Seyedmorteza SADAT

Ph.D. Student | Disney Research & ETH Zurich

+41 779 75 7708 @ ssadat@ethz.ch
linkedin.com/in/mortezasadat
Google Scholar
Stampfenbachstr. 48, 8092 Zurich

Education

2022 - present Ph.D. Computer Science, **Disney Research Studios & ETH Zurich**, Zurich, Switzerland
Topic : Scalable and Controllable Image Generation Using Diffusion Models
Thesis Supervisors : Prof. Dr. Otmar Hilliges, Dr. Romann M. Weber

2020 - 2022 M.Sc. Computer Science, **ETH Zurich**, Zurich, Switzerland
Thesis : Alternative perspectives on Generative Modeling, Major : Machine Intelligence
Thesis Supervisors : Dr. Romann M. Weber, Prof. Dr. Markus Gross
GPA : **5.78/6**

2015 - 2020 B.Sc. Electrical Engineering, **Sharif University of Technology**, Tehran, Iran
Major : Communication Systems
GPA : **19.19/20**

Technical Skills

Deep Learning Frameworks	PyTorch, JAX, Tensorflow, Keras, Accelerate
Machine Learning Frameworks	Scikit-Learn, XGboost
Computer Vision Frameworks	Diffusers, Detectron2, OpenCV
Programming Languages	Python, Java, C, C++, and Golang
Engineering Softwares	Matlab, Julia
Productivity Softwares	Latex, Microsoft Office

Doctoral Research Experience | Disney Research & ETH Zurich

[Spotlight at ICLR 2024] Improving the Diversity of Diffusion Models

- Identified a lack of sample diversity in diffusion models.
- Solved the diversity issue by injecting noise to the conditioning signal at inference.
- Achieved state-of-the-art FID on ImageNet using the DiT-XL/2 model.
- Tested the pipeline in PyTorch on several models including Stable Diffusion, DiT, and LDM.

[NeurIPS 2024] Efficient Autoencoders for Latent Diffusion Models

- Improved encoder efficiency in latent diffusion models.
- Reduced encoder size by 6× using wavelet transforms.
- Enhanced training with better adversarial setup and decoder design.
- Matched SD-VAE reconstruction quality with smaller model.
- Built pipeline with Diffusers, PyTorch, and Accelerate.

[ICLR 2025] Eliminating Oversaturation and Artifacts of Classifier-Free Guidance

- Addressed the oversaturation issue in diffusion models using the orthogonal complement of the CFG update.
- Enhanced robustness to oversaturation by incorporating adaptive rescaling into CFG.
- Verified the method across a wide range of models, including SD, SDXL, SD3, SDXL-Lightning, DiT, and EDM2.
- **Recognized by the community**, with planned integration into the Diffusers library.

[ICLR 2025] Rethinking Classifier-free Guidance in Diffusion Models

- Proposed a novel formulation for classifier-free guidance (CFG) that eliminates the need for an unconditional model.
- Extended CFG to unconditional models via time-step perturbation sampling.
- Improved efficiency by removing the CFG component during training.
- Verified both methods in PyTorch across models including SD, SDXL, DiT, and EDM2.

[TMLR] Efficient Distillation of Classifier-Free Guidance using Adapters

- Doubled the speed of CFG by integrating the guidance step directly into the model's forward pass.
- Improved guidance distillation efficiency by training lightweight adapters instead of full model fine-tuning.
- Enhanced distillation quality by supervising on CFG-guided trajectories.
- Validated the complete pipeline in PyTorch across multiple models, including DiT, SD, and SDXL.

[Under review] Guidance in the Frequency Domain Enables High-Fidelity Sampling at Low CFG Scales

- Introduced a novel reformulation of classifier-free guidance in the frequency domain using Laplacian pyramids.
- Demonstrated that low- and high-frequency components of guidance affect diversity and fidelity differently.
- Proposed asymmetric frequency scaling to enhance visual quality at low CFG scales without sacrificing diversity.
- Achieved consistent improvements across diffusion models (SD, SDXL, SD3, DiT, EDM2) with no extra training cost.

[Under review] HIGS : History-Guided Sampling for Plug-and-Play Enhancement of Diffusion Models

- Proposed a momentum-based inference technique that uses past model predictions to enhance diffusion sampling.
- Improved image fidelity and detail without retraining or additional computation.
- Achieved state-of-the-art FID of 1.61 on unguided ImageNet generation with only 30 sampling steps.
- Demonstrated universal compatibility across diffusion models (SDXL, SD3, DiT, SiT) under low-NFE and low-CFG regimes.

Publications [Google Scholar]

[1] **Seyedmorteza Sadat**, Jakob Buhmann, Derek Bradley, Otmar Hilliges, and Romann M. Weber. CADS : Unleashing the diversity of diffusion models through condition-annealed sampling. In *The Twelfth International Conference on Learning Representations (*Spotlight at ICLR 2024*)*, .

[2] **Seyedmorteza Sadat**, Jakob Buhmann, Derek Bradley, Otmar Hilliges, and Romann M. Weber. LiteVAE : Lightweight and efficient variational autoencoders for latent diffusion models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems (*NeurIPS 2024*)*, .

[3] **Seyedmorteza Sadat**, Manuel Kansy, Otmar Hilliges, and Romann M. Weber. No training, no problem : Rethinking classifier-free guidance for diffusion models. In *The Thirteenth International Conference on Learning Representations (*ICLR 2025*)*, .

[4] **Seyedmorteza Sadat**, Otmar Hilliges, and Romann M. Weber. Eliminating oversaturation and artifacts of high guidance scales in diffusion models. In *The Thirteenth International Conference on Learning Representations (*ICLR 2025*)*, .

[5] **Seyedmorteza Sadat**, Farnood Salehi, and Romann M. Weber. Higs : History-guided sampling for plug-and-play enhancement of diffusion models, 2025. URL `https://arxiv.org/abs/2509.22300`.

[6] Cristian Perez Jensen and **Seyedmorteza Sadat**. Efficient distillation of classifier-free guidance using adapters, 2025, (Under Review). URL `https://arxiv.org/abs/2503.07274`.

[7] **Seyedmorteza Sadat**, Tobias Vontobel, Farnood Salehi, and Romann M. Weber. Guidance in the frequency domain enables high-fidelity sampling at low cfg scales, 2025.

[8] Justin Studer, Dhruv Agrawal, Dominik Borer, **Seyedmorteza Sadat**, Robert W. Sumner, Martin Guay, and Jakob Buhmann. Factorized motion diffusion for precise and character-agnostic motion inbetweening. In *Proceedings of the 17th ACM SIGGRAPH Conference on Motion, Interaction, and Games (*MIG 2024*)*.

[9] Christoph Gebhardt, Robin Willardt, **Seyedmorteza Sadat**, Chih-Wei Ning, Andreas Brombach, Jie Song, Otmar Hilliges, and Christian Holz. Regressor-guided image editing regulates emotional response to reduce online engagement, 2025. URL `https://arxiv.org/abs/2501.12289`.

[10] Tobias Vontobel, **Seyedmorteza Sadat**, Farnood Salehi, and Romann M. Weber. Hiwave : Training-free high-resolution image generation via wavelet-based diffusion sampling, 2025.

[11] Javad Rajabi, Soroush Mehraban, **Seyedmorteza Sadat**, and Babak Taati. Token perturbation guidance for diffusion models, 2025.

Other Research Experiences

March 2022 – September 2022 | **Alternative perspectives on Generative Modeling, Disney Research Studios, Zurich, Switzerland**

- Master thesis on new perspectives for generative modeling in the computer vision domain
- Experimenting with novel modifications of GANs, Normalizing Flows, and Score-based models.

January 2022 – June 2021 | **Generative Modeling of Full-body Human Images, ETH Zurich, Switzerland**
- Preprocessing the dataset using DensePose segmentation and Human parsing modules
- Designing a GAN architecture for generating full-body human images in clothing
- Training the network and enhancing the artifacts in generated images by analyzing the failure cases
- Comparing the model with existing image generation pipelines

June 2021 – February 2021 | **Learning Individually Fair Representations in Vision based Systems, ETH Zurich, Switzerland**
- Extended the individually fair representation learning framework to high dimensional image data
- Trained a disentangled β-VAE model for encoding logical constraints on image domain, such as changes in smile, gender, or age
- Trained an encoder-classifier framework that could classify similar inputs to similar outputs
- Tested the effectiveness of the algorithm on learning fair representations using MNIST and Celeb-A

June 2021 – February 2021 | **Improving Exploration in Self-supervised Reinforcement Learning via DPPs, ETH Zurich, Switzerland**
- Analyzed the effectiveness of exploration in self-supervised Reinforcement Learning
- Developed a new algorithm based on determinantal point processes for increasing exploration
- Compared the algorithm with existing approaches in terms of the diversity of selected goals
- Demonstrated empirically that using DPPs can result in more exploration

February 2020 – July 2019 | **Computer and Communication Science Department, EPFL, Switzerland**
- Examined the depth effect on the generalization ability of neural networks in the infinite-width regime
- Studied the relation between the neural tangent kernel and depth
- Proved empirically that networks with higher depth have better generalization in the kernel regime

Research Interests

- Diffusion-based generative models for image and video
- Addressing classical computer vision and graphics problems using generative priors
- Multimodal understanding and generation leveraging large language models (LLMs)

Honors and Awards

since September 2020 **Direct Doctorate Scholarship** awarded by the Department of Computer Science, ETH Zürich
September 2019 **Research fellowship**, computer and communication science department, EPFL
September 2018 **Research fellowship**, Information theory department, The Chinese University of Hong Kong
September 2018 **Ranked 9** amongst 175 electrical engineering students, class of 2019, Sharif University
August 2015 **Silver medal** in the 9th **International Olympiad** on Astronomy and Astrophysics, Indonesia
September 2014 **Gold medal** in the 10th **National Olympiad** on Astronomy and Astrophysics, Iran

Invited Talks

- **Stability AI**, Presented "LiteVAE : Lightweight and efficient variational autoencoders for latent diffusion models"
- **Disney Research**, Presented "State-of-the-art guidance methods for diffusion-based generative models"

Academic Services

- Served as a reviewer for NeurIPS 2024, ACMM 2024, ICML 2025, ICCV 2025, and NeurIPS 2025
- Recognized as an outsanding reviewer for ICCV 2025 and NeurIPS 2025
- Supervised multiple Master's and Bachelor's projects at Disney Research and ETH
- Worked as a TA for several ETH courses (e.g., Machine Perception, Probabilistic AI, Introduction to Machine Learning).

Course Highlights

Machine Perception	Deep Learning	Computational Intelligence Lab
Advanced Machine Learning	Probabilistic AI	Reliable and interpretable AI
Convex Optimization	Information Theory	Signal Processing

Languages

- English : Fluent (TOEFL Score : 115, GRE Score : 159 Verbal, 169 Quantitative)
- Persian : Native
- German : Basic

References

- Dr. Romann M. Weber (Disney Research Studios)
- Prof. Dr. Otmar Hilliges (ETH Zurich)
- Prof. Dr. Rüdiger Urbanke (EPFL)