

Machine Learning IC Leader with 10+ post-PhD years defining ML strategy and delivering large-scale ML systems at Meta and Thumbtack, generating over \$100M in realized incremental revenue and serving millions of users daily. Deep expertise in Search, Ranking, and Personalization (Recommender Systems, Ads), with a track record of architecting complex systems using Multi-Task Learning, fine-tuned LLMs, and multi-shot prompting. Mentor of tens of ICs with publications in top applied ML and research venues (KDD, WWW, WSDM, CIKM, MS, ISR, MISQ).

#### RECENT KEY ACHIEVEMENTS

- ◇ **ML Strategy, Ranking (Thumbtack):** Defined ML strategy and architected the first multi-task ranker for the core funnel, driving a **2.1% incremental revenue** lift through a **2.5% increase in created projects**.
- ◇ **Generative AI Innovation (Thumbtack):** Pioneered and drove the standardization of Generative AI workflows, resulting in an **Innovation Award**, **two research papers**, and a **patent filing** for Thumbtack's first multimodal search engine.
- ◇ **Strategic ML Leadership (Meta):** Defined the ML roadmap and architected the solution for automating the deployment of launch-risk models, generating **\$94M** in incremental revenue.

#### 2024 - Current *Staff (L7) Applied Scientist* at **Thumbtack** (*Search and Ranking | Fulfillment*)

- ◇ Architected and deployed Thumbtack's first multi-task ranker, delivering an A/B tested **2.1%** increase in company revenue driven by a **2.5%** increase in projects created and a **9.8%** increase in hires. The complex ranking architecture optimizes for both engagement and satisfaction through Multi-Task Learning with LLM-based feature generation.
- ◇ Architected and deployed Thumbtack's first **multimodal search engine** (image, voice, natural language), earning the company's **Innovation Award** and resulting in both a **patent filing** and a **publication**. Multi-market A/B tests showed **28% lift in customer trust** and a **13% lift in project creation**. System integrates conversational prompting with fine-tuned LLMs using supervised (SFT) and preference-optimization (DPO) methods, classification models, and LLMs for feature generation and judging.
- ◇ Led the end-to-end design and deployment of Thumbtack's first **AI-powered match explanations** for both customers and professionals. Early market tests show a **+2% lift** in pro lead-acceptance rates and a **+15% lift** in customer service page views. System leverages prompting with lightweight RAG and LLM-as-a-judge evaluation.
- ◇ Pioneered and drove the strategic adoption of Generative AI across Thumbtack's hiring experience org, establishing highly cross-functional AI workflows and evaluation frameworks (LLM-as-a-Judge, Red Teaming, Human Evaluations) to define internal production standards for safety, quality, and alignment. Work resulted in a peer-reviewed **CIKM publication**.
- ◇ Drove the ML strategy and architectural design of the new dynamic pro-preferences system that infers the types of tasks pros intend to do, recommends new tasks, and increases transparency and pro control. Early iterations are live in experimental markets and supported a **patent filing**.
- ◇ **IC Leadership & XFN Strategy:** Drove ML vision and roadmap for the hiring experience organization. Provided technical leadership across Product, Engineering, and Design and **mentored 10+ applied scientists and engineers** on ML and GenAI systems.

#### 2022 - 2024 *Staff (IC6) Applied Scientist Tech Lead* at **Meta** (*Ads Infra and Ranking*)

- ◇ Architected and deployed an Ads Launch-Risk ML Framework to address bottlenecks in the model launch pipeline by detecting launch anomalies using supervised and unsupervised learning. The system safely automated low-risk launches and generated **\$94M in incremental revenue** in 2023 H2.
- ◇ Architected and drove cross-functional alignment across several infra and measurement teams to build a ROI framework for Ad's reliability systems. The framework prioritized 25+ systems and guided investment decisions, with projected impact of \$400M (2023–2026).
- ◇ Designed a statistical quality-assurance framework for automated Ads model launches, using stratified sampling and counterfactual impact estimation to enforce model rejection boundaries and prevent \$10M in monthly losses.

- ◇ **IC Leadership & XFN Strategy:** Provided VP-level strategic trade-off analysis using causal inference to guide Ads infra decisions. Mentored junior ICs on opportunity scoping, roadmapping, modeling, and experimentation, raising the technical bar across the organization.

#### 2015 - 2022 *Research Professor of Applied Machine Learning* at **Boston College**

- ◇ Led applied ML research in recommender systems, with publications at KDD, WWW, Management Science, and ISR.
- ◇ Designed and evaluated ML frameworks that achieved 20–40% predictive improvements over online labor-market baselines, projecting up to a **6% market-revenue increase** and 22% wage improvement for platform participants.
- ◇ Taught graduate-level applied machine learning and programming courses and advised 10+ theses, demonstrating deep experience in technical mentorship.
- ◇ Organized international workshops and conferences and collaborated with industry partners on Applied ML and Causal Inference.

Prior to 2015: *Machine Learning Engineer (intern and later full-time)* at **Upwork** and *Research Scientist (intern)* at **Microsoft Research**.

---

#### TOPICS AND METHODOLOGIES

- ◇ Predictive modeling, supervised and unsupervised learning, ranking systems, recommender systems, deep learning, reinforcement learning, sequence models, NLP, LLMs, prompt engineering, fine-tuning, RAG.
- ◇ Experimentation and Causal Inference, Difference-in-Differences, Panel Data, Instrumental Variables, Propensity Score Matching, survival analysis, counterfactual estimation.

#### PROGRAMMING LANGUAGES AND ML PRODUCTION TOOLS

- ◇ Python (Pandas, NumPy, Scikit-learn, PyTorch, Transformers/HF), R, SQL, Stata, Shell Scripting, Git. (*Prior to 2016: Java, C#, C++, C.*)
- ◇ ML and Experimentation Infrastructure Meta: Dataswarm, FB Learner, Manifold (HDFS), Deltoid3, Presto, Hive, Unidash.
- ◇ ML and Experimentation Infrastructure at Thumbtack: GCP, BigQuery, Databricks, MLflow, Airflow, DBT, AWS (SageMaker), Docker, Statsig.

---

#### EDUCATION

- ◇ **New York University.** *PhD* in Applied Machine Learning & Causal Inference (Information Systems)
- ◇ National Technical University of Athens. BS in Computer Engineering (5-year program)

---

#### SELECTED HONORS AND AWARDS

- ◇ Thumbtack Innovation Award
- ◇ INFORMS ISS Gordon B. Davis **Young Scholar Award** and ISS Nunamaker-Chen **Dissertation Award**
- ◇ INFORMS Data Mining **Best Paper Award**
- ◇ Represented Cyprus in the **International Physics Olympiad**
- ◇ Represented Cyprus in the **International Biology Olympiad**
- ◇ Several awards and honorary mentions in **Annual Mathematics Cyprus Olympiads**

---

#### SAMPLE OF RECENT APPLIED MACHINE LEARNING PUBLICATIONS – (GOOGLE SCHOLAR PROFILE.)

- ◇ **Kokkodis Marios**, Mahendran Purusoth and Boatwright Grace. 2025. Taming the Unicorn: Turning Generative AI Into a Workhorse. Conference on Information and Knowledge Management *CIKM*. (link to paper)
- ◇ **Kokkodis Marios**, Demsyn-Jones Richard and Raghavan Vijay. 2025. Beyond the Hype: Embeddings vs. Prompting for Multiclass Classification Tasks. (link to paper)
- ◇ **Kokkodis Marios** and Ipeirotis G. Panagiotis. 2023. The Good, the Bad, and the Unhirable: Recommending Job Applicants in Online Labor Markets. *Management Science*. (link to paper)
- ◇ **Kokkodis Marios** and Ipeirotis G. Panagiotis. 2021. Demand-Aware Career Path Recommendations: A Reinforcement Learning Approach. *Management Science*. (link to paper)