

Chen-Yi Lu

765-746-9977 | [Email](#) | [Website](#) | [Google Scholar](#) | [LinkedIn](#) | [GitHub](#)

Education

Purdue University , West Lafayette, IN Ph.D., Agricultural & Biological Engineering Thesis: <i>Towards Efficient and Robust Diffusion Models</i> Advisor: Prof. Somali Chaterji	Expected May 2027
M.S., Agricultural & Biological Engineering	May 2023
National Taiwan University , Taipei, Taiwan B.S., Biomechatronics Engineering	June 2018

Selected Publications

ECCV (Oral)	What CLIP Knows but Cannot Say: Recovering Negation from Frozen Intermediate Features C.-Y. Lu, Y.-S. Chen, S. Chaterji. 2026.	
ICCV	SKALD: Learning-Based Shot Assembly for Coherent Multi-Shot Video Creation C.-Y. Lu, M. M. Tanjim, I. Dasgupta, S. Sarkhel, G. Wu, S. Mitra, S. Chaterji. 2025.	
CVPR	Improving Semi-Supervised Semantic Segmentation with Sliced-Wasserstein Feature Alignment and Uniformity C.-Y. Lu, K. Derakhshandeh, S. Chaterji. 2025.	
ECCV	ReCon: Training-Free Acceleration for Text-to-Image Synthesis with Retrieval of Concept Prompt Trajectories C.-Y. Lu*, S. Agarwal*, M. M. Tanjim, K. Mahadik, A. Rao, S. Mitra, S. Saini, S. Bagchi, S. Chaterji. 2024.	
Preprint	PULSE: On-Device Adversarial Image Purification using Latent Super-Resolution Consistency Model C.-Y. Lu, Y.-S. Chen, S. Chaterji. Under review, 2026.	

First or co-first (*) author on all listed publications.

Professional Experience

ICAN Lab, Purdue University Efficient & Robust Vision-Language Models <i>Graduate Research Assistant</i>	West Lafayette, IN Aug 2021 – Present
--	--

- Developed PeakPatch (**ECCV 2026, Oral**), showing that CLIP's intermediate layers retain negation information that is discarded at the output projection. Recovered it with lightweight embedding- and score-correction networks over frozen features, with no backbone fine-tuning, raising NegBench accuracy from 56.5 to **74.3** over prior art.
- Developed SWSeg (**CVPR 2025**), a semi-supervised semantic segmentation method using an auxiliary sliced-Wasserstein loss on intermediate features to enforce feature-space alignment and uniformity; +11.8 mIoU on PASCAL VOC and +8.9 mIoU on ADE20K over supervised baseline.
- Built PULSE, an on-device adversarial purifier using a 1-step latent super-resolution consistency model to defend vision-language models, retaining **96% accuracy under APGD attack** versus 0% undefended, at 38.7 ms/image on Jetson AGX Orin with 173M parameters.

Adobe Research Ad Creative Intelligence & Video Generation <i>Research Scientist Intern</i>	San Jose, CA May 2026 – Aug 2026
---	-------------------------------------

- Designed a scene-level creative representation for click-through-rate (CTR) prediction that stays human-interpretable and directly editable, pairing decomposed scene features with a transformer CTR predictor trained on a custom ranking objective; reached **0.19** cross-validated rank correlation over 12K+ ad creatives across 40+ campaigns.
- Built an agentic video-editing pipeline that ingests trending topics, spawns planning subagents to draft structured edit specs, and orchestrates image- and video-generation models to re-cut marketing ads around the trend, closing the loop from creative attribution to generation; in a 200-ad study with 50 marketers, **66.3%** judged the edited variant more compelling and timely than the original.

Adobe Research Video Generation & Evaluation <i>Research Scientist Intern</i>	San Jose, CA May 2024 – Nov 2024
---	-------------------------------------

- Developed SKALD (**ICCV 2025**), a multi-shot video assembly method whose Learned Clip Assembly (LCA) score rates temporal and semantic coherence with a transformer encoder over frozen vision-encoder embeddings, reaching 69% agreement with human raters while training on labels for under 5% of clips via a joint contrastive and feature-regression objective.
- Formulated shot assembly as combinatorial search and developed LCA-guided beam search over candidate orderings, improving assembly IoU **+48.6% relative** over prior art at a 43% lower search cost.

Adobe Research Efficient Diffusion Models <i>Research Scientist Intern</i>	San Jose, CA May 2023 – Aug 2023
--	-------------------------------------

- Developed ReCon (**ECCV 2024**), a training-free diffusion acceleration method that decomposes prompts into visual concepts and retrieves stored denoising trajectories, with mid-trajectory noise aggregation to resume denoising from retrieved concepts; **40% fewer NFEs** at matched fidelity, and +1.8 Pick score over retrieval baselines across 3 benchmarks. **USPTO patent filed**; prototyped for integration into Adobe Firefly.
- Built the retrieval backend behind ReCon: a FAISS index over CLIP concept embeddings mapped to cached denoising trajectories, with batched GPU precomputation of the trajectory store and top-*k* lookup on the serving path; prototyped as a drop-in inference path for Adobe Firefly.