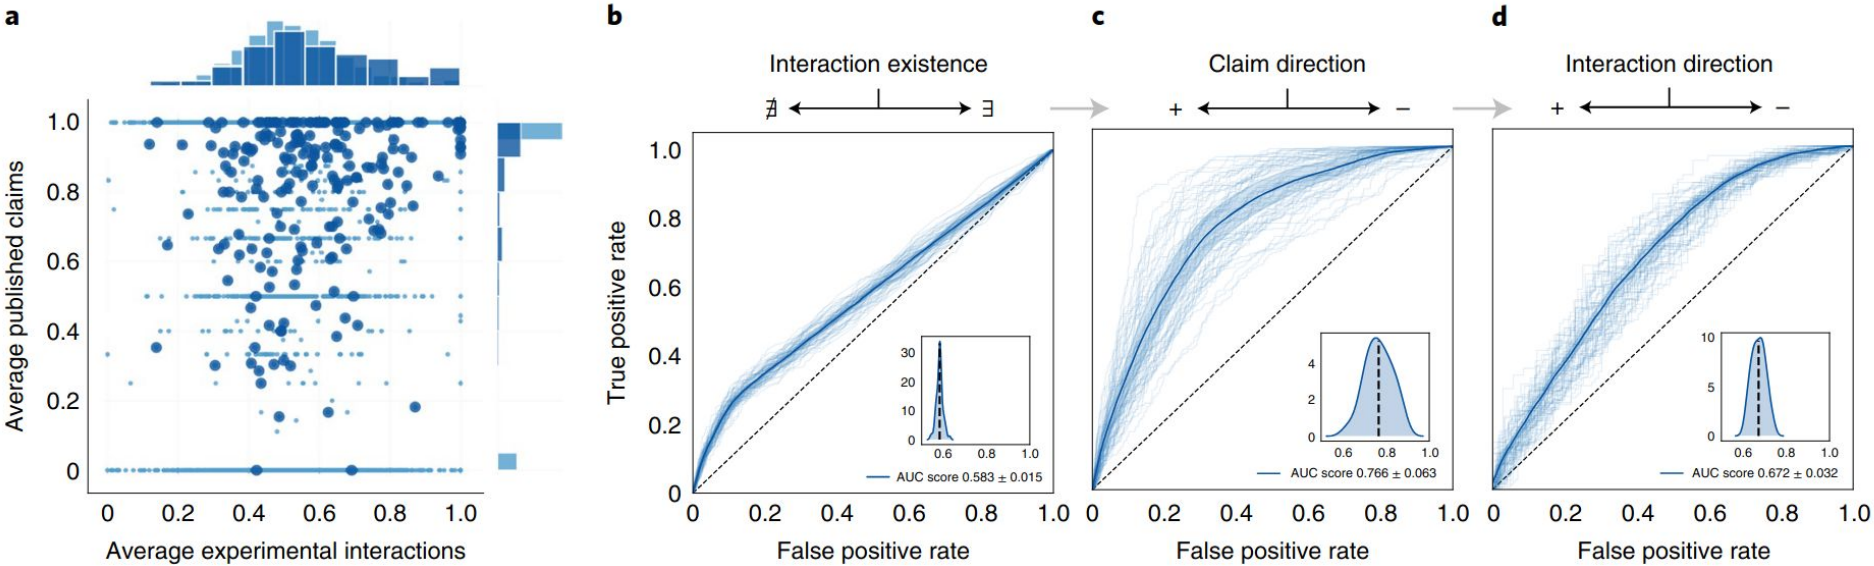

$$P(\pi_+^\alpha \mid \{(c_i^\alpha, f_i^\alpha)\}) \propto \prod_i P(c_i^\alpha, f_i^\alpha \mid \pi_+^\alpha) P(\pi_+^\alpha) \propto$$

$$\prod_i P(\pi_+^\alpha) \sum_{y_i^\alpha} P(c_i^\alpha \mid y_i^\alpha \pi_+^\alpha) P(y_i^\alpha \mid f_i^\alpha) \quad .$$

Policies Accelerating Consensus Convergence

Policy A: promote independence

Selecting subsamples with more communities improves AUCs of interaction prediction model

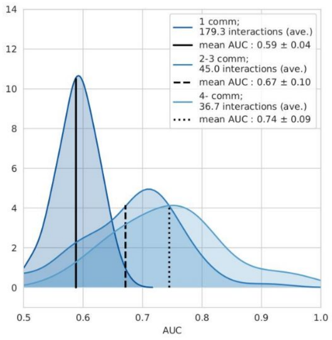

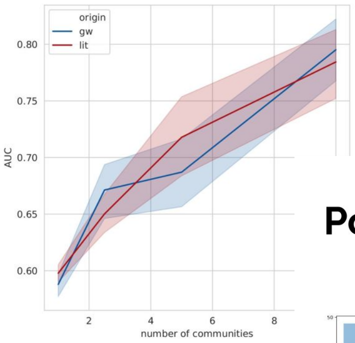

Policy B: altering the attention

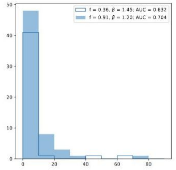

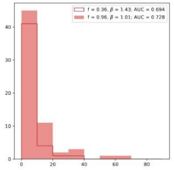

Flatter distributions result in higher information gain

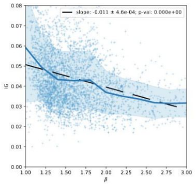

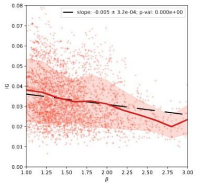

$$IG = \mathrm{ent}(p^{(0)}) - \frac{1}{k}\sum_{\alpha=1}^{\alpha=k} \mathrm{ent}(p^{\alpha})$$

$$\mathrm{ent}(p^{\alpha}) = -\sum p_i^{\alpha} \log p_i^{\alpha}$$

Emberloom: Predicting Career Transitions

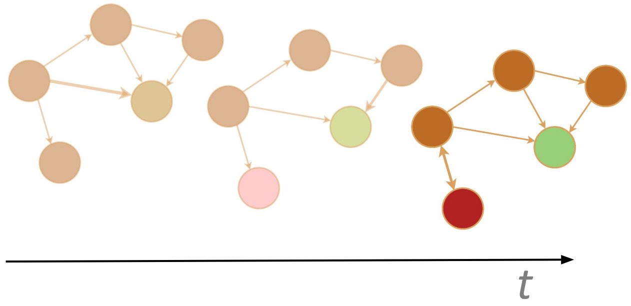

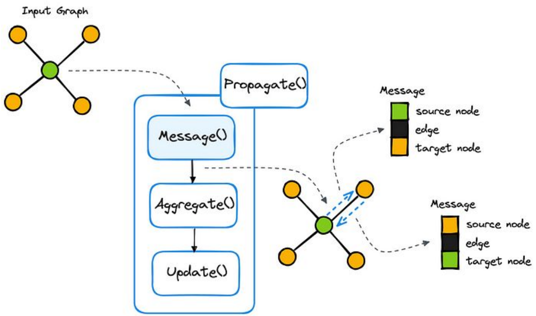

- History of a company with 40K+ employees over 36 months
- Positions do not form a total order: engineer track, management track
- Available features : email exchange frequency
- Use embeddings to represent employee latent state, GNN for modeling
- F1 score (harmonic mean of precision and recall): 0.761

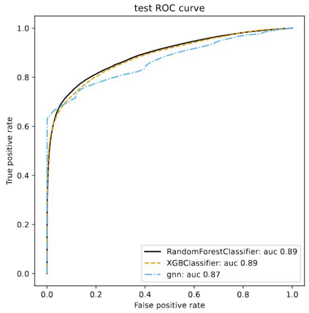

Reviewer Recommendation

Relationship schema:

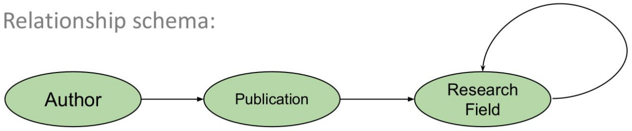

Given the publication records, who would be a good reviewer for paper?

Solutions

1. Recommend Reviewers
    a. construct **[coDomain] Author ↔ Author** from Author → Publication → ResearchField
    b. derive communities in **[coDomain] Author ↔ Author**
    c. For each new publication pick reviewers from the same **[coDomain]** comm id that have a different **[coAuth]** comm id

Publication -> Field

Author -> Publication

Field -> Field

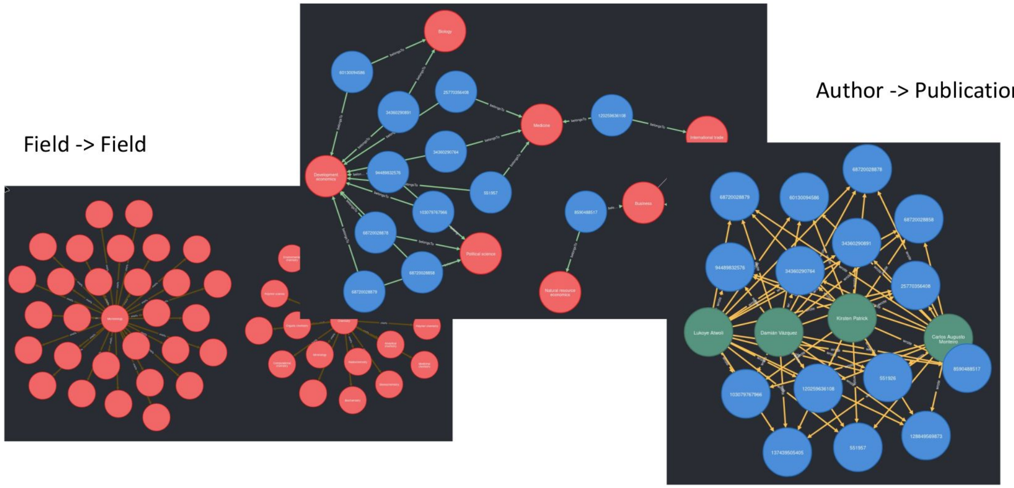

OntoCast

Ontology Assisted Agentic Transformation to Semantic Triples

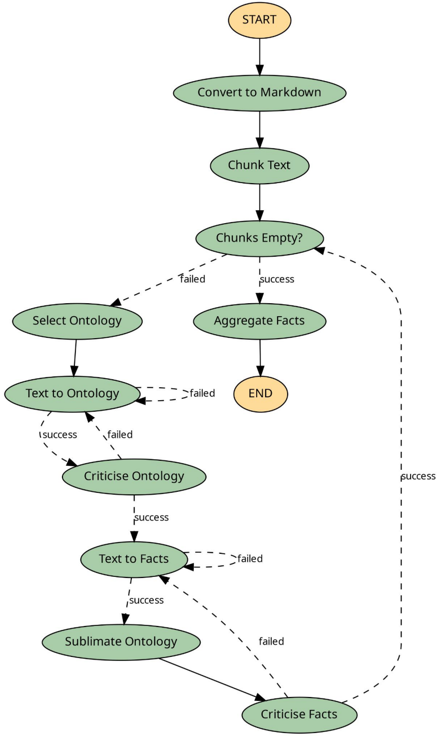

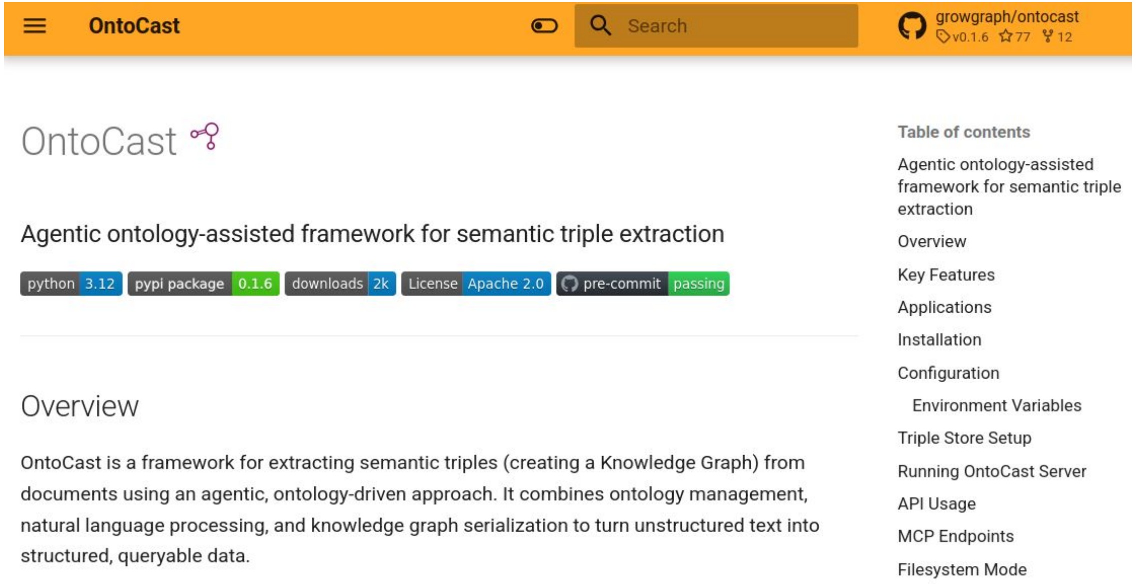

Projects:

- PerovKG with Lund University
  - HAMLET Physics Conference 2025, Denmark
  - MATSUS 2025 Conference Spain
- FcaOnt with AMU University (France)

https://growgraph.github.io/ontocast

GraFlo

Property graph DBs are the workhorses of graph modeling and analytics, they facilitate data prep for complex models.

Graflo is a multi-adapter that speeds up data ingestion in Property graphs. It features:

- Declarative transformations vs. custom ETL coding
- Adapters for Neo4j, ArangoDB and TigerGraph: Multi-database adapter reduces vendor lock-in
- Tested on graphs with billions of edges

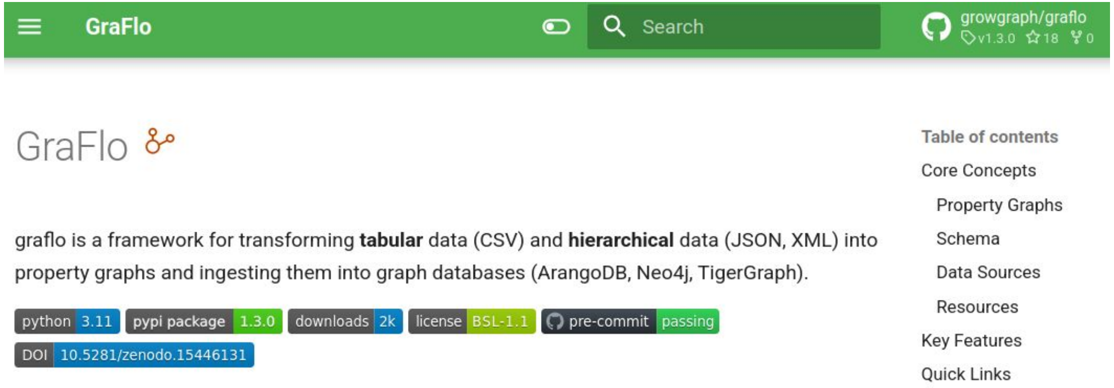
GraFlo
Search
growgraph/graflo
v1.3.0 18 0

GraFlo

graflo is a framework for transforming **tabular** data (CSV) and **hierarchical** data (JSON, XML) into property graphs and ingesting them into graph databases (ArangoDB, Neo4j, TigerGraph).

python 3.11 | pypi package 1.3.0 | downloads 2k | license BSL-1.1 | pre-commit passing
DOI 10.5281/zenodo.15446131

Table of contents

Projects:

- XSI
- DebKG
- Diversieve
- ReviewerRec

Being piloted at **BNP Paribas** for Infrastructure Digital Twin.

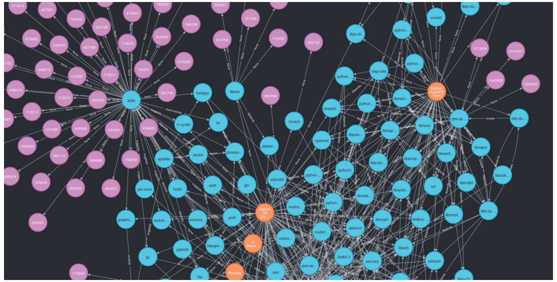

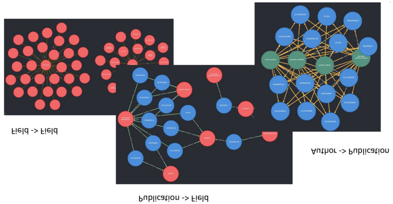

https://github.com/growgraph/graflo

OntoCast + GraFlo + Graph Optimization

🔧 **Ontocast: Ontology-Assisted KG Construction**

Automated Knowledge Graph Generation from Unstructured Data

- **Semantic Triple Extraction**: LLM-powered entity and relation extraction with domain-specific ontologies
- **Ontology Integration**: Pre-built schemas for research domains
- **Multi-Source Ingestion**: Pre-prints, Grant Proposals, Patent, Economic and Financial Data
- **Quality Assurance**: Automated entity linking, co-reference resolution, and fact verification

Key Benefits:

- Reduces KG construction time from months to days
- Ensures semantic consistency across heterogeneous data sources
- Maintains provenance and confidence scores for all extracted facts

📊 **Graph Optimization Engine: Signal Generation**

- **Resource Portfolio Optimization**: Integer programming over knowledge graphs for asset selection
- **Novel Hypothesis Generation:** The graph structure facilitates LLM-driven multi-hop reasoning to discover latent, non-obvious connections in the data, generating truly novel hypotheses that link previously isolated concepts.
- **Hypothesis Validity/Risk Modeling: Performs** Multi-hop Dependency Analysis and Graph Neural Network (GNN) inference to assess the systemic validity of a hypothesis, identifying claims that are structurally unsound or rely on known biased literature.

Classification of Digital Twins: Evolution of the Backbone

Traditional Twins

- **Traditional Data Stacks** (Relational Databases, Data Warehouses, Time Series DBs).
- Runs **Vanilla Machine Learning (ML)** models.

Agentic Twins

- **Data Org** is complemented by **Semantic Similarity** (often using Vector Databases and RAG indexes).
- Adds **Large Language Models (LLMs)** and coordinating agents to **ML**.

Graph-Enabled Agentic Twins (State-of-the-Art)

- **Knowledge Graph (KG)** as the unified, semantic, and structural backbone of the system.
- Employ **Agentic Optimization** running over the time series of the KGs, often utilizing Graph Neural Networks (GNNs) for inference and LLMs for high-level reasoning.

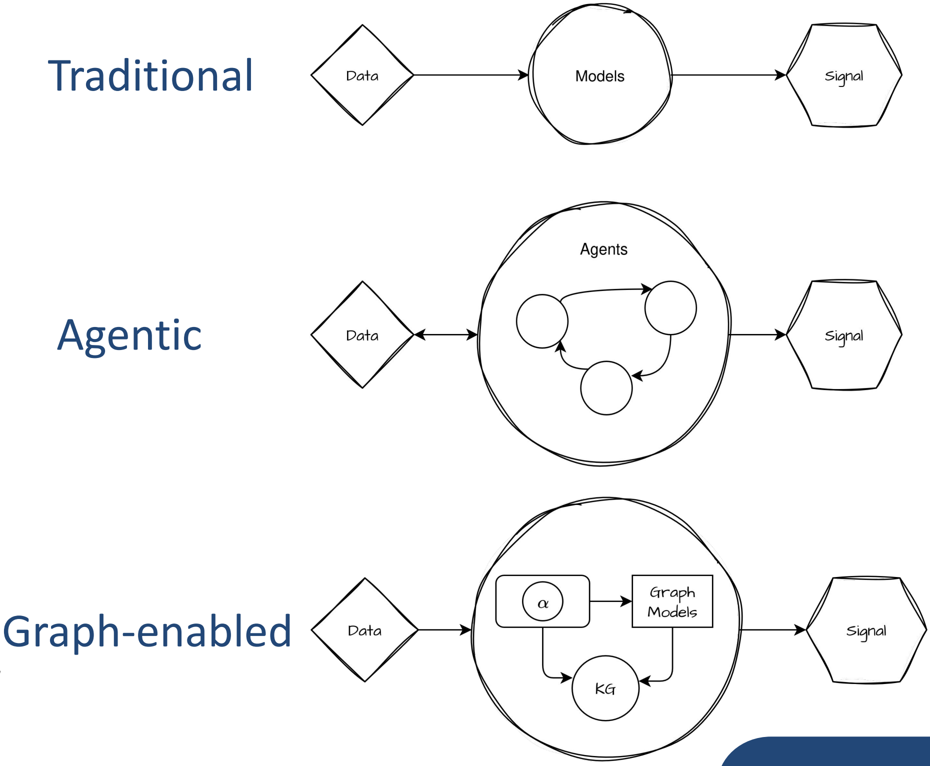

KG generation + Agentic Optimization

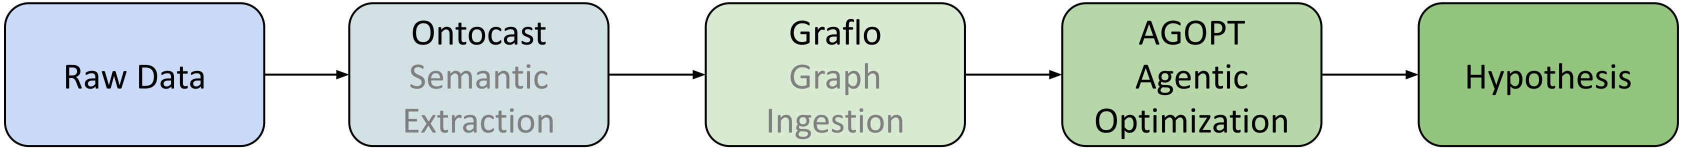

Ontocast extracts semantic triples from unstructured text
Graflo ingests these triples plus structured data into property graphs
AGOPT runs GNNs and agentic optimization over the unified knowledge graph [Symbolic AI]
KG provides content; LLM provides shape

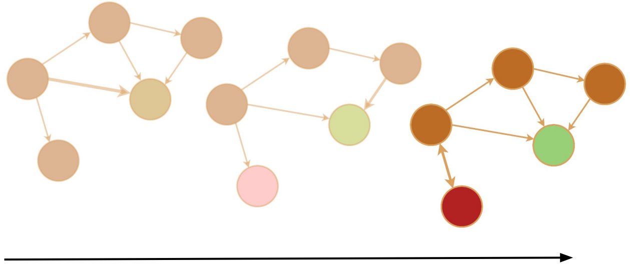

Conclusion

The Strategic Imperative. Research institutions face an unprecedented opportunity: while GenAI adoption accelerates globally, the gap between individual productivity gains and institutional transformation remains wide.
Digital Twins powered by Knowledge Graphs offer a systematic path to bridge this gap.

What We've Demonstrated

- Multi-dimensional optimization is not optional - single metrics create strategic blindness
- Knowledge Graphs uniquely enable transversal metrics across organizational silos
- Proven pilots (XSI, Emberloom, ReviewerRec) show measurable impact in real research workflows
- Graph-enabled Agentic Twins represent the state-of-the-art for strategic foresight

The Path Forward: Success requires three pillars.

1. **Start Small**: Identify high-value, well-defined use cases (reviewer discovery, funding optimization)
2. **Build Composably**: Each pilot must integrate into a unified, semantic infrastructure
3. **Optimize Portfolios**: Move from maximizing single KPIs to optimizing across your institution's true performance space

The Choice.
The Digital Twin is not about replacing human judgment - it's about augmenting strategic vision with computational foresight.