Learning Image Reconstruction with Domain Adaptation

Ultrasound is a non-ionizing, portable, low-cost, and real-time medical imaging modality, and therefore it is commonly used for many applications in the clinic. It operates by sending acoustic waves into the tissue, and converting the received echo signals into images. Typical clinical ultrasound imaging generates grayscale maps of reflected echo *magnitude*. Although informative for certain clinical tasks, these do not contain quantitative information and often yield low specificity, e.g., for differential diagnosis of cancer. Nevertheless, ultrasound data carries also other valuable information that may be extracted using novel methods. To that end, we presented the mapping of local tissue speed-of-sound (SoS) from differential measurements of acoustic arrivals viewed from different directions, by formulating this as an inverse problem of image reconstruction[1,2]. Such SoS images can show tissue compositions, e.g. cancerous tumors, in a quantifiable way and with much higher contrast than conventional ultrasound, cf. the Figure. We demonstrated preliminary clinical uses for breast cancer[3], breast density[4] (for risk assessment), and sarcopenia[5] (age-related muscle degeneration). Nevertheless, the ill-posed nature of this problem as well as noise and nonlinearities in the measurements have so far prevented a successful and wide-spread clinical application.

Deep learning (DL) may provide an attractive solution for the above image reconstruction problem, as DL techniques have been very successful in many signal processing and computer vision problems such as classification and segmentation, in particular in supervised learning settings. However, image reconstruction poses particular challenges for DL. With image reconstruction aiming the estimation of a spatially-varying (and often novel) tissue parameter, a main obstacle lies in the fact that true supervised learning is not possible, since the ground-truth, i.e. the actual precise spatial distribution of such parameter in human body, is often unknown. For supervised learning, synthetic simulations can be an alternative, however for complex imaging processes, there is often no comprehensive simulation system available that fully represents the imaging physics, nor the electrical imaging setup and potential in-vivo parameter distributions. Consequently, *DL-based reconstructions trained from simulations do not generalize well to real data*. Accordingly, this interdisciplinary project aims to develop theoretical and practical solutions for medical image reconstruction, with domain adaptation approaches from synthetic to real data.

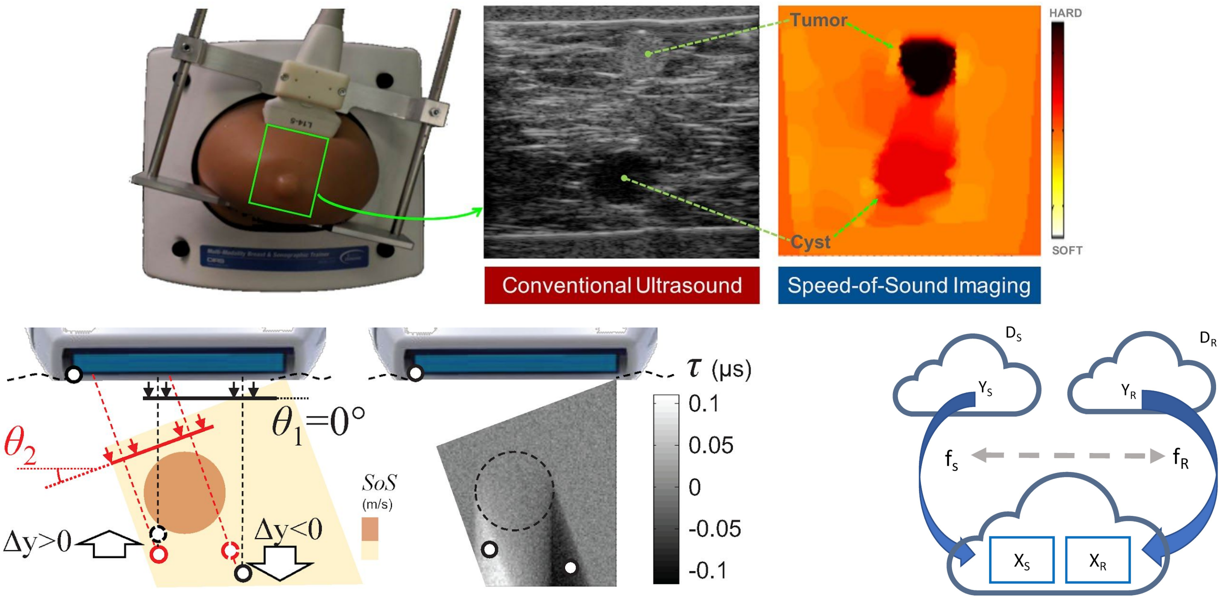

Figure: Speed-of-Sound Imaging (upper[1] and lower-left[2]). Illustration of domain adaptation from simulations (lower-right)

---

[1] Sanabria, Rominger, Goksel: "Speed-of-Sound Imaging Based on Reflector Delineation", *IEEE Trans Biomed Eng*, 2019.
[2] Sanabria, Ozkan, Rominger, Goksel: "Spatial Domain Reconstruction for Imaging Speed-of-Sound ...", *Physics in Med & Biology*, 2018.
[3] Ruby et al.: "Breast Cancer Assessment With Pulse-Echo Speed of Sound ...: Proof-of-Concept", *Investigative Radiology*, 2019.
[4] Sanabria et al.: "Breast-Density Assessment with Handheld Ultrasound: ...", *European Radiology*, 2018.
[5] Sanabria et al.: "Speed of sound ultrasound: a pilot study on a novel technique to identify sarcopenia in seniors", *European Rad*, 2019.

Project Description and Goals

There have been various image reconstruction approaches with DL[6], such as by end-to-end learning (e.g. with fully-connected layers as in AUTOMAP[7]), by learning projections of the measurement space (to correct inputs, e.g., to better fit an imaging model), by learning projections of reconstructions (for retrospective corrections in outputs), or using *variational networks* (VNs). Given supervised training, these methods were shown to work with some level of success. End-to-end approaches of learning blackbox functions to map measurements to reconstructions require supervision with huge numbers of training samples, and are sensitive to changes in data distributions, i.e. *domain shift*. With VNs for *loop unrolling*, neural network layers are used to represent iterations of a very general form of imaging inverse problem, in order to learn an optimal problem parametrization. Having first introduced VNs in the CT setting[8], we demonstrated their use for the first time for USCT as well as X-ray CT[9]. In recent work, multiple simulation domains were used for training VNs in order to increase their generalizability to real data[10]. Nevertheless, *it is yet an unsolved problem to develop readily generalizable methods for learned image reconstruction*.

Our project goal is the image reconstruction with real data, despite the lack of ground-truth for supervised learning; i.e. estimating a function f_R to map given measurement y_R to reconstructions x_R as shown in the Figure. Since we have access to simulated ground-truth x_S with the corresponding measurements y_S we will frame the task as a problem of domain adaptation[11] from simulations to real data. Although domain adaptation has been studied for classification and regression problems, there is much space and need for theoretical and applied research on it for image reconstruction problems. Although a specific case of the above problem can be treated as aligning the measurement domains y_S with y_R, note that the ultimate goal in a general setting is to infer the mapping f_R (e.g. based on the mapping f_S), without having access to labeled samples of x_R.

In contrast to typical domain adaptation problems in computer vision that often target photorealistic appearance, expectations of accuracy and confidence intervals are much higher in medical imaging, where the results may affect diagnostic and intervention decisions. We will accordingly investigate such application-specific concerns, such as learning bounds, quantifying confidence/uncertainty in reconstructions, and suitability of potential distance metrics (e.g. whether popular L^1, maximum mean discrepancy, or optimal transport based metrics such as the Wasserstein distance are appropriate for assessing domain distances in medical data).

We will approach the given problem from multiple aspects, and aim for unifying solutions as the project progresses: 1.) In one line of work, we will investigate *adversarial* approaches for domain adaptation[12,13] and unsupervised reconstruction[14]. 2.) In another line of work, we will investigate image-to-image (I2I) translation[15,16,17] and develop fundamental relations and analyze theoretical bounds for treating image reconstruction as a specific case of I2I. 3.) To unify the earlier two, we will investigate learning bounds and guarantees of domain adaptation theory[18,19].

Significance

The above ultrasound computed tomography (USCT) belongs to a group of image reconstruction problems called computed tomography (CT), which is relevant and commonly used for a wide range

---

[6] Ongie et al.: “Deep Learning Techniques for Inverse Problems in Imaging”, *IEEE J Selected Areas in Information Theory,* 2020.
[7] Zhu et al.: Image reconstruction by domain-transform manifold learning”, *Nature*, 2018.
[8] Vishnevskiy, Sanabria, Goksel: "Image Reconstruction via Variational Network for ... Sound-Speed Imaging", *MLMIR,* 2018.
[9] Vishnevskiy, Rau, Goksel: "Deep Variational Networks with Exponential Weighting for Learning Computed Tomography", *MICCAI,* 2019.
[10] Bernhardt, Vishnevskiy, Rau, Goksel: "Training Variational Networks with Multi-Domain Simulations ...", *IEEE TUFFC,* 2020.
[11] Wang and Deng: “Deep Visual Domain Adaptation: A Survey”, *Neurocomputing,* 2018.
[12] Ganin et al.: “Domain adversarial training of Neural Networks”, *J Machine Learning Research*, 2016.
[13] Hoffman et al.: “CyCADA: Cycle-Consistent Adversarial Domain Adaptation”, *ICML*, 2018.
[14] Pajot et al.: “Unsupervised Adversarial Image Reconstruction”, *ICLR* 2019.
[15] Murez et al.: “ Image to Image Translation for Domain Adaptation”, *CVPR*, 2018.
[16] Zhu et al.: “Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks”, ICCV, 2017.
[17] Liu et al.: “Unsupervised Image-to-Image Translation Networks”, *NIPS*, 2017.
[18] Zhang et al.: “Bridging Theory and Algorithm for Domain Adaptation”, *PMLR*, 2019.
[19] Redko et al.: “A Survey on Domain Adaptation Theory: Learning Bounds and Theoretical Guarantees”, 2020.

of applications, such as X-ray and PET imaging in medicine, as well as in geophysics and crystallography. Thus, studying the domain adaptation problem in USCT will have far-reaching consequences for other imaging modalities sharing the same fundamental challenges herein. In particular, a linear measurement / forward model for CT only assumes generic (positive definite) matrices, where several other popular computer vision problems can be written as its special cases; e.g. magnetic-resonance imaging (MRI) using a unitary (discrete Fourier) transform; image super resolution using a fat block-diagonal matrix; and image denoising using identity transform. Hence, our results for CT can provide solutions for domain adaptation using synthetic data in several other prominent problem settings.

Interdisciplinarity

Relying on expertise and know-how on signal processing, computer vision, engineering, medical physics, tissue biomechanics, and clinical imaging, this project is multi-faceted and interdisciplinary enterprise. The problem of domain adaptation for image reconstruction is rooted in mathematical foundations from probability theory, information theory, statistics, image processing, inverse problems, optimization, function analysis.

Not only will we develop a number of theoretical and mathematical techniques to address image reconstruction problems, but we will also apply these solutions in the CT reconstruction setting, in particular for our USCT imaging method. This project is a collaboration between the Computer-assisted Applications in Medicine at the Department of Information Technology, the Signal and Systems at the Department of Electrical Engineering, the Medtech Science and Innovation Centre, and the Uppsala University Hospital. With its close connections to the clinic, such project setup will facilitate the translation of the developed solutions to the application field.

A PhD student with a master's degree in Mathematics, Computer Science, Engineering, Image Analysis, Machine Learning, or a related field will be recruited, with a background and interests in mathematical foundations as well as experimental and programming aspects. A good command of mathematical optimization, probability theory, computer vision, and/or signal processing is a plus.

Advisors and Host Institute

Orcun Göksel[20] (main advisor), Computer-assisted Applications in Medicine http://goksel.org/
Department of Information Technology, Uppsala University &
Medtech Science and Innovation Centre, Uppsala University

Ayca Özcelikkale[21] (co-advisor), Signals and Systems https://sites.google.com/site/aycaozcelikkale
Department of Electrical Engineering, Uppsala University

Per Adolfsson (co-advisor), Head of Medical Technology, Hospital Physics and IT (MTF)
Uppsala University Hospital https://www.akademiska.se/om-oss/organisation/sjukhusledning/

The student will be co-financed by and hosted at the Department of Information Technology, while being cross-affiliated at the Medtech Science and Innovation Centre (https://medtech.uu.se/) co-located with MTF at the University Hospital; facilitating the clinical translation of project results.

[20] Project-related expertise and interests on inverse problems, image reconstruction, and computer vision in medical imaging, in particular for ultrasound, deep learning, probabilistic methods, uncertainty, and interpretability.

[21] Project-related expertise and interests on signal and image processing, statistical methods, compressed sensing and sparse recovery, sampling, information theory, generalization and performance bounds, reinforcement learning.