Physikalisch-Technische Bundesanstalt
Braunschweig und Berlin
Nationales Metrologieinstitut

Active learning for medical imaging using Fisher information

Type: Postdoc or PhD

Principal Investigators

- Jörg Martin (PTB, AG 8.42)
- Clemens Elster (PTB, AG 8.42)

Background

Active learning aims at selecting training data in a way such that the robustness of the trained net is maximized and/or the required number of training data is kept at a minimum. Active learning is a current line of research in deep learning (see [1-3]) and particularly relevant in cases where labeling is expensive or time consuming. This is particularly relevant in medical applications, where human experts are needed. State-of-the-art approaches are based on minimizing predictive uncertainty [2] or optimal exploration of feature space [3].

PTB-8.42 has implemented and tested current approaches of active learning for regression problems related to the contrast-detail curve determination in mammography image quality assessment. While clear benefits have been observed in some cases, the approaches often lack robustness, and in some cases even fail to improve training performance. Furthermore, the number of initial, randomly chosen, training samples has a large impact and strategies for its determination are still lacking.

Active learning is closely related to experimental design in statistics [4] for which the Fisher information matrix plays a key role. Current active learning approaches do not directly refer to those concepts, and the Fisher information matrix is rarely used in this context. One reason is that the calculation of the Fisher matrix is practically impossible for typical neural nets.

PTB-8.42 has successfully made use of the Fisher information matrix for detecting adversarial examples and out-of-distribution behavior [5,6]. The key idea was to use suitably selected projections of the Fisher matrix which scale well and can be efficiently calculated even for many millions of unknowns in the neural net.

Project Aim, Objectives and Program

The goal of the project is to design methods for active learning based on the Fisher information and to provide means for their efficient implementation. Specific objectives are to explore

- appropriate directions of projection of the Fisher information matrix for optimal active learning together with efficient calculation schemes,
- strategies for the choice of initial samples prior to starting the active learning strategy based on Fisher information,
- potential benefits of the Fisher information matrix approach by comparison with current active learning strategies for the two applications (i) mammography image quality assessment and (ii) ECG diagnostics.

Work program

- Preparation of databases for assessing active learning strategies for the two considered applications, and selection of network architectures to be considered.
- Development of suitable projections of the Fisher information matrix and efficient calculation schemes.
- Application of the developed active learning strategy based on the Fisher information for the two applications and comparison with results obtained by current approaches.

Available data

- Virtual mammography phantom data (>40 k, > 1 M possible), already available
- Labeled ECG data base published by PTB (> 20 k), publicly available

Candidate Requirements

- PhD and/or MSc in applied mathematics, computer science, or similar
- Experience in at least one of the fields of deep learning or statistics
- Software experience: Python (optimally in PyTorch) and joint development (Git)

References

[1] Mizokami S. *Deep active learning from the perspective of active learning theory*. Deep active learning. Springer, Singapore, 2018. 79-91.
[2] Gal Y, Islam R, Ghahramani Z. *Deep Bayesian Active Learning with Image Data*. ICML 2017: 1183-1192
[3] Wu D, Lin C-T, Huang J. *Active learning for regression using greedy sampling*. Inf. Sci. 474: 90-105 (2019)
[4] Pukelsheim F. *Optimal design of experiments*. Society for Industrial and Applied Mathematics, 2006.
[5] Martin J, Elster C. *Inspecting adversarial examples using the fisher information*. Neurocomputing 382: 80-86 (2020)
[6] Martin J, Elster C. *Detecting unusual input to neural networks*. Applied Intelligence (2020): 1-12.