A GPU-Accelerated Skeleton Discovery for Gaussian Distribution Models beyond GPU Device Memory Capacity

Motivation

The estimation of causal graphical models allows to solve important problems in many different domains. For example in genetics, gene regulatory networks can be seen as a practical embodiment of systems biology and can be used for drug design or diagnostics. While GPUs show significant speed-up compared to execution on CPU, the limited GPU memory capacity prevents processing arbitrarily high-dimensional datasets.

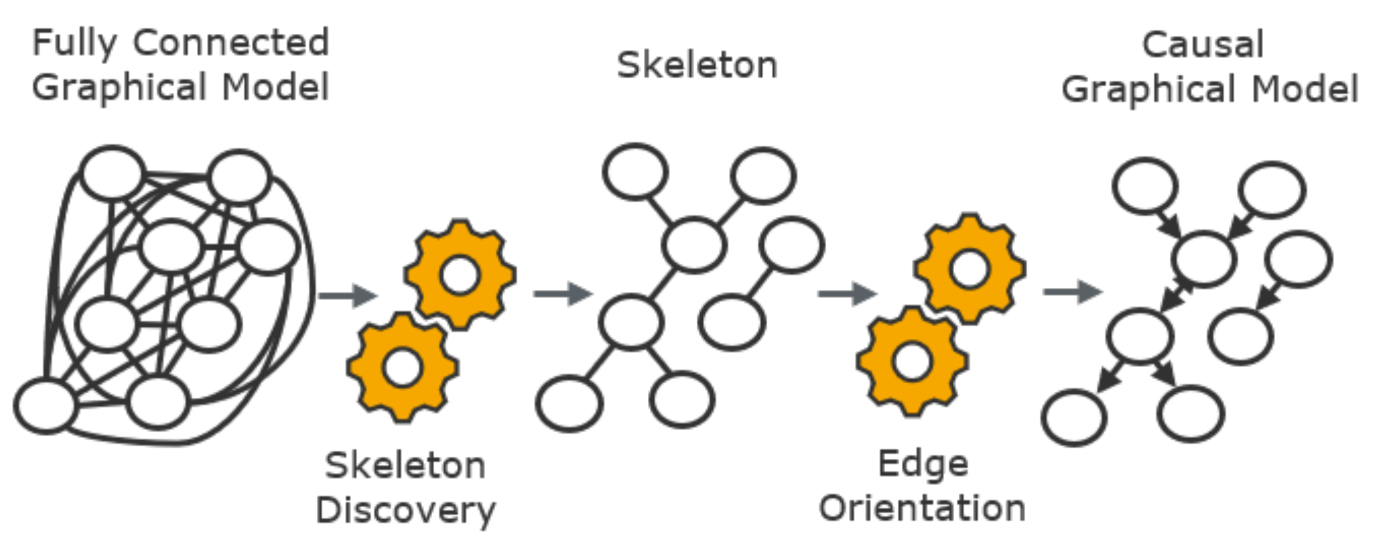

Figure: A schematic representation of the causal structure learning procedure

Causal Structure Learning Procedure

In the recent years, the notion of causality has grown from a nebulous concept into a mathematical theory[1]. A conceptual algorithm for learning the causal graphical model operates in two phases[2]:

- *Skeleton Discovery,* use conditional independence (CI) tests to receive information about the underlying relationships.
- *Edge Orientation,* determine the orientation of the detected relationships to construct a causal graphical model.

A Skeleton Discovery beyond GPU Memory Capacity

In previous work[3], a GPU-accelerated causal structure learning implementation has been proposed. While, it achieves significant speed-up, it is only applicable to datasets fitting into GPU memory. To overcome this limitation, we propose a block-wise skeleton discovery. The algorithm splits the input dataset into smaller blocks with size *bs* that fit into GPU memory. The impact of the block size on the execution time is shown in Table 1. Blocks should be large enough to allow for overlap of computation on CPU/GPU and data transfer.
The overhead of the block-wise approach compared to a non-block-wise implementation is marginal, compare Figure 1. Yet, one can see that the block-wise skeleton discovery allows to process datasets that exceed the GPU memory capacity.

bs	64	128	256	512	1024	2048	4096	8192
time in *s*	6042.6	5072.6	4909.5	4935.7	4862.1	4826.3	4788.2	5561.9

Table 1: Execution times with varying block sizes *bs* on a dataset with 20,000 variables.

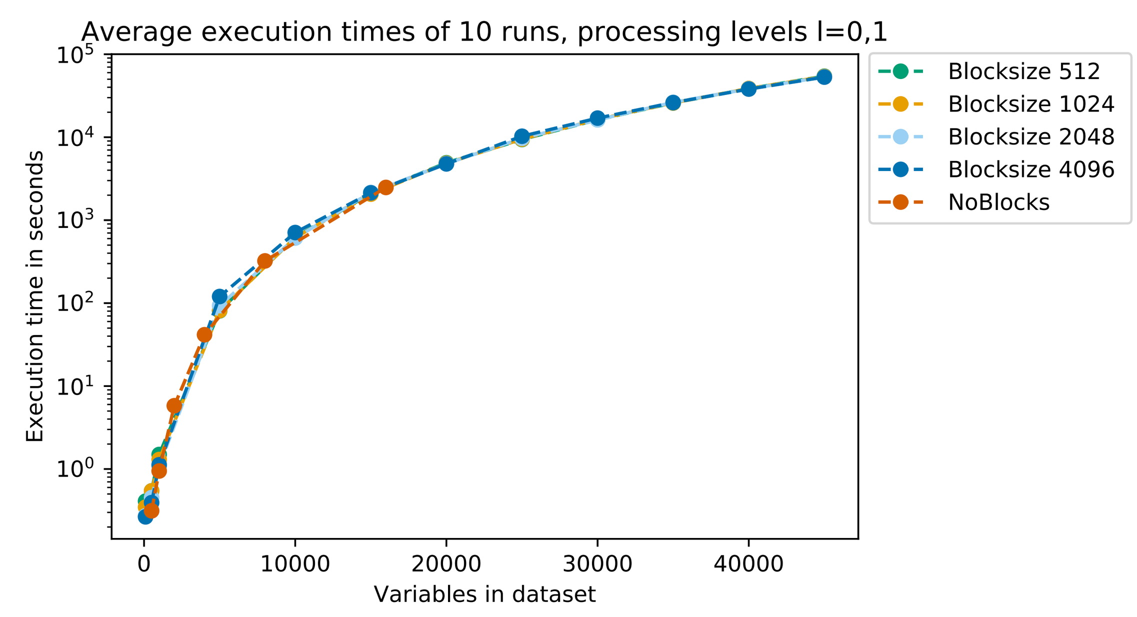

Figure 1: Execution times with increasing variables in dataset, assuming all CI tests are conducted.

References:

1) J. Pearl. Causality: Models, Reasoning and Inference. Cambridge University Press, New York, NY, USA, 2nd edition, 2009.
2) P. Spirtes, C. N. Glymour, and R. Scheines. Causation, prediction, and search. MIT press, 2000.
3) C. Schmidt, J. Huegle, and M. Uflacker. Order-independent constraint-based causal structure learning for gaussian distribution models using GPUs. In Proceedings of the 30th International Conference on Scientific and Statistical Database Management (SSDBM '18). ACM, 2018, New York, NY, USA, Article 19, 10 pages.

Projektbeteiligte

Christopher Schmidt, Johannes Huegle and Matthias Uflacker
Hasso-Plattner-Institute
Enterprise Platform and Integration Concepts
August-Bebel-Str. 88
14482 Potsdam, Germany
{christopher.schmidt, johannes.huegle, matthias.uflacker}@hpi.de

Hasso Plattner Institut
Digital Engineering · Universität Potsdam