MACHINE LEARNING · WEIGHT REUSE, TEST-TIME COMPUTE & AI FOR SCIENCE

Utku Umur Acikalin

uuacikalin@gmail.com | utkuumur.github.io | Ithaca, NY
Google Scholar | github.com/utkuumur | LinkedIn

Fourth-year Cornell CS Ph.D. candidate advised by Carla P. Gomes. Researches weight reuse, looped models, test-time compute scaling, and learned search: how shared parameters and iterative refinement let models go beyond their training depth and training distribution. Also builds neuro-symbolic models and 100M+-scale retrieval systems for molecular discovery.

EDUCATION

Cornell University · Ph.D. in Computer Science — 2023 - expected May 2028 (05/28)
Advisor: Carla P. Gomes

TOBB University of Economics and Technology — Ankara, Turkey
M.S., Computer Engineering, 2019 - 2022 · GPA 4.00/4.00
B.S., Computer Engineering, 2015 - 2019 · GPA 3.74/4.00

RESEARCH EXPERIENCE

Ph.D. Researcher · Cornell University, Gomes Lab · with the Schroeder Lab (chemistry) — 2023 - present

MoRE: Mixture of Reused Experts

Co-first author · COLM 2026 (accepted) · ICML 2026 Weight-Space Symmetries Workshop · arXiv · Code

- Co-designed a hybrid between standard MoEs and looped (Universal) Transformers: adjacent layers share one expert pool with per-layer routers and learnable depth embeddings, giving **~18x more routing choices with no added expert parameters**.
- Outperformed a DeepSeek-v3-style MoE and MoEUT across **114M–1.15B parameters** at matched parameters and FLOPs, requiring minimal changes to standard MoE pipelines.

X²GNN: two-level looped GNNs & test-time compute scaling

First author · ICLR 2025 · Paper · Code

- Built a two-level looped GNN (4-layer recurrent inner GIN/GAT block + outer stochastic output-space refinement loop) that mitigates recurrent “overthinking”: making the refinement deterministic raises the MIS gap from **0.43% to 4.97%**.
- Trained on a single refinement step, yet keeps improving with more inference loops, and exposes an explicit parallel-breadth (*C*) vs. sequential-depth (*T*) test-time compute trade-off across graph geometries.
- Strong size and OOD generalization: **1.2% from optimal** on large Max Clique graphs 4x larger than training (prior neural: >22%), and generalizes from 10^3-node training graphs to **10^6 nodes** on regular-graph MIS.

GEMS: learned iterative search over molecular isomer graphs

First author · under review at ICLR 2027 · ICML 2026 workshops: GenBio (spotlight), AI4Physics · arXiv

- Formulated structure elucidation as learned discrete search over formula- and valence-preserving bond edits; a single forward pass scores all valid edits from spectrum-predicted fingerprints, with shortest-path supervision enabling large-scale structural pretraining.
- Reached **29.6% Hit@1 on NPLIB1 and 30.2% on MassSpecGym** (target excluded from initialization), outperforming FRIGID on NPLIB1 in ~3.3 min vs. 26.2 h.

AIMe & MS²KOSMOS: neuro-symbolic modeling & 100M+-scale retrieval

Co-first author · bioRxiv 2026, under review at Nature · Preprint · Public system

- Co-designed **AIMe**, a neuro-symbolic MS/MS predictor (**0.83 cosine on NIST** vs. 0.75 strongest baseline), and built CPU/GPU retrieval over **105M structures and 800M+ spectra** with few-second interactive queries and synthesis-validated discoveries.

SELECTED PUBLICATIONS & PREPRINTS

* Equal contribution.

MoRE: Mixture of Reused Experts.
E. S. Qiu*, **U. U. Acikalin***, J. Lovelace, C. Belardi, A. B. Mulchandani, C. P. Gomes, K. Q. Weinberger.
COLM 2026, accepted. Earlier version: ICML 2026 Weight-Space Symmetries Workshop.

GEMS: Molecular Structure Elucidation as Learned Search over Isomer Graphs.
U. U. Acikalin, Y. Wang, G. J. Gouveia, T. Schwertfeger, F. C. Schroeder, C. P. Gomes.
arXiv 2026; under review at ICLR 2027. Earlier version: ICML 2026 workshops, GenBio (spotlight) and AI4Physics.

Charting the small-molecule universe from mass spectra with neuro-symbolic AI.
U. U. Acikalin*, D. Feng*, A. M. Ferber*, G. J. Gouveia*, et al.
bioRxiv 2026; under review at Nature.

Learning to Explore and Exploit with GNNs for Unsupervised Combinatorial Optimization.
U. U. Acikalin, A. M. Ferber, C. P. Gomes.
ICLR 2025.

Intellectual property protection lost and competition: An examination using large language models.
U. U. Acikalin, T. Caskurlu, G. Hoberg, G. M. Phillips.
Journal of Financial Economics 182:104306, 2026.

Unsupervised Combinatorial Probabilistic Reasoning: Probabilistic Coin Change Problem.
Z. Qu, Y. Wang, **U. U. Acikalin**, et al.
AAAI 2026.

Earlier work (11 papers, 2019 - 2025): combinatorial optimization (INFORMS Journal on Computing, Constraints, GECCO), algorithmic game theory, NLP (Natural Language Engineering, SIGIR PatentSemTech workshop), and economics. Full list at utkuumur.github.io.

TECHNICAL SKILLS

Machine learning: PyTorch, Hugging Face Transformers, scikit-learn; mixture-of-experts and looped transformers, graph neural networks, test-time compute & learned search policies, LLM fine-tuning and extraction.

Systems, optimization & HPC: CPU/GPU similarity search at 100M+ scale, MPI, AWS; integer programming (Gurobi, CPLEX), OR-Tools, A* and heuristic search.

Languages: Python, C++, CUDA, Java, JavaScript.

Molecular ML: RDKit; MS/MS spectrum prediction and retrieval, fragmentation modeling, molecular fingerprints (Morgan, MIST), isomer-graph search; NIST20, MassSpecGym, NPLIB1, CASMI, GNPS.

PRESENTATIONS, SERVICE & TEACHING

ICML 2026, Seoul: GEMS at GenBio (spotlight) and AI4Physics; MoRE at Weight-Space Symmetries.
ICLR 2025, Singapore: X²GNN.

Reviewer: NeurIPS, ICLR, AAAI; ICML 2026 GenBio Workshop; IEEE Transactions on Circuits and Systems I.

Teaching assistant: Introduction to Computer Programming; Design and Analysis of Algorithms; Social Networks; Combinatorics.

SELECTED AWARDS

TUBITAK 2210 M.S. Scholarship, 2020 - 2022; **TOBB ETU Graduate Special Merit Scholarship**, 2019 - 2021; **TUBITAK Undergraduate Research Project Contest**, 2nd prize, 2019.