OPTIMIZATION VIA SEPARATED REPRESENTATIONS AND THE CANONICAL TENSOR DECOMPOSITION

MATTHEW J REYNOLDS*, GREGORY BEYLKIN†, AND ALIREZA DOOSTAN*

ABSTRACT. We introduce a new, quadratically convergent algorithm for finding maximum absolute value entries of tensors represented in the canonical format. The computational complexity of the algorithm is linear in the dimension of the tensor. We show how to use this algorithm to find global maxima of non-convex multivariate functions in separated form. We demonstrate the performance of the new algorithms on several examples.

1. INTRODUCTION

Finding global extrema of a multivariate function is an ubiquitous task with many approaches developed to address this problem (see e.g. [19]). Unfortunately, no existing optimization method can guarantee that the results of optimization are true global extrema unless restrictive assumptions are placed on the function. Assumptions on smoothness of the function do not help since it is easy to construct an example of a function with numerous local extrema "hiding" the location of the true one. While convexity assumptions are helpful for finding global maxima, in practical applications there are many non-convex functions. For non-convex functions various randomized search strategies have been suggested and used but none can assure that the results are true global extrema (see, e.g., [27, 19]).

We propose a new approach to the problem of finding global extrema of a multivariate function under the assumption that the function has certain structure, namely, a separated representation with a reasonably small separation rank. While this assumption limits the complexity of the function, i.e. number of independent degrees of freedom in its representation, there is no restriction on its convexity. Furthermore, a large number of functions that do not appear to have an obvious separated representation do in fact possess one, as was observed in [5, 6, 4]. In particular, separated representations have been used recently to address curse of dimensionality issues that occur when solving PDEs with random data in the context of uncertainty quantification (see, e.g., [11, 13, 14, 15, 17, 22, 23, 24]).

We present a surprisingly simple algorithm that uses canonical tensor decompositions (CTDs) for the optimization process. A canonical decomposition of a tensor $\mathbf{U} \in \mathbb{R}^{M_1 \times M_2 \times \cdots \times M_d}$ is of the form,

$$\mathbf{U} = U\left(i_1 \ldots i_d\right) = \sum_{l=1}^{r} s_l u_{i_1}^l u_{i_2}^l \cdots u_{i_d}^l,$$

1

where $u_{i_j}^l$, $j = 1, \ldots, d$ are entries of vectors in $\mathbb{R}^{M_d}$ and r is the separation rank. Using the CTD format to find entries of a tensor with the maximum absolute value was first explored in [16] via an analogue of the matrix power method. Unfortunately, the algorithm in [16] has the same weakness as the matrix power method: its convergence can be very slow, thus limiting its applicability. Instead of using the power method, we introduce a quadratically convergent algorithm based on straightforward sequential squaring of tensor entries in order to find entries with the maximum absolute value.

For a tensor with even a moderate number of dimensions finding the maximum absolute value of the entries appears to be an intractable problem. A brute-force search of a d-directional tensor requires $\mathcal{O}\left(N^d\right)$ operations, an impractical computational cost. However, for tensors in CTD format the situation is different since the curse of dimensionality can be avoided.

Our approach to finding the entries with the maximum absolute value (abbreviated as "the maximum entry" where it does not cause confusion) is as follows: as in [16], we observe that replacing the maximum entries with 1's and the rest of the entries with zeros results in a tensor that has a low separation rank representation (at least in the case where the number of such extrema is reasonably small). To construct such a tensor, we simply square the entries via the Hadamard product, normalize the result and repeat these steps a sufficient number of times. The resulting algorithm is quadratically convergent as we raise the tensor entries to the power 2^n after n steps. Since the nominal separation rank of the tensor grows after each squaring, we use CTD rank reduction algorithms to keep the separation rank at a reasonable level. The accuracy threshold used in rank reduction algorithms limits the overall algorithm to finding the global maximum within the introduced accuracy limitation. However, in many practical applications, we can find the true extrema as this accuracy threshold is user-controlled. These operations are performed with a cost that depends linearly on the dimension d.

The paper is organized as follows. In Section 2 we briefly review separated representations of functions and demonstrate how they give rise to CTDs. In Section 3 we introduce the quadratically convergent algorithm for finding the maximum entry of a tensor in CTD format and discuss the selection of tensor norm for this algorithm. In Section 4 we use the new algorithm to find the maximum absolute value of continuous functions represented in separated form. In Section 5 we test both the CTD and separated representation optimization algorithms on numerical examples, including a parameterized linear elliptic PDE. We provide our conclusions and a discussion of these algorithms in Section 6.

2. Background

2.1. **Separated representation of functions and the canonical tensor decomposition.** Separated representations of multivariate functions and operators for computing in higher dimensions were introduced in [5, 6]. The separated representation of a function $u(x_1, x_2, \ldots, x_d)$ is a natural extension of separation of variables as we seek an approximation

$$(2.1) \qquad u(x_1, \ldots, x_d) = \sum_{l=1}^{r} s_l u_1^{(l)}(x_1) \cdots u_d^{(l)}(x_d) + \mathcal{O}\left(\epsilon\right),$$

where $s_l > 0$ are referred to as s-values. In this approximation the functions $u_j^{(l)}(x_j)$, $j = 1, \ldots, d$ are not fixed in advance but are optimized in order to achieve the accuracy goal with (ideally) a minimal *separation rank* r. In (2.1) we set $x_j \in \mathbb{R}$ while noting that in general the variables x_j may be complex-valued or low dimensional vectors. Importantly, a separated representation is not a projection onto a subspace, but rather a nonlinear method to track a function in a high-dimensional space using a small number of parameters. We note that the separation rank indicates just the nominal number of terms in the representation and is not necessarily minimal.

Any discretization of the univariate functions $u_j^{(l)}(x_j)$ in (2.1) with $u_{i_j}^{(l)} = u_j^{(l)}(x_{i_j})$, $i_j = 1, \ldots, M_j$ and $j = 1, \ldots, d$, leads to a d-dimensional tensor $\mathbf{U} \in \mathbb{R}^{M_1 \times \cdots \times M_d}$, a canonical tensor decomposition (CTD) of separation rank r_u,

$$(2.2) \qquad \mathbf{U} = U(i_1, \ldots, i_d) = \sum_{l=1}^{r_u} s_l^u \prod_{j=1}^{d} u_{i_j}^{(l)},$$

where the s-values s_l^u are chosen so that each vector $\mathbf{u}_j^{(l)} = \left\{ u_{i_j}^{(l)} \right\}_{i_j=1}^{M_j}$ has unit Euclidean norm $\|\mathbf{u}_j^{(l)}\|_2 = 1$ for all j, l. The CTD has become one of the key tools in the emerging field of numerical multilinear algebra (see, e.g., the reviews [9, 28, 20]). Given CTDs of two tensors $\mathbf{U}$ and $\mathbf{V}$ of separation ranks r_u and r_v, their inner product is defined as

$$(2.3) \qquad \langle \mathbf{U}, \mathbf{V} \rangle = \sum_{l=1}^{r_u} \sum_{l'=1}^{r_v} s_l^u s_{l'}^v \prod_{j=1}^{d} \left\langle \mathbf{u}_j^{(l)}, \mathbf{v}_j^{(l')} \right\rangle,$$

where the inner product $\langle \cdot, \cdot \rangle$ operating on vectors is the standard vector dot product. The Frobenius norm is then defined as $\|\mathbf{U}\|_F = \sqrt{\langle \mathbf{U}, \mathbf{U} \rangle}$. Central to our optimization algorithm is the Hadamard, or point-wise, product of two tensors represented in CTD format. We define this product as,

$$\mathbf{U} * \mathbf{V} = (U * V)(i_1, \ldots, i_d) = \sum_{l=1}^{r_u} \sum_{l'=1}^{r_v} s_l^u s_{l'}^v \prod_{j=1}^{d} \left(u_{i_j}^{(l)} \cdot v_{i_j}^{(l')} \right).$$

Common operations involving CTDs, e.g. sum or multiplication, lead to CTDs with separation ranks that may be larger than necessary for a user-specified accuracy. To keep computations with CTDs manageable, it is crucial to reduce the separation rank. The separation rank reduction operation for CTDs, here referred to as τ_ϵ, is defined as follows: given a tensor $\mathbf{U}$ in CTD format with separation rank r_u,

$$\mathbf{U} = U(i_1, \ldots, i_d) = \sum_{l=1}^{r_u} s_l^u \prod_{j=1}^{d} u_{i_j}^{(l)},$$

and user-supplied error ϵ, find a representation

$$\mathbf{V} = V(i_1, \ldots, i_d) = \sum_{l'=1}^{r_v} s_{l'}^v \prod_{j=1}^{d} v_{i_j}^{(l')},$$

with lower separation rank, $r_v < r_u$, such that $\|\mathbf{U} - \mathbf{V}\| < \epsilon \|\mathbf{U}\|$. The workhorse algorithm for the separation rank reduction problem is Alternating Least Squares

(ALS) which was introduced originally for data fitting as the PARAFAC (PARAllel FACtor) [18] and the CANDECOMP [10] models. ALS has been used extensively in data analysis of (mostly) three-way arrays (see e.g. the reviews [9, 28, 20] and references therein). For the experiments in this paper we use exclusively the randomized interpolative CTD tensor decomposition, CTD-ID, described in [8], as a faster alternative to ALS.

2.2. Power method for finding the maximum entry of a tensor.

The method for finding the maximum entry of a tensor in CTD format in [16] relies on a variation of the matrix power method adapted to tensors. To see how finding entries with maximum absolute value can be cast as an eigenvalue problem, we define the low rank tensor $\mathbf{Y}$ with 1s at the locations of such entries in $\mathbf{U}$, and set $\lambda = \max_{i_1,\ldots,i_d} |U(i_1, i_2, \ldots, i_d)|$. The Hadamard product,

$$(2.4) \qquad \mathbf{U} * \mathbf{Y} = \lambda \mathbf{Y},$$

reveals the underlying eigenvalue problem. We describe the algorithm of [16] as Algorithm 1 and use it for performance comparisons.

Algorithm 1

To find the entries with maximum absolute value of a tensor $\mathbf{U} \in \mathbb{R}^{N_1 \times N_2 \times \cdots \times N_d}$ in CTD format, given a tolerance $\epsilon > 0$ for the reduction algorithm τ_ϵ, we define

$\mathbf{Y}_0 = \left(\prod_{i=1}^d \frac{1}{N_i} \right) (\mathbf{1}_1 \circ \mathbf{1}_2 \circ \cdots \circ \mathbf{1}_d)$, where $\mathbf{1}_i = (1, 1, \ldots, 1)^T \in \mathbb{R}^{N_i}$,
and iterate (up to $k_{\max}$ times) as follows:

for $k = 1, 2, \ldots, k_{\max}$ **do**
    (1) $\mathbf{Q}_k = \mathbf{U} * \mathbf{Y}_{k-1}, \ \lambda_k = \langle \mathbf{Y}_{k-1}, \mathbf{Q}_k \rangle, \ \mathbf{Z}_k = \mathbf{Q}_k / \|\mathbf{Q}_k\|_F$
    (2) $\mathbf{Y}_k = \tau_\epsilon (\mathbf{Z}_k)$
end for

The resulting tensor $\mathbf{Y}_k$ ideally has a low separation rank with zero entries except for a few non-zeros indicating the location of the extrema. In the case of a single extremum the tensor $\mathbf{Y}_k$ has separation rank 1.

3. A QUADRATICALLY CONVERGENT ALGORITHM FOR FINDING MAXIMUM ENTRIES OF TENSORS

In order to construct $\mathbf{Y}$ from (2.4), we simply square all entries of the tensor, normalize the result, and repeat these steps a number of times. Fortunately, the CTD format allows us to square tensor entries using the Hadamard product. These operations increase the separation rank, so the squaring step is (usually) followed by the application of a rank reduction algorithm such as ALS, CTD-ID, or their combination.

Algorithm 2

To find the entries with maximum absolute value of a tensor $\mathbf{U} \in \mathbb{R}^{N_1 \times N_2 \times \cdots \times N_d}$ in CTD format, given a reduction tolerance $\epsilon > 0$ for the reduction algorithm τ_ϵ, we define

$\mathbf{Y}_0 = \mathbf{U} / \|\mathbf{U}\|_F$, and iterate (up to $k_{\max}$ times) as follows:

for $k = 1, 2, \ldots, k_{\max}$ **do**
    (1) $\mathbf{Q}_k = \mathbf{Y}_{k-1} * \mathbf{Y}_{k-1}$
    (2) $\mathbf{Y}_k = \tau_\epsilon(\mathbf{Q}_k)$, $\mathbf{Y}_k = \mathbf{Y}_k / \|\mathbf{Y}_k\|_F$
end for

If we identify the largest (by absolute value) entry as $a = |F(j_1, j_2, \ldots, j_d)|$ and the next largest as b, the ratio $b/a < 1$ decreases rapidly,

$$\left(\frac{b}{a}\right)^{2^j} \leq \epsilon$$

and the total number of iterations, j, to reach ϵ is logarithmic,

$$j \leq \log_2\left(\frac{\log(\epsilon)}{\log(b) - \log(a)}\right).$$

Since this ratio is expected to be reasonably small for most entries, the tensor will rapidly become sparse after only a small number of iterations j. The quadratic rate of convergence of Algorithm 2 is the key difference with the algorithm in [16] whose convergence rate is at best linear.

3.1. **Termination conditions for Algorithm 2.** We have explored three options to terminate the iteration in Algorithm 2. First, the simplest, is to fix the number of iterations in advance provided some knowledge of the gap between the entry with the maximum absolute value and the rest of the entries is available. The second option is to define $\lambda_k = \langle \mathbf{Y}_k, \mathbf{U} \rangle$ and terminate the iteration once the rate of its decrease becomes small, i.e., $(\lambda_k - \lambda_{k-1}) / \lambda_{k-1} < \delta$, where $\delta > 0$ is a user-supplied parameter. The third option is to observe the separation rank of $\mathbf{Y}_k$ and terminate the iteration when it becomes smaller than a fixed, user-selected value. In Section 5 we indicate which termination condition is used in our numerical examples.

3.2. **Selection of tensor norm.** Working with tensors, we have a rather limited selection of norms that are computable. The usual norm chosen for work with tensors is the Frobenius norm which may not be appropriate in all situations where the goal is to find the maximum absolute value entries. Unfortunately, the Frobenius norm is only weakly sensitive to changes of individual entries. The alternative s-norm (see [8, Section 4]), i.e., the largest s-value of the rank one separated approximation to the tensor in CTD format, is better in some situations. In particular, it allows the user to lower the tolerance to below single precision (within a double precision environment). The improved control over truncation accuracy makes this norm more sensitive to individual entries of the tensor. We note that while it is straightforward to compute the s-value of a rank one separated approximation via a convergent iteration, theoretically it is possible that this computed number is not

the largest s-value. However, we have not encountered such a situation using the s-norm in practical problems.

3.3. Numerical demonstration of convergence.

We construct an experiment using random tensors, with low separation rank ($r = 4$ in what follows), in dimension $d = 6$ with $M = 32$ samples in each direction. The input CTD is formed by adding a rank three CTD and a rank one CTD together. The rank three CTD is formed using random factors ($\mathbf{u}_j^{(l)}$ in (2.2)) whose samples were drawn from a uniform distribution on $[.9, 1]$ (the s-values were set to 1). The rank one CTD has all zeros except for a spike added in at a random location. The size of this spike is chosen to make the maximum entry have a magnitude of 3.5. For the experiments in this section we set the reduction tolerance of τ_ϵ in Algorithm 2 to $\epsilon = 10^{-6}$ and use the Frobenius norm.

We observe that as the iteration proceeds, the maximum entry of one of the CTD terms eventually becomes dominant relative to the rest. The iteration continues until the rest of the terms are small enough to be eliminated by the separation rank reduction algorithm τ_ϵ. This is illustrated in Figure 3.1 where we display the maximum entry from all rank one terms in the CTD for each iteration. Before the first iteration (iteration 0 in Figure 3.1), the gaps between the maxima of the rank one terms are not large. As we iterate, the maximum entry of one of the rank one terms rapidly separates itself from the others and, by iteration 6, the reduction algorithm removes all terms except one. This remaining term has an entry 1 corresponding to the location of the maximum entry of the original tensor and all other entries are 0.


FIGURE 3.1. Comparing the maximum absolute values of entries from rank one terms of a CTD. After 6 iterations Algorithm 2 converges to a rank 1 tensor.

Another interesting case occurs when a tensor has multiple maximum entry candidates, either close to or exactly the same in magnitude. To explore this case, we construct an experiment similar to the previous one, namely, find the maximum absolute value of a CTD of dimension $d = 6$ and $M = 32$ samples in each direction. This CTD is constructed with the same rank 3 random CTD as above, but this time we added in two spikes at random locations so that there are two maximal spikes of absolute value 3.5. In cases such as this, Algorithm 2 can be run for a predetermined number of iterations, after which the factors are examined and a subset of candidates can be identified and explicitly verified. For example, in Figure 3.2, we observe that by iteration 6 there are only two candidates for the location of the maximum. As a warning we note that allowing the code to continue to run many more iterations will leave only one term due to the finite reduction tolerance ϵ of τ_ϵ in Algorithm 2. Indeed, in Figure 3.2 we observe the maximum entries of the two rank one terms begin to split around iteration number 26, and by iteration 37 only one term remains. This example illustrates a potential danger of insisting on a single term termination condition.


FIGURE 3.2. In this example there are two maxima of the same size. The values at these locations converge to $1/\sqrt{2}$ and the rest of the entries converge to 0. If allowed to run long enough, the value at one of the locations becomes dominant due to the truncation error of the rank reduction algorithm. The other location is then eliminated by the algorithm and only one term remains.

4. CTD-BASED GLOBAL OPTIMIZATION OF MULTIVARIATE FUNCTIONS

We describe an approach using Algorithm 2 for the global optimization of functions that admit low rank separated representations. We consider two types of problems: first, where the objective function is already in a separated form, ready

to yield a CTD appropriate for use with Algorithm 2. The second, more general problem, is to construct a separated representation given data or an analytic expression of the function, from which a CTD can then be built.

In the first problem the objective function is already represented in separated form as in (2.1), and an interpolation scheme is associated with each direction. Such problems have been subject to growing interest due to the use of CTDs for solving operator equations, e.g., deterministic or stochastic PDE/ODE systems (see, for example, [5, 6, 13, 24, 11, 12]).

In the second problem the goal is to optimize a multivariate function where no separated representation is readily available. What typically is available is a data set of function values (scattered or on a grid). In such cases our recourse is to first construct a separated representation of the underlying multivariate objective function and then re-sample it in order to obtain a CTD of the form (2.2). The construction of a separated representation given a set of function values can be thought of as a regression problem (see [4, 15]). The algorithms in these papers use function samples to build a separated representation of form (2.1). Following [4], an interpolation scheme is set up in each direction and the ALS algorithm is used to reduce the problem to regression in a single variable. The univariate interpolation scheme may use a variety of possible bases, as well as nonlinear approximation techniques. Once a separated representation is constructed, the CTD for finding the global maxima is then obtained to satisfy the local interpolation requirement stated below. The resulting CTD is then used as input into Algorithm 2 to produce an approximation of the global maxima. An example of this approach is included in Section 5.3 below.

If the objective function is given analytically, it is preferable to use analytic techniques to approximate the original function via a separated representation. In such cases, sampling the objective function becomes an important issue when finding global maxima of a function via Algorithm 2. The key sampling requirement is to ensure that the objective function can be interpolated at an arbitrary point from its sampled values up to a desired accuracy. This implies that sampling rate must depend on the local behavior of the function, e.g., the sampling rate is greater near singularities. For example, if a function has algebraic singularities, an efficient method to interpolate is to use wavelet bases. In such a case we emphasize the importance of using interpolating scaling functions, i.e., scaling functions whose coefficients are function values or are simply related to those. Since the global maxima are typically searched for in finite domains, the multiresolution basis should also work well in an interval (or a box in higher dimensions). This leads us to suggest the use of Lagrange interpolating polynomials within an adaptive spatial subdivision framework. We note that rescaled Lagrange interpolating polynomials form an orthonormal basis on an interval and the spatial subdivision is available as multiwavelet bases (see e.g., [1, 2]).

Finding the maximum entries of a CTD yields only an approximation of the maxima and their locations for the function. However, due to the interpolation requirements stated above these are high-quality approximations so that, if needed, a local optimization algorithm can then be used to find the true global maxima. Since in this case we use local optimization, it is preferable to oversample the CTD representation of a function, i.e. to use large M_j's in (2.2). This will not result in a prohibitive cost as the rank reduction algorithms, e.g. ALS, that operate on CTDs

scale linearly with respect to the number of samples in each direction [6]. However, the total number of samples needed to satisfy the local interpolation requirement may be large and, thus, additional steps to accelerate the reduction algorithm τ_ϵ in Algorithm 2 may be required. One such technique consists of using the QR factorization as explained in [8, Section 2.4].

An example of applying analytic techniques to put a function in separated form and subsequently using Algorithm 2 to optimize the separated function is provided in Section 5.2.

5. NUMERICAL EXAMPLES

5.1. **Comparison with power method algorithm from Section 2.2.** To test Algorithm 2, we construct tensors from random factors and find the entries with the largest absolute value. For these examples we select dimension $d = 8$ and set $M = 32$ samples in each direction. The test CTDs are constructed by adding rank four and rank one CTDs together. The rank four CTDs are composed of random factors (vectors $\mathbf{u}_j^{(l)}$ in (2.2)) whose entries are drawn from the uniform distribution on $[.9, 1]$ and the corresponding s-values are set to 1. The rank one CTDs contain all zeros except for a magnitude 4 spike added in at a random location. The CTDs in these examples have a unique maximum entry, and we chose to terminate the algorithm when the output CTD reaches separation rank one. We run tests for both Algorithm 2 and Algorithm 1, using for the rank reduction operation τ_ϵ the CTD-ID algorithm with accuracy threshold $\epsilon = 10^{-6}$ in the s-norm.


FIGURE 5.1. Histograms illustrating the number of iterations required for convergence to a rank one CTD by the power method optimization (hatched pattern) and the squaring method optimization (solid gray).

A histogram displaying the required number of iterations to converge for 500 individual tests is shown in Figure 5.1. While both algorithms find the correct maximum location in all tests, Algorithm 2 outperforms Algorithm 1 in terms of the number of iterations required. While Algorithm 2 requires much fewer iterations than the power method for the same problem, a lingering concern may be that the squaring process in Algorithm 2 requires the reduction of CTDs with much larger separation ranks. In Figure 5.2 we show histograms of computation times. We observe that despite reducing larger separation rank CTDs, the computation times using Algorithm 2 are significantly smaller. These computation times are for our MATLAB code running on individual cores of a Intel Xeon 2.40 GHz processor.

FIGURE 5.2. Histograms illustrating the computation times, in seconds, required for convergence to a rank one CTD by the power method optimization (hatched pattern) and the squaring method optimization (solid gray).

5.2. Optimization example.

5.2. **Optimization example.** As an example of applying analytic techniques to put a function in separated form for the purpose of optimization, we consider Ackley's test function [3], commonly used for testing global optimization algorithms,

$$
(5.1) \qquad u\left(\mathbf{x}\right) = a\, e^{-b\left(\frac{1}{d}\sum_{i=1}^{d} x_i^2\right)^{1/2}} + e^{\frac{1}{d}\sum_{i=1}^{d}\cos(cx_i)},
$$

where d is the number of dimensions and a, b, and c are parameters. For our tests we set $d = 10$, $a = 20$, $b = 0.2$, and $c = 2\pi$. We choose this test function for two reasons: first, it has many local maxima, thus a local method without a good initial guess may not find the correct maximum. Second, one of the terms is not in separated form, hence we must construct a separated approximation of the function first.

The true maximum of (5.1) occurs at $\mathbf{x} = \mathbf{0}$ with a value of $u\left(\mathbf{0}\right) = a + e$. The first term in (5.1) is radially symmetric which we approximate with a linear combination of Gaussians using the approximation of e^{-xy} from [7, Section 2.2],

$$
(5.2) \qquad G_e\left(x\right) = \frac{hb}{2\sqrt{\pi d}}\sum_{j=0}^{R}\exp\left(-\frac{b^2}{4d}e^{s_j} - x^2 e^{s_j} + \frac{1}{2}s_j\right),
$$

where $s_j = s_{start} + jh$, and h, s_{start}, and R are parameters chosen such that given $\epsilon, \delta > 0$,

$$
(5.3) \qquad \left| G_e\left(x\right) - e^{-\frac{b}{\sqrt{d}}x} \right| \leq \epsilon
$$

for $0 < \delta < x < \infty$. We arrive at the approximation of the first term in (5.1),

$$(5.4) \qquad \left| e^{-b\left(\frac{1}{d}\sum_{i=1}^{d} x_i^2\right)^{1/2}} - \sum_{j=0}^{R} w_j \prod_{i=1}^{d} \exp\left(-x_i^2 e^{s_j}\right) \right| \le \epsilon$$

where $w_j = \frac{hb}{2\sqrt{\pi d}} \exp\left(-\frac{b^2}{4d} e^{s_j} + \frac{1}{2} s_j\right)$.

To correctly sample (5.1), we first consider two terms of the function separately, the radially-symmetric exponential term, and the exponential cosine term which is already in separated form, $\exp\left(\frac{1}{d}\sum_{i=1}^{d} \cos\left(cx_i\right)\right) = \prod_{i=1}^{d} \exp\left(\frac{1}{d} \cos\left(cx_i\right)\right)$. These terms require different sampling strategies. The sampling rate of the radially-symmetric exponential term near the origin should be significantly higher than away from it. On the other hand, uniform sampling is sufficient for the exponential cosine term. We briefly discuss the combined sampling of these terms to satisfy both requirements.

To sample the radially-symmetric exponential term in (5.1), we use the expansion by Gaussians (5.2) as a guide. We choose 10 equally-spaced points centered at zero and sample the Gaussian e^{-x^2}. For each Gaussian in (5.2), we scale these points by the factor $e^{s_j/2}$ in order to consistently sample each Gaussian independent of its scale. Since (5.2) includes a large number of rapidly decaying Gaussians, too many of the resulting samples will concentrate near the origin. Therefore, we select a subset as follows: first, we keep all samples from the sharpest Gaussian, x_{0l}, $l = 1, \ldots, 10$. We then take the samples from second sharpest Gaussian and keep only those outside the range the previous samples, i.e., we keep x_{1m} such that $|x_{1m}| > \max_l |x_{0l}|$, $m = 1, \ldots, 10$. Repeating this process for all terms $j = 0, 1, \ldots, R$ in expansion (5.2) yields samples for the radially-symmetric exponential term in (5.1).

We oversample the exponential cosine term using 16 samples per oscillation. Notice that this sampling strategy is not sufficient for the radially-symmetric exponential term in (5.1) (and vice-versa). To combine these samples, we first find the radius where the sampling rate of the radially-symmetric exponential term in (5.1) drops below the sampling rate chosen to sample the exponential cosine term. We then remove all samples whose coordinates are greater in absolute value than the radius, and replace them with samples at the chosen rate of 16 samples per oscillation. The resulting set of samples consists of $M = 1080$ points for each direction $j = 1, \ldots, 10$.

Using expansion (5.2) leads to a 77 term CTD representation for the chosen accuracy $\epsilon = 10^{-8}$ and parameter $\delta = 3 \cdot 10^{-6}$ in (5.3). An initial reduction of this CTD using the CTD-ID algorithm with accuracy threshold $\epsilon = 10^{-6}$ in the Frobenius norm produces a CTD with separation rank 11 representing the function $u\left(\mathbf{x}\right)$. We use this CTD as an input into Algorithm 2, yielding a rank one solution after 24 iterations. We set the reduction tolerance in Algorithm 2 for this experiment to $\epsilon = 10^{-6}$ and find the maximum entry of the tensor located at a distance 2.13×10^{-3} away from the function's true maximum location, the origin. Using the output as the initial guess in a local, gradient-free, optimization algorithm (in our case a compass search, see e.g. [21]) yields the maximum value with relative error of 1.03×10^{-7} located at a distance 1.84×10^{-6} away from the origin.

5.3. **PDE optimization example.** As an example of constructing a separated representation from function evaluations for the purpose of optimization, we consider the problem of finding the maximum of the solution $u(\mathbf{x}, \mathbf{z})$ to the parameterized linear elliptic PDE

$$(5.5) \qquad \begin{aligned} -\nabla \cdot (a(\mathbf{x}, \mathbf{z})\nabla u(\mathbf{x}, \mathbf{z})) &= 1, \quad \mathbf{x} \in \mathcal{D}, \\ u(\mathbf{x}, \mathbf{z}) &= 0, \quad \mathbf{x} \in \partial\mathcal{D}, \end{aligned}$$

defined on the unit square $\mathcal{D} = (0, 1) \times (0, 1)$ with boundary $\partial\mathcal{D}$. This problem represents a generic model for diffusive processes and is widely used to study the effect of parameter uncertainty in applications such as subsurface flows, heat conduction, and chemical transport.

The diffusion coefficient $a(\mathbf{x}, \mathbf{z})$ is parametric and following [26] is modeled by

$$(5.6) \qquad a(\mathbf{x}, \mathbf{z}) = a_0 + \sigma_a \sum_{k=1}^{d} \sqrt{\zeta_k} \varphi_k(\mathbf{x}) z_k,$$

where $\mathbf{z} = (z_1, \ldots, z_d)$ and the parameters z_k are distributed over the interval $[-1, 1]$, and we choose $a_0 = 0.1$, $\sigma_a = 0.01$, and $d = 20$. In (5.6), $\{\zeta_k\}_{k=1}^{d}$ are the d largest eigenvalues associated with $\{\varphi_k\}_{k=1}^{d}$, the $L_2(\mathcal{D})$-orthonormalized eigenfunctions of the kernel $C_{aa}(\mathbf{x}_1, \mathbf{x_2}) = \exp(-\|\mathbf{x}_1 - \mathbf{x}_2\|_1 / l_c)$, where l_c is set to $l_c = 2/3$. Given the choices of parameters a_0, σ_a, d, and l_c, the model in (5.6) leads to strictly positive realizations of $a(\mathbf{x}, \mathbf{z})$.

The objective in this example is to approximate $u(\mathbf{x}, \mathbf{z})$ at $\mathbf{x}^* = (0.4839, 0.4839)$ as a function of $\mathbf{z}$ and find the maximum value of $u(\mathbf{x}^*, \mathbf{z})$ to study the worst-case scenario response of the system in the presence of parameter variations. In the following experiments we take two different approaches to optimizing $u(\mathbf{x}^*, \mathbf{z})$: The first approach is to construct a separated representation of $u(\mathbf{x}^*, \mathbf{z})$, evaluate it on a tensor grid yielding a CTD, and find the maximum via Algorithm 2. For comparison, we also use a genetic algorithm, a representative approach to finding the maximum of non-convex functions. Both methods require evaluation of $u(\mathbf{x}^*, \mathbf{z})$ (at a fixed $\mathbf{z} \in [-1, 1]^{20}$) by solving (5.5). In this example we use a finite element solver with a triangular mesh of uniform size $h = 1/32$.

In our approach, we construct a separated representation by first computing 600 realizations of $u(\mathbf{x}^*, \mathbf{z})$ via solution of (5.5). Using the resulting data points, we estimate $u(\mathbf{x}^*, \mathbf{z})$ in separated form using the method from [4, 15]. The resulting function has separation rank $r = 5$, where the factors $u_j^{(l)}(z_j)$ in (2.1) are represented in a Legendre basis of polynomials of degree at most three. We then validate the separated representation by computing 3400 additional samples of $u(\mathbf{x}^*, \mathbf{z})$; the relative RMS validation error is 3.256×10^{-4}, which is within the accuracy of the finite element solver.

Using the separated representation, we construct a CTD with 32 samples equispaced in $[-1, 1]$ for each direction. We then use Algorithm 2 (with an ALS tolerance of $\epsilon = 10^{-3}$ to keep the separation rank relatively low) to find the maximum entry of the CTD. As a baseline for comparing the output from Algorithm 2, we use the genetic algorithm from the MATLAB Optimization Toolbox with the recommended settings.

To compare the performance of Algorithm 2 with the genetic algorithm, we conducted 500 tests of each approach since both algorithms have random components. Algorithm 2 uses random rank updates during the ALS reduction procedure, while

the genetic algorithm uses different initial populations for each test. The largest maximum found by Algorithm 2 in the 500 tests was 8.744×10^{-1}, with mean 8.731×10^{-1} and standard deviation 5.7×10^{-4}. For comparison, the largest maxima found by the genetic algorithm in these tests was 8.750×10^{-1}, with mean 8.747×10^{-1} and standard deviation 3.6×10^{-4}. Taking into account that the accuracy of solving (5.5) and the RMS validation error of the separated representation are both of the same order as the computed standard deviations, the results obtained from the two algorithms are consistent.

On the other hand, an important difference between the two methods is the number of function evaluations required to find the maximum, in our case solving (5.5) once for each function evaluation. While in the 500 experiments described above the genetic algorithm required on average 4.419×10^4 solutions of (5.5) to obtain the maximum, the separated representation of $u(\mathbf{x}^*, \mathbf{z})$ was constructed using only 600 solutions of (5.5). Algorithm 2 required on average of 34 iterations to compute the maximum, whereas the largest number of iterations required was 43.

It is well known that genetic algorithms require a large number of function evaluations. In practical applications computing a large number of solutions of PDEs may be impractical. In such applications it is quite common to construct so-called surrogate models to reduce computational costs (see, e.g. [25]). For example, one may use a separated representation as such a surrogate, and then optimize via genetic algorithm. However, even with this modification, the advantage of Algorithm 2 over genetic algorithms (and other global search methods) is linear scaling of its complexity with the dimension d.

6. Discussion and conclusions

Entries of a tensor in CTD format with the largest absolute value can be reliably identified by the new quadratically convergent algorithm, Algorithm 2. This algorithm, in turn, can be used for global optimization of multivariate functions that admit separated representations. Unlike algorithms based on random search strategies (see e.g. [19]), our approach allows the user to estimate the potential error of the result. The computational cost of Algorithm 2 is dominated by the cost of reducing the separation ranks of intermediate CTDs, not the dimension of the optimization problem. If ALS is used to reduce the separation rank and d is the dimension, r is the maximum separation rank (after reductions) during the iteration, and M is the maximum number of components in each direction, then the computational cost of each rank reduction step can be estimated as $\mathcal{O}\left(r^4 \cdot M \cdot d \cdot N_{iter}\right)$, where N_{iter} is the number of iterations required by the ALS algorithm to converge. The computational cost of the CTD-ID algorithm is estimated as $\mathcal{O}\left(r^3 \cdot M \cdot d\right)$ and, if it is used instead of ALS, the reduction step is faster by a factor of $\mathcal{O}\left(r \cdot N_{iter}\right)$ [8]. We note that, while linear in dimension, Algorithm 2 may require significant computational resources due to the cubic (or quartic for ALS) dependence on the separation rank r.

Finally, we note that the applicability of our approach to global optimization extends beyond the use of separated representations of multivariate functions or tensors in CTD format. In fact, any structural representation of a multivariate function that allows rapid computation of the Hadamard product of corresponding tensors and has an associated algorithm for reducing the complexity of the representation should work in a manner similar to Algorithm 2.

References

[1] B. Alpert. A class of bases in L^2 for the sparse representation of integral operators. *SIAM J. Math. Anal*, 24(1):246–262, 1993.

[2] B. Alpert, G. Beylkin, D. Gines, and L. Vozovoi. Adaptive solution of partial differential equations in multiwavelet bases. *J. Comput. Phys.*, 182(1):149–190, 2002.

[3] T. Bäck. *Evolutionary algorithms in theory and practice: evolution strategies, evolutionary programming, genetic algorithms.* Oxford University Press, 1996.

[4] G. Beylkin, J. Garcke, and M. J. Mohlenkamp. Multivariate regression and machine learning with sums of separable functions. *SIAM Journal on Scientific Computing*, 31(3):1840–1857, 2009.

[5] G. Beylkin and M. J. Mohlenkamp. Numerical operator calculus in higher dimensions. *Proc. Natl. Acad. Sci. USA*, 99(16):10246–10251, August 2002.

[6] G. Beylkin and M. J. Mohlenkamp. Algorithms for numerical analysis in high dimensions. *SIAM J. Sci. Comput.*, 26(6):2133–2159, July 2005.

[7] G. Beylkin and L. Monzón. Approximation of functions by exponential sums revisited. *Appl. Comput. Harmon. Anal.*, 28(2):131–149, 2010.

[8] D. J. Biagioni, D. Beylkin, and G. Beylkin. Randomized interpolative decomposition of separated representations. *Journal of Computational Physics*, 281:116–134, 2015.

[9] R. Bro. Parafac. Tutorial & Applications. In *Chemom. Intell. Lab. Syst., Special Issue 2nd Internet Conf. in Chemometrics (incinc'96)*, volume 38, pages 149–171, 1997. http://www.models.kvl.dk/users/rasmus/presentations/parafac_tutorial/paraf.htm.

[10] J. D. Carroll and J. J. Chang. Analysis of individual differences in multidimensional scaling via an N-way generalization of Eckart-Young decomposition. *Psychometrika*, 35:283–320, 1970.

[11] F. Chinesta, P. Ladeveze, and E. Cueto. A short review on model order reduction based on proper generalized decomposition. *Archives of Computational Methods in Engineering*, 18(4):395–404, 2011.

[12] H. Cho, D. Venturi, and G.E. Karniadakis. Numerical methods for high-dimensional probability density function equations. *Journal of Computational Physics*, 305:817–837, 2016.

[13] A. Doostan and G. Iaccarino. A least-squares approximation of partial differential equations with high-dimensional random inputs. *Journal of Computational Physics*, 228(12):4332–4345, 2009.

[14] A. Doostan, G. Iaccarino, and N. Etemadi. A least-squares approximation of high-dimensional uncertain systems. Technical Report Annual Research Brief, Center for Turbulence Research, Stanford University, 2007.

[15] A. Doostan, A. Validi, and G. Iaccarino. Non-intrusive low-rank separated approximation of high-dimensional stochastic models. *Comput. Methods Appl. Mech. Engrg.*, 263:42–55, 2013.

[16] M. Espig, W. Hackbusch, A. Litvinenko, H.G. Matthies, and E. Zander. Efficient analysis of high dimensional data in tensor formats. In *Sparse Grids and Applications*, pages 31–56. Springer, 2013.

[17] M. Hadigol, A. Doostan, H. Matthies, and R. Niekamp. Partitioned treatment of uncertainty in coupled domain problems: A separated representation approach. *Computer Methods in Applied Mechanics and Engineering*, 274:103–124, 2014.

[18] R. A. Harshman. Foundations of the Parafac procedure: model and conditions for an "explanatory" multi-mode factor analysis. Working Papers in Phonetics 16, UCLA, 1970. http://publish.uwo.ca/~harshman/wpppfac0.pdf.

[19] R. Horst and P. M. Pardalos. *Handbook of global optimization*, volume 2. Springer Science & Business Media, 2013.

[20] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.

[21] T. G. Kolda, R.M. Lewis, and V. Torczon. Optimization by direct search: New perspectives on some classical and modern methods. *SIAM review*, 45(3):385–482, 2003.

[22] A. Nouy. A generalized spectral decomposition technique to solve a class of linear stochastic partial differential equations. *Computer Methods in Applied Mechanics and Engineering*, 196(37-40):4521–4537, 2007.

[23] A. Nouy. Generalized spectral decomposition method for solving stochastic finite element equations: Invariant subspace problem and dedicated algorithms. *Computer Methods in Applied Mechanics and Engineering*, 197:4718–4736, 2008.

[24] A. Nouy. Proper generalized decompositions and separated representations for the numerical solution of high dimensional stochastic problems. *Archives of Computational Methods in Engineering*, 17:403–434, 2010.

[25] Y.S. Ong, P.B. Nair, and A.J. Keane. Evolutionary optimization of computationally expensive problems via surrogate modeling. *American Institute of Aeronautics and Astronautics Journal*, 41(4):687–696, 2003.

[26] M. Reynolds, A. Doostan, and G. Beylkin. Randomized alternating least squares for canonical tensor decompositions: application to a pde with random data. *SIAM J. Sci. Comput.*, accepted for publication (arXiv preprint arXiv:1510.01398), 2016.

[27] J.C. Spall. *Introduction to stochastic search and optimization: estimation, simulation, and control*, volume 65. John Wiley & Sons, 2005.

[28] G. Tomasi and R. Bro. A comparison of algorithms for fitting the PARAFAC model. *Comput. Statist. Data Anal.*, 50(7):1700–1734, 2006.

* Department of Aerospace Engineering Sciences, 429 UCB, University of Colorado at Boulder, Boulder, CO 80309, † Department of Applied Mathematics, UCB 526, University of Colorado at Boulder, Boulder, CO 80309