Kurtis Voris

Lead Applied Scientist · Personalization for Agentic Systems

User State & Action Models · Sequential Recommendation · Neural Search · Agent Evaluation

kurtisvoris@gmail.com · linkedin.com/in/kurtisvoris · Northern California

Lead applied scientist building personalization and evaluation systems for search, recommendation, and agentic experiences at Zillow and Amazon. Research and production work centers on sequential user state and action models, LLM-based user simulation, and causal and statistical evaluation. Currently leading Zillow's foundational user state action model for agentic experiences, from representation design and training objectives to offline evaluation and production deployment.

CORE EXPERTISE

Sequential User Modeling: Transformer architectures over multimodal interaction event streams for user state, intent inference, and next-action prediction.

Causal & Statistical Evaluation: Calibrated propensities, bias correction, causal inference, robust offline metrics, and efficient online experimentation for personalization.

LLM Fine-Tuning for User Simulation: Cost-efficient supervised fine-tuning on real and synthetic data to enable realistic offline evaluation and end-to-end agentic optimization.

Agentic Personalization: User representations and next-best-action policies that leverage user state to optimize long-term outcomes in conversational and proactive agent experiences.

EXPERIENCE

Zillow

Jan 2025 – Present

Applied Scientist, Agentic Foundations · Remote (Northern CA)

- **Foundational User State Action Model:** Proposed and led development of a transformer trained from scratch on multimodal home-shopping sequences, including search behavior, listing views, saves, and decision-relevant conversation. Defined tokenization, training objectives, and representation outputs. The model produces dense user embeddings and next-event, attribute, and long-horizon predictions that power intent modeling, personalized retrieval, and next-best-action in Zillow's agentic AI experiences. Blog.
- **Agent Simulation & Evaluation:** Fine-tuned LLMs to simulate realistic multi-turn shopper conversations; the simulator was adopted as the offline gate for new agent architectures and major prompt changes.

Amazon

Jan 2017 – Jan 2025

Applied Scientist, Alexa News — *Nov 2023 – Jan 2025*

- **Personalized News Recommendations:** Conceived and launched the Alexa News recommendation system in four months, progressing from proof of concept to production. Built the end-to-end stack, including content scoring, topic clustering, embedding retrieval, and diversity-aware ranking.
- **News Summarization:** Designed an LLM summarization pipeline with headline generation and custom evaluation metrics for numerical accuracy and claim validity.

Applied Scientist, FireTV Search — *Oct 2020 – Nov 2023*

- **Neural Vector Search:** Defined the technical strategy and led a 12-month initiative spanning model development, ANN retrieval, and SageMaker serving for plot- and quote-based voice search. Fine-tuned bi-encoders on multimodal catalog metadata and built an ANN index over the full video catalog.
- **Search Re-Ranker:** Developed a two-stage neural re-ranking system with training-data design and offline evaluation. Improved playback rate by 320 basis points with a pointwise model and by a further 150 basis points with pairwise ranking.

Research Scientist, Alexa Information — *Oct 2018 – Oct 2020*

- **Follow-up Question Recommendation:** Built a recommender using question co-occurrence, neural collaborative filtering, and BERT near-duplicate classification over sequential question flows.
- **Sports Interest Embeddings:** Trained word2vec on follow-action sequences to produce team embeddings for Sports Briefing and the interest picker.

Data Scientist, Amazon Studios · Consultant, AWS ProServe — *Jan 2017 – Oct 2018*

- **Content & Talent Models:** Built multimodal plot and actor models to predict show popularity and inform bid pricing; extracted entity-level review sentiment for creators.

HomeUnion

Feb 2015 – Jan 2017

Data Scientist

- **Automated Valuation & Forecasting:** Built gradient-boosted valuation models across 110 million U.S. homes, achieving 4.3% median rent error; published an ML home-price index methodology.

SynteractHCR

Jun 2012 – Feb 2015

Statistical Programmer & Biostatistician

EDUCATION

M.S. Statistics San Diego State University · 2015

B.S. Statistics California Polytechnic State University, San Luis Obispo · 2012

METHODS & TECHNOLOGY

PyTorch · Sequential Transformers · Multi-Task / Multi-Horizon Learning · Two-Tower Models · Learning-to-Rank · FAISS · Spark · Databricks · AWS/SageMaker · Docker · Python · SQL