FCAI Finnish Center for Artificial Intelligence

fcai.fi

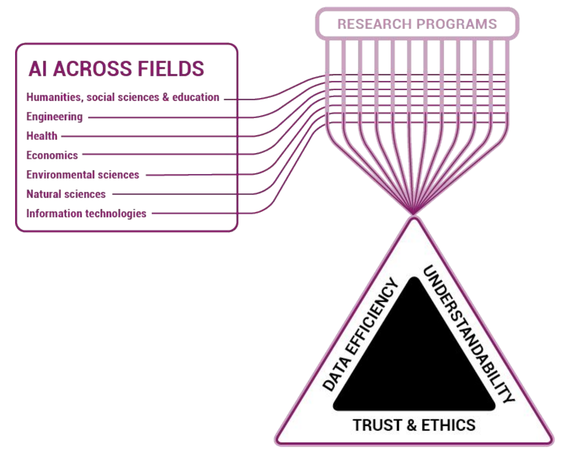

Agile probabilistic AI (Research Program R1)

FCAI Research Program R1 develops AI tools using probabilistic programming and Bayesian statistics

Program objectives

Our goal is to provide principled Bayesian workflow for data and model driven analysis of real world data. This includes three main objectives:

- Development of high-quality verified software for probabilistic programming
- Development of computationally efficient and accurate model-agnostic Bayesian inference algorithms
- Theory and tools for understanding the modelling workflow as a whole, integrating model specification, validation, refinement, and visualization

Methodologies

To reach our objectives, we apply a diverse set of methodologies:

- Analysing algorithmic convergence and robustness by formal mathematical analysis and extensive numerical simulations
- Benchmarking new algorithms and methods to existing state-of-the-art competitors
- Implementing developed methods in open-source software such as R, Python, and Stan
- Investigate the quality and the usability of developed software by means of empirical human-computer-interaction studies

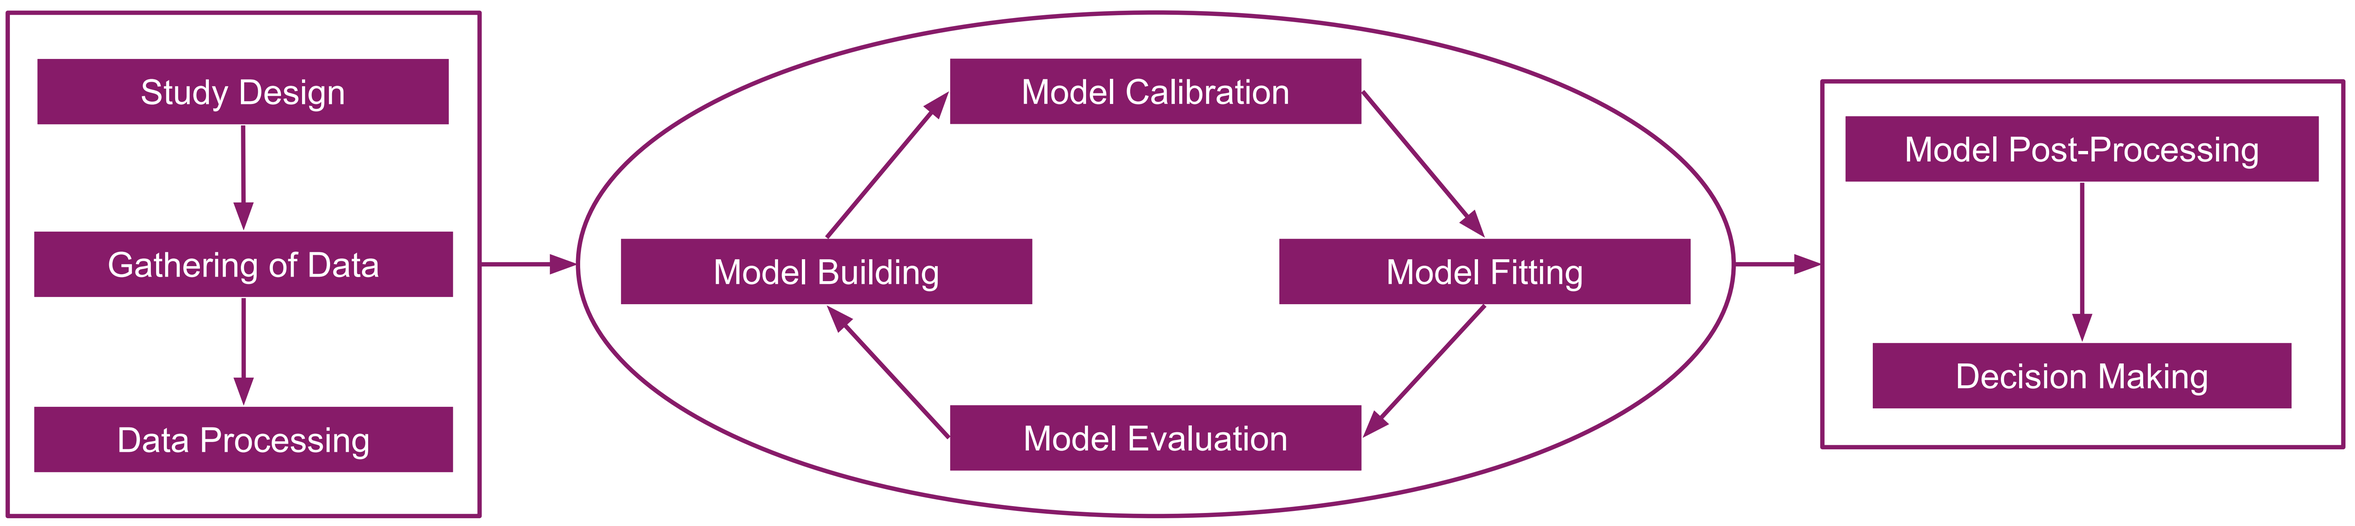

High level overview Bayesian workflow that can help to build better AI and AI can assist by automating some parts of the workflow.

Selected research results

- **Efficient and robust procedures for approximate Bayesian cross-validation** (Bürkner et al.: Approximate leave-future-out cross-validation for Bayesian time series models, *Journal of Statistical Computation and Simulation*, 2020; Paananen et al.: Implicitly adaptive importance sampling, *Statistic and Computing*, 2020, Magnusson et al. Leave-one-out cross-validation for Bayesian model comparison in large data, *AISTATS 2020*)
- **Efficient Bayesian variable selection in high dimensions** (Piironen et al: Projective inference in high-dimensional problems: prediction and feature selection. *Electronic Journal of Statistics*, 2020)
- **More reliable and faster approximate inference for probabilistic programming** (Dhaka et al. Robust, accurate stochastic optimization for variational inference. *NeurIPS 2020*; Margossian et al. Hamiltonian Monte Carlo using an embedded Laplace approximation. *NeurIPS 2020*, Vehtari et al. Rank-normalization, folding, and localization: An improved R-hat for assessing convergence of MCMC, *Bayesian Analysis*, 2020)
- **Interpretable AI** (Afrabandpey et al. Making Bayesian predictive models interpretable: A decision theoretic approach. *Machine Learning*, 2020)
- **Methods for understanding and improving decision making when using approximate inference** (Kusmierczyk et al. Variational Bayesian decision-making for continuous utilities, *NeurIPS 2019*; Kusmierczyk et al. Correcting predictions for approximate Bayesian inference, *AAAI 2020*).
- **Interactive prior elicitation for reliable and low-effort quantification of expert knowledge** (Hartmann et al. Flexible prior elicitation via the prior predictive distribution, *UAI 2020*)

Coordinating professor

Aki Vehtari

Associate professor in computational probabilistic modeling

Aalto University

aki.vehtari@aalto.fi