

Research Statement

Min Lee

School of Computing and Information Systems, Singapore Management University

Email: mhlee@smu.edu.sg

December, 2025

Recent advances in artificial intelligence (AI) have made it increasingly applicable to address challenges in every facet of our lives. Yet, deploying AI systems remains challenging and has received concerns about their harms. My work investigates to design and develop human-centered approaches that enable people to responsibly collaborate with AI/robotic systems and evaluate their impact on a task of high-stakes, diverse application domains, such as healthcare (e.g. stroke rehabilitation [1,2], improving vision-language models for medical image analysis, detecting stroke symptoms [10], screening head and neck cancer [12], and monitoring older adults [6]) and public services [7].

Human-AI/Robotic Collaborative Systems for Post-Stroke Rehabilitation

For an effective human-AI/robot collaboration, we have iteratively engaged with therapists and post-stroke patients to design, develop, and evaluate two human-AI collaborative systems to improve the practices of physical stroke rehabilitation: 1) an AI-based decision support system for therapists [1] and 2) a robotic exercise coach for post-stroke patients [2].

Human-AI Collaborative Decision Making

Our human-AI collaborative decision-making (HACDM) system for stroke rehabilitation assessment provides interactive explanations to assist users in (i) delegating tasks to AI (Figure 1a and Figure 1b), (ii) onboarding with AI (Figure 1c), and (iii) completing their rehabilitation assessment (Figure 2).

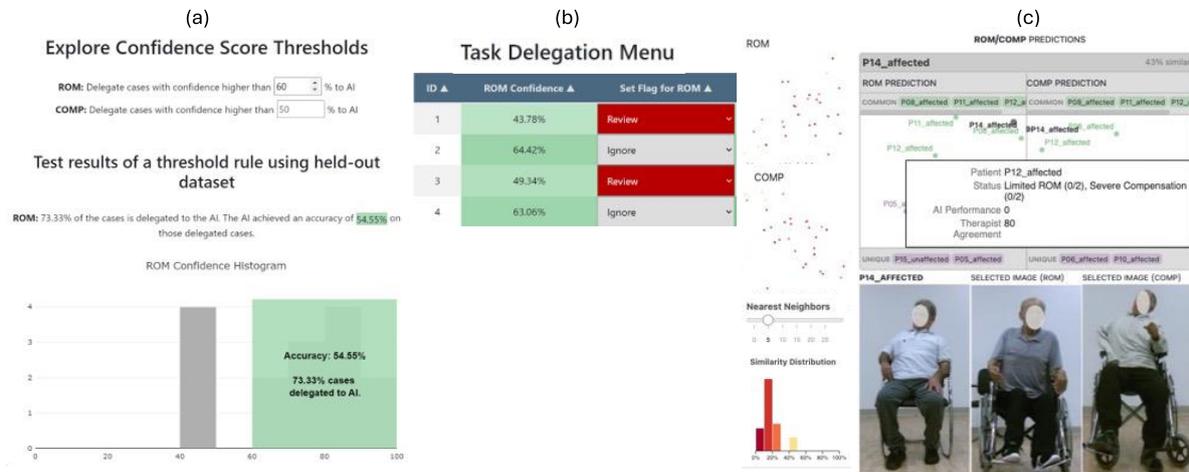


Figure 1. Interface of HACDM: (a) users can explore different thresholds of confidence scores, review AI performance on delegated cases to AI using a held-out dataset, and specify a threshold to delegate cases to AI; (b) users can review AI confidence scores to confirm cases that require an expert review; (c) users can review an embedding visualization of data and hover around embedding data points to review bench-markable information of AI (e.g. the status of a patient and AI performance) and images of a nearest neighbor post-stroke survivor.

To improve a health professionals' onboarding process with AI, we leveraged the existing guidelines of human-AI interaction to create onboarding materials for the system and conducted interviews with 12 therapists and 4 students in medicine and health to collect their feedback on an onboarding process with AI [9]. Drawing on the insights from the interviews that therapists desire to quickly review examples along with the AI performance, we developed interactive example-based explanations of the HACDM (Figure 1b) for onboarding

with AI [9]. The evaluation study with 16 domain experts and 15 novices showed that our interactive example-based explanations assisted participants to have a 9.3% higher ratio of changing their decisions into 'right' than those without interactive example-based explanations [9].

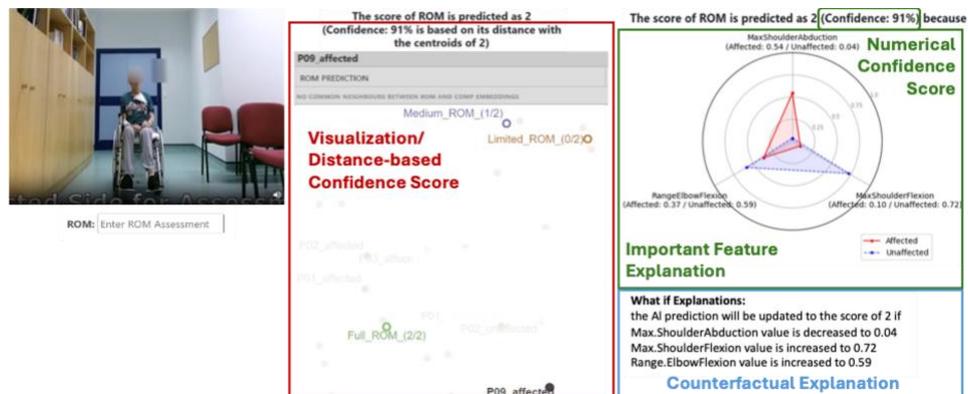


Figure 2: Interface of HAICDM presents a video of patient’s exercises, an AI predicted score, a distance-based visualization of a confidence score, and a patient-specific analysis with three most important features and counterfactual explanations.

To leverage the analytic strengths of AI and the expert knowledge of therapists, the HAICDM integrates an AI model with a rule-based model into a hybrid model for assessing patients’ quality of motion [3]. Instead of just presenting a numerical confidence score, the HAICDM provides an interactive distance-based visualization of a confidence score [11]: identifying the *k*-nearest neighbourhoods of a case and visualizing their embedding spaces along with the embedding spaces of the centroid of class labels (Figure 2). Also, the HAICDM automatically identifies salient frames [8] and salient features of assessment [3] to generate a patient-specific analysis along with counterfactual explanations for rehabilitation assessment (Figure 2) [5]. It also accommodates therapist’s feedback (e.g. feature relevance and relabelling) to iteratively refine the system and personalize rehabilitation assessment [1,2,4].

Studies with therapists show that the HAICDM system empowers therapists to have a richer understanding of patient’s performance and achieve significantly higher agreement on assessment (0.71 F1-score) than the traditional system (0.66 F1-score, $p < 0.01$) [1]. Our distance-based visualization of confidence scores with example-based explanations are effective to significantly reduce users to make ‘Wrong’ decisions (9.85% lower ratio, $p < 0.01$) after reviewing AI outputs [11]. In addition, our counterfactual explanations assisted therapists in reducing their overreliance on ‘wrong’ AI outputs by 21% than presenting only salient feature explanations [5]. After reviewing the analysis of the HAICDM, therapists can understand the capabilities of the HAICDM and provide feedback to significantly improve its performance to replicate the therapist’s assessment from 0.83 to 0.91 F1-score ($p < 0.01$) [1].

These results show that both domain experts and an AI system can learn from each other’s strengths over interactions and generate hybrid intelligence on a complex decision-making task with improved accuracy.

Human-Robot Collaborative Physical Therapy

My research also explores how an interactive robotic coach can collaborate with post-stroke patients to improve their engagement in physical rehabilitation therapy [2]. This system also applies our hybrid intelligence approach to automatically monitor and guide rehabilitation exercises of patients through social interactions (e.g. verbal, visual, and gesture-based feedback) (Figure 3). In contrast to prior work that utilizes pre-defined, generic feedback, this



Figure 3: An interactive robot coach monitors a patient’s exercise and provides personalized feedback

system tunes with a new patient's motions and generates transparent, personalized corrective feedback.

A study showed that our system can be tuned with patient's exercises to significantly improve its performance to provide personalized assessment and feedback from 0.74 to 0.82 F1-score ($p < 0.01$) [2]. In addition, a real-world evaluation with 10 participants showed that our system can adapt to new participants and achieved a 0.81 F1-score to assess their exercises, which is comparable to experts' performance and better than fine-tuning a feed-forward neural network (0.74 F1-score) [2].

Vision-Language Models (VLMs) for Medical Image Analysis

My research aims to addressing a key limitation of modern vision–language models (VLMs): their internal attention maps often misalign with clinically meaningful regions, spreading attention over irrelevant context and undermining trust in medical settings. To address this, we propose Activation Regulation Loss that combines a containment term (suppress attention outside the region of interest) and a spread term (control how attention is distributed within the region) and develop a controllable saliency framework for medical vision–language applications. Across 12 medical imaging datasets, our framework yields more concentrated, clinically aligned attention maps and suppressing spurious activations.

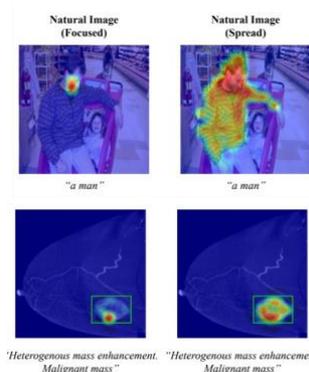


Figure 4: illustrating the ability to controllably produce focused versus spread attention.

Head and Neck Cancer Screening

We aim to design and develop computational techniques and practices to improve collaborations between clinicians and AI for the triage and diagnosis of HNC. Specifically, these techniques include a YOLO-based glottis detector, a classifier of patient's status using clinical endoscopy videos, uncertainty quantification techniques, and interactive Vision Language Model-based AI explanations, which will enable clinicians to understand AI outputs and provide their inputs to improve AI.

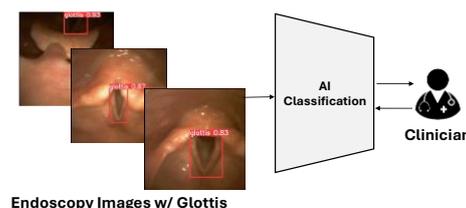


Figure 5: AI-Assisted System for Head and Neck Cancer Screening in Low-Resourced Settings.

Detecting Stroke Symptoms

We have explored a multimodal fusion-based deep learning model that utilizes unstructured data (i.e. images with facial line segments) and structured data (i.e. features of facial expressions) [10], which achieve an 0.96 F1-score using videos of 20 facial palsy patients to detect facial palsy.

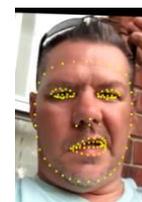


Figure 6: Overlaid facial landmarks for Facial Palsy Detection

Older Adult Care

We have designed and implemented an intelligent system for older adult care, leveraging the findings from a focus-group session with family caregivers [6]. The system (Figure 7) detects abnormal patterns of activities of older adults with wireless motion sensors and ML algorithms [6]. The system also provides inter-active dialogue responses to explain abnormal patterns of older adults to caregivers and allow older adults to proactively share their status with caregivers for an adequate intervention [6]. An evaluation study with older adults and professional caregivers showed that they appreciate the value of our system to provide a faster, personalized service while enabling older adults to proactively control what information shared through interactive dialogue responses [6].

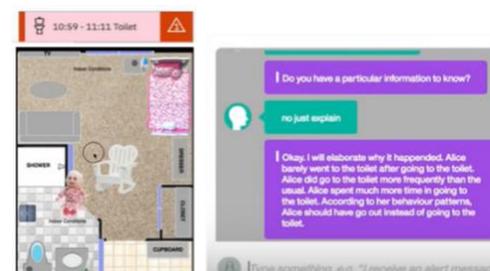


Figure 7: An intelligent system that detects abnormal behavior patterns and explains its analysis to a caregiver and an older adult.

Selected Publications and Outputs

1. **Min Hun Lee**, Daniel P. Siewiorek, Asim Smailagic, Alexandre Bernardino, and Sergi Bermúdez Badia. 2021. "A Human-AI Collaborative Approach for Clinical Decision Making on Rehabilitation Assessment". In Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '21).
2. **Min Hun Lee**, Daniel P. Siewiorek, Asim Smailagic, Alexandre Bernardino, and Sergi Bermúdez Badia. 2020. "Design, development, and evaluation of an interactive personalized social robot to monitor and coach post-stroke rehabilitation exercises." *User Modeling and User-Adapted Interaction* 33.2 (2023): 545-569.
3. **Min Hun Lee**, Daniel P. Siewiorek, Asim Smailagic, Alexandre Bernardino, and Sergi Bermúdez Badia. 2020. "Interactive hybrid approach to combine machine and human intelligence for personalized rehabilitation assessment." Proceedings of the ACM Conference on Health, Inference, and Learning. 2020.
4. **Min Hun Lee**, Daniel P. Siewiorek, Asim Smailagic, Alexandre Bernardino, and Sergi Bermúdez Badia. 2022. "Towards Efficient Annotations for a Human-AI Collaborative, Clinical Decision Support System: A Case Study on Physical Stroke Rehabilitation Assessment". In 27th International Conference on Intelligent User Interfaces (IUI '22).
5. **Min Hun Lee** and Chong Jun Chew. 2023. "Understanding the Effect of Counterfactual Explanations on Trust and Reliance on AI for Human-AI Collaborative Clinical Decision Making." Proceedings of the ACM on Human-Computer Interaction 7. CSCW2 (2023).
6. **Min Hun Lee**, Daniel P. Siewiorek, and Alexandre Bernardino. 2023. "Designing a Human-Centered Intelligent System to Monitor & Explain Abnormal Patterns of Older Adults." Proceedings of the 25th International ACM SIGACCESS Conference on Computers and Accessibility. 2023.
7. Logan Stapleton, **Min Hun Lee**, Diana Qing, Marya Wright, Alexandra Chouldechova, Ken Holstein, Zhiwei Steven Wu, and Haiyi Zhu. 2022. "Imagining new futures beyond predictive systems in child welfare: A qualitative study with impacted stakeholders." In 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)
8. **Min Hun Lee** and Yi Jing Choy. 2023. "Exploring a Gradient-based Explainable AI Technique for Time-Series Data: A Case Study of Assessing Stroke Rehabilitation Exercises". In ICLR 2023 Workshop on Time Series Representation Learning for Health.
9. **Min Hun Lee**, Renee Ng, Silvana Xin Yi Choo, and Shamala D. Thilarajah. 2024. "Interactive Example-based Explanations to Improve Health Professionals' Onboarding with AI for Human-AI Collaborative Decision Making". In 27th European Conference on Artificial Intelligence (ECAI 2024), 2024.
10. Oo, Heng Yim Nicole, **Min Hun Lee**, and Jeong Hoon Lim. "A Multimodal Fusion Model Leveraging MLP Mixer and Hand-crafted Features-based Deep Learning Networks for Facial Palsy Detection." Pacific-Asia Conference on Knowledge Discovery and Data Mining 2025, Special Session on Data Science: Foundations and Applications (DSFA). 2025
11. **Min Hun Lee** and M. Z. Tok. "Towards Uncertainty Aware Task Delegation and Human-AI Collaborative Decision-Making. In Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT). 2025.
12. **Min Hun Lee**, Lam, S.S.W., Liew, S.X.H., Dorosan, M., Graves, N., Karlström, J., Tan, H.K. and Lee, W.T., 2025, August. AI-assisted triage and decision support of head and neck cancer screening and diagnosis in low-resourced settings. In Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI). (pp. 9781-9789). 2025.