

## RESEARCH

My goal is to build learning systems that exhibit human-like reasoning capabilities. I study compositional generalization across model families (CLIP-like models, diffusion models, VLMs), investigating what general principles enable models to re-use knowledge from training when testing requires recombining it. On the theoretical side, I study idealized systems to identify necessary properties for compositional reasoning; on the practical side, I diagnose why current models fail and how to fix them. My recent work extends from single-object compositionality to multi-object understanding (binding properties to objects, spatial relationships, counting).

Throughout this line of work I have used synthetic environments and simulators for controlled evaluation. More recently, I am building compositionally diverse, procedural visual worlds (in the spirit of CLEVR/GQA) that enable step-by-step verification of visual reasoning in foundation-scale autoregressive VLMs, with the longer-term goal of using such environments for training.

## EDUCATION

- **University of Tübingen**, Germany · *Ph.D., Machine Learning* Sep 2023 - Present
- **University of Tübingen**, Germany · *M.Sc., Machine Learning* · GPA: 1.0 / 1.0 2020 - 2022
- **Kaunas University of Technology**, Lithuania · *B.Sc., Computer Science* 2015 - 2019

## RESEARCH EXPERIENCE

- **PhD Student**, University of Tübingen (advisor: Seong Joon Oh) Sep 2023 - Present
- **Research Assistant**, Max Planck Institute for Intelligent Systems (advisor: Georg Martius) Jun 2021 - Jun 2022
- **Research Assistant**, Kaunas University of Technology, ML group (advisor: Mantas Lukoševičius) and Cryptography group (advisor: Eligijus Sakalauskas) Jan 2018 - Dec 2019

## PUBLICATIONS

\* indicates equal contribution

- [1] **Uselis, Arnas**, Yujin Jeong, Yanpeng Zhao, Alexander Rubinstein, Seong Joon Oh, Yonatan Bitton, and Paul Gavrikov. "CRISP: Compositional Reasoning over Images via Stackable Programs for VLMs." Preprint, 2026.
- [2] **Uselis, Arnas\***, Darina Koishigarina\*, and Seong Joon Oh. "How Can Embedding Models Bind Concepts?" International Conference on Machine Learning (ICML), 2026. **Spotlight (top 2.2%)**.
- [3] Morelli, Fabian, **Arnas Uselis**, Ankit Sonthalia, and Seong Joon Oh. "[Sparse Autoencoders enable Robust and Interpretable Fine-tuning of CLIP models.](#)" arXiv preprint, 2026. [code](#)
- [4] Jung, Seokwon, Alexander Rubinstein, **Arnas Uselis**, Sangdoon Yun, and Seong Joon Oh. "[MEME: Multi-Entity & Evolving Memory Evaluation.](#)" arXiv preprint, 2026. [project](#)
- [5] Kargi, Bora, **Arnas Uselis**, and Seong Joon Oh. "[Half-Truths Break Similarity-Based Retrieval.](#)" arXiv preprint, 2026. [code](#)
- [6] Koishigarina, Darina, **Arnas Uselis**, and Seong Joon Oh. "[CLIP Behaves like a Bag-of-Words Model Cross-modally but not Uni-modally.](#)" International Conference on Learning Representations (ICLR), 2026. [code](#)
- [7] **Uselis, Arnas**, Andrea Dittadi, and Seong Joon Oh. "[Compositional Generalization Requires Linear, Orthogonal Representations in Vision Embedding Models.](#)" International Conference on Machine Learning (ICML), 2026. **Spotlight (top 2.2%)**. [code](#)
- [8] Jeong, Yujin, **Arnas Uselis**, Iro Laina, Seong Joon Oh, and Anna Rohrbach. "[When Do Diffusion Models learn to Generate Multiple Objects?](#)" International Conference on Machine Learning (ICML), 2026. [code](#)
- [9] Sonthalia, Ankit, **Arnas Uselis**, and Seong Joon Oh. "[On the Rankability of Visual Embeddings.](#)" Conference on Neural Information Processing Systems (NeurIPS), 2025. [code](#)
- [10] Jeong, Yujin\*, **Arnas Uselis\***, Seong Joon Oh, and Anna Rohrbach. "[Diffusion Classifiers Understand Compositionality, but Conditions Apply.](#)" Conference on Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track, 2025. [code](#)

[11] **Uselis, Arnas**, Andrea Dittadi, and Seong Joon Oh. "Does Data Scaling Lead to Visual Compositional Generalization?" International Conference on Machine Learning (ICML), 2025.

[12] **Uselis, Arnas**, and Seong Joon Oh. "Intermediate Layer Classifiers for OOD Generalization." International Conference on Learning Representations (ICLR), 2025. [code](#)

[13] **Uselis, Arnas**, Andrea Dittadi, and Seong Joon Oh. "Beyond Decodability: Linear Feature Spaces Enable Visual Compositional Generalization." Workshop on Spurious Correlation and Shortcut Learning at the International Conference on Learning Representations (ICLR), 2025.

[14] Lukoševičius, Mantas, and **Arnas Uselis**. "Efficient implementations of echo state network cross-validation." Cognitive Computation 15.5 (2023): 1470-1484. [code](#)

[15] Lukoševičius, Mantas, and **Arnas Uselis**. "Task-Synchronized Recurrent Neural Networks." arXiv preprint arXiv:2204.05192, 2023. [code](#)

[16] **Uselis, Arnas**, Mantas Lukoševičius, and Lukas Stasytis. "Localized convolutional neural networks for geospatial wind forecasting." Energies 13.13 (2020): 3440. [code](#)

[17] Sakalauskas, Eligijus, Aleksejus Mihalkovich, and **Arnas Uselis**. "Security analysis of KAP based on enhanced MPF." IET Information Security 14.4 (2020): 410-418.

[18] Lukoševičius, Mantas, and **Arnas Uselis**. "Efficient cross-validation of echo state networks." International Conference on Artificial Neural Networks (ICANN). Springer, Cham, 2019.

## THESIS SUPERVISION

---

- **Darina Koishigarina**, University of Tübingen · Understanding CLIP's cross-modal and uni-modal behavior [paper](#) 2024-2025
- **Bora Kargi**, University of Tübingen · Investigating language injection into captions in CLIP models at scale (LAION-400M) [paper](#) 2024-2025
- **Fabio Morelli** (co-supervised with Ankit Sonthalia), University of Tübingen · Understanding robust fine-tuning from the perspective of preserved features during fine-tuning *Ongoing*
- **Seokwon Jung** (co-supervised with Alex Rubinstein), KAIST · Evaluating memory in multi-turn conversations that require cascading events in LLMs *Ongoing*

## Review experience

---

**Conference reviewer:** ICLR 2024, ICLR 2025, WACV 2025, ICLR 2026, ICML 2026, CVPR 2026

**Recognitions:** **Outstanding Reviewer**, CVPR 2026 (top 5% of 17,491 reviewers); **Silver Reviewer**, ICLR 2026

## Skills

---

- **ML topics:** compositional generalization, representation learning, diffusion models, autoregressive models, vision-language models, OOD robustness, synthetic environments.
- **Frameworks:** PyTorch (primary), NumPy, SciPy, scikit-learn.
- **Programming:** Python, C++; Git, Linux, SLURM/HPC.

## Selected Projects

---

- **STND** · Built an experiment manager for SLURM clusters with Google Sheets integration. [code](#)
- **Minimal Transformer** · Implemented the original Transformer architecture from scratch in PyTorch. [code](#)
- **VAE and semi-supervised VAE** · Replicated core results of VAE and SSL-VAE papers on MNIST. [code](#)
- **RL tournament** · Implemented a reinforcement learning agent with self-play for a competitive multi-agent environment; placed 3rd. [report](#)

## HONORS & AWARDS

---

- 3rd Place, "**Dr. J. P. Kazicko**" National Programming Olympiad *2015*
- Recipient, **KTU Talent Scholarship** *2016 - 2019*
- Winner, **KTU Programming Competition** *2017 - 2018*
- **ACM ICPC:** NEERC Semi-Finals (2017), NEERC Western Subregion (2016-2018)