Chair of Scientific Computing in Computer Science
TUM School of Computation, Information and Technology
Technical University of Munich

Data-Driven Solver Selection for Sparse Linear Matrices at Scale

Hayden Liu Weng[§], Hans Joachim Bungartz[§], and Felix Dietrich[§]

[§] {h.liu, bungartz, f.dietrich}@tum.de, TUM School of Computation, Information and Technology, Technical University of Munich

Introduction

- (Sparse) Linear systems are everywhere
- Choosing appropriate solvers and preconditioners is very challenging, especially for a novice
- Non-optimal choices might not converge

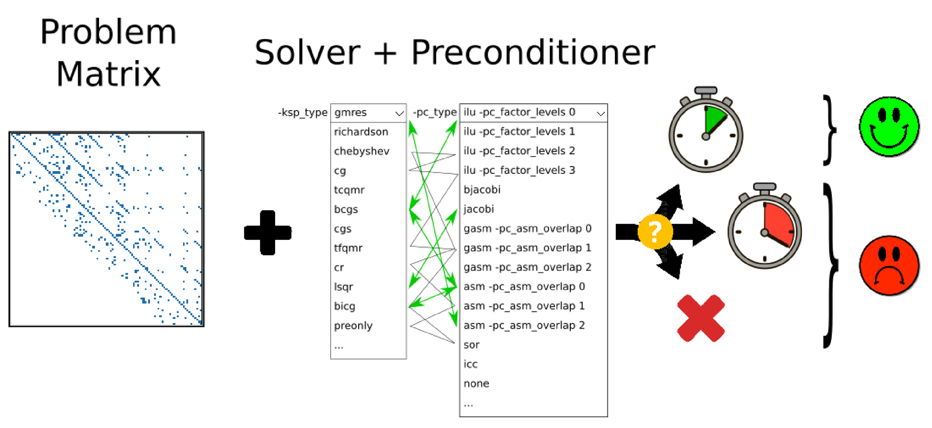

Figure 1 The choice of solver-preconditioner pair can yield quick, slow, or no convergence.

- SANS systems [1] attempt to use data instead of expert knowledge to automate the selection process
- Classical machine learning has been used for solver selection in, e.g., SALSA [2] and Lighthouse [3]
- Other recent approaches use neural networks [4]

Main Ideas

- Feature selection limits model complexity and reduces computation time
- Regression (wrt. runtime) maintains information on relative performance (vs. classification)
- Embedding lessens the impact of having unbalanced classes and limited data
- Misclassification error doesn't convey the impact of a wrong choice (slower convergence vs. no convergence!)
  - The Absolute Relative Error (ARE) better conveys how costly a wrong prediction is:

$$\text{ARE} = \frac{\|t_{\text{pred}} - t_{\text{best}}\|}{t_{\text{best}}}$$

Takeaway Messages

- The established method is not necessarily guaranteed to be the best
- There is no single-best solution
- It is worthwhile to look under the hood (beyond a black-box optimizer)
- This approach can be generalized and extended to further problems!

Methods and Results

This work adapts and extends the approach from Yeom [8], following a 3-step process:

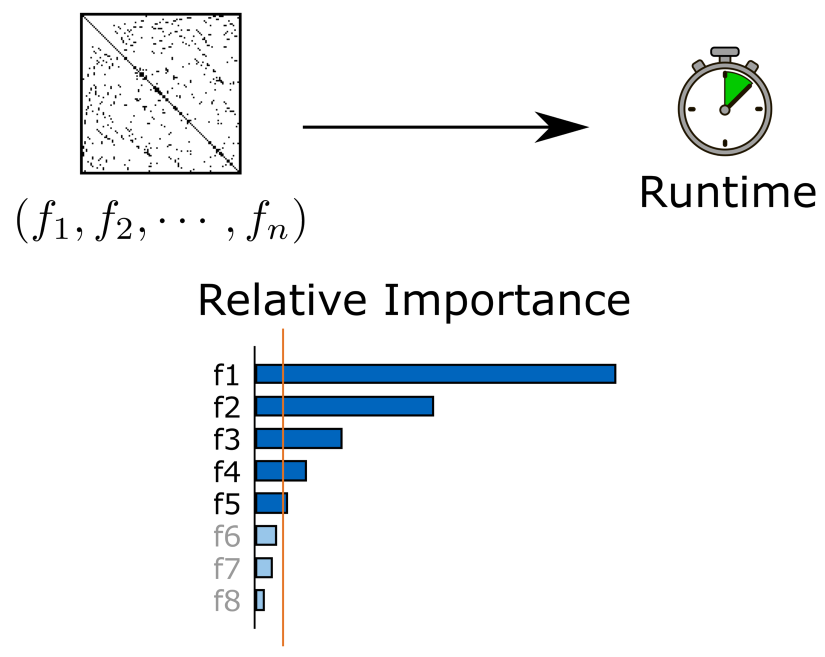

(a) Relevant feature selection via regression analysis

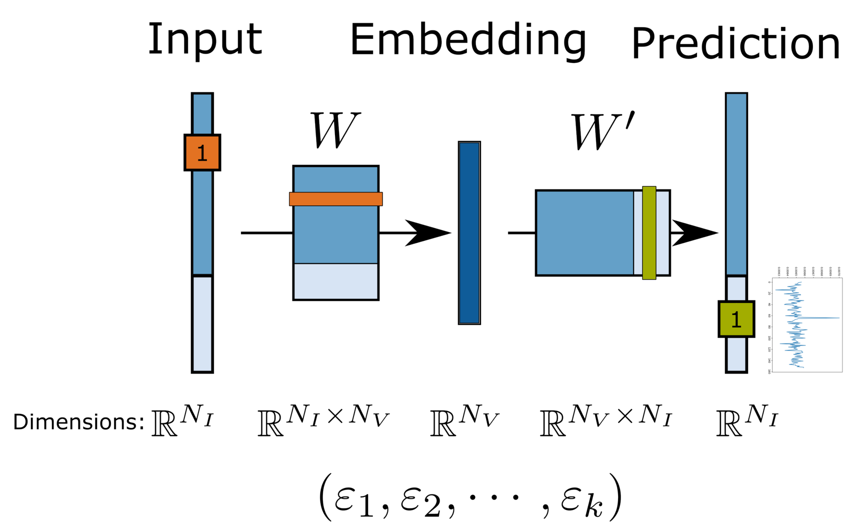

(b) Performance Vector Space construction via Word2vec

$$\min_{\alpha} \|\mathbf{f}' - \mathbf{F}\alpha\|_2^2$$
$$\text{s.t.} \quad \|\alpha\|_1 \leq \delta, \quad \alpha_i \geq 0$$
$$\mathbf{F}\alpha \rightarrow \mathbf{E}\alpha$$

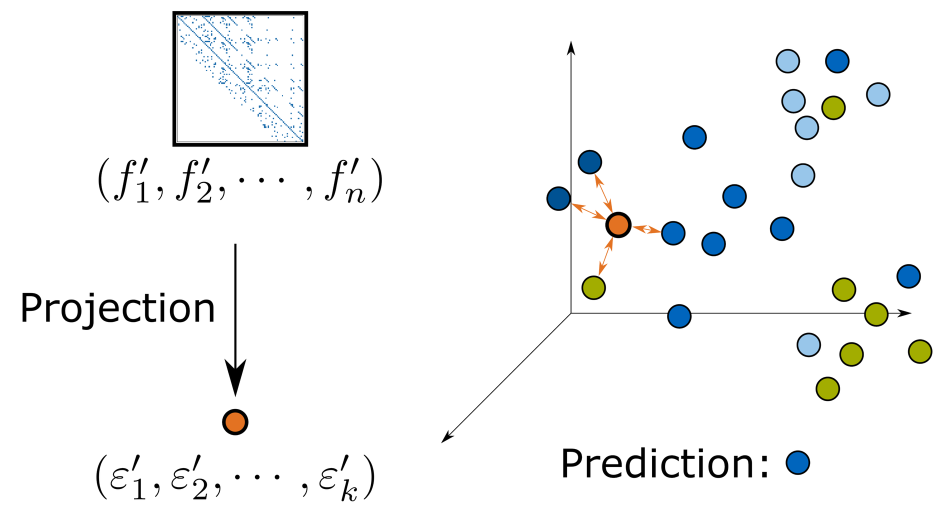

(c) Projection of new samples into the PVS and Solver Selection via k-Nearest Neighbors

Figure 2 Graphical representation of the individual steps.

Applying the method to matrix property and performance data on 775 square matrices with real values from the *SuiteSparse Matrix Collection* from [6] yielded the following results:

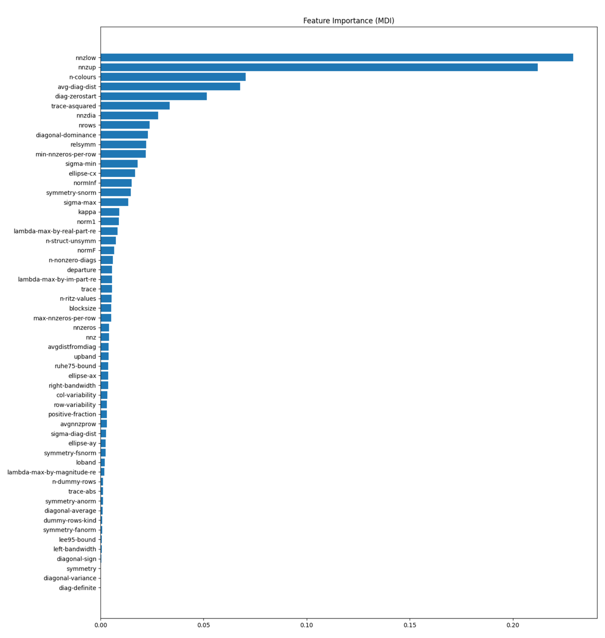

(a) Considered Features by importance

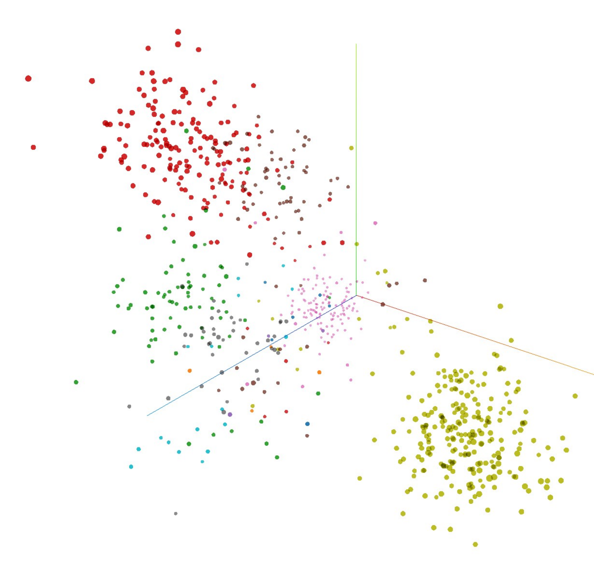

(b) PVS colored by best solver

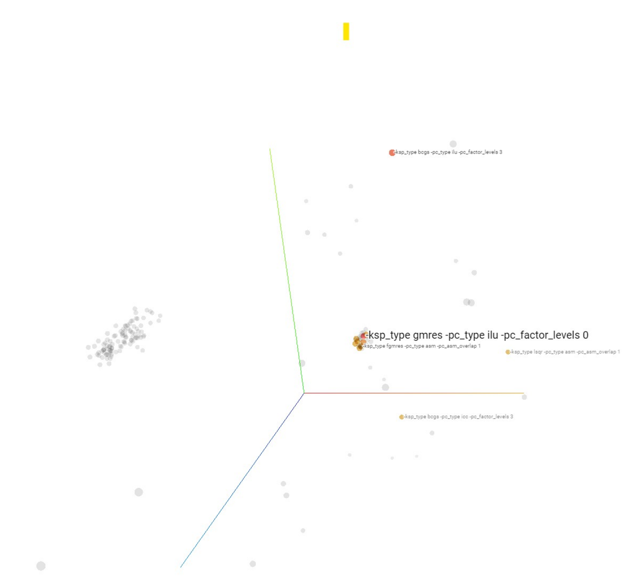

(c) Solver Embedding

Figure 3 Results for feature selection and Embedding of the SuiteSparse matrices

Outlook

- Beyond matrix properties, available software and hardware strongly affects performance (cf. [5])
- Behavior at scale differs from single-core performance and needs to be analyzed further (cf. [6])
- Modern systems increasingly allow using GPUs for computation, adding yet another layer to the problem (cf. [7])
- The current approach serves as proof-of-concept and can be adapted to other selection/tuning problems.

References

[1] J. Dongarra and V. Eijkhout, "Self-adapting numerical software for next generation applications," *The International Journal of High Performance Computing Applications*, vol. 17, no. 2, pp. 125–131, 2003.

[2] S. Bhowmick, V. Eijkhout, Y. Freund, E. Fuentes, and D. Keyes, "Application of machine learning to the selection of sparse linear solvers," *Int. J. High Perf. Comput. Appl*, 2006.

[3] E. Jessup, P. Motter, B. Norris, and K. Sood, "Performance-based numerical solver selection in the lighthouse framework," *SIAM Journal on Scientific Computing*, vol. 38, no. 5, pp. S750–S771, 2016.

[4] Y. Funk, M. Götz, and H. Anzt, "Prediction of optimal solvers for sparse linear systems using deep learning," in *Proceedings of the 2022 SIAM Conference on Parallel Processing for Scientific Computing*, pp. 14–24, SIAM, 2022.

[5] P. A. Motter, *Hardware awareness for the selection of optimal iterative linear solvers.* PhD thesis, University of Colorado at Boulder, 2017.

[6] K. Sood, *Iterative Solver Selection Techniques for Sparse Linear Systems.* PhD thesis, University of Oregon, 2019.

[7] H. Anzt, M. Gates, J. Dongarra, M. Kreutzer, G. Wellein, and M. Köhler, "Preconditioned krylov solvers on gpus," *Parallel Computing*, vol. 68, pp. 32–44, 2017.

[8] J.-S. Yeom, J. J. Thiagarajan, A. Bhatele, G. Bronevetsky, and T. Kolev, "Data-driven performance modeling of linear solvers for sparse matrices," in *2016 7th International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*, pp. 32–42, IEEE, 2016.