

Jinen Setpal

+1 (765) 490-1435 | jinen@setpal.net | jinen.setpal.net | github.com/jinensetpal | dagshub.com/jinensetpal

EDUCATION

Purdue University

Aug. 2024 – Present

Doctor of Philosophy, Electrical & Computer Engineering

West Lafayette, IN, USA

Relevant Coursework: Optimization for Deep Learning, Machine Learning Theory, Computational Optimal Transport & Deep Generative Models, Real Analysis & Measure Theory, Numerical Linear Algebra, Random Variables & Signals, Trustworthy ML, Modern Differential Geometry

Purdue University

Aug. 2021 – May 2024

Bachelor of Science in Data Science

West Lafayette, IN, USA

Relevant Coursework:

- **Graduate Level:** Deep Learning, Advanced Topics in Reasoning with Large Language Models, Applied Regression Analysis, System Security
- **Undergraduate Level:** Data Mining & Machine Learning, Large Scale Data Analysis, Intro to AR/VR, Embedded Systems, Data Structures & Algorithms, Linear Algebra, Statistical Theory, Multivariate Calculus, Probability, Statistical Programming

PUBLICATIONS

Improving Feature Alignment in ConvNets using ContrastiveCAMs and Core-Focused Cross-Entropy

Sep. 2025

Under peer review.

Purdue University, USA

- Derived *ContrastiveCAMs*, an approach for computing feature attribution maps with strong theoretical correctness guarantees in convnets.
- Spurious Correlations observed with *ContrastiveCAMs* motivated an alignment-motivated constraint to ERM, called *Core-Constrained Risk Minimization* – this enforces only target-relevant regions (e.g., foreground from semantic segmentation) being used to inform predictions.
- Developed a classification-calibrated modification to cross-entropy as a surrogate for solving the constrained optimization objective, which learns to suppress non-target regions while extracting predictive performance.
- Reported state-of-the-art performance on robustness datasets: Hard-ImageNet, Spawrious; notable alignment improvements for general binary, multiclass & multilabel settings: Oxford IIIT-Pets, PASCAL VOC; as well as improvements as a pre-training backbone in downstream segmentation tasks.

BoilerBot: A Reliable Task-Oriented Chatbot Enhanced with Large Language Models

Oct. 2023

2nd Proceedings of Alexa Prize TaskBot (Alexa Prize 2023). Hu, Setpal, et al.

Purdue University, USA

- Fine-tuned 8-bit quantized large language models using LoRA for downstream tasks such as task title augmentation and patching failures within speech recognition.
- First-author on the grant proposal; earned \$250,000 in funding and an AWS account with unrestricted compute.
- Extended Amazon's COBOT (Conversational Bot) Toolkit, integrating custom APLs and logic modules for constraint-based state management.

CutLang V2: Advances in a runtime-interpreted analysis description language for HEP data

Jul. 2021

Frontiers in Big Data, 4, 27. Ünel, Sekmen et al.

CERN, Switzerland

- Developed Interpreter Functions through lexical analysis using Flex & Bison (.cpp).
- Setup CI/CD Scripts w/ Automated Email Delivery using GitHub Actions & SendGrid.

ArchiMeDe @ DankMemes: A New Model Architecture for Meme Detection

Dec. 2020

7th Evaluation Campaign, Final Workshop, EVALITA 2020. Setpal, Sarti

Turin, Italy

- Achieved .7664 F1-Score on test dataset (+.2466 baseline) w/ Video Presentation during final workshop.
- Built a multimodal ensemble using transfer learning by fine-tuning AlexNet, DenseNet & ResNet.

EMPLOYMENT

Machine Learning Engineer

Jun. 2022 – Aug. 2024

*DagsHub**Tel Aviv, Israel*

- Built generalized auto-labelling workflows integrating Label Studio with MLflow, which automates inference from any model registered on DagsHub enabling easy setup of active-learning pipelines.
- Developed intelligent prefetching, automatic path-column and datatype detection, data streaming and automated tensorization to enable training on data indexed within the Data Engine.
- Developed and deployed DPT: a conversational agent that enables users to interact with DagsHub documentation, and debug machine learning projects incorporating tools integrated within the DagsHub stack.
- Developed a data streaming client by monkeypatching Python's `open()` and extending FUSE to lazily pull files from a specified remote using DagsHub's web APIs.
- Built trainer integrations (automatic data, model, experiment and artifact logging) with the PyCaret framework, HuggingFace's Transformers library, and YOLOv8, to simplify adding observability to training workflows.
- Implemented and deployed open-source data science projects reproducing and extending past research. Examples: CheXNet, Panoptic Deeplab, YOLOv6.

Systems Developer

Sep. 2020 – Jul. 2021

*Teachiq AB / exam.net**Stockholm, Sweden*

- Packaged custom security implementations by forking open source `xmodmap(.c)` utility to a node.js module for exam.net's linux-based exam-delivery kiosk application.
- Reproducibly exploited the assessment kiosk on exam.net's web client and recommended mitigations.

FUNDED RESEARCH

Drone Video Object Recognition

Jan. 2022 – May 2023

*NSF Award 21204301 – PI: Prof. Yung-Hsiang Lu**Purdue University, USA*

- Team lead over the Spring 2023 Semester. Our team attempted to leverage Gazebo, ROS2 & the previous year's scoring function to develop a multi-agent reinforcement learning approach to the sample solution. Won 2nd Place for the **Undergraduate Research Expo Award** under Purdue College of Science.
- Developed an architecture for split-confidence resolution, achieving .9937 test accuracy as part of the reference solution made for the IEEE international autonomous UAV competition. Bootloader patching and setup for linux-based drones with OpenVINO accelerated IoT.

A Systematic Study of Cryptographic Function Identification Approaches in Binaries

Aug. 2021 – Dec. 2022

*NSF Award 2047991 – PI: Prof. Christina Garman**Purdue University, USA*

- Employed rudimentary techniques within NLP to establish a baseline approach for reconstructing cryptographic functions from disassembler code used to generate corresponding binaries.
- Evaluated current state-of-art classification tools against rigorous benchmark scripts. Currently under peer-review.

TRADE SECRETS

Semi-Supervised Class Activation Mappings for Target Localization & Super-Resolution

Sep. 2021 – Apr. 2022

*Final Presentation, TE AI Cup 2022. Setpal, et al.**TE Corporate, UK*

- Won the **Best Innovation Award**, developing subclassed TensorFlow layers for accurate, efficient prediction over classes with minute differences.
- By evaluating feature vectors from the model's penultimate convolutional layer over a dynamic weight threshold, we generate a bounding box to localize the region of the image critical to the final classification.

Leveraging Latent Features for Modular Multiclass Classification

Sep. 2021 – Apr. 2022

*Final Presentation, TE AI Cup 2022. Setpal, et al.**TE Corporate, UK*

- Designed & developed a novel modular, scalable architecture for classification achieving .99846 real-data classification accuracy over a +.2466 baseline.
- Implemented a latent feature aggregator network to enable minimal re-training for appending and removing target connectors from the multi-class classifier.

 CONFERENCE PRESENTATIONS

<i>The Machine Learning Angle for Open Source Science</i>	25 th Oct. 2023
<i>The Linux Foundation Member Summit (LFMS) 2023</i>	Monterey, CA, USA
<i>Interpretability Tools as Feedback Loops</i>	30 th Nov. 2022
<i>Toronto Machine Learning Summit (TMLS) 2022</i>	Toronto, Canada

 PROJECTS

Optimized torch.cdists Kernel using Triton	Jun. 2025
<i>Independent</i>	Lafayette, IN, USA
<ul style="list-style-type: none"> Developed a Triton Kernel with an identical API to <code>torch.cdists</code>, leveraging similarities with matmul operation semantics to significantly reduce both memory and runtime. Includes a backward kernel with correctness checks to facilitate training NNs; compatible with <code>torch.compile</code>. 	
[Re] Graph R-CNN for Scene Graph Generation	Sep. 2023 – Nov. 2023
<i>DagsHub × ML@Purdue Hackathon Fall 2023</i>	West Lafayette, IN, USA
<ul style="list-style-type: none"> Reproduced the Graph R-CNN for Scene Graph Generation paper as the template repository for the Scene Graph Generation Challenge in the beginner section of the DagsHub × ML@Purdue Hackathon. Included functions for model training, inference, registration and data processing for the VisualGenome dataset. 	
Time-Series Modelling for Outbreak Prediction	Oct. 2021
<i>CERN's The Port Hackathon</i>	West Lafayette, IN, USA
<ul style="list-style-type: none"> Predicted <i>oidium</i> outbreaks within vineyards in Germany. Achieved test accuracy of 0.995 (± 0.0025) when predicting outbreak risk, trained on daily data from 2013 - 2020. 	
Embedded Realtime Semantic Segmentation	Feb. 2021 – Apr. 2021
<i>Independent</i>	Mumbai, India
<ul style="list-style-type: none"> Embedded DeeplabV3+ with a MobileNetsv3 backbone to an android application. Established a data conversion pipeline (NV21 → YUV_420_888 → JPEG → Bitmap → TensorImage), with an inference framerate of $\approx 25fps$ on a Qualcomm SM8150 Snapdragon 855 (7 nm) processor. 	

 TECHNICAL SKILLS

Languages: Python, C, C++, Triton, x86 Assembly, Java, Kotlin, Bash, JavaScript, MATLAB, R, SQL, ROS2
Frameworks: PyTorch, JAX, TensorFlow, Keras, Pandas, Pillow, ROOT, Matplotlib, FUSE, Node.js, Express.js
Tools: Git, MLFlow, DVC, Docker, Radare2, Ghidra, TravisCI, GitGuardian, Kubernetes, Gazebo
Cloud Utilities: Google Cloud Console (Compute, Networking, Storage), Amazon Web Services (Redshift, ECR, ECS, S3, Sagemaker, CodePipeline, CodeCommit, CloudWatch, CloudFormation, Lambda), Azure Pipeline, GitHub Actions

 TEACHING

Graduate Teaching Assistant	Aug. 2025 – May 2026
<i>ECE 57000 – Artificial Intelligence @ Purdue University</i>	West Lafayette, IN, USA
<ul style="list-style-type: none"> Developed questions for the course midterms and final, as well as programming and presentation-based assignments. Held office hours. Enrollment: ~ 300 students. Developed a web application for LLM-based assignment grading, with a submission gallery and a ranking system. 	
Course Instructor	Aug. 2022/2023 – Dec. 2022/2023
<i>CS 39000 – Web Application Development @ Purdue University</i>	West Lafayette, IN, USA
<ul style="list-style-type: none"> Curriculum design and course instructor for a two-credit course. Net enrollment: 100 students. Covered HTML/CSS, JavaScript, React, Node.js, Express.js, MongoDB, Web Security & Cloud Hosting. 	
Undergraduate Teaching Assistant	Feb. 2022 – May 2022
<i>STAT 190 – Topics in Statistics for Undergraduates @ Purdue University</i>	West Lafayette, IN, USA
<ul style="list-style-type: none"> Lab instructor for Purdue's Corporate Partner MISO, developing industry solutions using Data Science. Graded assignments, held office hours, conducted code review. Taught classes on git, CI/CD & Data Mining. 	

TECHNICAL PRESENTATIONS

Sliced Wasserstein Distance, Wasserstein GANs <i>ML@Purdue Reading Group</i>	17 th Apr. 2025 <i>West Lafayette, IN, US</i>
Introduction to Optimal Transport <i>ML@Purdue Reading Group</i>	10 th Apr. 2025 <i>West Lafayette, IN, US</i>
Neural Tangent Kernel <i>ML@Purdue Reading Group</i>	6 th Mar. 2025 <i>West Lafayette, IN, US</i>
Sparse AutoEncoders <i>ML@Purdue Reading Group</i>	13 th Feb. 2025 <i>West Lafayette, IN, US</i>
Deepseek R1 <i>ML@Purdue Reading Group</i>	30 th Jan. 2025 <i>West Lafayette, IN, US</i>
Average Gradient Outer Product as a Mechanism for Deep Neural Collapse <i>ML@Purdue Reading Group</i>	24 th Oct. 2024 <i>West Lafayette, IN, US</i>
Deep Neural Collapse <i>ML@Purdue Reading Group</i>	10 th Oct. 2024 <i>West Lafayette, IN, US</i>
Group Transformation Invariance & Equivariance in CNNs & MLPs <i>ML@Purdue Reading Group</i>	26 th Sep. 2024 <i>West Lafayette, IN, US</i>
Attention is All You Need <i>ML@Purdue Reading Group</i>	3 rd Sep. 2024 <i>West Lafayette, IN, US</i>
Omnipredictors <i>ECE ML Reading Group</i>	25 th Apr. 2024 <i>West Lafayette, IN, USA</i>
Direct Preference Optimization <i>ML@Purdue Reading Group</i>	28 th Mar. 2024 <i>West Lafayette, IN, USA</i>
Towards Monosemanticity (SAEs) <i>ML@Purdue Reading Group</i>	8 th Feb. 2024 <i>West Lafayette, IN, USA</i>
A Mathematical Framework for Transformer Circuits <i>ML@Purdue Reading Group</i>	1 st Feb. 2024 <i>West Lafayette, IN, USA</i>
Deduplicating Training Data Makes Language Models Better <i>CS 592-LLM — Advanced Topics in Reasoning with Large Language Models</i>	2 nd Nov. 2023 <i>West Lafayette, IN, USA</i>
Groups, Rings & Fields <i>ML@Purdue Reading Group</i>	25 th Oct. 2023 <i>West Lafayette, IN, USA</i>
OOD Generalization via Risk Extrapolation <i>ML@Purdue Reading Group</i>	28 th Mar. 2023 <i>West Lafayette, IN, USA</i>
Enforcing Group-Transformation Invariance in MLPs <i>ML@Purdue Reading Group</i>	24 th Feb. 2023 <i>West Lafayette, IN, USA</i>
Interpretability Tools as Feedback Loops <i>Boilermake X</i>	21 st Jan. 2023 <i>West Lafayette, IN, USA</i>
AlphaTensor <i>SIGAI Reading Group</i>	15 th Nov. 2022 <i>West Lafayette, IN, USA</i>

WORKSHOPS

BlueDot's AI Alignment Fundamentals <i>AI Safety Purdue</i>	13 th Sep. 2024 – 12 th Nov. 2024 <i>West Lafayette, IN, USA</i>
Utilizing HPC with SLURM <i>Catapult Hacks</i>	31 st Mar. 2024 <i>West Lafayette, IN, USA</i>
Parameter-Efficient LLM Fine-Tuning <i>Purdue Hackers × ML@Purdue</i>	9 th Nov. 2023 <i>West Lafayette, IN, USA</i>
Intro to Data Engine <i>DagsHub × ML@Purdue Hackathon Fall 2023</i>	21 st Oct. 2023 <i>West Lafayette, IN, USA</i>
Intro to DagsHub <i>DagsHub × ML@Purdue Hackathon Fall 2023</i>	14 th Oct. 2023 <i>West Lafayette, IN, USA</i>

Training Word Embeddings

Purdue Hackers × ML@Purdue

Model Registry and Deployment with MLFlow

DagsHub

Experiment Tracking for Machine Learning with MLFlow

DagsHub

Linear Regression

ML@Purdue Spring 2023 Workshops

MLOps for Research Reproducibility

SIGAI Fall 2022 Workshops

14th Sep. 2023

West Lafayette, IN, USA

6th Apr. 2023

West Lafayette, IN, USA

24th Mar. 2023

West Lafayette, IN, USA

8th Mar. 2023

West Lafayette, IN, USA

6th Oct. 2022

West Lafayette, IN, USA

OUTREACH

Peer Reviewer

NeurReps Workshop @ NeurIPS 2025

Program Chair

DagsHub × Purdue Hackathon Fall 2023

Technical Advisor, Officer

ML@Purdue / SIGAI

Sep. 2025

San Diego, CA, USA

Oct. 2023 – Nov. 2023

West Lafayette, IN, USA

Jan. 2022 – Present

West Lafayette, IN, USA