
Fast geometric learning with symbolic matrices

Jean Feydy*
Imperial College London
jfeidy@ic.ac.uk

Joan Alexis Glaunès*
Université de Paris
alexis.glaunes@parisdescartes.fr

Benjamin Charlier*
Université de Montpellier
benjamin.charlier@umontpellier.fr

Michael M. Bronstein
Imperial College London / Twitter
m.bronstein@imperial.ac.uk

Abstract

Geometric methods rely on tensors that can be encoded using a symbolic formula and data arrays, such as kernel and distance matrices. We present an extension for standard machine learning frameworks that provides comprehensive support for this abstraction on CPUs and GPUs: our toolbox combines a versatile, transparent user interface with fast runtimes and low memory usage. Unlike general purpose acceleration frameworks such as XLA, our library turns generic Python code into binaries whose performances are competitive with state-of-the-art geometric libraries – such as FAISS for nearest neighbor search – with the added benefit of flexibility. We perform an extensive evaluation on a broad class of problems: Gaussian modelling, K-nearest neighbors search, geometric deep learning, non-Euclidean embeddings and optimal transport theory. In practice, for geometric problems that involve 10^3 to 10^6 samples in dimension 1 to 100, our library speeds up baseline GPU implementations by up to two orders of magnitude.

1 Introduction

Fast numerical methods are the fuel of machine learning research. Over the last decade, the sustained development of the CUDA ecosystem has driven the progress in the field: though Python is the lingua franca of data science and machine learning, most frameworks rely on efficient C++ backends to leverage the computing power of GPUs [1, 86, 101]. Recent advances in computer vision or natural language processing attest to the fitness of modern libraries: they stem from the mix of power and flexibility that is provided by PyTorch, TensorFlow and general purpose accelerators such as XLA.

Nevertheless, important work remains to be done. Geometric computations present a clear gap in performances between Python and C++: notable examples are implementations of point cloud convolutions or of the nearest neighbor search [65, 76]. To scale up geometric computations to real-world data, a common practice is therefore to replace the compute-intensive parts of a Python code by handcrafted CUDA kernels [35, 60, 92]. These are expensive to develop and maintain, which leads to an unfortunate need to compromise between ease of development and scalability.

To address this issue, we present KeOps: an extension for PyTorch, NumPy, Matlab and R that combines the speed of a handcrafted CUDA kernel with the simplicity of a high level language. Our toolbox optimizes map-reduce operations on generalized point clouds and provides transparent support for distance-like matrices, as illustrated in Figure 1. The resulting computations are fully differentiable and have a negligible memory footprint. Their runtimes are competitive with state-of-the-art CUDA libraries when they exist, and peerless in the many use cases that are not covered by existing implementations. Our library fits seamlessly within existing codebases and provides a sizeable performance boost to a wide range of methods. Among other applications, we present optimal transport solvers and geometric operators in hyperbolic spaces which are orders of magnitude faster than the state-of-the-art. We believe that our library is an important addition to the existing arsenal of tools and will have a stimulating impact on machine learning research.




Figure 1: Machine learning frameworks understand variables as matrices, also known as tensors. (a) These are usually **dense** and encoded as explicit numerical arrays $(M_{i,j}) = (M[i, j]) \in \mathbb{R}^{N \times M}$ that can have a large memory footprint. (b) Alternatively, some operators can be encoded as **sparse matrices**: we store in memory the indices (i_n, j_n) and values $M_n = M_{i_n, j_n}$ that correspond to a small number of non-zero coefficients. Reduction operations are then implemented using indexing methods and scattered memory accesses. (c) We provide support for a third class of tensors: **symbolic matrices** whose coefficients are given by a formula $M_{i,j} = F(x_i, y_j)$ that is evaluated on data arrays (x_i) and (y_j) . Reduction operations are implemented using parallel schemes that compute the coefficients $M_{i,j}$ on-the-fly. We take advantage of the structure of CUDA registers to bypass costly memory transfers and achieve optimal runtimes on a wide range of applications.

2 Related work

Machine learning and data science applications often encounter the problem of computing the proximity or distance between data samples. Given x_1, \dots, x_N and $y_1, \dots, y_M \in \mathbb{R}^D$ two clouds of N and M points in dimension D , the bottleneck of many methods is an **interaction step** of the form:

$$a_i \leftarrow \square_{j=1}^M F(i, j, x_i, y_j), \quad \forall i \in \llbracket 1, N \rrbracket, \quad (1)$$

where F is a vector-valued formula and \square is an associative reduction operator, e.g. a sum or a minimum. This paper is part of a large body of work that lowers the $\mathcal{O}(NM)$ computational cost of such an operation: we now recall the main approaches to this problem.

Sparse matrices. A first strategy is to prune out negligible terms: for every index i , we perform the reduction (1) on a subset of neighbors $\mathcal{N}(i) \subset \llbracket 1, M \rrbracket$. As illustrated in Figure 1, this method is akin to using sparse matrices: the neighborhood structure is usually understood as a connectivity matrix that comes from a triangle mesh or a K -nearest neighbors (KNN) graph [16, 69, 114]. This method can be used whenever the operation F is local but has a major limitation: at a low level, truncated reductions rely on random memory accesses that do not stream well on GPUs [25, 82]. Consequently, speed-ups are only achieved if the neighborhoods $\mathcal{N}(i)$ are orders of magnitude smaller than the full set of indices $\llbracket 1, M \rrbracket$ – a condition that is often too restrictive and cannot be satisfied.

Nearest neighbor finders. Going further, the implementation of KNN queries is itself a geometric problem in the mould of (1). When the datasets (x_i) and (y_j) have a small intrinsic dimension, efficient schemes can outperform brute-force approaches [11, 31, 49, 53, 65, 77, 85]. Unfortunately, these methods rely on pre-computations that are too expensive to be performed at every iteration of a training loop. Reference implementations also tend to lack flexibility and only support a handful of metrics: for instance, in spite of a strong interest for hyperbolic embeddings in the machine learning literature [72, 83], Poincaré metrics are not supported out-of-the-box by standard libraries.

Approximated convolutions. When the reduction \square is a sum and $F(i, j, x_i, y_j) = k(x_i - y_j) = K_{i,j}$ is a translation-invariant kernel, the interaction (1) is understood as a discrete convolution. To speed up this operation, a first idea is to rely on low-rank decompositions of the kernel matrix $(K_{i,j})$ [90, 111, 116]. Multiscale schemes can be used to handle singular kernels [7, 8, 50, 52, 115] or compress generic operators [5, 17, 55]. Alternatively, semi-Eulerian methods rely on intermediate

grid representations to leverage fast Fourier transforms or convolution routines [32, 51, 76]. These approaches can achieve dramatic speed-ups but tend to require a significant amount of tuning for each kernel k . They work best when the latter is smooth or is defined on a space of dimension $D \leq 3$.

Acceleration frameworks. In contrast to mathematical approaches, several compilation frameworks have been designed to speed-up machine learning architectures. Modern toolboxes accelerate a wide range of operations but are not geared towards geometric problems: most of them keep a focus on distributed learning [63, 64, 97, 108] or image processing and dense tensor manipulations [22, 58, 74, 105]. TVM [22] and CuPy [84] are the two libraries which are closer to our work: they both provide partial support for symbolic tensors. However, they have limited support for automatic differentiation and require the use of a custom low-level syntax to produce efficient binaries.

3 Motivation

Requirements for geometric data analysis and learning. None of the aforementioned methods are fully suited for modern research in geometric data analysis and machine learning. Let us briefly explain why. First of all, some acceleration schemes do not stream well on GPUs or have to rely on expensive pre-computations: hierarchical matrices [55] or advanced nearest neighbor finders [77] can hardly be used in the training loop of a neural network. Other strategies make strong assumptions on the properties of the convolution filter k or on the dimension and geometry of the ambient feature space. These restrictions make existing tools cumbersome to use in deep learning, where one wishes to have modelling freedom, e.g. w.r.t. the choice of the embedding space geometry and dimension. Finally, most acceleration frameworks for Python expect users to be knowledgeable on GPU parallelism or do not support automatic differentiation. The bottomline is that most existing tools are not ready to be used by a majority of researchers in the community.

A gap in the literature. In order to tackle these issues, the developers of deep learning libraries have recently put an emphasis on just-in-time compilation for neural networks. For instance, the recent PyTorch JIT and XLA engines enable operator fusion and unlock performance speed-ups for research code [15, 86]. These *general purpose* compilers are fully transparent to users and show promise for a wide range of applications. Nevertheless, they fall short on geometric computations along the lines of (1). This is most apparent for nearest neighbor search [36, 60, 65], matrix-vector products with kernel matrices and message passing methods on point clouds [35, 36, 102], where one still has to develop and maintain custom CUDA kernels to achieve state-of-the-art performance.

Symbolic matrices. We notice that all the aforementioned methods rely on reductions of an N -by- M matrix $(M_{i,j}) = (F(i, j, x_i, y_j))$ that is often **too large to be stored in memory as a dense tensor**. Acknowledging the fact that memory management is a bottleneck for tensor programs, we choose to focus on the fundamental concept of symbolic matrices, illustrated in Figure 1. For the first time, we provide support for this abstraction on the GPU with all the desirable features of a deep learning library: a math-friendly interface, high performance, transparent support for batch processing and automatic differentiation. The example below is representative of our user interface:

```

1  from torch import rand, autograd      # NumPy, R and Matlab are also supported
2  from pykeops.torch import LazyTensor  # Symbolic wrapper for PyTorch Tensors
3
4  # Setup data on the CPU and/or GPU with shapes (N,D), (M,D), (M,E):
5  N, M, D, E = 10 ** 5, 10 ** 6, 50, 100
6  x, y, b = rand(N, D, requires_grad=True), rand(M, D), rand(M, E)
7
8  # Perform arbitrary symbolic computations:
9  x_i = LazyTensor(x.view(N, 1, D))     # (N,D) Tensor -> (N,1,D) Symbolic Tensor
10 y_j = LazyTensor(y.view(1, M, D))    # (M,D) Tensor -> (1,M,D) Symbolic Tensor
11 D_ij = ((x_i - y_j) ** 2).sum(dim=2)  # (N,M) Symbolic matrix of squared distances.
12 K_ij = (- D_ij).exp()                 # (N,M) Symbolic Gaussian kernel matrix.
13
14 # Come back to genuine torch Tensors with reductions on dimensions 0 and 1:
15 nn = D_ij.argmax(K=10, dim=1)        # K-NN search: (N,10) array of integer indices.
16 a = K_ij @ b                          # Kernel matrix-vector product: (N,M) * (M,E) = (N,E)
17 [g_x] = autograd.grad((a ** 2).sum(), [x]) # Seamless backpropagation.

```

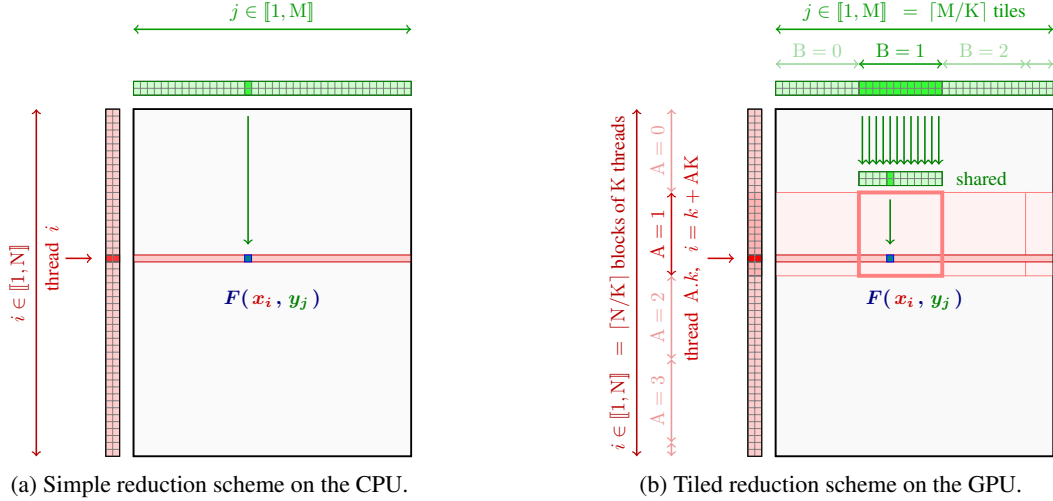


Figure 2: We rely on fast parallel schemes to compute reductions of symbolic matrices, as in Eq. (1). (a) On the CPU, each thread i computes a value a_i by looping over the reduction index j and consuming the values of F **on-the-fly**. (b) On the GPU, we cut (a) in K -by- K tiles (where K is the CUDA block size) to leverage the low latency of the shared memory buffer and block-wise memory accesses. This ensures an optimal management of the y_j 's: we refer to our online documentation (www.kernel-operations.io) and to the N-body simulation chapter of [82] for details.

4 Implementation

Creating a symbolic tensor. The entry point to our library is a `LazyTensor` wrapper that turns dense arrays into symbolic matrices (lines 9-10). We use standard operations to build up arbitrary formulas F : lines 11-12 update a symbolic representation of the matrices D_{ij} and K_{ij} . Our math engine is **lazy**, and defers the evaluation of formulae until reduction time. It is also **versatile**: we support batch dimensions, operator broadcasting and a wide range of elementary operations. Formulae can use an arbitrary number of variables x_i^1, x_i^2, \dots and y_j^1, y_j^2, \dots with varied dimensions.

Parallel reduction schemes. Numerical computations take place at the reduction time (lines 15-16), when an associative operator \square is called to perform an interaction step (1). To this end, we use a tiled map-reduce scheme that is detailed in Figure 2. Our C++ engine optimizes the use of registers to avoid costly memory transfers: the algorithm has $\mathcal{O}(NM)$ time complexity but does not allocate any buffer in the global device memory. For the sake of performance, we compile a specific binary for every new formula-reduction pair (F, \square) . These are stored in a cache directory for later use: the compilation is only done once. Unlike most acceleration frameworks, we do not expect the sizes of the data arrays $(x_i) \in \mathbb{R}^{N \times D}$ and $(y_j) \in \mathbb{R}^{M \times D}$ to be fully known at compile time: our binaries only rely on the feature dimensions D . This allows us to work with datasets of varying sizes, without having to re-compile binaries or sub-sample point clouds to an arbitrary point count. Notably, this allows our engine to process raw shape data on-the-fly.

Pruning negligible interactions. Our library provides support for alternative reduction strategies that improve numerical accuracy [67] or increase GPU usage when $N < M$. Going further, our symbolic tensors support the specification of block-wise sparsity masks as optional attributes. Block reduction tiles are encoded using a collection of $(i_{start}, i_{end}, j_{start}, j_{end})$ tuples of indices that allow our engine to focus on a subset of the full collection of interaction pairs $\llbracket 1, N \rrbracket \times \llbracket 1, M \rrbracket$. They can be used to prune out negligible terms from symbolic reductions, without giving up on the contiguous memory accesses that make or break the performance of CUDA kernels.

Among high-level computing frameworks, this feature is unique to our library. Block-diagonal sparsity masks allow us to provide seamless support for batch processing, even in heterogeneous situations that are of interest for point cloud and mesh processing. Assuming that the data arrays have been pre-sorted in contiguous clusters, block-sparsity masks also enable the GPU implementation of fast multiscale methods such as the Barnes-Hut algorithm [7].

Strengths and limitations. At its heart, KeOps leverages the low Kolmogorov complexity of symbolic arrays: it can be used when the computational bottleneck of a method is an interaction step of the form (1). As we show next in the benchmarks of Section 5, KeOps is likely to offer gains on runtime and memory usage when the numbers of samples N and M range from 10^3 to 10^7 .

The main limitation of KeOps stems from the overflow of CUDA registers in the reduction of Figure 2: these result in decreased performances on large feature vectors with $D > 100$. The problem is known as *register spilling*, with some documented work-arounds [25, 82]. Our priority for future developments is to improve performance for problems with $D \sim 1k$.

Another drawback is that we do not pre-ship binaries but instead rely on C++/CUDA compilers to run our kernels. To mitigate deployment issues and ease maintenance in the long run, we implement the core of the library in C++ and keep dependencies to a minimum [61]. In practice, compilation times range from 10 to 25 seconds when $D \leq 1,000$ and can be prohibitive for $D \geq 2,000$. Our library runs out-of-the-box on Linux or Mac configurations that provide a CUDA environment, e.g. fresh Google Colab sessions.

5 Experiments and Applications

Configuration. We now showcase our toolbox on a wide range of machine learning problems. All benchmarks were performed on a workstation equipped with 8 Intel Xeon Gold 6142 CPU @ 2.60GHz cores (16 threads), 128Gb of RAM and a Nvidia RTX 2080 Ti GPU with 11Gb of device memory. When relevant, we include comparisons with PyTorch JIT and JAX/XLA: just like our library, these two frameworks offer a transparent user interface and high performance on GPU. At the Python level, the implementations tested are rigorously equivalent to each other: we implement the same computations and only have to account for minor syntactic divergences. Whenever possible, we work with batches of samples and shapes to improve GPU usage.

Notations, datasets. In tables, “mem” stands for an out-of-memory error, “ ∞ ” for a time that was too high to be recorded and “—” denotes a lack of available implementation. L^1 and L^2 denote the Manhattan and Euclidean metrics, respectively, $\langle \cdot, \cdot \rangle$ denotes the cosine similarity and \mathbb{H}^D is the Poincaré metric on a hyperbolic space of dimension D [18, 21]. We perform numerical experiments with random normal samples and freely available datasets: digits from Scikit-Learn [87], Stanford dragon [26], ShapeNet [19], MNIST [73], SIFT [62], GloVe-25 and GloVe-100 [88] were taken from the ANN-benchmarks repository [4], while HyperE-10 and HyperE-50 are hyperbolic embeddings processed from WordNet datasets [95].

5.1 Kernel Methods and Clustering

Kernel methods. Accelerating kernel methods is one of the core strengths of KeOps [20]. The code snippet of Section 3 shows how to perform a Gaussian convolution on point clouds and can be adapted to any other standard kernel. As detailed in the Supplementary Materials, symbolic `LazyTensors` can be transparently interfaced with the standard iterative solvers of the `scipy.sparse.linalg` library [66, 104]. This allows us to solve large-scale Gaussian regression problems ($N > 100k-1M$) in seconds on a single GPU [57, 78]. Going forward, KeOps provides solid numerical foundations for recent methods in the field that rely on kernel matrices [45, 80].

Clustering. As detailed in the Supplementary Materials, KeOps can be used to implement K-means clustering with several metrics. Going further, we use our symbolic `LazyTensors` to implement the standard Expectation Maximisation (EM) algorithm on a Gaussian Mixture Model (GMM) [56]: we estimate the weight, mean and covariance matrix of each of the K components. To this end, the EM algorithm strives to maximize the likelihood that is computed from a data sample of N points in \mathbb{R}^D . The most costly computations involve sum reductions of a K -by- N `LazyTensor` that encodes interactions between the K means of the clusters and the N sample points. Our results are summarized in Table 1. We show that KeOps performs well when computations involve variables of modest size: it can scale to large datasets in seconds, without memory overflows. This is most apparent when covariance matrices are diagonal and encoded as vectors of size D . In the second half of the table, we handle full D -by- D covariance matrices and notice a slow-down when $D \geq 10$.

5.2 Nearest neighbor search with any metric

Syntax for a K-nearest neighbors query. Our library enables the efficient implementation of brute-force KNN finders for any metric that is known in closed form. Users can apply the `argKmin(K=...)` reduction to extract the indices for the K-smallest values along the lines of any symbolic tensor:

```
1 # Turn (N,D) and (M,D) dense arrays into symbolic tensors:
2 x_i, y_j = LazyTensor(x.view(N,1,D)), LazyTensor(y.view(1,M,D)) # (N,1,D), (1,M,D)
3
4 # Compute the (N,M) symbolic matrix of distances:
5 E_ij = ((x_i - y_j) ** 2).sum(dim=2) # Squared Euclidean metric
6 M_ij = (x_i - y_j).abs().sum(dim=2) # Manhattan distance
7 C_ij = 1 - (x_i | y_j) # Cosine similarity
8 H_ij = E_ij / (x_i[...0] * y_j[...0]) # Hyperbolic metric on the half-space
9
10 # Perform a K-NN search - in this case, for the hyperbolic metric:
11 indices = H_ij.argKmin(K=10, dim=1) # Dense (N,K) array of integer indices
```

In the example above, the hyperbolic metric is defined on the half-space model of dimension D , with a positive first coordinate $x[0] > 0$ for feature vectors $x \in \mathbb{R}^D$ [18]. The mapping $x \mapsto \operatorname{arcosh}(1 + x/2)$ is increasing, which allows us to work with the pseudo-distance $H(x_i, y_j) = \|x_i - y_j\|^2 / (x_i[0]y_j[0])$: similar acceleration tricks can be applied to e.g. the Poincaré ball metric.

Exact and approximate nearest neighbor search. The complexity of a KNN query depends on the number of points N and features D that are available for a given dataset: in the literature, two types of strategies have been designed to handle different scenarios. On the one hand, **exact** brute-force schemes compute all pairwise distances between query and data points. These methods stream well on GPUs, have little to no parameters to tune and require no pre-processing of the data. On the other hand, **approximate** schemes leverage the structure of the dataset to prune out useless computations whenever possible. These methods are usually implemented on CPUs and tend to rely on hierarchical decompositions of the dataset that must be pre-computed ahead of the first KNN query. As discussed in Sections 2 and 3, approximate methods are tailored for large-scale retrieval whereas brute-force schemes are generally more suited to geometric scenarios.

Benchmarks. To illustrate this behaviour, we compare our brute-force reduction scheme to a selection of common baselines: brute-force PyTorch and JAX implementations on the GPU; the popular FAISS library [65], with an approximate HNSW algorithm on the CPU [77] and two dedicated CUDA implementations on the GPU: an exact brute-force scheme and the approximate IVF-Flat method. To showcase the variety of settings in which KNN queries are used throughout the literature, we run these methods on a collection of *random* normal samples and *structured* datasets. We use $N = 10^4$ to 10^7 points with $D = 3$ to 784 features: these lower and upper bounds let us represent the range of problems that are encountered in geometric data analysis, from shape processing to the manipulation of word embeddings. We stress that for larger datasets, approximate methods would clearly outperform our brute-force approach.

Results are displayed in Table 2. A first observation is that on these medium-sized problems, our brute-force GPU implementation is faster than approximate CPU methods to build a full KNN graph: even when FAISS-HNSW outperforms KeOps in terms of queries per second, this comes at the cost of a significant pre-processing time. A second remark is that our **generic** reduction engine for brute-force computations is on par with the hand-crafted CUDA routines of the FAISS library: KeOps is less efficient than the FAISS-Brute-force routine when $D \geq 50$ but is up to $\times 3$ faster in smaller dimensions. Crucially, we also provide the only competitive routines for non-Euclidean geometries, such as the increasingly popular hyperbolic metrics [83].

Conclusion. Overall, the performances of our C++ engine hold up to scrutiny: KeOps provides respectable runtimes on small and medium-sized problems (up to $N = 10^6$ and $D = 100$) with the added benefit of **flexibility**. These results attest to the effectiveness of our optimization techniques and bode well for the other computations that are supported by the KeOps engine. Going forward, we note that our library could probably be used as an efficient GPU backend for approximate KNN methods [31, 54, 118] and intend to provide the relevant tools for researchers in the field.

Table 1: Fitting a Gaussian Mixture Model: we perform 10 iterations of the standard EM algorithm with N points and K components in dimension D.

Covariances	N	K	D	Sklearn	PyTorch	Ours
Diagonal	50k	100	5	2.9 s	19.0 ms	33 ms
Diagonal	500k	1k	5	286 s	0.94 s	0.22 s
Diagonal	5M	10k	5	∞	mem	11.38 s
Diagonal	500k	1k	50	∞	mem	2.96 s
Diagonal	5M	10k	50	∞	mem	245 s
Full	50k	100	5	6.9 s	0.23 s	0.04 s
Full	500k	1k	5	830 s	mem	0.362s
Full	50k	100	20	16.0 s	mem	0.84 s
Full	500k	1k	20	∞	mem	63 s

Table 2: KNN search: average queries per second with a dataset of N points in dimension D. We work with batches of 10k queries at a time and $K = 10$ neighbors. The first three columns correspond to schemes that are provided by the FAISS library; the approximate methods HNSW and IVF-Flat are tuned to provide a minimum recall of 90% and optimize runtimes (we use the parameters of the ANN-benchmarks website as a first reference); when relevant, the pre-processing time is reported in parenthesis. We stress that choosing optimal parameters for the HNSW and IVF-Flat routines is a fairly complex problem: as non-specialists, we cannot guarantee that our experiments reflect the best level of performance that can be reached by these impressive methods. This is especially true for the IVF-PQ routine that is provided by FAISS and combines the IVF algorithm with a quantization method: it certainly fares even better than IVF-Flat on large-scale problems, but we found it to be very complex to use and opted to not include our unreliable benchmarks in this Table. High number of queries and low pre-processing is better: we highlight the column with the fastest time to process N queries and thus build a KNN graph for the dataset. (*) performed with smaller batch sizes when necessary to avoid memory overflows

Dataset, metric	N	D	HNSW (CPU)	Bruteforce	IVF-Flat	PyTorch*	JAX*	Ours
Random, L^2	10k	3	1.5e6 (0.15s)	3.3e6	1.9e6	8.8e5	1.6e5	6.8e6
Random, L^2	1M	3	1.3e6 (25s)	5.0e4	1.6e6	6.8e3	5.5e2	1.8e5
Random, L^2	1M	10	3.1e5 (53s)	4.6e4	2.9e5	4.9e3	4.9e2	1.1e5
Random, L^2	1M	100	1.5e3 (540s)	3.1e4	mem	9.0e2	4.0e2	1.4e4
Random, L^2	10M	100	(∞)	3.2e3	mem	∞	∞	1.4e3
Random, L^1	1M	10	—	—	—	1.7e3	4.6e2	1.1e5
Random, L^1	1M	100	—	—	—	2.1e2	3.0e2	1.5e4
MNIST, L^2	60k	784	5.4e4 (14s)	1.5e5	2.2e5	4.7e4	3.4e3	2.5e4
MNIST, L^1	60k	784	—	—	—	5.7e2	2.0e3	2.5e4
GloVe-25, \langle, \rangle	1.2M	25	7.7e4 (130s)	4.2e4	1.6e5	2.7e3	4.1e2	4.8e4
GloVe-100, \langle, \rangle	1.2M	100	8.6e3 (480s)	2.6e4	3.9e4	6.8e2	3.3e2	1.3e4
Random, \langle, \rangle	1M	10	1.8e5 (77s)	4.6e4	2.7e5	5.2e3	5.3e2	1.5e5
Random, \mathbb{H}^D	1M	10	—	—	—	3.2e3	5.1e2	7.1e4

Table 3: Accelerating geometric deep learning architectures: shown is average training / inference time per shape on the ShapeNet dataset, with clouds of $N = 2,048$ points in dimension $D = 3$.

	PointCNN \rightarrow Ours	DGCNN \rightarrow Ours
training	254 ms \rightarrow 128 ms	170 ms \rightarrow 80 ms
inference	172 ms \rightarrow 43 ms	109 ms \rightarrow 20 ms

5.3 Geometric deep learning and geometric primitives

Geometric deep learning. Fast low-dimensional KNN search has important applications in the field of geometric deep learning [16], a popular class of algorithms now ubiquitous in 3D computer vision and graphics [110]. As a first illustration, we use KeOps to speed-up two popular networks for 3D point cloud segmentation: Point CNNs [75] and Dynamic Graph CNNs (DGCNN) [110]. We follow closely the original architectures for part segmentation on the ShapeNet dataset [19] and compare two different implementations: a reference PyTorch_Geometric code [36] and a hybrid PyTorch_Geometric+KeOps implementation, where KNN graphs are built as in Section 5.2. Results are summarized in Table 3: training denotes a full “forward + backward” pass through the network whereas inference is forward only, with a faster batch normalisation [59]. Note that the Deep Graph Library [109] relies on a bruteforce PyTorch implementation for KNN search and behaves like PyTorch_Geometric on these problems. In practice, the switch to KeOps provides $\times 2$ and $\times 5$ speedups for training and inference times, respectively: the construction of KNN graphs becomes a negligible overhead, while the majority of the network runtime is spent on MLP and batch normalization layers.

Geometric descriptors at all scales. Table 4 shows the results of using KeOps to compute geometric features on generalized point clouds. In the interaction step (1), these correspond to the case where the reduction \square is a (possibly weighted) average on a neighborhood of x_i and the formula F is a function of the difference $x_i - y_j \in \mathbb{R}^D$: the identity $x \mapsto x$, an outer product $x \mapsto xx^T$ or a multi-layer perceptron with H hidden neurons and O output features. Neighborhoods are either defined through a 40-nearest neighbors query or by weighting each pair of points with a window of radius σ such as:

$$k(x_i, y_j) = \exp(-\|x_i - y_j\|^2 / 2\sigma^2). \tag{2}$$

Our library is well suited for these computations and consistently outperforms the PyTorch and JAX baselines by up to two orders of magnitude. We remark that on clouds of $N = 2,048$ points, our bruteforce scheme (set radius) is up to an order of magnitude faster than the best sparse methods, which rely on scattered memory accesses to build neighborhoods as (N, K, D) arrays. A similar behavior is observed for chamfer and Energy distance computations [14, 99]. These results highlight the dichotomy between *contiguous* (“bruteforce”) and *scattered* (“sparse”) memory accesses, which are optimized separately by GPUs and are best suited to different types of computations. We stress that KeOps supports the specification of block-wise sparsity schemes, which let users implement tree-based pruning strategies to best leverage the structure of their problems.

Conclusion. Overall, our library enables the quick development of differentiable layers for geometric deep learning which are an order of magnitude faster than Python baselines. It factors out C++ boilerplate code for point cloud processing and lets researchers focus on their *models*. We thus believe that KeOps will be of utmost interest to the developers of shape analysis frameworks [12, 36, 60, 92, 103] and CUDA kernels for point cloud convolution [35, 102, 119]. Future developments may also be relevant to natural language processing: transformer architectures and attention layers fit a similar design pattern, albeit in higher-dimensional spaces [100, 106, 113].

Table 4: Accelerating geometric primitives: average time per shape on the ShapeNet dataset, with clouds of $N = 2,048$ points in dimension $D = 3$.

Primitive / Neighborhood	PyTorch	JAX		Ours	
	40-NN	40-NN	Set radius	40-NN	Set radius
Local mean vectors	686 μ s	1,052 μ s	469 μ s	121 μ s	12 μs
Local covariance matrices	700 μ s	1,093 μ s	1,259 μ s	138 μ s	23 μs
MLP features (H = O = 8)	737 μ s	1,180 μ s	4,089 μ s	192 μ s	75 μs
MLP features (H = O = 16)	775 μ s	1,253 μ s	7,043 μ s	240 μs	649 μ s
Chamfer loss	374 μ s		130 μ s		21 μs
Energy distance	486 μ s		378 μ s		31 μs

5.4 Easing the development of complex geometric programs

Dimension reduction and geometric embeddings. The uniform manifold approximation and projection (UMAP) algorithm [79] is a standard method for dimension reduction. It can be used with various input metrics and relies on the analysis of a KNN graph of the input data. To showcase the benefits of KeOps for data analysis, we benchmark three different implementations of this method: the reference CPU-only library [79]; CuML, a fast GPU implementation that relies on FAISS for nearest neighbor queries [65, 91]; and a custom CuML+KeOps implementation. As typical examples of UMAP for visualization, we embed several datasets in the Euclidean plane: the digits, SIFT and MNIST datasets (endowed with the Euclidean and Manhattan metrics); the Glove-25 dataset (cosine similarity); and the HyperE-10 and -50 embeddings (hyperbolic metric).

Since the initial construction of the KNN graph is the most compute-intensive part of the UMAP algorithm, results are similar to those of Section 5.2 (a full table of results is provided in the Supplementary Materials). KeOps provides a $\times 2-5$ speed-up for the visualization of low-dimensional datasets, but is outperformed by CuML+FAISS when the input dimension D exceeds 50. Crucially, we provide the only GPU method that can handle non-Euclidean metrics: for hyperbolic embeddings, our implementation is $\times 200$ faster than the baseline on the HyperE-10 and -50 datasets [72, 83].

Optimal transport. Optimal transport (OT) is a generalization of sorting to spaces of dimension $D > 1$ [81, 107]. It revolves around the resolution of a linear optimization problem whose value is usually known as the Wasserstein or Earth Mover’s distance between two point clouds (x_i) and (y_j) [68, 93]. As detailed in [37, 89], modern OT solvers rely on iterated reductions performed on a cost matrix $C(x_i, y_j) = \|x_i - y_j\|$ or $\frac{1}{2}\|x_i - y_j\|^2$ and stand to benefit greatly from our library. To demonstrate this, we benchmark an exact linear solver [13, 42] against several variants and implementations of the Sinkhorn algorithm: a vanilla Sinkhorn loop [27, 29, 34, 41, 44, 94, 98, 112, 117]; an accelerated algorithm with simulated annealing [71]; and a fast multiscale scheme with kernel truncation [10, 96]. For the sake of numerical stability, we use symmetric updates [70] and perform all computations in the logarithmic domain [23]. To reflect the varied use cases of OT in the machine learning literature, we tackle two different problems: a high-precision matching between deformations of the Stanford dragon in dimension $D = 3$, and a low-precision matching between deformations of the Glove-25 dataset in dimension $D = 25$. These two regimes correspond to typical applications in computer vision [24, 38] and statistics [43, 46, 48], respectively.

A table of results is provided in the Supplementary Materials. In all settings, the switch from a PyTorch implementation to KeOps provides a $\times 5 - 20$ speed-up with no loss of precision. Most importantly, the low memory footprint of symbolic `LazyTensors` allows us to scale up to large problems ($N, M > 100k$) without memory overflows. Our support for block-wise sparsity masks also lets us provide the first implementation of a multiscale solver for discrete OT on the GPU. In practice, this means that large OT problems with $N = 100k$ or $1M$ samples can now be solved with high precision in fractions of a second instead of minutes or hours [37]. This implementation is straightforward to generalize to stochastic settings [2, 3, 47] and non-Euclidean cost functions [9, 28, 30]: we believe that it will open new ranges of applications to the OT community [37, 39].

6 Conclusion

KeOps combines a user-friendly interface with a high level of performance: we believe that it fills an important niche in the machine learning toolbox. We hope that our library will stimulate research in the field, with a simple but powerful structure that makes it a good tool for research off the beaten track. We look forward to feedback from users and keep the door open for contributors.

In months to come, our priority will be to improve performances on high-dimensional vectors with the newly released GPU Tensor cores and add support for quantization or approximation schemes such as the Nyström and FFM methods [5, 111]. We also work towards easing the deployment of pre-compiled binaries and target support of the ONNX standard [6]. These improvements should allow KeOps to become a standard toolbox in the field, both as an efficient backend for high-level software [12, 40, 45, 80] and as a versatile prototyping tool for theoretical research.

Broader Impact

Our work targets a wide range of machine learning applications, from kernel methods to geometric deep learning. In these fields, our library lowers the barrier of entry to state-of-the-art performances: fast nearest neighbors queries or point cloud convolutions can now be implemented by researchers who have no background in parallel computing. We hope that this will empower small research teams and organizations who don't have access to dedicated teams of software engineers. More specifically, the flexibility of our library is ideally suited to the formulation of data-driven models for shape analysis and point cloud processing. Progress in these sectors can have a major impact in computer vision and medical imaging – topics that carry both risks and promises for society as a whole. We hope that our library will promote the growth of a diverse ecosystem of academic and industrial actors, and look forward to seeing applications of our work to e.g. computational anatomy.

Acknowledgments and Disclosure of Funding

The three first authors are the project leaders: they contributed equally to the library and its documentation. Michael Bronstein is supported by the ERC Consolidator grant No. 724228 (LEMAN). KeOps was first motivated by applications to computational anatomy, in collaboration with Alain Trouvé. We also thank Ghislain Durif, who develops the R bindings of KeOps, as well as François-David Collin, who helps us to maintain our hardware and testing infrastructure.

References

- [1] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 265–283, 2016.
- [2] J. Altschuler, F. Bach, A. Rudi, and J. Weed. Massively scalable Sinkhorn distances via the Nyström method. *arXiv preprint arXiv:1812.05189*, 2018.
- [3] J. Altschuler, J. Niles-Weed, and P. Rigollet. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. In *Advances in Neural Information Processing Systems*, pages 1964–1974, 2017.
- [4] M. Aumüller, E. Bernhardsson, and A. Faithfull. ANN-Benchmarks: A benchmarking tool for approximate nearest neighbor algorithms. *Information Systems*, 87:101374, 2020.
- [5] M. Aussal and M. Bakry. The Fast and Free Memory method for the efficient computation of convolution kernels. *arXiv preprint arXiv:1909.05600*, 2019.
- [6] J. Bai, F. Lu, K. Zhang, et al. ONNX: Open neural network exchange. <https://github.com/onnx/onnx>, 2019.
- [7] J. Barnes and P. Hut. A hierarchical $O(N \log N)$ force-calculation algorithm. *Nature*, 324(6096):446, 1986.
- [8] R. Beatson and L. Greengard. A short course on fast multipole methods. *Wavelets, multilevel methods and elliptic PDEs*, 1:1–37, 1997.
- [9] J.-D. Benamou, W. L. Ijzerman, and G. Rukhaia. An entropic optimal transport numerical approach to the reflector problem. working paper or preprint, Apr. 2020.
- [10] J.-D. Benamou and M. Martinet. Capacity constrained entropic optimal transport, Sinkhorn saturated domain out-summation and vanishing temperature. working paper or preprint, May 2020.
- [11] N. Bhatia et al. Survey of nearest neighbor techniques. *arXiv preprint arXiv:1007.0085*, 2010.
- [12] A. Bône, M. Louis, B. Martin, and S. Durrleman. Deformetrica 4: an open-source software for statistical shape analysis. In *International Workshop on Shape in Medical Imaging*, pages 3–13. Springer, 2018.
- [13] N. Bonneel, M. Van De Panne, S. Paris, and W. Heidrich. Displacement interpolation using Lagrangian mass transport. In *Proceedings of the 2011 SIGGRAPH Asia Conference*, pages 1–12, 2011.
- [14] G. Borgefors. Distance transformations in arbitrary dimensions. *Computer vision, graphics, and image processing*, 27(3):321–345, 1984.

- [15] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, and S. Wanderman-Milne. JAX: composable transformations of Python+NumPy programs. *v0.1.55*, 2018.
- [16] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: going beyond Euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.
- [17] L. Cambier and E. Darve. Fast low-rank kernel matrix factorization through skeletonized interpolation. *SIAM Journal on Scientific Computing*, 41(3):A1652–A1680, 2019.
- [18] J. W. Cannon, W. J. Floyd, R. Kenyon, W. R. Parry, et al. Hyperbolic geometry. *Flavors of geometry*, 31:59–115, 1997.
- [19] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- [20] B. Charlier, J. Feydy, J. A. Glaunès, F.-D. Collin, and G. Durif. Kernel operations on the GPU, with autodiff, without memory overflows. *arXiv preprint arXiv:2004.11127*, 2020.
- [21] É. Charpentier, E. Ghys, and A. Lesne. *The scientific legacy of Poincaré*, volume 36. American Mathematical Soc., 2010.
- [22] T. Chen, T. Moreau, Z. Jiang, L. Zheng, E. Yan, H. Shen, M. Cowan, L. Wang, Y. Hu, L. Ceze, et al. TVM: An automated end-to-end optimizing compiler for deep learning. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, pages 578–594, 2018.
- [23] L. Chizat, G. Peyré, B. Schmitzer, and F.-X. Vialard. Scaling algorithms for unbalanced optimal transport problems. *Mathematics of Computation*, 87(314):2563–2609, 2018.
- [24] H. Chui and A. Rangarajan. A new algorithm for non-rigid point matching. In *Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on*, volume 2, pages 44–51. IEEE, 2000.
- [25] S. Cook. *CUDA programming: a developer’s guide to parallel computing with GPUs*. Newnes, 2012.
- [26] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312. ACM, 1996.
- [27] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, pages 2292–2300, 2013.
- [28] J. Delon, J. Salomon, and A. Sobolevski. Local matching indicators for transport problems with concave costs. *SIAM Journal on Discrete Mathematics*, 26(2):801–827, 2012.
- [29] W. E. Deming and F. F. Stephan. On a least squares adjustment of a sampled frequency table when the expected marginal totals are known. *The Annals of Mathematical Statistics*, 11(4):427–444, 1940.
- [30] S. Di Marino, A. Gerolin, and L. Nenna. Optimal transportation theory with repulsive costs. *Topological optimization and optimal transport*, 17:204–256, 2017.
- [31] W. Dong, C. Moses, and K. Li. Efficient k-nearest neighbor graph construction for generic similarity measures. In *Proceedings of the 20th international conference on World wide web*, pages 577–586, 2011.
- [32] A. Dutt and V. Rokhlin. Fast Fourier transforms for nonequispaced data. *SIAM Journal on Scientific computing*, 14(6):1368–1393, 1993.
- [33] D. Eddelbuettel and R. François. Rcpp: Seamless R and C++ integration. *Journal of Statistical Software*, 40(8):1–18, 2011.
- [34] S. Erlander. *Optimal spatial interaction and the gravity model*, volume 173. Springer Science & Business Media, 1980.
- [35] M. Fey, J. Eric Lenssen, F. Weichert, and H. Müller. SplineCNN: Fast geometric deep learning with continuous b-spline kernels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 869–877, 2018.
- [36] M. Fey and J. E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- [37] J. Feydy. *Geometric data analysis, beyond convolutions*. PhD thesis, Université Paris-Saclay, 2020.

- [38] J. Feydy, B. Charlier, F.-X. Vialard, and G. Peyré. Optimal transport for diffeomorphic registration. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 291–299. Springer, 2017.
- [39] J. Feydy, P. Roussillon, A. Trouvé, and P. Gori. Fast and scalable optimal transport for brain tractograms. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 636–644. Springer, 2019.
- [40] J. Feydy, T. Séjourné, F.-X. Vialard, S.-i. Amari, A. Trouvé, and G. Peyré. Interpolating between optimal transport and MMD using Sinkhorn divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2681–2690, 2019.
- [41] S. E. Fienberg et al. An iterative procedure for estimation in contingency tables. *The Annals of Mathematical Statistics*, 41(3):907–917, 1970.
- [42] R. Flamary and N. Courty. POT python optimal transport library, 2017.
- [43] A. Forrow, J.-C. Hütter, M. Nitzan, P. Rigollet, G. Schiebinger, and J. Weed. Statistical optimal transport via factored couplings. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2454–2465, 2019.
- [44] A. Galichon and B. Salanié. Matching with trade-offs: Revealed preferences over competing characteristics. *CEPR Discussion Paper No. DP7858*, 2010.
- [45] J. Gardner, G. Pleiss, K. Q. Weinberger, D. Bindel, and A. G. Wilson. GPyTorch: Blackbox matrix-matrix Gaussian process inference with GPU acceleration. In *Advances in Neural Information Processing Systems*, pages 7576–7586, 2018.
- [46] A. Genevay, L. Chizat, F. Bach, M. Cuturi, and G. Peyré. Sample complexity of Sinkhorn divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1574–1583, 2019.
- [47] A. Genevay, M. Cuturi, G. Peyré, and F. Bach. Stochastic optimization for large-scale optimal transport. In *Advances in Neural Information Processing Systems*, pages 3440–3448, 2016.
- [48] A. Genevay, G. Peyré, and M. Cuturi. Learning generative models with Sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pages 1608–1617, 2018.
- [49] F. Gieseke, J. Heineremann, C. Oancea, and C. Igel. Buffer kd trees: processing massive nearest neighbor queries on GPUs. In *International Conference on Machine Learning*, pages 172–180, 2014.
- [50] L. Greengard. *The rapid evaluation of potential fields in particle systems*. MIT press, 1988.
- [51] L. Greengard and J.-Y. Lee. Accelerating the nonuniform fast Fourier transform. *SIAM review*, 46(3):443–454, 2004.
- [52] L. Greengard and J. Strain. The fast Gauss transform. *SIAM Journal on Scientific and Statistical Computing*, 12(1):79–94, 1991.
- [53] M. Greenspan and M. Yurick. Approximate kd tree search for efficient ICP. In *Fourth International Conference on 3-D Digital Imaging and Modeling, 2003. 3DIM 2003. Proceedings.*, pages 442–448. IEEE, 2003.
- [54] R. Guo, P. Sun, E. Lindgren, Q. Geng, D. Simcha, F. Chern, and S. Kumar. Accelerating large-scale inference with anisotropic vector quantization. In *International Conference on Machine Learning*, 2020.
- [55] W. Hackbusch. *Hierarchical matrices: algorithms and analysis*, volume 49. Springer, 2015.
- [56] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer series in statistics. Springer, 2009.
- [57] J. Hensman, N. Fusi, and N. D. Lawrence. Gaussian processes for big data. In *Uncertainty in Artificial Intelligence*, page 282. Citeseer, 2013.
- [58] Intel AI. PlaidML: a framework for making deep learning work everywhere. v0.7.0, 2019.
- [59] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [60] K. J., E. Smith, J.-F. Lafleche, C. Fuji Tsang, A. Rozantsev, W. Chen, T. Xiang, R. Lebedev, and S. Fidler. Kaolin: A PyTorch library for accelerating 3d deep learning research. *arXiv:1911.05063*, 2019.

- [61] W. Jakob, J. Rhineland, and D. Moldovan. PyBind11 – seamless operability between C++11 and python, 2017. <https://github.com/pybind/pybind11>.
- [62] H. Jegou, M. Douze, and C. Schmid. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence*, 33(1):117–128, 2010.
- [63] Z. Jia, S. Lin, C. R. Qi, and A. Aiken. Exploring hidden dimensions in parallelizing convolutional neural networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2018.
- [64] Z. Jia, M. Zaharia, and A. Aiken. Beyond data and model parallelism for deep neural networks. In *Proceedings of the 2nd Conference on Machine Learning and Systems (MLSys)*, 2018.
- [65] J. Johnson, M. Douze, and H. Jégou. Billion-scale similarity search with GPUs. *arXiv preprint arXiv:1702.08734*, 2017.
- [66] E. Jones, T. Oliphant, P. Peterson, et al. SciPy: Open source scientific tools for Python, 2001. [Online; accessed 17th November 2019].
- [67] W. Kahan. Further remarks on reducing truncation errors. *Communications of the ACM*, 8(1):40, 1965.
- [68] L. Kantorovich. On the transfer of masses (in Russian). *Doklady Akademii Nauk*, 37(2):227–229, 1942.
- [69] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [70] P. A. Knight, D. Ruiz, and B. Uçar. A symmetry preserving algorithm for matrix scaling. *SIAM journal on Matrix Analysis and Applications*, 35(3):931–955, 2014.
- [71] J. J. Kosowsky and A. L. Yuille. The invisible hand algorithm: Solving the assignment problem with statistical physics. *Neural networks*, 7(3):477–490, 1994.
- [72] J. Lamping, R. Rao, and P. Piroli. A focus+context technique based on hyperbolic geometry for visualizing large hierarchies. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 401–408, 1995.
- [73] Y. LeCun and C. Cortes. The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- [74] T.-M. Li, M. Gharbi, A. Adams, F. Durand, and J. Ragan-Kelley. Differentiable programming for image processing and deep learning in Halide. *ACM Transactions on Graphics (TOG)*, 37(4):1–13, 2018.
- [75] Y. Li, R. Bu, M. Sun, W. Wu, X. Di, and B. Chen. PointCNN: Convolution on X-transformed points. In *Advances in neural information processing systems*, pages 820–830, 2018.
- [76] Z. Liu, H. Tang, Y. Lin, and S. Han. Point-Voxel CNN for efficient 3d deep learning. In *Advances in Neural Information Processing Systems*, pages 963–973, 2019.
- [77] Y. A. Malkov and D. A. Yashunin. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE transactions on pattern analysis and machine intelligence*, 2018.
- [78] G. Matheron. Principles of geostatistics. *Economic geology*, 58(8):1246–1266, 1963.
- [79] L. McInnes, J. Healy, and J. Melville. UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- [80] G. Meanti, L. Carratino, L. Rosasco, and A. Rudi. Kernel methods through the roof: handling billions of points efficiently. *arXiv preprint arXiv:2006.10350*, 2020.
- [81] G. Monge. Mémoire sur la théorie des déblais et des remblais. *Histoire de l’Académie Royale des Sciences*, pages 666–704, 1781.
- [82] H. Nguyen. *Gpu gems 3*. Addison-Wesley Professional, 2007.
- [83] M. Nickel and D. Kiela. Poincaré embeddings for learning hierarchical representations. In *Advances in neural information processing systems*, pages 6338–6347, 2017.
- [84] R. Okuta, Y. Unno, D. Nishino, S. Hido, and C. Loomis. CuPy: A NumPy-compatible library for Nvidia GPU calculations. In *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*, 2017.

- [85] S. M. Omohundro. *Five balltree construction algorithms*. International Computer Science Institute Berkeley, 1989.
- [86] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in PyTorch. *NIPS 2017 Workshop Autodiff*, 2017.
- [87] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [88] J. Pennington, R. Socher, and C. D. Manning. GloVe: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014.
- [89] G. Peyré and M. Cuturi. Computational optimal transport. *arXiv:1610.06519*, 2017.
- [90] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Advances in neural information processing systems*, pages 1177–1184, 2008.
- [91] RAPIDS Development Team. *RAPIDS: Collection of Libraries for End to End GPU Data Science*, 2018.
- [92] N. Ravi, J. Reizenstein, D. Novotny, T. Gordon, W.-Y. Lo, J. Johnson, and G. Gkioxari. PyTorch3D. <https://github.com/facebookresearch/pytorch3d>, 2020.
- [93] Y. Rubner, C. Tomasi, and L. J. Guibas. The earth mover’s distance as a metric for image retrieval. *International Journal of Computer Vision*, 40(2):99–121, Nov. 2000.
- [94] L. Ruschendorf et al. Convergence of the iterative proportional fitting procedure. *The Annals of Statistics*, 23(4):1160–1174, 1995.
- [95] F. Sala, C. De Sa, A. Gu, and C. Re. Representation tradeoffs for hyperbolic embeddings. In J. Dy and A. Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4460–4469, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.
- [96] B. Schmitzer. Stabilized sparse scaling algorithms for entropy regularized transport problems. *SIAM Journal on Scientific Computing*, 41(3):A1443–A1481, 2019.
- [97] N. Shazeer, Y. Cheng, N. Parmar, D. Tran, A. Vaswani, P. Koanantakool, P. Hawkins, H. Lee, M. Hong, C. Young, R. Sepassi, and B. Hechtman. Mesh-TensorFlow: Deep learning for supercomputers. In *Neural Information Processing Systems*, 2018.
- [98] R. Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *Ann. Math. Statist.*, 35:876–879, 1964.
- [99] G. J. Székely and M. L. Rizzo. Testing for equal distributions in high dimension. *InterStat*, 5(16.10), 2004.
- [100] Y. Tay, M. Dehghani, D. Bahri, and D. Metzler. Efficient transformers: A survey. *arXiv preprint arXiv:2009.06732*, 2020.
- [101] Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, abs/1605.02688, May 2016.
- [102] H. Thomas, C. R. Qi, J.-E. Deschaud, B. Marcotegui, F. Goulette, and L. J. Guibas. KPConv: Flexible and deformable convolution for point clouds. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [103] J. Tierny, G. Favelier, J. A. Levine, C. Gueunet, and M. Michaux. The topology toolkit. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):832–842, 2017.
- [104] S. Van Der Walt, S. C. Colbert, and G. Varoquaux. The NumPy array: a structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2):22, 2011.
- [105] N. Vasilache, O. Zinenko, T. Theodoridis, P. Goyal, Z. DeVito, W. S. Moses, S. Verdoolaege, A. Adams, and A. Cohen. Tensor comprehensions: Framework-agnostic high-performance machine learning abstractions. *arXiv preprint arXiv:1802.04730*, 2018.
- [106] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.

- [107] C. Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.
- [108] M. Wang, C.-c. Huang, and J. Li. Supporting very large models using automatic dataflow graph partitioning. In *Proceedings of the Fourteenth EuroSys Conference 2019*, pages 1–17, 2019.
- [109] M. Wang, L. Yu, D. Zheng, Q. Gan, Y. Gai, Z. Ye, M. Li, J. Zhou, Q. Huang, C. Ma, et al. Deep graph library: Towards efficient and scalable deep learning on graphs. *arXiv preprint arXiv:1909.01315*, 2019.
- [110] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon. Dynamic graph CNN for learning on point clouds. *ACM Transactions on Graphics (TOG)*, 38(5):1–12, 2019.
- [111] C. K. Williams and M. Seeger. Using the Nyström method to speed up kernel machines. In *Advances in neural information processing systems*, pages 682–688, 2001.
- [112] A. G. Wilson. The use of entropy maximising models, in the theory of trip distribution, mode split and route split. *Journal of Transport Economics and Policy*, pages 108–126, 1969.
- [113] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush. HuggingFace’s Transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771, 2019.
- [114] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and S. Y. Philip. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [115] C. Yang, R. Duraiswami, N. A. Gumerov, and L. Davis. Improved fast Gauss transform and efficient kernel density estimation. In *Proceedings of the Ninth IEEE International Conference on Computer Vision*, volume 2, page 464. IEEE Computer Society, 2003.
- [116] T. Yang, Y.-F. Li, M. Mahdavi, R. Jin, and Z.-H. Zhou. Nyström method vs random Fourier features: A theoretical and empirical comparison. In *Advances in neural information processing systems*, pages 476–484, 2012.
- [117] G. U. Yule. On the methods of measuring association between two attributes. *Journal of the Royal Statistical Society*, 75(6):579–652, 1912.
- [118] W.-L. Zhao. Approximate k-NN graph construction: a generic online approach. *arXiv preprint arXiv:1804.03032*, 2018.
- [119] Y. Zhou, C. Wu, Z. Li, C. Cao, Y. Ye, J. Saragih, H. Li, and Y. Sheikh. Fully convolutional mesh autoencoder using efficient spatially varying kernels. In *arxiv*, 2020.