

COMPUTING TENSOR Z -EIGENVECTORS WITH DYNAMICAL SYSTEMS*

AUSTIN R. BENSON[†] AND DAVID F. GLEICH[‡]

Abstract. We present a new framework for computing Z -eigenvectors of general tensors based on numerically integrating a dynamical system that can only converge to a Z -eigenvector. Our motivation comes from our recent research on spacey random walks, where the long-term dynamics of a stochastic process are governed by a dynamical system that must converge to a Z -eigenvector of a transition probability tensor. Here, we apply the ideas more broadly to general tensors and find that our method can compute Z -eigenvectors that algebraic methods like the higher-order power method cannot compute.

Key words. tensors, eigenvectors, dynamical systems

AMS subject classifications. 15A69, 65F15, 65P99

DOI. 10.1137/18M1229584

1. Preliminaries on tensor eigenvectors. Computing matrix eigenvalues is a classic problem in numerical linear algebra and scientific computing. Given a square matrix \mathbf{A} , the goal is to find a vector-scalar pair (\mathbf{x}, λ) with $\mathbf{x} \neq 0$ satisfying

$$(1.1) \quad \mathbf{Ax} = \lambda\mathbf{x}.$$

The pair (\mathbf{x}, λ) is called the eigenpair, \mathbf{x} the eigenvector, and λ the eigenvalue. After several decades of research and development, we have, by and large, reliable methods and software for computing *all* eigenpairs of a given matrix \mathbf{A} . (Experts will, of course, be aware of exceptions, but we hope they would agree with the general sentiment of the statement.)

In numerical *multilinear* algebra, there are analogous eigenvector problems (note the plurality). For example, given a three-mode cubic tensor $\underline{\mathbf{T}}$ (here meaning that $\underline{\mathbf{T}}$ is a multidimensional $n \times n \times n$ array with entries $\underline{T}_{i,j,k}$, $1 \leq i, j, k, \leq n$), the two most common tensor eigenvector problems are the following:

$$\begin{array}{ll} Z\text{-eigenvectors (Qi, 2005)} & H\text{-eigenvectors (Qi, 2005)} \\ l^2\text{-eigenvectors (Lim, 2005)} & l^k\text{-eigenvectors (Lim, 2005)} \\ \sum_{jk} \underline{T}_{i,j,k} x_j x_k = \lambda x_i, 1 \leq i \leq n & \sum_{jk} \underline{T}_{i,j,k} x_j x_k = \lambda x_i^2, 1 \leq i \leq n \\ \|\mathbf{x}\|_2 = 1. & \mathbf{x} \neq 0. \end{array}$$

We use the “ Z ” and “ H ” terminology instead of “ l^2 ” and “ l^k .” Both Z - and H -eigenvectors are defined for tensors with the dimension equal in all modes (such a

*Received by the editors November 29, 2018; accepted for publication (in revised form) by M. Mohlenkamp August 5, 2019; published electronically November 12, 2019.

<https://doi.org/10.1137/18M1229584>

Funding: The work of the first author was partly supported by NSF award DMS-1830274 and ARO award W911NF-357 19-1-0057. The work of the second author was supported by NSF award CCF-1149756, IIS-1422918, IIS-1546488, the NSF Center for Science of Information STC, CCF-0939370, DARPA SIMPLEX, NASA, DOE DE-SC0014543, and the Sloan Foundation.

[†]Department of Computer Science, Cornell University, Ithaca, NY 14853 (arb@cs.cornell.edu).

[‡]Department of Computer Science, Purdue University, West Lafayette, IN 47907 (dgleich@purdue.edu).

¹In true pure mathematical parlance, this is the definition of a hypermatrix, not a tensor (see Lim (2013) for the precise definitions.) However, “tensor” has become synonymous with a multidimensional array of numbers (Kolda and Bader, 2009), and we adopt this terminology here.

tensor is called *cubic* (Comon et al., 2008)). The definitions can be derived by showing that the eigenpairs are KKT points for a generalization of a Rayleigh quotient to tensors (Lim, 2005). One key difference between the types is that H -eigenvectors are scale-invariant, while Z -eigenvectors are not—this is why we put a norm constraint on the vector. Specifically, if we ignore the norm constraint and scale \mathbf{x} by a constant, the corresponding Z -eigenvalue would change; for H -eigenpairs, this is not the case. If $\underline{\mathbf{T}}$ is symmetric, then it has a finite set of Z -eigenvalues, and, moreover, there must be a real eigenpair when the order of the tensor (i.e., the number of modes or indices) is odd (Cartwright and Sturmfels, 2013).

This paper presents a new framework for computing Z -eigenpairs. Tensor Z -eigenvectors show up in a variety of applications, including evolutionary biology (Bini, Meini, and Poloni, 2011; Meini and Poloni, 2011), low-rank factorizations and compression (De Lathauwer, De Moor, and Vandewalle, 2000; Kofidis and Regalia, 2002; Anandkumar et al., 2014), signal processing (De Lathauwer, 1997; Kofidis and Regalia, 2001), quantum geometry (Wei and Goldbart, 2003; Hu, Qi, and Zhang, 2016), medical imaging (Qi, Wang, and Wu, 2008), and data mining (Benson, Gleich, and Leskovec, 2015; Gleich, Lim, and Yu, 2015; Wu, Benson, and Gleich, 2016; Benson, 2019). All real eigenpairs can be computed with a Lasserre-type semidefinite programming hierarchy (Cui, Dai, and Nie, 2014; Nie and Wang, 2014; Nie and Zhang, 2017), but the problem of computing them remains NP-hard (Hillar and Lim, 2013), and the scalability of such methods is limited (see subsection 3.3 for experiments). Thus, we still lack robust and scalable general-purpose methods for computing these eigenvectors.

We introduce two special cases of tensor contractions that will be useful:

1. The *tensor apply* takes a cubic tensor and a vector and produces a vector, akin to Qi’s notation (Qi, 2005):

$$\begin{array}{ll} \text{three-mode tensor} & \mathbf{y} = \underline{\mathbf{T}}\mathbf{x}^2 & y_i = \sum_{j,k} \underline{T}_{i,j,k} x_j x_k \\ \text{\textit{m}-mode tensor} & \mathbf{y} = \underline{\mathbf{T}}\mathbf{x}^{m-1} & y_i = \sum_{i_2, \dots, i_m} \underline{T}_{i, i_2, \dots, i_m} x_{i_2} \cdots x_{i_m}. \end{array}$$

2. The *tensor collapse* takes a cubic tensor and a vector and produces a matrix:

$$\begin{array}{ll} \text{three-mode tensor} & \mathbf{Y} = \underline{\mathbf{T}}[\mathbf{x}] & \mathbf{Y} = \sum_k \underline{T}_{:, :, k} x_k \\ & & Y_{ij} = \sum_k \underline{T}_{i, j, k} x_k \\ \text{\textit{m}-mode tensor} & \mathbf{Y} = \underline{\mathbf{T}}[\mathbf{x}]^{m-2} & \mathbf{Y} = \sum_{i_3, \dots, i_m} \underline{T}_{:, :, i_3, \dots, i_m} x_{i_3} \cdots x_{i_m} \\ & & Y_{ij} = \sum_{i_3, \dots, i_m} \underline{T}_{i, j, i_3, \dots, i_m} x_{i_3} \cdots x_{i_m}. \end{array}$$

For the tensor collapse operator, the “:” symbol signifies taking all entries along that index, so $\underline{T}_{:, :, k}$ is a square matrix. The tensor may not be symmetric, but we are always contracting onto the first mode (tensor apply) or first and second modes (tensor collapse); we assume that $\underline{\mathbf{T}}$ has been permuted in the appropriate manner for the problem at hand. With this notation, the Z -eigenvector problem can be written as

$$(1.2) \quad \underline{\mathbf{T}}\mathbf{x}^{m-1} = \lambda\mathbf{x}, \quad \|\mathbf{x}\|_2 = 1.$$

The crux of our computational method is based on the following observation that relates tensor and matrix eigenvectors.

Observation 1.1. A tensor Z -eigenvector \mathbf{x} of an m -mode tensor must be a *matrix eigenvector* of the collapsed matrix $\underline{\mathbf{T}}[\mathbf{x}]^{m-2}$, i.e.,

$$(1.3) \quad \underline{\mathbf{T}}\mathbf{x}^{m-1} = \lambda\mathbf{x} \iff \underline{\mathbf{T}}[\mathbf{x}]^{m-2}\mathbf{x} = \lambda\mathbf{x}.$$

The catch, of course, is that the matrix itself depends on the tensor eigenvector we want to compute, which we do not know beforehand. Therefore, we still have a nonlinear problem.

2. A dynamical system’s framework for computing Z-eigenvectors. Observation 1.1 provides a new perspective on the tensor Z-eigenvector problem. Specifically, tensor Z-eigenvectors are matrix eigenvectors, just for some unknown matrix. Our computational approach is based on the following continuous-time dynamical system, for reasons that we will make clear in subsection 2.2:

$$(2.1) \quad \frac{d\mathbf{x}}{dt} = \Lambda(\underline{\mathbf{T}}[\mathbf{x}]^{m-2}) - \mathbf{x}.$$

Here, Λ is some fixed map that takes as input a matrix and produces as output some prescribed eigenvector of the matrix with unit norm. For example, on an input \mathbf{M} , Λ could be defined to compute several objects:

1. the eigenvector of \mathbf{M} with k th smallest/largest magnitude eigenvalue;
2. the eigenvector of \mathbf{M} with k th smallest/largest algebraic eigenvalue;
3. the eigenvector of \mathbf{M} closest in distance to a given vector \mathbf{v} .

We resolve the ambiguity in the sign of the eigenvector by picking the sign based on the first element. In the case of multiple eigenvectors sharing an eigenvalue, we propose using the closest eigenvector to \mathbf{x} , although we have not evaluated this technique.

PROPOSITION 2.1. *Let Λ be a prescribed map from a matrix to one of its eigenvectors. Then, if the dynamical system in (2.1) converges to a nonzero solution, it must converge to a tensor Z-eigenvector.*

Proof. If the dynamical system converges, then it converges to a stationary point. Any stationary point has zero derivative, so

$$\begin{aligned} \frac{d\mathbf{x}}{dt} = 0 &\iff \Lambda(\underline{\mathbf{T}}[\mathbf{x}]^{m-2}) = \mathbf{x} \iff \underline{\mathbf{T}}[\mathbf{x}]^{m-2}\mathbf{x} = \lambda\mathbf{x} \text{ for some } \lambda \text{ that depends on } \Lambda \\ &\iff \underline{\mathbf{T}}\mathbf{x}^{m-1} = \lambda\mathbf{x}. \quad \square \end{aligned}$$

One must be a bit careful with the input and output values of Λ . If $\underline{\mathbf{T}}$ is not symmetric, then $\underline{\mathbf{T}}[\mathbf{x}]^{m-2}$ might not be diagonalizable, and we may have to deal with complex eigenvalues. To keep the dynamical system real-valued, one could always modify the map Λ to output the real part. However, the tensor need not be symmetric (nor $\underline{\mathbf{T}}[\mathbf{x}]^{m-2}$ normal for all \mathbf{x}) for the dynamical system to maintain real values. In fact, our motivation for this dynamical system comes from a tensor that is not necessarily symmetric, which we will discuss in subsection 2.2.

Proposition 2.1 leads to a broad framework for computing Z-eigenvectors:

1. choose a map Λ and a numerical integration scheme;
2. numerically integrate (2.1).

Different choices of Λ may converge to different Z-eigenvectors, and different numerical integration schemes may lead to different convergence properties. Figure 1 shows a concrete example, where Λ picks the eigenvector corresponding to eigenvalue with largest magnitude real part, along with the forward Euler numerical integration scheme.

The dynamical system in (2.1) has no dependence on the time t . Thus, the system might be a good candidate for an explicit solution; however, this would require integrating the map Λ , for which an explicit solution is unclear in general. Thus, we focus on numerical integration.

```

1  using LinearAlgebra
2
3  function tensor_apply(T::Array{Float64,3}, x::Vector{Float64})
4      n = length(x)
5      y = zeros(Float64, n)
6      for k in 1:n; y += T[:, :, k] * x * x[k]; end
7      return y
8  end
9
10 function tensor_collapse(T::Array{Float64,3}, x::Vector{Float64})
11     n = length(x)
12     Y = zeros(Float64, n, n)
13     for k in 1:n; Y += T[:, :, k] * x[k]; end
14     return Y
15 end
16
17 function dynsys_forw_eul(T::Array{Float64,3}, h::Float64, niter::Int64)
18     function dx_dt(u::Vector{Float64}) # Derivative
19         F = eigen(tensor_collapse(T, u))
20         ind = sortperm(abs.(real(F.values)))[1]
21         v = F.vectors[:, ind]
22         return sign(v[1]) * v - u # sign consistency
23     end
24     x = normalize(ones(Float64, size(T, 1)), 1) # starting point
25     eval_hist = [x' * tensor_apply(T, x)]
26     for _ = 1:niter
27         x += h * dx_dt(x) # forward Euler
28         push!(eval_hist, x' * tensor_apply(T, x)) # Rayleigh quotient
29     end
30     return (x, eval_hist) # guess at evec and history of evals
31 end

```

FIG. 1. *Julia implementation of the dynamical system for a 3-mode tensor with a map Λ that picks the largest magnitude real eigenvalue and numerical integration with the forward Euler method. Code snippet is available at <https://gist.github.com/arbenson/f28d1b2de9aa72882735e1be24d05a7f>. A more expansive code is available at <https://github.com/arbenson/TZE-dynsys>.*

2.1. Forward Euler and diagonal tensors. As an illustrative example, we consider the special case of using the forward Euler numerical integration scheme for computing the tensor eigenvalues of an n -dimensional, m -mode diagonal tensor $\underline{\mathbf{T}}$. Without loss of generality, assume that the diagonal entries of $\underline{\mathbf{T}}$ are decreasing in order so that $\underline{T}_{i,\dots,i} < \underline{T}_{j,\dots,j}$ if $i > j$. This tensor has at least n Z -eigenpairs: $(\mathbf{e}_i, \underline{T}_{i,\dots,i})$ for $1 \leq i \leq n$, where \mathbf{e}_i is the i th standard basis vector. Suppose that we want to compute the i th eigenvector and set Λ to select the unit-norm eigenvector closest to \mathbf{e}_i in angle. Since $\underline{\mathbf{T}}[\mathbf{x}]^{m-2}$ is diagonal, its eigenvectors are the standard basis vectors, and $\Lambda(\underline{\mathbf{T}}[\mathbf{x}]^{m-2}) = \mathbf{e}_i$. Let $\mathbf{r}_k = \mathbf{x}_k - \mathbf{e}_i$ be the residual at the k th iteration. If the step size is h , then

$$\begin{aligned} \|\mathbf{r}_{k+1}\| &= \|\mathbf{x}_{k+1} - \mathbf{e}_i\| = \|\mathbf{x}_k + h(\mathbf{e}_i - \mathbf{x}_k) - \mathbf{e}_i\| \\ &= (1-h)\|\mathbf{x}_k - \mathbf{e}_i\| = (1-h)\|\mathbf{r}_k\| = (1-h)^k\|\mathbf{r}_0\|. \end{aligned}$$

Thus, the forward Euler scheme converges if $h \leq 1$ and converges in one step if $h = 1$. Figure 2 (left) illustrates the dynamics for an example tensor $\underline{\mathbf{T}}$ ($n = 3$, $m = 3$) with $\underline{T}_{1,1,1} = 5$, $\underline{T}_{2,2,2} = 2$, and $\underline{T}_{1,1,1} = 1$. In this case, the entire surface of the three-dimensional sphere is a basin of attraction for this eigenvector, which is consistent with our convergence analysis. In fact, for this specific case, we have the closed-form solution

$$\mathbf{x}(t) = e^{-t}[\mathbf{x}(0) - \mathbf{e}_i] + \mathbf{e}_i,$$

and we have exponential convergence to a solution, consistent with Figure 2 (left).




FIG. 2. Vector field of dynamical systems from (2.1) of points on the surface of the three-dimensional unit sphere for a three-dimensional, third-order diagonal tensor with entries 5, 2, and 1 along the diagonal. Fixed points are marked with a blue “X.” Sample trajectories of the dynamical system are illustrated with darker red points corresponding to earlier time and lighter yellow corresponding to later time (computed with forward Euler with step size 0.01). (Left) When the map Λ is the eigenvector closest in angle to the third standard basis vector, the entire surface is a basin of attraction for this vector, which is a Z -eigenvector of the tensor. The trajectories head toward this eigenvector. (Right) When the map Λ is the eigenvector corresponding to the smallest algebraic eigenvalue, the Z -eigenvector corresponding to the third standard basis vector has no basin of attraction. The map is undefined at the attracting eigenvector where $(x_1, x_2) \approx (0.18, 0.44)$ because the eigenspace corresponding to the smallest algebraic eigenvalue has dimension greater than one.

However, our analysis relies on the particular choice of Λ . Suppose instead that we choose Λ to select the eigenvector corresponding to the smallest algebraic eigenvalue, and we are trying to compute the eigenvector \mathbf{e}_3 of a $3 \times 3 \times 3$ diagonal tensor $\underline{\mathbf{T}}$ with strictly decreasing diagonal entries. Moreover, suppose we have a starting iterate $\mathbf{x}_0 = [\varepsilon/2 \ \varepsilon/2 \ 1 - \varepsilon]^T$, which is close to the Z -eigenvector \mathbf{e}_3 . Then

$$\Lambda(\underline{\mathbf{T}}[\mathbf{x}_0]) = \Lambda \left(\begin{bmatrix} \frac{\varepsilon}{2} \underline{\mathbf{T}}_{1,1,1} & 0 & 0 \\ 0 & \frac{\varepsilon}{2} \underline{\mathbf{T}}_{2,2,2} & 0 \\ 0 & 0 & (1 - \varepsilon) \underline{\mathbf{T}}_{3,3,3} \end{bmatrix} \right) = \mathbf{e}_2$$

if ε is sufficiently small. Forward Euler integration with step size h gives the next iterate $\mathbf{x}_1 = \mathbf{x}_0 + h(\Lambda(\underline{\mathbf{T}}[\mathbf{x}_0]) - \mathbf{x}_0) = [(1 - h)\varepsilon \ (1 - h)\varepsilon + h \ 1 - \varepsilon - h]^T$, which is further away from the Z -eigenvector \mathbf{e}_3 than \mathbf{x}_0 . Thus, there is no basin of attraction for the Z -eigenvector \mathbf{e}_3 with this particular choice of map Λ .

It turns out that this is a case where the dynamical system does not converge. The system is ill-defined for some points, and, moreover, these points are attractors (specifically, the eigenvector $(x_1, x_2, x_3) \approx (0.18, 0.44, 0.88)$ in the example above; see Figure 2, right). In general, for some time, the dynamical system will evolve in the direction of \mathbf{e}_i , where $\underline{\mathbf{T}}_{i,i,i}x_i < \min_j \underline{\mathbf{T}}_{j,j,j}x_j$ for $i \neq j$, $i, j \in \{1, 2, 3\}$. Along this direction, the i th coordinate of the vector increases until $\underline{\mathbf{T}}_{i,i,i}x_i = \underline{\mathbf{T}}_{j,j,j}x_j$ for some $j \neq i$. At this point, the map is ill-defined since the eigenspace corresponding to the smallest eigenvalue of $\underline{\mathbf{T}}[\mathbf{x}]$ has dimension at least two. Since the diagonal tensor entries are distinct by assumption, this is not a fixed point. We can disambiguate the map at these ambiguous points. However, any way of doing so besides artificially mapping the vector to a Z -eigenvector of $\underline{\mathbf{T}}$ would result in immediate attraction back to one of these ambiguous points.

2.2. Spacey random walks motivation for the dynamical system. The motivation for the dynamical system comes from our previous analysis of a stochas-

tic process known as the “spacey random walk” that relates tensor eigenvectors of a particular class of tensors to a stochastic process (Benson, Gleich, and Lim, 2017). Specifically, the class of tensors are irreducible *transition probability tensors* (any irreducible tensor \underline{P} with $\sum_{i_1=1}^n P_{i_1, i_2, \dots, i_m} = 1$ for $1 \leq i_2, \dots, i_m \leq n$). For simplicity, we discuss a three-mode transition probability tensor \underline{P} , where the entries can be interpreted as coming from a second-order Markov chain—the entry $P_{i,j,k}$ is the probability of transitioning to state i given that the last two states were j and k . Due to the theory of Li and Ng (2014), there exists a tensor Z -eigenvector \mathbf{x} with eigenvalue 1 satisfying

$$(2.2) \quad \underline{P}\mathbf{x}^2 = \mathbf{x}, \quad \sum_{i=1}^n x_i = 1, \quad x_i \geq 0.$$

The vector \mathbf{x} is stochastic, but it does *not* represent the stationary distribution of a Markov chain. Instead, we showed that \mathbf{x} is the limiting distribution of a non-Markovian, generalized vertex-reinforced random walk (Benaïm, 1997) that we called the spacey random walk (Benson, Gleich, and Lim, 2017). In the n th step of a spacey random walk, after the process has visited states X_1, \dots, X_n , it *spaces out* and forgets its second last state (that is, the state X_{n-1}). It then invents a new history state Y_n by randomly drawing a past state X_1, \dots, X_n . Finally, it transitions to X_{n+1} via the second-order Markov chain represented by \underline{P} as if its last two states were X_n and Y_n ; i.e., it transitions to X_{n+1} with probability P_{X_{n+1}, X_n, Y_n} . (In contrast, a true second-order Markov chain would transition with probability $P_{X_{n+1}, X_n, X_{n-1}}$.)

Using results from Benaïm (1997), we showed that the long-term dynamics of the spacey random walk for an m -mode transition probability tensor are governed by the following dynamical system (Benson, Gleich, and Lim, 2017):

$$(2.3) \quad \frac{d\mathbf{x}}{dt} = \Pi(\underline{P}[\mathbf{x}]^{m-2}) - \mathbf{x},$$

where Π is a map that takes a column-stochastic transition matrix and maps it to the Perron vector of the matrix. In other words, if the spacey random walk converges, it must converge to an attractor of the dynamical system in (2.3). The dynamical system in (2.3) is a special case of the more general system in (2.1), where the map Λ picks the eigenvector with the largest algebraic eigenvalue (the Perron vector) and the tensor has certain structural properties (it is an irreducible transition probability tensor).

To summarize, our prior work studied a specific case of the general dynamical system in (2.1) to understand the stochastic process behind principal Z -eigenvectors of transition probability tensors. The general dynamical system provides a new framework for computing general tensor eigenvectors—if the dynamical system in (2.1) converges, then it converges to a tensor Z -eigenvector. The dynamical system may not have an attractor (Peterson, 2018), but it usually does in practice (see section 3).

2.3. Relationship to the Perron iteration. Bini, Meini, and Poloni derived a Perron iteration to compute the minimal nonnegative solution of the equation

$$(2.4) \quad \mathbf{x} = \mathbf{a} + \underline{B}\mathbf{x}^2,$$

where \mathbf{a} and \underline{B} are nonnegative and the all-ones vector \mathbf{e} is a (nonminimal) nonnegative solution (Bini, Meini, and Poloni, 2011; Meini and Poloni, 2011). The Perron iteration for computing the minimal nonnegative solution is

$$(2.5) \quad \mathbf{x}_{k+1} = \Pi[\underline{F} + \underline{B}[\mathbf{e}] - \underline{B}[\mathbf{x}_k]],$$

where $\mathbf{F} = \sum_j \underline{\mathbf{B}}_{:,j,:}$, \mathbf{e} is the vector of all ones, and Π maps a nonnegative matrix to its Perron vector with unit 1-norm. Suppose that $\mathbf{x}_0 \geq 0$ and $\mathbf{e}^T \mathbf{x}_0 = 1$. Then every iterate \mathbf{x}_k is stochastic and

$$(2.6) \quad \mathbf{F} = \underline{\mathbf{W}}[\mathbf{x}_k], \underline{W}_{i,j,\ell} = F_{i,j}, \quad \underline{\mathbf{B}}[\mathbf{e}] = \underline{\mathbf{Z}}[\mathbf{x}_k], \underline{Z}_{i,j,\ell} = [\underline{\mathbf{B}}[\mathbf{e}]]_{i,j}.$$

Thus, we can rewrite the Perron iteration in (2.5) as $\mathbf{x}_{k+1} = \Pi(\underline{\mathbf{T}}[\mathbf{x}_k])$, where $\underline{\mathbf{T}} = \underline{\mathbf{W}} + \underline{\mathbf{Z}} - \underline{\mathbf{B}}$. These iterates are equivalent to forward Euler integration of the dynamical system in (2.1) with unit step size and eigenvector map $\Lambda = \Pi$.

Meini and Poloni (2018) derived a similar Perron iteration for the solution to (2.2) for the case of a 3-mode transition probability tensor. The algorithm first computes a minimal substochastic nonnegative vector \mathbf{m} satisfying $\mathbf{m} = \underline{\mathbf{P}}\mathbf{m}^2$ using a Newton method. The Perron iteration for the transition probability tensor is then $\mathbf{x}_{k+1} = \Pi(\underline{\mathbf{T}}[\mathbf{x}_k])$ for $\underline{T}_{i,j,\ell} = \underline{P}_{i,j,\ell}(1 + m_j + m_\ell)$. The iterates are again equivalent to forward Euler integration of (2.1) with unit step size.

2.4. Relationship to the shifted higher-order power method. The shifted higher-order power method (Kolda and Mayo, 2011) can be derived by noticing that

$$(2.7) \quad (1 + \gamma)\lambda\mathbf{x} = \underline{\mathbf{T}}\mathbf{x}^{m-1} + \gamma\lambda\mathbf{x}$$

for any eigenpair. This yields the iteration

$$(2.8) \quad \mathbf{x}_{k+1} = \frac{\frac{1}{1+\gamma} (\underline{\mathbf{T}}\mathbf{x}_k^{m-1} + \gamma\mathbf{x}_k)}{\left\| \frac{1}{1+\gamma} (\underline{\mathbf{T}}\mathbf{x}_k^{m-1} + \gamma\mathbf{x}_k) \right\|_2}$$

for any shift parameter γ (the case where $\gamma = 0$ is just the classical “higher-order power method” (De Lathauwer, De Moor, and Vandewalle, 2000; Regalia and Kofidis, 2000; Kofidis and Regalia, 2002)). Kolda and Mayo showed that when $\underline{\mathbf{T}}$ is symmetric, the iterates in (2.8) converge monotonically to a tensor eigenvector given an appropriate shift γ .

If $\underline{\mathbf{T}} = \underline{\mathbf{P}}$ for some transition probability tensor $\underline{\mathbf{P}}$ and we are interested in the case when $\lambda = 1$ and we normalize via $\|\mathbf{x}\|_1 = 1$, then one can also derive these iterates by the dynamical system

$$(2.9) \quad \frac{d\mathbf{x}}{dt} = \underline{\mathbf{P}}\mathbf{x}^{m-1} - \mathbf{x}$$

(cf. (2.3)). If this dynamical system converges ($d\mathbf{x}/dt = 0$), then $\mathbf{x} = \underline{\mathbf{P}}\mathbf{x}^{m-1}$, and \mathbf{x} is a tensor Z -eigenvector with eigenvalue 1. If we numerically integrate (2.9) using the forward Euler method with step size $h = 1/(1 + \gamma)$ and any starting vector \mathbf{x}_0 satisfying $\mathbf{x}_0 \geq 0$ and $\|\mathbf{x}_0\|_1 = 1$, then the iterates are

$$(2.10) \quad \mathbf{x}_{k+1} = \mathbf{x}_k + \frac{1}{1 + \gamma} (\underline{\mathbf{P}}\mathbf{x}_k^{m-1} - \mathbf{x}_k)$$

$$(2.11) \quad = \frac{1}{1 + \gamma} (\underline{\mathbf{P}}\mathbf{x}_k^{m-1} + \gamma\mathbf{x}_k) = \frac{\frac{1}{1+\gamma} (\underline{\mathbf{P}}\mathbf{x}_k^{m-1} + \gamma\mathbf{x}_k)}{\left\| \frac{1}{1+\gamma} (\underline{\mathbf{P}}\mathbf{x}_k^{m-1} + \gamma\mathbf{x}_k) \right\|_1},$$

which are the same as the shifted higher-order power method iterates in (2.8). The last equality follows from the fact that $\|\mathbf{x}_k\|_1 = 1$ and $\mathbf{x}_k \geq 0$, which is true by a simple induction argument. The base case holds by the initial conditions and

$$(2.12) \quad \|\underline{\mathbf{P}}\mathbf{x}_k^{m-1} + \gamma\mathbf{x}_k\|_1 = \|\underline{\mathbf{P}}\mathbf{x}_k^{m-1}\|_1 + \gamma = 1 + \gamma$$

since $\underline{\mathbf{P}}\mathbf{x}_k^{m-1}$ and \mathbf{x}_k are both stochastic vectors.

With a general tensor $\underline{\mathbf{T}}$, we can either enforce normalization by evolving the dynamical system over, say, a unit sphere, or we can let the vector \mathbf{x} be unnormalized. The latter case gives a more direct connection to the shifted symmetric higher-order power method (SS-HOPM). In this case, any vector \mathbf{x} where $\underline{\mathbf{T}}\mathbf{x}^{m-1} = \|\mathbf{x}\|_2\mathbf{x}$ is a tensor Z -eigenvector. This leads to the following dynamical system:

$$(2.13) \quad \frac{d\mathbf{x}}{dt} = \underline{\mathbf{T}}\mathbf{x}^{m-1} - \|\mathbf{x}\|_2\mathbf{x}.$$

If $d\mathbf{x}/dt = 0$, then $\|\mathbf{x}\|_2\mathbf{x} = \underline{\mathbf{T}}\mathbf{x}^{m-1}$, so \mathbf{x} is a Z -eigenvector of $\underline{\mathbf{T}}$ with eigenvalue $\|\mathbf{x}\|_2$. Now suppose that we numerically integrate the dynamical system in (2.13) by

1. taking a forward Euler step to produce the iterate \mathbf{x}'_{k+1} ;
2. projecting \mathbf{x}'_{k+1} onto the unit sphere by $\mathbf{x}_{k+1} = \mathbf{x}'_{k+1}/\|\mathbf{x}'_{k+1}\|_2$.

If the step size of the forward Euler method is $h = 1/(1 + \gamma)$, then

$$(2.14) \quad \mathbf{x}'_{k+1} = \mathbf{x}_k + \frac{1}{1 + \gamma}(\underline{\mathbf{T}}\mathbf{x}_k^{m-1} - \|\mathbf{x}_k\|_2\mathbf{x}_k) = \frac{1}{1 + \gamma}(\underline{\mathbf{T}}\mathbf{x}_k^{m-1} + \gamma\mathbf{x}_k)$$

since $\|\mathbf{x}_k\|_2 = 1$. The projection onto the unit sphere then gives the shifted higher-order power method iterates in (2.8).

3. Numerical examples. We now show that our method works on two test tensors used in prior work. Subsection 3.1 shows that our approach can compute all eigenvalues of a specific tensor, while the (shifted) higher-order power method cannot compute all of the eigenvalues. Subsection 3.2 verifies that our approach can compute all eigenvalues of a tensor whose eigenvalues were found with semidefinite programming (SDP). Finally, subsection 3.3 shows that our method is faster than SS-HOPM and the SDP method.

3.1. Example 3.6 from Kolda and Mayo (2011). Our first test case is a $3 \times 3 \times 3$ symmetric tensor from Kolda and Mayo (2011, Example 3.6):

$$\underline{\mathbf{T}}_{::,1} = \begin{bmatrix} -0.1281 & 0.0516 & -0.0954 \\ 0.0516 & -0.1958 & -0.179 \\ -0.0954 & -0.179 & -0.2676 \end{bmatrix}, \quad \underline{\mathbf{T}}_{::,2} = \begin{bmatrix} 0.0516 & -0.1958 & -0.179 \\ -0.1958 & 0.3251 & 0.2513 \\ -0.179 & 0.2513 & 0.1773 \end{bmatrix},$$

$$\underline{\mathbf{T}}_{::,3} = \begin{bmatrix} -0.0954 & -0.179 & -0.2676 \\ -0.179 & 0.2513 & 0.1773 \\ -0.2676 & 0.1773 & 0.0338 \end{bmatrix}.$$

The tensor has 7 eigenvalues, which Kolda and Mayo classify as “positive stable,” “negative stable,” or “unstable” (see Figure 3, top), corresponding to positive definiteness, negative definiteness, or indefiniteness of the projected Hessian of the Lagrangian of their optimization function (Kolda and Mayo, 2011). (Since the tensor has an odd number of modes, we only consider eigenvalues up to sign.) Kolda and Mayo showed that their SS-HOPM, a generalization of the symmetric higher-order power method (S-HOPM) (De Lathauwer, De Moor, and Vandewalle, 2000; Regalia and Kofidis, 2000; Kofidis and Regalia, 2002), only converges to eigenvectors of the positive or negative stable eigenvalues. An adaptive version of SS-HOPM has the same shortcoming (Kolda and Mayo, 2014). A recently proposed Newton iteration can converge to eigenpairs where the projected Hessian has eigenvalues bounded away from 0 (Jaffe, Weiss, and Nadler, 2018).

λ	Type	S-HOPM	SS-HOPM	V1	V2	V3	V4	V5
0.0180	Neg. stable	0	18	0	25	0	100	0
0.4306	Neg. stable	38	29	38	0	45	0	0
0.8730	Neg. stable	62	40	62	0	47	0	0
0.0006	Pos. stable	0	13	0	19	8	0	0
0.0018	Unstable	0	0	0	25	0	0	32
0.0033	Unstable	0	0	0	35	0	0	37
0.2294	Unstable	0	0	0	0	0	0	31




FIG. 3. (Top) The 7 eigenvalues of the test tensor from Kolda and Mayo (2011, Example 3.6) and the number of random trials (out of 100) that converge to the eigenvalue for (i) the symmetric higher-order power method, S-HOPM (De Lathauwer, De Moor, and Vandewalle, 2000; Regalia and Kofidis, 2000; Kofidis and Regalia, 2002); (ii) the shifted symmetric higher-order power method, SS-HOPM (Kolda and Mayo, 2011); and (iii) 5 variations of our dynamical systems approach. Variation 1 selects the largest magnitude eigenvalue, Variation 2 selects the smallest magnitude eigenvalue, Variation 3 selects the largest algebraic eigenvalue, Variation 4 selects the smallest algebraic eigenvalue, and Variation 5 selects the second smallest algebraic eigenvalue. Results for (i) and (ii) are from Kolda and Mayo (2011). Our algorithm is the only one that is able to compute all of the eigenvalues, including those which are “unstable,” the eigenvectors to which SS-HOPM and S-HOPM cannot converge (Kolda and Mayo, 2011). (Bottom) Convergence plots for the three unstable eigenvalues from variation 5 of our algorithm in terms of the Rayleigh quotient, where \mathbf{x}_k is the k th iterate.

Of the 7 eigenpairs for the above tensor, 3 are unstable. Our dynamical systems approach can compute all 7 eigenpairs, using 5 variations of the dynamical system:

1. Λ maps \mathbf{M} to the eigenvector with largest magnitude eigenvalue;
2. Λ maps \mathbf{M} to the eigenvector with smallest magnitude eigenvalue;
3. Λ maps \mathbf{M} to the eigenvector with largest algebraic eigenvalue;
4. Λ maps \mathbf{M} to the eigenvector with smallest algebraic eigenvalue;
5. Λ maps \mathbf{M} to the eigenvector with second smallest algebraic eigenvalue.

We used the forward Euler method with step size set to 0.5 in order to compute the eigenvalues. Empirically, convergence is fast, requiring fewer than 10 iterations (Figure 3, bottom row). One can also compute these eigenvectors with semidefinite programming (Cui, Dai, and Nie, 2014), although the scalability of such methods is limited (see subsection 3.3). We next provide numerical results from a tensor in this literature.

3.2. Example 4.11 from Cui, Dai, and Nie (2014). Our second test case is a $5 \times 5 \times 5$ symmetric tensor from Cui, Dai, and Nie (2014, Example 4.11):

$$(3.1) \quad T_{i,j,k} = \frac{(-1)^i}{i} + \frac{(-1)^j}{j} + \frac{(-1)^k}{k}, \quad 1 \leq i, j, k \leq 5.$$

The tensor has 3 eigenvalues (again, the tensor has an odd number of modes, so

λ	SDP	V1	V2	V3	V4	V5
9.9779	✓	94	0	0	100	0
4.2876	✓	6	0	100	0	0
0.0000	✓	0	100	0	0	100




FIG. 4. (Top) The 3 eigenvalues of the test tensor from Cui, Dai, and Nie (2014, Example 4.11) and the number of random trials (out of 100) that converge to the eigenvalue for 5 variations of our dynamical systems approach. Variation 1 selects the largest magnitude eigenvalue, variation 2 selects the smallest magnitude eigenvalue, variation 3 selects the largest algebraic eigenvalue, variation 4 selects the smallest algebraic eigenvalue, and variation 5 selects the second smallest algebraic eigenvalue. Our algorithm is able to compute all of the eigenvalues, which the SDP approach is guaranteed to compute. (Bottom) Convergence plots for the three eigenvalues from different variations of our algorithm in terms of the Rayleigh quotient, where \mathbf{x}_k is the k th iterate.

the eigenvalues are only defined up to sign). We use the same 5 variations of our algorithm to compute the eigenpairs (Figure 4). Again, we are able to compute all of the eigenvalues of the tensor, and convergence is rapid.

3.3. Scalability experiments. Now we compare the performance of our algorithm to both SS-HOPM, as implemented in the Tensor Toolbox for MATLAB.² (Bader and Kolda, 2006; Bader et al., 2017)), and the SDP method, also implemented in MATLAB.³ Our implementation is written in Julia and is also publicly available.⁴

The following order- m , n -dimensional tensor (which is a generalization of (3.1)) serves as the test case for our experiments:

$$(3.2) \quad \underline{\mathcal{T}}_{i_1, \dots, i_m} = \sum_{r=1}^m \frac{(-1)^{i_r}}{i_r}, \quad 1 \leq i_1, \dots, i_m \leq n.$$

With SS-HOPM, we use a tolerance of 10^{-6} , a shift of 1, and $100n$ random initializations. With the SDP method, we use the default parameter settings. With our dynamical systems method, we use a stopping tolerance of 10^{-6} , the forward Euler

²<https://www.tensortoolbox.org>

³<http://www.math.ucsd.edu/~njw/CODES/reigsyntensor/areigstrweb.html>

⁴<https://github.com/arbenson/TZE-dynsys>




FIG. 5. Running time for SS-HOPM (Kolda and Mayo, 2011), the SDP method (Cui, Dai, and Nie, 2014), and our dynamical system method on the test tensor in (3.2). While the SDP method does not scale, the approach has guarantees on the eigenvalues it can compute. Our dynamical systems approach uses the map for the eigenvectors of the k th largest algebraic and magnitude (for $k = 1, \dots, n$), along with the forward Euler numerical integration scheme with step size 0.5. We terminated the SDP method if it did not complete within 24 hours, so only results for dimension up to 10 appear in the plot on the far right.

integration scheme with step size = 0.5, and maps Λ corresponding to k th largest algebraic and magnitude eigenvalue, $k = 1, \dots, n$, each with 50 trials of random initial starting points. With this setup, SS-HOPM and our approach use the same number of randomly initialized trials. We performed all experiments on a 3.1-GHz Intel Core i7 MacBook Pro with 16 GB of RAM.

Figure 5 shows the running times of the algorithms for $m = 3, 4, 5$ and $n = 5, 6, \dots, 15$. The main takeaway is that the SDP method is much slower than the other two methods—this is the price we pay for being able to compute all of the real eigenvalues and dealing with NP-hardness. Our dynamical systems approach is faster than SS-HOPM, which is somewhat surprising since we require an eigendecomposition of an $n \times n$ matrix at each iteration. However, the performance difference is a result of rapid convergence, as observed in Figures 3 and 4. Finally, although the tensors here are relatively small, our method has been used in recent work to compute eigenvectors of tensors of orders 3, 4, and 5 with dimensions in the tens of thousands (Benson, 2019).

4. Stochastics as a guide. Scalable methods for computing tensor eigenvectors remain a challenge. Our new framework for computing Z -eigenvectors offers insights through three observations. First, a tensor Z -eigenvector is a matrix eigenvector of *some* matrix, where the matrix is obtained by applying the tensor collapse operator with the Z -eigenvector itself. Second, for a certain class of tensors where eigenvectors have a stochastic interpretation, the dynamical system in (2.1) is the one that governs the long-term dynamics of the stochastic process. Third, the same type of dynamical system seems to work for more general tensors. This framework can compute tensor eigenvectors that other scalable methods, such as SS-HOPM, cannot.

The dynamical system framework is a flexible setup to create solvers for tensor eigenvector problems, and dynamical systems have also been used in matrix eigenvector problems (Chu, 1984; Golub and Liao, 2006). The difference between SS-HOPM and our proposed framework, for instance, is essentially that SS-HOPM takes a single step of the power method on the matrix $\underline{T}[\mathbf{x}]$, whereas we converge to an eigenvector of $\underline{T}[\mathbf{x}]$. There is a rich space to interpolate between these positions. Straightforward ideas include low-degree polynomial filters that target specific eigenvectors.

Indeed, one major challenge is knowing what map Λ to choose—different choices lead to different eigenvectors, and there is no immediate relationship between them for general tensors. A second class of open questions relates to convergence theory. At the




FIG. 6. Vector field of dynamical systems from (2.1), along with trajectories of the dynamical system, akin to Figure 2. However, here we normalize the vector to have unit 2-norm after each step of numerical integration of the dynamical system. This helps converge to the attractor, where the smallest algebraic eigenvalue map is undefined. A better understanding of why this helps is an avenue for future research.

moment, we have demonstrated that there are both convergent and nonconvergent cases (Figure 2). We can alleviate this problem by normalizing the iterates of the integration scheme to have unit 2-norm after each iteration (Figure 6); however, we do not have a good theory for why this works. Finally, our method is not immediately applicable to H -eigenvectors because Observation 1.1 no longer holds. Adapting our methodology to this class of eigenvectors is an area for future research.

Our framework came from relating tensor eigenvectors to stochastic processes. This is quite different from the core ideas in the tensor literature, which are firmly rooted in algebraic generalizations. We hope that these results encourage further development of the relationships between stochastics and tensor problems.

Acknowledgments. We thank Brad Nelson for providing valuable feedback.

REFERENCES

- A. ANANDKUMAR, R. GE, D. HSU, S. M. KAKADE, AND M. TELGARSKY (2014), *Tensor decompositions for learning latent variable models*, J. Mach. Learn. Res., 15, pp. 2773–2832, <http://dl.acm.org/citation.cfm?id=2627435.2697055>.
- B. W. BADER AND T. G. KOLDA (2006), *Algorithm 862: MATLAB tensor classes for fast algorithm prototyping*, ACM Trans. Math. Software, 32, pp. 635–653, <https://doi.org/10.1145/1186785.1186794>.
- B. W. BADER, T. G. KOLDA, ET AL. (2017), *MATLAB Tensor Toolbox Version 3.0-dev.*, <https://www.tensortoolbox.org>.
- M. BENAÏM (1997), *Vertex-reinforced random walks and a conjecture of Pemantle*, Ann. Probab., 25, pp. 361–392, <http://projecteuclid.org/euclid.aop/1024404292>.
- A. R. BENSON (2019), *Three hypergraph eigenvector centralities*, SIAM J. Math. Data Sci., 1, pp. 293–312, <https://doi.org/10.1137/18m1203031>.
- A. R. BENSON, D. F. GLEICH, AND J. LESKOVEC (2015), *Tensor spectral clustering for partitioning higher-order network structures*, in Proceedings of the 2015 SIAM International Conference on Data Mining, Vancouver, British Columbia, Canada, pp. 118–126, <https://doi.org/10.1137/1.9781611974010.14>.
- A. R. BENSON, D. F. GLEICH, AND L.-H. LIM (2017), *The spacey random walk: A stochastic process for higher-order data*, SIAM Rev., 59, pp. 321–345, <https://doi.org/10.1137/16M1074023>.

- D. A. BINI, B. MEINI, AND F. POLONI (2011), *On the solution of a quadratic vector equation arising in Markovian binary trees*, Numer. Linear Algebra Appl., 18, pp. 981–991, <https://doi.org/10.1002/nla.809>.
- D. CARTWRIGHT AND B. STURMFELS (2013), *The number of eigenvalues of a tensor*, Linear Algebra Appl., 438, pp. 942–952, <https://doi.org/10.1016/j.laa.2011.05.040>.
- M. T. CHU (1984), *The generalized Toda flow, the QR algorithm and the center manifold theory*, SIAM J. Algebraic Discrete Methods, 5, pp. 187–201, <https://doi.org/10.1137/0605020>.
- P. COMON, G. GOLUB, L.-H. LIM, AND B. MOURRAIN (2008), *Symmetric tensors and symmetric tensor rank*, SIAM J. Matrix Anal. Appl., 30, pp. 1254–1279, <https://doi.org/10.1137/060661569>.
- C.-F. CUI, Y.-H. DAI, AND J. NIE (2014), *All real eigenvalues of symmetric tensors*, SIAM J. Matrix Anal. Appl., 35, pp. 1582–1601, <https://doi.org/10.1137/140962292>.
- D. F. GLEICH, L.-H. LIM, AND Y. YU (2015), *Multilinear PageRank*, SIAM J. Matrix Anal. Appl., 36, pp. 1507–1541, <https://doi.org/10.1137/140985160>.
- G. H. GOLUB AND L.-Z. LIAO (2006), *Continuous methods for extreme and interior eigenvalue problems*, Linear Algebra Appl., 415, pp. 31–51, <https://doi.org/10.1016/j.laa.2005.01.009>.
- C. J. HILLAR AND L.-H. LIM (2013), *Most tensor problems are NP-hard*, J. ACM, 60, pp. 1–39, <https://doi.org/10.1145/2512329>.
- S. HU, L. QI, AND G. ZHANG (2016), *Computing the geometric measure of entanglement of multipartite pure states by means of non-negative tensors*, Physical Rev. A, 93, <https://doi.org/10.1103/physreva.93.012304>.
- A. JAFFE, R. WEISS, AND B. NADLER (2018), *Newton correction methods for computing real eigenpairs of symmetric tensors*, SIAM J. Matrix Anal. Appl., 39, pp. 1071–1094, <https://doi.org/10.1137/17m1133312>.
- E. KOFIDIS AND P. A. REGALIA (2001), *Tensor approximation and signal processing applications*, Contemp. Math., 280, pp. 103–134.
- E. KOFIDIS AND P. A. REGALIA (2002), *On the best rank-1 approximation of higher-order supersymmetric tensors*, SIAM J. Matrix Anal. Appl., 23, pp. 863–884, <https://doi.org/10.1137/s0895479801387413>.
- T. G. KOLDA AND B. W. BADER (2009), *Tensor decompositions and applications*, SIAM Rev., 51, pp. 455–500, <https://doi.org/10.1137/07070111X>.
- T. G. KOLDA AND J. R. MAYO (2011), *Shifted power method for computing tensor eigenpairs*, SIAM J. Matrix Anal. Appl., 32, pp. 1095–1124, <https://doi.org/10.1137/100801482>.
- T. G. KOLDA AND J. R. MAYO (2014), *An adaptive shifted power method for computing generalized tensor eigenpairs*, SIAM J. Matrix Anal. Appl., 35, pp. 1563–1581, <https://doi.org/10.1137/140951758>.
- L. DE LATHAUWER (1997), *Signal Processing Based on Multilinear Algebra*, Ph.D. thesis, Katholieke Universiteit Leuven, Leuven, Belgium.
- L. DE LATHAUWER, B. DE MOOR, AND J. VANDEWALLE (2000), *On the best rank-1 and rank- (r_1, r_2, \dots, r_N) approximation of higher-order tensors*, SIAM J. Matrix Anal. Appl., 21, pp. 1324–1342, <https://doi.org/10.1137/s0895479898346995>.
- W. LI AND M. K. NG (2014), *On the limiting probability distribution of a transition probability tensor*, Linear Multilinear Algebra, 62, pp. 362–385, <https://doi.org/10.1080/03081087.2013.777436>.
- L.-H. LIM (2005), *Singular values and eigenvalues of tensors: A variational approach*, in 1st IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing, Puerto Vallarta, Mexico, pp. 129–132.
- L.-H. LIM (2013), *Tensors and hypermatrices*, in Handbook of Linear Algebra, 2nd ed., L. Hogben, ed., Chapman and Hall/CRC Press, Boca Raton, FL, pp. 231–260.
- B. MEINI AND F. POLONI (2011), *A Perron iteration for the solution of a quadratic vector equation arising in Markovian binary trees*, SIAM J. Matrix Anal. Appl., 32, pp. 248–261, <https://doi.org/10.1137/100796765>.
- B. MEINI AND F. POLONI (2018), *Perron-based algorithms for the multilinear PageRank*, Numer. Linear Algebra Appl., 25, e2177, <https://doi.org/10.1002/nla.2177>.
- J. NIE AND L. WANG (2014), *Semidefinite relaxations for best rank-1 tensor approximations*, SIAM J. Matrix Anal. Appl., 35, pp. 1155–1179, <https://doi.org/10.1137/130935112>.
- J. NIE AND X. ZHANG (2017), *Real eigenvalues of nonsymmetric tensors*, Comput. Optim. Appl., 70, pp. 1–32, <https://doi.org/10.1007/s10589-017-9973-y>.
- J. PETERSON (2018), *Personal communication*.
- L. QI (2005), *Eigenvalues of a real supersymmetric tensor*, J. Symbolic Comput., 40, pp. 1302–1324, <https://doi.org/10.1016/j.jsc.2005.05.007>.

- L. QI, Y. WANG, AND E. X. WU (2008), *D*-eigenvalues of diffusion kurtosis tensors, *J. Comput. Appl. Math.*, 221, pp. 150–157, <https://doi.org/10.1016/j.cam.2007.10.012>.
- P. A. REGALIA AND E. KOFIDIS (2000), *The higher-order power method revisited: Convergence proofs and effective initialization*, in Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, IEEE, Istanbul, Turkey, <https://doi.org/10.1109/icassp.2000.861047>.
- T.-C. WEI AND P. M. GOLDBART (2003), *Geometric measure of entanglement and applications to bipartite and multipartite quantum states*, *Phys. Rev. A*, 68, <https://doi.org/10.1103/physreva.68.042307>.
- T. WU, A. BENSON, AND D. F. GLEICH (2016), *General tensor spectral co-clustering for higher-order data*, in Advances in Neural Information Processing Systems, pp. 2559–2567, Barcelona, Spain, <http://papers.nips.cc/paper/6376-general-tensor-spectral-co-clustering-for-higher-order-data>.