

FRAPPE: Fast Rank Approximation with Explainable Features for Tensors

William Shiao^[0000-0001-5813-2266] and Evangelos E. Papalexakis^[0000-0002-3411-8483]

University of California Riverside
wshia002@ucr.edu, epapalex@cs.ucr.edu

Abstract Tensor decompositions have proven to be effective in analyzing the structure of multidimensional data. However, most of these methods require a key parameter: the number of desired components. In the case of the CANDECOMP/PARAFAC decomposition (CPD), the ideal value for the number of components is known as the canonical rank and greatly affects the quality of the decomposition results. Existing methods use heuristics or Bayesian methods to estimate this value by repeatedly calculating the CPD, making them extremely computationally expensive. In this work, we propose FRAPPE, the first method to estimate the canonical rank of a tensor without having to compute the CPD. This method is the result of two key ideas. First, it is much cheaper to generate synthetic data with known rank compared to computing the CPD. Second, we can greatly improve the generalization ability and speed of our model by generating synthetic data that matches a given input tensor in terms of size and sparsity. We can then train a specialized single-use regression model on a synthetic set of tensors engineered to match a given input tensor and use that to estimate the canonical rank of the tensor—all without computing the expensive CPD. FRAPPE is over 24 \times faster than the best-performing baseline, and exhibits a 10% improvement in MAPE on a synthetic dataset. It also performs as well as or better than the baselines on real-world datasets.

Keywords: Rank estimation · Tensor decomposition · Self-supervised · CANDECOMP/PARAFAC

1 Introduction

Tensors are widely used in the field of computer science and can be used to model a variety of datasets. One popular family of methods that can be applied to tensors is tensor decompositions. The performance of these methods typically depends on a key parameter—the number of components in the decomposition, also known as the rank of the decomposition. If our decomposition rank is too low, we may fail to capture key information, but if it is too high, we may capture too much noise. In the case of the popular CANDECOMP/PARAFAC decomposition, (CPD) [9], the ideal value for the rank of the decomposition is known as the canonical rank. However, finding this is a challenging task. A similar problem, finding the *exact rank* of a tensor (the minimum number of CPD components required to represent the tensor exactly), is known to be NP-hard over the set of rational numbers [11, 12].

As such, various heuristic-based approaches have been proposed to solve this task. One such method, proposed by Bro and Kiers [5], uses the core consistency diagnostic (CORCONDIA) to evaluate the “appropriateness” of a PARAFAC model. AutoTen [20] uses CORCONDIA to select a suitable number of components automatically. Normalized SVD (NSVD) [24] also attempts to provide a heuristic for the rank of a tensor, but it is not fully automatic and currently requires human supervision. Fernandes et al. [8] propose NORMO, which computes the CPD of a tensor at various ranks, looks at the mean correlation between each of its components, and selects the maximum number of components with no redundancy between the components.

Another set of approaches is the Bayesian approaches to this problem. Zhao et al. [27] attempt to tackle this problem by using Bayesian inference to automatically determine the decomposition’s rank. Mørup and Hansen [17] use a Bayesian framework with Automatic Relevance Determination (ARD) to estimate the rank of a Tucker decomposition. This can easily be adapted since the CPD can be formulated as a restricted Tucker decomposition [25]. Zhou et al. [28] compute the CPD and use a convolutional neural network (CNN) on the factors in an effort to determine the canonical rank of the tensor. However, they primarily focus on and evaluate their method within the image domain.

One common theme across all of these methods is that they *all require exhaustively computing the CPD multiple times across all the candidate ranks*. This is computationally expensive and impractical for large tensors, so we would like to avoid this step. However, doing so poses several challenges: (a) it is difficult to directly train a supervised machine learning model due to the scarcity of tensors with known canonical rank (which can typically only be determined via manual human inspection), and (b) it is difficult to develop a model capable of handling tensors from various domains and of various sizes.

To solve (a), we propose using synthetic tensors, which we can create with known canonical rank. To solve (b), we propose using hand-picked size-agnostic features extracted from each tensor. However, we find that supervised regression models trained using that data (see Sec. 4.2) fail to generalize well to real-world tensors. To overcome this limitation, we propose FRAPPE: a fully self-supervised model that uses synthetic data specifically tailored for a given input tensor. FRAPPE outperforms existing baselines on synthetic data and is on par with the best baselines on real-world data. It is also $\sim 24\times$ faster than the best-performing baseline (NORMO).

We evaluate the performance of both models on three classes of tensors: synthetic tensors, real tensors that have been extensively studied in existing literature to determine their rank, and the weight tensor of a convolutional neural network (to find the ideal CPD rank for compressing the weights while maintaining good classification accuracy). Finally, we also use these models to analyze which features are the best and worst predictors of a tensor’s rank.

Our contributions include the following:

- *Novel method:* We propose FRAPPE—the first method to estimate the canonical rank of generic tensor *without having to compute the CPD*.

	Proposed	Baseline Methods			
		NORMO [8]	ARD-Tucker [7]	AutoTen [20]	TRN-PD [28]
Avoids Computing CPD	✓	✗	✗	✗	✗
Avoids Human Supervision	✓	✓	✓	✓	✓
Avoids Strictly Labelled Dataset	✓	✓	✓	✓	✗
Explainable	✓	○ ¹	✗	✗	✗
Works for General Tensors	✓	✓	✓	✓	✗ ²
Eval. Speed	24×	1×	12×	18×	†

Table 1. Comparison of different tensor rank estimation methods. All speeds are relative to the evaluation time of the slowest baseline method. Exact evaluation times are shown in Table 3. †: We did not benchmark against TRN-PD because it is designed for image data and image denoising, which is not the focus of this work.

- *Speed*: We show that FRAPPE offers a $\sim 24\times$ improvement in evaluation time over the best-performing baseline method and is $\sim 1.6\times$ faster than the fastest baseline.
- *Extensive Evaluation*: We evaluate our methods and baseline methods on synthetically generated tensors, real tensors with known rank, and the weight tensor of a convolutional neural network. We show that FRAPPE outperforms the best baseline by about 10% MAPE on a synthetic dataset and performs as well as or better than baselines on real-world datasets.
- *Explainability*: We examine the importance of each of our hand-picked features to see which features contribute the most to classification results (see Fig. 4).
- *Reproducible*: We provide the Python code for our method online at <https://github.com/willshiao/frappe>. We also provide the code to reproduce our baseline results.

2 Background

Here we provide an overview of the notation and background for this work. We denote a set, tensor, matrix, vector, and scalar as \mathcal{X} , \mathbf{X} , \mathbf{x} , x , respectively. Let $\mathcal{T}_{(i,j,k)}$ denote the value in the i -th, j -th, and k -th entry along the 1st, 2nd, and 3rd modes, respectively, of the 3rd order tensor \mathcal{T} . Let $(:)$ refer to all possible indexes in the given mode when used in tensor indexing (e.g., $\mathcal{T} = \mathcal{T}_{(:, :, :)}$ for a third-order tensor \mathcal{T}).

We write the Frobenius norm of \mathcal{T} as $\|\mathcal{T}\|_F$, which is defined as follows for a third-order tensor of size $I \times J \times K$:

$$\|\mathcal{T}\|_F = \sqrt{\sum_{i=1}^I \sum_{j=1}^J \sum_{k=1}^K \mathcal{T}_{i,j,k}^2}. \quad (1)$$

We also define several functions. Let $\text{ZEROS}(\mathcal{T})$ return the number of zero values in a tensor \mathcal{T} , and $\text{RAND}(i, j)$ return a random matrix of size $i \times j$. Finally, let $\mathcal{T} \sim \mathcal{N}_{i \times j \times k}(\mu, \sigma)$ represent sampling a tensor \mathcal{T} of size $i \times j \times k$ from a normal distribution with a mean of μ and standard deviation of σ . \circ represents the tensor outer product.

2.1 CANDECOMP/PARAFAC

One of the most common tensor decompositions is the canonical polyadic or CANDECOMP/PARAFAC decomposition (CPD) [6, 10]. The CPD decomposes a tensor \mathcal{T} into the sum of R outer products of vectors. Formally, for a third-order tensor:

$$\mathcal{T} \approx \sum_{r=1}^R \mathbf{a}_r \circ \mathbf{b}_r \circ \mathbf{c}_r \quad (2)$$

We can then attempt to iteratively minimize the Frobenius norm of the error: $\left\| \mathcal{T} - \sum_{r=1}^R \mathbf{a}_r \circ \mathbf{b}_r \circ \mathbf{c}_r \right\|_F$. A common method for this is to use the alternating least squares algorithm. By convention, we refer to the factors as matrices $\mathbf{A}, \mathbf{B}, \mathbf{C}$, which consist of the corresponding vectors horizontally stacked. For example, the i -th column of \mathbf{A} would be \mathbf{a}_i . A key component of this method is R —the *rank of the decomposition*. As described above, choosing a good value for R is extremely difficult. The ideal value for R is the canonical rank of \mathcal{T} , and estimating that value is the focus of this work.

3 Problem Formulation

We tackle the problem of predicting the canonical rank of a tensor. Formally, the problem can be stated as follows:

Problem Definition

Given a tensor $\mathcal{X} = \mathcal{L} + \mathcal{N}$, where \mathcal{L} is a tensor with low rank and \mathcal{N} is tensor of high-rank noise, find the canonical rank R of \mathcal{X} such that we minimize $\left\| \mathcal{L} - \sum_{r=1}^R \mathbf{a}_r \circ \mathbf{b}_r \circ \mathbf{c}_r \right\|_F$ where $\mathbf{A}, \mathbf{B}, \mathbf{C}$ are factors from the CPD of \mathcal{X} .

Note that R is different from the *exact rank* of the tensor (the minimum number of components needed to exactly represent \mathcal{X}), which is out of the scope of this work.

Another secondary goal of our method is to identify which features are important in predicting the rank of a tensor, which could help us develop/improve rank prediction methods in the future. We define this task as follows:

Subproblem: Explainability

Given a tensor \mathcal{X} and a predicted rank \hat{R} , compute the importance scores \mathbf{t} of different attributes of \mathcal{X} in the regression result.

4 Proposed Method

We first describe which hand-picked features we extract from the tensor and the reasoning behind selecting each one. We then train a regression model on the dataset and show that it has some key limitations: (a) it performs well on the synthetic dataset but fails to generalize well to real data, and (b) it requires a relatively expensive training step. We then propose our method, FRAPPE, which addresses both of these issues.

4.1 Tensor Features

Given a tensor $\mathcal{T} \in \mathbb{R}^{I \times J \times K}$, we extract 112 features (4D or higher tensors require some adjustment, as described below):

- **Tensor Dimensions (#1–3):** I, J, K . We include these features so that the underlying model can adjust the expectations for other parameters relative to the size of the tensor. This allows our model to generalize to different input sizes.
- **Non-zeros (#4–13):** The total number of non-zero values in the tensor, as well as the minimum, median, and maximum number of non-zero values over all possible slices of the tensor. That is, we extract the number of non-zeros from $\mathcal{T}_{(i,:,:)} \forall i \in [1, I]$, $\mathcal{T}_{(:,j,:)} \forall j \in [1, J]$, and $\mathcal{T}_{(:,:,k)} \forall k \in [1, K]$. The intuition behind this feature is that, generally, the rank of a tensor increases as it becomes more sparse. Measuring this for all possible slices allows us to essentially measure the sparsity along each dimension.
- **Ranks of Slices (#14–22):** The minimum, median, and maximum rank over all possible slices of the tensor. The intuition behind this feature is that if slices along any given two dimensions are generally of low rank (implying that fewer components are needed in the SVD), fewer components may also be required in the CPD.
- **Thresholded Non-zeros (#23–103):** The minimum, median, and maximum number of values in the *normalized* (we use min-max normalization) tensor less than or equal to each of the following thresholds: $[0.1, 0.2, \dots, 0.9]$ across all possible slices of the tensor. This is an adaptation of the number of non-zeros (features #4–13) for dense tensors and allows us to find the number of low-value entries, which could be a result of noise and be of lesser importance to the canonical rank of the tensor.
- **Correlations (#104–112):** The minimum, median, and maximum correlation over all possible pairs of slices of the tensor. The intuition behind

this feature is that higher correlation between slices generally indicates that fewer components would be needed along a given axis, resulting in a lower canonical rank.

Note that the number of features is constant, *regardless of the size of the tensor*. This allows it to easily scale to large tensors. We then train a LightGBM [13] model on these features, with the goal of predicting the rank of the target tensor.

Why Correlation? It may not be immediately apparent why we choose to use the correlation between different slices. We choose to do so because the distribution of pairwise correlation of vectors in a matrix is a good indicator of its rank. We show this theoretically (under some assumptions) below. We hypothesize that this remains true in the tensor case; while difficult to formally prove, we also show that this feature is effective in our model in Fig. 4.

Claim. The rank of a matrix \mathbf{M} decreases as the mean pairwise correlation of its factors increases. Let $\mathbf{M} \in \mathbb{R}^{m \times n}$. We assume, for the sake of simplicity, that a set \mathcal{C} contains all $\frac{n}{2}$ non-overlapping pairs of columns of \mathbf{M} , and all other pairs have a correlation of 0.

Proof. We first define the rank of a matrix \mathbf{M} as k , where k is the largest value such that the k -th singular value $\sigma_k > \epsilon$. Let us assume that any two vectors \mathbf{x}, \mathbf{y} has an outer product $\mathbf{x}\mathbf{y}^T$ that can be written as $\mathbf{a}\mathbf{b}^T + \mathbf{N}$, where $\mathbf{N} = \gamma\sigma_1\mathbf{u}_1\mathbf{v}_1^T + \gamma^2\sigma_2\mathbf{u}_2\mathbf{v}_2^T + \dots + \gamma^k\sigma_k\mathbf{u}_k\mathbf{v}_k^T$ such that $\sigma_i < \epsilon \forall i$, where ϵ is an arbitrarily small number and $\text{CORR}(\mathbf{x}, \mathbf{y}) \rightarrow 1 \implies \gamma \rightarrow 0$. Let $c = \mathbb{E}[\text{CORR}(\mathbf{M}^{(i)}, \mathbf{M}^{(j)})] \forall (\mathbf{M}^{(i)}, \mathbf{M}^{(j)}) \in \mathcal{C}$. Let $\mathbf{M}^{(i)}$ be the i -th column of \mathbf{M} . Then, we can write \mathbf{M} as $\mathbf{a}_1\mathbf{b}_1^T + \mathbf{a}_2\mathbf{b}_2^T + \dots + \mathbf{a}_{n/2}\mathbf{b}_{n/2}^T + \mathbf{N}_1 + \mathbf{N}_2 + \dots + \mathbf{N}_{n/2}$. Let $\mathbf{N}_\Sigma = \mathbf{N}_1 + \mathbf{N}_2 + \dots + \mathbf{N}_{n/2}$. By our assumption above, we have at most $\frac{n}{2}$ duplicate components so the maximum value for a singular value of \mathbf{N} is $\frac{n}{2}\gamma c \epsilon$, which is still less than ϵ for $\gamma < \frac{2}{n}$, and subsequently true for sufficiently large c . Therefore, the rank of \mathbf{M} decreases as our mean pairwise correlation c increases.

Computational Complexity. Features #1–13 and #23–103 are extremely quick to calculate, requiring only a single pass over the tensor. Features #14–22 and #104–112 are significantly more expensive since they require computing the SVD and the correlation, respectively, over various slices of the tensor. However, these operations are still much faster than the baseline methods, which require repeatedly computing the full CPD over the whole tensor. Table 3 compares the speed of our method against the baselines.

Let $n = \max(I, J, K)$. Features #1–13 take a linear amount of time to calculate with respect to the number of non-zero values in the tensor, with a worst-case of $\mathcal{O}(n^3)$ (in the case of a dense tensor). Features #14–22 take $\mathcal{O}((I + J + K)s) = \mathcal{O}(ns)$ time, where s is the runtime of sparse SVD on a slice of our tensor. In the case of a fully dense tensor, this takes $\mathcal{O}((I + J + K)n^3) = \mathcal{O}(n^4)$ time to calculate.

Features #23–103 require the tensor to be min-max normalized first, which can be done in $\mathcal{O}(n^3)$ time. It also takes $\mathcal{O}(n^3)$ to count the number above/below each threshold. Features #104–112 take $\mathcal{O}(n^2)$ time to calculate over $\frac{n(n-1)}{2}$ pairs, resulting in a runtime of $\mathcal{O}(n^4)$.

Overall, these features take a *maximum* of $\mathcal{O}(n^4)$ time to calculate for dense tensors, where n is the size of the largest dimension in the tensor. It is worth noting that it would take significantly less time for sparse tensors (e.g., a multi-view graph, knowledge base, word cooccurrence). This is also a massive improvement compared to the runtime of computing the CPD once, let alone the repeated runs required by existing methods. We show a comparison of the runtimes below in Table 3.

Extension to Higher-Order Tensors While the features described above work well for 3rd-order tensors, they require some modification for 4th or higher-order tensors. For example, the rank features are computed over all slices of the tensor. In the 3rd order case, these slices are matrices, but in the 4th (or higher) order case, these slices are 3rd (or higher) order sub-tensors. Finding the rank of a matrix can be done in polynomial time via the SVD, but finding the rank of a 3D tensor is the problem we are tackling in this paper. Due to this issue and the relatively low importance of the rank features in Fig. 4, we omit the rank features for the 4D model and find that it still performs fairly well. This is shown empirically in our experiments on a 4D AlexNet weight tensor below in Sec. 5.4. As with existing work in canonical rank estimation [8, 20], we focus on primarily on 3rd-order tensors and reserve detailed evaluation for future work.

4.2 A First Attempt

We can now train a regression model using the features defined in Sec. 4.1 above. To do this, we first generate a synthetic dataset with tensors of known rank (see Sec. 5.2). We then train a LightGBM [13] model on this synthetic dataset and evaluate the model on each of the evaluation datasets in Sec. 5 below. However, we find that the LightGBM model is unable to generalize well beyond the synthetic datasets and incorrectly predicts the rank on all of the real-world datasets (described in Sec. 5.3). This is likely because real-world tensors come from a variety of different domains.

Another limitation of this approach is that we have to train our model on a very large dataset to ensure that we account for different ranks, sparsities, and noise levels. While this makes predicting the rank of a new tensor very quick, it greatly increases the amount of time required to train a regression model. To address this, we propose FRAPPE.

4.3 FRAPPE

The primary issue we need to solve with FRAPPE compared to the approach described in Sec. 4.2 is the poor generalization of the naive regression model.

Algorithm 1. Pseudocode for FRAPPE: a self-supervised algorithm to estimate the rank of a given tensor. Note that comments precede the line that they describe.

Input: An input tensor $\mathcal{T} \in \mathbb{R}^{i \times j \times k}$, ρ (the maximum rank to consider), τ (the number of times to sample from each rank), and α (the maximum amount of noise to add in synthetic samples).

Output: \hat{R} , an estimate for the rank of the input tensor \mathcal{T} .

```

1: function FRAPPE( $\mathcal{T}$ ,  $\rho$ ,  $\tau$ ,  $\alpha$ )
2:   data  $\leftarrow [ ]$ 
3:   ranks  $\leftarrow [ ]$ 
4:    $\triangleright$  Calculate the sparsity of our tensor
5:    $s \leftarrow \frac{\text{ZEROS}(\mathcal{T})}{i \times j \times k}$ 
6:   for  $x \in [1, \tau]$  do
7:     for  $r \in [1, \rho]$  do
8:        $\mathbf{A} \leftarrow \text{SPARSIFY}(\text{RAND}(i, r), s)$ 
9:        $\mathbf{B} \leftarrow \text{SPARSIFY}(\text{RAND}(j, r), s)$ 
10:       $\mathbf{C} \leftarrow \text{SPARSIFY}(\text{RAND}(k, r), s)$ 
11:       $\triangleright$  Construct a synthetic tensor  $\mathcal{S}$  from  $\mathbf{A}, \mathbf{B}, \mathbf{C}$ 
12:       $\mathcal{S} \leftarrow \sum_{y=1}^r \mathbf{a}_y \circ \mathbf{b}_y \circ \mathbf{c}_y$ 
13:       $\triangleright$  Select a random noise ratio.
14:       $n \sim [0, \alpha]$ 
15:       $\triangleright$  Sample an  $i \times j \times k$  Gaussian noise tensor
16:       $\mathcal{N} \sim \mathcal{N}_{i \times j \times k}(0, 1)$ 
17:       $\triangleright$  Scale the noise and add it to  $\mathcal{S}$ 
18:       $\mathcal{S} \leftarrow \mathcal{S} + (n \frac{\|\mathcal{S}\|_F}{\|\mathcal{N}\|_F}) \mathcal{N}$ 
19:       $\triangleright$  Extract features from  $\mathcal{S}$ , as described in Sec. 4.1
20:      data.append(EXTRACTFEATURES( $\mathcal{S}$ ))
21:      ranks.append( $r$ )
22:   Train a LightGBM regression model on data and ranks
23:   Evaluate the model on EXTRACTFEATURES( $\mathcal{T}$ ) to get our prediction  $\hat{r}$ 
24:   return  $\hat{R}$ 

```

Input: An input matrix $\mathbf{M} \in \mathbb{R}^{a \times b}$ and the desired sparsity level s .

Output: The sparsified matrix.

```

1: function SPARSIFY( $\mathbf{M}$ ,  $s$ )
2:   Select  $s \times a \times b$  random indices of  $\mathbf{M}$ 
3:   Set each of these indices to 0
4:   return  $\mathbf{M}$ 

```

This poor generalization is fundamentally caused by a distribution difference between the training data and the input tensor. We can address this issue by ensuring that the synthetic training data closely mimics a given input tensor. We do this by attempting to match the size and sparsity of the input tensor when generating the synthetic dataset. We can then train a specialized LightGBM model *for each tensor being considered*. This specialized model only performs evaluation on the input tensor once before being discarded.

Generating training data engineered to mimic an input tensor also solves the issue of long training times since less data is required to achieve the same results.

This brings down training times significantly, allowing us to train a disposable regression model for each input tensor while maintaining a speed advantage over the other baseline models. Note that the “training” happens at *evaluation* time, and no external training step is required. FRAPPE is formally described in detail in Algorithm 1.

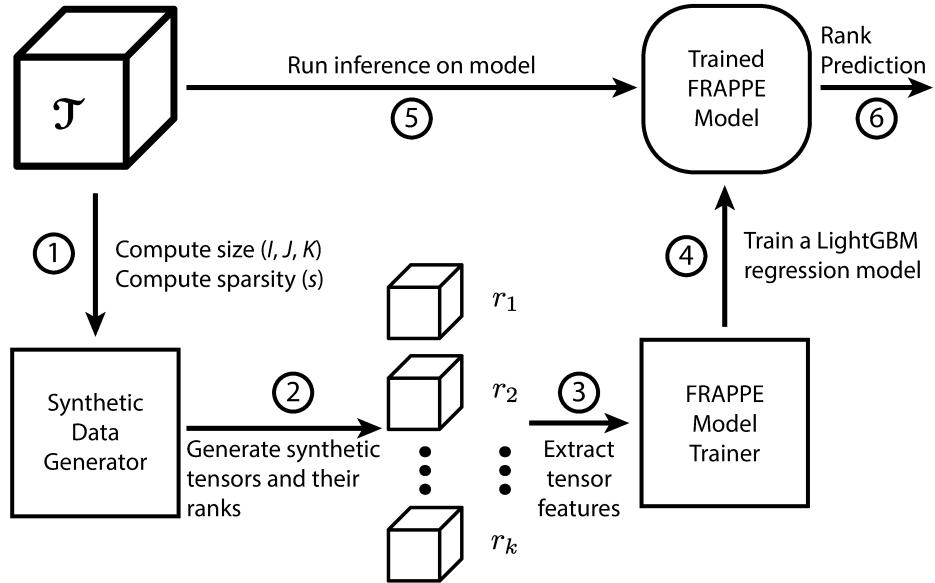


Figure 1. Diagram of FRAPPE. We first generate synthetic tensors based on the size and sparsity of the input tensor, extract features, train the model, and use that model to predict the rank of the original input tensor.

5 Experimental Evaluation

Evaluating the effectiveness of our methods is difficult because there are very few tensors with known rank so we cannot solely rely on real-world data. As such, we choose to evaluate our methods on three different types of data:

1. **Synthetic Tensors:** Predicting the rank of tensors that we generate to have known rank (described in Sec. 5.2).
2. **Real-world Tensors:** Predicting the rank of real-world tensors that have been previously studied and have known rank (described in Sec. 5.3).
3. **CNN Weights:** Compressing the 4D weights of an AlexNet [16] model’s final layer (described in Sec. 5.4).

We compare the performance and runtime of our methods to other current state-of-the-art methods. All of the timing experiments are conducted on a

c3-highcpu-176 Google Cloud Platform instance with 176 CPU cores and 352 GB of RAM.

5.1 Baseline Methods

We compare FRAPPE against three baselines: NORMO [8], AutoTen [20] and ARD-Tucker [17]. These methods require a maximum rank to consider. Since this greatly affects the runtime of the algorithm, we used double the true rank of the tensor for the two algorithms. This theoretically bounds the maximum error, but setting it higher would greatly increase both the runtime and computational cost of running the baselines. We also use this same bound for FRAPPE.

We had to slightly modify AutoTen to only use the Frobenius-norm-based evaluation because we have negative values in our dataset, which causes the KL-divergence-based evaluation to fail. This slightly reduces the accuracy of AutoTen but greatly improves its runtime. We also had to modify NORMO and AutoTen to add support for 4D tensors in order to run it on our CNN weight tensor. We chose to use the suggested value of $\delta = 0.7$ for NORMO.

Dataset	Rank	FRAPPE	NORMO	AutoTen	ARD
amino_acid [3]	3 [3, 14]	3	4	3	3
dorrit [23]	4 [23]	4	4	4	4
sugar [4]	3 [4]	3	4	2	3
reality_mining [7]	6? ³ [19]	9	2	2	18
enron [22]	4/7? [2, 21]	3	1	4	12

Table 2. Performance of the rank estimation methods on real-world datasets. The first column is the known rank of the datasets, and correct estimates are bolded. We can see that FRAPPE and AutoTen generally perform the best overall.

5.2 Synthetic Dataset

To generate a synthetic dataset, we first generate the CPD factor matrices $\mathbf{A}, \mathbf{B}, \mathbf{C}$ by drawing values from a uniform random distribution. We ensure that we generate a diverse dataset by generating a set percentage of tensors from each of the following subsets (each of which corresponds to a common real-world tensor):

¹ NORMO is partially explainable in that higher correlation in the factors indicates a bad choice for a given rank. However, NORMO does not provide any indication as to which attributes in the *original* tensor contribute to higher/lower predictions.

² TRN-PD is designed for images and was originally evaluated for CPD-based denoising.

³ The number of communities in this dataset was explored in [19] and estimated to be 6. The number of communities loosely corresponds to the canonical rank of the tensor, but this is not a guarantee like with the fluorescence datasets.

1. **Sparse tensors with sparse noise:** Tensors that have sparsity introduced in the factors and Gaussian noise added to non-zero elements in the tensor. *Example:* a weighted time-evolving graph.
2. **Sparse tensors with dense noise:** Tensors that have sparsity introduced in the factors and Gaussian noise added to all elements in the tensor. *Example:* a time-evolving grid of sensor readings.
3. **Dense tensors:** Tensors that have fully dense tensors and dense noise added. *Example:* a tensor of fluorometer readings.

Each of these cases represents a different type of tensor that can be found in real-world datasets. We add 2%–10% noise in each case and scale the amount of noise by the norm of the noise tensor compared to the generated tensor. More formally, where α is the noise ratio, \mathcal{N} is our noise tensor and \mathcal{G} is our generated tensor: $\mathcal{G} = \mathcal{G} + \alpha \frac{\|\mathcal{G}\|_F}{\|\mathcal{N}\|_F} \mathcal{N}$. This is similar to the noise addition process in the Tensor Toolbox’s [1] `create_problem` function.

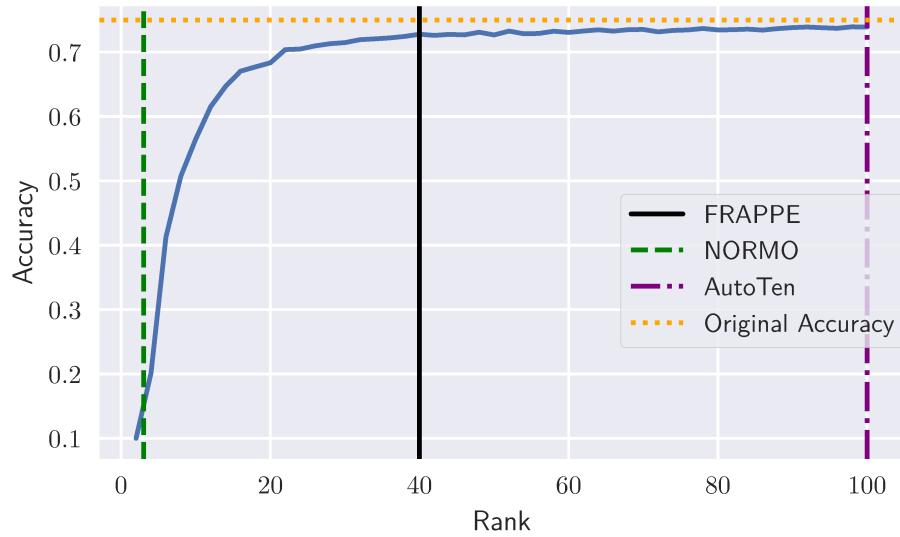


Figure 2. Plot of AlexNet performance with the final CNN layer compressed at various ranks. Also shown are the rank estimates produced by FRAPPE and the baselines. NORMO predicts a very low value and AutoTen predicts a very high value. FRAPPE and FRAPPE pick what appear to be reasonable values. ARD-Tucker did not terminate in time and is not included.

5.3 Real-World Datasets

We evaluate the effectiveness of our methods and baseline methods on datasets of known rank. These include three fluorescence datasets: `amino_acid` [3],

`dorrit` [23], and `sugar` [4]. These datasets contain naturally tri-linear data with a known number of chemicals, each of which corresponds to a component in the CPD. We also evaluate our methods on the `reality_mining` [7] and `enron` [22] datasets. The canonical rank of these tensors is less clear, as we have to rely on existing literature that has manually analyzed the data and determined the approximate rank.

From the results displayed in Table 2, we can see FRAPPE is able to perfectly estimate the rank of the fluorescence datasets. ARD is also able to correctly estimate the rank of the datasets. NORMO and AutoTen perform fairly well but are slightly off. Only AutoTen is able to successfully estimate the rank for the `enron` dataset, but FRAPPE is also very close. Unfortunately, we are limited in the amount of data we can use for this evaluation because there are very few real-world tensors with known ranks. While it is difficult to draw any concrete conclusions due to the limited number of datasets, we can see that FRAPPE performs just as well as or slightly better than other existing rank estimation methods.

5.4 CNN-Based Evaluation

We evaluate the performance of our network in selecting the correct number of components when decomposing a 4D weight tensor for a convolutional neural network (CNN). We focus specifically on the last layer ($\mathbb{R}^{256 \times 256 \times 3 \times 3}$) of an AlexNet [16] model trained on the CIFAR-10 [15] image dataset.

The goal of this task is to select the correct number of components such that we compress the weights as much as possible and maintain as high accuracy as possible. This rank would be the “knee” of the rank-accuracy curve. A plot of this curve with the values selected by the various rank estimation methods is shown in Fig. 2.

NORMO predicts a low estimate of 3 for the rank of the weight tensor, and AutoTen predicts the maximum estimate of 100. Both of these appear to be poor choices for this task since the accuracy is very low at rank 3 and the gains are heavily diminished by rank 100. FRAPPE estimates the rank to be 40, which is not perfect, but appears to be fairly reasonable estimates near the “knee” of the plot. ARD-Tucker did not terminate within 24 hours and is therefore excluded from the plot.

5.5 Synthetic Dataset Results

Table 3 shows that FRAPPE outperforms the baseline methods, with a $\sim 10\%$ improvement in Mean Absolute Percentage Error (MAPE) over NORMO (the best-performing baseline), ~ 4 rank improvement in Mean Absolute Error (MAE), and an improvement of 200 in Mean Squared Error (MSE). It is also $\sim 24\times$ faster than NORMO, and $\sim 1.6\times$ faster than AutoTen (the fastest baseline).

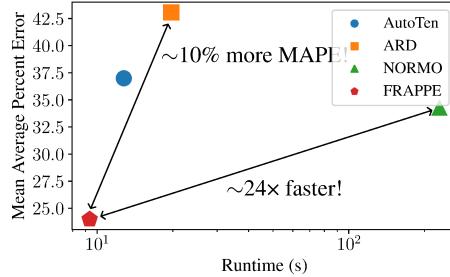


Figure 3. Runtime/performance comparison between the different rank estimation methods on a synthetic dataset. The closer to the lower left, the better.

Model	MAPE	MAE	MSE	Runtime
AutoTen ⁴ [20]	36.99%	11.55	314.80	12.77s
ARD [17]	43.06%	13.52	414.70	19.72s
NORMO [8]	34.36%	10.44	282.38	229.34s
FRAPPE	23.99%	6.43	81.40	9.34s

Table 3. Mean Absolute Percentage Error (MAPE), Mean Absolute Error (MAE), and Mean Square Error (MSE), and mean runtime of different models on a synthetic dataset.

5.6 Feature Importance Results

In Fig. 4, we plot the importance of the different feature groups (described in Sec. 4.1 above). We compute the importance of a feature based on the number of times it is used in the split of a LightGBM tree.

We can see that the most important features are the correlation over different slices. This makes intuitive sense because higher correlation between slices along a given mode means that components are more related along that mode, leading to fewer components being required. This also happens to be the intuition behind NORMO, one of the baseline methods, which looks instead at the correlation between different factors of the CPD.

Another interesting set of features are the thresholded non-zero counts. In Fig. 4, we can see that the importance increases from a threshold of 0.1 to 0.2, but steadily decreases after that until 0.9, which is almost irrelevant. This is likely because lower relative values (since we normalize the values) are more likely to be noise artifacts or simply unimportant values in the decomposition. Lower values are also more likely to be excluded in an R -rank decomposition, simply because the incremental error from excluding that value is smaller than it would be for a larger-valued element.

6 Additional Related Work

On top of the CPD rank estimation methods described above, there have also been similar methods developed for different tensor factorizations. Vesselinov et al. [26] propose NTF k , which is an unsupervised method to perform Non-negative Tensor Factorization (NTF). NTF k uses a custom clustering algorithm (similar to k -means) to cluster the columns of a factor matrix and evaluates their quality with silhouette score. It then uses a combination of those scores and the reconstruction error to choose a good target for the number of components.

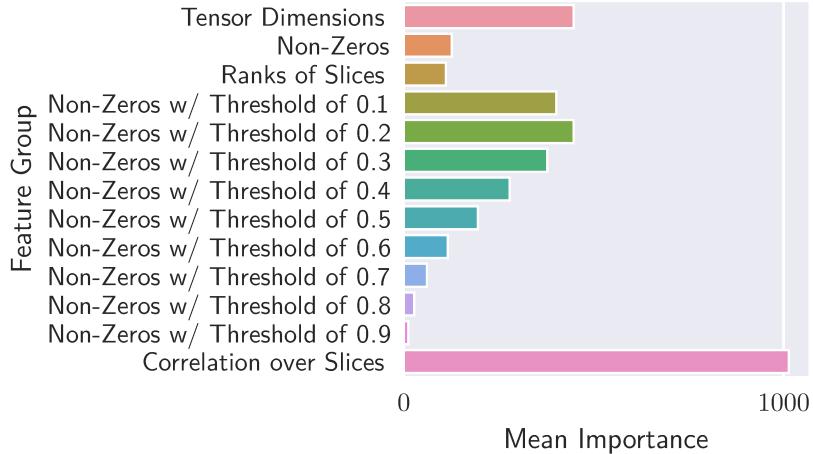


Figure 4. Chart of feature group importance (higher is more important), as output by a LightGBM model trained on a synthetic dataset. The most important feature group is the correlation over slices of the tensor, followed by the number of non-zeros thresholded at 0.2. We calculate feature group importance by how often a feature is used in splits. The feature groups are described in more detail in Sec. 4.1.

Nebgen et al. [18] attempt to solve a similar problem for non-negative matrix factorization (NMF), which decomposes a matrix \mathbf{V} into $\mathbf{V} = \mathbf{W}\mathbf{H}$.

They first repeatedly calculate the NMF on data perturbed by noise from a random uniform distribution. Then, they use k -means to cluster the columns of \mathbf{W} . They then use the Akaike information criterion (AIC) and silhouette score to calculate the quality of the clusters. Finally, they use a neural network over a sliding window of the AIC and cluster silhouette values to determine the ideal number of components.

7 Conclusion

In this work, we propose FRAPPE, a novel method for estimating the canonical rank of a tensor in a self-supervised manner without having to compute the expensive CPD. The key insight of this approach is that it is much cheaper to generate synthetic data of known rank than to compute the actual CPD. Following this idea, we also propose a set of hand-crafted features that can effectively determine a given tensor’s rank.

We find that naively training a supervised regression model on this data results in poor generalization to real-world tensors and requires a long training step. To alleviate this issue, we propose FRAPPE, which first generates synthetic tensors (with known rank) from an input tensor and trains a new model on these synthetic tensors, improving the generalization ability of the model, reducing the amount of training data required, and speeding up the model training.

We extensively evaluate FRAPPE on three different types of data: synthetic tensors, real-world tensors, and the weight tensor of a CNN. We find that our model generally outperforms existing baselines across all three datasets, is over $24\times$ faster than the best-performing baseline (NORMO), and $\sim 1.6\times$ faster than the fastest baseline (AutoTen).

Finally, we examine the features we used to make these predictions and why we believe these features are important in estimating a tensor’s rank. We provide both theoretical reasoning and empirical evidence for why the pair-wise correlation between slices is a good indicator of the tensor’s rank. We believe this opens up an interesting line of research about estimating the number of components of a tensor decomposition without ever having to compute a single decomposition.

8 Ethical Statement

One important ethical consideration regarding this work can arise if it is used on sensitive information like financial or healthcare data. As with all machine learning methods, our method is not infallible and may produce incorrect results. In certain mission-critical systems, that can potentially lead to severe consequences if care is not taken to validate model results.

Another ethical concern is the possibility of adversarial attacks on the model. In our work, we do not consider the case of adversarially crafted input tensors designed to trick the model into performing poorly or predicting an extremely high rank, which could potentially result in maliciously altered results.

Acknowledgments

This research was supported by the National Science Foundation under CAREER grant no. IIS 2046086 and CREST Center for Multidisciplinary Research Excellence in Cyber-Physical Infrastructure Systems (MECIS) grant no. 2112650. This research was also sponsored by the Combat Capabilities Development Command Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-13-2-0045 (ARL Cyber Security CRA).

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Combat Capabilities Development Command Army Research Laboratory, the United States Air Force, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

References

1. Bader, B., Kolda, T.: Matlab tensor toolbox version 2.2. Albuquerque, NM, USA: Sandia National Laboratories (2007)

2. Bader, B.W., Harshman, R.A., Kolda, T.G.: Pattern analysis of directed graphs using dedicom: an application to enron email. Tech. rep., Sandia National Laboratories (SNL), Albuquerque, NM, and Livermore, CA ... (2006)
3. Bro, R.: Parafac. tutorial and applications. *Chemometrics and Intelligent Laboratory Systems* **38**(2), 149–171 (1997), ISSN 0169-7439, [https://doi.org/https://doi.org/10.1016/S0169-7439\(97\)00032-4](https://doi.org/https://doi.org/10.1016/S0169-7439(97)00032-4)
4. Bro, R.: Exploratory study of sugar production using fluorescence spectroscopy and multi-way analysis. *Chemometrics and Intelligent Laboratory Systems* **46**, 133–147 (1999)
5. Bro, R., Kiers, H.A.: A new efficient method for determining the number of components in parafac models. *Journal of chemometrics* **17**(5), 274–286 (2003)
6. Carroll, J.D., Chang, J.J.: Analysis of individual differences in multidimensional scaling via an n-way generalization of “eckart-young”decomposition. *Psychometrika* **35**(3), 283–319 (1970), <https://doi.org/10.1007/BF02310791>
7. Eagle, N., (Sandy) Pentland, A.: Reality mining: Sensing complex social systems. *Personal Ubiquitous Comput.* **10**(4), 255–268 (mar 2006), ISSN 1617-4909, <https://doi.org/10.1007/s00779-005-0046-3>
8. Fernandes, S., Fanaee-T, H., Gama, J.: Normo: A new method for estimating the number of components in cp tensor decomposition. *Engineering Applications of Artificial Intelligence* **96**, 103926 (2020), ISSN 0952-1976, <https://doi.org/https://doi.org/10.1016/j.engappai.2020.103926>
9. Harshman, R.: Foundations of the parafac procedure: Models and conditions for an “explanatory” multimodal factor analysis (1970)
10. Harshman, R.A., et al.: Foundations of the parafac procedure: Models and conditions for an “explanatory” multimodal factor analysis. *UCLA Working Papers in Phonetics* (1970)
11. Håstad, J.: Tensor rank is np-complete. *Journal of Algorithms* **11**(4), 644–654 (1990)
12. Hillar, C.J., Lim, L.H.: Most tensor problems are np-hard. *Journal of the ACM (JACM)* **60**(6), 45 (2013)
13. Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., Liu, T.Y.: Lightgbm: A highly efficient gradient boosting decision tree. *NeurIPS* **30** (2017)
14. Kiers, H.: A three-step algorithm for candecomp/parafac analysis of large data sets with multicollinearity. *Journal of Chemometrics - J CHEMOMETR* **12**, 155–171 (05 1998), [https://doi.org/10.1002/\(SICI\)1099-128X\(199805/06\)12:33.3.CO;2-X](https://doi.org/10.1002/(SICI)1099-128X(199805/06)12:33.3.CO;2-X)
15. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images. *Citeseer* (2009)
16. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Pereira, F., Burges, C.J.C., Bottou, L., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems*, vol. 25, Curran Associates, Inc. (2012)
17. Mørup, M., Hansen, L.K.: Automatic relevance determination for multi-way models. *Journal of Chemometrics* **23**(7-8), 352–363 (2009)
18. Nebgen, B.T., Vangara, R., Hombrados-Herrera, M.A., Kuksova, S., Alexandrov, B.S.: A neural network for determination of latent dimensionality in non-negative matrix factorization. *Machine Learning: Science and Technology* **2**, 25012 (2020), <https://doi.org/10.1088/2632-2153/aba372>, URL <https://doi.org/10.1088/2632-2153/aba372>
19. Papalexakis, E.E., Akoglu, L., Ience, D.: Do more views of a graph help? community detection and clustering in multi-graphs. In: *FUSION '16*, pp. 899–905 (2013)
20. Papalexakis, Evangelos E: Automatic unsupervised tensor mining with quality assessment. In: *SIAM SDM* (2016)

21. Pasricha, R., Gujral, E., Papalexakis, E.E.: Identifying and alleviating concept drift in streaming tensor decomposition. CoRR **abs/1804.09619** (2018), URL <http://arxiv.org/abs/1804.09619>
22. Priebe, C.E., Conroy, J.M., Marchette, D.J., Park, Y.: Scan statistics on enron graphs. Computational & Mathematical Organization Theory **11**(3), 229–247 (Oct 2005), ISSN 1572-9346, <https://doi.org/10.1007/s10588-005-5378-z>
23. Riu, J., Bro, R.: Jack-knife technique for outlier detection and estimation of standard errors in parafac models. Chemometrics and Intelligent Laboratory Systems **65**(1), 35–49 (2003), ISSN 0169-7439, [https://doi.org/https://doi.org/10.1016/S0169-7439\(02\)00090-4](https://doi.org/https://doi.org/10.1016/S0169-7439(02)00090-4)
24. Tsitsikas, Y., Papalexakis, E.E.: NSVD : N ormalized S ingular V alue D eviation Reveals Number of Latent Factors in Tensor Decomposition . In: Proceedings of the 2020 SIAM International Conference on Data Mining, pp. 667–675, Society for Industrial and Applied Mathematics (1 2020), <https://doi.org/10.1137/1.9781611976236.75>
25. Tucker, L.: Some mathematical notes on three-mode factor analysis. Psychometrika **31**(3), 279–311 (1966)
26. Vesselinov, V.V., Mudunuru, M.K., Karra, S., O’Malley, D., Alexandrov, B.S.: Unsupervised machine learning based on non-negative tensor factorization for analyzing reactive-mixing. CoRR **abs/1805.06454** (2018), URL <http://arxiv.org/abs/1805.06454>
27. Zhao, Q., Zhang, L., Cichocki, A.: Bayesian CP factorization of incomplete tensors with automatic rank determination. CoRR **abs/1401.6497** (2014), URL <http://arxiv.org/abs/1401.6497>
28. Zhou, M., Liu, Y., Long, Z., Chen, L., Zhu, C.: Tensor rank learning in cp decomposition via convolutional neural network. Signal Processing: Image Communication **73**, 12–21 (2019), ISSN 0923-5965, <https://doi.org/10.1016/j.image.2018.03.017>, tensor Image Processing