



Research Round-Up

September 2026

Discover our latest insights
in natural and artificial
intelligence research

-  Researchers including Institute Investigator Yilun Du and Affiliate Faculty Mengyu Wang **developed** a representation-guided method to improve medical image interpretation by multimodal language models.
-  Graduate Fellow William Tong and collaborators **show** how gating mechanisms influence the balance between memorization and in-context learning in state-space models.
-  A team including Institute Investigator SueYeon Chung and Chi-Ning Chou **demonstrate** how neural manifolds can be optimized to support representational learning in an olfactory memory network.

Featured Formula

$$W_1(P, Q) = \inf_{\pi \in \Pi(P, Q)} \mathbb{E}_{(x, y) \sim \pi} [d(x, y)]$$

Formula from the preprint "[Curriculum Learning as Transport: Understanding Curricula with Wasserstein Geodesics](#)," that captures the minimum cost of moving training probability from one curriculum stage to another, forming the basis for smooth easy-to-hard learning paths.

Featured Figure




Figure from the article "[Representational learning by optimization of neural manifolds in an olfactory memory network](#)," showing how training separates odor representations, making familiar and related odors more distinct and easier to classify in the brain.

Featured Artifacts

[PlayTrain: An Efficient Reinforcement Learning Framework for LLM-Generated Adaptable JavaScript Games](#)

Truong, Ryan, Lance Ying, Samuel J. Gershman, and Kazuki Irie (2026).

[Get the Code](#) ►

[MAxBench: A Multinomial Concept Recovery Benchmark.](#)

Appapogu, Divya, Freya Behrens, Yonatan Belinkov, and Aaron Mueller (2026).

[Get the Code](#) ►



AI Innovation

Discrete Beckmann Transport Models for One-Step Language Modeling and Reasoning.

Tang, Sophia, and Shiyi Wang. arXiv:2609.15903 (2026).

Curriculum Learning as Transport: Understanding Curricula with Wasserstein Geodesics.

Shin, Changho, and David Alvarez-Melis. arXiv:2609.09099 (2026).

PlayTrain: An Efficient Reinforcement Learning Framework for LLM-Generated Adaptable JavaScript Games.

Truong, Ryan, Lance Ying, Samuel J. Gershman, and Kazuki Irie. arXiv:2609.09059 (2026).

Thinking at the Right Size: Amortized Distillation Across Post-Trained LLMs.

Zhou, Yan, Sara Kangaslahti, Jonathan Geuter, Nihal V. Nayak, Marco Fumero, Francesco Locatello, and David Alvarez-Melis. arXiv:2608.22854 (2026).

Interpretability and AI Theory

Directing Large Language Models to Follow the Letter or Spirit of the Law.

Peng, Qian, Andrew Li, Sam Chen, Sonia K. Murthy, Yonatan Belinkov, and Tomer D. Ullman. arXiv:2609.23083 (2026).

On the Importance of Gating: Memorization vs. In-Context Learning in State Space Models.

Tong, William L., Aryo Lotfi, Emmanuel Abbe, Vaggelakos Kostas, Vishnu Banna, Etai Littwin, Josh Susskind, Cengiz Pehlevan, and Eran Malach. arXiv:2609.16540 (2026).

MAxBench: A Multinomial Concept Recovery Benchmark.

Appapogu, Divya, Freya Behrens, Yonatan Belinkov, and Aaron Mueller. arXiv:2609.13072 (2026).

Do LLMs Choose Like Humans? Using Cognitive Theory to Evaluate LLM Decision-Making.

Sun, Johnathan, Andrei Shleifer, and Yonatan Belinkov. arXiv:2609.22225 (2026).

Applications of AI

Representation-guided in-context learning for medical image interpretation with multimodal large language models.

Zhao, Minda, Fangyu Hu, Yan Luo, Yutong Yang, Jiahui Cai, Kaichen Zhou, Manling Li, et al. arXiv:2609.24057 (2026).

World models for biomedicine.

Noori, Ayush, Nic Fishman, Ada Fang, Lukas Fesser, and Marinka Zitnik. Cell (2026).

A Brain-inspired Hierarchical Framework for Zero-Shot Robot Task Reasoning and Execution.

Wang, Guangming, Pengfei Ye, Qizhen Ying, Yixiong Jing, Yuxiang Ma, Haonan Chen, Haibing Wu, Olaf Wysocki, Molong Duan, and Brian Sheil. arXiv:2609.05985 (2026).

Deriving OCT-Equivalent Retinal Nerve Fiber Layer Thickness Maps from Fundus Photographs with Deep Learning Improves Glaucoma Diagnosis.

Shi, Lily, Min Shi, In Young Chung, Louis R. Pasquale, Lucy Q. Shen, and Mengyu Wang. Ophthalmology Science (2026).

Neuroscience & NeuroAI

Representational learning by optimization of neural manifolds in an olfactory memory network.

Hu, Bo, Nesibe Z. Temiz, Chi-Ning Chou, Peter Rupprecht, Claire Meissner-Bernard, Benjamin Titze, SueYeon Chung, and Rainer W. Friedrich. Nature Neuroscience (2026).

Note: This is a partial list of articles and preprints published by Kempner-affiliated researchers in the last month. Papers are listed by topic and publication/upload date, with the most recent first.



View this report and other Kempner Research-Round-ups:
kempnerinstitute.harvard.edu/kempner-community/research-roundup/