Pinar Yanardag

Assistant Professor of Computer Science, Virginia Tech · Director, GEM Lab
pinary@vt.edu · pinguar.org · gemlab-vt.github.io · Google Scholar

Generative models can imagine almost anything. I make them controllable.

My research focuses on control, alignment, and personalization in generative AI, with an emphasis on efficient long-horizon video generation and interactive world models. I lead the Generative Modeling Lab (GEM Lab) at Virginia Tech. Since Fall 2023, my lab has produced 25+ papers at CVPR, NeurIPS, ICCV/ECCV and AAAI. My work is supported by an NSF CAREER award on generative cinematography and by industry partners including Google, Adobe, Snap and Qualcomm. I serve as an Area Chair for CVPR, ICCV and ECCV, and was selected as an Outstanding Area Chair for CVPR 2026.

Academic appointments

Virginia Tech — Aug 2023 – present
Tenure-track Assistant Professor, Department of Computer Science · Director, Generative Modeling Lab (GEM Lab)

Boğaziçi University, Istanbul, Turkey — 2019 – 2022
Visiting Professor, Department of Computer Engineering · Director, Creative AI Technologies Lab (CATLAB)

Massachusetts Institute of Technology — 2016 – 2018
Postdoctoral Research Associate, MIT Media Lab, Cambridge, Massachusetts

Education

Purdue University — Aug 2010 – May 2016
Ph.D. in Computer Science · Advisors: SVN Vishwanathan (primary) and Jennifer Neville (secondary)

Boğaziçi University — Feb 2009 – Jun 2010
M.Sc. in Computer Engineering · Advisor: Ethem Alpaydın

Çanakkale Onsekiz Mart University — Sep 2003 – Jun 2007
B.S. in Computer Engineering

Awards & honors

2026	NSF CAREER Award
2026	Outstanding Area Chair, CVPR
2025	Best Paper Award, Visual Concepts Workshop at CVPR
2022	Best Paper Honorable Mention, ACM Creativity & Cognition
2021	Best Paper Award, CTRL-GEN Workshop at NeurIPS
2020	Primetime Emmy nomination for Creative Director in Outstanding Main Title Design
2019	TÜBİTAK 2232 Outstanding Research Fellow
2019	Finalist, MIT $15K Creative Arts Competition
2018	Technology Frontiers Award, MIT Breastfeeding Hackathon
2010	Fulbright Ph.D. Fellowship
2010	Purdue Graduate Tuition Scholarship
2010	Google Anita Borg Memorial Scholarship, EMEA
2007	Linux Users Association of Turkey Annual Award

Publications

2026

NeurIPS '26 **VideoMLA: Low-Rank Latent KV Cache for Minute-Scale Autoregressive Video Diffusion**
Hidir Yesiltepe, Jiazhen Hu, Tuna Han Salih Meral, Adil Kaan Akan, Kaan Oktay, Hoda Eldardiry, **Pinar Yanardag**
Advances in Neural Information Processing Systems (NeurIPS), 2026.

NeurIPS '26 **From Zero to Hero: Training-Free Custom Concept Spawning in World Models**
Kiymet Akdemir, **Pinar Yanardag**
Advances in Neural Information Processing Systems (NeurIPS), 2026.

NeurIPS '26 **AdaState: Self-Evolving Anchors for Streaming Video Generation**
Yusuf Dalva, **Pinar Yanardag**
Advances in Neural Information Processing Systems (NeurIPS), 2026.

NeurIPS '26 **HiPhy: Hierarchical Alignment for Physically-Plausible Multi-Principle Video Generation**
Tahira Kazimi, Shubhankar Borse, Munawar Hayat, Fatih Porikli, **Pinar Yanardag**
Advances in Neural Information Processing Systems (NeurIPS), 2026.

NeurIPS '26 **RAVEL: Rare Concept Generation and Editing via Graph-driven Relational Guidance**
Kavana Venkatesh, Yusuf Dalva, Ismini Lourentzou, **Pinar Yanardag**
Advances in Neural Information Processing Systems (NeurIPS), 2026.

ECCV '26 **Learn Once, Edit Anywhere: Visual Direction Transfer for Diffusion Models**
Yusuf Dalva, Hidir Yesiltepe, **Pinar Yanardag**
European Conference on Computer Vision (ECCV), 2026.

ECCV '26 **RankT2I: A Submodular Framework for Discovering Interpretable and Diverse Semantics in Text-to-Image Models**
Ritika Allada, **Pinar Yanardag**
European Conference on Computer Vision (ECCV), 2026.

CVPR '26 **∞-RoPE: Action-Controllable Infinite Video Generation Emerges From Autoregressive Self-Rollout**
Hidir Yesiltepe, Tuna Han Salih Meral, Adil Kaan Akan, Kaan Oktay, **Pinar Yanardag**
Conference on Computer Vision and Pattern Recognition (CVPR), 2026.

CVPR '26 **Diverse Video Generation with Determinantal Point Process-Guided Policy Optimization**
Tahira Kazimi, Connor Dunlop, **Pinar Yanardag**
Conference on Computer Vision and Pattern Recognition (CVPR), 2026.

CVPR '26 **DTG-Restore: Training-Free Diffusion Refinement for Generative Video Super-Resolution**
Hidir Yesiltepe, Koutilya PNVR, Gaurav Pathak, Navaneeth Bodla, Bharat Singh, **Pinar Yanardag***, Jinrong Xie*
Conference on Computer Vision and Pattern Recognition (CVPR), 2026. *Equal supervision.

CVPR '26 Findings **Harmonized Multi-Layer Text-to-Image Generation with Generative Priors**
Yusuf Dalva, Yijun Li, Qing Liu, Nanxuan Zhao, Jianming Zhang, Zhe Lin, **Pinar Yanardag**
Conference on Computer Vision and Pattern Recognition (CVPR) Findings, 2026.

SIGGRAPH '26 **Canvas2Image: Compositional Image Generation through Multi-Task Canvas**
Yusuf Dalva, Guocheng Gordon Qian, Maya Goldenberg, Tsai-Shien Chen, Kfir Aberman, Sergey Tulyakov, **Pinar Yanardag**, Kuan-Chieh Jackson Wang
SIGGRAPH, 2026.

AAAI '26 **MotionFlow: Attention-Driven Motion Transfer in Video Diffusion Models**
Tuna Han Salih Meral, Hidir Yesiltepe, Connor Dunlop, **Pinar Yanardag**
AAAI Conference on Artificial Intelligence (AAAI), 2026.

AAAI '26 **Plot'n Polish: Zero-shot Story Visualization and Disentangled Editing with Text-to-Image Diffusion Models**
Kiymet Akdemir, Jing Shi, Kushal Kafle, Brian Price, **Pinar Yanardag**
AAAI Conference on Artificial Intelligence (AAAI), 2026.

2025

NeurIPS '25 **Personalized Image Editing in Text-to-Image Diffusion Models via Collaborative Direct Preference Optimization**
Connor Dunlop, Matthew Zheng, Kavana Venkatesh, **Pinar Yanardag**
Advances in Neural Information Processing Systems (NeurIPS), 2025.

NeurIPS '25 **LoRAShop: Training-Free Multi-Concept Image Generation and Editing with Rectified Flow Transformers**
Yusuf Dalva, Hidir Yesiltepe, **Pinar Yanardag**
Advances in Neural Information Processing Systems (NeurIPS), 2025. Spotlight.

NeurIPS '25 **Dynamic View Synthesis as an Inverse Problem**
Hidir Yesiltepe, **Pinar Yanardag**
Advances in Neural Information Processing Systems (NeurIPS), 2025.

NeurIPS '25 **CREA: A Collaborative Multi-Agent Framework for Creative Image Editing and Generation**
Kavana Venkatesh, Connor Dunlop, **Pinar Yanardag**
Advances in Neural Information Processing Systems (NeurIPS), 2025.

ICCV '25 **CLoRA: A Contrastive Approach to Compose Multiple LoRA Models**
Tuna Han Salih Meral, Enis Simsar, Federico Tombari, **Pinar Yanardag**
International Conference on Computer Vision (ICCV), 2025. Highlight.

ICCV '25 **LoRAverse: A Submodular Framework to Retrieve Diverse Adapters for Diffusion Models**
Mert Sonmezer, Matthew Zheng, **Pinar Yanardag**
International Conference on Computer Vision (ICCV), 2025.

ICML '25 **ConceptAttention: Diffusion Transformers Learn Highly Interpretable Features**
Alec Helbling, Tuna Han Salih Meral, Ben Hoover, **Pinar Yanardag**, Duen Horng (Polo) Chau
International Conference on Machine Learning (ICML), 2025. Oral. Best Paper Award, Visual Concepts Workshop at CVPR 2025.

CVPR '25 **FluxSpace: Disentangled Semantic Editing in Rectified Flow Transformers**
Yusuf Dalva, Kavana Venkatesh, **Pinar Yanardag**
Conference on Computer Vision and Pattern Recognition (CVPR), 2025.

CVPR '25 **Explaining in Diffusion: Explaining a Classifier through Hierarchical Semantics with Text-to-Image Diffusion Models**
Tahira Kazimi, Ritika Allada, **Pinar Yanardag**
Conference on Computer Vision and Pattern Recognition (CVPR), 2025.

CVPR '25 **LoRACLR: Contrastive Adaptation for Customization of Diffusion Models**
Enis Simsar, Thomas Hofmann, Federico Tombari, **Pinar Yanardag**
Conference on Computer Vision and Pattern Recognition (CVPR), 2025.

2024

NeurIPS '24 D&B **Stylebreeder: Exploring and Democratizing Artistic Styles through Text-to-Image Models**
Matthew Zheng, Enis Simsar, Hidir Yesiltepe, Federico Tombari, Joel Simon, **Pinar Yanardag**
Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track, 2024.

CVPR '24 **NoiseCLR: A Contrastive Learning Approach for Unsupervised Discovery of Interpretable Directions in Diffusion Models**
Yusuf Dalva, **Pinar Yanardag**
Conference on Computer Vision and Pattern Recognition (CVPR), 2024. Oral.

CVPR '24 **CONFORM: Contrast is All You Need For High-Fidelity Text-to-Image Diffusion Models**
Tuna Han Salih Meral, Enis Simsar, Federico Tombari, **Pinar Yanardag**
Conference on Computer Vision and Pattern Recognition (CVPR), 2024.

CVPR '24 **RAVE: Randomized Noise Shuffling for Fast and Consistent Video Editing with Diffusion Models**
Ozgur Kara, Bariscan Kurtkaya, Hidir Yesiltepe, James M. Rehg, **Pinar Yanardag**
Conference on Computer Vision and Pattern Recognition (CVPR), 2024. Highlight.

ICCC '24 **ORACLE: Leveraging Mutual Information for Consistent Character Generation with LoRAs in Diffusion Models**
Kiymet Akdemir, **Pinar Yanardag**
International Conference on Computational Creativity (ICCC), 2024.

Prior to Virginia Tech

WACV '23 **Fantastic Style Channels and Where to Find Them: A Submodular Framework for Discovering Diverse Directions in GANs**
Enis Simsar, Umut Kocasari, Ezgi Gulperi Er, **Pinar Yanardag**
Winter Conference on Applications of Computer Vision (WACV), 2023.

WACV '23 **Text and Image Guided 3D Avatar Generation and Manipulation**
Zehranaz Canfes, Furkan Atasoy, Alara Dirik, **Pinar Yanardag**
Winter Conference on Applications of Computer Vision (WACV), 2023. Spotlight.

ECCV '22 **FairStyle: Debiasing StyleGAN2 with Style Channel Manipulations**
Cemre Efe Karakas, Alara Dirik, Eylül Yalcinkaya, **Pinar Yanardag**
European Conference on Computer Vision (ECCV), 2022.

ACM C&C '22 **MIDISpace: Finding Linear Directions in Latent Space for Music Generation**
Meliksah Turker, Alara Dirik, **Pinar Yanardag**
ACM Creativity & Cognition, 2022. Best Paper Honorable Mention.

WACV '22 **StyleMC: Multi-Channel Based Fast Text-Guided Image Generation and Manipulation**
Umut Kocasari, Alara Dirik, Mert Tiftikci, **Pinar Yanardag**
Winter Conference on Applications of Computer Vision (WACV), 2022.

ICCV '21 **LatentCLR: A Contrastive Learning Approach for Unsupervised Discovery of Interpretable Directions**
Oguz Kaan Yuksel, Enis Simsar, Ezgi Gulperi Er, **Pinar Yanardag**
International Conference on Computer Vision (ICCV), 2021.

NeurIPS WS '21 **Controlled Cue Generation for Play Scripts**
Alara Dirik, Hilal Donmez, **Pinar Yanardag**
Controllable Generative Modeling in Language and Vision (NeurIPS Workshop), 2021. Best Paper Award.

ICCC '21 **Nightmare Machine: A Large-scale Study to Induce Fear using Artificial Intelligence**
Pinar Yanardag, Nick Obradovich, Manuel Cebrian, Iyad Rahwan
International Conference on Computational Creativity (ICCC), 2021.

ACM C&C '21 **Shelley: A Crowd-sourced Collaborative Horror Writer**
Pinar Yanardag, Manuel Cebrian, Iyad Rahwan
ACM Creativity & Cognition, 2021.

EMNLP '16 **WordRank: Learning Word Embeddings via Robust Ranking**
Shihao Ji, Hyokun Yun, **Pinar Yanardag**, Shin Matsushima, SVN Vishwanathan
Conference on Empirical Methods in Natural Language Processing (EMNLP), 2016.

NeurIPS '15 **A Structural Smoothing Framework For Robust Graph Comparison**
Pinar Yanardag, SVN Vishwanathan
Advances in Neural Information Processing Systems (NeurIPS), 2015.

KDD '15 **Deep Graph Kernels**
Pinar Yanardag, SVN Vishwanathan
ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2015.

CLOUD '14

Crowdsourced Resource Sizing of Virtual Appliances
Pinar Yanardag, Rean Griffith, Anne Holler, K. Shankari, Xiaoyun Zhu, Ravi Soundararajan, Adarsh Jagadeeshwaran, Pradeep Padala
IEEE International Conference on Cloud Computing (IEEE CLOUD), 2014.

PATENT

US Patent

Crowd-sourced Operational Metric Analysis of Virtual Appliances
Rean Griffith, Anne Holler, Shankari Kalyanaraman, Xiaoyun Zhu, Vijayaraghavan Soundararajan, Adarsh Jagadeeshwaran, Pradeep Padala, **Pinar Yanardag** Delul
U.S. Patent 10,432,699 B2, granted October 1, 2019.

Research funding

RESEARCH GRANTS

National Science Foundation – CAREER (CISE/RI) 2026 · $562,245
Toward Generative Cinematography with Controllable Video Synthesis · PI · Award no. 2543524

Children's National Hospital 2024 · $25,000
Controlled Synthesis of Training Data for Syndromic Diseases for Children · PI

Turkish NSF (TÜBİTAK) 2019 – 2022 · $450,000
Generative Adversarial Networks for Medical Image Synthesis · Sole PI

INDUSTRY GIFTS

The lab has received over $500K in gift funds from industry partners to support its research.

fal AI 2026
Reinforcement Learning for Improved Video Synthesis · Sole PI

Qualcomm Research 2025
Physics-Guided Video Synthesis · Sole PI

fal AI 2025
Autoregressive Video Model Distillation · Sole PI

Snap Research 2025
Compositional Image Generation with Multimodal Controls · Sole PI

Deloitte Research 2024
Contrastive Mechanistic Interpretability for LLM Steerability · Sole PI

Google Research 2024
Test-time Personalization of T2I Diffusion Models · Sole PI

Adobe Research 2024
Harmonized Multi-Layer Text-to-Image Generation with Generative Priors · Sole PI

UNIVERSITY & INSTITUTIONAL AWARDS

College of Agriculture and Life Sciences, Virginia Tech 2025 · $75,000
Transforming Plant Disease Diagnosis in Virginia through Generative AI and Hyperspectral Sensing · Co-PI

Commonwealth Cyber Initiative (CCI) 2025 · $100,000
GenA(eye)ris: Security for Generative Biometric Models · Co-PI

College of Engineering Major Grants Initiative, Virginia Tech 2025 · $125,000
Advancing Generative AI for Intelligent XR · Co-PI

Institute of Creativity & Arts 2025 · $40,000
Reimagining Filmmaking with Generative AI · Co-PI

Center for Advanced Innovation in Agriculture 2024 · $125,000
Generative AI for Plant Stress Phenotypes · Co-PI

Institute for Critical Technology and Applied Science 2024 · $15,000
Generative AI for Improved Hyperspectral Imaging in Agriculture · PI

Advising & mentoring

CURRENT DOCTORAL STUDENTS, VIRGINIA TECH

Hidir Yesiltepe
Tuna Han Salih Meral
Kiymet Akdemir
Yusuf Dalva
Tahira Kazimi
Connor Dunlop

DOCTORAL CO-ADVISING

Enis Simsar, Ph.D. student, ETH Zurich; co-advised with Thomas Hofmann and Federico Tombari.

FORMER GRADUATE RESEARCH ADVISEES

Ritika Allada, master's thesis, 2025, Virginia Tech · Matthew Zheng, master's thesis, 2025, Virginia Tech · Alara Dirik, master's thesis, 2022 (co-advisor), Boğaziçi University

FORMER UNDERGRADUATE RESEARCH ADVISEES, BOĞAZIÇI UNIVERSITY

Enis Simsar · ETH Zurich
Kerem Zaman · UNC Chapel Hill
Merve Rabia Barin · Purdue
Oguz Kaan Yuksel · EPFL
Mert Yuksekgonul · Stanford
Berkay Doner · EPFL
Dilara Gokay · DeepMind
Eylul Yalcinkaya · Meta
Elif Sema Balcioglu · EPFL
Umut Kocasari · TUM
Cemre Efe Karakas · Amazon
Ramiz Dundar · Brown

Teaching experience

Virginia Tech – Department of Computer Science 2023 – present
Instructor: Learning-based Computer Vision (graduate) · Fantastic Generative AI Models and Where to Find Them (graduate) · Creative AI Capstone (undergraduate) · Entrepreneurship with Generative AI Capstone (undergraduate)

Boğaziçi University – Computer Engineering 2019 – 2023
Instructor: Advanced Topics in Deep Learning (graduate) · How to Generate (Almost) Anything (graduate)

London College of Fashion, University of the Arts London Feb – Apr 2020
Instructor: AI & Fashion. Student projects were showcased at London Fashion Week 2020.

Massachusetts Institute of Technology Sep 2016
Teaching Assistant: MAS.S67 Machine Learning, Society & Autonomy. Prepared and delivered deep-learning tutorials for more than 50 students; held office hours and graded projects.

Purdue University – Department of Statistics
Co-instructor: STAT598Z Introduction to Computing for Statisticians, with SVN Vishwanathan. Statistical machine learning with an emphasis on optimization.

Selected talks

INVITED TALK **Rethinking Memory: From Long-Horizon Video to Interactive World Models** Oct 2026
Georgia Tech, School of Computational Science & Engineering

KEYNOTE **The Art of Taming Video Models: Leveraging Internal Representations for Training-Free Control** Jul 2026
From Frames to Stories Workshop, ICML

KEYNOTE **The Art of Taming Generative AI: From Prompting to Creative Direction** Jun 2026
AI for Visual Arts Workshop, CVPR

KEYNOTE **Leveraging Hidden Priors for Training-free Control: From Long-Horizon Video to Interactive World Models** Jun 2026
AI-assisted Long Video Creation Workshop, CVPR

KEYNOTE **The Other Half of Personalization: From Customizing Concepts to Personalizing Taste** Jun 2026
Personalization in Generative AI Workshop, CVPR

INVITED TALK **Foundations of Controllable and Personalized Generative AI** Apr 2026
University of Michigan, Ann Arbor

INVITED TALK **The Next Generation of Generative AI** 2025
Virginia Academy of Science, Engineering and Medicine

INVITED TALK **Creative AI** 2024
Tech Girls Summer Camp and Bee VT Summer Camp, Blacksburg, Virginia

INVITED TALK **Interpretable Deep Generative Models** 2022
MIT EECS Computer Vision Group

INVITED TALK **Interpretable Deep Generative Models** 2022
Carnegie Mellon University, Robotics Institute

INVITED TALK **Human-AI Collaboration: The New Art Duo** 2019
NVIDIA GTC Conference, San Jose, California

INVITED TALK **How to Generate (Almost) Anything** 2019
Frontiers in Science, Technology, and the Arts, MIT

KEYNOTE **Humans and Machines Stronger Together** 2019
North American Chilean Institute, Santiago, Chile

Professional service

Conference leadership	Area Chair: CVPR, ICCV, ECCV, AAAI, WACV. Lead Organizer: Personalization in Generative AI Workshop, ICCV 2025 and CVPR 2026. Co-organizer: Responsible AI Workshop, CVPR 2025.
Research evaluation	Panelist, NSF CISE Robust Intelligence Program.
University service	Virginia Tech: Faculty Search Committee; Ph.D. Qualifier Exam Committee; Faculty Senate; Publicity Committee; Research Incentive Committee.
Community mentoring	Google Summer of Code Mentor, 2008 and 2009.

Industry & earlier research experience

AI Fiction · Founder and CEO Jul 2018 – 2022
Led an AI studio developing projects for film and entertainment, including work on the Westworld main titles nominated for a 2020 Primetime Emmy.

GLITCH · Co-founder Jan 2019 – Sep 2022
Co-founded a fashion venture combining generative adversarial networks with designer-developed products; finalist in the 2019 MIT $15K Creative Arts Competition.

Amazon · Machine Learning Scientist Intern, Search & Discovery (P13N), Seattle Jul – Sep 2015
Developed video recommendations and a deep-learning framework for e-book similarity.

VMware · R&D Intern, Distributed Resource Management, Palo Alto Jun – Aug 2014
Developed regression-based methods for vertical scaling of network resources; work published at RADIO 2014.

VMware · R&D Intern, Distributed Resource Management & Performance, Palo Alto Jun – Aug 2013
Developed machine-learning methods for virtual-machine sizing; work published at IEEE CLOUD 2014.

Purdue University · Graduate Research Assistant 2012 – 2016
Research appointments in Computer Science, Statistics, and the Office of the Vice President for Information Technology.

Boğaziçi University · Researcher, Perceptual Intelligence Lab Dec 2009 – Jul 2010
Built a web-based framework for running and comparing classification, dimensionality-reduction, and evaluation methods.

TÜBİTAK · Researcher and Team Leader Nov 2007 – Nov 2009
Core developer of Pardus GNU/Linux and lead of its security team; coordinated vulnerability fixes and security advisories with Linux vendors.

Selected projects

How to Generate (Almost) Anything · Creator Aug 2019
Led a Boston-based collective of researchers and artists collaborating on generative AI projects.

Deep Empathy · Co-creator Oct 2017
Collaborated with UNICEF to investigate whether AI-generated depictions of disaster-affected cities could foster empathy. Selected media: New Scientist, Fast Co Design, Digital Trends, Le Monde, Der Spiegel.

Shelley · Creator Oct 2017
Developed a recurrent-neural-network system for collaborative horror-story writing with online participants. Selected media: TIME, Washington Post, The Guardian, Huffington Post, CBS, New Scientist, Fortune.

Nightmare Machine · Creator Oct 2016
Conducted a large-scale study of fear elicited by AI-generated imagery; collected over one million votes from more than 100 countries. Selected media: Washington Post, The Atlantic, Forbes, BBC, NPR, Boston Globe, NBC, Vice.

Career notes

Three career breaks for maternity leave in 2018, 2021, and 2022, totaling two years.

Last updated: October 2026