TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Electronic Systems Group - Department of Electrical Engineering

Maximizing Performance of Signal Processing Algorithms using Approximate Bayesian Inference

Martin Roa Villescas (m.roa.villescas@tue.nl)

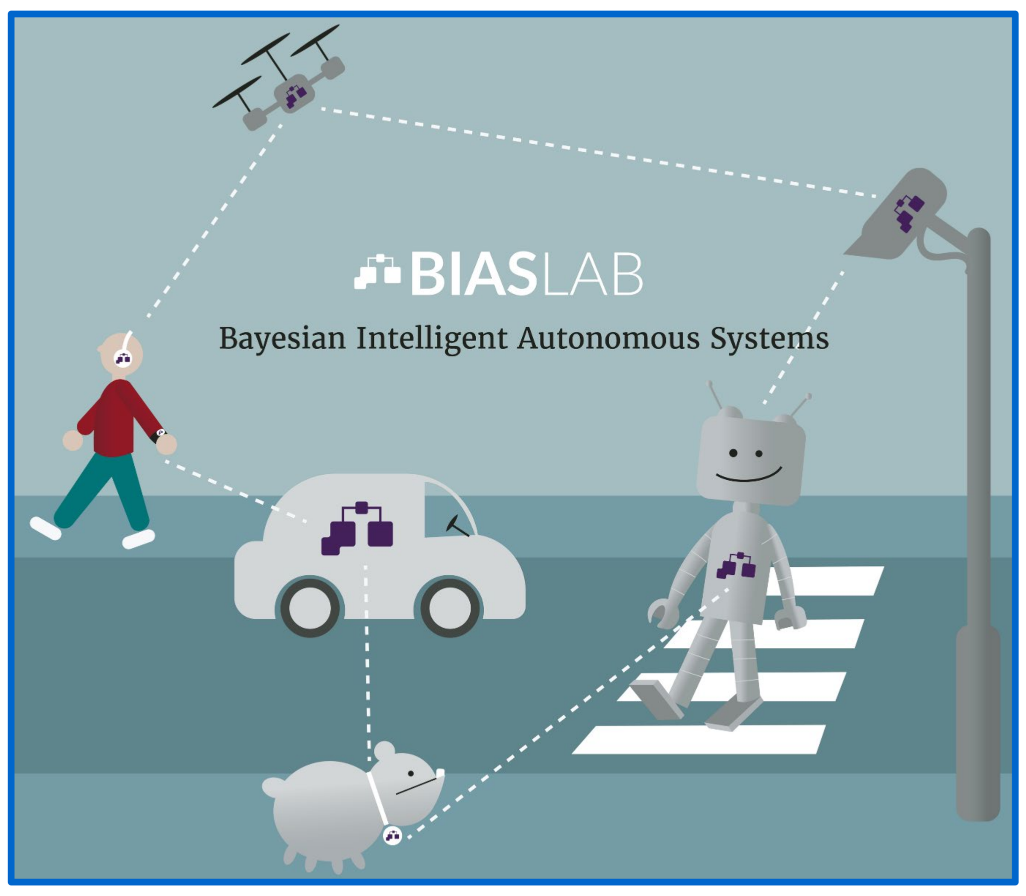

Research Problem - Developing a principled design method for balancing the *performance vs. resource availability* trade-off for signal processing algorithms.

Research Context - This research effort is executed in the context of developing high-performance algorithms for extremely constrained platforms such as IoT and wearable devices.

Applications - An important application is to automate the design of audio processing algorithms for hearable devices, but we are also interested in agent-based learning of gestures, motion analysis, etc.

Cooperation - Large part of the ongoing research is carried out in cooperation with BIASlab, an academic research team in the Signal Processing Systems group.

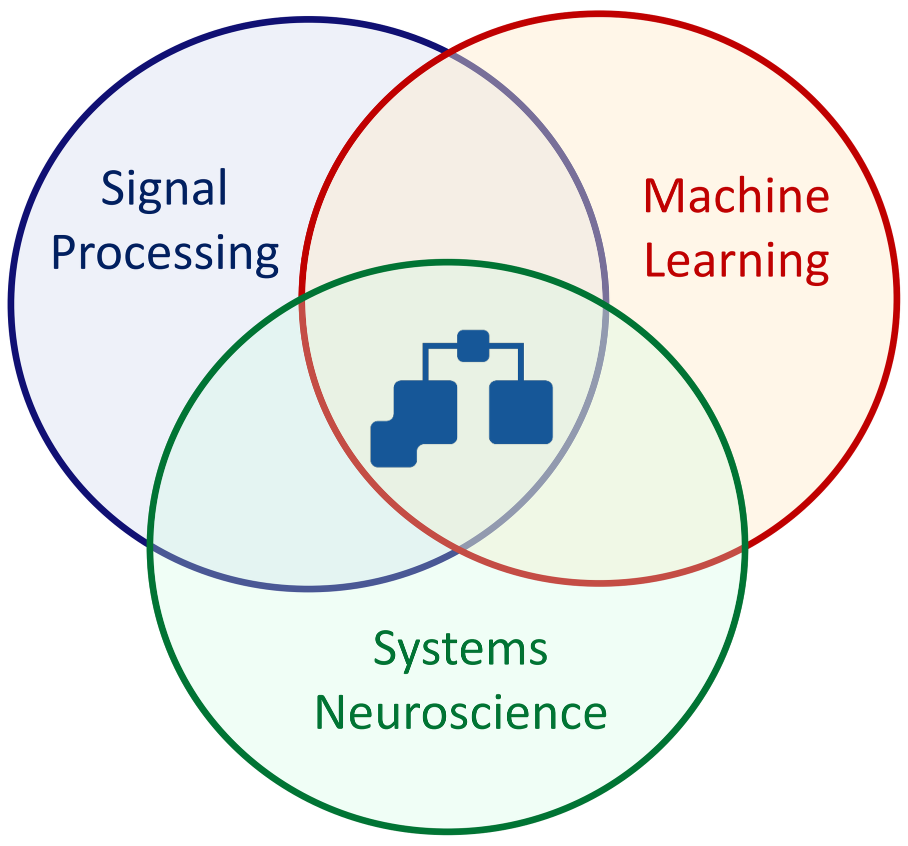

Underlying Foundations

Signal Processing Algorithms as Inference Tasks in Graphical Models

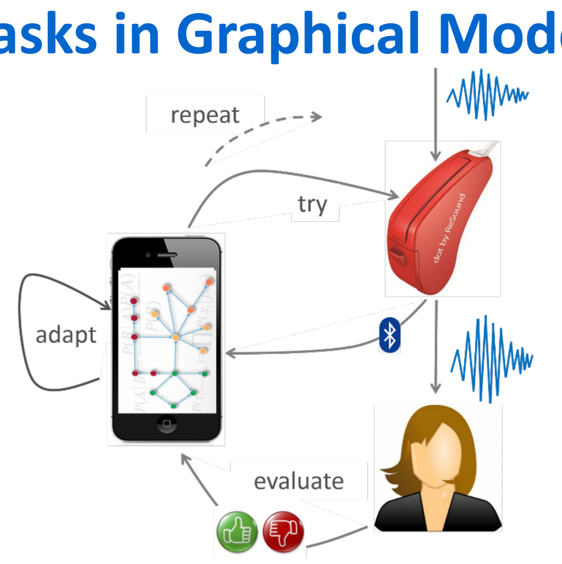

Signal processing algorithms are framed as inference tasks in generative graphical models. This involves: 1. Defining the model, 2. Inferring the hidden states and 3. Comparing the performance of various models and selecting the best one.

Approximate Inference via Variational Bayes

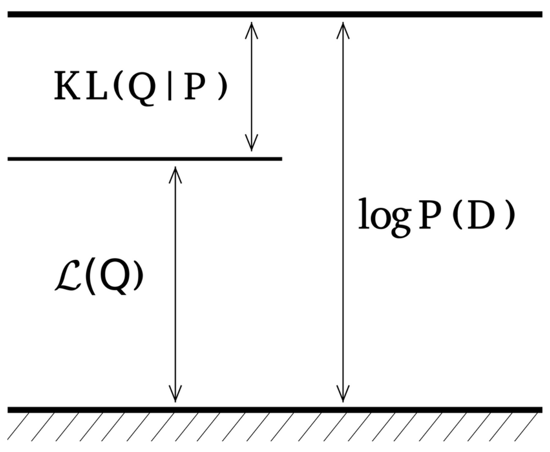

Exact Bayesian inference often involves intractable numerical computations. We therefore resort to an approximate inference framework called variational Bayes that turns the inference problem into a simpler optimization problem: minimization of free energy.

Message Passing Algorithms in Forney Style Factor Graphs (FFG)

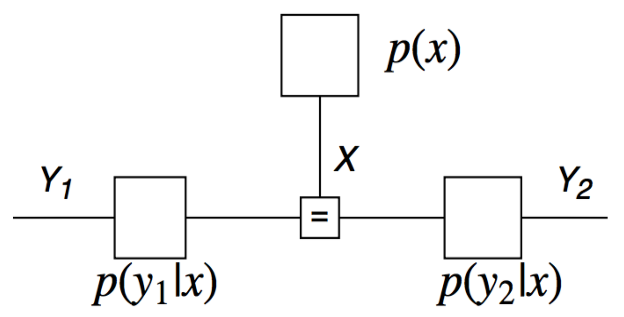

All inference tasks are carried out as message passing algorithms in FFGs. This has the advantage of enabling the automation of the signal processing design. In effect, designers have faster design cycles where the main focus is placed on the solution of the problem at hand.

The Hierarchical Gaussian Filter (HGF)

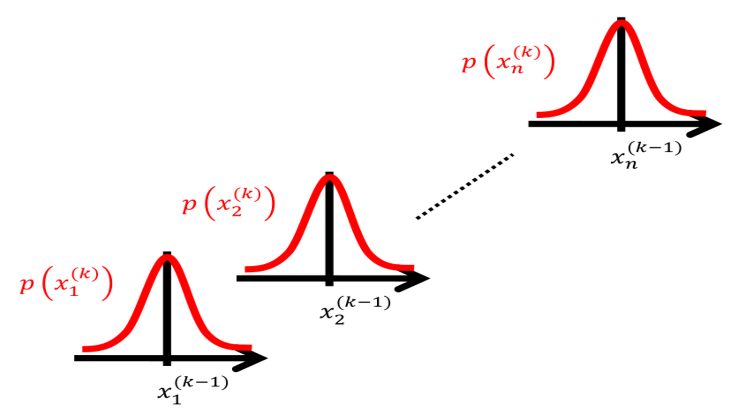

We use the HGF (Mathys, 2012) as a driver case for our investigation. The HGF is a popular model in the computational neuroscience literature that aims to explain mechanisms of adaptive behavior. The model represents a hierarchy of Gaussian random walks, coupled through their variance.

Research Idea

Inspired by the functional behavior of the brain, we claim that the amount of energy that an agent should spend on learning the model of a system, should be based on how well that agent can explain the incoming observations of such system. In other words, an agent should not spend much energy if it can predict incoming observations with sufficient accuracy. On the other hand, if the incoming observations are not in line with the agent's beliefs about the behavior of the system (i.e. there is a high degree of surprise in the observations), then the model has to take measures that target the minimization of the surprise.

Autonomous Acoustic Systems (AAS)

This research effort is executed in the context of the AAS project. The goal of the AAS project is to further develop the next generation of embedded systems by offering adaptive signal processing algorithms that perform optimally within a stringent energy budget for autonomous operation.

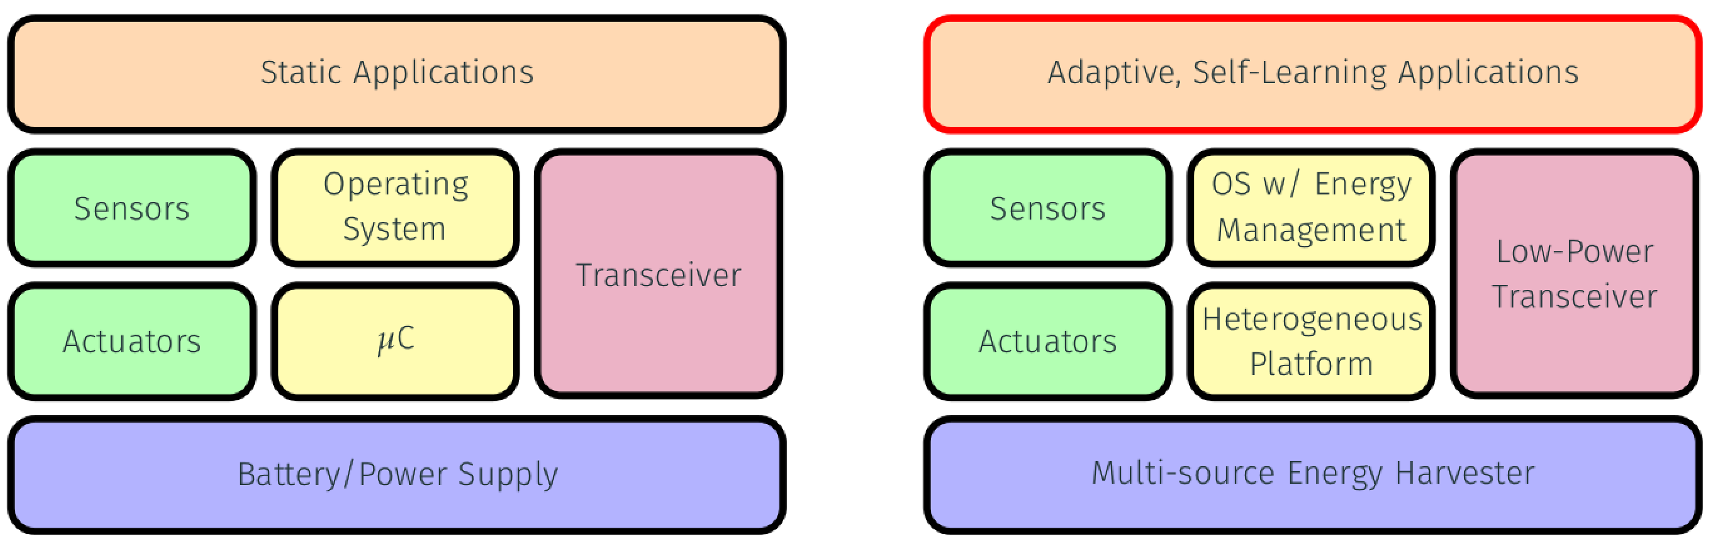

(a) Traditional Embedded System (b) ZERO-Energy Embedded System

Henk Corporaal Bert de Vries Sander Stuijk Patrick Wijnings Martin Roa Villescas