technische universität dortmund

SFB 876 Providing Information by Resource-Constrained Data Analysis

Project A1

Data Mining for Ubiquitous System Software

Prof. Dr. Katharina Morik, Prof. Dr. Olaf Spinczyk

Problem

Machine Learning under Resource-Constraints

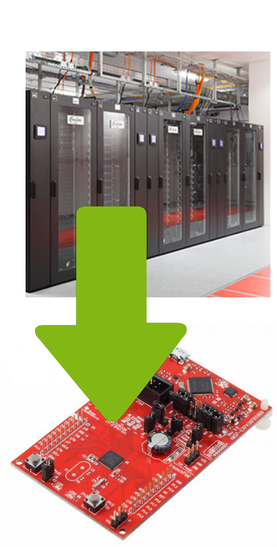

Exponential family models are . . .

- Derived from first principles
- Well-structured

$$\mathbb{P}(\boldsymbol{x}) = \exp(\langle \boldsymbol{\theta}, \phi(\boldsymbol{x}) \rangle - \ln Z(\boldsymbol{\theta}))$$

- Flexible

- Able to aggregate data set $\mathcal{D}$ into finite vector whose dimension is independent of $|\mathcal{D}|$.

- Adapt **model** to hardware architecture
- Need to **reduce arithmetic**, **parameter**, and **computational** complexity **with guarantees**

Resource-aware OS Data Acquisition

Existing tools for Operating System Data Acquisition are . . .

- Highly specialized for their use-case,
- Allow direct access to data structures via imperative programming, or
- Only cover either state or events.

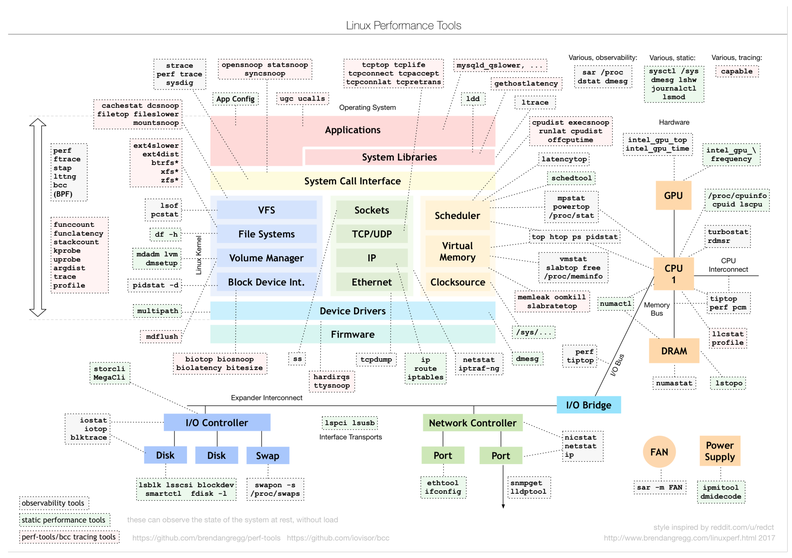

No **common** and **safe interface** to operating system data **and** operating system events exist.

No **unified** way to query **user-space** and **kernel** data simultaneously.

Methodology

Stochastic Quadrature

Probabilistic inference via **discrete Clenshaw-Curtis quadrature** [Piatkowski; Morik; ICML; 2016], [Piatkowski; Morik; UAI; 2018]

$$\hat{\psi}_k(\boldsymbol{x}) \approx \psi_k(\boldsymbol{x}) = \exp(\langle \boldsymbol{\theta}, \phi(\boldsymbol{x}) \rangle)$$

$$\hat{Z}_k(\boldsymbol{\theta}) = \int_{\mathcal{X}} \hat{\psi}_k(\boldsymbol{x}) \, \mathrm{d}\, \nu(\boldsymbol{x}) = \tau \mathbb{E}_{I,\boldsymbol{J}} \left[\sigma(\boldsymbol{c}_I) \prod_{r=0}^{I} \boldsymbol{\theta}_{\boldsymbol{J}(r)} \right]$$

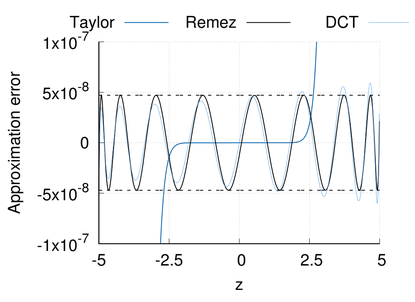

- **Numerical** instead of **structural** approximation keeps conditional independences intact
- New approach to **normalization** and **marginalization** with **provable approximation guarantees**
- Trade **resource-consumption** against **quality** via polynomial degree

Integer Exponential Families

Undirected probabilistic models with **integer-only arithmetic** [Piatkowski; Lee; Morik; Neurocomputing; 2016]

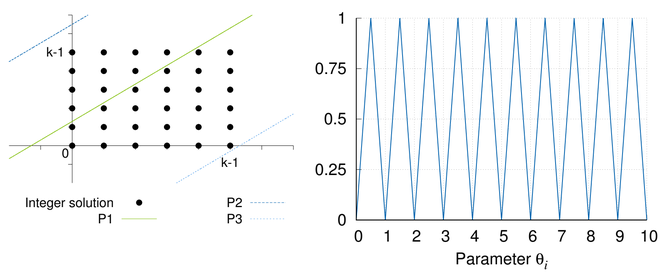

- **Integer-regularization enforces integrality while minimizing objective over $\mathbb{R}^d$**
- **Integer gradient descent algorithm** is **guaranteed** to converge
- Probabilistic inference via **bit-length propagation** (no floating point OPs)

Kernel Continuous Query Language

Relational data stream interface to **kernel** [Streicher; Lochmann; Spinczyk; TRIOS@SOSP; 2015]

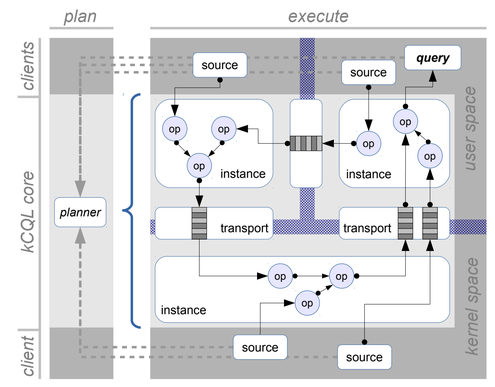

- Merge asynchronous **streams** and **relations** from operating system events and state
- Supports queries **across address spaces**

Results

SQM-based for ML Estimation

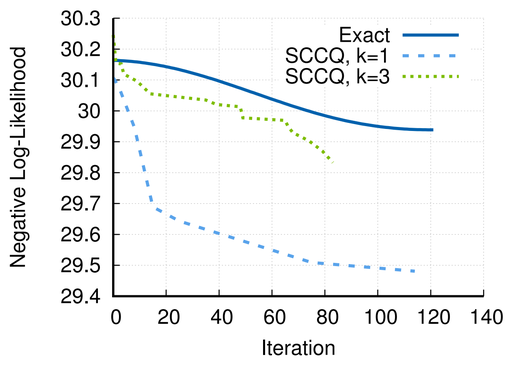

Fast Monte-Carlo Sampler

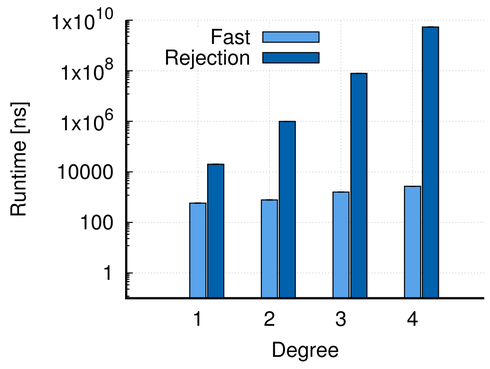

Runtime of Integer Inference

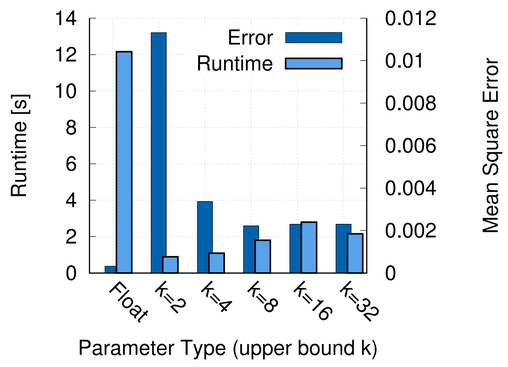

Lightweight Data Acquisition

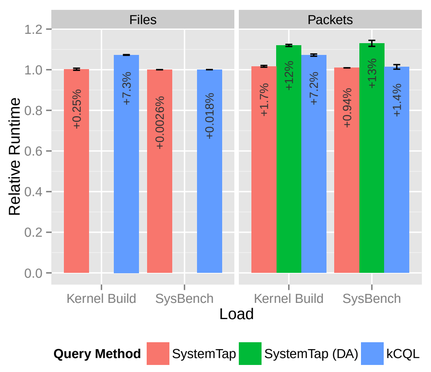

Learning on Devices

- Reproducible load tests with **trace-based benchmarks** [Lochmann; Bruckner; Spinczyk; LTB@ICPE; 2017]
- Decision trees and random forest for **field programmable gate arrays** [Buschjaeger; Morik; CIS-I; 2018]
- Approximate inference via **bit-length propagation** on **ultra-low-power** MCU (250× speed-up) and FPGA [Piatkowski; Diss.; 2018]
- Summary extraction on **data streams** for **embedded systems** [Buschjaeger; Morik; ...; IOTSTREAMING@ECMLPKDD; 2017]

Transfer

- **Dynamic route planning** with **STRF** [Liebig; Piatkowski; Bockermann; Morik; InformationSystems; 2017] B4 C3
- **Vehicle classification** via **radio-fingerprinting** [Sliwa; Piatkowski; ...; Wietfeld; ITSC; 2018] B4
- **Estimation of uplink transmission power** [Falkenberg; Sliwa; Piatkowski; Wietfeld; VTC; 2018] A4 B4
- **Deconvolution algorithms** for Cherenkov astronomy [Bunse; Piatkowski; Ruhe; Rhode; Morik; DSAA; 2018] C3

Methods

- **Proximal alternation linearized minimization** for **Boolean** matrix factorization [Hess; Morik; Piatkowski; DAMI; 2017] C1
- Controlling the **false discovery rate** in Boolean matrix factorization [Hess; Piatkowski; Morik; SDM; 2018] C1
- Approximate **deadline miss probabilities** for **embedded systems** [v.d.Brueggen; Piatkowski; Chen; Chen; Morik; ECRTS; 2018] B2
- Interpretable domain adaptation via **optimization over the Stiefel manifold** [Poelitz; Duivesteijn; Morik; ML; 2016]

Artificial Intelligence Group

CS 12 computer science 12