Cuong Nguyen

website: www.cuongnd.com email: n.cuongg@gmail.com phone: (503) 706-9932

OVERVIEW

I build agentic AI systems that turn frontier model capabilities into structured, editable, and reliable creative workflows. I am currently building my startup Scrambo, where I'm tackling the hard problem of making multi-agent systems reliable enough to be creative companions.

Previously, I was a Senior Research Scientist at Adobe Research, working at the intersection of HCI and applied AI. My research led to multiple successful technology transfers, including VR video-editing tools in Premiere Pro and After Effects that received a Technology & Engineering Emmy Award in 2025, and two motion-graphics features in Adobe Express. I have published more than 25 papers at top HCI conferences and filed 39 patents, 19 issued.

RELEVANT FOCUS

- Creative agentic systems: develop agents capable of proposing alternatives for creative exploration and exposing their structure for inspection, comparison, and conversational refinement.
- Reliable agent workflow: build deterministic verifiers grounded in what users consider correct and usable, and design generate–validate–repair loops that turn unreliable generations into dependable, executable outputs.
- Evaluation & post-training: held-out benchmarks for agentic systems (loop pass rate, execution success), synthetic SFT data, QLoRA fine-tuning in PyTorch, teacher distillation.
- Production AI system building with Python, TypeScript, Modal, Hugging Face, and frontier LLM APIs.

EXPERIENCE

Founder Scrambo
www.scrambo.dev *June 2026–present*
Building a multi-agent SDK that helps developers create reliable agentic video products. Own the design, architecture, and engineering end to end.

- **Harness for self-repair editing loop.** Build generate–validate–repair workflows that freeze accepted timeline decisions, isolate failures, and provide repair agents with structured state, targeted feedback, and domain-specific tools so iteration does not regress completed work.
- **Customizable evaluation layers.** Combine rubric-based LLM/VLM judges for perceptual questions with deterministic validators for fast, auditable constraints.
- **Harness-informed model improvement.** Turn failures and successful repairs captured by Scrambo into training data, allowing smaller models to learn repair strategies that would otherwise require repeated orchestration. Used a Gemini teacher model to generate examples of successful video captions and step-by-step repairs based on verifier feedback, then used those examples to QLoRA-tune Qwen3-VL-8B. In a 100-task held-out pilot, the tuned model completed 71% of tasks within four attempts, compared with 13% for base Qwen and 86% for teacher.
- **Multi-agent orchestration.** Architect agent teams that hand off cleanly across long, error-prone video pipelines, with structured context, tool-error recovery, and clear seams for embedding domain expertise.

Senior Research Scientist Adobe Research
601 Townsend St, San Francisco, CA 94103 *30 Oct 2017–June 2026*
Key roles

- AI product research and engineering: Invented model-assisted authoring experiences guided by user-centered HCI principles. Led agentic motion-graphics research where LLMs generated, revised, and reasoned over structured motion programs while preserving user control and predictable editing behavior.
- Prompt, context, and evaluation design: Develop prompts, structured representations, prototype workflows, user studies, and quantitative metrics to understand model behavior and assess launch readiness across output quality, execution reliability, performance, and user experience.
- Product launch support: Drive technology transfer with product, design, and engineering teams by building proof-of-concept prototypes directly in product codebases, adapting AI models and commercial LLM APIs for production constraints, and collaborating on performance, reliability, UX, and maintainability goals.

- Knowledge transfer and mentoring: Organize research seminars for the HCI community within Adobe, serve on the Immersive technology patent committee, document research and evaluation methods for partner teams, and supervise interns from top research universities.

Notable product impacts

- *Animate All: Auto animate design documents*. Tech transfer shipped in Adobe Express in 2024. Increased animation usage by approximately 30% post-launch, reaching approximately 1.8M monthly uses as of October 2025.
- *Face-aware video captioning*. Tech transfer shipped in Adobe Express in 2023, bringing content-aware video understanding into an end-user creative workflow.
- *Project Pronto*. Adobe MAX Sneaks live demo in 2019 demonstrating rapid AR authoring workflows that translated research prototypes into product-facing concepts.
- *360 video editing in VR*. Tech transfer shipped in Premiere Pro and After Effects. Received the Technology & Engineering Emmy Award in 2025.

Research Intern Adobe Research
San Francisco, CA *06 June 2016–09 Dec 2016*
Worked with Stephen DiVerdi and Aaron Hertzmann in two back-to-back research internships to explore novel interfaces for Virtual Reality.

- Accepted paper and best paper honorable mention award at ACM CHI 2017 for Vremiere, and accepted paper at ACM UIST 2017 for CollaVR.
- Selected to demo at Adobe MAX 2016.
- Media coverage from CNET, UploadVR, and RoadToVR.
- Became the foundation for Premiere Pro's VR features, which granted the Premiere Pro team the Red Shark industry award.

Graduate Research Assistant Portland State University
Portland, OR *15 Sep 2012–16 Sep 2017*
Researched and developed new interfaces and techniques to enable novel video interaction experiences.

Notable projects: *Video Summagator*, a 3D interface that lets users quickly understand video content through interactive volume visualization; *Responsive Software Tutorial*, a video player for software tutorials that lets users directly interact with the software shown in the video to quickly locate meaningful content; and *GazeNoter*, a gaze-based video player for lecture video that tracks user gaze to highlight interesting lecture notes and automatically control video playback.

SELECTED AI SYSTEMS

Visual Lyrics: Built a human-agent system for expressive, audio-driven animated typography. Multiple LLM agents generate animated text and backgrounds as web code, while the interface exposes the structure for inspection and conversational refinement. In a 10-participant user study, users created outputs they found expressive and novel and responded positively to the system's support for creative iteration.

MotionCoder: Designed a generation pipeline for high-quality, editable motion effects using SVG-based visual composition, retrieval over curated motion-code examples, and code analysis that exposes meaningful parameters for real-time iteration. Evaluated on 42 curated prompts using both execution checks and an LLM judge that measured whether the rendered animation expressed the requested motion; retrieval substantially improved motion-alignment scores over zero-shot generation.

DepthScape: Led research on a proactive design agent for depth-based composition. The system analyzes an input image, identifies salient objects and surfaces, and uses code generation, a small DSL, and in-context examples to create parametric 3D anchors for spatial design. In evaluation, it generated at least one parametric anchor for 97 of 100 diverse stock images, with 89% of generated programs compiling into interactive anchors.

PUBLICATIONS

Role-Aware Virtual Agents for Navigational Interaction guided by a Multimodal Large Language Model
Minyoung Kim, Changyang Li, Cuong Nguyen, Lap-Fai Yu
SIGGRAPH 2026

DepthScape: Authoring 2.5D Designs with Depth Estimation, Semantic Understanding, and Geometry Extraction
Xia Su, Cuong Nguyen, Matheus Gadelha, Jon E. Froehlich

In Proceedings of ACM DIS 2026

Visual Lyrics: Automatic Generation of Animated Lyric Videos
David Chuan-En Lin, Cuong Nguyen, Valentina Shin, Nikola Martelaro
In Proceedings of ACM IUI 2026. Project webpage

SweeperBot: Making 3D Browsing Accessible through View Analysis and Visual Question Answering
Chen Chen, Cuong Nguyen, Alexa Siu, Ding Li, Nadir Weibel
International Journal of Human-Computer Interaction 2025

MemoVis: A GenAI-Powered Tool for Creating Companion Reference Images for 3D Design Feedback
Chen Chen, Cuong Nguyen, Thibault Groueix, Vladimir G. Kim, Nadir Weibel
ACM TOCHI. Presented at UIST 2024

PaperToPlace: Transforming Instruction Documents into Spatialized and Context-Aware Mixed Reality Experiences
Chen Chen, Cuong Nguyen, Jane Hoffswell, Jennifer Healey, Trung Bui, Nadir Weibel
In Proceedings of UIST 2023

GestureCanvas: Prototyping Compound Freehand Interaction in VR through Programming by Demonstration
Anika Sayara, Emily Lynn Chen, Cuong Nguyen, Robert Xiao, Dongwook Yoon
In Proceedings of UIST 2023

PoseVEC: Authoring Adaptive Pose-aware Effects Using Visual Programming and Demonstrations
Yongqi Zhang, Cuong Nguyen, Rubaiat Habib Kazi, Lap-Fai Yu
In Proceedings of UIST 2023

VideoDoodles: Hand-Drawn Animations on Videos with Scene-Aware Canvases
Emilie Yu, Kevin Blackburn-Matzen, Cuong Nguyen, Oliver Wang, Rubaiat Habib Kazi, Adrien Bousseau
ACM Transactions on Graphics, SIGGRAPH 2023

Rapido: Prototyping Interactive AR Experiences through Programming by Demonstration
Germán Leiva, Jens Emil Sloth Grønbæk, Clemens Nylandsted Klokmose, Cuong Nguyen, Paul Asente, Rubaiat Habib Kazi
In Proceedings of UIST 2021

Pronto: Rapid Augmented Reality Video Prototyping Using Sketches and Enaction
Germán Leiva, Cuong Nguyen, Rubaiat Habib Kazi, Paul Asente
In Proceedings of CHI 2020

TutoriVR: A Video-based Tutorial System for Design Applications in Virtual Reality
Balasaravanan Thoravi Kumaravel, Cuong Nguyen, Stephen DiVerdi, Bjoern Hartmann
In Proceedings of CHI 2019

Vremiere: In-Headset Virtual Reality Video Editing
Cuong Nguyen, Stephen DiVerdi, Aaron Hertzmann, Feng Liu
In Proceedings of CHI 2017. Best Paper Honorable Mention Award

Making Software Tutorial Video Responsive
Cuong Nguyen and Feng Liu
In Proceedings of CHI 2015. Best Paper Honorable Mention Award

EDUCATION

Portland State University 2012–2017
Ph.D. in Computer Science Advised by Dr. Feng Liu

Portland State University 2010–2012
Bachelor of Science in Computer Science

Ho Chi Minh City University of Science 2007–2009
Bachelor of Science in Computer Science

ISSUED PATENTS

P8804-US	Dynamically Rendering 360-Degree Videos Using View-Specific-Filter Parameters
P6845-GB	Collaborative Review of Virtual Reality Video
P6606-US	Facilitating editing of virtual-reality content using a virtual-reality headset
P6845-US2	Collaborative Virtual Reality Anti-Nausea and Video Streaming Techniques
P6845-US1	Collaborative Interaction with Virtual Reality Video
P7511-US	Dynamically Modifying Virtual and Augmented Reality Content to Reduce Depth Conflict Between User Interface Elements and Video Content
P8321-US	Interfaces and Techniques to Retarget 2D Screencast Videos into 3D Tutorials in Virtual Reality
P9279-US	Augmented Video Prototyping
P8425-US	Selecting Objects Within A Three-Dimensional Point Cloud Environment
P8321-US-DIV1	3D Simulation of a 3D Drawing in Virtual Reality
P10088-US	Systems for Augmented Reality Authoring of Remote Environments
P11028-US	Responsive Video Canvas Generation

HONORS & AWARDS

Technology & Engineering Emmy Award 2025
Best Paper Honorable Mention Award, ACM CHI 2017
Best Paper Honorable Mention Award, ACM CHI 2015
Best Paper Honorable Mention Award, ACM CHI 2013

ACTIVITIES

Invited talks

- *Building GenAI-Powered Tools for Editable Visual Effects*. Stanford University, November 2025. Details.
- *Computational text-based interactions for spatial creativity*. New York University, April 2025.
- *On-device creative authoring for virtual and augmented reality*. George Mason University's Distinguished Lecture Series, November 2020; Yale University, April 2021; Portland State University, April 2021.
- *Project Pronto: Augmented Reality Video Prototyping*. University of British Columbia, April 2021.
- *Designing In-Headset Authoring Tools for Virtual Reality Video*. UC Berkeley, November 2018; Stanford University, November 2018.

Service

- Associate Chair for ACM CHI 2020, ACM CHI 2022, and ACM UIST 2024.
- Reviewer for IEEE VR, ACM UIST, ACM ISS, ACM CHI, ACM Multimedia, IEEE Transactions on Multimedia, Information Visualization, and ACM SIGGRAPH Asia.

INTERNS SUPERVISED

Bala Kumaravel, UC Berkeley; Roberto Antonio Montano Murillo, University of Sussex; Megha Nawhal, Simon Fraser University; German Leiva, Université Paris-Saclay; Jeremy Hartmann, University of Waterloo; Zach Wang, Yale University; Rawan Alghofaili, George Mason University; Yongqi Zhang, George Mason University; Mary Zhu, Stanford University; Cory Ilo, Virginia Tech; Chen Chen, UC San Diego; Claudia Osorio, University of Southern California; David Lin, Carnegie Mellon University; Xia Su, University of Washington; Yunxiang Zhang, New York University.

SKILLS

- **Languages and systems**: Python, TypeScript, web applications, backend prototypes, product-codebase prototyping.
- **AI workflows**: OpenAI and Gemini APIs, prompt engineering, context design, retrieval over curated examples, code generation, small DSLs, model fine-tuning.
- **Code and program tooling**: code generation, code analysis, small DSLs, automated generate–validate–repair loops, executable-output verification.
- **Post-training and evaluation**: QLoRA SFT, teacher distillation, verifier-guided generation, executable-output testing, LLM-based evaluations, quantitative metrics, failure analysis, launch-readiness assessment.
- **Collaboration**: Research-to-product technology transfer, cross-functional product engineering, technical communication, evaluation documentation, intern mentorship, team enablement.