Curriculum Vitae – Jinho Chang

Personal Information

Jinho Chang
Seoul, South Korea
Personal Webpage
https://github.com/jinhojsk515
www.linkedin.com/in/jinho-chang-09a2a1284
Google Scholar https://scholar.google.com/citations?user=GrHZicUAAAAJ

Ph.D. student at KAIST, currently a member of the BioImaging-Signal Processing & Learning Lab(BISPL). I am interested in diffusion and flow-matching models for their variational applications and inference-time techniques, and multi-modal models for informative representations and cross-modal generation. Also, leveraging my chemistry background, I aim to apply state-of-the-art AI techniques to various important chemical challenges with minimal adaptation costs.

Research Interests

-AI for chemistry : I'm interested in various molecular tasks that are hard to be solved in principle, including conditional molecule design and editing, synthesizability and retro-synthesis prediction, intrinsic property prediction, etc. Incorporating and aligning different information such as structure, biochemical properties, and natural text into the representation has enabled the improved the performance of different generative and foundation models.

-Diffusion/flow model and its applications : Beyond their strong generative performance, analytic and variational analyses on diffusion and flow-matching models have enabled a wide range of inference-time sampling techniques. I'm working on the potential training-free guidance and editing methods for continuous and discrete diffusion/flow models, including reward-guided editing, negative sampling guidance, score distillation, molecule optimization, etc.

-Multimodal learning : Leveraging various multimodal learning of a shared data space or representation alignment, I've worked for a better cross-modal understanding and generation for a vision-language capabilities for medical imaging, and molecular structure and its properties.

Publications

More details can be found in Google Scholar Page.

AI for Chemistry.

- LDMol : Text-Conditioned Molecule Diffusion Model Leveraging Chemically Informative Latent Space [link]
  Jinho Chang and Jong Chul Ye *ICML, 2025*
  We present a new latent diffusion model LDMol for molecule generation with complex conditions like natural texts. Leveraging a chemically informative latent space obtained with contrastive learning, LDMol not only outperforms previous text-to-molecule generative models, but can also be applied to downstream tasks like retrieval and molecule editing.

- Bidirectional Generation of Structure and Properties Through a Single Molecular Foundation Model [link]
  Jinho Chang and Jong Chul Ye *Nature Communications, 2024*
  We introduce a multimodal molecular pre-trained model that integrates molecular structures and their biochemical properties, which enables various multimodal and unimodal tasks like conditional molecule generation, property prediction, molecule classification, and reaction prediction.

- Drug-likeness Scoring Based on Unsupervised Learning [link]
  Kyunghoon Lee, Jinho Chang, Seonghwan Seo, and Woo Youn Kim *Chemical Science, 2022*
  We proposed a novel unsupervised learning model that could quantify the drug-likeness of a given molecule, by adopting a language model trained via unsupervised learning. Our model showed more consistent performance across different datasets and gave more reasonable scores.

Image/Text Generative Models.

- LLM-CXR : Instruction-Finetuned LLM for CXR Image Understanding and Generation [link]
  Suhyeon Lee, Wonjun Kim, Jinho Chang, and Jong Chul Ye *ICLR, 2024*
  Taking inspiration from previous work on the transformer and VQ-GAN combination for bidirectional image and text generation, we developed a method for instruction-tuning a text-only LLM to gain vision-language capabilities for medical images. We let LLM understand tokenized visual inputs by instructing it to answer questions about image inputs and generate image outputs.

- DreamMotion : Space-Time Self-Similarity Score Distillation for Zero-Shot Video Editing [link]
  Hyeonho Jeong, Jinho Chang, Geon Yeong Park, and Jong Chul Ye *ECCV, 2024*
  We present a score distillation sampling for video editing to circumvent the standard reverse diffusion process. We propose to match the space-time self-similarities during the score distillation, demonstrating its superiority in altering appearances while accurately preserving the original structure and motion.

- Training-Free Reward-Guided Image Editing via Trajectory Optimal Control [link]
  Jinho Chang, Jaemin Kim, and Jong Chul Ye *In Review, 2025*
  Prior reward-guided sampling methods in diffusion/flow-matching largely focused on the generation task and remain underexplored for image editing. We propose a training-free editing framework that casts the reward-guided editing as a trajectory optimal control problem, and iteratively updates adjoint states to steer the trajectory endpoint into a reward-optimized output. Our method surpasses inversion-based training-free guidance baselines across diverse editing tasks, achieving stronger rewards with higher source image fidelity without reward hacking.

- Contrastive CFG : Improving CFG in Diffusion Models by Contrasting Positive and Negative Concepts [link]
  Jinho Chang, Changsun Lee, Hyungjin Chung, and Jong Chul Ye *In Review, 2025*
  We point out the drawback of naively negating the Classifier-Free Guidance (CFG) term, and propose a novel approach to enhance negative CFG in conditional diffusion models using contrastive loss. By aligning or repelling the denoising direction based on the given condition, our method overcomes the distortions caused by traditional negative CFG, achieving superior sample quality and effective removal of unwanted features across various scenarios.

- Ground-A-Score : Scaling Up the Score Distillation for Multi-Attribute Editing [link]
  Hangeol Chang, Jinho Chang, and Jong Chul Ye *arXiv, 2024*
  We propose a simple yet powerful model-agnostic image editing method that tackles complex image editing queries in a manner of divide-and conquer. Moreover, the selective application with a new penalty coefficient helps to precisely target editing areas while preserving the integrity of the objects in the source image.

Reviewer Participation

Nature Communications	*2024 -*
Communications Chemistry	*2024 -*
International Conference on Learning Representations	*2025, 2026*
Conference on Computer Vision and Pattern Recognition	*2025*
International Conference on Computer Vision	*2025*

Experiences

Internship as AI developer HITS (https://hits.ai/) *January 2021 - June 2021*
Development of a reinforcement learning model to generate molecules with specific chemical or pharmaceutical properties guided by the given RL reward.

Education

Ph.D. student in Artificial Intelligence *September 2022 - present*
Korea Advanced Institute of Science and Technology (KAIST)

B.S. in Computer Science and Chemistry (double major) *March 2018 - August 2022*
Korea Advanced Institute of Science and Technology (KAIST)