Pengyu (Steven) Kan

Contact Information

Email : pkan2@jhu.edu | Cell : 608-421-8013 | Google Scholar : profile | https://pkan2.github.io

Education

Johns Hopkins University Sep. 2021 - Current
Ph.D. candidate in Computer Science.
Advisors : Prof. Kenichi Oishi, Prof. Craig Jones

University of Wisconsin, Madison Sep. 2016 - May 2020
Bachelor of Science in Computer Science, Mathematics and Economics. Certificate in Physics.
GPA : 3.99/4.0

My research lies at the intersection of deep learning and medical imaging, with a focus on brain MRI analysis. Specifically, I work on transformer-based models for brain age estimation and diffusion-based approaches for unsupervised neonatal lesion detection and inpainting. More broadly, my interests span multimodal representation learning, clinically informed model design, and improving the generalizability of medical imaging models. My long-term research vision is to advance methods for multimodal information representation through vision–language models and foundation models, and to leverage deep learning to bridge AI and clinical practice, with the goal of deepening our knowledge of brain cognition, neurodevelopment, and neurodegeneration.

Awards

Emerging Scholar Presenter, Best Poster Award, Distinct Abstract Award & Travel Award in American Neurological Association	2025
Phi Beta Kappa Honor Society	2019
UW-Madison Juli Plant Grainger Scholarship in Economics	2019
UW-Madison Undergraduate Scholarship for Summer Study	2018, 2019
UW-Madison Meek Bishop Scholarship in Economics	2018
Mathematical Contest in Modeling 2018 Honorable Mention	2018
Dean's List	2016 - 2019
College of Letters & Science General Scholarship	2017

Publications

[1] Clinically Interpretable Transformer Model Reveals Neurodevelopmental Delay in Preterm and NICU Infants Using Diffusion MRI
Yifan Shuai*, **Pengyu Kan***, Xingyu Chen, Kenichi Oishi
Accepted ISMRM 2026.

[2] OpenMAP-BrainAge: Generalizable and Interpretable Brain Age Predictor
Pengyu Kan, Craig Jones, Kenichi Oishi
Radiology Advances, 2026.

[3] TransMAD: Transformer-based Masked Adversarial Defenses to Patch Attacks
Pengyu Kan*, Laura Alexandra Daza*, Pablo Arbelaez, René Vidal
Pre-print.

[4] Towards Diffeomorphism Invariant Convolution Neural Networks
Pengyu Kan, Rudrasis Chakraborty, Vishnu Suresh Lokhande, Vikas Singh
Differential Geometry in Computer Vision and Machine Learning workshop. 2021.

[5] A Guided Latent Dirichlet Allocation Approach to Investigate Real-time Latent Topics of Twitter Data during Hurricane Laura
Sulong Zhou, **Pengyu Kan**, Qunying Huang, Janet Silbernagel
Journal of Information Science. 2021.

[6] Application of Image Segmentation in Surface Water Extraction of Freshwater Lakes using Radar Data
Sulong Zhou, **Pengyu Kan**, Janet Silbernagel, Jiefeng Jin
ISPRS Int. J. Geo-Inf. 2020, 9, 424.

[7] Conditional Gradient Optimizer
Pengyu Kan, Vishnu Lokhande
Tensorflow Addons.

Diffusion-based Neonatal HIE Lesion Detection and Inpainting June 2025 - Present

- Advisors : Prof. Kenichi Oishi, Prof. Craig Jones
- Applied diffusion-based generative models to reconstruct healthy neonatal DTI MRI and design unsupervised lesion detection
- Used reconstruction error and accumulative distance along the diffusion trajectory for HIE lesion localization to handle low-contrastive anomaly scenario
- Employed ControlNet to preserve anatomical consistency during inpainting and reconstruction

Neonatal Brain Age Prediction from Diffusion MRI March 2025 – October 2025

- Advisor Prof. Kenichi Oishi
- Applied multiview DTI encoder with query-based processing for fast inference on neonatal scans
- Aligned and fused information across Fractional Anisotropy (FA), Mean Diffusivity (MD) and B0 maps
- Estimated biological brain age as a marker of neurodevelopmental maturity, with attention-based interpretability linking predictions to clinical risk in preterm and NICU-treated infants
  Accepted ***ISMRM 2026****.* [Project page]

Transformer-based Generalizable and Interpretable Brain Age Predictor May. 2024 - May 2025

- Advisors : Prof. Kenichi Oishi, Prof. Craig Jones
- Applied transformer-based architecture to predict biological brain age through T1w MRI scans
- Achieved linear computational cost with respect to input through the stem and trunk design
- Fused information across different anatomical views to balance between input size and information
- Strong correlation with cognitive scores and brain disorders compared to other state-of-the-art models ; reported MAE of 3.65 years on ADNI/OASIS and 3.54 years on external AIBL validation
  ANA 2025 (American Neurological Association) : Poster Award, Abstract of Distinction, oral presentation (one of five talks selected from hundreds of abstracts), Travel Award, and Emerging Scholar
  Accepted ***Radiology Advances 2026****.* [Preprint] [Project page] [News]

Transformer-based Masked Adversarial Defenses to Patch Attacks Sep. 2021 - July. 2023

- Advisor : Prof. René Vidal
- Studied adversarial robustness of attention mechanism and transformer-based architectures in image classification and object detection tasks
- Compared with convolution-based models, including ResNet and Faster-RCNN
- Designed Masked Adversarial Defenses with randomness specifically targeting adversarial patch attacks
- Achieved comparable adversarial robustness to state-of-the-art defenses, in a more efficient training process
- Reached top performance in DARPA GARD project on CARLA object detection for street scene [Pre-print]

Diffeomorphism Invariant Layer for Neural Network Aug. 2020 - May. 2021

- Advisors : Prof. Vikas Singh and Dr. Rudrasis Chakraborty
- Designed and implemented equiconv and invariant layers for Neural Network to handle diffeomorphism deformation in input images
- Applied over ResNet and ConvNet and tested on Fashion MNIST, CIFAR10, SVHN, STL10 and CelebA datasets
- Compared with baseline models, including the Spatial Transformation model, CPAB Spatial Transformation model and Rotational Equivariance CNN model
- Test accuracy outperformed baseline models in cases of small scale affine and homography deformation by around 10% on CIFAR10 and around 20% on STL10
  [Accepted DiffCVML 2021 CVPR workshop][Published paper]

Latent Dirichlet Allocation (LDA) Method for Analyzing Twitter Data on Disaster Management Aug. 2020 - Dec. 2020

- Collaborator : Dr. Sulong Zhou
- Applied Natural Language Processing (NLP) and the LDA method to extract topic and situational awareness (SA) information from Twitter data during the 2020 Hurricane Laura
- Proposed general workflow to access SA information from Twitter with four main steps, including data collection, preprocessing, LDA with guided approach and latent topic clusters
- Used the NLTK package for data preprocessing, including tokenizing, removing stop words and lemmatizing
- Evaluated the selected latent topics with the Coherence Value model and PyLDAvis Visualization tool
- Found 7 main topics among these tweets and their temporal trends during the 10-day period of the Hurricane Laura
  [Published Paper] [Code]

Adversarial Deep Metric Learning April 2020 - Jan. 2021

- Advisors : Prof. Earlence Fernandes and Prof. Somesh Jha
- Formulated adversarial attack toward Deep Metric Learning (DML) model
- Addressed the issue of data dependence for the distance-based loss function for DML model
- Tested robustness of traditionally trained and adversarial trained DML model toward PGD, FGSW, and Carlini & Wagner's attack on ResNet over CUB200, CARS196 and SOP datasets
- Applied attacks separately targeting on the positive anchor points, positive data points and negative datapoints in the triplet loss for the DML
- Observed weak robustness of traditionally trained DML model and higher accuracy of adversarial trained model under these attacks
- Applied high dimensional Gaussian samples with our proposed adversarial trained model and embedded them into 2d space to illustrate the effect of adversarial training
  [Pre-print]

Frank Wolfe Method for Area Under ROC (AUROC) Problem Feb. 2020 - Aug. 2020

- Advisors : Prof. Vikas Singh and Prof. Sathya N. Ravi
- Reformed AUROC problem based on the Riemann sum and as a constrained saddle point optimization problem, under 1 - norm constraint for the Lagrangian multiplier
- Compared performance of Frank Wolfe method with conventional SGD solution toward this saddle point optimization problem over MNIST, CIFAR10, CIFAR100 dataset on ResNet, Deep CNN and Kernel method
- Observed higher AUROC score over the validation set by 0.04 on CIFAR10 and ResNet with using Frank Wolfe method than conventional SGD solution for this constrained saddle point problem
- Applied Block Frank Wolfe method and added calibration loss term to reduce the convergence gap between our proposed method and the conventional SGD solution

The Otsu Method for Water Imaging Classification Dec. 2019 - May 2020

- Collaborator : Dr. Sulong Zhou
- Applied the Otsu Method on radar reading and satellite images to classify water
- Tested on the radar values of the Poyang Lake over years to classify the water regions, with an evaluation accuracy of 91%
- Analyzed the changing pattern of the freshwater area of Poyang Lake over years to help for freshwater resource preservation
  [Published paper]

Conditional Gradient (CG) Optimizer March 2019 - Jan. 2020

- Advisors : Prof. Vikas Singh and Prof. Sathya N. Ravi
- Implemented and contributed CG Optimizer with Frobenius norm and Nuclear norm to TensorFlow Addons Package
- Provided an alternative optimizer for constrained optimization problems and alternative approach of regularization for deep learning, through considering hard penalty of constraints
- Wrote tutorial for CG Optimizer on MNIST Dataset and compared its performance with SGD
- Took responsibility of maintenance and improvement of the CG Optimizer
  [Documentation] [Conditional Gradient Optimizer] [Tutorial]

Coding Projects

AI Shopping Agent for Older Adults June, 2025 - Present

- Developed an LLM-driven shopping agent (LangGraph, LangSmith, OpenRouter) with voice and text interfaces, replacing dense mobile-app navigation with natural-language search and product evaluation
- Combined structured product catalogs with semantic retrieval over a Redis vector store for fast, relevant matches to user intent
- Optimized the LangGraph workflow and LLM context design, reducing typical end-to-end latency from about 120 s to about 60 s
- Built an age-friendly scoring model that rates product fit to health conditions, mobility constraints, and preferences (footwear, supplements, assistive devices, and related categories)
- Delivered an end-to-end prototype covering intent parsing, retrieval, ranking, and personalized recommendations, with continued expansion of product categories and evaluation metrics
- Deployed on AWS (EC2 compute, S3 storage) for hosting and scaling the agent pipeline [Website]

Teaching Experience

Head Teaching Assistant, Johns Hopkins University Jan. 2024 - May 2024
Vision as Bayesian Inference (601.783) : Assisted on students learning for computer vision ; Prepared assignments and solutions ; Managed teaching assistant team

Teaching Assistant, Johns Hopkins University Sep. 2023 - Dec. 2023
Computer Vision (601.661) : Assisted on students learning for computer vision ; provided lectures on introduction to transformer [Slides]

Peer Mentor, University of Wisconsin - Madison Sep. 2018 - May 2020
Algorithms (CS 577) : Held office hours and explained concepts of Algorithms to students ; provided feedback on students' homework

Artificial Intelligence (CS 540) : Held office hours to explain concepts and ideas related to machine learning and provided examples for better understanding

Academic Mentor, Center of Academic Excellence Sep. 2017 - May 2019
Helped first - year minority students get used to college study ; Tutored college students with Java and object-oriented programming ; Provided extra coding exercises and examples from real life to improve students' interest

Volunteer Tutor, AmeriCorps Sep. 2016 - May 2017
Tutored high school students in Geometry and Algebra ; Provided weekly feedback of students success and adjusted challenge level of exercise

Mentoring

Stelios Antoniadis, Johns Hopkins Computer Science Jan. 2026 - Present

Skills

Programming languages : Python, Java, C++, C
Deep learning framework : PyTorch, Tensorflow
AI agent framework : LangGraph, LangSmith, OpenRouter
Languages : English, Chinese

Relevant Coursework

Machine Learning, Deep Learning, Computer Vision, Vision as Bayesian Inference, Non-linear Optimizations, Matrix Analysis, Bayesian Statistics, Applied Statistics, Monte Carlo Methods, Linear Programming, Algorithms, Numerical Analysis, Analysis, Number Theory, Stochastic Process