

Fast Similarities in Factorized Tensors

Michael E. HOULE¹

Hisashi KASHIMA²

Michael NETT^{1,2}

¹ National Institute of Informatics, Japan
² The University of Tokyo, Japan

SUMMARY

Low-rank factorizations of higher-order tensors have become an invaluable tool for researchers from many scientific disciplines. Tensor factorizations have been successfully applied for moderately sized multi-modal data sets involving a small number of modes. However, a significant hindrance to the full realization of the potential of tensor methods is a lack of scalability on the client side: even when low-rank representations are provided by an external agent possessing the necessary computational resources, client applications are quickly rendered infeasible by the space requirements for explicitly storing a (dense) low-rank representation of the input tensor.

We consider the problem of efficiently computing common similarity measures between entities expressed by fibers (vectors) or slices (matrices) within a given factorized tensor. We show that after appropriate preprocessing, inner products can be efficiently computed independently of the dimensions of the input tensor.

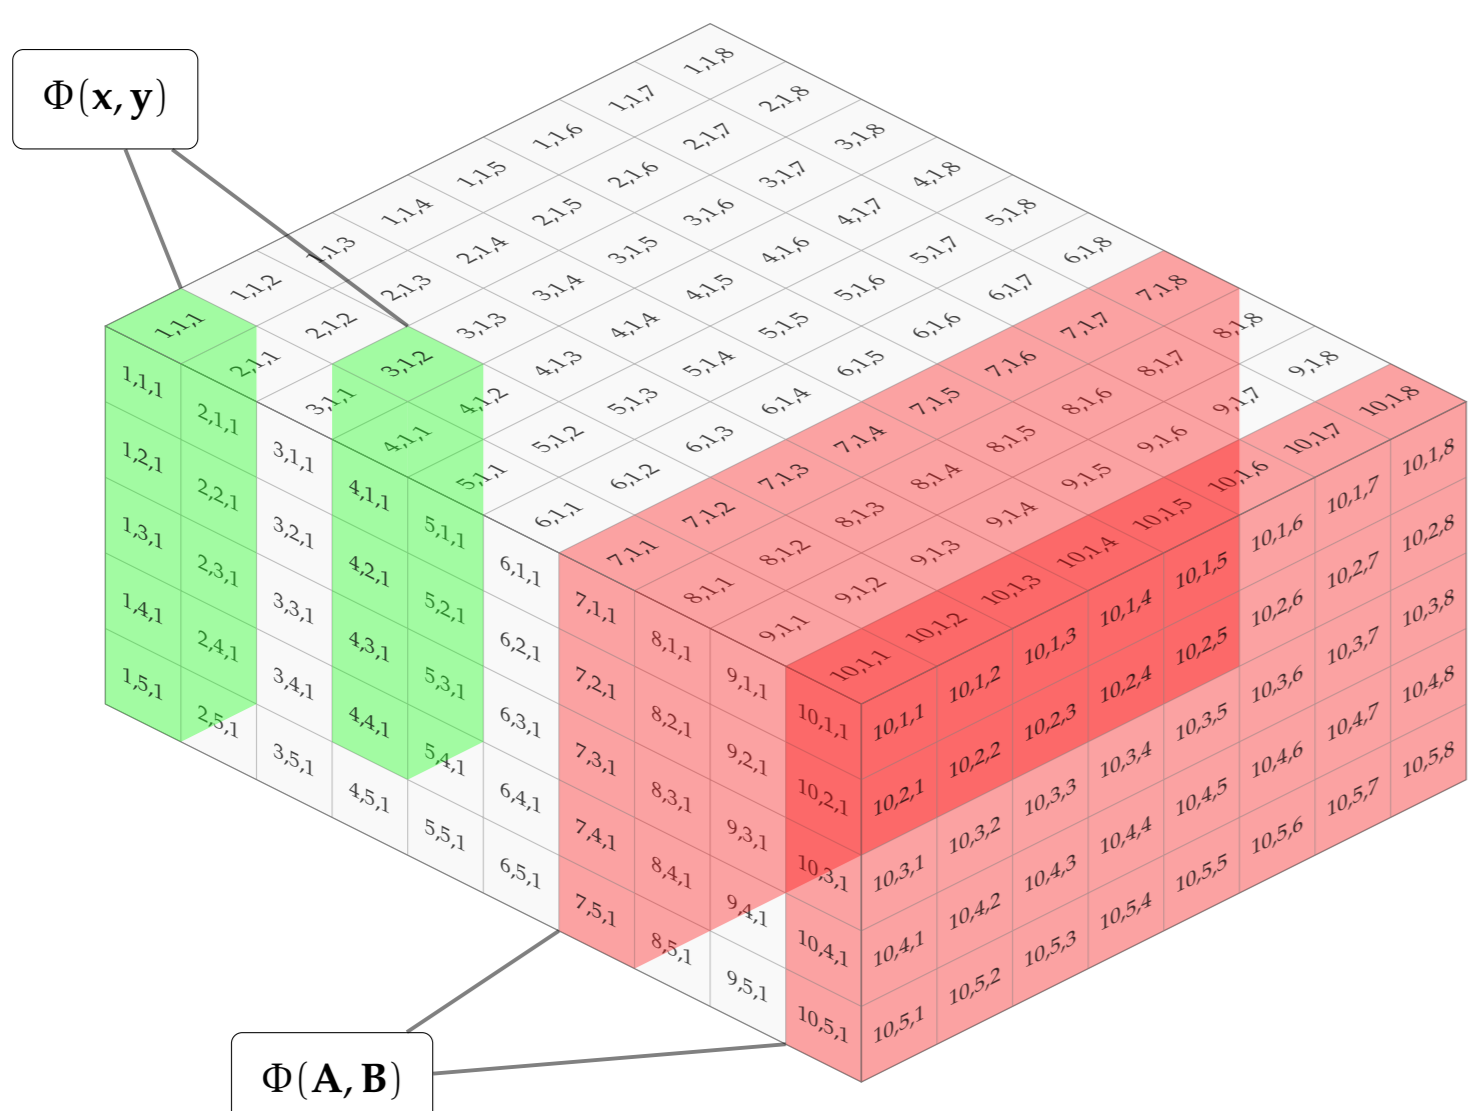
MOTIVATION

Within many emerging areas of computing, such as data mining, recommendation systems, security, and multimedia, applications of similarity search naturally arises in the context of such fundamental tasks as clustering, classification, matching and detection. Multi-modal data can be naturally represented in the form of a tensor (also known as a multiway array), a higher-dimensional extension of the matrix representation.

Tensor-based data modeling is particularly appealing whenever the data dynamics can be captured by truncated (low-rank) representations, in terms of a small number of latent variables. However, due to the complexity inherent in managing multi-modal data, effective and scalable strategies for similarity search are crucial to the overall performance of such systems.

TENSORS AND SIMILARITIES

Tensors are a suitable means to represent multi-object relationships, such as ratings given by customers to restaurants at certain times.



User profiles can be represented as slices (matrices) and rating profiles as fibers (vectors). The need to efficiently search for similar rating- or user profiles arises in applications such as, for example, recommender systems.

FACTORIZATION MODELS

- The *CP model* describes a tensor \mathfrak{X} as a conical combination of rank-1 tensors $\mathfrak{X}^{[1]}, \dots, \mathfrak{X}^{[m]}$. The individual $\mathfrak{X}^{[i]}$ may be expressed in terms of vector products.

$$\mathfrak{X} = \begin{bmatrix} 4 & 6 & 11 \\ 2 & 2 & 5 \\ 3 & 3 & 4 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 2 \\ 1 & 0 & 2 \\ 0 & 1 & 0 \end{bmatrix} + \begin{bmatrix} 3 & 6 & 9 \\ 1 & 2 & 3 \\ 1 & 2 & 3 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 2 & 1 & 1 \end{bmatrix}$$

$$= \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \circ [1 \text{ } 0 \text{ } 2] + \begin{bmatrix} 3 \\ 1 \\ 1 \end{bmatrix} \circ [1 \text{ } 2 \text{ } 3] + \begin{bmatrix} 0 \\ 0 \\ 2 \end{bmatrix} \circ [2 \text{ } 1 \text{ } 1]$$

- The *Tucker model* is a form of higher-order SVD. It represents a tensor \mathfrak{X} as the ‘mode-wise’ product of a core tensor \mathfrak{C} and p factor matrices.

$$\mathfrak{C} = \begin{bmatrix} 1 & 1 \\ 1 & 0 \\ 1 & 0 \end{bmatrix}, \mathbf{U}^{[1]} = \begin{bmatrix} 1 & 2 \\ 0 & 1 \\ 2 & 1 \end{bmatrix}, \mathbf{U}^{[2]} = \begin{bmatrix} 1 & 2 \\ 0 & 2 \\ 0 & 1 \end{bmatrix}, \mathfrak{X} = \begin{bmatrix} 5 & 2 & 1 \\ 1 & 0 & 0 \\ 7 & 4 & 2 \end{bmatrix}$$

$$\diamond_{31} = c_{11}u_{31}^{[1]}u_{11}^{[2]} + c_{12}u_{31}^{[1]}u_{12}^{[2]} + c_{21}u_{32}^{[1]}u_{11}^{[2]} + c_{22}u_{32}^{[1]}u_{12}^{[2]}$$

- The *pairwise interaction model* expresses a tensor \mathfrak{X} by independent interaction between individual pairs of factors $\mathbf{U}^{[i,j]}$ and $\mathbf{U}^{[j,i]}$.

$$\mathfrak{X} = \begin{bmatrix} 1 & 3 & 0 \\ 2 & 7 & 2 \\ 1 & 3 & 0 \end{bmatrix}, \mathbf{U}^{[1,2]} = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 1 & 0 \end{bmatrix}, \mathbf{U}^{[2,1]} = \begin{bmatrix} 1 & 3 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

PREPROCESSING

Depending on the factorization model, we precompute the following values.

- Given a rank- m CP model, store inner products of the form $\Phi(\mathbf{u}_k^{[i]}, \mathbf{u}_k^{[j]})$, for each mode $1 \leq k \leq p$ with $1 \leq i \leq j \leq m$.
- Given a rank- (m_1, \dots, m_p) Tucker model, store $\Phi(\mathbf{u}_i^{[k]}, \mathbf{u}_j^{[k]})$ for each mode $k \in \{1, \dots, p\}$ and $1 \leq i \leq j \leq m_k$.
- Given a PITF model, store inner products of the form $\Phi(\mathbf{u}_a^{[i,j]})$ and $\Phi(\mathbf{u}_b^{[i,j]}, \mathbf{u}_b^{[i,k]})$ for any choice of $a, b \in \{1, \dots, m\}$ and $i, j, k \in \{1, \dots, p\}$.

FAST SIMILARITY COMPUTATION

When computing similarities between substructures (such as fibers or slices) we can avoid the full model size (n_1, \dots, n_p) by using the previously computed inner product values. For example, given a rank- m CP model of a tensor, we can compute the similarity of two fibers using:

$$\Phi_{CP}(\hat{\mathfrak{X}}_{\square i_2 \dots i_p}, \hat{\mathfrak{X}}_{\square j_2 \dots j_p})$$

$$= \sum_{k=1}^{n_1} \hat{\mathfrak{X}}_{ki_2 \dots i_p} \hat{\mathfrak{X}}_{kj_2 \dots j_p}$$

$$= \sum_{k=1}^{n_1} \left(\sum_{a=1}^m w_a u_{1k}^{[a]} u_{2i_2}^{[a]} \dots u_{pi_p}^{[a]} \sum_{b=1}^m w_b u_{1k}^{[b]} u_{2j_2}^{[b]} \dots u_{pj_p}^{[b]} \right)$$

$$= \sum_{a,b=1}^m w_a w_b u_{2i_2}^{[a]} \dots u_{pi_p}^{[a]} u_{2j_2}^{[b]} \dots u_{pj_p}^{[b]} \sum_{k=1}^{n_1} u_{1k}^{[a]} u_{1k}^{[b]}$$

$$= \sum_{a,b=1}^m w_a w_b u_{2i_2}^{[a]} \dots u_{pi_p}^{[a]} u_{2j_2}^{[b]} \dots u_{pj_p}^{[b]} \Phi(\mathbf{u}_1^{[a]}, \mathbf{u}_1^{[b]})$$

The process is similar for higher-order substructures (slices, ...).

COMPLEXITIES

Complexities of the different factorization models, as compared to operations on the original unfactorized tensor.

Complexity	Original	Tucker	CP	PITF
Size	$\prod_i n_i$	$\prod_i m_i + \sum_i m_i n_i$	$m \sum_i n_i$	$mp \sum_i n_i$
Access	1	$p \prod_i m_i$	mp	mp^2
Prepr.	1	$\sum_i m_i^2$	$m^2 p$	$m^2 p^2$
Naïve Sim.	n_k	$n_k p \prod_i m_i$	$n_k m p$	$n_k m p^2$
Fast Sim.	—	$p m_k \prod_i m_i$	$m^2 p$	$m^2 p^4$

PERFORMANCE




Figure: Speed-ups achieved for a rank- m CP factorization of an order- p tensor.

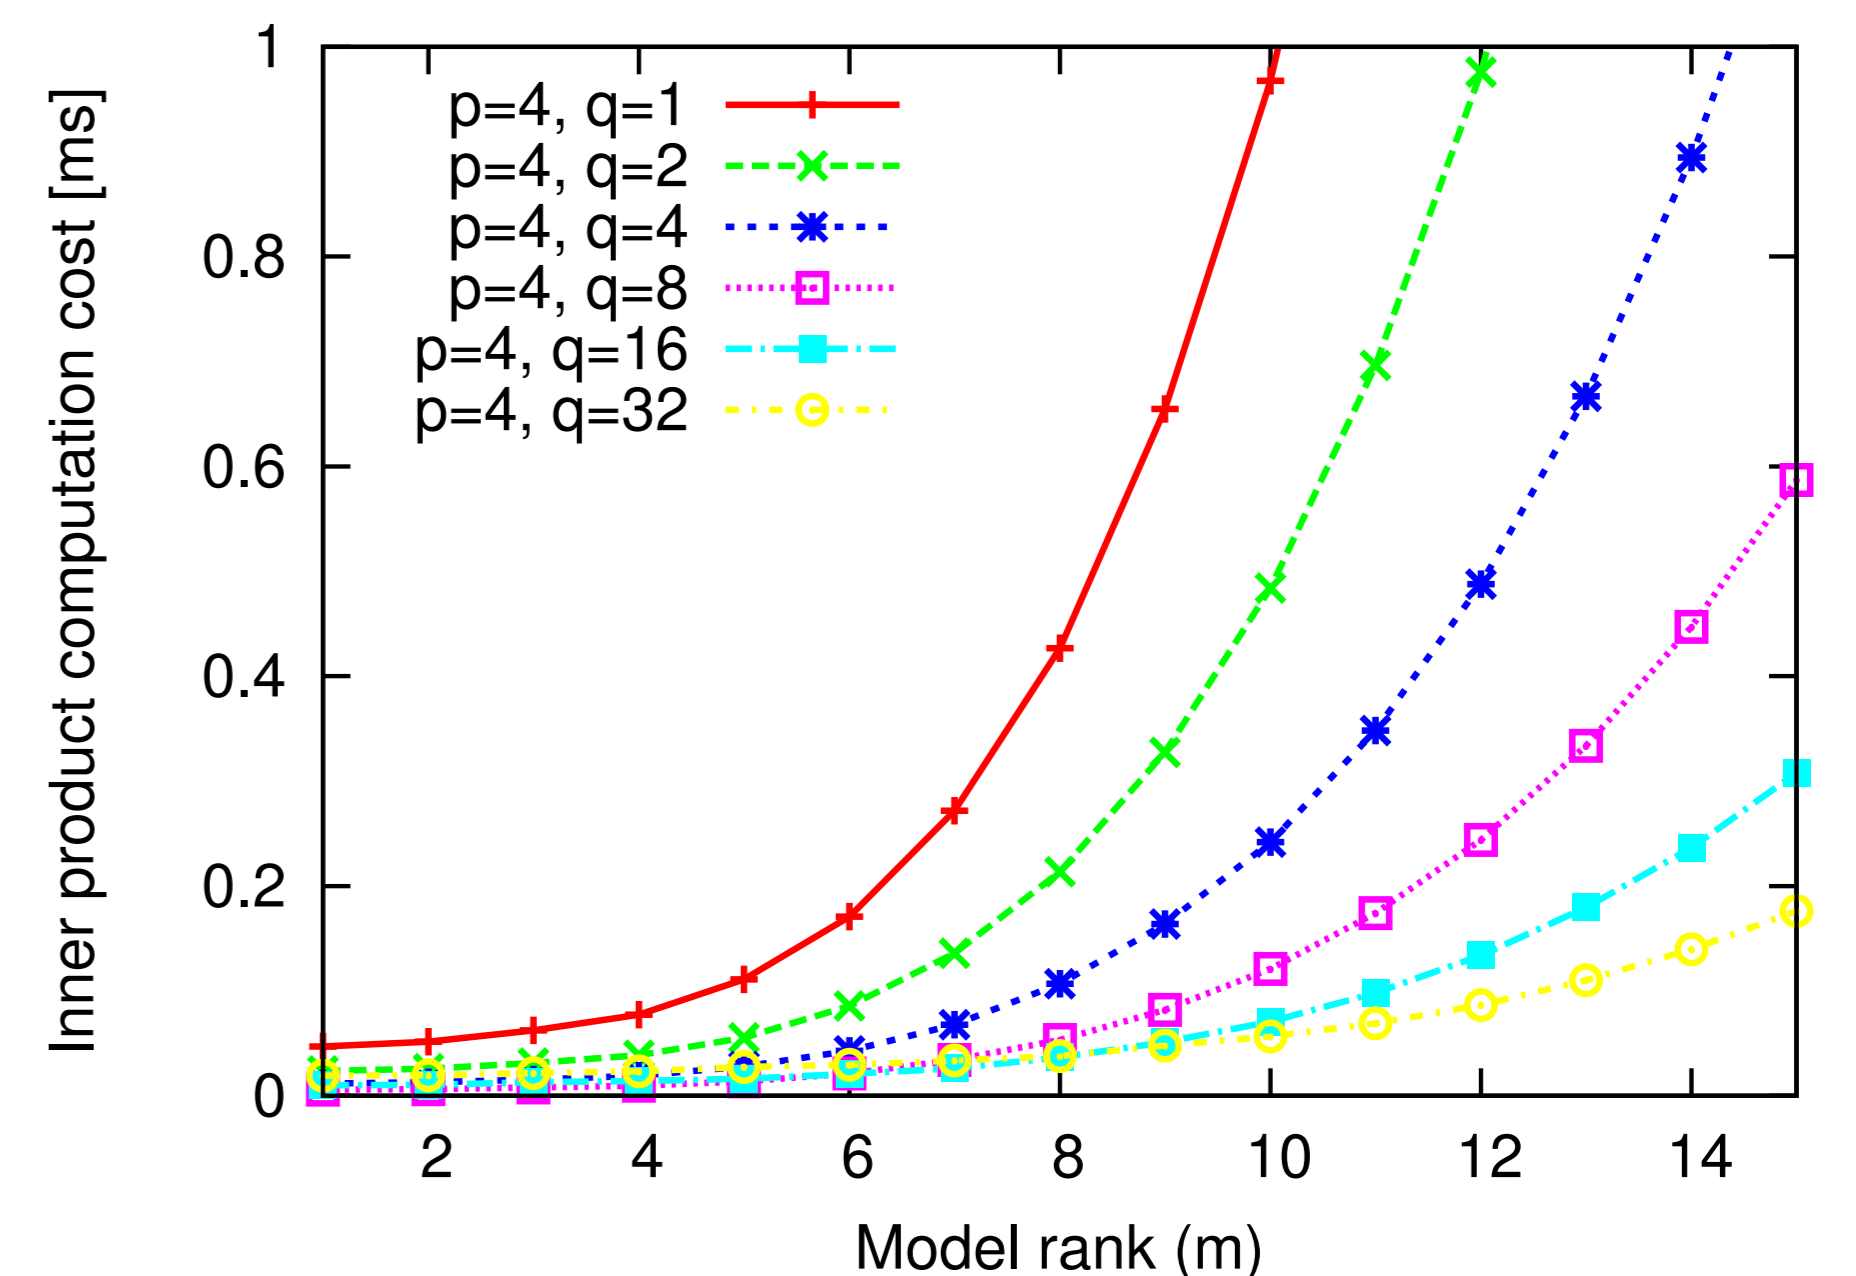


Figure: Speed-ups achieved for a rank- $(m \times \dots \times m)$ Tucker factorization of an order- p tensor. The computation is distributed across q processors.

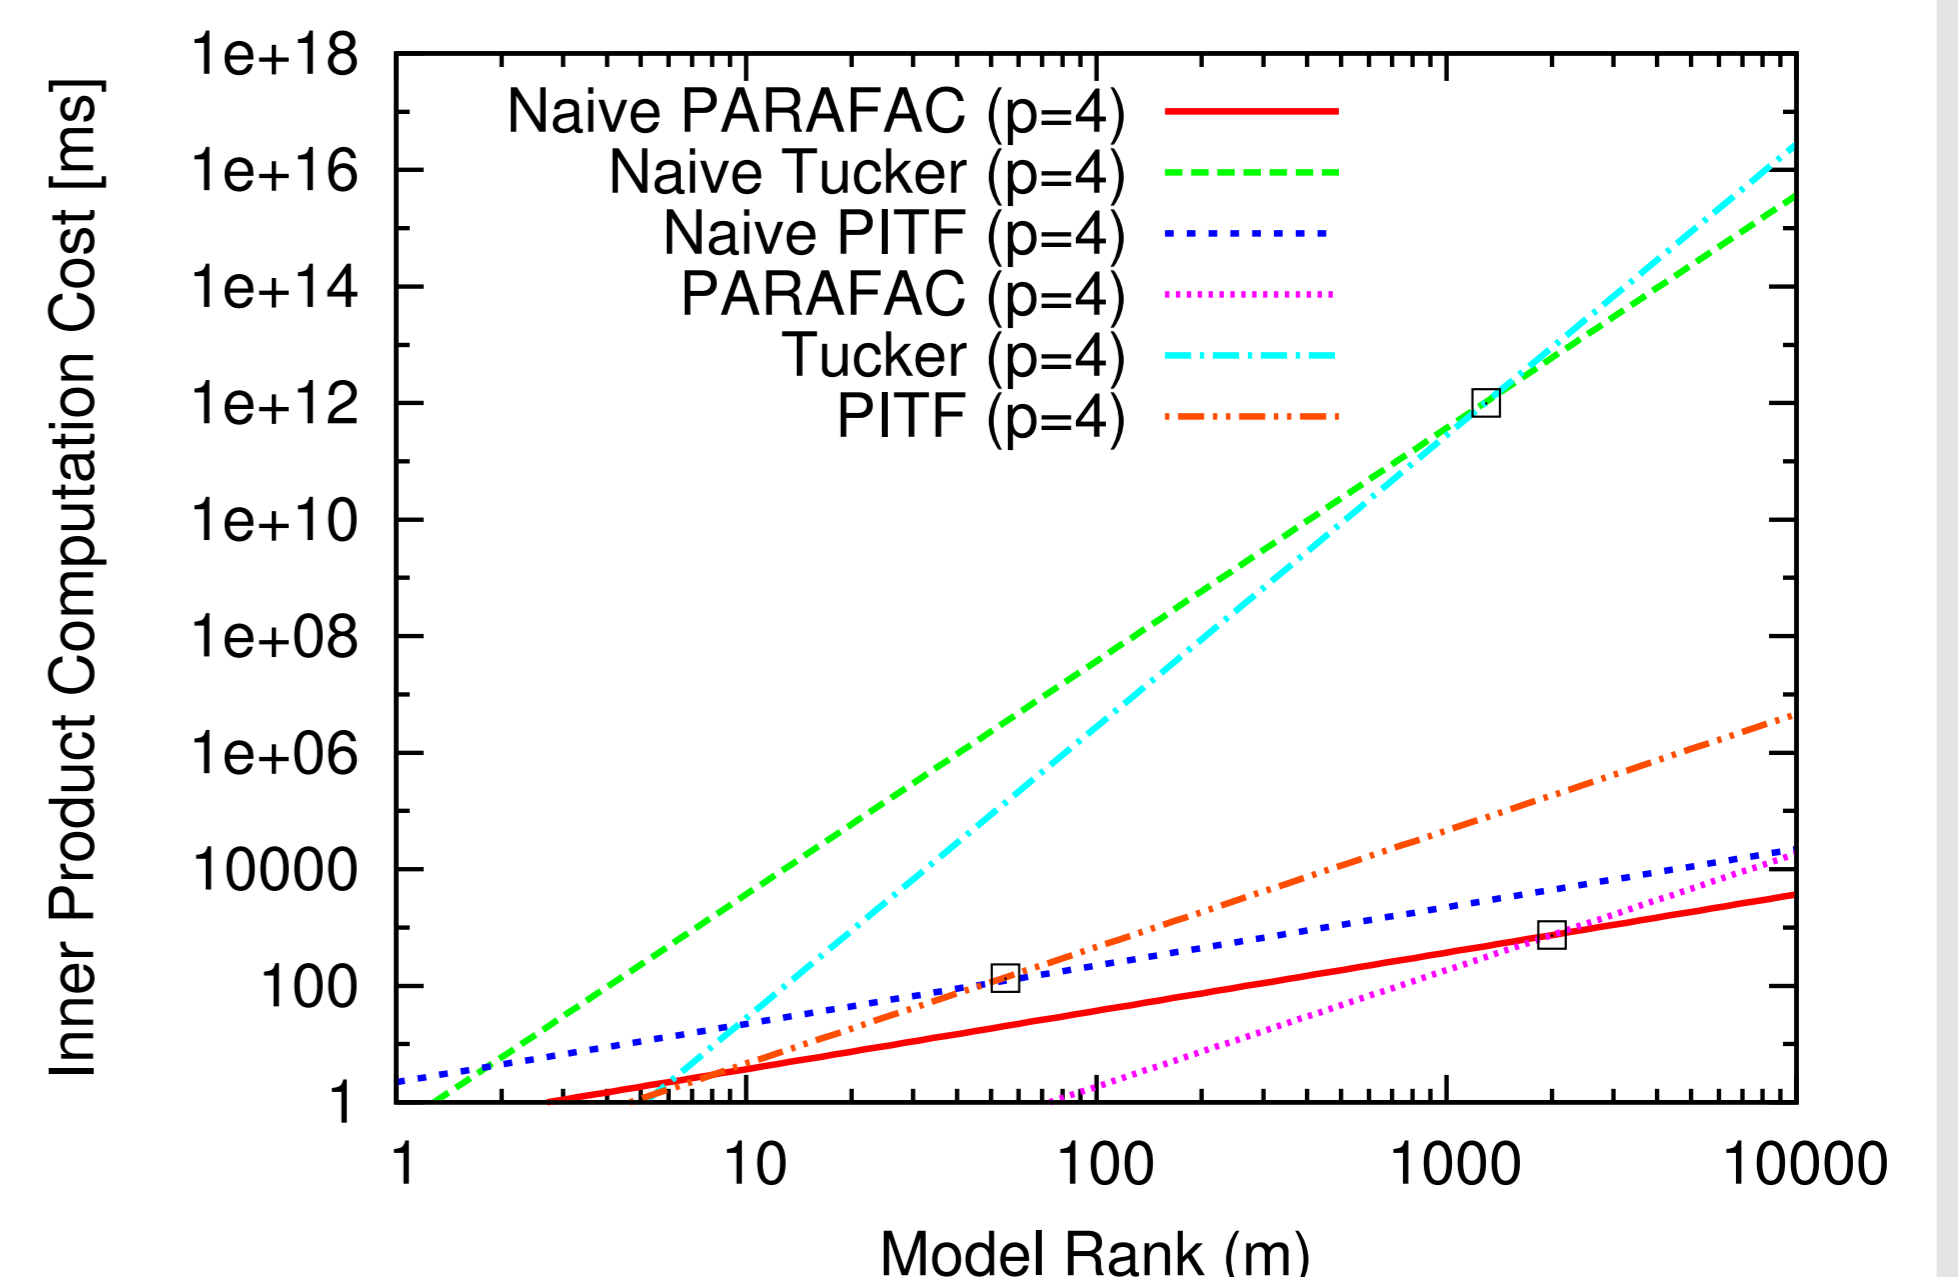


Figure: Break-even points of naïve and fast similarity computations.

ADDITIONAL MATERIAL

- Technical Report
- Poster
- Implementation
- Documentation

