MEDICAL IMAGING SYMPOSIUM FOR PHDS AND POSTDOCS

MISP² 2026

schedule
keynote speaker
book of abstracts

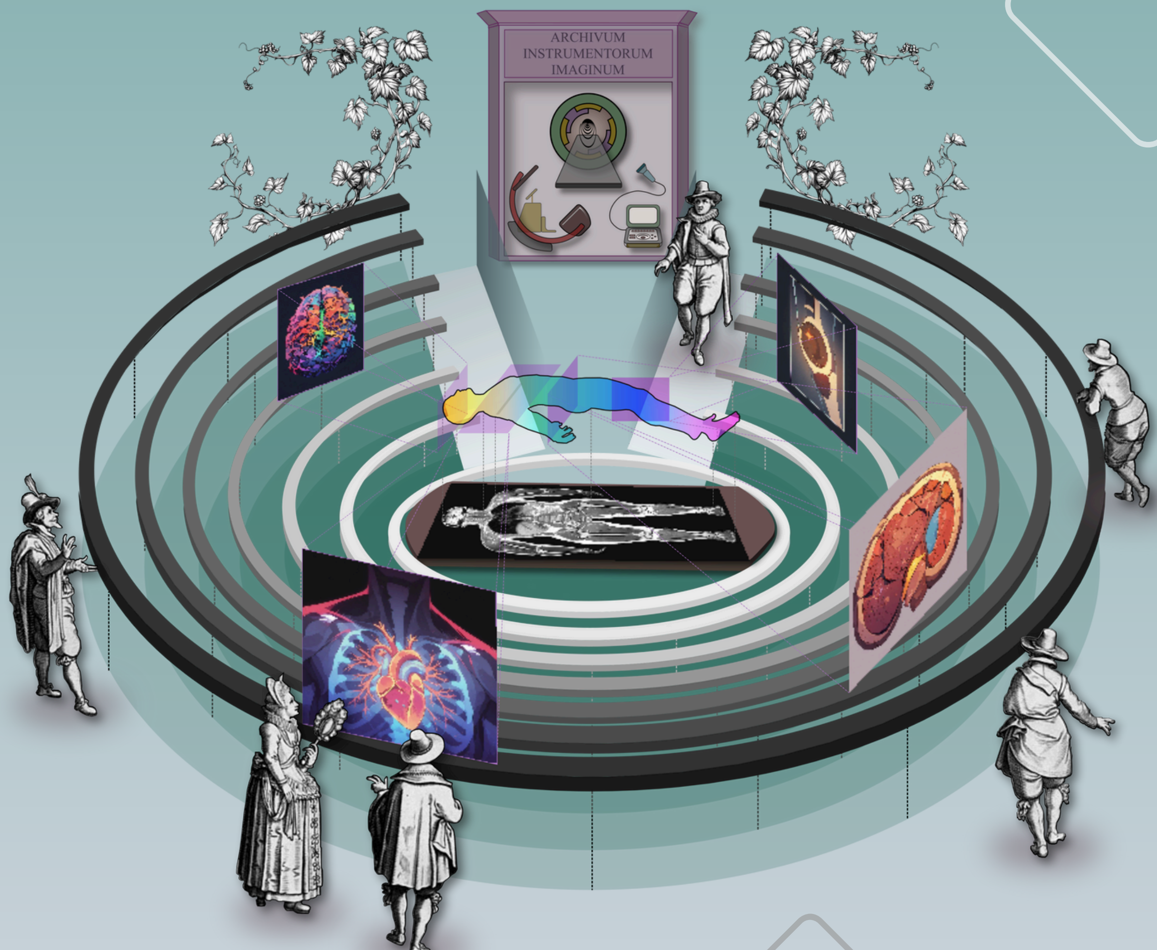

MISP² | 27 January 2026 | Leiden University Medical Center | Leiden

SCHEDULE

9:00 - 9:45	Registration
9:45 - 10:00	Opening
10:00 - 10:55	Keynote talk I
11:00 - 12:00	Session I
12:00 - 13:00	Lunch
13:00 - 13:55	Keynote talk II
14:00 - 15:30	Poster session
15:30 - 16:30	Session II
16:30 - 17:30	Award ceremony + Informal drinks

KEYNOTE SPEAKER

Willem Grootjans | LUMC

Willem Grootjans is an assistant professor and technical physician at the department of Radiology at the Leiden University Medical Center. He graduated cum laude from the Technical University of Twente at the faculty of Science and Technology in Technical Medicine with a specialization in Robotics and Imaging. In 2012, he started his PhD at the department of Radiology and Nuclear Medicine of the Radboud University Medical Center in Nijmegen, focusing on the use of positron emission tomography for personalizing clinical management of patients with lung cancer.

In 2016 he obtained his PhD at the Radboud faculty of Medical Sciences and started as a postdoctoral fellow at the Leiden University Medical Center at the department of Radiology. During his postdoc, he continued his work in nuclear imaging sciences, as well as computed tomography and magnetic resonance imaging, for the purpose of personalizing clinical management of sarcoma patients. In 2018 he formed a specialized image processing group, at the department of Radiology of the Leiden University Medical Center, responsible for extracting relevant information from medical images for diagnosis and image-guided treatment decisions, also known as the 'Imaging Services Group' (ISG). Since 2019 he is head of the ISG which is now responsible for delivering high quality information to radiologists by using state-of-the art image processing software (including AI) that can be used for radiological reporting in clinical routine. Willem continues his mission to research and implement new technological innovations to improve the quality of healthcare, while keeping in mind the needs of the patient, healthcare professional, and society at the same time.

Keynote talk:
AI in Radiology: From Concept to Clinical Impact

LUMC

SESSION I

1. Zero-shot evaluation of promptable foundation models for MSK CT segmentation

Caroline Magg
Amsterdam University Medical Center

Promptable Foundation Models (FMs) have revolutionized interactive segmentation for medical imaging, as they can serve different purposes, from accelerating annotations to being fine-tuned to specific datasets. Selecting the most suitable model depends on the task at hand, the available resources, and the required accuracy. With the growing number of models, evaluations are typically conducted in isolation, making direct comparison challenging. Although broad benchmarks have been proposed to demonstrate generalization across diverse datasets, clinical application often demand solutions for specific tasks. Thus, we evaluated 11 promptable FMs (SAM, SAM2, Med-SAM, Med-SAM2, SAM-Med2D, SAM-Med3D, ScribblePrompt, MedicoSAM, SegVol, Vista3D, nnInteractive) in the context of musculoskeletal (MSK) CT segmentation. A total of 370 slices were extracted from private and public datasets, covering 18 bone and implant classes across four anatomical regions: shoulder, wrist, lower leg, and hip. Non-iterative 2D and 3D prompting strategies, automatically extracted from the reference mask, were used. The models were categorized based on their prediction dimensionality (2D vs. 3D), and evaluated in the Pareto sense to avoid prioritizing a single metric (i.e., Dice Similarity Coefficient, 95-percentile Hausdorff distance, Normalized Surface Dice) while also incorporating computational resources. The two main insights are: The bounding box prompt performs well in 2D and 3D; Within the pool of models, medical dedicated FMs, such as MedicoSAM2D, Med-SAM2, and nnInteractive can perform on-par or outperform SAM and SAM2.1. In future work, to further understand the most effective use of promptable FMs, we aim to investigate the prompt robustness of Pareto-optimal models using human-generated prompts.

2. Implicit Neural Representations for continuous image reconstruction (in diffusion MRI)

Tom Hendriks
Eindhoven University of Technology

Implicit neural representations (INRs) provide a way to create a continuous, spatially correlated representation of an image or a volume, by using a neural network to map encoded coordinates to a desired output. INRs have been used across a variety of fields, finding success in data compression, novel-view synthesis, image reconstruction, and physics-informed modelling. In the field of diffusion magnetic resonance imaging (dMRI), we have used this technique to make spatially regularized representations of the acquisitions, which showed great noise robustness and interpolation accuracy. Furthermore, we have used INRs to fit (multi-shell multi-tissue) constrained spherical deconvolution, and the Standard Model of White Matter, which improved the results compared to the conventional voxel-based approaches.

The usage of INRs is becoming increasingly prevalent in the field of medical image analysis. In this work we look at how INRs can be used for a variety of tasks (in the field of dMRI) and talk about their strengths and limitations. Additionally, we will show some recently published work and look into exciting ongoing work.

3. Combining Generalized Additive Models and Deep Learning: An Interpretable Convolutional Neural Network for Alzheimer's Diagnosis and Prognosis

Wenjie Kang
Erasmus Medical Center

Machine learning methods, particularly deep learning models applied to neuroimaging data—have shown great promise for improving early diagnosis of Alzheimer's disease (AD). However, the limited interpretability of these black-box models remains a major barrier to their integration into clinical workflows. Generalized Additive Models (GAMs) offer strong interpretability but are constrained to low-dimensional inputs and cannot directly process high-dimensional data such as images. In this study, we propose an end-to-end Interpretable Convolutional Neural Network (ICNN) that integrates convolutional neural networks (CNNs) with the transparency of GAMs for AD diagnosis and prognosis. We compare the ICNN with state-of-the-art black-box deep learning models and conventional machine learning approaches. Our model achieves performance comparable to leading deep learning methods. Using the Alzheimer's Disease Neuroimaging Initiative (ADNI) dataset, the ICNN obtained an AUC of 0.942 for AD vs. control classification and 0.705 for predicting conversion from mild cognitive impairment (MCI) to AD. Furthermore, we validated the ICNN on two external testing sets PND and OASIS3. A key strength of the ICNN is its inherent interpretability: it provides regional feature importance at both individual and group levels, offering transparent insights into the model's decision-making process and highlighting the contributions of specific brain regions to AD diagnosis.

4. Neural Surrogates for Hemodynamics Estimation in Cardiovascular Diseases

Patryk Rygiel
University of Twente

Computational fluid dynamics (CFD) is a commonly used method of in-silico blood flow modelling in the assessment of cardiovascular diseases (CVDs). CFD, a numerical approach to solving the Navier-Stokes equations that govern fluid dynamic, can provide hemodynamic factors like time-averaged wall shear stress (TAWSS) or oscillatory shear index (OSI), that have been found to correlate with the development and progression of various CVDs. However, CFD simulations are known to be computationally demanding, limiting their seamless integration in clinical practice. Hence, in recent years, geometric deep learning methods, operating directly on 3D shapes, have been proposed as compelling surrogates, estimating hemodynamic parameters in just a few seconds.

As reference data acquisition for model training is costly, improving neural surrogate data efficiency is essential. To address this challenge, we propose a geometric neural surrogate that leverages symmetry preservation and physics-based regularization to provide strong inductive biases that enhance data efficiency. We demonstrate that our model generalizes to patient-specific inflow conditions and diverse vascular topologies while being trained on only 100 geometries. Furthermore, using physics-guided active learning, we can reduce the required training set by an additional 50% without compromising model quality. We evaluate our approach across multiple vascular datasets, ranging from aortas to coronary arteries, and analyze clinically relevant hemodynamic metrics such as wall shear stress, velocity, and pressure changes. Our results highlight the robustness and generalizability of the proposed neural surrogate and underscore its potential for improving hemodynamic parameter estimation in clinical practice.

5. Generalizable, Cross-sequence Physics-Informed Quantitative MRI Super-resolution

Alireza Samadifardheris
Erasmus Medical Center

Quantitative MRI (qMRI) provides tissue-specific parameters (PD, T1, T2) but is typically acquired at low resolution to keep scan times practical. Conventional weighted images, however, offer higher spatial detail. Existing qMRI super-resolution methods depend on high-resolution ground-truth maps, which are difficult to obtain in practice. We propose PHIRE-Q, a physics-informed, self-supervised framework that uses high-resolution weighted images to guide qMRI super-resolution without requiring high-resolution qMRI during training. The method is trained on synthetic QRAPMASTER data and evaluated for cross-sequence generalization on 3D MUPA acquisitions.

qMRI maps from 27 volunteers were used to generate synthetic high-resolution weighted images via an analytic signal model, then degraded to simulate low-resolution acquisitions. A ResNet-based model was trained to predict high-resolution qMRI by minimizing two losses: (1) a reconstruction loss enforcing agreement between degraded simulated contrast images and those derived from predicted qMRI, and (2) a physics-based loss encouraging synthesized weighted images from predictions to resemble high-resolution guides. Performance was quantified using SSIM and HFEN. Generalization was tested by applying PHIRE-Q to low-resolution MUPA qMRI with real T1w/T2w images as guides.

PHIRE-Q improved structural fidelity and reduced high-frequency error for all parameter maps in the synthetic validation set. When applied to MUPA data, the method enhanced PD, T1, and T2 maps toward higher-sampled references and produced synthesized weighted images that matched real guides more closely than those derived from low-resolution inputs.

PHIRE-Q enables self-supervised, physics-informed qMRI super-resolution and generalizes across acquisition sequences.

KEYNOTE SPEAKERS

Peter Kitslaar & Prerak Mody | Medis

Pieter Kitslaar has worked for more than 15 years with cardiac imaging, using both classical and AI techniques. Cardiac CT software co-developed by him has been used in 150+ publications together with clinical collaborators across the world. He is currently head of the Applied Research department at Medis.

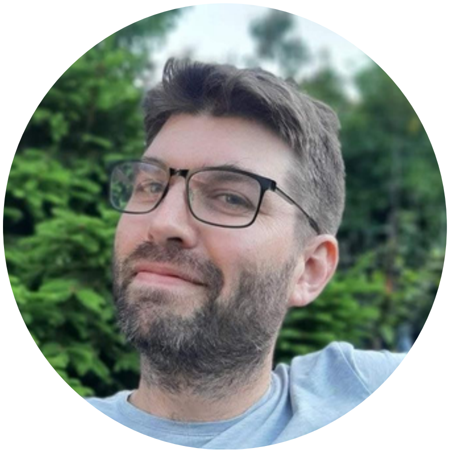

Prerak Mody received his PhD in artificial intelligence for medical imaging at Leiden University Medical Center. He has over 6 years of experience applying deep learning techniques in the field of medical imaging, both in academia and industry. He is currently a researcher at Medis Medical Imaging where he explores tools to help better understand data at scale.

Keynote talk:
AI Data Strategies in Cardia Image Processing

Medis MEDICAL IMAGING

SESSION II

1. MIMIC: Multiple Instance learning for Identification of Mycosis fungoides in Cutaneous biopsies

Siemen Brussee
Leiden University Medical Center

The diagnosis of early-stage mycosis fungoides (eMF) is challenging due to substantial morphological overlap with benign inflammatory dermatoses (BIDs), resulting in a median diagnostic delay of 36 months and considerable resource utilization. To address this challenge, we developed MIMIC (Multiple Instance-learning for Identification of Mycosis fungoides In Cutaneous biopsies), a weakly supervised deep learning model aimed to accurately identify MF cases among cases with suspicion for MF.

MIMIC was trained on 3,893 H&E-stained whole slide images (WSIs) collected from six European centers. Model performance was assessed using geographic validation across four independent international sites (n=383 WSIs) and a multi-reader study involving 11 dermatopathologists and pathologists (n=172 WSIs). Diagnostic performance and confidence were compared between MIMIC and human readers using a Graded Response Model.

In geographic validation, MIMIC demonstrated strong transportability, with a macro-average AUROC of 0.90. In the reader study, MIMIC performed comparably to expert pathologists in discriminating eMF from BIDs (AUROC 0.81 vs. 0.80). Human readers assigned MF diagnoses more conservatively, compared to MIMICs more confident MF predictions. Some false-positive classifications occurred among complex BID cases, reflecting the diagnostic difficulty of the task.

Given the inherent challenges of eMF diagnosis, these findings support the potential of AI-assisted triage. By functioning as a digital safety net during initial H&E review, MIMIC may improve diagnostic efficiency, reduce delays, and enable more targeted use of specialized diagnostic resources. Prospective validation, multimodal integration and modeling of spatial cellular relationships are important next steps for model improvement and clinical implementation.

2. Ranking Explainable AI Methods for Head and Neck Cancer Outcome Prediction

Baoqiang Ma
University Medical Center Utrecht

For head and neck cancer (HNC) patients, prognostic outcome prediction can support personalized treatment strategy selection. Improving prediction performance of HNC outcomes has been extensively explored by using advanced artificial intelligence (AI) techniques on PET/CT data. However, the interpretability of AI remains a critical obstacle for its clinical adoption. Unlike previous HNC studies that empirically selected explainable AI (XAI) techniques, we are the first to comprehensively evaluate and rank 13 XAI methods across 24 metrics, covering faithfulness, robustness, complexity and plausibility. Experimental results on the multi-center HECKTOR challenge dataset show large variations across evaluation aspects among different XAI methods, with Integrated Gradients (IG) and DeepLIFT (DL) consistently obtained high rankings for faithfulness, complexity and plausibility. This work highlights the importance of comprehensive XAI method evaluation and can be extended to other medical imaging tasks.

3. Interactive Physically Based Visualization of MRI Brain Anatomy and Tractography

Bram Kraaijeveld
Eindhoven University of Technology

Cinematic rendering aims to improve the visual understanding of medical imaging data, yet current (closed-source) tools remain limited in interactivity and performance. We propose an open-source framework for real-time, physically based visualization that unifies volumetric brain anatomy, derived from T1/T2 MRI with diffusion MRI tractography. Building on ideas explored in our open-source vibrant-tractography project (https://github.com/as-the-crow-flies/vibrant-tractography), the envisioned system will support interactive manipulation of lighting, transfer functions, cutting planes, and viewing parameters while maintaining photorealistic quality.
The central idea is to enable researchers, clinicians, and educators to explore neuroanatomy and white-matter pathways dynamically, gaining clearer insight into spatial relationships that are difficult to convey through static images or traditional 2D slice-based views. By combining physically based anatomical rendering and tractography in a single interactive environment, this proposed tool has the potential to enhance neuroscience research, medical education, and pre-surgical planning.

4. Bilingual expertise impacts brain age: region-wise study in bilinguals, translators, and interpreters

Mingshi Chen
Amsterdam University Medical Center

Bilingualism has been linked to structural brain change that may enhance neural reserve and delay Alzheimer disease. The Dynamic Restructuring Model proposes stage wise adaptations across bilinguals, translators, and interpreters. Brain age models estimate biological age, and the brain age gap ΔBAG is the age bias corrected predicted age minus chronological age. We asked whether ΔBAG and its regional associations differ across these stages

We studied 103 native Czech speakers with advanced English, including 42 bilinguals, 35 translators, and 26 interpreters. MRI on a 3T Siemens Prisma included T1 weighted imaging at 0.7 millimeter resolution and diffusion tensor imaging. FastSurfer provided regional volumes and intracranial volume, and TBSS yielded tract wise fractional anisotropy. Seven brain age models were evaluated with age bias correction. Group differences used analysis of variance. ROI volumes were residualized for intracranial volume, age, and sex, and ROI to ΔBAG Pearson correlations used false discovery rate correction.

After age bias correction, only the BrainAge model showed a significant group difference, with interpreters younger relative to age than bilinguals. BrainAge also revealed expertise specific regional patterns. In bilinguals, higher ΔBAG related positively to ventricular and cerebrospinal fluid volume. Translators showed minimal ROI level associations. Interpreters showed negative associations in orbitofrontal cortex and a positive association in right ventral diencephalon.

Debiased brain age modeling differentiates bilingual expertise stages and supports a nonlinear trajectory from early expansion to stabilization to targeted renormalization across cortical and subcortical systems. Replication in older and longitudinal cohorts is needed..

5. **Pathology Report Generation and Multimodal Representation Learning for Cutaneous Melanocytic Lesions**

Ruben Lucassen
University Medical Center Utrecht & Eindhoven University of Technology

Millions of melanocytic skin lesions are examined by pathologists each year, the majority of which concern common nevi (i.e., ordinary moles). While most of these lesions can be diagnosed in seconds, writing the corresponding pathology report is much more time-consuming. Automating part of the report writing could, therefore, alleviate the increasing workload of pathologists. In this work, we develop a vision-language model specifically for the pathology domain of cutaneous melanocytic lesions. The model follows the Contrastive Captioner framework and was trained and evaluated using a melanocytic lesion dataset of 42,512 H&E-stained whole slide images and 19,645 corresponding pathology reports. Our results show that the quality scores of model-generated reports were on par with pathologist-written reports for common nevi, assessed by an expert pathologist in a reader study. While report generation revealed to be more difficult for rare melanocytic lesion subtypes, the cross-modal retrieval performance for these cases was considerably better.

POSTERS

1. Automated Classification of Left Atrial Appendage Contrast Filling Defects in Computed Tomography

Laura Angeja, Amsterdam University Medical Center

2. Automated Method Design for Cancer Image Classification by Differential Evolution and Ensembling

Natalia Oviedo Acosta, Erasmus Medical Center

3. Beyond Artifacts: Dual-Stage MAR leveraging Wavelet Representations along with CT Domain Translation

Mubashara Rehman, Università degli Studi di Udine

4. Beyond end-to-end – Exploring Extremes in Medical Continual Learning

Amin Ranem, University of Twente

5. Beyond Nodules: AI for Comprehensive Lung Tumor Segmentation on CT

Valerio Pugliese, The Netherlands Cancer Institute

6. Deep learning for vessel occlusion classification using CT perfusion maps in acute ischemic stroke

Odysseas Papakyriakou, Amsterdam University Medical Center

7. Deep learning-based post-treatment brain MRI generation for patients with glioma from pre-RT priors

Selena Huisman, Amsterdam University Medical Center

8. Explainable AI and Vision Transformers for diagnosis of early diabetic retinopathy in retina color images

Maciej Szymkowski, Bialystok University of Technology

9. Federated Fine-tuning of Medical Foundation Models for MRI-based Dementia Classification

Kaouther Mouheb, Erasmus Medical Center

10. From Sections to Patient: AI-Based Patient-Level Gleason Grading from Whole Slide Image

Nefise Uysal, Radboud University Medical Center

11. Generalization of low-field 3D MRI acceleration via the CIRIM network across knee, spine and brain

Daisy van den Berg, Amsterdam University Medicl Center

12. In Silico Experimentation allows optimization of Photoacoustic Imaging Hardware for Thyroid Nodule Risk Estimation

Max Rietberg, University of Twente

13. Leveraging ECG Foundation Models in Critical Care for Sinus Rhythm and Atrial Fibrillation Classification

Maria Galanty, University of Amsterdam

14. Leveraging the wearable 1-lead ECG signal: From cardiac rhythm to cardiac function assessment

Viktor van der Valk, Leiden University Medical Center

15. Multimodal learning with whole slide images and genomics for DCIS risk prediction

Marek Oerlemans, The Netherlands Cancer Institute

16. Non-invasive estimation of hemodynamic parameters in pulmonary hypertension with a deep learning approach integrating all B-mode cine loops in an echocardiographic exam

Li-Hsin Cheng, Leiden University Medical Center

17. Normative Deviations of Turbulent Dynamics for Predicting Depressive Symptoms

Jerke van den Berg, Amsterdam University Medical Center

18. Optimizing Photon Counting CT imaging for Interstitial Lung Disease assessment

Cristina Cretu, Erasmus Medical Center

19. Overcoming the Spatial-Temporal Trade-off in DCE MRI Using Recurrent Inference Machines

Dilara Tank, Amsterdam University Medical Center

20. Quantitative Mapping From Conventional MRI Using Self-Supervised Physics-Guided Deep Learning: Applications To A Large-Scale, Heterogeneous Clinical Dataset

Jelmer van Lune, University Medical Center Utrecht

21. Self-Supervised Deep Learning for Label-Free Brain Metastasis Detection in Clinical MR Imaging

Anne Rückert, University Medical Center Utrecht

22. Towards Automated Quality Control of Neonatal 7T MRSI using Unsupervised Algorithms

Dennis van de Sande, Eindhoven University of Technology

23. Towards high-performance breast diffusion MRI: evaluation of temperature, nerve stimulation, and diffusion measurements with a prototype breast gradient insert

Leon Arends, University Medical Center Utrech

ORANISER

The Division of Image Processing, in Dutch abbreviated as the LKEB (Laboratorium voor Klinische en Experimentele Beeldverwerking), is a research group within the Department of Radiology, LUMC. The director of LKEB is Prof. Boudewijn P.F. Lelieveldt PhD.

We perform fundamental and applied research in the area of biomedical image processing and analysis. We execute extensive validation studies of the developed techniques, both technically and clinically. We aim to impact the healthcare system by bringing research results close to the clinic, through collaboration with clinicians as well as with industry. While most of our research focuses on medical imaging, we are also interested in biological and genetic data. We have worked in particular on CT, MRI, IVUS and OCT imaging, in the brain, chest, heart, abdomen, vasculature and bones. In addition, we investigate highly heterogeneous data such as omics, imaging and mass-spectrometry combined, e.g. to produce new insights about structural and functional organization of the brain.

Methodologically we develop artificial intelligence and machine learning technologies such as deep learning. We combine these with general computer science methods and mathematical modeling. Our work encompasses segmentation, registration, quantification, visualization, radiomics, etc. To bring our research results close to a clinical end-user we develop high quality software. Our scientific programmers therefore work through a formalized Software Development Process (SDP), often in close collaboration with industry, in particular with Medis medical imaging systems BV. This has led to a number of commercially successful products, and spinoff to a number of large medical imaging vendors. The open source image registration software elastix is also maintained at LKEB.

LKEB

LUMC

SPONSOR

For more than 30 years, Medis has been a trusted provider of high-quality quantitative analysis solutions for cardiovascular imaging. Its internationally recognized software seamlessly combines simplicity with precision, enabling ease of use while supporting clinical excellence for patients worldwide.

Quality and accuracy are central to Medis' approach. The company rigorously tests and validates its software to ensure the highest standards are consistently met, without compromise. Through meaningful innovation, Medis has contributed to significant improvements in the diagnosis and treatment of patients globally, underscoring its ongoing commitment to advancing healthcare outcomes.

For over 30 years, Medis has provided high-quality quantitative analysis solutions for cardiovascular imaging to the medical community. Founded in 1989 as a spin-off from Leiden University Medical Centre (LUMC, the Netherlands), the company was established by its founder and current CSO, Hans Reiber. From the outset, Medis focused on X-ray–based cardiac image analysis, driven by the mission of making advanced analytical tools accessible to medical researchers and clinical specialists worldwide.

The company's heritage and core values are rooted in this strong purpose of contributing to a healthier society by delivering reliable and innovative tools to cardiologists, radiologists, researchers, and industry partners.

Medis
MEDICAL IMAGING

Improving the lives of patients by empowering medical professionals with the ultimate cardiovascular imaging solutions.

MISP^2 | 27 January 2026 | Leiden University Medical Center | Leiden