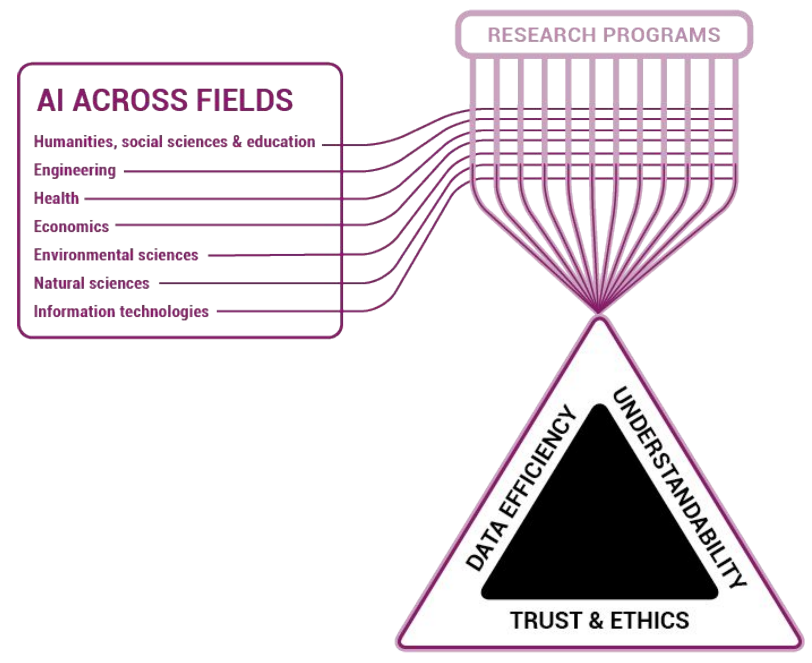

Next-generation data-efficient deep learning (Research Program R3)

FCAI Research Program R3 develops reliable methods which harness the power of deep learning while achieving good results with less training data and less human supervision

Program objectives

Deep learning has been the driver of the remarkable recent progress in multiple domains such as speech recognition, image processing, machine translation, reinforcement learning. In this program, we build the next generation of deep learning algorithms which allow:

- **Learning in a data-efficient manner**. Current deep learning relies on large training data sets which are often difficult to obtain in many practical applications.
- **Transparency and reliability**. While not strengths of conventional deep learning, these are key to practical applicability. We develop deep learning methods with probabilistic principles aiding uncertainty quantification and model robustness.
- **Understandability**. Current deep learning still lacks understanding of the modelled domain, it has little interpretability, can easily be fooled by adversarial examples, and has no reasoning capabilities.

Methodologies

1. Sample-efficient deep learning

We develop deep learning methods that are less data-hungry. Our approach includes semi-supervised learning, few-shot learning (learning to learn new tasks from few examples), active learning (*e.g.*, efficient strategies for data labeling), combining domain knowledge (*e.g.*, rules, first principles or simulators) with data-driven models, use of auxiliary data, and both model-based and model-free reinforcement learning (RL) methods.

2. Deep learning with probabilistic principles

We combine probabilistic methods with deep learning. This includes aspects of Bayesian deep learning, links to non-parametric modelling, and developing model architectures which are more robust to outliers and missing training samples, and capable of quantifying the uncertainty of the model predictions.

3. Domain understanding

Current deep learning is good at producing statistically probable answers but it does not develop proper “understanding” of how the world works. It does not know causal relations between concepts, it cannot reason and explain its decisions in a logical manner. We work on deep learning algorithms that combine perception with reasoning capabilities.

4. Generative design

Generative design is an iterative design process that involves a computer program generating design recommendations to a human designer. By design here we mean a process of creating new materials, floor-plans, engines, games, and so on. We aim at developing deep learning methods that can be in the core of such a design recommender system.

5. Deep learning for new domains

We develop deep learning algorithms that can be applicable in new domains where data are less structured and scarce. Examples include medical records, financial documents, equipment specifications, computer logs, *etc.*

Recent research highlights

L. Meronen, C. Irwanto, A. Solin. Stationary Activations for Uncertainty Calibration in Deep Learning, NeurIPS 2020.

This work introduces a new family of non-linear neural network activation functions that mimic the properties induced by the widely-used Matérn family of kernels in Gaussian process (GP) models. Matérn activation functions result in similar appealing properties to their counterparts in GP models, and shows both good performance and uncertainty calibration in Bayesian deep learning tasks, such as out-of-distribution (OOD) uncertainty.

For the proposed method, uncertain known and certain unknown samples feel intuitive (good calib.), while the results seem arbitrary for the ReLU.

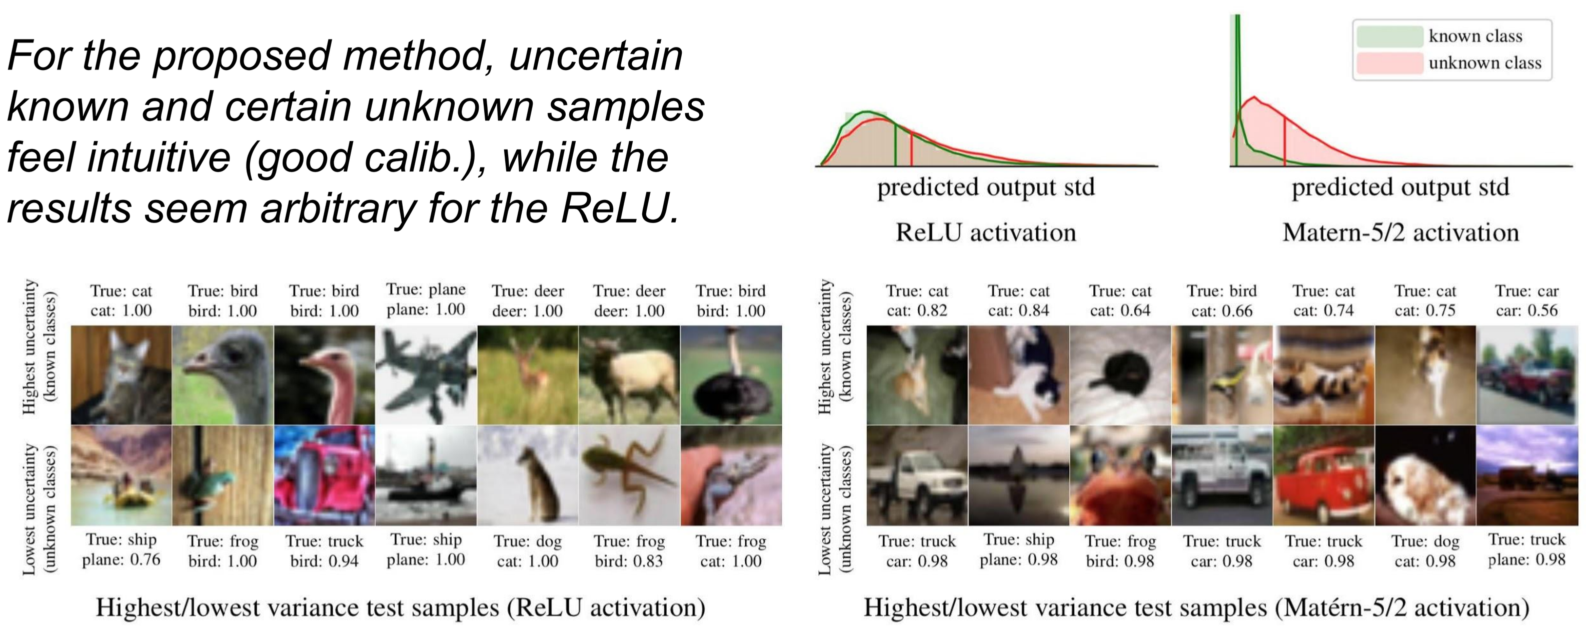

X. Li, S. Wang, Y. Zhao, J. Verbeek, J. Kannala. Hierarchical Scene Coordinate Classification and Regression for Visual Localization, CVPR 2020.

Visual localization is critical to many applications in computer vision and robotics. We present a new hierarchical scene coordinate network to predict pixel scene coordinates in a coarse-to-fine manner from a single RGB image. This hybrid approach outperforms existing scene coordinate regression methods and reduces significantly the performance gap w.r.t. explicit feature matching methods.

A. Heljakka, Y. Hou, J. Kannala, A. Solin. Deep Automodulators, NeurIPS 2020.

This paper introduces a new category of generative autoencoders called automodulators. These networks can faithfully reproduce individual real-world input images like regular autoencoders, but also generate a fused sample from a combination of such images, allowing instantaneous ‘style-mixing’ and other new applications.

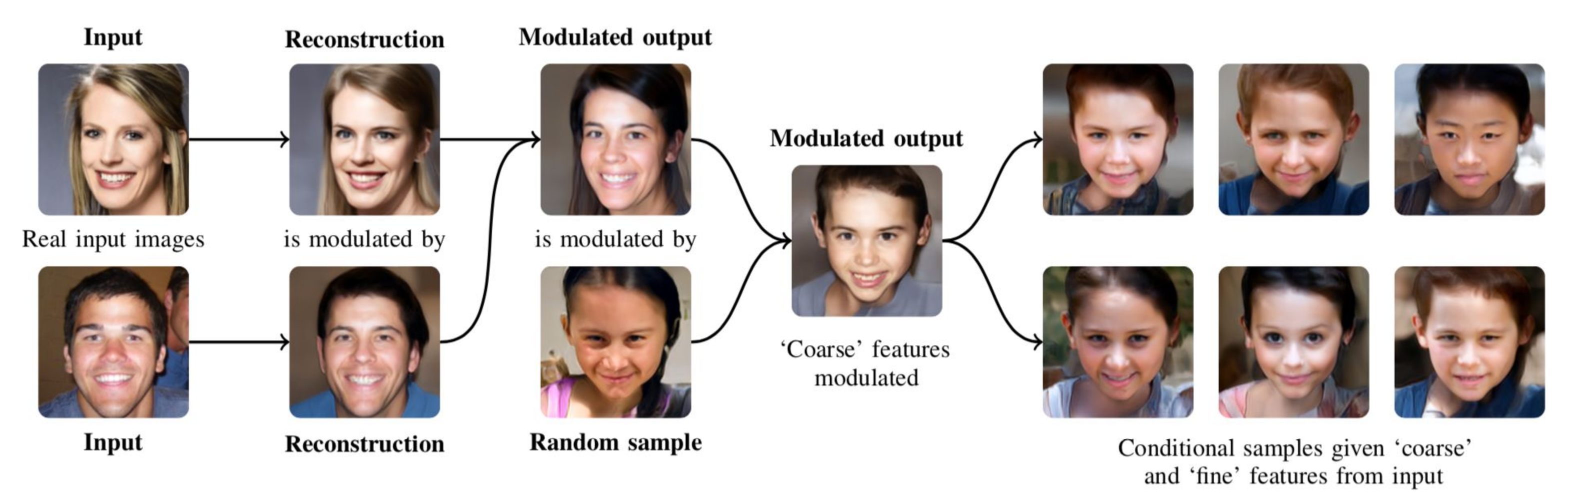

Coordinating professor

Arno Solin

Assistant professor, Machine learning

Aalto University

arno.solin@aalto.fi