A Dual Academic–Industry Perspective on Enterprise AI

AI for Science, Cornell University, 2026

Serdar Kadıoğlu

Chief AI Officer, Northwestern Mutual

Adj. Assoc. Professor, Dept. of Computer Science, Brown

skadio.github.io

Navigating the Career Landscape

🎓 Full-Time Academia

Research is primary mandate. The incentive structure is clear: publication is how careers advance. Teaching and service are present, but **scholarly** output defines identity.

🔬 Research Labs

Fundamental and/or applied and timeboxed research. **Closer incentives to academia** but shipping matters too. The pressure to **demonstrate** near-term value is real.

🏭 Industry (Non-Research)

Expected to **ship products for customers.**
Research is not on the roadmap.
You must create your own mandate.

🔁 The Dual Hybrid Path

Bridging both. Hard but uniquely rewarding and increasingly common. **Requires active, and sustained effort**, no guaranteed payoff.

Across Two Worlds

🎓 Academia

Brown University — Adjunct Associate Professor in CS

Teaching, mentoring, service and publishing at the frontier.

Collaborations: CMU, USC, AI community, etc.

🏢 Industry

Northwestern Mutual — Chief AI Officer

Fidelity Investments — AI Center
Oracle – Adobe.

Partnerships: Amazon, NVIDIA, HBS, Sloan, McKinsey, etc.

This talk is about the **stories behind the dual career**, not about the results, but the **motivation**, **tension**, and successful **patterns**.

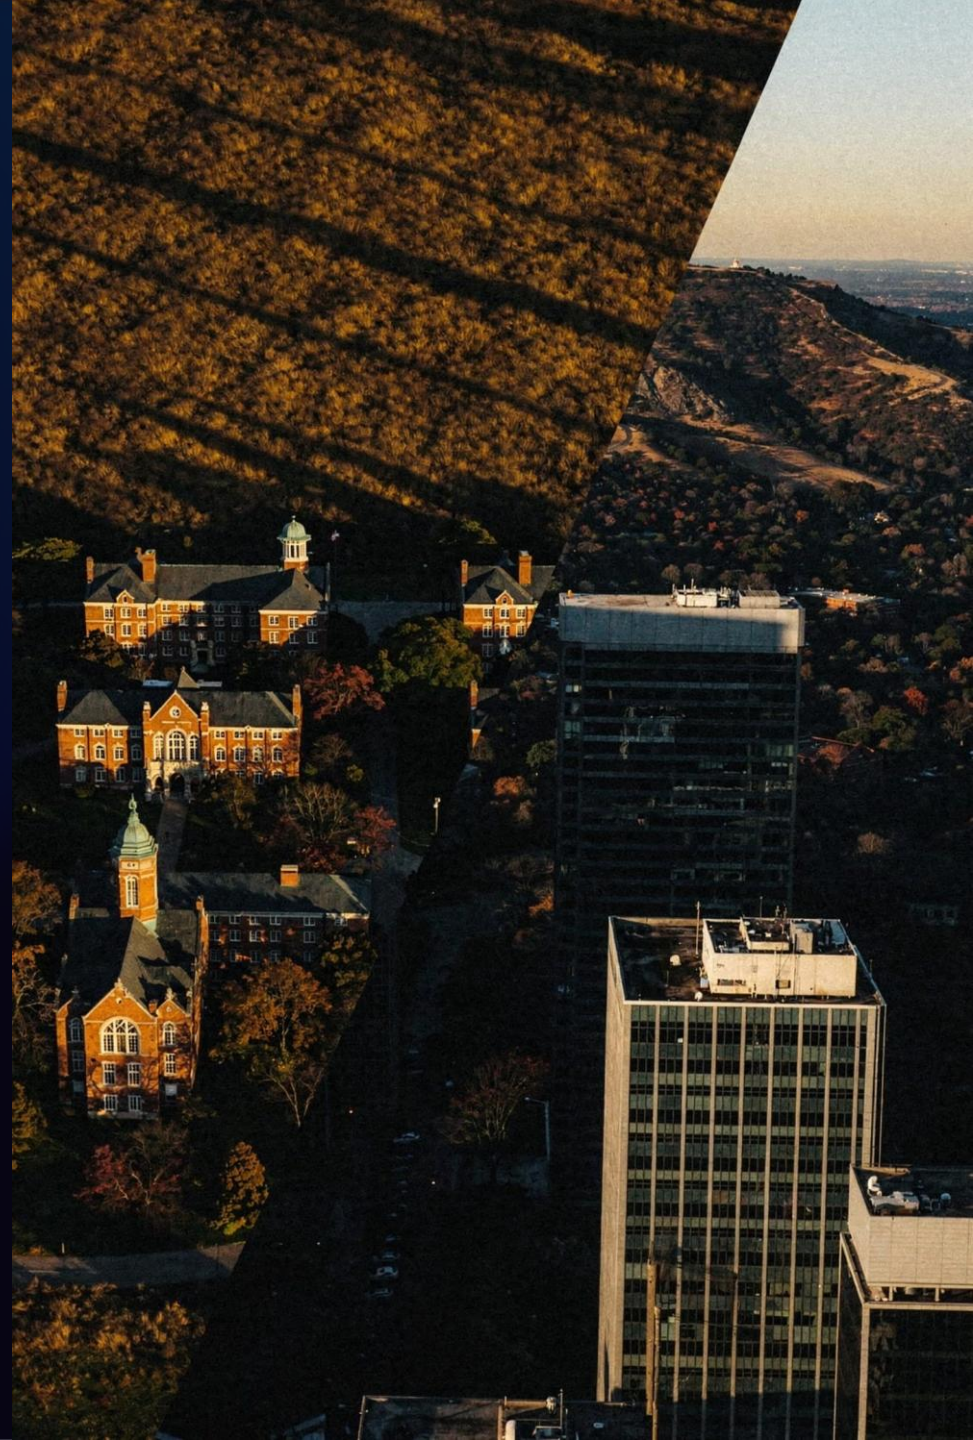

Career path is itself a constraint problem

The objectives evolves. The constraints shift. No closed-form solution exists. Search for your own path.

- **Be Intentional & choose problems that matter**
- **Translate ideas into systems**
- **Contribute back**
- **Build people around you**

Being Intentional in a Non-Research Role

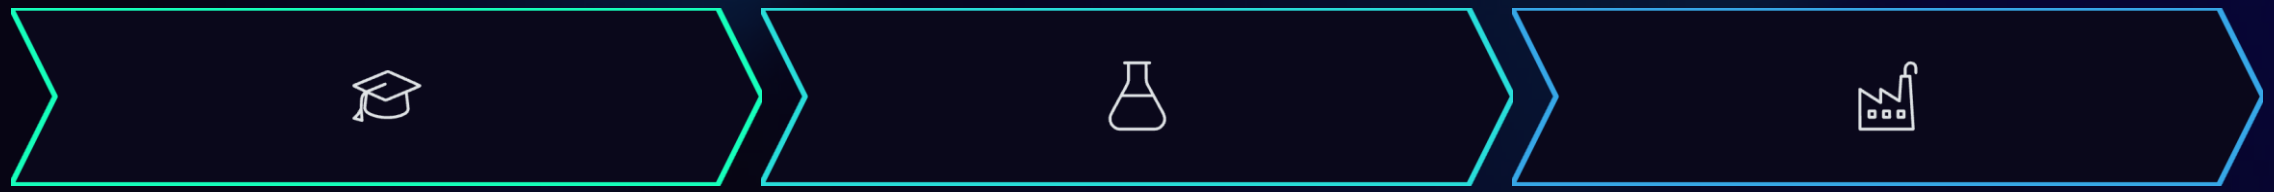

Academia

Research is the primary mandate. **Incentive** structure is clear.

Research Labs

Research is **expected**. Demonstratable value matters too.

Industry (Non-Research)

Expected to ship products. Research is **nobody's job**

Unless you make it yours.

The **hardest discipline** to sustain in any dual career:
carving out space for research when your **environment isn't designed for it**.

Research Without a Mandate Is the Hardest

> ***No one is going to ask you to do research****. The only way it works is* ***YOU*** *choose it, deliberately, repeatedly, and with* ***no guarantee of reward.***

In a non-research industry role, there is **no clear incentive to publish**. No protected time. No performance review credit. No guarantee your efforts will ever be relevant to the next product cycle. Small, **consistent efforts accumulate** but they do so slowly, with **no promise of return**.

- **Intentionally** choosing research when **it isn't measured** is itself a skill worth building deliberately. The payoff, if it comes, arrives slowly but the **good news is it compounds.**

Today's Journey

What We'll Cover

1

CP/SAT/OR Foundations

Why CP/SAT/OR is a uniquely positioned for durable intellectual fundamentals

2

Stories from the Field

Modularity, reproducibility, responsible AI, theory vs. practice, teaching, and partnership

3

Key Takeaways

Principles for building a career that compounds

THE FOUNDATION

Grateful to CP/SAT/OR: A Transferable Training

Training in CP/SAT/OR has been a foundation layer that has **transferred most broadly** across my career.

The CP curriculum offers **intellectual breadth** and requires speaking to **learning & reasoning** communities.

This trains you to move fluently across **theory, algorithms, and implementation** and these skills compound throughout the years.

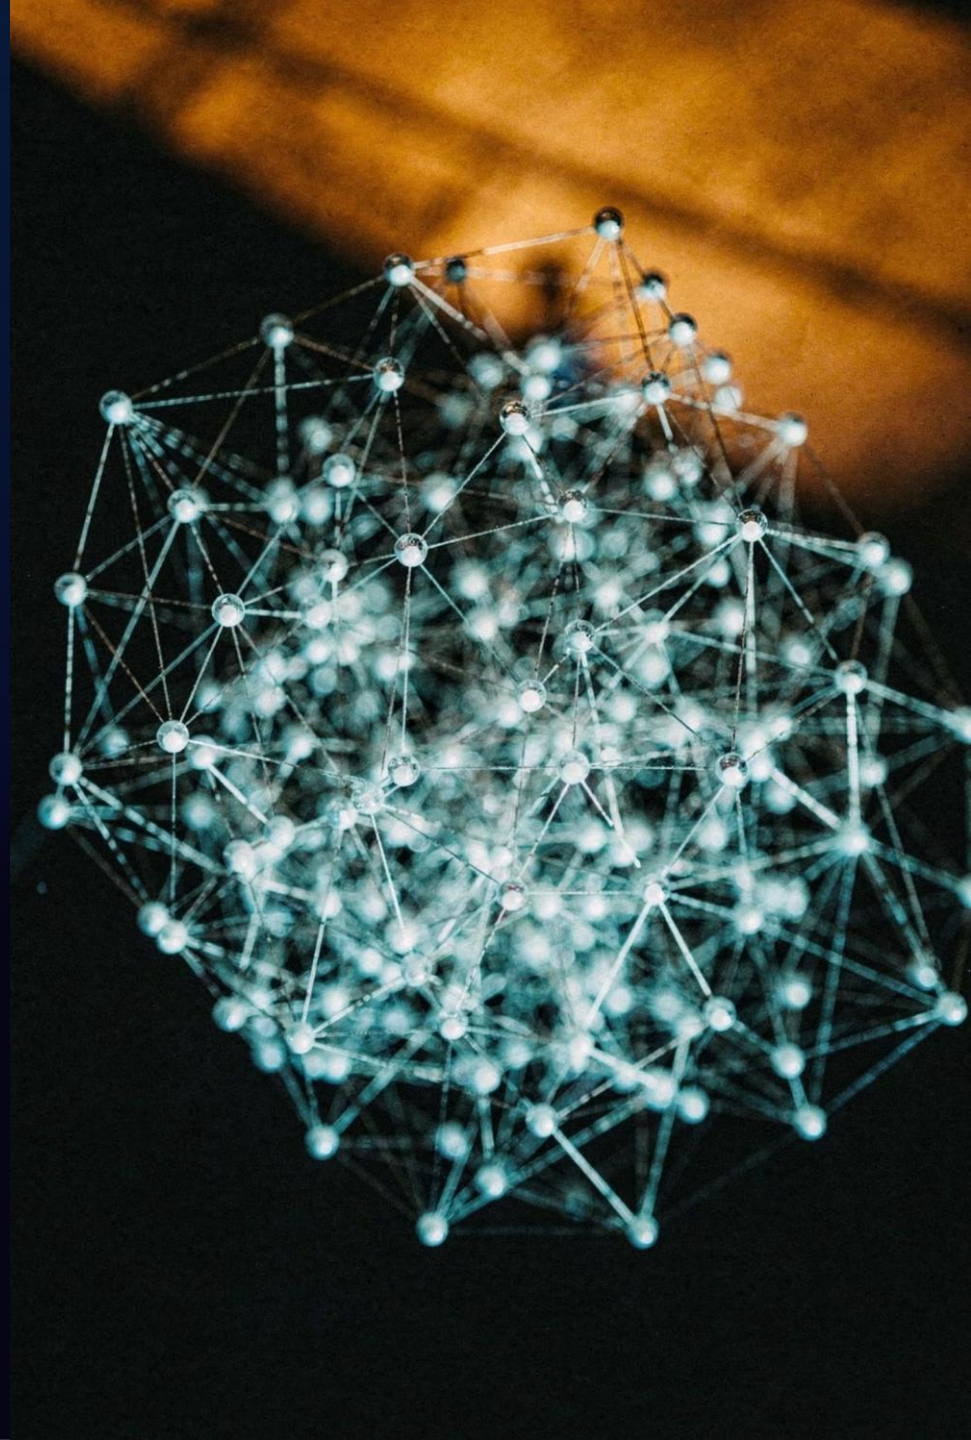

What CP/SAT/OR Gave Me

Foundation across domains and transfers in ways you won't fully appreciate until years later.

Formal Languages (NLP)

Grammar constraints, Chomsky hierarchy, automata, CYK parsing and NLP from first principles, invaluable years later.

Inference (Theory)

Graph-theoretic invariants and inference methods that generalize far beyond constraint propagation.

Learning-Guided Portfolios (ML)

Learning-guided solvers, portfolios, train/test protocols became Data Science, AutoML, and today's recursive self-improvement loops.

Solver Implementation (Software Craft)

Runtime optimization, amortized complexity, incremental design. These are best practices no textbook fully captures.

Things I Wish I Had Learned Earlier

Data-Driven AI & Unstructured Data

Strong model-driven intuitions and abstraction.

Modern systems learn from raw data representations. Unstructured data curation, labeling pipelines, and bias.

Uncertainty & Statistical Thinking

Deterministic assumptions. stochastic real deployments

Intuitions for uncertainty, variance, and distribution shift are essential for production AI.

Distributed Training & Large-Scale Compute

GPU distributed training is fundamentally different

Memory models, comm overhead, and debugging workflows are unlike classical programming.

Today's Journey

What We'll Cover

1

CP Foundations

Why Constraint Programming is a uniquely positioned for durable intellectual fundamentals

2

Stories from the Field

Modularity, reproducibility, responsible AI, theory vs. practice, teaching, and partnership

3

Key Takeaways

Principles for building a career that compounds

Build
modular software
for others to use,
including *future you.*

The Power of Modularity & Compounding Value

Composable by Design

Each module is technically deep on its own. Combined, they produce value greater than the sum of their parts.

Improvement Propagates

When one component improves, the entire pipeline benefits. Progress never stops.

Independent Yet Collaborative

Modules can be reviewed, updated, and handed off independently enables cross-team contribution without tight coupling.

Go the Last Mile

Packaging for others forces rigorous API design, testing, and deployment. Discipline makes your future work dramatically faster.

Recurring Success Pattern for the Last Mile

→ Internal Tool & Shared Development

→ Production Platform & Experiment Accelerator

→ Research Accelerator & Community Resource

→ Community Collaboration & Future Research

Building Higher-Order Abstractions from the Components of Recommender Systems

Serdar Kadıoğlu[1, 2], **Bernard Kleynhans**[1]
[1] AI Center of Excellence, Fidelity Investments, Boston, USA
[2] Department of Computer Science, Brown University, Providence, USA
serdar.kadioglu@fmr.com, bernard.kleynhans@fmr.com

Mab2Rec (AAAI'24)

pip install mab2rec

BALANS: Multi-Armed Bandits-based Adaptive Large Neighborhood Search for Mixed-Integer Programming Problems

Junyang Cai[1], **Serdar Kadıoğlu**[2,3], **Bistra Dilkina**[1]
[1]Department of Computer Science, University of Southern California
[2]AI Center of Excellence, Fidelity Investments
[3]Department of Computer Science, Brown University
{caijunya, dilkina}@usc.edu, serdark@cs.brown.edu

Balans (IJCAI'25)

pip install balans

Mab2Rec: Recommender System (AAAI'24)

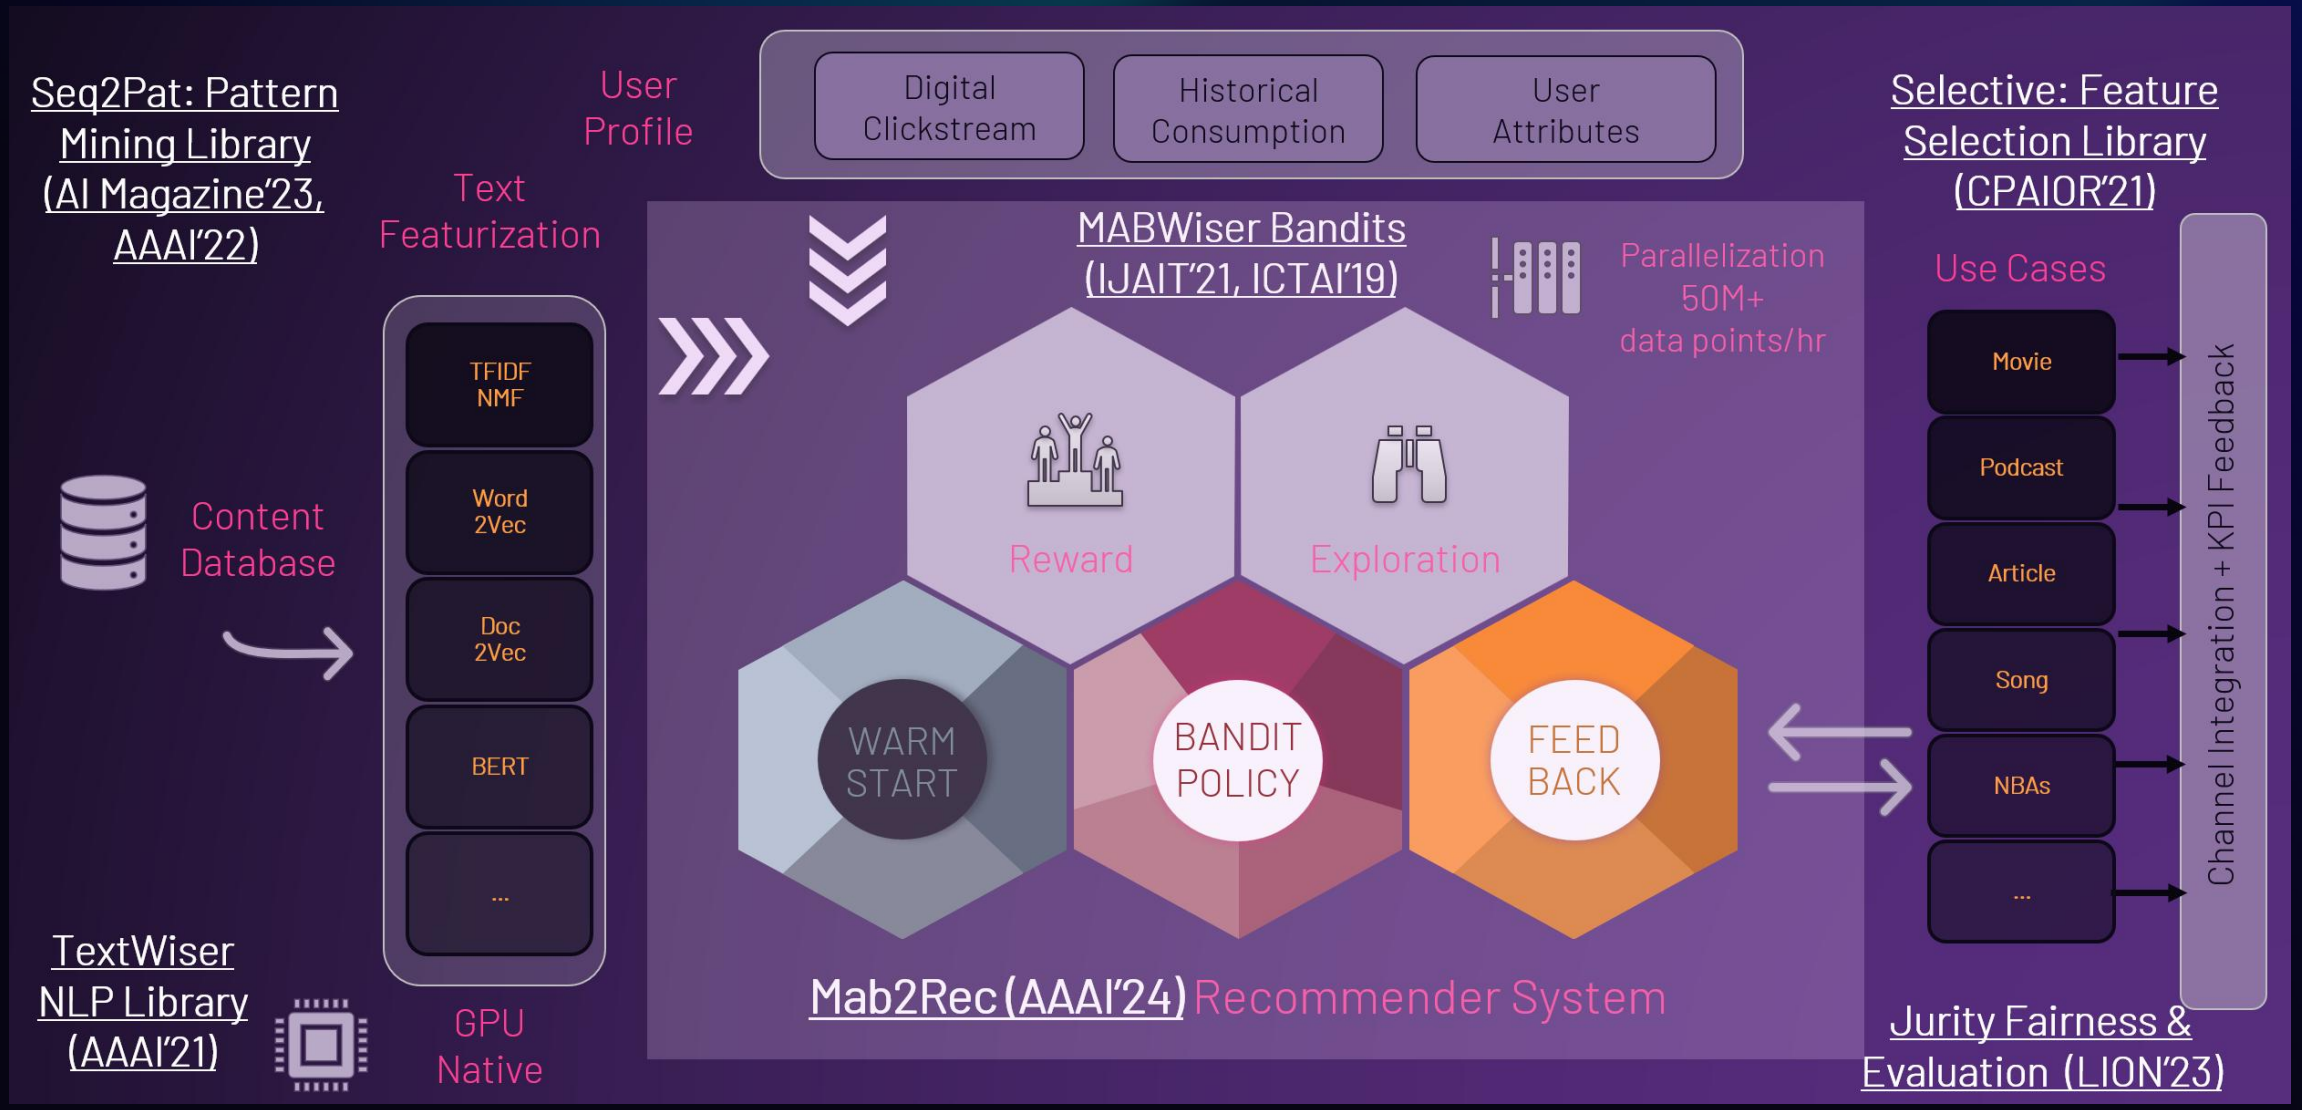

https://fidelity.github.io/mab2rec

Balans: Online Meta-Solver for MIPs (IJCAI'25)

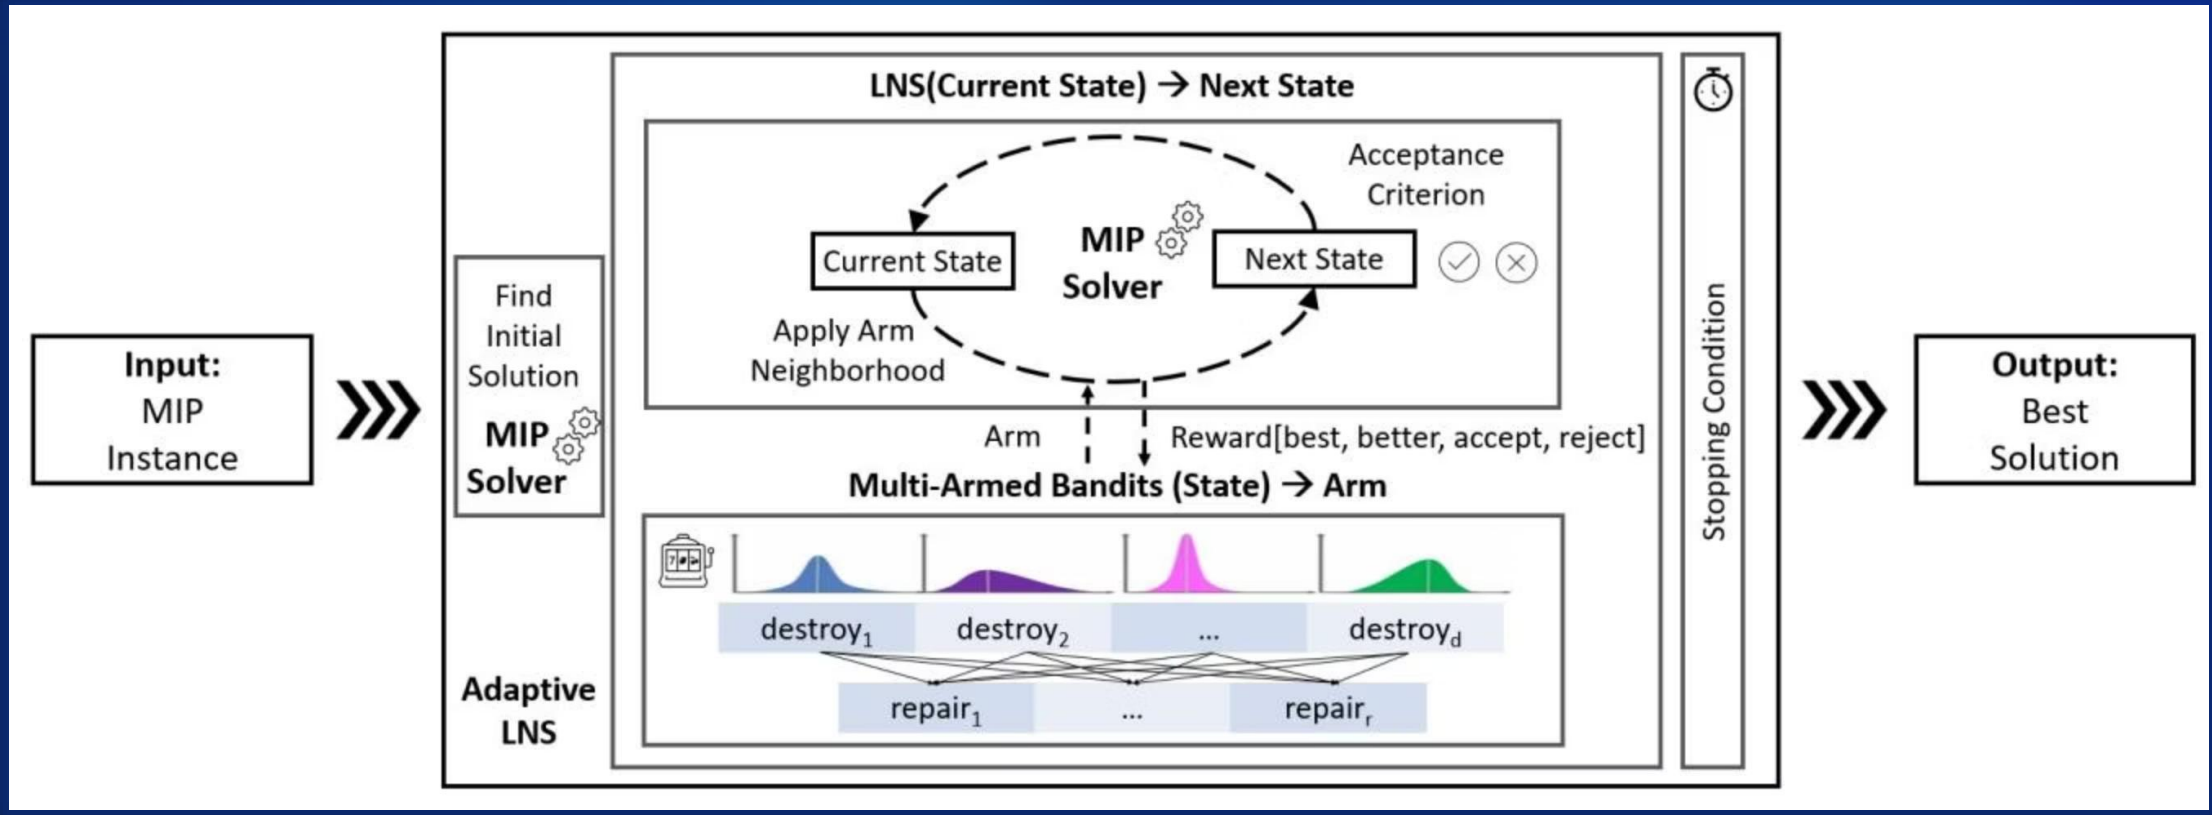

https://coin-or.github.io/balans

REPRODUCIBILITY

The Bug That Took a Year

Non-Deterministic Behavior of Thompson Sampling with Linear Payoffs and How to Avoid It

Doruk Kilitçioğlu *doruk.kilitcioglu@fmr.com*
Fidelity Investments
245 Summer St
Boston MA 02210

Serdar Kadıoğlu *serdar.kadioglu@fmr.com*
Fidelity Investments
245 Summer St
Boston MA 02210

Transactions of Machine Learning Research (TMLR'22)

The Bug That Took a Year

The Anomaly

Unit test failed on one machine only.

The Discovery

Found decade-old algorithm non-determinism bug.

The Fix

Traced inconsistency deep in a dependency.

TMLR Certificate

Proved fix, published reference, earned reproducibility certificate.

A classical recommendation algorithm, deployed for over a decade and used in the industry, made **arbitrarily different recommendations** to the same input.

The Academic Reaction
and why it misses the point ...

> *"But do the average results still hold?"*

Yes, sometimes, but not always, hence incorrect claims in papers.

But that entirely misses the point. A non-deterministic algorithm makes **arbitrary recommendations** to different **people**, or to the same person at different times.

In a regulated financial environment, this is a a **compliance and fairness risk**.
Type of work you do when no one is watching *because it is the right thing to do.*

RESPONSIBLE AI

When Fairness Metrics Meet Real-World

ACM TELO'25

RESEARCH-ARTICLE | OPEN ACCESS

Surrogate Modeling to Address the Absence of Protected Membership Attributes in Fairness Evaluation

Authors: Serdar Kadioğlu, Melinda Thielbar | Authors Info & Claims

ACM Transactions on Evolutionary Learning, Volume 5, Issue 3 • Article No.: 22, Pages 1 - 25 • https://doi.org/10.1145/3700145

Published: 29 August 2025 Publication History

pip install jurity

https://fidelity.github.io/mab2rec/#jurity

When Fairness Metrics Meet the Real World

The Academic Assumption

Ground truth labels are known. Protected attributes are observable. Fairness can be measured directly against a known target distribution.

The Enterprise Reality

Ground truth is rarely available. Protected attributes **cannot be legally collected**. Some academic proposals even suggest predicting these, which is legally fraught and problematic.

Fairness Platform

Developed practical, assumption-light fairness methods. Deployed a **fairness platform** bridging research and operations

RESPONSIBLE AI

Explainable AI

AAAI'2025

BoolXAI: Explainable AI Using Expressive Boolean Formulas

Serdar Kadıoğlu[1,2], Elton Yechao Zhu[3], Gili Rosenberg [4], John Kyle Brubaker[4], Martin J. A. Schuetz[4,5], Grant Salton [4,5,6], Zhihuai Zhu [4], Helmut G. Katzgraber [4]

[1] AI Center of Excellence, Fidelity Investments, MA, USA
[2] Department of Computer Science, Brown University, RI, USA
[3] Fidelity Center for Applied Technology, FMR LLC, MA, USA
[4] Amazon Quantum Solutions Lab, WA, USA
[5] AWS Center for Quantum Computing, CA, USA
[6] Institute for Quantum Information and Matter, California Institute of Technology, CA, USA
{serdark@cs.brown.edu}

pip install boolxai

https://fidelity.github.io/mab2rec/#boolxai

From Lunch Conversations to BoolXAI

Internal Seminar

Small reading groups at Fidelity AI Center exploring XAI formulations. Open culture of sharing.

Collaboration with Amazon

BoolXAI paper published at AAAI 2025.

Cross-Business Connection

Lunches spark explainability ideas. Informal lunches with Fidelity's Quantum Computing group sparked new problem framings around explainable AI.

BoolXLLM

Interactive LLM platform guided by stakeholders. Usage patterns drove new demand. Business stakeholders driving research direction.

PRACTICALITY

Learning-Based Methods Face Practical Limitations

Heavy Training Dependency

Training is computationally costly and depends on **carefully curating datasets** with desired properties and distributions.

Limited Generalization

Adapting learning-based methods to **new distributions and domains** remains a significant challenge in the field.

Solver Dependency Paradox

Ironically, **training depends on optimization solvers** to create labeled datasets, defeating the purpose of improving solving for hard instances.

REPRESENTATION LEARNING FOR COMBINATORIAL OPTIMIZATION

Forge: Foundational Optimization Representations from Graph Embeddings

Representational learning for combinatorial optimization. Generate embeddings from MIP instances, pre-train models, and fine-tune for downstream tasks uch as predicting integral gap, search guidance, backdoor prediction, and solver configuration.

Serdar Kadıoğlu[1,2] Zohair Shafi[2,3]

[1] Department of Computer Science, Brown University [2] AI Center of Excellence, Fidelity Investments [3] Northeastern University

Paper | Code | Slides | Models | Dataset

https://skadio.github.io/forge

LLM Copilots for Text-to-Model Translation

A suite of LLM modeling copilots, datasets, fined-tuned models, demos, interactive editor, and online leaderboard for translating natural language text into formal combinatorial constraint models.

Serdar Kadıoğlu[1,2] Karthik Uppuluri[2]

[1] Department of Computer Science, Brown University [2] AI Center of Excellence, Fidelity Investments

Paper | Copilots | Slides | Leaderboard | Interactive Editor | Collections

Tutorial on LLMs for Constraint Modeling	IJCAI'26	Workshop on LLM-Solve	CP'26	Master Class on LLMs for CP/OR	CPAIOR'26
Tutorial on LLMs for Optimization	AAAI'26	Workshop on LLM-Solve	CP'25	Workshop on LLMs for OR	AAAI'24

pip install text2model

https://skadio.github.io/text2model

TEACHING

A Late-Night Call and a Lecture

I met a colleague in New York, and we talked about his new teaching assignment. Six months later:

> *"Can you please prepare a lecture on Branch-and-Bound solvers for tomorrow?"*

I spent the night building slides.

Framed the topic differently than a typical textbook: *"**Here's the theory and here's what really happens when you build the system.**"*

Intentional, over the years paying attention to that gap and open problems live.

I could deliver that lecture because I've kept a foot in community and **took notes**.

Some of the best collaborations **happen informally**.

Partnership Mentality

You achieve larger impact as part of a bigger team.

Building the people around you is a soft skill but also a **career and research strategy**.

Build the People Around You & Institutionalize

Internal Champions

Find internal advocates at academia, industry. Partnerships need champions on both sides. Many of my co-authors were publishing for the very first time.

Research Seminars

Recurring forums at Oracle and Fidelity's AI Center. Consistent, low-barrier sessions compound over years into strong intellectual culture.

Academic Bridges

Non-profits, co-ops, internships, and university labs create the neutral ground where collaborations begin.

External Partnerships

The CP Doctoral Programme is exactly this kind of high-leverage partnership.

Main Takeaways

Be Intentional

Protect time for research regardless of role. Small, **deliberate efforts compound**. No guarantees, but no shortcuts either.

Choose Problems That Matter

Edge of theoretically interesting and practically relevant. Enterprise scale surfaces them but **you have to be paying attention.**

Modularity & Rigor & Last Mile

Build deep, composable work. Go the last mile. **Strong theory + serious implementation + principled experimentation**. Details that look like engineering are often the research.

Openness & Community

Contribute back, open source, help next generation and institutionalize. Grow people around you. The community you invest in will invest back in ways you cannot predict.

Strategic Pillars of Enterprise AI

AI Learning from

Robust, scalable, reproduc
from structured, unstruct
semi-structured datasets

Selective, TextWiser,

AI for Automated

Extraction and translation
tasks and intents for huma

Text2Model, Learn2Zin

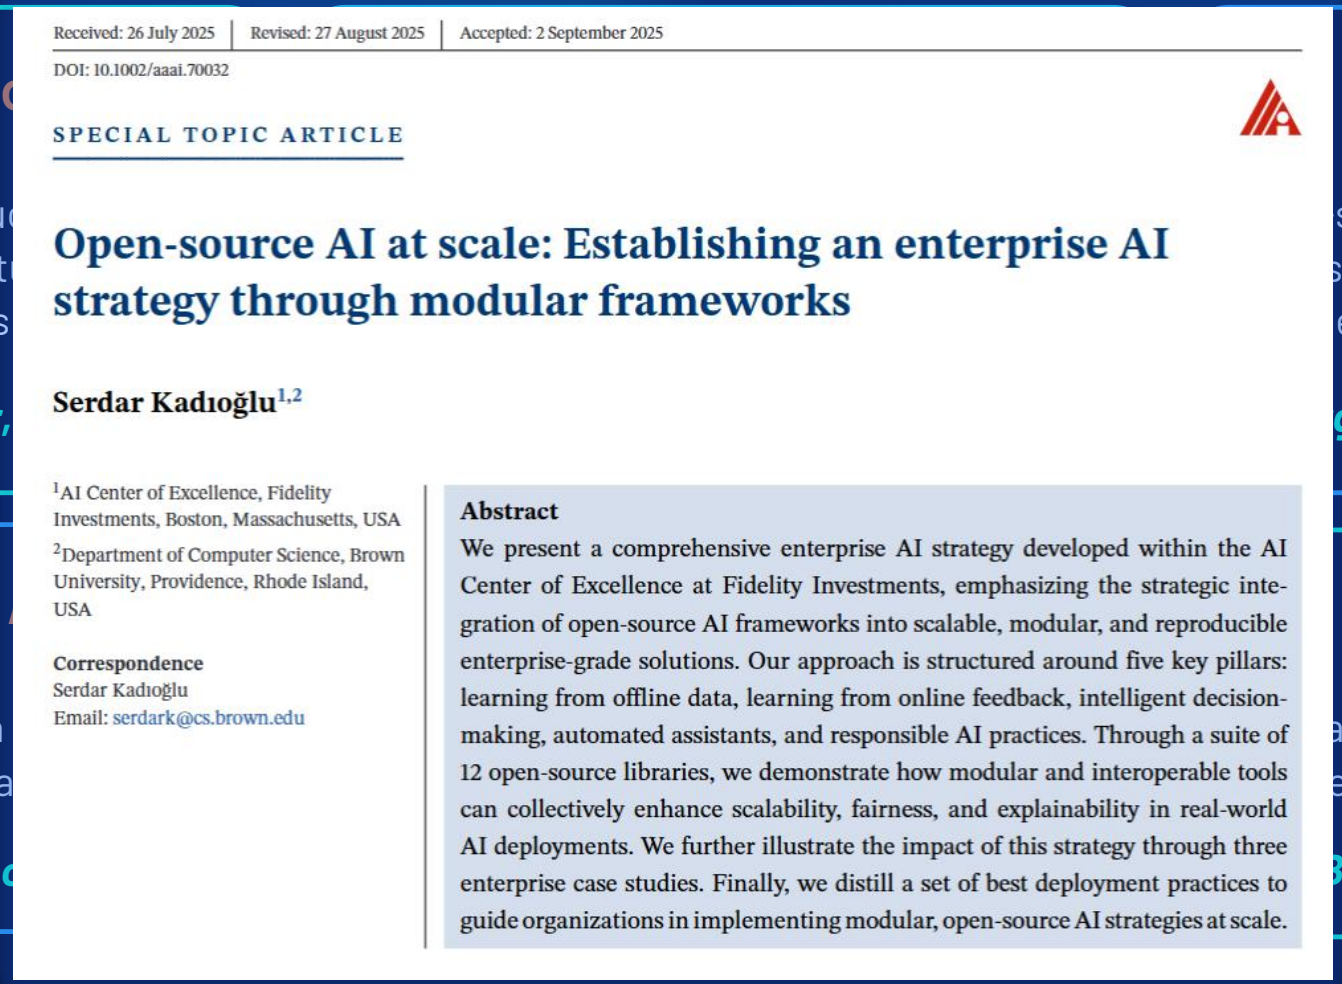

Received: 26 July 2025 | Revised: 27 August 2025 | Accepted: 2 September 2025

DOI: 10.1002/aaai.70032

SPECIAL TOPIC ARTICLE

Open-source AI at scale: Establishing an enterprise AI strategy through modular frameworks

Serdar Kadıoğlu[1,2]

[1]AI Center of Excellence, Fidelity Investments, Boston, Massachusetts, USA

[2]Department of Computer Science, Brown University, Providence, Rhode Island, USA

Correspondence
Serdar Kadıoğlu
Email: serdark@cs.brown.edu

Abstract

We present a comprehensive enterprise AI strategy developed within the AI Center of Excellence at Fidelity Investments, emphasizing the strategic integration of open-source AI frameworks into scalable, modular, and reproducible enterprise-grade solutions. Our approach is structured around five key pillars: learning from offline data, learning from online feedback, intelligent decision-making, automated assistants, and responsible AI practices. Through a suite of 12 open-source libraries, we demonstrate how modular and interoperable tools can collectively enhance scalability, fairness, and explainability in real-world AI deployments. We further illustrate the impact of this strategy through three enterprise case studies. Finally, we distill a set of best deployment practices to guide organizations in implementing modular, open-source AI strategies at scale.

AI for Decision Making

scale, integrated, (meta)
s for resource
ement and optimization.

ge, Balans, PathFinder

ability, evaluation, fairness,
ems.

BoolXLMM

Open-Source AI at Scale: Establishing an Enterprise AI Strategy [AI Magazine'25]

skadio.github.io