Amin Shabani

Senior Machine Learning Researcher, RBC Borealis · Vancouver, Canada

aminshabaany@gmail.com shabani.me Google Scholar aminshabani aminshabani

Experience

Senior Machine Learning Researcher Jul 2026 – Present
RBC Borealis, Vancouver
Research for quantitative finance: models that stay accurate, auditable, and privacy-compliant under the data-governance constraints of a regulated bank.

- **What on-policy distillation learns** (ESC, arXiv 2026; code). Entropy-shape agreement predicts the best LLM teacher.

Machine Learning Researcher Apr 2024 – Jun 2026
RBC Borealis, Vancouver
LLM systems for finance, from prototype to production. Two lines of work, both published and open-sourced:

- **Agentic table summarization** (FACTS, ACL Findings 2026; code). Recast query-focused summarization as offline template generation: an agentic workflow synthesizes guiding questions, SQL, and Jinja2 templates, each validated by an LLM-council ensemble. Only schemas reach the LLM and reusable templates keep runtime flat in table size. 97% intent match, 94% SQL accuracy; best or second-best on every metric across FeTaQA, QTSumm, QFMTS.
- **Efficient context-aware reranking** (EBCAR, ICLR 2026; code). A 126M-parameter reranker over pre-computed passage embeddings; document-ID and position encodings with hybrid full/document-masked attention capture cross-document evidence and coreference. Best average nDCG@10 (64.92) on ConTEB at 29.3 queries/s on one A100, an order of magnitude faster than LLM listwise rerankers.
- **1st place**, Financial Agentic Retrieval Grand Challenge (FinAgentBench), AI for Finance Symposium 2025.

Research Scientist/Engineer Intern Jun 2023 – Nov 2023
Adobe Research, San Jose
Developed Visual Layout Composer (CVPR 2024): an image diffusion model and a vector diffusion model over element boxes denoise in lockstep through cross-domain attention, so element appearance informs geometry while the output stays a layered, editable design. State of the art on multiple layout benchmarks; filed as a US patent application.

Machine Learning Research Intern Jan 2022 – May 2022
Borealis AI (now RBC Borealis), Vancouver
Developed Scaleformer (ICLR 2023), a model-agnostic multi-scale framework for time-series forecasting: coarse-to-fine refinement with shared weights, cross-scale normalization against distribution shift, and an adaptive loss against error accumulation. 5.5–38.5% lower error on five transformer backbones at minimal overhead; filed as a US patent application.

Education

Ph.D. in Computing Science Sep 2019 – Aug 2025
Simon Fraser University, Canada · advisor Yasutaka Furukawa
Thesis: diffusion models for structured geometry — vector floorplans, spatial puzzles, design layouts. summit.sfu.ca/item/39603

M.S. in Electrical and Computer Engineering Sep 2017 – Aug 2019
Seoul National University, Korea · advisor Kyoung Mu Lee

B.S. in Computer Science Oct 2012 – Aug 2017
Sharif University of Technology, Iran

Publications

[1] Shiping Yang, Chuanqi Tang, Hossein Hajimirsadeghi, Layla El Asri, **Amin Shabani**, Angel X. Chang, *Distilling the Shape of Uncertainty: What On-Policy Distillation Learns and How to Predict It*. arXiv preprint 2026.

[2] Ye Yuan, **Amin Shabani**, Siqi Liu, *FACTS: Table Summarization via Offline Template Generation with Agentic Workflows*. ACL Findings 2026.
Reusable SQL + Jinja2 templates; schema-only LLM exposure; runtime flat in table size.

[3] Ye Yuan, **Amin Shabani**, Siqi Liu, *Embedding-based Context-aware Reranker*. ICLR 2026.
Embedding-space reranker with hybrid full/masked attention; SOTA nDCG@10 on ConTEB at 29 queries/s.

[4] Ye Yuan, **Amin Shabani**, Siqi Liu, *Fast, Accurate, and Privacy-Compliant Table Summarization via Offline Template Generation*. NeurIPS Workshop on Generative AI in Finance 2025.

[5] **Amin Shabani**, Zhaowen Wang, Difan Liu, Nanxuan Zhao, Jimei Yang, Yasutaka Furukawa, *Visual Layout Composer: Image-Vector Dual Diffusion Model for Design Layout Generation*. CVPR 2024.

Coupled image–vector diffusion; appearance-aware yet fully editable layouts.

[6] Sepidehsadat Hosseini, **Amin Shabani**, Saghar Irandoust, Yasutaka Furukawa, *PuzzleFusion: Unleashing the Power of Diffusion Models for Spatial Puzzle Solving*. NeurIPS 2023 (Spotlight).
Jigsaw and room-arrangement solving as diffusion over SE(2) piece transforms; new Voronoi-jigsaw and MagicPlan benchmarks.

[7] **Amin Shabani**, Sepidehsadat Hosseini, Yasutaka Furukawa, *HouseDiffusion: Vector Floorplan Generation via a Diffusion Model with Discrete and Continuous Denoising*. CVPR 2023.
Discrete/continuous denoising over polygonal loops with graph-conditioned attention; exact parallelism, orthogonality, corner-sharing.

[8] **Amin Shabani**, Amir Abdi, Lili Meng, Tristan Sylvain, *Scaleformer: Iterative Multi-scale Refining Transformers for Time Series Forecasting*. ICLR 2023.
Architecture-agnostic coarse-to-fine refinement; 5.5–38.5% error reduction over five transformer backbones.

[9] Weilian Song, Mahsa Maleki Abyaneh, **Amin Shabani**, Yasutaka Furukawa, *Vectorizing Building Blueprints*. ACCV 2022.

[10] **Amin Shabani**, Weilian Song, Makoto Odamaki, Hirochika Fujiki, Yasutaka Furukawa, *Extreme Structure from Motion for Indoor Panoramas without Visual Overlaps*. ICCV 2021.
Poses for non-overlapping indoor panoramas via a learned realism score over top-down layouts; sub-meter error on 47% of 286 houses.

[11] **Amin Shabani**, Laleh Samadfam, Mohammad Amin Sadeghi, *Local Visual Microphones: Improved Sound Extraction from Silent Video*. BMVC 2017 (Spotlight).
Aggregates local vibration phases for higher audio quality and sound direction; real-time, 2–3 orders of magnitude faster.

[12] M. Abdolmaleki, J. P. Hutchinson, S. Gh. Ilchi, E. S. Mahmoodian, N. Matsumoto, **Amin Shabani** (alphabetical), *On Uniquely k-List Colorable Planar Graphs, Graphs on Surfaces, and Regular Graphs*. Graphs and Combinatorics 2018.

Patents

- Parishad BehnamGhader, **Amin Shabani**, Jiawei He, Leonid Sigal, Siqi Liu. *Systems, apparatuses, methods, and non-transitory computer-readable storage devices for information retrieval*. US 19/429,732, Royal Bank of Canada, 2025.
- Zhaowen Wang, Nanxuan Zhao, Jimei Yang, Difan Liu, **Amin Shabani**. *Generating visually aware design layouts using a multi-domain diffusion neural network*. US 18/641,137, Adobe, 2024.
- **Amin Shabani**, Tristan Sylvain, Lili Meng, Amir Abdi. *Multi-scale artificial neural network and a method for operating same for time series forecasting*. US 18/197,197, Royal Bank of Canada, 2023.

Talks & Service

- **Talks & panels:** Panelist, GroundLM (EMNLP 2026 workshop), Budapest, Oct 2026 · Upper Bound 2026 (Amii), Edmonton, May 2026 · Samsung AI Center, Toronto — diffusion models for layout generation, Jan 2024 · IPM (Institute for Research in Fundamental Sciences), Tehran — decomposing complete tripartite graphs into 5-cycles, 2017.
- **Reviewer:** CVPR 2026, ECCV 2024, NeurIPS 2023, SIGGRAPH 2023, ICCV 2023, CVPR 2023, ACCV 2022, ICCV 2021.

Honors & Awards

- 1st place, Financial Agentic Retrieval Grand Challenge (FinAgentBench), AI for Finance Symposium — 2025
- One RBC Team Award, RBC — 2025
- Spotlight presentations, NeurIPS 2023 and BMVC 2017
- Helmut & Hugo Eppich Family Graduate Scholarship, SFU — 2021, 2023
- SFU Computing Science Graduate Fellowship — 2020 – 2022
- Graduate Scholarship for Excellent Foreign Students, Seoul National University — 2017 – 2019
- Bronze medal, 20th Iranian National Olympiad in Informatics — 2010

Teaching

- **Simon Fraser University** — TA for ten courses (2020–2023): Computer Vision (×2), Data Structures (×2), Introduction to Programming (×2), Software Development Methods, Networks, Computing Systems, Communication Strategies.
- **Sharif University of Technology** — TA for Image Processing, Data Structures, Discrete Mathematics, Computer Programming (2013–2016). **NODET high schools** — Informatics Olympiad and C++ courses (2011–2017).

Technical Skills

- **Methods:** diffusion and generative models, transformers, structured/vector output prediction, retrieval-augmented generation and reranking, agentic LLM workflows, contrastive learning, time-series forecasting, 3D vision.
- **Languages & frameworks:** Python, PyTorch, C++, SQL, TensorFlow, MATLAB.