

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Bringing UMAP Closer to the Speed of Light with GPU Acceleration

Corey J. Nolet

NVIDIA

Univ. of Maryland, Baltimore County
cjnolet@gmail.com

Victor Lafargue

NVIDIA

viclafargue@nvidia.com

Edward Raff

Booz Allen Hamilton

Univ. of Maryland, Baltimore County
raff_edward@bah.com

Thejaswi Nanditale

NVIDIA

snanditale@nvidia.com

Tim Oates

Univ. of Maryland, Baltimore County
oates@umbc.edu

John Zedlewski

NVIDIA

zjedlewski@nvidia.com

Joshua Patterson

NVIDIA

joshuap@nvidia.com

Abstract

The Uniform Manifold Approximation and Projection (UMAP) algorithm has become widely popular for its ease of use, quality of results, and support for exploratory, unsupervised, supervised, and semi-supervised learning. While many algorithms can be ported to a GPU in a simple and direct fashion, such efforts have resulted in inefficient and inaccurate versions of UMAP. We show a number of techniques that can be used to make a faster and more faithful GPU version of UMAP, and obtain speedups of up to 100x in practice. Many of these design choices/lessons are general purpose and may inform the conversion of other graph and manifold learning algorithms to use GPUs. Our implementation has been made publicly available as part of the open source RAPIDS cuML library (<https://github.com/rapidsai/cuml>).

1 Introduction

Like other manifold learning algorithms, the Uniform Manifold Approximation and Projection algorithm (UMAP) [1] relies upon the manifold hypothesis [2] to preserve local neighborhood structure by modeling high-dimensional data in a low-dimensional space. This is in contrast to linear dimensionality reduction techniques like PCA, which aim only to preserve global Euclidean structure [3]. UMAP produces low-dimensional embeddings that are useful for both visual analytics and downstream machine learning tasks. Unlike other manifold learning algorithms, such as IsoMap [4], Locally Linear Embeddings (LLE) [5], Laplacian Eigenmaps [6], and t-Distributed Stochastic Neighbor Embeddings (T-SNE) [7], UMAP has native support for supervised, unsupervised, and semi-supervised metric learning. Since its introduction in 2018, it has found use in exploratory data analysis applications [8–11], as well as bioinformatics, cancer research [12], single-cell genomics [13–15],

and the interpretation of highly non-linear models like deep neural networks [16]. This combination of features and quality of results has made UMAP a widely used and popular tool.

The wide array of applications and use of UMAP makes it desirable to produce faster versions of the algorithm. This is compounded by an increasing demand for exploratory and interactive visualization [16–18, 10], that necessitates a lower latency in results. GPUs are a strong candidate for achieving faster implementation by trading per-core clock speeds for significantly more cores, provided that they can be utilized effectively in parallel computation. A direct conversion of UMAP to the GPU already exists in the GPUMAP [19] project but, due to technical details, is not always faithful in reproducing the same quality of results. In this work, we show that applying a few general techniques can produce a version that is both faster and faithful in its results.

This paper contributes three components to the growing ecosystem [20–23] of GPU-accelerated tools for data science in Python. First, we contribute a near drop-in replacement of the UMAP-learn Python library, which has been GPU-accelerated end-to-end in CUDA/C++. The speedup of this new implementation is evaluated against the current state-of-the-art, including UMAP-learn on the CPU and an existing GPU port of UMAP. Second, we contribute a GPU-accelerated near-drop-in replacement of the trustworthiness score, which is often used to evaluate the extent to which manifold learning algorithms preserve local neighborhood structure. Finally, we contribute a distributed version of UMAP and provide empirical evidence of its effectiveness. In Appendix A, we describe two additional features in our implementation that we have found to further close the performance gap for exploratory and interactive workflows.

The remainder of our paper is organized as follows. We will discuss related work in section 2, with a brief review of UMAP in section 3. Our approach to implementing UMAP for the GPU is detailed in section 4, with code available as part of the RAPIDS cuML library (<https://github.com/rapidsai/cuml>). Our results show up to $100\times$ speedups in section 5, followed by our conclusions in section 6.

2 Related Work

UMAP’s ability to shortcut the need for computing n^2 pairwise distances by defining local neighborhoods with the k -nearest neighbors around each data point is like other manifold learning algorithms. T-SNE predates UMAP and has found popularity in many of the same communities that UMAP has now become considered the state of the art [11]. T-SNE models point distances as probability distributions, constructing a students-t kernel from training data and minimizing the Kullback-Liebler divergence against the low-dimensional representation. While originally intractable for datasets containing more than a few thousand points, GPU-accelerated variants have recently breathed new life into the algorithm [24]. Still, T-SNE has not been shown to work well for downstream machine learning tasks and lacks support for supervised learning.

The reference implementation of UMAP is built on top of the Numba[25] library and uses just-in-time (JIT) compilation to make use of parallel low-level CPU optimizations. The GPUMAP library [19] is a direct port of the reference library to the GPU, using Numba’s `cuda.jit` feature, along with the CuPy library, to directly replace many SciPy library invocations with CUDA-backed implementations. Like other libraries that require fast nearest neighbors search on GPUs [24], GPUMAP uses the FAISS library [22]. Our implementation also uses the FAISS library. GPUMAP invokes FAISS through the Python API, missing opportunities for zero-copy exchanges of memory pointers on device [21] that our implementation leverages.

Manifold learning algorithms typically use the trustworthiness [26] score to evaluate a trained model’s preservation of local neighborhood structure. The trustworthiness score penalizes divergences in the nearest neighbors between the algorithm’s input and output, ranking the similarities of the neighborhoods. Scikit-learn [27] provides an implementation of trustworthiness, but the computational costs and memory footprint associated with computing the entire n^2 pairwise distance matrix makes it prohibitively slow to evaluate datasets greater than a couple thousand samples. To the best of our knowledge, there are no existing ports of the trustworthiness score to the GPU. We fill this gap with our new batchable implementation, which we demonstrate can scale well over 100s of thousands of samples on a single GPU with reasonable performance.

3 Uniform Manifold Approximation and Projection

Like many manifold learning algorithms, the UMAP algorithm can be decomposed into three major stages, which we briefly describe in this section. For explanations, derivations, and further details, we refer the reader to the official UMAP paper [1].

In the first stage, a k -nearest neighbors (k -NN) graph is constructed using a distance metric, $d(x, y)$. The second stage weights the closest neighbors around each vertex in the nearest neighbors graph, converting them to fuzzy sets and combining them into a fuzzy union. The fuzzy set membership function learns a locally adaptive exponential kernel that smoothes the distances in each local neighborhood of the k -NN graph by finding a smoothing normalizer σ_i such that Equation 1 is satisfied. ρ in this equation contains the distances to the closest non-zero neighbor around each vertex. The triangular conorm [28, 29] in Equation 2 combines the matrix of individual fuzzy sets, A , into a fuzzy union by symmetrizing the graph and adding the element-wise (Hadamard) product.

$$\sum_{j=i}^k \exp(-\max(0, d(x_i, x_{i_j}) - \rho_i) \sigma_i^{-1}) = \log_2(k) \quad (1)$$

$$B = (A + A^T) + (A \circ A^T) \quad (2)$$

In the third and final stage, the embeddings are laid out in the topological space using stochastic gradient descent. An initial layout is performed either by sampling embeddings from a uniform distribution or computing a spectral embedding over the fuzzy union. The cross-entropy objective function $-\sum_{a,b \in B} (\log(\Phi(a, b)) + \sum_{c \in B}^m \log(1 - \Phi(a, c)))$ is minimized over the edges of the fuzzy union, B , from Equation 2. This is done with negative sampling where m is the number of negative samples per edge. Φ (Equation 3) is the current membership strength in the newly embedded space and min_dist controls the minimum separation distance between points in the embedded space. We use the approximate form of Φ in this paper for simplicity. The $\log(\Phi)$ term in the objective is computed using the source and destination vertices on each edge and $\log(1 - \Phi)$ is computed using the source vertex with negative sampling.

$$\Phi(x, y) \approx \begin{cases} 1 & \|x - y\|_2 \leq min_dist \\ \exp(-\|x - y\|_2 - min_dist) & otherwise \end{cases} \quad (3)$$

When training labels are provided, an additional step in the neighborhood weighting stage adjusts the membership strengths of the fuzzy sets based on their labels. In addition to its learned parameters, the trained UMAP model keeps a reference to the k -NN index computed on the training data. This is used to create a mapping of the trained embeddings to a new set of vertices during inference.

4 GPU-Accelerating UMAP

Our implementation is primarily written in C++, which is wrapped in a Python API through the Cython library. Data can be passed into our Python API using common formats like Numpy [30] or Pandas [31], as well as GPU array libraries such as CuPy [20], Numba [25], or RAPIDS cuDF [21]. When necessary, the data is automatically copied onto the device (e.g., when a Numpy array is passed in). Columnar memory layouts, such as those used in Apache Arrow [32], tend to exploit the optimized memory access patterns on GPUs, such as coalesced accesses [33]. Like UMAP-learn, our Python API maintains compatibility with the Scikit-learn [27] API. We used the RAPIDS memory manager (RMM) to create a single memory pool for each process to avoid device synchronization from allocations and deallocations of temporary device memory. When available, we made use of existing libraries with optimized CUDA primitives, such as Thrust [34], cuSparse [35, 36], cuGraph, and cuML [37, 21].

Our implementation begins with a straightforward port of UMAP-learn to CUDA/C++, diverging from the design of UMAP-learn only where we found a significant benefit to performance or the memory footprint. The prior GPUMAP implementation attempted a direct conversion of the code design, using Numba CUDA-JIT [11] functions and CuPy, without any significant diversions. While hypothetically easier to maintain, we will show this produces results that do not always match the original implementation, and does not deliver meaningful speedups. In each section below, we will

detail some of the major design choices that make our GPU implementation faster and more faithful to the original results.

Copying data between host memory and a GPU device comes at a high cost and can quickly become a bottleneck. Transferring 1MB of data can take several hundreds of microseconds. Even when memory is copied asynchronously between device and host, the underlying CUDA stream needs to be synchronized before the data can be used on the host, further increasing this latency. We reduce the need to transfer between host and device as much as possible, even if that means running code on the GPU that has little to no speedup over the CPU [38]. Standards like the `__cuda_array_interface__` [21] and `dlpack` [39] enable Python libraries to share CUDA memory pointers directly, without the need for copies [21]. This further reduces the need for transfers to host, and like the standard `__array_interface__` [40], enables implicit conversion between types even as data is passed between different libraries.

4.1 GPU Architecture

The NVIDIA General-purpose GPU computing architecture [41, 42] enables parallelism through the single-instruction multiple data design paradigm [21]. A single GPU device contains several banks of global memory that are accessible from a grid containing thousands of instruction processing cores called thread-blocks. Each thread-block has its own faster but smaller bank of memory, called shared memory, which is available to a series of threads that it executes concurrently. Each thread has a series of registers available, which it uses to store local variables. The amount of registers and shared memory available to each will depend on the number of threads configured for each thread-block.

GPUs provide the most performance gains when memory access patterns are able to take advantage of features like shared memory, uniform conditional branching, and coalesced memory accesses. Though core-for-core typically not as fast as a CPU, parallelizing operations over GPU threads can still provide significant performance gains even when memory access patterns are not efficient [38].

4.2 Constructing the World k -NN Graph

The UMAP-learn library utilizes nearest neighbors descent [43] for construction of an approximate nearest neighbors graph, however no known GPU-accelerated versions of this algorithm exist at the time of writing. Tree-based approximate variants, such as the algorithms available in Scikit-learn, also don't have a straightforward port to the GPU [44]. This is a direct result of the iterative nature of traversal, as well as the storage and representation requirements for the trees after they are constructed.

Our implementation of UMAP makes use of the FAISS library [22] for fast nearest neighbors search on GPUs. Other GPU-accelerated manifold implementations have used this same approach (e.g., t-SNE [24]). FAISS provides both exact and approximate methods to nearest neighbors search, the former being used by default in our implementation. We use the exact search provided by FAISS since it is performant and doesn't require underlying device memory be copied during the hand-off.

For smaller datasets of a few hundred thousand samples and a few hundred features, we found the quadratic scale of the exact k -NN graph computation to comprise 26% of the total time UMAP spends in compute, making it the second largest performance bottleneck next to the optimization of the embeddings. However, as we demonstrate in Figure 2, the k -NN graph can quickly become the largest bottleneck as the number of samples and features increase to larger sizes, while the optimization stage consistently maintains high performance. This is not a surprising find, since the brute force approach requires exhaustive distances to be computed along with a heap, which is required to maintain the sorted order of closest neighbors. An additional cause of significant performance degradation during this stage is FAISS' incompatibility with outside memory managers, causing unavoidable and expensive synchronous device memory allocations and deallocations for temporary scratch space. As described in subsection A.1, we modified the API to accept a k -NN graph that has already been computed. This modification provides a strategy to avoid these expensive synchronizations in exploratory environments where a model might be trained several times on the same data.

4.3 Handling Sparse Data

Once computed, many operations are performed over the sparse k -NN graph. This is common in many modern manifold learning approaches as well as network analysis problems, where these

performance optimizations may be reused. The edges of the k -NN graph are unidirectional and the index and distance arrays represent the column and data arrays of a sparse format. The fixed degree makes the row array implicit and allows the use of dense matrix operations until the construction of the fuzzy union, where the degree is no longer fixed.

We use the COOrdinate (COO), or edge list format for efficient out-of-order parallel construction and subsequent element-wise operations. We have found sorting an out-of-order COO can take up to 3% of the total time spent in compute. When it is efficient to do so, we sort the COO arrays by row and create a Compressed Sparse Row (CSR) index into the column and data arrays, enabling both efficient row- and element-wise parallelism so long as the sorted order is maintained. See Figure 1 for a diagram.




Figure 1: Example of our CSR index being used to index into a sorted COO index.

While our implementation makes use of libraries like cuSparse and cuGraph for common operations on sparse data, we built several reusable primitives for operations such as sparse L_1 and L_∞ normalization, removal of nonzeros, and the triangular conorm, where existing implementations of these operations were not available. Aside from custom kernels that don't have much potential for reuse outside of UMAP, such as those described in the following three sections, reusable primitives comprise a large portion of the algorithm.

4.4 Neighborhood Weighting

The neighborhood weighting step begins with constructing the ρ and σ arrays, with one element for each vertex of the k -NN graph. ρ contains the distance to the closest neighbor of each vertex and σ contains the smoothing approximator to the fuzzy set membership function for the local neighborhoods of each source vertex in the k -NN graph. The operations for computing these two arrays are fused into a single kernel, which maps each source vertex of the k -NN graph in CSR format to a separate CUDA thread. The computations in this kernel are largely similar to corresponding Python code in the reference implementation and comprise less than 0.1% of the total time spent in compute.

The k -NN distances are weighted by applying the fuzzy set membership function from the previous step to the COO matrix containing the edges of each source vertex in the k -NN graph. Since this computation requires no dependencies between the edges in the neighborhood graph, the CUDA kernel maps each neighbor to their own thread individually.

As described in section 3, the final step of the neighborhood weighting stage combines all the fuzzy sets, using the triangular conorm to build a fuzzy union. We implemented this step by fusing both symmetrization sum and product steps together into a single kernel, using the CSR indptr we introduced in subsection 4.3 as a Compressed Sparse Column (CSC) indptr to look up the the transposed value and apply the triangular conorm to each element in parallel. This step comprises less than 0.2% of the total time spent in compute.

Larger kernels composed of smaller fused operations, such as computing the mean, min, and iterating for the adaptive smoothing parameters, allowed us to make use of registers where the alternative required intermediate and more expensive storage. We found a 12-15 \times speedup for the adaptive smoothing operations when compared to separate kernels that require intermediate results to be stored in global memory. The end-to-end neighborhood weighting stage exploits parallelism at the expense of potential thread divergence from non-uniform conditional branching, and help the kernels to stay compute-bound.

4.5 Embedding Updates

The first step of the embeddings optimization stage initializes the array of output embeddings. We provide both random and spectral initialization strategies. While the reference implementation uses a spectral embedding of the fuzzy union through the nearest-neighbors variant of the Laplacian eigenmaps [6] algorithm, we use the spectral clustering implementation from cuGraph [45], setting

the number of clusters to 1 and removing the lowest eigenpairs. We have found spectral clustering to be sufficient for maintaining comparable trustworthiness in our experiments while comprising less than 0.1% of the total time spent in compute.

The optimization step performs stochastic gradient descent over the edges of the fuzzy union, minimizing the cross entropy described in section 3. The gradient computation and update operations have been fused into a single kernel and parallelized so that each thread processes one edge of the fuzzy union. The CUDA kernel is scheduled iteratively for n_epochs to compute and apply the gradient updates to the embeddings in each epoch. The dependencies between the vertices in the updating of the gradients makes this step non-trivial to parallelize efficiently, which decreases potential for coalesced memory access and creates the need for atomic operations when applying gradient updates. As a result, we have seen this kernel take up to 30% of the total time spent in compute for datasets of a few hundred thousand samples with a few hundred features. When the k -NN graph is pre-computed, as discussed in subsection A.1, this step can comprise up to 50% of the remaining time spent in compute. The dependencies between vertices also create challenges to reproducibility, which we describe in subsection 4.6.

Both the source and destination vertices are updated for each edge during training. Since the trained embeddings should remain unchanged, only the destination vertex is updated during inference. In addition, both training and inference require the source vertex be updated for some number of randomly sampled vertices. Each source vertex will perform $n_components * (n_negative_samples + 1)$ atomic writes in each thread plus an additional write for the destination vertex during training.

When $n_components$ is small enough, such as a few hundred, we use shared memory to create a small local cache per compute thread, accumulating the updates for each source vertex from multiple negative samples before writing the results atomically to global memory. When shared memory can be used, this reduces atomic updates per thread by a factor of $n_components * n_negative_samples$. We have measured performance gains of 10% for this stage when $n_components = 2$ to 56% when $n_components = 16$ and expect the performance benefits to continue increasing in proportion to $n_components$. For these cases where $n_components$ is very small, such as $n_components = 2$, these updates can be accumulated right in the registers, providing a speedup of 49% for this stage. We suspect these strategies, and any future optimizations, will be useful broadly given the many algorithms (e.g., word2vec [46]) that make use of negative sampling.

4.6 Reproducibility

Following the original implementation of UMAP, the user can provide a seed to control the random initialization of weights to increase the reproducibility. This does not eliminate all inconsistency when working with parallel updates made from multiple threads. When using a limited number of CPU cores (≤ 40 in most circumstances), this effect is minimal. However, with a GPU that has thousands of parallel threads, even subtle timing differences between the thread-blocks can have a large impact on the consistency of results. In addition, large numbers of updates can become queued waiting to be performed atomically. A similar issue is observed with the Hogwild algorithm even when atomic updates are used [47–52], but at a larger scale. This problem is further exacerbated by small divergences in the processing of instructions that results from non-uniform conditional branching across threads.

Our use of a local cache to accumulate updates as described in subsection 4.5 alleviates this by decreasing the number of global atomic writes, helping to reduce the potential for thread divergence and resulting in higher quality solutions. While this minimizes the writes significantly, we still found the potential for inconsistencies to increase in proportion to the number of vertices in the dataset, the number edges in the fuzzy union, and the number of components being trained.

The results are made fully repeatable with exact precision by optionally using a 64-bit float array to accumulate the updates to the embeddings and applying the updates at the end of each epoch. The additional precision avoids the numerical instabilities created by repeatedly summing small values in a finite range while the single application of the updates removes the potential for race conditions between reads and writes [53]. We found the performance impact to increase with the number of components, from an end-to-end slowdown of $11\times$ with $n_components = 2$ to $20\times$ with $n_components = 16$ on a Volta GV100 GPU.

4.7 Distributed Inference

Because of its ability to embed out-of-sample data points [54], we scaled the UMAP algorithm to support datasets larger than a single GPU by training a model on a random sample of the training dataset, sending the trained embeddings to a set of workers, each mapped to its own GPU, and performing inference in parallel on the remaining data samples. Our implementation minimizes the use of host memory during communication by using CUDA IPC [55] to support fast communication over NVLink [56] internal to a physical machine and GPUDirect RDMA [57] to communicate across machine boundaries. We use the Dask library, which has been GPU-accelerated [21] and optimized with the Unified-Communications-X library (UCX) [58] to support CUDA IPC and GPUDirect transports automatically, without the need to invoke the aforementioned transports directly.

As demonstrated in Appendix B, we have found our distributed implementation to scale linearly with the number of GPUs. We find it can preserve structure for a dataset containing 23 thousand data points when trained on as little as 3% of the data with less than a 1% drop in trustworthiness. Further, we find only a 0.05% drop in trustworthiness when we embed the remaining 97% of the dataset over 16 separate workers.

5 Experiments

We compare the execution time and correctness of GPUMAP and our implementation against the multi-core implementation of UMAP-learn on CPU. The datasets are summarized in Table 1, showing the number of rows, columns, and classes. We evaluated the execution times of unsupervised training on each dataset for all three implementations and recorded the resulting times, in seconds. Where classes were provided, we also evaluated the supervised training mode. All experiments were conducted on a single DGX1 containing 8 Nvidia GV100 GPUs with Dual Intel Xeon 20-core CPUs. UMAP-learn was configured to take advantage of all the available threads on the machine.

We use trustworthiness to rank the degree to which local neighborhood structure is preserved between input and embedded spaces. Scikit-learn provides an implementation of this score, but the execution time and memory requirement of computing the pairwise distance matrix make it prohibitive on some of the datasets used in this paper. We implemented a batched GPU-accelerated version of trustworthiness that provides reasonably low execution times for datasets up to 1M samples. Table 2 contains the execution times of computing the trustworthiness score on various different numbers of samples in both Scikit-learn and cuML UMAP. This contribution was necessary to perform our evaluations and was used to evaluate the correctness of each implementation. Because users often select the result with the highest trustworthiness score, we report the max score in results.

We begin by demonstrating the speedups obtained by our new cuML UMAP implementation of UMAP in the standard unsupervised scenario. The timing results with standard deviation from 4 runs can be found in Table 3, with the trustworthiness score on the right. cuML UMAP dominates all other implementations in speed, with $17\times$ speedups compared to UMAP-learn on the smallest datasets, and increasing to up to $104.9\times$ on moderate scale datasets like MNIST. Similar results can be seen in the supervised case in Table 4. cuML UMAP is also $2.65 - 15.6\times$ faster

Table 1: Datasets used in experiments

Dataset	Rows	Columns	Classes
Digits [59]	1797	64	10
Shuttle [60]	58k	9	7
Fashion MNIST [61]	60k	784	10
MNIST [62]	60k	784	10
CIFAR-100 [63]	60k	1024	20
COIL-20 [64]	1440	16384	20
scRNA [13]	64.5k	5k	N/A
GoogleNews Word2vec [65]	3M	300	N/A

Table 2: Execution times for computing the Trustworthiness score with UMAP’s default of $n_neighbors = 15$. The first column shows the number of samples used, and the right two columns present run time in seconds. The number of features was fixed to 1024. Best results are in **bold**. Done using isotropic blobs

Samples	Scikit-learn	cuML UMAP
2k	0.33	0.13
5k	2.06	0.18
10k	8.64	0.24
20k	35.76	0.54
50k	303.08	2.07
100k	FAIL	5.74
1M	FAIL	446.26

than the prior GPUMAP, with an average $7.29\times$ advantage. This is biased towards GPUMAP’s favor by the fact that it has regressions in solution quality, as measured by trustworthiness, on 6/7 datasets.

Table 3: Each result shows mean \pm variance, followed by max trustworthiness score, of each implementation of UMAP for the unsupervised case with default parameters. Fastest result in **bold**.

Dataset	UMAP-Learn		GPUMAP		cuML UMAP	
	$\mu \pm \sigma^2$	Trust%	$\mu \pm \sigma^2$	Trust%	$\mu \pm \sigma^2$	Trust%
digits	6.328 ± 2.897	98.79	2.483 ± 1.058	95.58	0.3583 ± 0.0111	98.77
fashion mnist	45.87 ± 10.23	97.81	4.158 ± 1.800	97.50	0.455 ± 0.006	97.73
mnist	52.575 ± 1.1677	95.94	10.6071 ± 0.45444	94.43	0.70781 ± 0.0088	95.74
cifar100	105.85 ± 2.482	84.72	6.186 ± 1.770	84.01	1.009 ± 0.0188	83.42
coil20	11.210 ± 2.571	99.36	2.582 ± 0.0050	95.67	0.757 ± 0.5752	99.28
shuttle	38.88 ± 8.039	100.0	9.064 ± 3.431	97.78	0.5825 ± 0.0252	100.0
scRNA	223.9 ± 9.071	62.38	10.89 ± 1.604	94.35	4.103 ± 0.0601	97.81

The trustworthiness and speedups show the value of our contributions in section 4, and we in addition note that the CPU based UMAP-learn has it’s own regression on the scRNA dataset. This is a known issue caused by the lack of synchronized updates in it’s implementation¹, following the hog-wild style update of parameters [48]. This shows the importance of our register accumulation strategy introduced in subsection 4.6, allowing us to obtain better quality results in these extreme cases.

Table 4: Each result shows mean \pm variance, followed by max trustworthiness score, of each implementation of UMAP for the supervised case with default parameters. Fastest result in **bold**.

Dataset	UMAP-Learn		GPUMAP		cuML UMAP	
	$\mu \pm \sigma^2$	Trust%	$\mu \pm \sigma^2$	Trust%	$\mu \pm \sigma^2$	Trust%
digits	6.756 ± 0.1109	98.76	2.553 ± 1.095	95.55	0.4063 ± 0.0135	98.80
fashion mnist	53.09 ± 6.183	97.81	6.477 ± 0.0632	96.97	1.0370 ± 0.0002	97.76
mnist	89.1877 ± 6.4658	95.85	23.905 ± 7.057	94.69	0.9175 ± 0.00297	95.74
cifar100	98.42 ± 2.273	84.91	5.954 ± 0.0236	83.01	1.0816 ± 0.0003	83.82
coil20	12.34 ± 0.0217	98.68	8.210 ± 0.0275	93.33	0.3695 ± 0.0066	98.70
shuttle	50.17 ± 17.55	100.0	17.15 ± 23.92	96.67	0.5560 ± 0.0111	100.0

The original GPUMAP implementation has, at times, had runtime failures where no results are produced, and significant time was spent attempting to re-compile/fix these issues without success. This prevented it’s use on our largest dataset, Google-News word2vec embeddings. We use the Google-News corpus in particular as a large-scale experiment to show the value of our results, compared to UMAP-learn up to a time limit of 3 hours. The runtime comparing many cores with UMAP-learn to our cuML UMAP on a single GPU is shown in Figure 2. We can clearly see that cuML UMAP continues to dominate runtime with no loss in quality, obtaining $\geq 30\times$ speedups across $n = 1,024$ samples all the way up to the full 3M samples while UMAP-learn reached its time limit after 2.3M samples.

Our cuML UMAP’s 9.5 minutes to process all 3 million datapoints of Google-News is already a significant advantage in runtime. In addition, we note that all non k -NN work of the UMAP algorithm took only 9.3 seconds of that total time. This is important for the interactive and hyperparameter tuning scenarios. Our implementation allows computing the k -NN once, and then other hyperparameter settings can be adjusted with results obtained in seconds. subsection 4.7 briefly




Figure 2: Google-News Results showing runtime (y-axis) as more of the dataset is sampled (x-axis).

¹see <https://umap-learn.readthedocs.io/en/latest/reproducibility.html>

discusses our distributed UMAP, with preliminary results demonstrating an ability to embed 10M points across 8 GPUs in just under 5 seconds with only a marginal impact to trustworthiness. These results can speedup tuning and visualization by orders of magnitude, and is enabled by the optimizations we have contributed.

6 Conclusion

The UMAP algorithm is becoming a widely popular tool, which increases the demand and utility of a faster implementation. We have detailed a number of techniques that are easy to apply in code, and allow us to obtain a solution that is faster and more accurate, even at times compared to the original CPU-based implementation. This obtains up to $100\times$ speedups, and by eliminating all non k -NN calculations to $\leq 2\%$ of runtime, we enable interactive exploration and parameter tuning use-cases that were previously untenable.

Broader Impact Statement

This paper describes an end-to-end GPU-enabled version of the UMAP algorithm that is significantly faster than existing implementations, including GPU-enabled implementations, without sacrificing solution quality. The broader societal impacts of this work stem from the lessons learned in developing our implementation of UMAP, and the scientific advances that a significantly faster UMAP may enable.

UMAP is, for many, the default choice of a manifold learning algorithm. This is in part because, in contrast to other popular algorithms like LLE and t-SNE, it supports unsupervised, supervised, and semi-supervised metric learning. UMAP has seen applications in exploratory data analysis, and medical research in areas as diverse as bioinformatics, cancer, and single cell genomics. Faster and more accurate tools in these areas can speed the pace of innovation. Another application domain of UMAP is interpretation of deep neural networks. Again, the ability to scale up both in terms of the size and number of networks explored in this way makes it possible to peer more deeply inside these black boxes and advance both the science of deep learning and the application domains where these models are used.

Despite UMAP’s popularity and the existence of GPU-enabled implementations, ours represents a significant end-to-end speedup. We believe that the methodology we used can serve as inspiration for those attempting end-to-end GPU implementations of other important machine learning algorithms, again speeding up the pace of scientific advancement in the fields where they are used.

It is difficult to imagine intrinsic negative societal impacts of this work (as opposed to, say, work in face recognition where they are easy to imagine).

Acknowledgement

We extend our sincerest gratitude to all of those who helped enable our research, especially Philip Hynsu Cho and Dante Gama Dessavre from the RAPIDS cuML team. We would also like to thank the Clara Genomics team at Nvidia, especially Avantika Lal, Johnny Israeli, Raghav Mani, and Neha Tadimeti.

References

- [1] L. McInnes, J. Healy, and J. Melville, “Umap: Uniform manifold approximation and projection for dimension reduction,” *arXiv preprint arXiv:1802.03426*, 2018.
- [2] C. Fefferman, S. Mitter, and H. Narayanan, “Testing the manifold hypothesis,” *Journal of the American Mathematical Society*, vol. 29, no. 4, pp. 983–1049, 2016.
- [3] X. He, D. Cai, S. Yan, and H.-J. Zhang, “Neighborhood preserving embedding,” in *Tenth IEEE International Conference on Computer Vision (ICCV’05) Volume 1*, vol. 2. IEEE, 2005, pp. 1208–1213.

- [4] J. B. Tenenbaum, V. De Silva, and J. C. Langford, “A global geometric framework for nonlinear dimensionality reduction,” *science*, vol. 290, no. 5500, pp. 2319–2323, 2000.
- [5] S. T. Roweis and L. K. Saul, “Nonlinear dimensionality reduction by locally linear embedding,” *science*, vol. 290, no. 5500, pp. 2323–2326, 2000.
- [6] M. Belkin and P. Niyogi, “Laplacian eigenmaps and spectral techniques for embedding and clustering,” in *Advances in neural information processing systems*, 2002, pp. 585–591.
- [7] L. v. d. Maaten and G. Hinton, “Visualizing data using t-sne,” *Journal of machine learning research*, vol. 9, no. Nov, pp. 2579–2605, 2008.
- [8] C. Ordun, S. Purushotham, and E. Raff, “Exploratory analysis of covid-19 tweets using topic modeling, umap, and digraphs,” *arXiv preprint arXiv:2005.03082*, 2020.
- [9] L. Wander, A. Vianello, J. Vollertsen, F. Westad, U. Braun, and A. Paul, “Exploratory analysis of hyperspectral ftir data obtained from environmental microplastics samples,” *Analytical Methods*, vol. 12, no. 6, pp. 781–791, 2020.
- [10] B. Obermayer, M. Holtgrewe, M. Nieminen, C. Messerschmidt, and D. Beule, “Scelvis: exploratory single cell data analysis on the desktop and in the cloud,” *PeerJ*, vol. 8, p. e8607, 2020.
- [11] L. Oden, “Lessons learned from comparing c-cuda and python-numba for gpu-computing,” in *2020 28th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP)*. IEEE, 2020, pp. 216–223.
- [12] N. Andor, B. T. Lau, C. Catalanotti, V. Kumar, A. Sathe, K. Belhocine, T. D. Wheeler, A. D. Price, M. Song, D. Stafford *et al.*, “Joint single cell dna-seq and rna-seq of gastric cancer reveals subclonal signatures of genomic instability and gene expression,” *bioRxiv*, p. 445932, 2018.
- [13] K. J. Travaglini, A. N. Nabhan, L. Penland, R. Sinha, A. Gillich, R. V. Sit, S. Chang, S. D. Conley, Y. Mori, J. Seita *et al.*, “A molecular cell atlas of the human lung from single cell rna sequencing,” *bioRxiv*, p. 742320, 2019.
- [14] E. Becht, C.-A. Dutertre, I. W. Kwok, L. G. Ng, F. Ginhoux, and E. W. Newell, “Evaluation of umap as an alternative to t-sne for single-cell data,” *BioRxiv*, p. 298430, 2018.
- [15] N. Clara-Parabricks, “clara-parabricks/rapids-single-cell-examples: Examples of single-cell genomic analysis accelerated with rapids,” <https://github.com/clara-parabricks/rapids-single-cell-examples>, (Accessed on 06/03/2020).
- [16] S. Carter, Z. Armstrong, L. Schubert, I. Johnson, and C. Olah, “Exploring neural networks with activation atlases,” 2019.
- [17] N. Pezzotti, A. Mordvintsev, T. Holtt, B. P. Lelieveldt, E. Eisemann, and A. Vilanova, “Linear tsne optimization for the web,” *arXiv preprint arXiv:1805.10817*, 2018.
- [18] A. Chatzimparmpas, R. M. Martins, and A. Kerren, “t-visne: Interactive assessment and interpretation of t-sne projections,” *arXiv preprint arXiv:2002.06910*, 2020.
- [19] “p3732/gpumap: Gpu-parallelized umap,” <https://github.com/p3732/gpumap>, (Accessed on 06/05/2020).
- [20] 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.
- [21] S. Raschka, J. Patterson, and C. Nolet, “Machine learning in python: Main developments and technology trends in data science, machine learning, and artificial intelligence,” *Information*, vol. 11, no. 4, p. 193, 2020.
- [22] J. Johnson, M. Douze, and H. Jégou, “Billion-scale similarity search with gpus,” *IEEE Transactions on Big Data*, 2019.

- [23] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems*, 2019, pp. 8024–8035.
- [24] D. M. Chan, R. Rao, F. Huang, and J. F. Canny, “t-sne-cuda: Gpu-accelerated t-sne and its applications to modern data,” in *2018 30th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*. IEEE, 2018, pp. 330–338.
- [25] S. K. Lam, A. Pitrou, and S. Seibert, “Numba: A llvm-based python jit compiler,” in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, 2015, pp. 1–6.
- [26] J. Venna and S. Kaski, “Local multidimensional scaling,” *Neural Networks*, vol. 19, no. 6-7, pp. 889–899, 2006.
- [27] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg *et al.*, “Scikit-learn: Machine learning in python,” *the Journal of machine Learning research*, vol. 12, pp. 2825–2830, 2011.
- [28] E. Klement, R. Mesiar, and E. Pap, “Triangular norms,” *Tatra Mountains Math. Publ*, vol. 13, pp. 169–193, 1997.
- [29] D. Dubois and H. Prade, “A class of fuzzy measures based on triangular norms a general framework for the combination of uncertain information,” *International Journal of General Systems*, vol. 8, no. 1, pp. 43–61, 1982.
- [30] S. v. d. Walt, S. C. Colbert, and G. Varoquaux, “The numpy array: a structure for efficient numerical computation,” *Computing in Science & Engineering*, vol. 13, no. 2, pp. 22–30, 2011.
- [31] W. McKinney *et al.*, “pandas: a foundational python library for data analysis and statistics,” *Python for High Performance and Scientific Computing*, vol. 14, no. 9, 2011.
- [32] K. Nuggets, “Apache arrow and apache parquet: Why we needed different projects for columnar data, on disk and in-memory,” 2017.
- [33] J. W. Davidson and S. Jinturkar, “Memory access coalescing: a technique for eliminating redundant memory accesses,” *Acm Sigplan Notices*, vol. 29, no. 6, pp. 186–195, 1994.
- [34] N. Bell and J. Hoberock, “Thrust: A productivity-oriented library for cuda,” in *GPU computing gems Jade edition*. Elsevier, 2012, pp. 359–371.
- [35] M. Naumov, L. Chien, P. Vandermersch, and U. Kapasi, “Cuspars library,” in *GPU Technology Conference*, 2010.
- [36] A. Li, H. Mazhar, R. Serban, and D. Negrut, “Comparison of spmv performance on matrices with different matrix format using cusp, cuspars and viennacl,” Technical Report TR-2015-02–<http://sbel.wisc.edu/documents/TR-2015-02.pdf> ..., Tech. Rep., 2015.
- [37] A. Ocsa, “Sql for gpu data frames in rapids accelerating end-to-end data science workflows using gpus,” 2019.
- [38] M. Harris, “Nvidia developer blog,” Dec 2012. [Online]. Available: <https://devblogs.nvidia.com/how-optimize-data-transfers-cuda-cc/>
- [39] “dmlc/dlpack: Rfc for common in-memory tensor structure and operator interface for deep learning system,” <https://github.com/dmlc/dlpack>, (Accessed on 06/02/2020).
- [40] “The array interface — numpy v1.18 manual,” <https://numpy.org/doc/stable/reference/arrays.interface.html>, (Accessed on 06/02/2020).
- [41] J. D. Owens, M. Houston, D. Luebke, S. Green, J. E. Stone, and J. C. Phillips, “Gpu computing,” *Proceedings of the IEEE*, vol. 96, no. 5, pp. 879–899, 2008.
- [42] D. Luebke, “Cuda: Scalable parallel programming for high-performance scientific computing,” in *2008 5th IEEE international symposium on biomedical imaging: from nano to macro*. IEEE, 2008, pp. 836–838.

- [43] 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*, 2011, pp. 577–586.
- [44] P. Wieschollek, O. Wang, A. Sorkine-Hornung, and H. Lensch, “Efficient large-scale approximate nearest neighbor search on the gpu,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 2027–2035.
- [45] A. Fender, “Parallel solutions for large-scale eigenvalue problems arising in graph analytics,” Ph.D. dissertation, 2017.
- [46] T. Mikolov, G. Corrado, K. Chen, and J. Dean, “Efficient Estimation of Word Representations in Vector Space,” *Proceedings of the International Conference on Learning Representations (ICLR 2013)*, pp. 1–12, 2013. [Online]. Available: <http://arxiv.org/pdf/1301.3781v3.pdf>
- [47] H. Zhang, C.-J. Hsieh, and V. Akella, “HogWild++: A New Mechanism for Decentralized Asynchronous Stochastic Gradient Descent,” in *IEEE International Conference on Data Mining (ICDM)*. ICDM, 2016.
- [48] B. Recht, C. Re, S. Wright, and F. Niu, “Hogwild: A Lock-Free Approach to Parallelizing Stochastic Gradient Descent,” in *Advances in Neural Information Processing Systems 24*, J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2011, pp. 693–701. [Online]. Available: <http://papers.nips.cc/paper/4390-hogwild-a-lock-free-approach-to-parallelizing-stochastic-gradient-descent.pdf>
- [49] C.-J. Hsieh, H.-F. Yu, and I. S. Dhillon, “PASSCoDe: Parallel Asynchronous Stochastic Dual Co-ordinate Descent,” in *Proceedings of the 32Nd International Conference on Machine Learning - Volume 37*, ser. ICML’15. JMLR.org, 2015, pp. 2370–2379. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3045118.3045370>
- [50] E. Raff and J. Sylvester, “Linear Models with Many Cores and CPUs: A Stochastic Atomic Update Scheme,” in *2018 IEEE International Conference on Big Data (Big Data)*. IEEE, 12 2018, pp. 65–73. [Online]. Available: <https://ieeexplore.ieee.org/document/8622172/>
- [51] K. Tran, S. Hosseini, L. Xiao, T. Finley, and M. Bilenko, “Scaling Up Stochastic Dual Coordinate Ascent,” in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD ’15*. New York, New York, USA: ACM Press, 8 2015, pp. 1185–1194. [Online]. Available: <https://www.microsoft.com/en-us/research/publication/scaling-up-stochastic-dual-coordinate-ascent/http://dl.acm.org/citation.cfm?doid=2783258.2783412>
- [52] W.-S. Chin, Y. Zhuang, Y.-C. Juan, and C.-J. Lin, “A Fast Parallel Stochastic Gradient Method for Matrix Factorization in Shared Memory Systems,” *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 6, no. 1, pp. 2:1–2:24, 3 2015. [Online]. Available: <http://doi.acm.org/10.1145/2668133>
- [53] O. Villa, D. Chavarria-Miranda, V. Gurumoorthis, A. Márquez, and S. Krishnamoorthy, “Effects of floating-point non-associativity on numerical computations on massively multithreaded systems,” in *Proceedings of Cray User Group Meeting (CUG)*, 2009, p. 3.
- [54] Y. Bengio, J. François Païement, P. Vincent, O. Delalleau, N. L. Roux, and M. Ouimet, “Out-of-sample extensions for lle, isomap, mds, eigenmaps, and spectral clustering,” in *Advances in Neural Information Processing Systems 16*, S. Thrun, L. K. Saul, and B. Schölkopf, Eds. MIT Press, 2004, pp. 177–184. [Online]. Available: <http://papers.nips.cc/paper/2461-out-of-sample-extensions-for-lle-isomap-mds-eigenmaps-and-spectral-clustering.pdf>
- [55] S. Potluri, H. Wang, D. Bureddy, A. K. Singh, C. Rosales, and D. K. Panda, “Optimizing mpi communication on multi-gpu systems using cuda inter-process communication,” in *2012 IEEE 26th International Parallel and Distributed Processing Symposium Workshops & PhD Forum*. IEEE, 2012, pp. 1848–1857.
- [56] A. Li, S. L. Song, J. Chen, J. Li, X. Liu, N. R. Tallent, and K. J. Barker, “Evaluating modern gpu interconnect: Pcie, nvlink, nv-sli, nvswitch and gpudirect,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 31, no. 1, pp. 94–110, 2019.

- [57] A. Venkatesh, H. Subramoni, K. Hamidouche, and D. K. Panda, "A high performance broadcast design with hardware multicast and gpudirect rdma for streaming applications on infiniband clusters," in *2014 21st International Conference on High Performance Computing (HiPC)*. IEEE, 2014, pp. 1–10.
- [58] P. Shamis, M. G. Venkata, M. G. Lopez, M. B. Baker, O. Hernandez, Y. Itigin, M. Dubman, G. Shainer, R. L. Graham, L. Liss *et al.*, "Ucx: an open source framework for hpc network apis and beyond," in *2015 IEEE 23rd Annual Symposium on High-Performance Interconnects*. IEEE, 2015, pp. 40–43.
- [59] M. D. Garriss, J. L. Blue, G. T. Candela *et al.*, "Nist form-based handprint recognition system," in *Technical Report NISTIR 5469 and CD-ROM, National Institute of Standards and Technology*. Citeseer, 1994.
- [60] "Uci machine learning repository: Shuttle dataset." [Online]. Available: [https://archive.ics.uci.edu/ml/datasets/Statlog\(Shuttle\)](https://archive.ics.uci.edu/ml/datasets/Statlog(Shuttle))
- [61] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms," *arXiv preprint arXiv:1708.07747*, 2017.
- [62] L. Deng, "The mnist database of handwritten digit images for machine learning research [best of the web]," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 141–142, 2012.
- [63] A. Krizhevsky, G. Hinton *et al.*, "Learning multiple layers of features from tiny images," 2009.
- [64] S. A. Nene, S. K. Nayar, H. Murase *et al.*, "Columbia object image library (coil-20)," 1996.
- [65] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.
- [66] F. A. Wolf, P. Angerer, and F. J. Theis, "Scanpy for analysis of large-scale single-cell gene expression data," *BioRxiv*, p. 174029, 2017.
- [67] F. Chollet *et al.*, "Keras: The python deep learning library," *Astrophysics Source Code Library*, 2018.
- [68] B. Tasic, Z. Yao, L. T. Graybuck, K. A. Smith, T. N. Nguyen, D. Bertagnolli, J. Goldy, E. Garren, M. N. Economou, S. Viswanathan *et al.*, "Shared and distinct transcriptomic cell types across neocortical areas," *Nature*, vol. 563, no. 7729, pp. 72–78, 2018.

A API Enhancements

While libraries like Scanpy [66] invoke private Python functions internal to the reference implementation, we maintain compatibility only with UMAP’s Scikit-learn’s public estimators [21] interface. For the remainder of this section, we briefly discuss two enhancements to the standard Scikit-learn API that enable interactive data analysis and visualization workflows.

A.1 Pre-computed k -NN Graph

As mentioned in subsection 4.2 and demonstrated in section 5, the k -NN graph construction stage can quickly become the largest bottleneck to the end-to-end algorithm, eclipsing the remaining stages by orders of magnitude. When many UMAP models need to be trained with different parameters, such as in cluster analysis and hyperparameter-tuning environments, it can be very wasteful to recompute the k -NN graph when neither the training data, distance metric, nor the $n_neighbors$ parameter have changed.

We diverge from the reference API and provide an additional *knn_graph* parameter to *fit*, *fit_transform*, and *transform*. This new parameter allows the k -NN graph to be computed externally and passed into our API, therefore bypassing the computation altogether. This enhancement also makes our implementation more flexible and extensible, since new k -NN libraries can be used, even with distance metrics that are not yet supported.

A.2 Training Callbacks

Inspired by deep learning frameworks like Keras [67], the UMAP API has been enhanced to accept a custom Python function that will be invoked during each epoch of the embeddings optimization stage. This enhancement provides an opportunity to introspect and potentially manipulate the array of actual embeddings in GPU device memory during training. We have found this to be a useful feature that enable interactive visualization tools to provide visual feedback, such as animations, during training.

B Distributed UMAP Experiments

We tested the trustworthiness of our distributed UMAP implementation against the TASIC2018 [68] dataset, which includes approx. 23k cells with 1k genes. We trained a UMAP model on a single GPU using a random sample of the dataset and performed inference over partitions of the remaining data points. Figure B demonstrates that a reasonable trustworthiness can be achieved by training on only 3% of the dataset. Further, the increased variance when the number training samples decreases below 1% appears to create the formation of more dense and tightly packed clusters. Still, we find a marginal impact to trustworthiness as the number of training samples is decreased and as the number of partitions is increased.

We executed performance tests for our distributed UMAP implementation against 10M randomly generated samples using Dask with 1, 2, 4, and 8 workers on a DGX1 containing 8x GV100 GPUs. For each experiment, we trained a UMAP model on 1% of the data and started a timer. The trained model was broadcast to all of the workers in the Dask cluster and inference was performed in parallel. We stopped the timer when the data being inferenced was gathered back on the client. For UMAP-Learn experiments, we set the number of Numba threads to 80 and for cuML UMAPexperiments we mapped each worker to their own GPU. Figure 8 contains the results of this experiment. While both cuML UMAP and UMAP-learn achieve near-linear speedups as workers are added, cuML UMAP dominated with a $255\times$ speedup on a single worker and $100\times$ speedup on 8 workers. UMAP-learn would require 160 CPUs across 80 workers to achieve comparable performance.


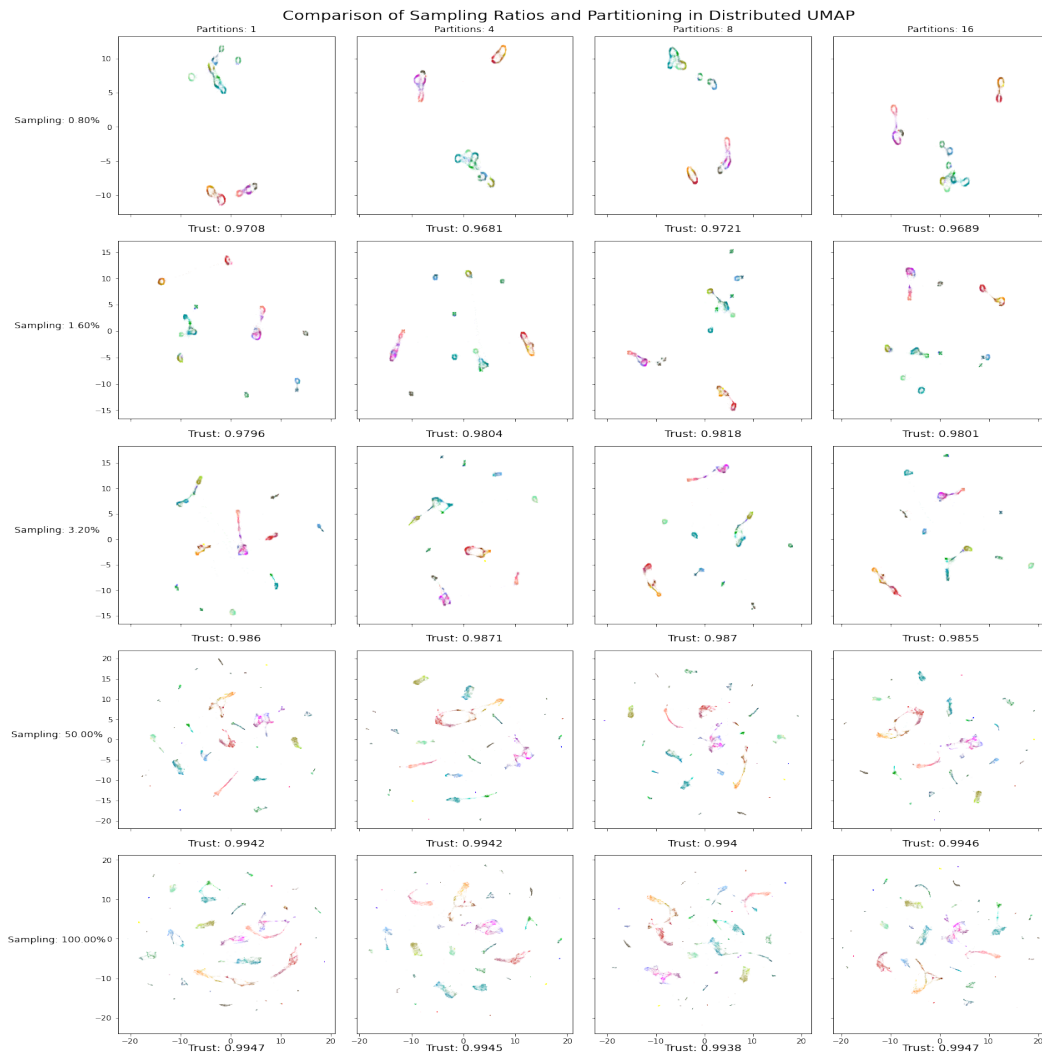


Figure 3: Multi-GPU Scaling



C Experimenting with Neighborhood Sizes

In addition to the unsupervised training experiments conducted with default values Table 3, we tested the three UMAP implementations with extreme values of $n_neighbors = 5$ and $n_neighbors = 50$. Following the experiments in section 5, these were also performed on a DGX1 containing $8 \times 32\text{gb}$ V100 GPUs with $2 \times$ Intel Xeon 20-core CPUs.

Table 5: Each result shows mean \pm variance, followed by max trustworthiness score, of each implementation of UMAP for the unsupervised case with $n_neighbors = 5$. Fastest result in **bold**.

Dataset	UMAP-Learn		GPUMAP		cuML UMAP	
	$\mu \pm \sigma^2$	Trust%	$\mu \pm \sigma^2$	Trust%	$\mu \pm \sigma^2$	Trust%
digits	5.251 ± 2.8944	99.10	2.543 ± 1.559	96.69	0.3764 ± 0.0109	99.23
fashion mnist	29.82 ± 2.7041	98.19	3.932 ± 1.913	97.28	0.5432 ± 0.0001	97.73
mnist	33.63 ± 0.7027	96.30	5.029 ± 1.872	94.70	0.6712 ± 0.0044	96.10
cifar100	66.99 ± 1.307	86.87	4.984 ± 1.861	84.12	0.8252 ± 0.0187	84.42
coil20	9.384 ± 0.001	99.67	3.121 ± 1.317	96.23	0.3274 ± 0.0178	99.44
shuttle	29.88 ± 5.204	96.01	12.73 ± 2.974	93.29	0.6337 ± 0.2727	96.80
scRNA	161.22 ± 6.435	99.85	10.66 ± 2.311	99.88	3.8772 ± 0.0108	99.87

Table 6: Each result shows mean \pm variance, followed by max trustworthiness score, of each implementation of UMAP for the unsupervised case with $n_neighbors = 50$. Fastest result in **bold**.

Dataset	UMAP-Learn		GPUMAP		cuML UMAP	
	$\mu \pm \sigma^2$	Trust%	$\mu \pm \sigma^2$	Trust%	$\mu \pm \sigma^2$	Trust%
digits	7.922 ± 3.0801	98.01	2.423 ± 1.433	96.75	0.6380 ± 0.5867	98.02
fashion mnist	88.06 ± 11.679	96.69	8.428 ± 1.895	97.54	1.0485 ± 0.0029	97.53
mnist	119.96 ± 1.3089	95.65	9.906 ± 2.005	95.46	1.0521 ± 0.0022	95.27
cifar100	222.52 ± 18.213	84.20	13.08 ± 3.214	83.19	1.2524 ± 0.0223	84.11
coil20	12.364 ± 2.651	97.40	FAIL \pm FAIL	FAIL	0.4217 ± 0.0009	97.27
shuttle	11.48 ± 0.0156	97.34	FAIL \pm FAIL	FAIL	0.3486 ± 0.008	97.20
scRNA	392.687 ± 15.024	69.49	FAIL \pm FAIL	FAIL	4.1645 ± 0.008	66.83

D Figures

Function *ScalableTrust* Matrix X , Matrix $embed$, Integer k , Integer $num_batches$

```

forall batch  $batch$  in  $num\_batches$  do
     $nei\_orig = PairwiseDists(X[batch,:]);$ 
     $nei\_embed = KNearestNeighbors(embed, k);$ 
     $t = 0;$ 
    forall row  $c$  in  $nei\_orig$  do
         $t = t + Rank(nei\_orig, nei\_embed, k);$ 
    end
end
return  $t(1 - k(2/(nk * 2n - 3k) - 1));$ 
end

```

Algorithm 1: The pairwise distance computations in our GPU-accelerated trustworthiness implementation are batched to preserve memory.




Figure 4: Fashion MNIST embedded with UMAP-Learn



Figure 5: Fashion MNIST embedded with GPUMAP




Figure 6: Fashion MNIST embedded with cuML UMAP




Figure 7: The GPU Architecture contains global device memory that is accessible by several thread-blocks. Each thread-block contains shared memory which can be accessed by their internal threads. Threads each contain a set of registers.

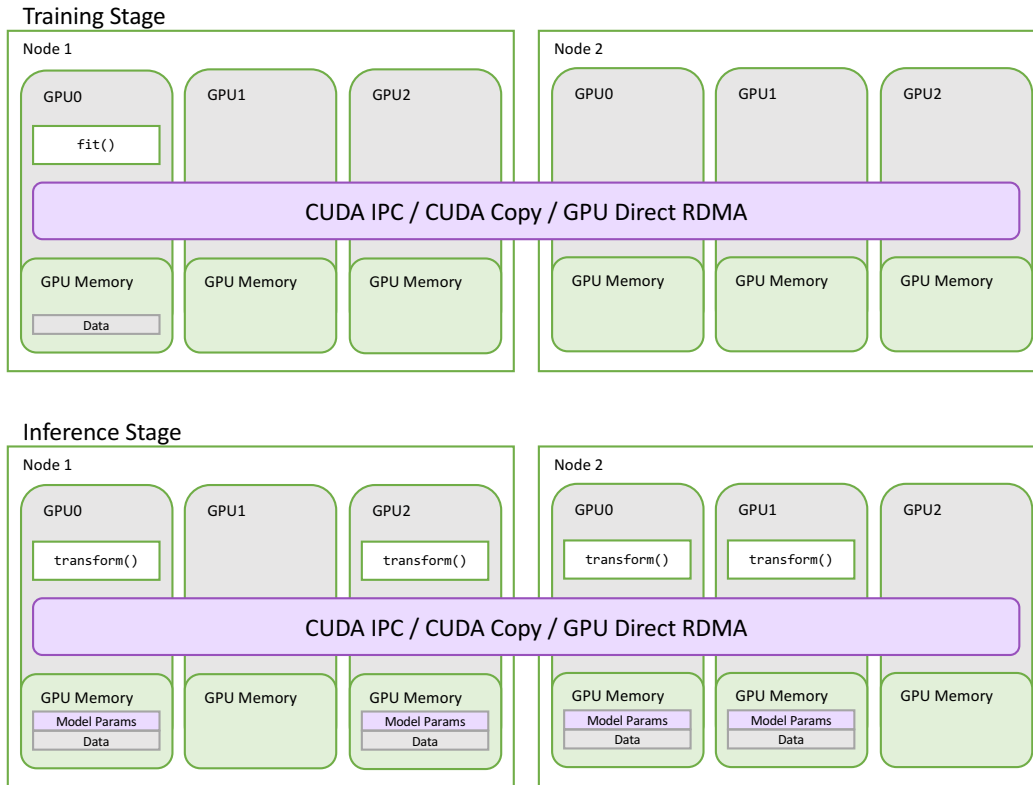


Figure 8: Distributed UMAP is executed on a cluster of workers, each mapped to a single GPU. A subsampling of the training data is used for training the model on a single worker and the model is scattered to workers containing data for out-of-sample prediction. UCX is used to transport GPU memory across the workers.