Advances in Digital Imaging and Computer Vision (TP283)

Prof. Kostas Marias

Special Topic: Deep Learning in Medical Imaging - Course Project Descriptions

Questions: trivizakis@hmu.gr

SUMMARY

Type I Dermatoscopy: Dermatological Image Analysis (ISIC Dataset)

Project 1: Uncertainty-Aware Medical Image Segmentation[1]

Project 2: Semantic Discovery & Latent Space Alignment for Clinical Dermatology[2]

Project 3: Rigorous Benchmarking of Explainable AI Methods in Skin Lesion Classification[3]

Type II Radiology: Lung Cancer Radiogenomics (NSCLC Radiogenomics Dataset[4])

Project 4: Deep Learning Algorithmic Fairness & Stratified Bias Auditing

Project 5: Cross-Modality Radiotrasnscriptomic Translation via Transformer Architectures

Course Assessment Framework & Timelines

- Weight Distribution of the reports: 40% Literature Review | 60% Deep Learning Analysis & Implementation
- Milestone 1 (Literature Review & Methodology Proposal): Due 31/05/2026
- Milestone 2 (Final Project Code & Analysis Report): Due late June 2026

---

1 https://challenge.isic-archive.com/data/#2018
2 https://challenge.isic-archive.com/data/#2017
3 https://challenge.isic-archive.com/data/#2024
4 Bakr, S., Gevaert, O., Echegaray, S., Ayers, K., Zhou, M., Shafiq, M., Zheng, H., Zhang, W., Leung, A., Kadoch, M., Shrager, J., Quon, A., Rubin, D., Plevritis, S., & Napel, S. (2017). Data for NSCLC Radiogenomics (Version 4) [Data set]. The Cancer Imaging Archive. https://doi.org/10.7937/K9/TCIA.2017.7hs46erv

Project 1: Uncertainty-Aware Medical Image Segmentation

- **Objective:** Implement and evaluate a segmentation model uncertainty mechanism on the ISIC skin lesion dataset to differentiate between reliable predictions and boundaries requiring expert human audit.
- **Scope of Work:** Develop or adapt a robust baseline semantic segmentation network (e.g., U-Net, Attention U-Net, or CIS-UNet) to delineate lesion boundaries.
    - Integrate an explicit uncertainty quantification mechanism. Students may choose an **inference-based approach** (e.g., Monte Carlo Dropout to simulate posterior distributions, or Test-Time Augmentation metrics) or a **training-based approach** (e.g., Evidential Deep Learning or deep ensembles modeling heteroscedastic aleatoric loss).
    - Map and evaluate the resulting pixel-wise uncertainty maps against ground-truth segmentation edge variance.
    - **Keywords:** medical image segmentation uncertainty quantification, uncertainty in deep learning, probabilistic medical image segmentation, Monte Carlo Dropout segmentation, heteroscedastic loss segmentation

Project 2: Semantic Discovery & Latent Space Alignment for Clinical Dermatology

- **Objective:** Establish a mathematically interpretable latent space by bridging the hidden layers of a deep neural network directly to explicit, established dermatological features.
- **Scope of Work:**
    - Train a representation learning model (e.g. Variational Autoencoder [VAE], Contrastive Learning network, or standard CNN backbone) on the ISIC dataset.
    - Map, isolate, or align vectors within the latent space to categorical dermatological attributes present in the metadata, specifically targeting clinical features like *pigment networks, streaks, milia-like cysts, globules, or negative networks*.
    - Perform latent space traversal or regression analysis to demonstrate that shifting along specific latent dimensions corresponds predictably to the presence, density, or structural changes of these medical attributes.
    - **Keywords:** latent space disentanglement medical imaging, concept bottleneck models dermatology, representation learning neural networks, Variational Autoencoders dermatological feature mapping, contrastive learning semantic alignment, latent space traversal dermatology, interpretable visual concepts

Project 3: Rigorous Benchmarking of Explainable AI (XAI) Methods in Skin Lesion Classification

- **Objective:** Build a robust diagnostic deep learning classifier and construct a standardized benchmark evaluating the functional discrepancies, stability, and diagnostic reliability of modern visual attribution methods.
- **Scope of Work:**
    - Train a high-performing Deep Learning model (e.g., ResNet, EfficientNet, or Vision Transformer) to classify skin lesions into multi-class diagnostic categories.
    - Implement and extract feature attributions across all primary classes of XAI techniques[5]:
        - *Gradient/Activation-based:* Grad-CAM, Grad-CAM++, Integrated Gradients.
        - *Perturbation-based:* Occlusion Sensitivity Mapping, RISE.
        - *Game-Theoretic:* SHAP (Shapley Additive exPlanations) adapted for pixel/voxel attribution.
    - Critically report, quantify, and benchmark the visual and mathematical variations among these methods, analyzing their sensitivity to noise, localization accuracy against expert masks, and consistency across patient cohorts.
    - **Keywords**: explainable AI (XAI) medical imaging benchmark, saliency map evaluation dermatology, faithfulness metrics feature attribution, Grad-CAM++ skin lesion classification, Integrated Gradients medical imaging, SHAP pixel attribution evaluation, RISE perturbation saliency maps, localization accuracy visual explanations, XAI sensitivity analysis deep learning

5 Koutoulakis, E., Trivizakis, E., Markodimitrakis, E., Agelaki, S., Tsiknakis, M., & Marias, K. (2026). A critical review of explainable deep learning in lung cancer diagnosis. Artificial Intelligence Review, 59(1), 28.

Project 4: Deep Learning Algorithmic Fairness & Stratified Bias Auditing

- **Objective:** Design, evaluate, and mitigate systemic model biases across protected patient sub-populations within non-small cell lung cancer diagnostic workflows.
- **Scope of Work:**
    - Train a deep learning network to predict disease recurrence utilizing the NSCLC Radiogenomics data.
    - Perform a formal algorithmic fairness audit[6] by stratifying performance metrics (AUC, sensitivity, specificity, calibration curves) across a selected protected category or molecular division:
        - ***Option A:*** Patient demographic factors (Ethnicity)
        - ***Option B:*** Patient demographic factors (Gender)
        - ***Option C:*** Tumor molecular subtypes (EGFR mutation-positive vs. Wild-Type)
        - ***Option D:*** Tumor molecular subtypes (KRAS mutation-positive vs. Wild-Type)
    - Calculate formal fairness constraints and implement mitigation strategies (based on the IBM's AI-Fairness-360) to achieve equitable model performance by retraining an ML model with latent space features of the deep model.
    - **Keywords:** algorithmic fairness medical imaging, bias deep learning oncology, subpopulation shift lung cancer AI, equalized odds medical diagnostic models, adversarial debiasing radiomics, demographic parity clinical prediction, fair representation learning survival models

6 https://research.ibm.com/blog/ai-fairness-360

Project 5: Cross-Modality Radiotrasnscriptomic Translation via Transformer Architectures

- **Objective:** Engineer an attention-driven generative or translation architecture capable of mapping phenotypic macroscopic imaging features directly to microscopic transcriptomic gene expression profiles.
- **Scope of Work:**
  - Extract standardized quantitative image phenotypes from CT volumes using classical computational Radiomics pipelines (PyRadiomics) or deep representation-learning encoders[7].
  - Build a Transformer-based cross-modality mapping model designed to translate these input multi-scale visual/radiomic feature vectors into matched bulk or spatial transcriptomic expression profiles (RNA-Seq vectors[8]).
  - Leverage the global cross-attention mechanisms of the Transformer encoder-decoder blocks to identify and report which specific sub-regions or quantitative radiomic features exhibit the strongest mathematical dependencies with localized gene expression patterns.
  - **Keywords**: radiogenomics cross-modality translation, imaging to transcriptomics mapping, multi-modal fusion medical imaging bioinformatics, transformer architecture radiomics translation, cross-attention radiogenomics, vision-to-sequence deep learning oncology

7 Trivizakis, Eleftherios, et al. "Radiotranscriptomics of non-small cell lung carcinoma for assessing high-level clinical outcomes using a machine learning-derived multi-modal signature." BioMedical Engineering OnLine 22.1 (2023): 125. https://doi.org/10.1186/s12938-023-01190-z

8 Transcriptomic Data Link (GEO Archive): National Center for Biotechnology Information (NCBI) Gene Expression Omnibus, Accession Number: GSE103584. https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE103584