
Structured adaptive and random spinners for fast machine learning computations

Mariusz Bojarski¹
NVIDIA
mbojarski@nvidia.com

Francois Fagan^{1,2}
Columbia University
ff2316@columbia.edu

Nourhan Sakr^{1,4}
Columbia University
n.sakr@columbia.edu

Anna Choromanska¹
NYU Tandon School of Engineering, ECE
ac5455@nyu.edu

Cédric Gouy-Pailler^{1,3}
CEA, LIST
cedric.gouy-pailler@cea.fr

Tamas Sarlos¹
Google Research
stamas@google.com

Krzysztof Choromanski¹
Google Brain Robotics
kchoro@google.com

Anne Morvan^{1,3}
CEA, LIST,
Univ. Paris-Dauphine, CNRS, LAMSADE
anne.morvan@cea.fr

Jamal Atif
Univ. Paris-Dauphine, CNRS, LAMSADE
jamal.atif@dauphine.fr

Abstract

We consider an efficient computational framework for speeding up several machine learning algorithms with almost no loss of accuracy. The proposed framework relies on projections via structured matrices that we call *Structured Spinners*, which are formed as products of three structured matrix-blocks that incorporate rotations. The approach is highly generic, i.e. i) structured matrices under consideration can either be fully-randomized or learned, ii) our structured family contains as special cases all previously considered structured schemes, iii) the setting extends to the non-linear case where the projections are followed by non-linear functions, and iv) the method finds numerous applications including kernel approximations via random feature maps, dimensionality reduction algorithms, new fast cross-polytope LSH techniques, deep learning, convex optimization algorithms via Newton sketches, quantization with random projection trees, and more. The proposed framework comes with theoretical guarantees characterizing the capacity of the structured model in reference to its unstructured counterpart and is based on a general theoretical

principle that we describe in the paper. As a consequence of our theoretical analysis, we provide the first theoretical guarantees for one of the most efficient existing LSH algorithms based on the $\mathbf{HD}_3\mathbf{HD}_2\mathbf{HD}_1$ structured matrix [Andoni et al., 2015]. The exhaustive experimental evaluation confirms the accuracy and efficiency of structured spinners for a variety of different applications.

1 Introduction

A striking majority of machine learning frameworks performs projections of input data via matrices of parameters, where the obtained projections are often passed to a possibly highly nonlinear function. In the case of randomized machine learning algorithms, the projection matrix is typically Gaussian with i.i.d. entries taken from $\mathcal{N}(0, 1)$. Otherwise, it is learned through the optimization scheme. A plethora of machine learning algorithms admits this form. In the randomized setting, a few examples include variants of the Johnson-Lindenstrauss Transform applying random projections to reduce data dimensionality while approximately preserving Euclidean distance [Ailon and Chazelle, 2006, Liberty et al., 2008, Ailon and Liberty, 2011], kernel approximation techniques based on random feature maps produced from linear projections with Gaussian matrices followed by nonlinear mappings [Rahimi and Recht, 2007], [Le et al., 2013, Choromanski and Sindhwani, 2016, Huang et al., 2014], [Choromanska et al., 2016], LSH-based schemes [Har-Peled et al., 2012, Charikar, 2002, Terasawa and Tanaka, 2007], including the fastest known variant of the cross-polytope LSH [Andoni et al., 2015], algorithms solving convex optimization problems with random sketches of

¹equal contribution

²supported by Bloomberg LP grant

³CEA, LIST, 91191 Gif-sur-Yvette, France. Partly supported by *DGA* (French Ministry of Defense)

⁴partly supported by NSF grant CCF-1421161

Hessian matrices [Pilanci and Wainwright, 2015, Pilanci and Wainwright, 2014], quantization techniques using random projection trees, where splitting in each node is determined by a projection of data onto Gaussian direction [Dasgupta and Freund, 2008], and many more.

The classical example of machine learning nonlinear models where linear projections are learned is a multi-layered neural network [LeCun et al., 2015, Goodfellow et al., 2016], where the operations of linear projection via matrices with learned parameters followed by the pointwise nonlinear feature transformation are the building blocks of the network’s architecture. These two operations are typically stacked multiple times to form a deep network.

The computation of projections takes $\Theta(mn|\mathcal{X}|)$ time, where $m \times n$ is the size of the projection matrix, and $|\mathcal{X}|$ denotes the number of data samples from a dataset \mathcal{X} . In case of high-dimensional data, this comprises a significant fraction of the overall computational time, while storing the projection matrix frequently becomes a bottleneck in terms of space complexity.

In this paper, we propose the remedy for both problems, which relies on replacing the aforementioned algorithms by their “structured variants”. The projection is performed by applying a structured matrix from the family that we introduce as *Structured Spinners*. Depending on the setting, the structured matrix is either learned or its parameters are taken from a random distribution (either continuous or discrete if further compression is required). Each structured spinner is a product of three matrix-blocks that incorporate rotations. A notable member of this family is a matrix of the form $\mathbf{HD}_3\mathbf{HD}_2\mathbf{HD}_1$, where \mathbf{D}_i s are either random diagonal ± 1 -matrices or adaptive diagonal matrices and \mathbf{H} is the Hadamard matrix. This matrix is used in the fastest known cross-polytope LSH method introduced in [Andoni et al., 2015].

In the structured case, the computational speedups are significant, i.e. projections can be calculated in $o(mn)$ time, often in $O(n \log m)$ time if Fast Fourier Transform techniques are applied. At the same time, using matrices from the family of structured spinners leads to the reduction of space complexity to sub-quadratic, usually at most linear, or sometimes even constant.

The key contributions of this paper are:

- The family of structured spinners providing a highly parametrized class of structured methods and, as we show in this paper, with applications in various randomized settings such as: kernel approximations via random feature maps, dimensionality reduction algorithms, new fast cross-polytope LSH techniques, deep learning, convex optimization algorithms via Newton sketches, quantization with random projection trees, and more.
- A comprehensive theoretical explanation of the effectiveness of the structured approach based on structured spinners. Such analysis was provided in the literature before for a strict subclass of a very general family of structured matrices that we consider in this paper, i.e. the proposed family of structured spinners contains all previously considered structured matrices as special cases, including the recently introduced P -model [Choromanski and Sindhwani, 2016]. To the best of our knowledge, we are the first to theoretically explain the effectiveness of structured neural network architectures. Furthermore, we provide first theoretical guarantees for a wide range of discrete structured transforms, in particular for the fastest known cross-polytope LSH method [Andoni et al., 2015] based $\mathbf{HD}_3\mathbf{HD}_2\mathbf{HD}_1$ discrete matrices.

Our theoretical methods in the random setting apply the relatively new Berry-Esseen type Central Limit Theorem results for random vectors.

Our theoretical findings are supported by empirical evidence regarding the accuracy and efficiency of structured spinners in a wide range of different applications. Not only do structured spinners cover all already existing structured transforms as special instances, but also many other structured matrices that can be applied in all aforementioned applications.

2 Related work

This paper focuses on structured matrices, which were previously explored in the literature mostly in the context of the Johnson-Lindenstrauss Transform (JLT) [Johnson and Lindenstrauss, 1984], where the high-dimensional data is linearly transformed and embedded into a much lower dimensional space while approximately preserving the Euclidean distance between data points. Several extensions of JLT have been proposed, e.g. [Liberty et al., 2008, Ailon and Liberty, 2011, Ailon and Chazelle, 2006, Vybíral, 2011]. Most of these structured constructions involve sparse [Ailon and Chazelle, 2006, Dasgupta et al., 2010] or circulant matrices [Vybíral, 2011, Hinrichs and Vybrál, 2011] providing computational speedups and space compression.

More recently, the so-called Ψ -regular structured matrices (Toeplitz and circulant matrices belong to this wider family of matrices) were used to approximate angular distances [Choromanska et al., 2016] and signed Circulant Random Matrices were used to approximate Gaussian kernels [Feng et al., 2015]. Another work [Choromanski and Sindhwani, 2016] applies structured matrices coming from the so-called P -model, which further generalizes the Ψ -regular fam-

ily, to speed up random feature map computations of some special kernels (angular, arc-cosine and Gaussian). These techniques did not work for discrete structured constructions, such as the $\mathbf{HD}_3\mathbf{HD}_2\mathbf{HD}_1$ matrices, or their direct non-discrete modifications, since they require matrices with low (polylog) chromatic number of the corresponding coherence graphs.

Linear projections are used in the LSH setting to construct codes for given datapoints which speed up such tasks as approximate nearest neighbor search. A notable set of methods are the so-called cross-polytope techniques introduced in [Terasawa and Tanaka, 2007] and their aforementioned discrete structured variants proposed in [Andoni et al., 2015] that are based on the Walsh-Hadamard transform. Before our work, they were only experimentally verified to produce good quality codes.

Furthermore, a recently proposed technique based on the so-called *Newton Sketch* provides yet another example of application for structured matrices. The method [Pilanci and Wainwright, 2015, Pilanci and Wainwright, 2014] is used for speeding up algorithms solving convex optimization problems by approximating Hessian matrices using so-called *sketch matrices*. Initially, the sub-Gaussian sketches based on i.i.d. sub-Gaussian random variables were used. The disadvantage of the sub-Gaussian sketches lies in the fact that computing the sketch of the given matrix of size $n \times d$ requires $O(mnd)$ time, where $m \times n$ in the size of the sketch matrix. Thus the method is too slow in practice and could be accelerated with the use of structured matrices. Some structured approaches were already considered, e.g. sketches based on randomized orthonormal systems were proposed [Pilanci and Wainwright, 2015].

All previously considered methods focus on the randomized setting, whereas the structured matrix instead of being learned is fully random. In the context of adaptive setting, where the parameters are being learned instead, we focus in this paper on multi-layer neural networks. We emphasize though that our approach is much more general and extends beyond this setting. Structured neural networks were considered before, for instance in [Yang et al., 2015], where the so-called *Deep Fried Neural Convnets* were proposed. Those architectures are based on the adaptive version of the Fastfood transform used for approximating various kernels [Le et al., 2013], which is a special case of structured spinner matrices.

Deep Fried Convnets apply adaptive structured matrices for fully connected layers of the convolutional networks. The structured matrix is of the form: $\mathbf{SHG\Pi\mathbf{H}\mathbf{B}}$, where \mathbf{S} , \mathbf{G} , and \mathbf{B} are adaptive diagonal matrices, $\mathbf{\Pi}$ is a random permutation matrix, and \mathbf{H} is the Walsh-Hadamard matrix. The method reduces the storage and computational costs of matrix multi-

plication step from, often prohibitive, $\mathcal{O}(nd)$ down to $\mathcal{O}(n)$ storage and $\mathcal{O}(n \log d)$ computational cost, where d and n denote the size of consecutive layers of the network. At the same time, this approach does not sacrifice the network’s predictive performance. Another work [Moczulski et al., 2016] that offers an improvement over Deep Fried Convnets, looks at a structured matrix family that is very similar to $\mathbf{HD}_3\mathbf{HD}_2\mathbf{HD}_1$ (however is significantly less general than the family of structured spinners). Their theoretical results rely on the analysis in [Huhtanen and Perämäki, 2015].

The Adaptive Fastfood approach elegantly complements previous works dedicated to address the problem of huge overparametrization of deep models with structured matrices, e.g. the method of [Denil et al., 2013] represents the parameter matrix as a product of two low rank factors and, similarly to Adaptive Fastfood, applies both at train and test time, [Sainath et al., 2013] introduces low-rank matrix factorization to reduce the size of the fully connected layers at train time, and [Li, 2013] uses low-rank factorizations with SVD after training the full model. These methods, as well as approaches that consider kernel methods in deep learning [Cho and Saul, 2009, Mairal et al., 2014, Dai et al., 2014, Huang et al., 2014], are conveniently discussed in [Yang et al., 2015].

Structured neural networks are also considered in [Sindhwani et al., 2015], where low-displacement rank matrices are applied for linear projections. The advantage of this approach over Deep Fried Convnets is due to the high parametrization of the family of low-displacement rank matrices allowing the adjustment of the number of parameters learned based on accuracy and speedup requirements.

The class of structured spinners proposed in this work is more general than Deep Fried Convnets or low displacement rank matrices, but it also provides much easier structured constructions, such as $\mathbf{HD}_3\mathbf{HD}_2\mathbf{HD}_1$ matrices, where \mathbf{D}_i s are adaptive diagonal matrices. Furthermore, to the best of our knowledge we are the first to prove theoretically that structured neural networks learn good quality models, by analyzing the capacity of the family of structured spinners.

3 The family of *Structured Spinners*

Before introducing the family of structured spinners, we explain notation. If not specified otherwise, matrix \mathbf{D} is a random diagonal matrix with diagonal entries taken independently at random from $\{-1, +1\}$. By $\mathbf{D}_{t_1, \dots, t_n}$ we denote the diagonal matrix with diagonal equal to (t_1, \dots, t_n) . For a matrix $\mathbf{A} = \{a_{i,j}\}_{i,j=1, \dots, n} \in \mathbb{R}^{n \times n}$, we denote by $\|\mathbf{A}\|_F$ its Frobenius norm, i.e. $\|\mathbf{A}\|_F = \sqrt{\sum_{i,j \in \{1, \dots, n\}} a_{i,j}^2}$, and by $\|\mathbf{A}\|_2$ its spectral norm, i.e. $\|\mathbf{A}\|_2 = \sup_{\mathbf{x} \neq 0} \frac{\|\mathbf{A}\mathbf{x}\|_2}{\|\mathbf{x}\|_2}$. We denote by \mathbf{H}

the L_2 -normalized Hadamard matrix. We say that \mathbf{r} is a random Rademacher vector if every element of \mathbf{r} is chosen independently at random from $\{-1, +1\}$.

For a vector $\mathbf{r} \in \mathbb{R}^k$ and $n > 0$ let $\mathbf{C}(\mathbf{r}, n) \in \mathbb{R}^{n \times nk}$ be a matrix, where the first row is of the form $(\mathbf{r}^T, 0, \dots, 0)$ and each subsequent row is obtained from the previous one by right-shifting in a circulant manner the previous one by k . For a sequence of matrices $\mathbf{W}^1, \dots, \mathbf{W}^n \in \mathbb{R}^{k \times n}$ we denote by $\mathbf{V}(\mathbf{W}^1, \dots, \mathbf{W}^n) \in \mathbb{R}^{nk \times n}$ a matrix obtained by vertically stacking matrices: $\mathbf{W}^1, \dots, \mathbf{W}^n$.

Each structured matrix $\mathbf{G}_{struct} \in \mathbb{R}^{n \times n}$ from the family of structured spinners is a product of three main structured components/blocks, i.e.:

$$\mathbf{G}_{struct} = \mathbf{M}_3 \mathbf{M}_2 \mathbf{M}_1, \quad (1)$$

where matrices $\mathbf{M}_1, \mathbf{M}_2$ and \mathbf{M}_3 satisfy conditions:

Condition 1: Matrices: \mathbf{M}_1 and $\mathbf{M}_2 \mathbf{M}_1$ are $(\delta(n), p(n))$ -balanced isometries.

Condition 2: $\mathbf{M}_2 = \mathbf{V}(\mathbf{W}^1, \dots, \mathbf{W}^n) \mathbf{D}_{\rho_1, \dots, \rho_n}$ for some (Δ_F, Δ_2) -smooth set: $\mathbf{W}^1, \dots, \mathbf{W}^n \in \mathbb{R}^{k \times n}$ and some i.i.d sub-Gaussian random variables ρ_1, \dots, ρ_n with sub-Gaussian norm K .

Condition 3: $\mathbf{M}_3 = \mathbf{C}(\mathbf{r}, n)$ for $\mathbf{r} \in \mathbb{R}^k$, where \mathbf{r} is random Rademacher/Gaussian in the random setting and is learned in the adaptive setting.

Matrix \mathbf{G}_{struct} is a structured spinner with parameters: $\delta(n), p(n), K, \Lambda_F, \Lambda_2$. We explain the introduced conditions below.

Definition 1 ($(\delta(n), p(n))$ -balanced matrices)

A randomized matrix $\mathbf{M} \in \mathbb{R}^{n \times m}$ is $(\delta(n), p(n))$ -balanced if for every $\mathbf{x} \in \mathbb{R}^m$ with $\|\mathbf{x}\|_2 = 1$ we have: $\mathbb{P}[\|\mathbf{M}\mathbf{x}\|_\infty > \frac{\delta(n)}{\sqrt{n}}] \leq p(n)$.

Remark 1 One can take as \mathbf{M}_1 a matrix $\mathbf{H}\mathbf{D}_1$ since, as we will show in the Supplement, matrix $\mathbf{H}\mathbf{D}_1$ is $(\log(n), 2ne^{-\frac{\log^2(n)}{8}})$ -balanced.

Definition 2 ((Δ_F, Δ_2) -smooth sets) A deterministic set of matrices $\mathbf{W}^1, \dots, \mathbf{W}^n \in \mathbb{R}^{k \times n}$ is (Δ_F, Δ_2) -smooth if:

- $\|\mathbf{W}_1^i\|_2 = \dots = \|\mathbf{W}_n^i\|_2$ for $i = 1, \dots, n$, where \mathbf{W}_j^i stands for the j^{th} column of \mathbf{W}^i ,
- for $i \neq j$ and $l = 1, \dots, n$ we have: $(\mathbf{W}_l^i)^T \cdot \mathbf{W}_l^j = 0$,
- $\max_{i,j} \|(\mathbf{W}^j)^T \mathbf{W}^i\|_F \leq \Delta_F$ and $\max_{i,j} \|(\mathbf{W}^j)^T \mathbf{W}^i\|_2 \leq \Delta_2$.

Remark 2 If the unstructured matrix \mathbf{G} has rows taken from the general multivariate Gaussian distribution with diagonal covariance matrix $\Sigma \neq \mathbf{I}$ then one needs to rescale vectors \mathbf{r} accordingly. For clarity, we assume here that $\Sigma = \mathbf{I}$ and we present our theoretical results for that setting.

All structured matrices previously considered are special cases of a wider family of structured spinners (for clarity, we will explicitly show it for some important special cases). We have:

Lemma 1 The following matrices: $\mathbf{G}_{circ} \mathbf{D}_2 \mathbf{H}\mathbf{D}_1$, $\sqrt{n} \mathbf{H}\mathbf{D}_3 \mathbf{H}\mathbf{D}_2 \mathbf{H}\mathbf{D}_1$ and $\sqrt{n} \mathbf{H}\mathbf{D}_{g_1, \dots, g_n} \mathbf{H}\mathbf{D}_2 \mathbf{H}\mathbf{D}_1$, where \mathbf{G}_{circ} is Gaussian circulant, are valid structured spinners for $\delta(n) = \log(n)$, $p(n) = 2ne^{-\frac{\log^2(n)}{8}}$, $K = 1$, $\Lambda_F = O(\sqrt{n})$ and $\Lambda_2 = O(1)$. The same is true if one replaces \mathbf{G}_{circ} by a Gaussian Hankel or Toeplitz matrix.

3.1 The role of three blocks $\mathbf{M}_1, \mathbf{M}_2$, and \mathbf{M}_3

The role of blocks $\mathbf{M}_1, \mathbf{M}_2, \mathbf{M}_3$ can be intuitively explained. Matrix \mathbf{M}_1 makes vectors “balanced”, so that there is no dimension that carries too much of the L_2 -norm of the vector. The balanceness property was already applied in the structured setting [Ailon and Chazelle, 2006].

The role of \mathbf{M}_2 is more subtle and differs between adaptive and random settings. In the random setting, the cost of applying the structured mechanism is the loss of independence. For instance, the dot products of the rows of a circulant Gaussian matrix with a given vector \mathbf{x} are no longer independent, as it is the case in the fully random setup. Those dot products can be expressed as a dot product of a fixed Gaussian row with different vectors \mathbf{v} . Matrix \mathbf{M}_2 makes these vectors close to orthogonal. In the adaptive setup, the “close to orthogonality” property is replaced by the independence property.

Finally, matrix \mathbf{M}_3 defines the capacity of the entire structured transform by providing a vector of parameters (either random or to be learned). The near-independence of the aforementioned dot products in the random setting is now implied by the near-orthogonality property achieved by \mathbf{M}_2 and the fact that the projections of the Gaussian vector or the random Rademacher vector onto “almost orthogonal directions” are “close to independent”. The role of the three matrices is described pictorially in Figure 1.

3.2 Stacking together Structured Spinners

We described structured spinners as square matrices, but in practice we are not restricted to those, i.e. one can construct an $m \times n$ structured spinner for $m \leq n$ from the square $n \times n$ structured spinner by taking its first m rows. We can then stack vertically these independently constructed $m \times n$ matrices to obtain an $k \times n$ matrix for both: $k \leq n$ and $k > n$. We think about m as another parameter of the model that tunes the “structuredness” level, i.e. larger values of m indicate more structured approach while smaller values lead to more random matrices ($m = 1$ case is the fully




Figure 1: Pictorial explanation of the role of three matrix-blocks in the construction of the structured spinner. Left picture: \mathbf{M}_1 rotates \mathbf{v} such that the rotated version \mathbf{v}_r is balanced. Middle picture: \mathbf{M}_2 transforms vectors $\mathbf{v}, \mathbf{w}, \mathbf{u}$ such that their images $\mathbf{v}_r, \mathbf{w}_r, \mathbf{u}_r$ are near-orthogonal. Right picture: The projections of the random vector \mathbf{r} onto such two near-orthogonal vectors \mathbf{v}, \mathbf{w} are near-independent.

unstructured one).

4 Theoretical results

We now show that structured spinners can replace their unstructured counterparts in many machine learning algorithms with minimal loss of accuracy.

Let $\mathcal{A}_{\mathcal{G}}$ be a machine learning algorithm applied to a fixed dataset $\mathcal{X} \subseteq \mathbb{R}^n$ and parametrized by a set \mathcal{G} of matrices $\mathbf{G} \in \mathcal{R}^{m \times n}$, where each \mathbf{G} is either learned or Gaussian with independent entries taken from $\mathcal{N}(0, 1)$. Assume furthermore, that $\mathcal{A}_{\mathcal{G}}$ consists of functions f_1, \dots, f_s , where each f_i applies a certain matrix \mathbf{G}_i from \mathcal{G} to vectors from some linear space \mathcal{L}_i of dimensionality at most d . Note that for a fixed dataset \mathcal{X} function f_i is a function of a random vector

$$\mathbf{q}_{f_i} = ((\mathbf{G}_i \mathbf{x}^1)^T, \dots, (\mathbf{G}_i \mathbf{x}^{d_i})^T)^T \in \mathbb{R}^{d_i \cdot m},$$

where $\dim(\mathcal{L}_i) = d_i \leq d$ and $\mathbf{x}^1, \dots, \mathbf{x}^{d_i}$ stands for some fixed basis of \mathcal{L}_i .

Denote by f'_i the structured counterpart of f_i , where \mathbf{G}_i is replaced by the structured spinner (for which vector \mathbf{r} is either learned or random). We will show that f'_i s “resemble” f_i s distribution-wise. Surprisingly, we will show it under very weak conditions regarding f_i s. In particular, they can be nondifferentiable, even non-continuous.

Note that the above setting covers a wide range of machine learning algorithms. In particular:

Remark 3 In the kernel approximation setting with random feature maps one can match each pair of vectors $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ to a different $f = f_{\mathbf{x}, \mathbf{y}}$. Each f computes the approximate value of the kernel for vectors \mathbf{x} and \mathbf{y} . Thus in that scenario $s = \binom{|\mathcal{X}|}{2}$ and $d = 2$ (since one can take: $\mathcal{L}_{f(\mathbf{x}, \mathbf{y})} = \text{span}(\mathbf{x}, \mathbf{y})$).

Remark 4 In the vector quantization algorithms using random projection trees one can take $s = 1$ (the algorithm \mathcal{A} itself is a function f outputting the partitioning

of space into cells) and $d = d_{\text{intrinsic}}$, where $d_{\text{intrinsic}}$ is an intrinsic dimensionality of a given dataset \mathcal{X} (random projection trees are often used if $d_{\text{intrinsic}} \ll n$).

4.1 Random setting

We need the following definition.

Definition 3 A set \mathcal{S} is b -convex if it is a union of at most b pairwise disjoint convex sets.

Fix a function $f_i : \mathbb{R}^{d_i \cdot m} \rightarrow \mathcal{V}$, for some domain \mathcal{V} . Our main result states that for any $\mathcal{S} \subseteq \mathcal{V}$ such that $f_i^{-1}(\mathcal{S})$ is measurable and b -convex for b not too large, the probability that $f_i(\mathbf{q}_{f_i})$ belongs to \mathcal{S} is close to the probability that $f'_i(\mathbf{q}_{f'_i})$ belongs to \mathcal{S} .

Theorem 1 (structured random setting) Let \mathcal{A} be a randomized algorithm using unstructured Gaussian matrices \mathbf{G} and let s, d and f_i s be as at the beginning of the section. Replace the unstructured matrix \mathbf{G} by one of structured spinners defined in Section 3 with blocks of m rows each. Then for n large enough, $\epsilon = o_{md}(1)$ and fixed f_i with probability p_{succ} at least:

$$1 - 2p(n)d - 2 \binom{md}{2} e^{-\Omega(\min(\frac{\epsilon^2 n^2}{K^4 \Lambda_F^2 \delta^4(n)}, \frac{\epsilon n}{K^2 \Lambda_2 \delta^2(n)})} \quad (2)$$

with respect to the random choices of \mathbf{M}_1 and \mathbf{M}_2 the following holds for any \mathcal{S} such that $f_i^{-1}(\mathcal{S})$ is measurable and b -convex:

$$|\mathbb{P}[f_i(\mathbf{q}_{f_i}) \in \mathcal{S}] - \mathbb{P}[f'_i(\mathbf{q}_{f'_i}) \in \mathcal{S}]| \leq b\eta,$$

where the probabilities in the last formula are with respect to the random choice of \mathbf{M}_3 , $\eta = \frac{\delta^3(n)}{n^{\frac{5}{3}}}$, and $\delta(n), p(n), K, \Lambda_F, \Lambda_2$ are as in the definition of structured spinners from Section 3.

Remark 5 The theorem does not require any strong regularity conditions regarding f_i s (such as differentiability or even continuity). In practice, b is often a

small constant. For instance, for the angular kernel approximation where f_i s are non-continuous and for S -singletons, we can take $b = 1$ (see Supplement).

Now let us think of f_i and f'_i as random variables, where randomness is generated by vectors \mathbf{q}_{f_i} and $\mathbf{q}_{f'_i}$ respectively. Then, from Theorem 1, we get:

Theorem 2 Denote by F_X the cdf of the random variable X and by ϕ_X its characteristic function. If f_i is convex or concave in respect to \mathbf{q}_{f_i} , then for every t the following holds: $|F_{f_i}(t) - F_{f'_i}(t)| = O(\frac{\delta^3(n)}{n^{\frac{5}{2}}})$. Furthermore, if f_i is bounded then: $|\phi_{f_i}(t) - \phi_{f'_i}(t)| = O(\frac{\delta^3(n)}{n^{\frac{5}{2}}})$.

Theorem 1 implies strong accuracy guarantees for the specific structured spinners. As a corollary we get:

Theorem 3 Under assumptions from Theorem 1 the probability p_{succ} from Theorem 1 reduces to: $1 - 4ne^{-\frac{\log^2(n)}{8}}d - 2\binom{md}{2}e^{-\Omega(\frac{\epsilon^2 n}{\log^4(n)})}$ for the structured matrices $\sqrt{n}\mathbf{HD}_3\mathbf{HD}_2\mathbf{HD}_1$, $\sqrt{n}\mathbf{HD}_{g_1, \dots, g_n}\mathbf{HD}_2\mathbf{HD}_1$ as well as for the structured matrices of the form $\mathbf{G}_{struct}\mathbf{D}_2\mathbf{HD}_1$, where \mathbf{G}_{struct} is Gaussian circulant, Gaussian Toeplitz or Gaussian Hankel matrix.

As a corollary of Theorem 3, we obtain the following result showing the effectiveness of the cross-polytope LSH with structured matrices $\mathbf{HD}_3\mathbf{HD}_2\mathbf{HD}_1$ that was only heuristically confirmed before [Andoni et al., 2015].

Theorem 4 Let $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ be two unit L_2 -norm vectors. Let $\mathbf{v}_{\mathbf{x}, \mathbf{y}}$ be the vector indexed by all $(2m)^2$ ordered pairs of canonical directions $(\pm \mathbf{e}_i, \pm \mathbf{e}_j)$, where the value of the entry indexed by (\mathbf{u}, \mathbf{w}) is the probability that: $h(\mathbf{x}) = \mathbf{u}$ and $h(\mathbf{y}) = \mathbf{w}$, and $h(\mathbf{v})$ stands for the hash of \mathbf{v} . Then with probability at least: $p_{success} = 1 - 8ne^{-\frac{\log^2(n)}{8}} - 2\binom{2m}{2}e^{-\Omega(\frac{\epsilon^2 n}{\log^4(n)})}$ the version of the stochastic vector $\mathbf{v}_{\mathbf{x}, \mathbf{y}}^1$ for the unstructured Gaussian matrix \mathbf{G} and its structured counterpart $\mathbf{v}_{\mathbf{x}, \mathbf{y}}^2$ for the matrix $\mathbf{HD}_3\mathbf{HD}_2\mathbf{HD}_1$ satisfy: $\|\mathbf{v}_{\mathbf{x}, \mathbf{y}}^1 - \mathbf{v}_{\mathbf{x}, \mathbf{y}}^2\|_\infty \leq \log^3(n)n^{-\frac{2}{5}} + c\epsilon$, for n large enough, where $c > 0$ is a universal constant. The probability above is taken with respect to random choices of \mathbf{D}_1 and \mathbf{D}_2 .

For angles in the range $[0, \frac{\pi}{3}]$ the result above leads to the same asymptotics of the probabilities of collisions as these in Theorem 1 of [Andoni et al., 2015] given for the unstructured cross-polytope LSH.

The proof for the discrete structured setting applies Berry-Esseen-type results for random vectors (details are in the Supplement) showing that for n large enough ± 1 random vectors \mathbf{r} act similarly to Gaussian vectors.

4.2 Adaptive setting

The following theorem explains that structured spinners can be used to replace unstructured fully connected

neural network layers performing dimensionality reduction (such as hidden layers in certain autoencoders) provided that input data has low intrinsic dimensionality. These theoretical findings were confirmed in experiments that will be presented in the next section. We will use notation from Theorem 1.

Theorem 5 Consider a matrix $\mathbf{M} \in \mathbb{R}^{m \times n}$ encoding the weights of connections between a layer l_0 of size n and a layer l_1 of size m in some learned unstructured neural network model. Assume that the input to layer l_0 is taken from the d -dimensional space \mathcal{L} (although potentially embedded in a much higher dimensional space). Then with probability at least

$$1 - 2p(n)d - 2\binom{md}{2}e^{-\Omega(\min(\frac{t^2 n^2}{K^4 \Lambda_F^2 \delta^4(n)}, \frac{tn}{K^2 \Lambda_2 \delta^2(n)})} \quad (3)$$

for $t = \frac{1}{md}$ and with respect to random choices of \mathbf{M}_1 and \mathbf{M}_2 , there exists a vector \mathbf{r} defining \mathbf{M}_3 (see: definition of the structured spinner) such that the structured spinner $\mathbf{M}^{struct} = \mathbf{M}_3\mathbf{M}_2\mathbf{M}_1$ equals to \mathbf{M} on \mathcal{L} .

5 Experiments

In this section we consider a wide range of different applications of structured spinners: locality-sensitive hashing, kernel approximations, and finally neural networks. Experiments with Newton sketches are deferred to the Supplement. Experiments were conducted using Python. In particular, NumPy is linked against a highly optimized BLAS library (Intel MKL). Fast Fourier Transform is performed using numpy.fft and Fast Hadamard Transform is using ffht from [Andoni et al., 2015]. To have a fair comparison, we have set up: OMP_NUM_THREADS = 1 so that every experiment is done on a single thread. Every parameter of the structured spinner matrix is computed in advance, such that obtained speedups take only matrix-vector products into account. All figures should be read in color.

5.1 Locality-Sensitive Hashing (LSH)

In the first experiment, we consider cross-polytope LSH. In Figure 2, we compare collision probabilities for the low dimensional case ($n = 256$), where for each interval, collision probability has been computed for 20000 points. Results are shown for one hash function (averaged over 100 runs). We report results for a random 256×64 Gaussian matrix \mathbf{G} and five other types of matrices from a family of structured spinners (descending order of number of parameters): $\mathbf{G}_{circ}\mathbf{K}_2\mathbf{K}_1$, $\mathbf{G}_{Toeplitz}\mathbf{D}_2\mathbf{HD}_1$, $\mathbf{G}_{skew-circ}\mathbf{D}_2\mathbf{HD}_1$, $\mathbf{HD}_{g_1, \dots, g_n}\mathbf{HD}_2\mathbf{HD}_1$, and $\mathbf{HD}_3\mathbf{HD}_2\mathbf{HD}_1$, where \mathbf{K}_i , $\mathbf{G}_{Toeplitz}$, and $\mathbf{G}_{skew-circ}$ are respectively a Kronecker matrix with discrete entries, Gaussian Toeplitz and Gaussian skew-circulant matrices.

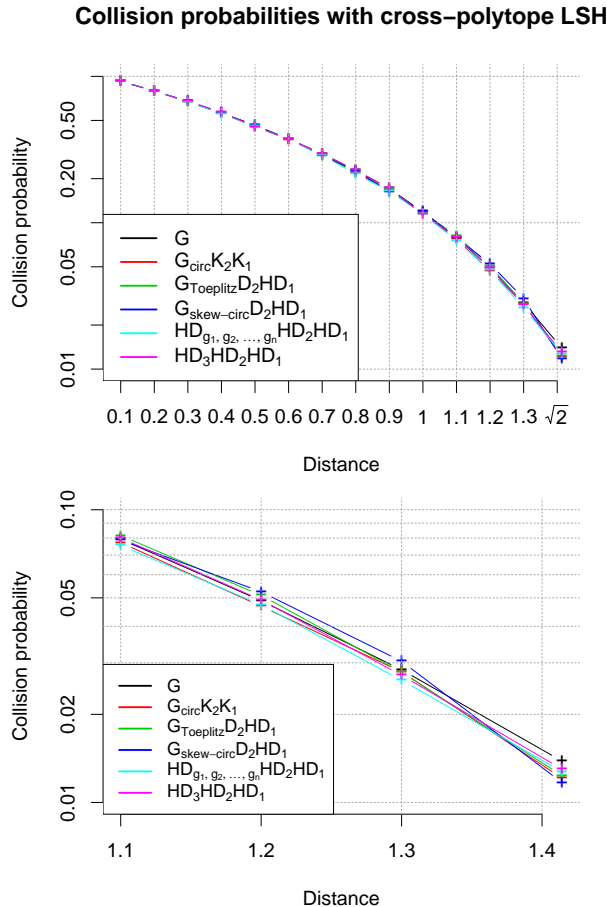


Figure 2: Cross-polytope LSH - collision probabilities. (bottom) A zoom on higher distances enables to distinguish the curves which are almost superposed.

All matrices from the family of structured spinners show high collision probabilities for small distances and low ones for large distances. As theoretically predicted, structured spinners do not lead to accuracy losses. All considered matrices give almost identical results.

5.2 Kernel approximation

In the second experiment, we approximate the Gaussian and angular kernels using Random Fourier features. The Gaussian random matrix (with i.i.d. Gaussian entries) can be used to sample random Fourier features with a specified σ . This Gaussian random matrix is replaced with specific matrices from a family of structured spinners for Gaussian and angular kernels. The obtained feature maps are compared. To test the quality of the structured kernels' approximations, we compute Gram-matrix reconstruction error as in [Choromanski and Sindhvani, 2016]: $\frac{\|\mathbf{K} - \tilde{\mathbf{K}}\|_F}{\|\mathbf{K}\|_F}$, where $\mathbf{K}, \tilde{\mathbf{K}}$ are respectively the exact and approximate Gram-matrices, as a function of the number of random features. When number of random features k is greater than data dimensionality n , we apply block-mechanism described in 3.2.

For the Gaussian kernel, $\mathbf{K}_{i,j} = e^{-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|_2^2}{2\sigma^2}}$ and for the angular kernel, $\mathbf{K}_{i,j} = 1 - \frac{\theta}{\pi}$ with $\theta = \cos^{-1}\left(\frac{\mathbf{x}_i^T \mathbf{x}_j}{\|\mathbf{x}_i\|_2 \|\mathbf{x}_j\|_2}\right)$. For the approximation, $\tilde{\mathbf{K}}_{i,j} = \frac{1}{\sqrt{d'}} s(\mathbf{A}\mathbf{x}_i)^T \frac{1}{\sqrt{d'}} s(\mathbf{A}\mathbf{x}_j)$ where $s(x) = e^{-\frac{ix}{\sigma}}$ and $\tilde{\mathbf{K}}_{i,j} = 1 - \frac{d_H(s(\mathbf{A}\mathbf{x}_i), s(\mathbf{A}\mathbf{x}_j))}{d'}$ where $s(x) = \text{sign}(x)$ respectively. In both cases, function s is applied pointwise. d_H stands for the Hamming distance and x_i, x_j are points from the dataset.

We used two datasets: G50C (550 points, $n = 50$) and USPST (test set, 2007 points, $n = 256$). The results for the USPST dataset are given in the Supplement. For Gaussian kernel, bandwidth σ is set to 17.4734 for G50C and to 9.4338 for USPST. The choice of σ comes from [Choromanski and Sindhvani, 2016] in order to have comparable results. The results are averaged over 10 runs and the following matrices have been tested: Gaussian random matrix \mathbf{G} , $\mathbf{G}_{\text{circ}} \mathbf{K}_2 \mathbf{K}_1$, $\mathbf{G}_{\text{Toepplitz}} \mathbf{D}_2 \mathbf{HD}_1$, $\mathbf{G}_{\text{skew-circ}} \mathbf{D}_2 \mathbf{HD}_1$, $\mathbf{HD}_{g_1, \dots, g_n} \mathbf{HD}_2 \mathbf{HD}_1$ and $\mathbf{HD}_3 \mathbf{HD}_2 \mathbf{HD}_1$.

Figure 5 shows results for the G50C dataset. In case of G50C dataset, for both kernels, all matrices from the family of structured spinners perform similarly to a random Gaussian matrix. $\mathbf{HD}_3 \mathbf{HD}_2 \mathbf{HD}_1$ performs better than all other matrices for a wide range of sizes of random feature maps. In case of USPST dataset (see: Supplement), for both kernels, all matrices from the family of structured spinners again perform similarly to a random Gaussian matrix (except $\mathbf{G}_{\text{circ}} \mathbf{K}_2 \mathbf{K}_1$ which gives relatively poor results) and $\mathbf{HD}_3 \mathbf{HD}_2 \mathbf{HD}_1$ is giving the best results. Finally, the efficiency of structured spinners does not depend on the dataset.

Table 1 shows substantial speedups obtained by the structured spinner matrices. The speedups are computed as $\text{time}(\mathbf{G})/\text{time}(\mathbf{T})$, where $\text{time}(\mathbf{G})$ and $\text{time}(\mathbf{T})$ are the runtimes for respectively a random Gaussian matrix and a structured spinner matrix.

5.3 Neural networks

Finally, we performed experiments with neural networks using two different network architectures. The first one is a fully-connected network with two fully connected layers (we call it MLP), where we refer to the size of the hidden layer as h , and the second one is a convolutional network with following architecture:

- Convolution layer with filter size 5×5 , 4 feature maps + ReLU + Max Pooling (region 2×2 and step 2×2)
- Convolution layer with filter size 5×5 , 6 feature maps + ReLU + Max Pooling (region 2×2 and step 2×2)
- Fully-connected layer (h outputs) + ReLU
- Fully-connected layer (10 outputs)
- LogSoftMax.

Experiments were performed on the MNIST data set. In both experiments, we re-parametrized each matrix

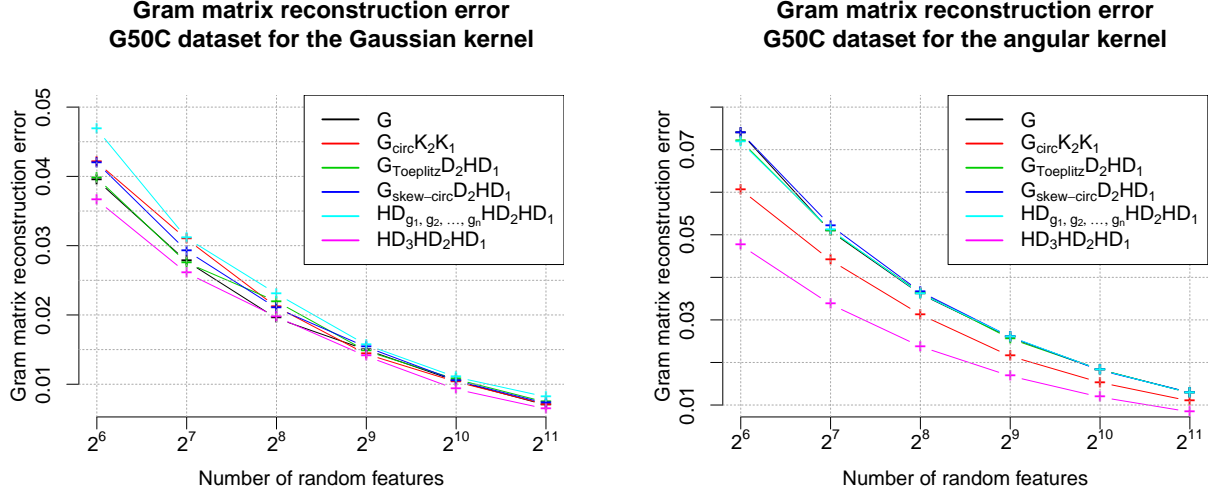


Figure 3: Accuracy of random feature map kernel approximation for the G50C dataset.

MATRIX DIM.	2^9	2^{10}	2^{11}	2^{12}	2^{13}	2^{14}	2^{15}
$G_{Toepplitz} D_2 HD_1$	x1.4	x3.4	x6.4	x12.9	x28.0	x42.3	x89.6
$G_{skew-circ} D_2 HD_1$	x1.5	x3.6	x6.8	x14.9	x31.2	x49.7	x96.5
$HD_{g_1, \dots, g_n} HD_2 HD_1$	x2.3	x6.0	x13.8	x31.5	x75.7	x137.0	x308.8
$HD_3 HD_2 HD_1$	x2.2	x6.0	x14.1	x33.3	x74.3	x140.4	x316.8

Table 1: Speedups for Gaussian kernel approximation via structured spinners.

h	2^4	2^5	2^6	2^7	2^8	2^9	2^{10}	2^{11}	2^{12}
unstructured	42.9	51.9	72.7	99.9	163.9	350.5	716.7	1271.5	2317.4
$HD_3 HD_2 HD_1$	109.2	121.3	109.7	114.2	117.4	123.9	130.6	214.3	389.8

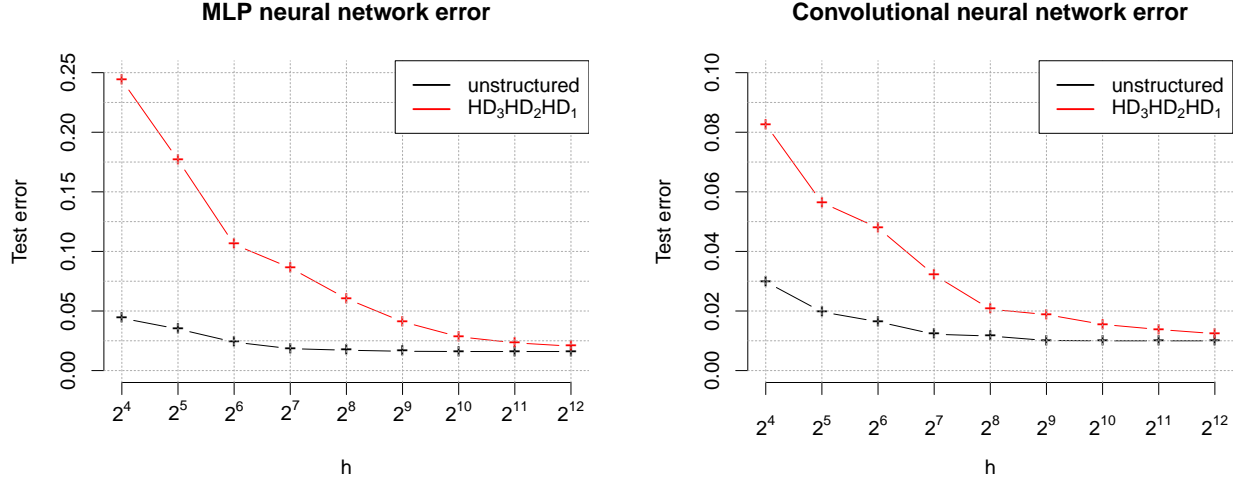
 Table 2: Running time (in $[\mu s]$) for the MLP - unstructured matrices vs structured spinners.


Figure 4: Test error for MLP (top) and convolutional network (bottom).

of weights of fully connected layers with a structured $HD_3 HD_2 HD_1$ matrix from a family of structured spinners. We compare this setting with the case where the unstructured parameter matrix is used. Note that in case when we use $HD_3 HD_2 HD_1$ only linear number of parameters is learned (the Hadamard matrix is deterministic and even does not need to be explicitly stored, instead Walsh-Hadamard transform is used). Thus the network has significantly less parameters than in the unstructured case, e.g. for the MLP network we have $\mathcal{O}(h)$ instead of $\mathcal{O}(\text{input size} \times h)$ parameters.

In Figure 4 and Table 2 we compare respectively the test error and running time of the unstructured and structured approaches. Figure 4 shows that for large enough h , neural networks with structured spinners achieve similar performance to those with unstructured projections, while at the same time using structured spinners lead to significant computational savings as shown in Table 2. As mentioned before, the $HD_3 HD_2 HD_1$ -neural network is a simpler construction than the Deep Friend Convnet, however one can replace it with any structured spinner to obtain compressed neural network architecture of a good capacity.

References

- [Ailon and Chazelle, 2006] Ailon, N. and Chazelle, B. (2006). Approximate nearest neighbors and the fast Johnson-Lindenstrauss transform. In *STOC*.
- [Ailon and Liberty, 2011] Ailon, N. and Liberty, E. (2011). An almost optimal unrestricted fast Johnson-Lindenstrauss transform. In *SODA*.
- [Andoni et al., 2015] Andoni, A., Indyk, P., Laarhoven, T., Razenshteyn, I. P., and Schmidt, L. (2015). Practical and optimal LSH for angular distance. In *NIPS*.
- [Bentkus, 2003] Bentkus, V. (2003). On the dependence of the Berry–Esseen bound on dimension. *Journal of Statistical Planning and Inference*, 113(2):385–402.
- [Brent et al., 2014] Brent, R. P., Osborn, J. H., and Smith, W. D. (2014). Bounds on determinants of perturbed diagonal matrices. *arXiv:1401.7084*.
- [Charikar, 2002] Charikar, M. (2002). Similarity estimation techniques from rounding algorithms. In *STOC*.
- [Cho and Saul, 2009] Cho, Y. and Saul, L. K. (2009). Kernel methods for deep learning. In *NIPS*.
- [Choromanska et al., 2016] Choromanska, A., Choromanski, K., Bojarski, M., Jebara, T., Kumar, S., and LeCun, Y. (2016). Binary embeddings with structured hashed projections. In *ICML*.
- [Choromanski and Sindhvani, 2016] Choromanski, K. and Sindhvani, V. (2016). Recycling randomness with structure for sublinear time kernel expansions. In *ICML*.
- [Dai et al., 2014] Dai, B., Xie, B., He, N., Liang, Y., Raj, A., Balcan, M.-F., and Song, L. (2014). Scalable kernel methods via doubly stochastic gradients. In *NIPS*.
- [Dasgupta et al., 2010] Dasgupta, A., Kumar, R., and Sarlos, T. (2010). A sparse johnson: Lindenstrauss transform.
- [Dasgupta and Freund, 2008] Dasgupta, S. and Freund, Y. (2008). Random projection trees and low dimensional manifolds. In *STOC*.
- [Denil et al., 2013] Denil, M., Shakibi, B., Dinh, L., Ranzato, M., and Freitas, N. D. (2013). Predicting parameters in deep learning. In *NIPS*.
- [Feng et al., 2015] Feng, C., Hu, Q., and Liao, S. (2015). Random feature mapping with signed circulant matrix projection. In *IJCAI*.
- [Goodfellow et al., 2016] Goodfellow, I., Bengio, Y., and Courville, A. (2016). Deep learning. Book in preparation for MIT Press.
- [Har-Peled et al., 2012] Har-Peled, S., Indyk, P., and Motwani, R. (2012). Approximate nearest neighbor: Towards removing the curse of dimensionality. *Theory of Computing*, 8(14):321–350.
- [Hinrichs and Vybrál, 2011] Hinrichs, A. and Vybrál, J. (2011). Johnson-lindenstrauss lemma for circulant matrices. *Random Struct. Algorithms*, 39(3):391–398.
- [Huang et al., 2014] Huang, P.-S., Avron, H., Sainath, T., Sindhvani, V., and Ramabhadran, B. (2014). Kernel methods match deep neural networks on timit. In *ICASSP*.
- [Huhtanen and Perämäki, 2015] Huhtanen, M. and Perämäki, A. (2015). Factoring matrices into the product of circulant and diagonal matrices. *Journal of Fourier Analysis and Applications*, 21(5):1018–1033.
- [Johnson and Lindenstrauss, 1984] Johnson, W. and Lindenstrauss, J. (1984). Extensions of Lipschitz mappings into a Hilbert space. In *Conference in modern analysis and probability*, volume 26 of *Contemporary Mathematics*, pages 189–206.
- [Le et al., 2013] Le, Q., Sarlós, T., and Smola, A. (2013). Fastfood-computing hilbert space expansions in loglinear time. In *ICML*.
- [LeCun et al., 2015] LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521(7553):436–444.
- [Li, 2013] Li, J. (2013). Restructuring of deep neural network acoustic models with singular value decomposition. In *Interspeech*.
- [Liberty et al., 2008] Liberty, E., Ailon, N., and Singer, A. (2008). Dense fast random projections and lean Walsh transforms. In *RANDOM*.
- [Mairal et al., 2014] Mairal, J., Koniusz, P., Harchaoui, Z., and Schmid, C. (2014). Convolutional kernel networks. In *NIPS*.
- [Moczulski et al., 2016] Moczulski, M., Denil, M., Appleyard, J., and de Freitas, N. (2016). Acdc: A structured efficient linear layer. In *ICLR*.
- [Pilanci and Wainwright, 2014] Pilanci, M. and Wainwright, M. J. (2014). Randomized sketches of convex programs with sharp guarantees. In *ISIT*.
- [Pilanci and Wainwright, 2015] Pilanci, M. and Wainwright, M. J. (2015). Newton sketch: A linear-time optimization algorithm with linear-quadratic convergence. *CoRR*, abs/1505.02250.

- [Rahimi and Recht, 2007] Rahimi, A. and Recht, B. (2007). Random features for large-scale kernel machines. In *NIPS*.
- [Sainath et al., 2013] Sainath, T. N., Kingsbury, B., Sindhvani, V., Arisoy, E., and Ramabhadran, B. (2013). Low-rank matrix factorization for deep neural network training with high-dimensional output targets. In *ICASSP*.
- [Sindhvani et al., 2015] Sindhvani, V., Sainath, T. N., and Kumar, S. (2015). Structured transforms for small-footprint deep learning. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 3088–3096.
- [Terasawa and Tanaka, 2007] Terasawa, K. and Tanaka, Y. (2007). Spherical LSH for approximate nearest neighbor search on unit hypersphere. In *WADS*.
- [Vybíral, 2011] Vybíral, J. (2011). A variant of the Johnson-Lindenstrauss lemma for circulant matrices. *Journal of Functional Analysis*, 260(4):1096–1105.
- [Yang et al., 2015] Yang, Z., Moczulski, M., Denil, M., de Freitas, N., Smola, A., Song, L., and Wang, Z. (2015). Deep fried convnets. In *ICCV*.