Kempner INSTITUTE

For the Study of Natural & Artificial Intelligence at Harvard University

Research Round-Up

August 2026

Discover our latest insights in natural and artificial intelligence research

Highlights

- Associate Faculty Haim Sompolinsky and his team **investigate** how sparse prototype-based representations support classification and prediction across different modalities, exploring how intelligent systems organize concepts.
- A team including Kempner Institute Investigator Yilun Du introduces **Equilibrium Forcing**, an adaptive video generation method that produces more coherent long-duration videos without explicit noise conditioning.
- Postdoctoral Fellow Ryan Badman and his collaborators **find** that mice assign intrinsic value to information, represented in the orbitofrontal cortex.

Featured Formula

$$\mathrm{EIG}(q_t) = H(p_t) - \sum_a P(a \mid q_t)\, H\left(p_{t+1}^{(a)}\right) = I(Y; A_t \mid q_t)$$

Expected information gain measures how much uncertainty a question is expected to remove about a hidden target, from the preprint **"Do LLMs Know What to Ask and When? Evaluating Multi-Turn Information Seeking."** The more effectively a question narrows the possible answers, the higher its information gain is.

Featured Figure

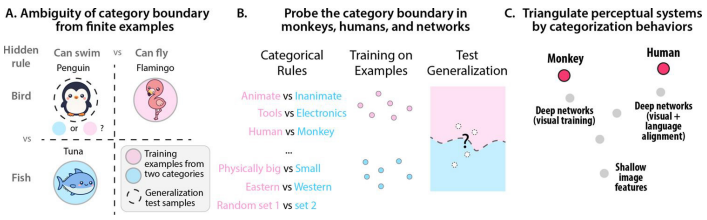

Diagram from the article **"How Monkeys Carve Up the Visual World"**. Given familiar labeled images in the first column, monkeys, humans, and neural networks infer visual category boundaries and classify novel images, revealing how each system organizes the visual world.

Featured Artifacts

Do LLMs Know What to Ask and When? Evaluating Multi-Turn Information Seeking
Yepeng Huang, Jiawen Zhang, Michelle Dai, Xiaorui Su, Shanghua Gao, Zi Wang, and Marinka Zitnik (2026).

Get the Code

Improving Outdoor Navigation for People with Blindness Using an AI-Driven Smartphone Application and Personalized Audio Guidance
Raymond Liu and Patrick Slade (2026).

Get the Code

AI Innovation

Do LLMs Know What to Ask and When? Evaluating Multi-Turn Information Seeking.
Huang, Yepeng, Jiawen Zhang, Michelle Dai, Xiaorui Su, Shanghua Gao, Zi Wang, and Marinka Zitnik. arXiv:2608.14808 (2026).

Equilibrium Forcing: Adaptive Video Generation Without Noise Conditioning.
Lillemark, Hansen Jin, Alex Rojas, Zachary Novack, Runqian Wang, Yilun Du, Yian Ma, Taylor Berg-Kirkpatrick, and Rose Yu. arXiv:2608.14706 (2026).

Beckmann Transport Models: From Autonomous Flows to One-Step Maps.
Cheuk-Kit, Lee, Florentin Coeurdoux, Peter Potaptchik, Yilun Du, Michael Samuel Albergo, and Eric Vanden-Eijnden. arXiv:2608.01692 (2026).

Bridging Compute-and Data-Optimal Pretraining.
Qin, Tian, Kimia Hamidieh, and David Alvarez-Melis. arXiv:2607.25271 (2026).

Interpretability and AI Theory

Shared Circuits for Shared Grammar: Tracing Subject-Verb Agreement Across Languages.
Gidi, Isabella, Antonio AlmudÄŠvar, Core Francisco Park, Naomi Saphra, and Ricard Marxer. arXiv:2608.18545 (2026).

Sparse Prototype Code Underlies Classification and Prediction Across Modalities.
Avidan, Yehonatan, Daniel D. Lee, and Haim Sompolinsky. arXiv:2608.15632 (2026).

Applications of AI

Improving outdoor navigation for people with blindness using an AI-driven smartphone application and personalized audio guidance.
Liu, Raymond and Patrick Slade. Nature Biomedical Engineering (2026).

Theory-Driven Computational Psychiatry and Its Challenges: An Optimistic Outlook.
Fisher, Elizabeth, Noor Sajid, Laura Convertino, and Jakob Hohwy. Neuroscience & Biobehavioral Reviews (2026).

Neuroscience & NeuroAI

Sparse input representations explain odor discrimination in complex, concentration-varying mixtures.
McCalmon, Hannah, George Cai, Constantine Tsibouris, Noé Hamou, Jo Hoskins, Farhad Pashakhanloo, SueYeon Chung, Vikrant Kapoor, and Venkatesh N. Murthy. Cell Reports (2026).

How monkeys carve up the visual world.
Wang, Binxu. eLife 15: e112456 (2026).

Inferring brain-wide interactions using data-constrained recurrent neural network models.
Perich, Matthew G., Charlotte Arlt, Sofia Soares, Siyan Zhou, Manuel Beiran, Aaron S. Andalman, Tyler Benster et al. Neuron (2026).

Representations of the intrinsic value of information in mouse orbitofrontal cortex.
Bussell, Jennifer J., Ryan P. Badman, David Márton, Ethan S. Bromberg-Martin, L. F. Abbott, Kanaka Rajan, and Richard Axel. Nature Neuroscience (2026).

When Branch-Local Shunting Helps: A Gain-Load-Alignment Principle for Dendritic E/I Networks.
Safaai, Houman, Maceo Richards, Naeem Khoshnevis, and Bernardo L. Sabatini. arXiv:2607.24990 (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/