

Building Scalable Tensor Regression Models: Linear Solvers and Beyond

Hengrui Luo, Anna Ma, Akira Horiguchi, Li Ma

Abstract

The regression problem $y = f(\mathbf{X}) + \epsilon$, where $\mathbf{X} \in \mathbb{R}^{n \times d_1 \times d_2}$, $y \in \mathbb{R}^1$ is a tensor input variable, presents unique challenges at the intersection of numerical linear algebra and statistical modeling. This extended abstract synthesizes recent advancements in solving both linear and non-linear variants of this problem, focusing on scalable methods that leverage tensor structures.

In the linear case, we consider the tensor regression model $f(\mathbf{X}) = \mathbf{B} \circ \mathbf{X}$, where \circ denotes a tensor product. In [1], we focus on a generalization of matrix multiplication to tensors. This formulation leads to large-scale tensor linear systems that are computationally challenging to solve using traditional methods. Our primary contribution in this domain is the development of frontal slice approaches for iteratively solving these systems.

Our innovation in [1] lies in adapting classical iterative methods from numerical linear algebra to the tensor domain. By focusing on frontal slices of the coefficient tensor \mathbf{B} , we develop a family of iterative algorithms that significantly reduce the computational burden compared to direct solvers. Our frontal slice methods come in cyclic, block, and randomized variants, each offering different trade-offs between convergence speed and computational efficiency. These approaches draw inspiration from classical numerical linear algebra techniques such as Gauss-Seidel, block iterative methods, and randomized algorithms.

Another aspect of our work is the rigorous convergence analysis provided for these methods. We establish conditions for convergence and derive bounds on the convergence rate in terms of tensor properties. This analysis not only validates the proposed methods but also offers insights into algorithm behavior, guiding practitioners in method selection and parameter tuning.

The computational advantages of our approach are particularly striking for large-scale problems. We achieve a computational complexity of $O(d_2 \cdot \max(d_1, n)^3)$ for n samples and tensor dimensions d_1 and d_2 , comparing favorably to the $O(d_2^2 \cdot \max(d_1, n)^3)$ complexity of naive iterative solver via gradient descent, when d_2 is large.

Extending beyond linear models, $f(\mathbf{X})$ are assumed to be nonlinear. We address the challenge of non-linear tensor regression through the development of *tensor-input tree* (TT) models in [2]. Our model innovation generalize decision trees to handle tensor inputs, offering a flexible, non-parametric approach to modeling complex relationships in tensor data. The TT framework can be viewed as approximating the non-linear function $f(\mathbf{X})$ with a piecewise linear tensor function, where each piece corresponds to a leaf in the decision tree.

The development of TT models required innovative solutions to several challenges at the interface of numerical linear algebra and statistical learning. Our key contribution in [2] is the design of splitting criteria that effectively capture the multi-dimensional structure of tensor inputs. We propose criteria based on both variance reduction and low-rank approximation errors, leveraging tensor algebraic concepts to inform the tree-building process. To address the computational challenges of finding optimal splits in high-dimensional tensor spaces, we introduce two randomized numerical linear algebra techniques: leverage score sampling and branch-and-bound optimization. The complexity for generating the tree structure (via exhaustive search) is $O(n^2 \cdot d_1^2 \cdot d_2^2 \cdot \log k)$, where k is the number of nodes in the tree, in contrast to $O(n^3)$ in a tensor Gaussian process [3]. This approach aligns with theoretical results on tensor decompositions and proves highly effective in practice, especially for datasets exhibiting intrinsic low-rank structure.

The theoretical analysis provided for both our linear solvers and nonlinear TT models offers valuable insights into their behavior and limitations. For the linear case, we provide detailed convergence analysis, including results for both consistent and inconsistent systems. In the context of TT models, we establish consistency guarantees for coefficient estimates and derive oracle bounds for prediction errors.

Empirical evaluations on diverse datasets demonstrate the effectiveness of our approaches. In the linear case, we show the efficiency of frontal slice methods in applications like image deblurring. For non-linear regression, we compare TT models against state-of-the-art approaches like tensor Gaussian Processes, highlighting scenarios where TT offers superior performance or significant computational advantages.

From a numerical linear algebra perspective, our work opens several intriguing avenues for future research. The frontal slice methods suggest possibilities for developing tensor analogues of classical iterative methods. The efficient splitting criteria used in TT models could inspire new preconditioning techniques for tensor computations. Furthermore, the integration of randomized algorithms in both our linear solvers and regression models points to a broader trend of leveraging stochastic methods to tackle high-dimensional problems.

In conclusion, our research demonstrates the power of combining insights from numerical linear algebra with advanced statistical modeling techniques to address the challenges of tensor regression. By developing scalable methods for both linear and non-linear tensor regression (including tensor-on-tensor and tensor-on-scalar), we contribute to a comprehensive toolkit for tensor-based data analysis, laying the groundwork for future advancements in high-dimensional data analysis and scientific computing.

References

- [1] Hengrui Luo, and Anna Ma. “Frontal Slice Approaches for Tensor Linear Systems.” arXiv preprint arXiv:2408.13547 (2024).
- [2] Hengrui Luo, Akira Horiguchi, and Li Ma. “Efficient Decision Trees for Tensor Regressions.” arXiv preprint arXiv:2408.01926 (2024).
- [3] Rose Yu, Guangyu Li, and Yan Liu. “Tensor regression meets Gaussian processes.” International Conference on Artificial Intelligence and Statistics. PMLR, 2018.