



Automated Tensor-Relational Decomposition for Large-Scale Sparse Tensor Computation

Yuxin Tang, Zhiyuan Xin, Zhimin Ding, Xinyu Yao, Daniel Bourgeois, Tirthak Patel, Chris Jermaine
Rice University
Houston, Texas, USA
{yt33,zx58,zd21,xy38,dcb10,tp53,cmj4}@rice.edu

ABSTRACT

A *tensor-relational* computation is a relational computation where individual tuples carry vectors, matrices, or higher-dimensional arrays. An advantage of tensor-relational computation is that the overall computation can be executed on top of a relational system, inheriting the system’s ability to automatically handle very large inputs with high levels of sparsity while high-performance kernels (such as optimized matrix-matrix multiplication codes) can be used to perform most of the underlying mathematical operations. In this paper, we introduce upper-case-lower-case EinSum, which is a tensor-relational version of the classical Einstein Summation Notation. We study how to automatically rewrite a computation in Einstein Notation into upper-case-lower-case EinSum so that computationally intensive components are executed using efficient numerical kernels, while sparsity is managed relationally.

PVLDB Reference Format:

Yuxin Tang, Zhiyuan Xin, Zhimin Ding, Xinyu Yao, Daniel Bourgeois, Tirthak Patel, Chris Jermaine. Automated Tensor-Relational Decomposition for Large-Scale Sparse Tensor Computation. PVLDB, 19(6): 1101 - 1114, 2026.
doi:10.14778/3797919.3797921

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/yuxineverforever/upper-case-lower-case-einstein-notation>.

1 INTRODUCTION

In machine learning, a tensor is a specialized relation, mapping a key specifying a position in a multidimensional array to a scalar. Given the close relationship between tensors and relations, almost all computations over tensors can be implemented on top of a relational database system.

For example, consider the matrix multiplication chain $(X \times Y)^T \times Z$. This multiplication is used to implement a message passing scheme in a graph neural network [41, 50]. That is, each column in X stores an embedding for a vertex, Y encodes the (directional) edges between vertices, and the goal is to sum all embeddings being passed over edges to a given vertex, and then to compute a linear transformation (represented by Z) over the sums at each vertex. If

all input matrices are stored as (row, col, value) triples, the SQL corresponding to this computation is as follows:

```
SELECT XtimesY.row, Z.col,
       SUM (XtimesY.value * Z.value) AS value
FROM (SELECT X.row AS col, Y.col AS row,
           SUM (X.value * Y.value) AS value
      FROM X, Y WHERE X.col = Y.row
      GROUP BY X.row, Y.col) AS XtimesY, Z
WHERE XtimesY.col = Z.row
GROUP BY XtimesY.row, Z.col
```

Is the relational implementation effective? Unfortunately, this computation is unlikely to run well on a relational system. Imagine that a large, sparse graph represented by Y has 100 million vertices, each with 1000 neighbors, on average. The embeddings of each vertex (stored in X) are of size 8192, and the transformation Z results in a size 8192 vector. In this case, the first join of X and Y will produce approximately 820 trillion intermediate tuples, which are then aggregated down to 820 billion tuples. This result is then joined with Z , resulting in *6,000 trillion tuples*, a debilitating number.

One could instead run this computation directly using tensors on top of a conventional deep learning system (for example, PyTorch [30]). This would facilitate use of accelerators such as GPU, but the results are likely to be unsatisfactory. There is the problem that storing X requires 3.2TB of GPU RAM (at half precision). This requires 40 GPUs (assuming 80GB of storage each). Then there is the problem of computation. GPUs are not meant for sparse computation, and for very sparse matrices, even state-of-the-art, hand-tuned sparse-dense GPU kernels operate at around 0.1% compute utilization on an A100 GPU [42]. At 0.1% compute utilization, the 1.6×10^{15} floating point operations required for the first matrix multiply would require more than an hour of A100 GPU time.

A sparse tensor-relational implementation. Instead, one could run this computation *tensor-relationally* [46], decomposing the problem so that a relational system can take advantage of the sparsity inherent in the computation, but efficient dense CPU (or GPU) kernels can be used where appropriate:

```
X (col INT, value VECTOR[8192])
Y (row INT, col INT, value DOUBLE)
Z (value MATRIX[8192, 8192])

SELECT XtimesY.row,
       vec_mat_mult (XtimesY.value, Z.value) AS value
FROM (SELECT Y.col AS row,
           SUM (vec_sca_mult (X.value, Y.value)) AS value
      FROM X, Y WHERE X.col = Y.row
      GROUP BY Y.col) AS XtimesY, Z
```

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 19, No. 6 ISSN 2150-8097.
doi:10.14778/3797919.3797921

Kernels (`vec_mat_mult()` and `vec_sca_mult()`) will be very efficient, and because tuples are grouped into VECTOR and MATRIX structures, the number of intermediate tuples is radically reduced. The join of **X** and **Y** will now produce only 100 billion tuples, and the second join will produce only 100 billion tuples.

Further, by running this computation on top of a relational system, memory and parallelization/distribution will be automatically managed. A distributed database system can automatically parallelize the computation across multiple machines, managing memory so that the “Out-Of-Memory” errors ubiquitous in modern machine learning programming are avoided.

Sparse tensor-relational decompositions via upper-case-lower-case EinSum. In the above example, the choice was made to decompose **X** into row vectors, **Y** into scalars, and to leave **Z** undecomposed. The question we address in this paper is: how can we take an arbitrary tensor computation and decompose it into a tensor-relational computation where computationally intensive portions of the computation are handled using efficient kernels (such as `vec_mat_mult()`) but where sparsity is handled relationally?

To address this challenge, we propose a new tensor-relational variant of Einstein summation notation [4], referred to as upper-case-lower-case EinSum. Classic EinSum is a standard tensor calculus used to express machine learning and numerical computations. In upper-case-lower-case EinSum, expressions explicitly specify which parts of the computation are handled relationally and which parts are executed using efficient numerical kernels.

Given a directed acyclic graph of EinSum computations, the task of producing an equivalent, optimized tensor-relational computation can be reduced to automatically rewriting expressions into equivalent upper-case-lower-case EinSum expressions that maximize performance. In this paper, we propose an algorithm to perform this rewrite, called SPARSEEINSUM. The SPARSEEINSUM algorithm combines a simple cost model that evaluates the quality of an upper-case-lower-case EinSum expression under sparsity with a dynamic programming approach to optimize the rewrite.

We demonstrate experimentally that SPARSEEINSUM produces high-performing tensor-relational computations. Our experimental testbed includes a variety of sparse tensor workloads, such as large-scale graph neural networks, graph-based attention computations, and quantum circuit simulation.

2 THE UPPER-CASE-LOWER-CASE EINSUM APPROACH

Assume we are given a (potentially) complex numerical computation specified as a directed acyclic graph (DAG) of EinSum expressions, and our goal is to produce an optimized computation that can be executed on top of a relational system.

EinSum is described in detail in Section 3. At a high level, it asks the programmer to specify a simple declarative expression describing how the entries of an output tensor are computed. For example, consider the matrix multiplication $\mathbf{W} \leftarrow \mathbf{X} \times \mathbf{Y}$. In EinSum, this computation can be expressed as:

$$\forall i, k \mathbf{W}_{i,k} \leftarrow \sum_j \mathbf{U}_{i,j} \times \mathbf{V}_{j,k} \quad (1)$$

By allowing arbitrary aggregation and scalar functions, as well as arbitrary indexing into the input tensors, the output tensor, and the aggregation domain, EinSum can express a very broad class of tensor computations. As a result, EinSum is now widely used in machine learning. Popular frameworks such as NumPy [20], JAX [18], PyTorch [30], and Dask [13] all provide implementations of EinSum. If it is unreasonable to require programmers to write EinSum expressions directly, they may instead specify computations using a PyTorch-like syntax, in which high-level operators (e.g., `matmul`) act as thin wrappers around the underlying EinSum representation.

There is a close relationship between EinSum and SQL/relational databases, in that every EinSum expression can be translated into SQL and implemented as a sequence of joins followed by an aggregation [5]. However, the standard EinSum-to-SQL translation implicitly assumes that all data are stored in fully relational form, with individual scalar values stored in tuples. If not handled carefully, such a decomposition can result in poor performance.

To this end, we introduce the concept of upper-case-lower-case EinSum notation, which serves as the target representation for our SPARSEEINSUM rewrite. In this notation, indices written in upper case are handled relationally (i.e., they are “promoted”), whereas indices written in lower case are handled by tensor indexing (i.e., they are “demoted”). For example, consider the task of multiplying two 8192×8192 matrices. If we write Equation 1 as:

$$\forall i, k \mathbf{W}_{i,k} \leftarrow \sum_j \mathbf{U}_{i,j} \times \mathbf{V}_{j,k}$$

it corresponds to the following, tensor-relational implementation:

```
U (J INT, valU VECTOR[8192])
V (J INT, valV VECTOR[8192])

SELECT SUM (outer_prod (U.valU, V.valV)) AS valW
FROM U, V
WHERE U.J = V.J
```

Note that since *J* is the only index in upper-case notation, it is the only one to appear in the SQL code. The indices *i* and *k* are handled implicitly by the `outer_prod()` kernel function. The output tensor **W** is then stored as a single tuple storing the complete output 8192×8192 matrix. As an alternative, we can write:

$$\forall I, K \mathbf{W}_{I,K} \leftarrow \sum_j \mathbf{U}_{I,j} \times \mathbf{V}_{j,K}$$

This corresponds to the following:

```
U (I INT, valU VECTOR[8192])
V (K INT, valV VECTOR[8192])

SELECT U.I, V.K, inner_prod (U.valU, V.valV) AS valW
FROM U, V
```

Again, since *I* and *K* are handled relationally and *j* is not, *j* does not appear anywhere in the corresponding SQL; it is handled by the kernel function `inner_prod()` which loops over the two vectors `valU` and `valV` along the *j* dimension. The output matrix **W** is stored as a set of scalars indexed by *I* and *K*.

Given the upper-case-lower-case notation, the core problem then becomes: how to re-write a DAG of EinSum expressions into a DAG of upper-case-lower-case EinSum expressions that leverages

a relational engine as appropriate to handle sparsity, and leverages high-performance kernels as appropriate to leverage density?

To this end, the remainder of the paper is organized as follows. The next section of the paper describes EinSum more formally. Then Section 4 describes the upper-case-lower-case notation and how it can take advantage of sparsity. Section 6 describes how upper-case-lower-case EinSum expressions can be translated into SQL and thus executed on existing relational systems. Sections 7 and 8 describe the two major components of the SPARSEEINSUM rewrite algorithm: a cost model for predicting the utility of a re-write into upper-case-lower-case EinSum, and then a simple dynamic programming procedure for optimizing the cost of the re-write. Experiments in Section 9 evaluate the utility of the approach.

3 EINSUM BACKGROUND

The first few paragraphs of this section introduce EinSum and notions such as bounds and label projections, are taken from prior work [8].

We first define the notion of a *tensor*. We use bold upper-case (for example, \mathbf{U}) to denote a tensor. Define the *bound* vector for \mathbf{U} , denoted \mathbf{b}_U , to be a vector of integers of length r . r stands for “rank”. Matrices are rank-2 tensors. Next, define $\mathcal{I}(\mathbf{b}_U)$ to be the set $\{0 \dots \mathbf{b}_U[0] - 1\} \times \{0 \dots \mathbf{b}_U[1] - 1\} \times \dots \times \{0 \dots \mathbf{b}_U[r - 1] - 1\}$. This is the set of all indices or keys that obey the bound. A tensor \mathbf{U} is then a function from $\mathcal{I}(\mathbf{b}_U)$ to the set of real numbers.

Simple EinSum examples. We start with the classic example: matrix multiplication. Let \mathbf{U} and \mathbf{V} be matrices with bounds $[100, 200]$ and $[200, 50]$, respectively. Then matrix multiplication is written as: $\forall i, k \in \mathcal{I}([100, 50])$:

$$\mathbf{W}_{i,k} \leftarrow \sum_{j \in \mathcal{I}([200])} \mathbf{U}_{i,j} \times \mathbf{V}_{j,k} \quad (2)$$

We often drop the subscript on the aggregation as being unnecessary, as indices from the input tensors that are not used to index into the output tensor are, by definition, aggregated out. The bound vector for the output tensor is likewise implied by the bound vectors of the input, and can also be dropped.

This is an *extended* EinSum notation because we need not necessarily multiply items from the input tensors, and we need not necessarily use summation. For example, we can compute the maximum squared difference in each row between any two items in that row as: $\mathbf{W}_i \leftarrow \max(\mathbf{U}_{i,j} - \mathbf{V}_{i,j})^2$.

General form of EinSum. To describe the most general form of EinSum, we define the notion of a *label*, which is a symbol that can be bounded to a value. We use ℓ_U to denote a list (vector) of labels used to index into tensor \mathbf{U} .

We also need to define the *projection* and *permutation* of a bound vector. Given two lists of labels ℓ_1 and ℓ_2 , and a bound vector \mathbf{b} , define $\mathbf{b}[\ell_1; \ell_2]$ to be a vector of length $|\ell_1|$, where the i th entry is $\mathbf{b}[j]$ iff $\ell_1[i] = \ell_2[j]$. As an example, let $\mathbf{b} = [2, 3, 4]$ and let $\ell_1 = [k, i]$ and $\ell_2 = [i, j, k]$. Then $\mathbf{b}[\ell_1; \ell_2] = [4, 2]$.

Given this, binary EinSum expressions take the general form:

$$\forall \ell_W \in \mathcal{I}(\mathbf{b}_W) : \mathbf{W}_{\ell_W} \leftarrow \bigoplus_{\ell_{\text{agg}} \in \mathcal{I}(\mathbf{b}_{UV}[\ell_{\text{agg}}; \ell_{UV}])} \bigotimes (\mathbf{U}_{\ell_U}, \mathbf{V}_{\ell_V}) \quad (3)$$

Here, \bigoplus is the aggregation operator and \bigotimes is the scalar function applied to joined values (EinSum is *extended* as it allows for arbitrary \bigoplus and \bigotimes operations). In the above expression, to denote the concatenation of two label lists ℓ_U and ℓ_V , we use ℓ_{UV} . \mathbf{b}_{UV} similarly denotes the concatenation of two bound vectors.

Consider a more complicated EinSum expression. Assume that we have two tensors \mathbf{U} and \mathbf{V} with bound vectors $\mathbf{b}_U = [10, 100, 20]$ and $\mathbf{b}_V = [100, 20, 2000]$. We wish to transpose \mathbf{U} to obtain a tensor with bound $[20, 10, 100]$, then transpose \mathbf{V} to obtain a new tensor with bound $[20, 100, 2000]$, and do a batch matrix multiply for the two resulting tensors, and then sum out the batch dimension.

In EinSum, this is expressed in the single expression:

$$\forall i, k \in \mathcal{I}([10, 2000]), \mathbf{W}_{i,k} \leftarrow \sum_{b, j \in \mathcal{I}([20, 100])} \mathbf{U}_{i,j,b} \times \mathbf{V}_{j,b,k}$$

Considering the general form, we have $\ell_U = [i, j, b]$, $\ell_V = [j, b, k]$, $\ell_{\text{agg}} = [b, j]$ and $\mathbf{b}_{UV} = [10, 100, 20, 100, 20, 2000]$. The bound vector for the aggregation is computed as $\mathbf{b}_{UV}[\ell_{\text{agg}}; \ell_{UV}]$. How? ℓ_{agg} has two labels: b and j . As b occupies the third (and fifth) position in ℓ_{UV} , and j occupies the second (and fourth) position in ℓ_{UV} , we select the third (or fifth) item in \mathbf{b}_{UV} , and the second (or fourth) item results in \mathbf{b}_{UV} . This results in the bound vector $[20, 100]$.

Multi-Layer Perceptron in EinSum. Using EinSum, we are able to specify very complex computations, including modern transformers, by composing individual EinSum operators together into a directed, acyclic graph of operators. For a simple example of such a DAG, the following specifies a fully-connected neural network with two weight matrices, $\mathbf{W}_{i,j}^{(1)}$ and $\mathbf{W}_{i,j}^{(2)}$ used to connect the input to a hidden layer, and then a softmax:

$$\begin{aligned} \mathbf{A}_i &\leftarrow \sum_j \mathbf{W}_{i,j}^{(1)} \times \mathbf{X}_j; & \mathbf{B}_i &\leftarrow \sum_0 \text{ReLU}(\mathbf{A}_i) & \mathbf{C}_i &\leftarrow \sum_j \mathbf{W}_{i,j}^{(2)} \times \mathbf{B}_j; \\ \mathbf{D}_0 &\leftarrow \sum_i \exp(\mathbf{C}_i); & \mathbf{E}_i &\leftarrow \sum_0 \frac{\exp(\mathbf{C}_i)}{\mathbf{D}_0} \end{aligned}$$

4 SPARSE EINSUM

4.1 Why Worry About Sparsity?

EinSum is agnostic as to the implementation of the underlying tensors, but this implementation can have a significant practical effect on the efficiency of real-life calculations, especially if the underlying tensors are sparse. Consider the following two matrices:

$$\mathbf{U} \equiv \begin{bmatrix} 1.4 & 2.2 & 0 & 2.1 \\ 0 & 0 & 0 & 0 \\ 1.4 & 0 & 1.1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad \mathbf{V} \equiv \begin{bmatrix} 3.2 & 0 & 1.3 & 0 \\ 0 & 0 & 0.6 & 0 \\ 0 & 0 & 1.2 & 0 \\ 1.2 & 0 & 2.1 & 0 \end{bmatrix}$$

If we store these matrices in a classic, dense format such as row-major or column-major, most classical implementations of the matrix multiplication specified by the EinSum $\mathbf{W}_{i,k} \leftarrow \sum \mathbf{U}_{i,j} \times \mathbf{V}_{j,k}$, would perform 64 pairwise scalar multiplications. However, we could store these two matrices *relationally*, and have a much more efficient implementation of the matrix multiplication, at least in terms of the number of multiplications required. Consider representing \mathbf{U} as a set of vectors, where any vectors containing only zero values are not represented explicitly:

$$\tilde{\mathbf{U}} \equiv \{(0, [1.4 \ 2.2 \ 0 \ 2.1]), (2, [1.4 \ 0 \ 1.1 \ 0])\}$$

\bar{U} is implemented as a relation with schema $(i, \text{val}U)$ where $\text{val}U$ is of type `vector[4]`. V is implemented as a set of vectors:

$$\bar{V} \equiv \left\{ \left(0, \begin{bmatrix} 3.2 \\ 0 \\ 0 \\ 1.2 \end{bmatrix} \right), \left(2, \begin{bmatrix} 1.3 \\ 0.6 \\ 1.2 \\ 2.1 \end{bmatrix} \right) \right\} \quad (4)$$

with schema $(k, \text{val}V)$. The same matrix multiplication can now be implemented as a simple SQL query that performs a join with a dot product over the cross product of the two relations:

```
SELECT U.i, V.k,
       dot_product (U.valU, V.valV) AS valW
FROM  $\bar{U}$  AS U,  $\bar{V}$  AS V
```

Note that the result is a fully relational (fully decomposed) implementation of the resulting matrix, with schema $(i, k, \text{val}W)$:

$$W \equiv \{(0, 0, 7), (0, 2, 7.55), (2, 0, 4.48), (2, 2, 3.14)\}$$

Crucially, there are only 16 pairwise scalar multiplications required to perform this implementation, four of which will be invoked by the call to the `dot_product()` operation. Thus, only 25% as many scalar multiplications are required as the traditional way.

4.2 Promoting and Demoting Labels

For any relational implementation, there will be an overhead associated with pushing each tuple through the system, and this overhead may render the original matrix multiplication the best option. Still, there is clearly *some* level of sparsity for which this decomposed representation is preferred due to the lower computational load.

For any tensor, many decompositions into *tensor relations* of this form are possible. Consider an `EinSum` expression over tensor U with bound \mathbf{b}_U and label list ℓ_U . We describe a *decomposition* of U by a partitioning of ℓ_U into two lists, $\uparrow\ell_U$ and $\downarrow\ell_U$:

- (1) The labels in $\uparrow\ell_U$ are *promoted* so that they index into the resulting relation.
- (2) The labels in $\downarrow\ell_U$ are *demoted* so that they index into the tensors within the relation.

The decomposition of U induced by $\uparrow\ell_U$ and $\downarrow\ell_U$ is a database relation with schema

$$\bar{U}(\text{INT } \uparrow\ell_U[1], \text{INT } \uparrow\ell_U[2], \dots, \text{ TENSOR } \text{val}U)$$

In each tuple, $\text{val}U$ is a tensor with bound $\mathbf{b}_U[\downarrow\ell_U; \ell_U]$, and the i th attribute $\uparrow\ell_U[i]$ in the relation takes the integer values from 0 through $\mathbf{b}_U[\uparrow\ell_U; \ell_U][i] - 1$. Crucially, assuming that \oplus and \otimes form a semiring, if $\text{val}U$ is filled only with the special zero value of the semiring, the corresponding tuple can be eliminated from the resulting tensor relation, facilitating sparsity.

The set of all decompositions of a tensor U in an `EinSum` expression defines an equivalence class of all *tensor-relational* implementations for U in the expression, in that all of the tensor-relational implementations describe same mapping from $\mathcal{I}(\mathbf{b}_U)$ to the real numbers. Let \bar{U} be a decomposition of U , defined by $\uparrow\ell_U$ and $\downarrow\ell_U$. Then, for any index $\mathbf{i} \in \mathcal{I}(\mathbf{b}_U)$, the following relational query returns a single tuple containing a scalar with the value at $U_{\mathbf{i}}$:

```
SELECT (ISNULL (valU $\mathbf{i}[\downarrow\ell_U; \ell_U]$ , zero))
FROM (SELECT valU FROM  $\bar{U}$  WHERE  $\uparrow\ell_U = \mathbf{i}[\uparrow\ell_U; \ell_U]$ )
LEFT OUTER JOIN ZERO
```

Here, we assume a relation `ZERO(zero)` that contains a single tuple with the zero value for the semiring. The `WHERE` clause of the inner query uses the set of labels promoted to the relational part of the representation to select the tuple corresponding to \mathbf{i} , and then the outer `SELECT` clause uses the remaining labels to index into the tensor $\text{val}U$ to obtain the desired scalar. If the inner query returns no results, then $\text{val}U$ will be `NULL` and the query result will be the zero value for the associated semiring.

For a concrete example of this, re-consider V , its (possible) tensor-relational representation \bar{V} from this section, and the `EinSum` $W_{i,k} \leftarrow \sum U_{i,j} \times V_{j,k}$. In this case, \bar{V} corresponds to the decomposition $\uparrow\ell_V = k, \downarrow\ell_V = j$. Consider the indexing vector $\mathbf{i} = [3, 2]$. In this case, $\mathbf{i}[\downarrow\ell_V; \ell_V] = \mathbf{i}[j; jk] = 3$ and $\mathbf{i}[\uparrow\ell_V; \ell_V] = \mathbf{i}[k; jk] = 2$. Thus, to obtain $V[3, 2]$ in \bar{V} , we would use the query:

```
SELECT (ISNULL (valV3, zero))
FROM (SELECT valV FROM  $\bar{V}$  WHERE k = 2)
LEFT OUTER JOIN ZERO
```

4.3 “Upper-Case-Lower-Case” Notation

Rather than explicitly listing $\uparrow\ell_U$ and $\downarrow\ell_U$, we suggest the convention that in an `EinSum` expression over tensor relations, any labels in *upper-case* are in $\uparrow\ell_U$, and any labels in *lower-case* are in $\downarrow\ell_U$. This defines the *upper-case-lower-case* notation. So, consider the following `EinSum` expression for a tensor-relational matrix multiplication:

$$W_{i,K} \leftarrow \sum U_{i,J} \times V_{J,K}$$

In this case, labels J, K are used to index into relations, and label i into tensors. So this corresponds to the implementation where U is represented as the set:

$$\left\{ \left(0, \begin{bmatrix} 1.4 \\ 0 \\ 1.4 \\ 0 \end{bmatrix} \right), \left(1, \begin{bmatrix} 2.2 \\ 0 \\ 0 \\ 0 \end{bmatrix} \right), \left(2, \begin{bmatrix} 0 \\ 0 \\ 1.1 \\ 0 \end{bmatrix} \right), \left(3, \begin{bmatrix} 2.1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \right) \right\}$$

and V is represented as the set:

$$\{(0, 0, 3.2), (0, 2, 1.3), (1, 2, 0.6), (2, 2, 1.2), (3, 0, 1.2), (3, 2, 2.1)\}$$

while the output W is again represented as a set of column vectors, as in Equation 4.

For another example, consider the decomposition from the Introduction (the chain of two multiplications used to implement message passing in a graph neural network). Following the convention that those indexes promoted to a relational representation are given in upper case, this can be specified via the `EinSum`:

$$T_{K,i} \leftarrow \sum X_{i,J} \times Y_{J,K}; U_{I,k} \leftarrow \sum T_{I,j} \times Z_{j,k}$$

5 DECOMPOSING EINSUM

We can decompose any binary `EinSum` expression into an equivalent computation over tensor relations. This equivalence gives a roadmap to implement tensor-relational `EinSum` over sparse inputs relationally. Using the convention of the last section of the paper, \bar{U} and \bar{V} are used to denote the tensor-relational decompositions of U and V , respectively. Now, assume that we have an associated *kernel function* \mathcal{K} . This is a function that implements the `EinSum`:

$$W_{\downarrow\ell_W} \leftarrow \bigoplus \bigotimes (U_{\uparrow\ell_U}, V_{\downarrow\ell_V})$$

Re-write equation 3 as:

$$\mathbf{W}_{\ell_W} \leftarrow \bigoplus_{\downarrow \ell_{agg}, \uparrow \ell_{agg}} \bigotimes \left((\sigma_{\uparrow \ell_U}(\bar{\mathbf{U}}))_{\downarrow \ell_U}, (\sigma_{\uparrow \ell_V}(\bar{\mathbf{V}}))_{\downarrow \ell_V} \right) \quad (5)$$

A bit of additional notation is introduced here. Given a binding of the labels in $\uparrow \ell_U$ to values in the vector \mathbf{i} during the execution of the EinSum expression, $\sigma_{\uparrow \ell_U}(\bar{\mathbf{U}})$ represents the application of relational selection equivalent to running the query:

```
SELECT ISNULL (valU, zero)
FROM (SELECT valU FROM  $\bar{\mathbf{U}}$  WHERE  $\uparrow \ell_U = \mathbf{i}$ )
LEFT OUTER JOIN ZERO
```

Note that zero is a tensor, filled with zeros, with bound $\mathbf{b}[\downarrow \ell_U; \ell_U]$. Further, $\uparrow \ell_{agg}$ and $\downarrow \ell_{agg}$ represent the aggregation labels in the original expression partitioned into two lists, depending upon whether they appear in $\uparrow \ell_U, \uparrow \ell_V$ or in $\downarrow \ell_U, \downarrow \ell_V$.

Given this, Equation 5 can be re-written as:

$$\mathbf{W}_{\ell_W} \leftarrow \bigoplus_{\uparrow \ell_{agg}} \left(\left(\sigma_{\uparrow \ell_W, \uparrow \ell_{agg}} (\bowtie (\bar{\mathbf{U}}, \bar{\mathbf{V}}, \text{valW} \leftarrow \mathcal{K}(\text{valU}, \text{valV}))) \right)_{\downarrow \ell_W} \right)$$

Here, $\bowtie (\bar{\mathbf{U}}, \bar{\mathbf{V}}, \text{valW} \leftarrow \mathcal{K}(\cdot))$ performs a natural join of the relations $\bar{\mathbf{U}}$ and $\bar{\mathbf{V}}$ and computes additional attribute valW using the kernel function \mathcal{K} . Define $\bigoplus_{\uparrow \ell_{agg}}$ relationally so that it aggregates the tensors stored in valW by applying the \oplus operator pairwise to each item in a pair of tensors, then this last expression can be re-written as $\mathbf{W}_{\ell_W} \leftarrow$

$$\left(\sigma_{\uparrow \ell_W} \left(\bigoplus_{\uparrow \ell_{agg}} (\bowtie (\bar{\mathbf{U}}, \bar{\mathbf{V}}, \text{valW} \leftarrow \mathcal{K}(\text{valU}, \text{valV}))) \right) \right)_{\downarrow \ell_W} \quad (6)$$

This suggests a tensor-relational implementation for any binary EinSum expression over tensor relations:

```
SELECT  $\uparrow \ell_W$ , SUM ( $\mathcal{K}(\text{valU}, \text{valV})$ ) AS valW
FROM  $\bar{\mathbf{U}}$  NATURAL JOIN  $\bar{\mathbf{V}}$ 
GROUP BY  $\uparrow \ell_W$ 
```

6 COMPILING EINSUM INTO TENSOR-RELATIONAL SQL

This suggests a framework for compiling a DAG of EinSum expressions into tensor-relational algebra or tensor-relational SQL.

- (1) For each vertex (EinSum) expression in the DAG, re-write the EinSum into upper-case-lower-case EinSum.
- (2) For each resulting expression, generate SQL according to Equation 6. That is, generate a SELECT-FROM-WHERE-GROUP BY query where (a) the attributes chosen in the SELECT are those labels in $\uparrow \ell_W$ as well as “SUM($\mathcal{K}(\text{valU}, \text{valV})$) as valW”; (b) the WHERE clause performs a natural join of the two input relations; (c) we group by those attributes corresponding to labels in $\uparrow \ell_W$.
- (3) If there are mismatches in tensor-relational decomposition across decomposed EinSum expressions, generate tensor-relational algebra/tensor-relational SQL to change representation.

The last point bears some discussion. Imagine, for example, that we generated the following decomposition:

$$\mathbf{T}_{K,i} \leftarrow \sum \mathbf{X}_{i,J} \times \mathbf{Y}_{J,K} \quad \mathbf{U}_{i,k} \leftarrow \sum \mathbf{T}_{i,J} \times \mathbf{Z}_{J,k}$$

SQL for the first expression would be:

```
SELECT K, SUM ( $\mathcal{K}_1(\text{valX}, \text{valY})$ ) AS valT
FROM  $\bar{\mathbf{X}}$  NATURAL JOIN  $\bar{\mathbf{Y}}$ 
GROUP BY K
```

SQL for the second expression would be:

```
SELECT SUM( $\mathcal{K}_2(\text{valT}, \text{valZ})$ ) AS valU
FROM  $\bar{\mathbf{T}}$  NATURAL JOIN  $\bar{\mathbf{Z}}$ 
```

The output decomposition of \mathbf{T} from the first EinSum (a set of row vectors, as the first [row] index is upper-case) does not match the input decomposition of \mathbf{T} to the second EinSum (a set of column vectors, as the second [column] index is upper case). Thus, code must be generated to repartition \mathbf{T} between the operations. Assume the output of the first EinSum expression—partitioned according to $\mathbf{T}_{K,i}$, named as $\bar{\mathbf{T}}_{in}$, and we want to repartition according to $\mathbf{T}_{i,j}$ to form $\bar{\mathbf{T}}$. We would first decompose in the first dimension:

```
CREATE VIEW  $\bar{\mathbf{T}}_{int}(i, j, \text{valT})$  AS
SELECT  $\bar{\mathbf{T}}_{in}.k$ , A.index,  $\bar{\mathbf{T}}_{in}.\text{valT}[A.index]$ 
FROM ALLINTS (0,  $\mathbf{b}_T[1]$ ) AS A,  $\bar{\mathbf{T}}_{in}$ 
```

In this code, $\bar{\mathbf{T}}_{in}.\text{valT}[A.index]$ extracts the sub-tensor (here a scalar) at position A.index from $\bar{\mathbf{T}}_{in}.\text{valT}$. ALLINTS is a table function returning all values in the specified range. And then we re-compose on the second dimension:

```
CREATE VIEW  $\bar{\mathbf{T}}(j, \text{valT})$  AS
SELECT  $\bar{\mathbf{T}}_{int}.j$ , STACK ( $\bar{\mathbf{T}}_{int}.\text{valT}$ ,  $\bar{\mathbf{T}}_{int}.i$ , 0,  $\mathbf{b}_T[0]$ )
FROM  $\bar{\mathbf{T}}_{int}$ 
GROUP BY  $\bar{\mathbf{T}}_{int}.j$ 
```

Here, STACK is an aggregation function that can be used to stack scalars (or tensors) to create a higher-dimensional tensor. It accepts: (1) the value to stack, (2) the position of this value in the higher-dimensional tensor, (3) the new dimension to create, and (4) the maximum value of this dimension.

In the general case, the repartition from tensor relation $\bar{\mathbf{T}}_{in}$ to $\bar{\mathbf{T}}$ happens in two steps. First, there is a decomposition of the result of the first EinSum expression so that it is partitioned on *all* of the labels in either $\uparrow \ell_{T_{in}}$ or $\uparrow \ell_T$. Then there is an “un-partitioning” or aggregation to obtain $\bar{\mathbf{T}}$. Note that if $\uparrow \ell_{T_{in}}$ is a subset of $\uparrow \ell_T$, the aggregation can be skipped, and if $\uparrow \ell_T$ is a subset of $\uparrow \ell_{T_{in}}$ the decomposition can be skipped. If $\uparrow \ell_{T_{in}} = \uparrow \ell_T$, then the entire repartition can be skipped, as the partitionings are the same.

7 THE SPARSE EINSUM COST MODEL

This leaves two final questions, as our goal is to find the best decomposition. (1) How do we cost a given decomposition in sparsity-aware fashion? (2) How to efficiently search the space of all decompositions? In this section, we consider the first question.

7.1 Estimating Sizes Under Sparsity

When developing a cost model, the first question we must address is: given a (potentially sparse) tensor \mathbf{U} with label set ℓ_U and a tensor-relational decomposition $\bar{\mathbf{U}}$ defined by $\uparrow \ell_U$ and $\downarrow \ell_U$, how

many tuples are in \bar{U} ? That is, if we take a tensor and decompose it into a tensor relation, can we estimate how many tuples will be contained in the relation? We define this quantity using the notation $T(\bar{U})$.

We begin by defining two statistics that are sufficient to accurately estimate $T(\bar{U})$. First, for a label l in ℓ_U , let $V(l, \mathbf{U})$ denote the number of sub-tensors induced by l that have at least one non-zero entry, if we were to use $\uparrow\ell_U = \{l\}$. In other words, $V(l, \mathbf{U})$ is the answer to the question: If we were to decompose \mathbf{U} by representing label l relationally, how many tuples would we have?

To be precise, let ℓ'_U denote ℓ_U with l removed. Then $V(l, \mathbf{U})$ can be computed using the following EinSum expressions:

$$\begin{aligned} \mathbf{V}_l &\leftarrow \sum_{\ell'_U} \text{if}(\mathbf{U}_{\ell_U} == 0, 0, 1) \\ V(l, \mathbf{U}) &\leftarrow \sum_l \text{if}(\mathbf{V}_l == 0, 0, 1) \end{aligned}$$

Here, the function “if (bool, val1, val2)” returns val1 if bool is true, and val2 otherwise.

Note that $V(l, \mathbf{U})$ is in some sense equivalent to the notion of the number of distinct values for an attribute that is commonly used in relational query optimization. Why? Imagine that \mathbf{U} were represented purely relationally—that is, we perform a tensor-relational decomposition of \mathbf{U} into \bar{U} so that $\uparrow\ell_U = \ell_U$ and $\downarrow\ell_U = \{\}$, and remove any tuples storing the scalar value zero. Now consider attribute l . If l has d distinct values in \bar{U} , it means that there were to group all of the tuples having the same l value into a sub-tensor of \bar{U} , there would be exactly d sub-tensors, and none of those would be composed entirely of zeros (any sub-tensors composed purely of zeros would not contribute any tuples to \bar{U} , and hence would not contribute to the count d). Thus, $V(\mathbf{U}, l)$ would be d .

Further, assume we have $T(\mathbf{U})$, which is the number of non-zero entries in tensor \mathbf{U} . Again, this is analogous to the number of tuples in a relation. If we perform a pure relational decomposition of \mathbf{U} and remove tuples storing zero values, then the number of non-zero entries in the tensor \mathbf{U} would be the same as the number of tuples in such a pure relational decomposition.

Let $n(\uparrow\ell_U)$ denote the number of possible tuples in some tensor-relational decomposition \bar{U} for \mathbf{U} , induced by the promoted label set $\uparrow\ell_U$. This can be computed as:

$$n(\uparrow\ell_U) = \prod_{l \in \uparrow\ell_U} V(l, \mathbf{U})$$

Then the number of tuples in the tensor-relational decomposition \bar{U} induced by the set of promoted labels $\uparrow\ell_U$ can be estimated as:

$$T(\uparrow\ell_U) = n(\uparrow\ell_U) \left(1 - \exp\left(\frac{-T(\mathbf{U})}{n(\uparrow\ell_U)}\right) \right) \quad (7)$$

The intuition is that there are $n(\uparrow\ell_U)$ possible tuples in the tensor relation \bar{U} . $T(\uparrow\ell_U)$ estimates the number of tuples that have tensors with at least one non-zero value, if the $T(\mathbf{U})$ non-zero values are spread uniformly-at-random throughout the tensor relation.

7.2 The Cost Model

When implementing a generic, binary EinSum expression of the form of Equation 3 tensor-relationally, there are two operations, a

join and an aggregation. There is a third operation—repartition—that may be necessary to change tensor-relational decompositions between the output of one operation, and the input of the next. We consider how to cost each of them.

Costing the join. The set of labels participating in the relational join ℓ_{join} is the list of labels in both $\uparrow\ell_U$ and $\uparrow\ell_V$. The number of tuples resulting from the tensor-relational join can be estimated as:

$$T_{\text{join}}(\uparrow\ell_U, \uparrow\ell_V) = \frac{T(\uparrow\ell_U)T(\uparrow\ell_V)}{\prod_{l \in \ell_{\text{join}}} \max(V(\mathbf{U}, l), V(\mathbf{V}, l))}$$

This is an application of the standard, textbook estimator for the number of tuples resulting from a join. Note that if there are no labels common to $\uparrow\ell_U$ and $\uparrow\ell_V$ so that ℓ_{join} is empty, it is assumed that the denominator takes the value 1. Then the cost of the tensor-relational join is estimated as:

$$\begin{aligned} C_{\text{join}}(\uparrow\ell_U, \uparrow\ell_V) = \\ T_{\text{join}}(\uparrow\ell_U, \uparrow\ell_V) \times [(\text{sizeof}(\bar{U}) + \text{sizeof}(\bar{V}))C_{\text{xfer}} + C_{\text{kernel}} + C_{\text{fixed}}] \end{aligned}$$

Here, `sizeof` computes the size of a tuple from the tensor relation, in bytes, C_{xfer} is the per-byte cost to move a tuple from machine to machine (assuming a distributed environment), C_{kernel} is the cost to invoke the kernel \mathcal{K} (typically this is the number of floating-point operations required by \mathcal{K} , times a constant), and C_{fixed} is a fixed, per-tuple processing cost. Overall, the cost is then: (a) the cost to move tuples from both inputs into the join; plus (b) the cost to run the kernel, plus (b) a fixed cost per tuple.

A rule-of-thumb is to choose the ratio C_{xfer} by checking how long it takes to transfer a byte from machine to machine (or from disk to machine in a single-machine environment), to choose C_{kernel} by checking how long it takes to run a floating point computation on the execution platform, and to choose C_{fixed} by running a single-machine computation that produces and aggregates a large number of tuples (say, by running a cross product of two relations and aggregating the result) and dividing the runtime by the number of tuples produced to obtain the per-tuple overhead.

Costing the aggregation. The number of tuples resulting from the subsequent aggregation can be estimated as

$$\begin{aligned} T_{\text{agg}}(\uparrow\ell_U, \uparrow\ell_V) = \min \left(\frac{1}{2} T_{\text{join}}(\uparrow\ell_U, \uparrow\ell_V), \right. \\ \left. \prod_{l \in \uparrow\ell_{\text{agg}}} \begin{cases} \text{if } l \in \uparrow\ell_U \wedge l \in \uparrow\ell_V : \min(V(\mathbf{U}, l), V(\mathbf{V}, l)) \\ \text{else if } l \in \uparrow\ell_U : V(\mathbf{U}, l) \\ \text{else} : V(\mathbf{V}, l) \end{cases} \right) \end{aligned}$$

This is a variation on the standard, textbook estimator for the number of tuples resulting from an aggregation. Note the three cases in the second option of the “min”. The number of distinct values for each promoted grouping label is either (a) the minimum of the number of distinct values from both \mathbf{U} and \mathbf{V} if the label is a join attribute; (b) the number of distinct values from \mathbf{U} if the label comes only from \mathbf{U} ; (c) the number of distinct values from \mathbf{V} , otherwise. Then the cost of the aggregation is estimated as:

$$\begin{aligned} C_{\text{agg}}(\uparrow\ell_U, \uparrow\ell_V) = \\ (T_{\text{join}}(\uparrow\ell_U, \uparrow\ell_V) - T_{\text{agg}}(\uparrow\ell_U, \uparrow\ell_V)) \times [\text{sizeof}(\bar{\mathbf{W}})C_{\text{xfer}} + C_{\text{add}} + C_{\text{fixed}}] \end{aligned}$$

This formula mirrors the formula for C_{join} , the one difference being that the output of the aggregation is $T_{\text{agg}}(\uparrow\ell_U, \uparrow\ell_V)$ “aggregation groups”; reducing the $\frac{T_{\text{join}}(\uparrow\ell_U, \uparrow\ell_V)}{T_{\text{agg}}(\uparrow\ell_U, \uparrow\ell_V)}$ tuples in a aggregation group down to a single tuple requires moving and adding $\frac{T_{\text{join}}(\uparrow\ell_U, \uparrow\ell_V)}{T_{\text{agg}}(\uparrow\ell_U, \uparrow\ell_V)} - 1$ tuples, or $T_{\text{join}}(\uparrow\ell_U, \uparrow\ell_V) - T_{\text{agg}}(\uparrow\ell_U, \uparrow\ell_V)$ moves-and-adds overall.

Costing the repartition. Finally, there is the cost to repartition tensor relation $\bar{\mathbf{W}}$ when the output partition of the expression producing $\bar{\mathbf{W}}$ does not match the input to the consumer of $\bar{\mathbf{W}}$.

Let $\uparrow\ell_{\mathbf{W}_{\text{in}}}$ be the set of promoted labels in the input to the repartition, and let $\uparrow\ell_{\mathbf{W}}$ be the set of promoted labels in the target, consuming EinSum expression.

As illustrated in the SQL given in the prior section, the typical implementation of a repartition first decomposes the result of the first EinSum expression so that it is partitioned on *all* of the labels in either $\uparrow\ell_{\mathbf{W}_{\text{in}}}$ or $\uparrow\ell_{\mathbf{W}}$ (while both $\uparrow\ell_{\mathbf{W}_{\text{in}}}$ and $\uparrow\ell_{\mathbf{W}}$ are lists and not sets, we abuse notation and denote the list containing all labels in both as $\uparrow\ell_{\mathbf{W}_{\text{in}}} \cup \uparrow\ell_{\mathbf{W}}$). The number of tuples in the resulting tensor relation can be estimated using $T(\uparrow\ell_{\mathbf{W}_{\text{in}}} \cup \uparrow\ell_{\mathbf{W}})$. Note that this quantity cannot be less than either $T(\uparrow\ell_{\mathbf{W}})$ or $T(\uparrow\ell_{\mathbf{W}_{\text{in}}})$, as the intermediate decomposition uses a set of promoted labels that is a superset of both the input and output set of promoted labels.

Then it is possible to compute an expression for the cost C_{repart} :

$$\begin{aligned} C_{\text{repart}}(\uparrow\ell_{\mathbf{W}_{\text{in}}}, \uparrow\ell_{\mathbf{W}}) &= \\ & (T(\uparrow\ell_{\mathbf{W}_{\text{in}}} \cup \uparrow\ell_{\mathbf{W}}) - T(\uparrow\ell_{\mathbf{W}_{\text{in}}}))(\text{sizeof}(\bar{\mathbf{W}}_{\text{int}})C_{\text{xfer}} + C_{\text{fixed}}) + \\ & (T(\uparrow\ell_{\mathbf{W}_{\text{in}}} \cup \uparrow\ell_{\mathbf{W}}) - T(\uparrow\ell_{\mathbf{W}}))(\text{sizeof}(\bar{\mathbf{W}}_{\text{int}})C_{\text{xfer}} + C_{\text{fixed}}) \\ & = (2T(\uparrow\ell_{\mathbf{W}_{\text{in}}} \cup \uparrow\ell_{\mathbf{W}}) - T(\uparrow\ell_{\mathbf{W}_{\text{in}}}) - T(\uparrow\ell_{\mathbf{W}})) \times \\ & (\text{sizeof}(\bar{\mathbf{W}}_{\text{int}})C_{\text{xfer}} + C_{\text{fixed}}) \end{aligned}$$

C_{repart} has two parts. First is the cost of creating and processing the $T(\uparrow\ell_{\mathbf{W}_{\text{in}}} \cup \uparrow\ell_{\mathbf{W}}) - T(\uparrow\ell_{\mathbf{W}_{\text{in}}})$ new tuples during the first, decomposition step. And second is aggregating those tuples into $T(\uparrow\ell_{\mathbf{W}})$ inputs to the next operation; this requires moving around $T(\uparrow\ell_{\mathbf{W}_{\text{in}}} \cup \uparrow\ell_{\mathbf{W}}) - T(\uparrow\ell_{\mathbf{W}})$ intermediate tuples.

8 OPTIMIZING THE DECOMPOSITION

We now define a dynamic programming algorithm that computes the lowest cost, sparsity-aware, decomposition of a DAG of EinSum expressions. This algorithm re-writes each expression in the DAG into upper-case-lower-case notation. Combined with the results from Section 5, this allows such an EinSum program to be compiled into a high-performance, tensor-relational computation.

8.1 Dynamic Programming

Consider a directed graph $G = \langle V, E \rangle$, where each vertex in the vertex set V represents an EinSum operation. V has two types of vertices: (1) those with no inputs (these correspond to input tensors) and (2) those with two inputs (these correspond to binary EinSum expressions). For our optimization algorithm, the important features of each EinSum operation are the sets of labels ℓ_U, ℓ_V, ℓ_W present in the expression, as well as the bound vectors of the input and output tensors. These will allow us to make use of the cost models from the previous section.

Each directed edge in E is a triple consisting of (1) the vertex/EinSum expression producing a tensor, (2) the vertex/EinSum

Algorithm 1 COMPUTEMINCOSTFORTREE

```

1: topologically sort the vertices of the input EinSum graph  $G$ 
2: initialize lookup cost table  $C$  so that for any vertex  $v$  with no
   inputs,  $C[v, \uparrow\ell_{\mathbf{W}}] = 0$  for all  $\uparrow\ell_{\mathbf{W}} \in \ell_{\mathbf{W}}$ , where  $\ell_{\mathbf{W}}$  is the list of
   labels for the tensor  $\mathbf{W}$  produced by  $v$ 
3: for each vertex  $v \in V$ , in sorted order do
4:   % input tensors always have cost 0
5:   if  $v$  has no inputs then
6:     continue
7:   end if
8:   find  $v_U$  s.t.  $(v_U, v, \text{left}) \in E$ 
9:   find  $v_V$  s.t.  $(v_V, v, \text{right}) \in E$ 
10:   $\text{cost}_{\min} \leftarrow \infty$ 
11:  extract  $\ell_U, \ell_V, \ell_W$  from vertex  $v$ 
12:  % Here it will consider every output decomp for  $\mathbf{W}$ 
13:  for each subset  $\uparrow\ell_{\mathbf{W}} \in \ell_{\mathbf{W}}$  do
14:    %  $\uparrow\ell_U$  is consistent with  $\uparrow\ell_{\mathbf{W}}$  if, for each  $l$  in  $(\ell_U \cap \ell_{\mathbf{W}})$ ,
15:    %  $l$  is in  $\uparrow\ell_U$  iff it is also in  $\uparrow\ell_{\mathbf{W}}$ 
16:    for each subset  $\uparrow\ell_U \in \ell_U$  that is consistent with  $\uparrow\ell_{\mathbf{W}}$  do
17:       $\text{cost}_U \leftarrow \min_{\uparrow\ell_{U_{\text{in}}}} C_{\text{repart}}(\uparrow\ell_{U_{\text{in}}}, \uparrow\ell_U) + C[v_U, \uparrow\ell_{U_{\text{in}}}]$ 
18:      for each subset  $\uparrow\ell_V \in \ell_V$  that is consistent with  $\uparrow\ell_{\mathbf{W}}$  do
19:         $\text{cost}_V \leftarrow \min_{\uparrow\ell_{V_{\text{in}}}} C_{\text{repart}}(\uparrow\ell_{V_{\text{in}}}, \uparrow\ell_V) + C[v_V, \uparrow\ell_{V_{\text{in}}}]$ 
20:         $\text{cost}_W \leftarrow \text{cost}_U + \text{cost}_V +$ 
21:           $C_{\text{join}}(\uparrow\ell_U, \uparrow\ell_V) + C_{\text{agg}}(\uparrow\ell_U, \uparrow\ell_V)$ 
22:        if  $\text{cost}_W < \text{cost}_{\min}$  then
23:           $\text{cost}_{\min} \leftarrow \text{cost}_W$ 
24:        end if
25:      end for
26:    end for
27:    % remember the low cost to compute  $\mathbf{W}$  via decomp  $\uparrow\ell_{\mathbf{W}}$ 
28:     $C[v, \uparrow\ell_{\mathbf{W}}] \leftarrow \text{cost}_{\min}$ 
29:  end for
30: let  $v$  be the last vertex in  $V$ 
31: return  $\min_{\uparrow\ell_{\mathbf{W}}} C[v, \uparrow\ell_{\mathbf{W}}]$ 

```

expression consuming that same tensor, and (3) a constant that is either *left* or *right* indicating if this tensor is the left or right input into the consuming vertex.

The dynamic programming algorithm works its way through the graph, considering vertices in topologically-sorted order. As it does so, it populates a lookup table C . For a vertex v producing a tensor with label list $\ell_{\mathbf{W}}$, $C[v, \uparrow\ell_{\mathbf{W}}]$ will be computed for every possible set of promoted labels $\uparrow\ell_{\mathbf{W}}$ —this is all possible subsets of $\ell_{\mathbf{W}}$. $C[v, \uparrow\ell_{\mathbf{W}}]$ stores the minimum possible cost to produce tensor \mathbf{W} —considering all of the EinSum expressions necessary to compute \mathbf{W} —subject to the constraint that the tensor-relational decomposition of the output tensor is described by $\uparrow\ell_{\mathbf{W}}$.

It is possible to populate C in a single pass through the graph via a dynamic programming algorithm because as long as each tensor is consumed by exactly one EinSum expression, we do not care about the details of how the output associated with a vertex v is computed; we only care about (a) the output decomposition resulting from v , and (b) the cost to compute v . As long as C stores the lowest cost to

Algorithm 2 GETSTATS

- 1: **compile** input EinSum DAG into SQL over relations storing only scalars—that is, let $\uparrow\ell_U = \ell_U$ and $\downarrow\ell_U = \{\}$ for each tensor U , and compute statistics assuming that all tuples storing zero-valued scalars are removed from the resulting relations
 - 2: **compile** purely-relational SQL into relational algebra using classical methods
 - 3: **cost** relational algebra using classical methods, to compute/estimate the number of tuples for each input and intermediate relation; for relation \bar{U} corresponding to tensor U , let $T(U)$ denote this tuple count
 - 4: **cost** relational algebra using classical methods, to compute/estimate the number of distinct values for each attribute; for attribute l and relation \bar{U} , let $V(U, l)$ denote this value count
-

compute all decompositions of an EinSum expression corresponding to vertex v , to optimize the decomposition of a consumer of v , we need only look at all entries of the form $C[v, \cdot]$.

The dynamic programming algorithm is given as Algorithm 1. For each vertex v there is a triply-nested loop. At the outer level, we are computing the optimal cost for every possible tensor-relational decomposition of the output of v , described by the set of promoted labels $\uparrow\ell_W$. The goal is to compute the lowest cost to produce $\uparrow\ell_W$ for v . The two inner loops consider all possible input decompositions that are *consistent* with $\uparrow\ell_W$. If $\uparrow\ell_W$ promotes a label l so that it is handled relationally, then both inputs must promote that label. By considering all of these consistent input combinations, the optimal value for $C[v, \uparrow\ell_W]$ is computed.

8.2 Additional Considerations

Pre-computing the required statistics. In practice, running Algorithm 1 requires that for each tensor U , we have $T(U)$ (the number of non-zero entries in U), and that we have $V(U, l)$ for each label l (the number of non-zero tensors induced by label l). Obtaining these statistics can be achieved by running the GETSTATS algorithm: Algorithm 2. At the highest level, this algorithm relies on generating SQL corresponding to a fully-relational version of the input EinSum program, and then using classical relational costing methods to compute the required statistics.

When results are consumed more than once. The prior dynamic programming algorithm cannot work in this case, as the algorithm relies on the set of entries $C[v, \cdot]$ fully describing all information necessary to optimally use the EinSum expression corresponding to vertex v . However, data in machine learning computations (both input and intermediate) are often used multiple times. For example, during backpropagation, intermediate data are used both during the forward and backward passes. In both the graph convolutional neural network and the attention computation in our experiments, input data are used in two EinSum expressions. If v is used more than once, an optimal algorithm would need to worry about both uses of v , and co-optimize those.

The approach we implement is to decompose an EinSum program into a set of “trees”. That is, we first choose a subgraph of program G where no EinSum result is used more than once (in general, this subgraph should be “maximal” in the sense that if we can add

Dataset	Classes	Features	Nodes	Edges
Cora	7	1,433	2,485	5,069
Planetoid.Cora [43]	7	1,433	2,708	10,556
Planetoid .CiteSeer [43]	6	3,703	3,327	9,104
CitationFull .CiteSeer [7]	6	602	4,230	10,674
Amazon.Photo [34]	8	745	7,650	0.2M
Amazon .Computers [34]	10	767	13,752	0.4M
ogbn-arxiv* [21]	40	128	169,343	1.1M
ogbn-products* [21]	47	100	2.4M	61.8M
ogbn- papers100M* [21]	172	128	111M	1.6B
friendster* [9]	100	128	65M	1.8B

Table 1: Statistics of the datasets after graph standardization. Datasets marked with * are used for distributed execution.

another vertex and it is still the case that no vertex is used more than once, we should do so). Then, we optimize this subgraph, using Algorithm 1. At this point, the tensor-relational decomposition of each tensor that has been optimized is treated as being fixed. This process is then repeated, and another subgraph of program G where no EinSum result is used more than once is chosen and optimized, until all decompositions have been chosen.

This will not arrive at a globally optimal solution, and it could be the case that some tensor-relational decompositions that are used twice are not chosen so as to be globally optimal. However, as the cost of this is only a sub-optimal, additional repartition, the expected cost of this non-optimality will hopefully be inconsequential.

Memory constraints. Often, in terms of performance, the optimal schema design is to use only massive tensors, and not to decompose the tensors at all. In practice, this will lead to memory problems. In our implementation, we address this by constraining the search to reflect memory limitations—maximum sizes for the total bytes required for the inputs and output of a kernel call, for example.

Non-Binary EinSum expressions. All of the algorithms and cost models described thus far in the paper can easily be extended to unary EinSum expressions, and we do not consider this issue further.

Higher-degree EinSum expressions are more challenging since the contraction path is critical for performance. While rare in ML computations, they do occur, particularly in other application domains. Our suggested tactic is to use optimization libraries like `opt_einsum` [12] to compute the contraction path and then break such expressions as a series of binary EinSum expressions. More sophisticated approaches are a topic for future work.

9 EXPERIMENTS

9.1 Experimental Overview

Experimentally, we wish to answer the question: *Can upper-case-lower-case Einstein notation generated by SPARSEEINSUM show better scalability over traditional tensor or purely relational approaches?*

Dataset and num machines	Tensor-based		SPARSEEINSUM			
	DGL	ALI	Opt	2nd	3rd	
ogbn-arxiv	1	4.3	12.9	8.3	13.4	21.8
	2	1.6	6.9	5.2	8.8	12.6
	4	0.9	3.4	2.0	4.1	6.3
	8	0.6	2.0	1.2	3.6	5.7
ogbn-products	1	20.3	96.7	31.4	35.7	37.2
	2	17.9	51.1	16.5	22.6	28.1
	4	9.6	31.7	8.4	11.8	15.9
ogbn-papers100M	1	OOM	OOM	289.5	306.1	351.9
	2	OOM	OOM	171.4	168.4	207.4
	4	92.5	OOM	82.1	93.5	95.1
	8	59.3	OOM	42.4	52.0	61.1
friendster	1	OOM	OOM	473.9	686.0	721.1
	2	OOM	OOM	324.0	413.1	556.1
	4	OOM	OOM	156.9	277.6	381.3
	8	103.1	OOM	94.5	186.4	206.1

Table 2: Running time comparison (in seconds) for one iteration of GCN training across different cluster sizes. DGL (PyTorch) and AliGraph are compared with the top three SPARSEEINSUM decompositions.

9.2 Evaluating SPARSEEINSUM Decompositions

Distributed graph convolution neural networks (GCNs) [26].

We compare SPARSEEINSUM-generated tensor-relational computations with DGL v2.4.0¹ [39] running on a PyTorch backend [30]—this is a fully tensor-based system—as well as AliGraph [51]. These systems are chosen as they represent the most well-known systems for graph neural network computation. The distributed engine used to run the SPARSEEINSUM-generated tensor-relational computation is PlinyCompute² [52], a distributed relational computation engine. We use TACO³ compiler-generated kernels [11, 27, 28] to efficiently support tensor operations of arbitrary shape within the tensor-relational computations generated by SPARSEEINSUM.

Computational problem. The core computation in a graph convolutional neural network is given as the following EinSum:

$$\begin{aligned} \mathbf{T}_{i,k}^0 &\leftarrow \sum_j \hat{\mathbf{D}}_{i,j} \times \hat{\mathbf{A}}_{j,k}; & \mathbf{T}_{i,l}^1 &\leftarrow \sum_k \mathbf{T}_{i,k}^0 \times \hat{\mathbf{D}}_{k,l} \\ \mathbf{T}_{i,m}^2 &\leftarrow \sum_l \mathbf{T}_{i,l}^1 \times \mathbf{H}_{l,m}^{(l)}; & \mathbf{T}_{i,n}^3 &\leftarrow \sum_m \mathbf{T}_{i,m}^2 \times \mathbf{W}_{m,n}^{(l)} \\ \mathbf{H}_{i,n}^{(l+1)} &\leftarrow \sum_0 \sigma(\mathbf{T}_{i,n}^3) \end{aligned}$$

Here, $\hat{\mathbf{A}}$ is the adjacency matrix with self-loops added, $\hat{\mathbf{D}}$ is the degree matrix, $\mathbf{H}^{(l)}$ is the node embedding matrix at layer l , and $\mathbf{W}^{(l)}$ is the weight matrix at layer l . This expression computes the node embeddings for the next layer by: (1) normalizing the adjacency matrix using the degree matrices ($\hat{\mathbf{D}}^{-\frac{1}{2}} \hat{\mathbf{A}} \hat{\mathbf{D}}^{-\frac{1}{2}}$), (2) aggregating neighbor features through this normalized adjacency matrix, and (3) applying a learned transformation through $\mathbf{W}^{(l)}$ followed by a non-linear activation function σ .

¹2.4.x branch until commit c6c874b

²<https://github.com/riceplinygroup/plinycompute>

³<https://github.com/tensor-compiler/taco>




Figure 1: Example of quantum circuits partitioned into multiple layers from 1 to 12.

Dataset	Qubits	Depth	Width	Gate Count
seca_n11	11	37	11	1275
multiplier_n13	13	7	13	496

Table 3: Dataset statistics for quantum benchmarks.

Circuit	Num Machines	SPARSEEINSUM		
		Opt	2nd Best	3rd Best
seca_n11	1	36.124	26.7	32.013
	2	21.214	27.512	27.752
	4	13.817	14.347	15.424
	8	10.094	12.531	12.309
multiplier_n13	1	102.632	105.316	116.229
	2	55.247	64.731	71.315
	4	36.651	42.325	47.426
	8	21.922	25.127	30.98

Table 4: Time (secs) vs. machines for quantum circuits.

We implement a two-layer GCN under transductive setting with the following hyperparameters: Adam optimizer with learning rate $\eta = 0.1$, dropout rate $\gamma = 0.5$, and hidden layer dimension $D = 256$.

All distributed experiments were conducted on servers equipped with Intel® Xeon® Silver 4214 CPU, 128GB DDR4 memory, and 1TB SSD storage. The nodes are interconnected via 10 Gbps Ethernet cable. Experiments were performed using clusters from 1 to 8 machines. Statistics describing the data sets are given in Table 1.

Results and discussion. For large graphs that cannot fit into the memory of a single machine, we use ParMETIS [1] to assist with partitioning when running on DGL and AliGraph. The execution times are summarized in Table 2. To evaluate how accurate the cost model used by the SPARSEEINSUM re-write algorithm is, we execute the best upper-case-lower-case decomposition (**Opt**), and the second (**2nd**) and third-best (**3rd**) as well.

For the smallest graph (ogbn-arxiv), DGL (running on PyTorch) is fastest. This is not surprising, as PlinyCompute is built for scalability, not low latency. For the next larger graph (ogbn-products) the two approaches are about the same. For the larger two graphs (1B+ edges), the SPARSEEINSUM-generated computation has obvious advantages. It is the one that can run in all cases—DGL and AliGraph are both plagued by out-of-memory problems, despite manual tuning. In the cases when DGL can run, it is slower than the SPARSEEINSUM-generated computation by 8.6% to nearly 40%.

Note that the tensor-relational implementation scales well across different number of machines. SPARSEEINSUM-on-PlinyCompute shows a speed-up of 5.3X, 6.8X, and 5.0X going from one to eight machines on ogbn-products, ogbn-papers100M, and friendster, respectively. DGL-on-Pytorch fails on the two larger graphs, but has a speedup of only 3.2X on ogbn-products.

Method	Datasets					
	Cora	P. Cora	P.Cite Seer	C.Cite Seer	A. Photo	A. Comp
Sparse Attention Computation						
Hyper	866	901	1497	40	2416	7801
PstGSQL	8017	7990	12156	619	18429	OOD
SQLite	22052	22084	TLE	34875	TLE	OOD
SEINSUM	4.5	4.7	8.0	7.7	192	265
Dense Attention Computation						
Hyper	2194	2185	3457	75	2486	8123
PstGSQL	8130	8186	12432	636	18464	OOD
SQLite	16424	15180	TLE	39470	TLE	TLE
SEINSUM	609	602	950	54	OOD	OOD

Table 5: Running time (in seconds) for attention computations. P stands for Planetoid, C: CitationFull, A: Amazon. OOD is Out of Disk, as the database wrote too much data.

Method	Datasets					
	Cora	P. Cora	P.Cite Seer	C.Cite Seer	A. Photo	A. Comp
Graph Conv. Neural Network						
Hyper	10.7	10.8	21.021	4.4	197.5	390.4
PstGSQL	16.0	16.0	30.9	7.7	436.1	830.3
SQLite	TLE	TLE	TLE	TLE	TLE	TLE
SEINSUM	7.3	8.4	22.9	3.3	19.3	37.6

Table 6: Graph convolutional neural network running times (in seconds). TLE indicates Time Limit Exceeded.

Single-machine convolution neural networks. Single-machine experiments focus on comparing SPARSEEINSUM-generated tensor-relational implementations with equivalent, purely-relational implementations. These are generated using the PortableSQL EinSum compiler [5], chosen as it is the standard method in the database literature for converting EinSum to SQL. SPARSEEINSUM-generated tensor-relational computations are run on top of PostgreSQL, again using TACO-generated kernels. The purely relational implementation is run on SQLite [3], PostgreSQL [15], and Hyper [24].

Results and Discussion. Running times are shown in Table 6. For the four smallest data sets, there is not an appreciable difference between options: the tensor-relational implementation is generally fastest, but it cuts at most 25% off the running time of “pure relational + Hyper”. For the larger two data sets, the tensor-relational implementation is ten times faster than “pure relational + Hyper”, and more than 40 times faster than “pure relational + Postgres”.

In Table 8, we show the schema generated by the SPARSEEINSUM compiler for two of the input graphs. Interestingly, the computations are both implemented quite relationally. However, as one might expect, the exception is that the features and embeddings (which are dense) are implemented as vectors on a per-vertex basis, and the linear transformations are implemented as dense matrices.

Distributed quantum circuit simulation [47]. We adopt the widely used QASMbench [29] suite as the quantum circuit simulation benchmark. Qiskit 1.2.4 [40] and cuda-quantum 24.03.0 [25] are used to help with quantum circuit generation and execution. We

UCLC EinSum	Planetoid.CiteSeer
$\mathbf{T}_{I,k}^0 \leftarrow \sum_M \mathbf{X}_{I,M} \times \mathbf{W}_{M,k}^Q$	$\mathbf{X}(\text{INT } i, \text{INT } m, \text{DOUBLE } \text{val})$ $\mathbf{W}^Q(\text{INT } m, \text{ TENSOR}[1024] \text{ val})$
$\mathbf{T}_{J,k}^1 \leftarrow \sum_N \mathbf{X}_{J,N} \times \mathbf{W}_{N,k}^K$	$\mathbf{X}(\text{INT } j, \text{INT } n, \text{DOUBLE } \text{val})$ $\mathbf{W}^K(\text{INT } n, \text{ TENSOR}[1024] \text{ val})$
$\mathbf{T}_{I,J,k}^2 \leftarrow \sum_0 \mathbf{T}_{I,k}^0 \times \mathbf{A}_{I,J}$	$\mathbf{T}^0(\text{INT } i, \text{ TENSOR}[1024] \text{ val})$ $\mathbf{A}(\text{INT } i, \text{INT } j, \text{DOUBLE } \text{val})$
$\mathbf{T}_{I,J}^3 \leftarrow \sum_k \mathbf{T}_{I,J,k}^2 \times \mathbf{T}_{J,k}^1$	$\mathbf{T}^2(\text{INT } i, \text{INT } j,$ $\text{ TENSOR}[1024] \text{ val})$ $\mathbf{T}^1(\text{INT } j, \text{ TENSOR}[1024] \text{ val})$
$\text{Attn}_{I,J} \leftarrow \sum_0 \frac{1}{\sqrt{d_k}} (\mathbf{T}_{I,J}^3)$	$\mathbf{T}^3(\text{INT } i, \text{INT } j, \text{DOUBLE } \text{val})$

Table 7: Upper-case-lower-case Einstein notation produced by the SPARSEEINSUM for Planetoid.CiteSeer.

select `seca_n11` and `multiplier_n13` as large scale benchmarks. Characteristics for those benchmarks are shown in Table 3.

Computational problem. The tensor contraction for quantum gate operations can be written as:

$$\psi'_{i_n} = \sum_{i_{n-1}} \mathbf{U}_{i_n, i_{n-1}}^{(n)} \sum_{i_{n-2}} \mathbf{U}_{i_{n-1}, i_{n-2}}^{(n-1)} \cdots \sum_{i_1} \mathbf{U}_{i_2, i_1}^{(2)} \sum_{i_0} \mathbf{U}_{i_1, i_0}^{(1)} \psi_{i_0} \quad (8)$$

where $\mathbf{U}^{(k)}$ represents the unitary gate operation matrix at layer k , ψ is the initial quantum state vector, and ψ' is the final quantum state vector after applying all gate operations.

Each quantum circuit is partitioned into multiple layers where multiple gates can be executed in parallel within a layer. For each layer, we generate one unitary matrix by tensoring all gates within the layer. The whole simulation process passes qubits vector states into a unitary matrix for each layer one by one as shown in Figure 1.

Results and Discussion. The results for running large-scale benchmarks on a distributed environment are shown in Table 4. We evaluate the best, second-best, and third-best decompositions produced, according to the cost model used by the SPARSEEINSUM compiler. In almost all cases, the cost model has ordered the various decompositions correctly.

Scaling efficiency is reasonable: eight machines is 3.6X faster than one on `seca_n11`, and 4.6X faster on `multiplier_n13`. Scaling efficiency is worse than in the case of the GCN, as data movement overhead is high and the computation cost is relatively small.

Single-machine attention computation [45, 48]. Sparse attention resembles the classic dense attention, but it allows an arbitrary graph to indicate which items should be attending to which items.

Computational problem. The core computation in a graph transformer is sparse attention computation:

$$\begin{aligned} \mathbf{T}_{i,k}^0 &\leftarrow \sum_M \mathbf{X}_{i,m} \times \mathbf{W}_{m,k}^Q; & \mathbf{T}_{j,k}^1 &\leftarrow \sum_n \mathbf{X}_{j,n} \times \mathbf{W}_{n,k}^K \\ \mathbf{T}_{i,j,k}^2 &\leftarrow \sum_0 \mathbf{T}_{i,k}^0 \times \mathbf{A}_{i,j}; & \mathbf{T}_{i,j}^3 &\leftarrow \sum_k \mathbf{T}_{i,j,k}^2 \times \mathbf{T}_{j,k}^1 \\ \text{Attn}_{i,j} &\leftarrow \sum_0 \frac{1}{\sqrt{d_k}} (\mathbf{T}_{i,j}^3) \end{aligned}$$

where \mathbf{X} stores node features, \mathbf{W}_Q and \mathbf{W}_K are learnable query and key projection matrices respectively, \mathbf{A} is the adjacency matrix

Upper-Case-Lower-Case EinSum	CitationFull.CiteSeer	Amazon.Computers
$\mathbf{T}_{I,K}^0 \leftarrow \sum_J \hat{\mathbf{D}}_{I,J} \times \hat{\mathbf{A}}_{J,K}$	$\hat{\mathbf{D}}(\text{INT } i, \text{INT } j, \text{DOUBLE val})$ $\hat{\mathbf{A}}(\text{INT } j, \text{INT } k, \text{DOUBLE val})$	$\hat{\mathbf{D}}(\text{INT } i, \text{INT } j, \text{DOUBLE val})$ $\hat{\mathbf{A}}(\text{INT } j, \text{INT } k, \text{DOUBLE val})$
$\mathbf{T}_{I,L}^1 \leftarrow \sum_K \mathbf{T}_{I,K}^0 \times \hat{\mathbf{D}}_{K,L}$	$\mathbf{T}^0(\text{INT } i, \text{INT } k, \text{DOUBLE val})$ $\hat{\mathbf{D}}(\text{INT } k, \text{INT } l, \text{DOUBLE val})$	$\mathbf{T}^0(\text{INT } i, \text{INT } k, \text{DOUBLE val})$ $\hat{\mathbf{D}}(\text{INT } k, \text{INT } l, \text{DOUBLE val})$
$\mathbf{T}_{I,m}^2 \leftarrow \sum_L \mathbf{T}_{I,L}^1 \times \mathbf{H}_{L,m}^{(l)}$	$\mathbf{T}^1(\text{INT } i, \text{INT } l, \text{DOUBLE val})$ $\mathbf{H}^{(l)}(\text{INT } l, \text{ TENSOR}[602] \text{ val})$	$\mathbf{T}^1(\text{INT } i, \text{INT } l, \text{DOUBLE val})$ $\mathbf{H}^{(l)}(\text{INT } l, \text{ TENSOR}[767] \text{ val})$
$\mathbf{T}_{I,n}^3 \leftarrow \sum_m \mathbf{T}_{I,m}^2 \times \mathbf{W}_{m,n}^{(l)}$	$\mathbf{T}^2(\text{INT } i, \text{ TENSOR}[602] \text{ val})$ $\mathbf{W}^{(l)}(\text{ TENSOR}[602][256] \text{ val})$	$\mathbf{T}^2(\text{INT } i, \text{ TENSOR}[767] \text{ val})$ $\mathbf{W}^{(l)}(\text{ TENSOR}[767][256] \text{ val})$
$\mathbf{H}_{I,n}^{(l+1)} \leftarrow \sum_0 \sigma(\mathbf{T}_{I,n}^3)$	$\mathbf{T}^3(\text{INT } i, \text{ TENSOR}[256] \text{ val})$	$\mathbf{T}^3(\text{INT } i, \text{ TENSOR}[256] \text{ val})$

Table 8: Upper-case-lower-case Einstein notation produced by the SPARSEEINSUM compiler for the GCN computation, along with the corresponding relational schemas, for two of the test datasets.

Method	Datasets					
	Cora	P. Cora	P.Cite Seer	C.Cite Seer	A. Photo	A. Comp
Sparse Attention Computation						
SEINSUM	4.5	4.7	8.0	7.7	192	265
Greedy	21	23	52	17	769	884
Dense Attention Computation						
SEINSUM	609	602	950	54	OOD	OOD
Greedy	2198	2124	4107	315	OOD	OOD

Table 9: Comparing pure greedy runtime with the full dynamic programming solution for the attention computation.

indicating the graph structure, d_k is the dimension of the key vectors, and $\text{Attn}_{i,j}$ gives the attention scores before normalization. Following prior work [45, 48], we evaluate both dense attention (where each node attends to all other nodes) and sparse attention (where nodes only attend to their local neighborhood in the graph).

Results and Discussion. Running time results are given in Table 5. The surprising finding here is that *even though relational computations are inherently sparse*, the pure relational implementations are not much faster when using sparse attention, compared to dense attention. “Pure relational + Hyper” best utilizes the sparsity, and is up to 2X as fast in the sparse case compared to the dense case. In comparison, SPARSEEINSUM can be more than 100X as fast when computing sparse attention rather than dense attention, and is 30X to 100X as fast as “Pure relational + Hyper” in the sparse case.

Table 7 shows the generated schema for one of the data sets. The node embeddings are stored as vectors, but the weight matrices are stored into vectors, due to memory limitations. The input data itself is sparse and has been decomposed into purely relational tuples.

9.3 Ablations and Other Experiments

Effectiveness of the dynamic programming. It is reasonable to ask: Is it necessary to run an expensive and complicated dynamic programming search during the SPARSEEINSUM decomposition? We evaluate a simpler, greedy search: along a path through the graph of EinSum expressions, greedily pick the decomposition that minimizes the cost; do not remember decompositions with different output tensor formats that are suboptimal. We compare greedy to SPARSEEINSUM in Table 9 for attention on one machine.

Component	Datasets					
	Cora	P. Cora	P.Cite Seer	C.Cite Seer	A. Photo	A. Comp
Sparse Attention Computation						
Exec.	4.5	4.7	8.0	7.7	192	265
Schema	253	96	428	86	344	239
SQL	46	11	32	2.0	5.4	3.4
Kernels	0.04	0.04	0.04	0.08	0.04	0.04
Dense Attention Computation						
Exec	609	602	950	54	OOD	OOD
Schema	146	127	533	221	862	1690
SQL	51	24	58	44	130	401
Kernels	0.06	0.06	0.06	0.06	0.06	0.06

Table 10: Running time for the various system components, attention computation, in seconds. “Exec.” is the actual execution time, “Schema” is the time to pre-process the data and generate the schema, “SQL” is the time to generate the SQL including the data loading code, and “Kernels” is the time to generate the kernels using TACO.

None (as-is)	Perturbation					
	$C_{\text{xfer}} \times \frac{1}{10}$	$C_{\text{xfer}} \times 10$	$C_{\text{kernel}} \times \frac{1}{10}$	$C_{\text{kernel}} \times 10$	$C_{\text{fixed}} \times \frac{1}{10}$	$C_{\text{fixed}} \times 10$
192	130	253	294	1027	1027	305

Table 11: Effect of perturbing the various cost parameters (C_{xfer} , C_{kernel} , C_{fixed}) on the running time for the sparse attention computation on the Amazon.Photo data set.

Robustness to cost model errors. Is an accurate cost model necessary? We perturb the cost estimates by multiplying each cost estimate by a Gamma(α, θ) random variate, where $\alpha \times \theta = 1$ so that on expectation, the cost value is unchanged, but errors are possible. As we decrease α and increase θ the variance of the estimate (and hence the potential for error) increases. Results on the single-machine graph neural network computation are given in Table 12. Running time is averaged over ten different perturbed runs. Here we see that with $\alpha = \theta = 1$, there is no difference in the running time—as this is a significant level of noise, it suggests some robustness. The higher perturbation levels lead to slower times. Interestingly, $\alpha = \frac{1}{3}, \theta = 3$ is worse than $\alpha = \frac{1}{10}, \theta = 10$.

Robustness to cost parameters. We hold two of the three cost parameters (C_{xfer} , C_{kernel} , C_{fixed}) constant, and perturb the third,

Error	Datasets					
	Cora	P. Cora	P.Cite Seer	C.Cite Seer	A. Photo	A. Comp
	Graph Conv. Neural Network					
None	7.3	8.4	22.9	3.3	19.3	37.6
$\Gamma(1, 1)$	7.2	8.3	21.5	3.4	19.5	37.1
$\Gamma(\frac{1}{3}, 3)$	16.1	16.0	31.9	8.0	435.9	839.9
$\Gamma(\frac{1}{10}, 10)$	12.0	11.2	25.1	7.4	68.0	47.4

Table 12: Graph neural network running times (in seconds) with erroneous/noisy cardinality estimates. Here, $\Gamma(\alpha, \theta)$ means each estimate is multiplied by a random number generated using a Gamma(α, θ) distribution.

either cutting it by a factor of ten, or multiplying it by a factor of ten. Single-machine results are given in Table 12. A poor choice in cost parameter can be quite damaging. Interestingly, cutting C_{xferr} by a factor of ten actually reduced the running time of the resulting decomposition, suggesting that the C_{xferr} we chose was not optimal.

Running time breakdown. Table 10 compares the execution time for the actual attention computation with the time taken by the decomposition itself. In addition to the execution time, three other times are given: (1) Schema: time to process the input data, compute the necessary statistics, and run the dynamic programming. (2) SQL: time to generate the SQL, which is dominated by the time necessary to generate the INSERT INTO statements to load all of the data into PostgreSQL. (3) Kernels: time to generate the TACO kernels. We see that these times can be quite significant, though our code is written in Python, and for the single-machine attention experiments, all data are stored in text files. Further, to ensure portability across engines all data loading is handled using SQL, and is verbose. During deployment or training the attention computation would be run hundreds of times or more, so time would be amortized across runs.

10 RELATED WORK

EinSum was proposed by Einstein [16], and is now supported by machine learning/AI frameworks such as NumPy [20], JAX [18], PyTorch [30] and Dask [13]. However, upper-case-lower-case EinSum, where EinSum is modified to facilitate sharding specification, is unique. A general problem in this space is how to specify sharding—see <https://jax-ml.github.io/scaling-book/sharding/> for an example of such an effort—but upper-case-lower-case EinSum is unique in its ability to combine a specification of the computation with the way in which the computation is decomposed.

We propose a dynamic programming algorithm for the problem of re-writing a DAG of EinSum expressions (where paths between any connected vertices are unique) to a DAG of upper-case-lower-case EinSum. Dynamic programming algorithms of this variety have been around for a long time; our algorithm is strongly reminiscent of Felsenstein’s algorithm [17] from computational biology, and other efforts in sharding of compute graphs for machine learning have used dynamic programming [8, 49]. Likewise there exist cost models for sparse tensor computations [2, 35].

There are a number of kernel compilers for tensor computations, such as TVM [10], Triton [38], and TACO [11, 27, 28]. Both TACO and TVM (via SparseTIR [44]) provide explicit support for sparsity. Of the various tensor compilers with support for sparsity, Galley

[14] is closest to our own proposal, using many ideas from relational system design (such as applying relational methods for sparsity estimation; see Section 7 of this paper). This existing work targets the generation of CPU or GPU kernels or programs, not scalable relational computations that are meant to run on top of a relational system. The core idea in the SPARSEEINSUM re-write approach is to determine which parts of the computation should be run relationally and which parts should utilize high-performance CPU or GPU kernels generated by these other systems. In that sense, SPARSEEINSUM sits *on top* of these tools and is a consumer of the kernels that they produce. In fact, our SPARSEEINSUM prototype uses TACO to generate dense kernels.

There has been other recent work that explores the close relationship between tensors and relations in order to optimize tensor computations. Schleich et al. propose STOREL [31], which accepts a user-supplied storage format for each input tensor, and then optimizes the tensor computation. In contrast, SPARSEEINSUM accepts a declaratively-specified computation and then optimizes the representation format, choosing how to decompose the tensors into relations so as to optimize performance. The SPARSEEINSUM approach relies on the underlying relational system to optimize the computation.

In the Duck’s Brain [32], Schüle et al. are the most recent to argue for the use of relational systems augmented with tensors for modern ML computations—using precisely the sort of tensor-relational representations described in this paper. The Duck’s Brain is an example of a *relational platform* on which the output of SparseEinSum-generated tensor-relational computation could be run. The idea of combined tensor and relational processing has been explored in other works [22, 23, 37, 46], including SystemML [19] and SystemDS [6], which extended SQL-style operations with matrix primitives. However, these prior works do not consider how to automatically generate optimized tensor-relational representations. Like in the case of the Duck’s Brain, SPARSEEINSUM could be used to generate the decomposition used by these systems.

Another approach explored in the literature is the use of pure relational implementations of ML or AI computations (see, for example Pytond [33] or Portable SQL [5] for two recent works). Another recent work proposes an algorithm to switch from dense to sparse dynamically during runtime [36].

11 CONCLUSIONS

We introduced an algorithm that decomposes tensor computation into equivalent tensor-relational computation, as well as SPARSEEINSUM, a system that takes EinSum computation as input, and compiles computation into upper-case-lower-case Einstein notation. This computation can be implemented on virtually any database system with multidimensional array support. The compilation process incorporates sparsity estimation techniques and automatic cost-based schema optimization, enabling generated tensor-relational computations to achieve significant performance improvements.

ACKNOWLEDGMENTS

This research was supported by NSF grants 1918651, 2008240, 2131294, and 2212557, NIH CTSA #UM1TR004906, and US DOT Tier-1 UTC CYBER-CARE grant #69A3552348332.

REFERENCES

- [1] [n.d.]. ParMETIS. <https://github.com/KarypisLab/ParMETIS>.
- [2] Willow Ahrens, Fredrik Kjolstad, and Saman Amarasinghe. 2022. Autoscheduling for sparse tensor algebra with an asymptotic cost model. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. 269–285.
- [3] Grant Allen, Mike Owens, Grant Allen, and Mike Owens. 2010. Introducing SQLite. *The Definitive Guide to SQLite* (2010), 1–16.
- [4] Alan H Barr. 1991. The Einstein summation notation. *An Introduction to Physically Based Modeling (Course Notes 19)*, pages E 1 (1991), 57.
- [5] Mark Blacher, Julien Klaus, Christoph Staudt, Sören Laue, Viktor Leis, and Joachim Giesen. 2023. Efficient and Portable Einstein Summation in SQL. *Proc. ACM Manag. Data* 1, 2, Article 121 (jun 2023), 19 pages. <https://doi.org/10.1145/3589266>
- [6] Matthias Boehm, Iulian Antonov, Sebastian Baunsgaard, Mark Dokter, Robert Ginthör, Kevin Innerebner, Florijan Klezin, Stefanie Lindstaedt, Arnab Phani, Benjamin Rath, et al. 2019. SystemDS: A declarative machine learning system for the end-to-end data science lifecycle. *arXiv preprint arXiv:1909.02976* (2019).
- [7] Aleksandar Bojchevski and Stephan Günnemann. 2017. Deep gaussian embedding of graphs: Unsupervised inductive learning via ranking. *arXiv preprint arXiv:1707.03815* (2017).
- [8] Daniel Bourgeois, Zhimin Ding, Dimitrije Jankov, Jiehui Li, Mahmoud Sleem, Yuxin Tang, Jiawen Yao, Xinyu Yao, and Chris Jermaine. 2024. EinDecomp: Decomposition of Declaratively-Specified Machine Learning and Numerical Computations for Parallel Execution. *arXiv preprint arXiv:2410.02682* (2024).
- [9] Danah Michele Boyd. 2004. Friendster and publicly articulated social networking. In *CHI'04 extended abstracts on Human factors in computing systems*. 1279–1282.
- [10] Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Yan, Haichen Shen, Meghan Cowan, Leyuan Wang, Yuwei Hu, Luis Ceze, et al. 2018. {TVM}: An automated {End-to-End} optimizing compiler for deep learning. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*. 578–594.
- [11] Stephen Chou, Fredrik Kjolstad, and Saman Amarasinghe. 2018. Format abstraction for sparse tensor algebra compilers. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 123 (Oct. 2018), 30 pages. <https://doi.org/10.1145/3276493>
- [12] G Daniel, Johnnie Gray, et al. 2018. Opt_einsum—a python package for optimizing contraction order for einsum-like expressions. *Journal of Open Source Software* 3, 26 (2018), 753.
- [13] Jesse Daniel. 2019. *Data science with Python and Dask*. Simon and Schuster.
- [14] Kyle Deeds, Willow Ahrens, Magdalena Balazinska, and Dan Suciu. 2025. Galley: Modern Query Optimization for Sparse Tensor Programs. *Proceedings of the ACM on Management of Data* 3, 3 (2025), 1–24.
- [15] Korry Douglas and Susan Douglas. 2003. *PostgreSQL: a comprehensive guide to building, programming, and administering PostgreSQL databases*. SAMS publishing.
- [16] Albert Einstein, Leopold Infeld, and Banesh Hoffmann. 1938. The gravitational equations and the problem of motion. *Annals of mathematics* (1938), 65–100.
- [17] Joseph Felsenstein. 1981. Evolutionary trees from DNA sequences: a maximum likelihood approach. *Journal of molecular evolution* 17, 6 (1981), 368–376.
- [18] Roy Frostig, Matthew James Johnson, and Chris Leary. 2018. Compiling machine learning programs via high-level tracing. *Systems for Machine Learning* 4, 9 (2018).
- [19] Amol Ghoting, Rajasekar Krishnamurthy, Edwin Pednault, Berthold Reinwald, Vikas Sindhwani, Shirish Tatikonda, Yuanyuan Tian, and Shivakumar Vaithyanathan. 2011. SystemML: Declarative machine learning on MapReduce. In *2011 IEEE 27th International conference on data engineering*. IEEE, 231–242.
- [20] Charles R Harris, K Jarrod Millman, Stéfán J Van Der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. 2020. Array programming with NumPy. *Nature* 585, 7825 (2020), 357–362.
- [21] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. 2020. Open graph benchmark: Datasets for machine learning on graphs. *Advances in neural information processing systems* 33 (2020), 22118–22133.
- [22] Dimitrije Jankov, Shangyu Luo, Binhang Yuan, Zhuhua Cai, Jia Zou, Chris Jermaine, and Zekai J Gao. 2019. Declarative recursive computation on an rdbms, or, why you should use a database for distributed machine learning. *arXiv preprint arXiv:1904.11121* (2019).
- [23] Dimitrije Jankov, Binhang Yuan, Shangyu Luo, and Chris Jermaine. 2021. Distributed numerical and machine learning computations via two-phase execution of aggregated join trees. *Proceedings of the VLDB Endowment* 14, 7 (2021).
- [24] Alfons Kemper and Thomas Neumann. 2011. HyPer: A hybrid OLTP&OLAP main memory database system based on virtual memory snapshots. In *2011 IEEE 27th International Conference on Data Engineering*. IEEE, 195–206.
- [25] Jin-Sung Kim, Alex McCaskey, Bettina Heim, Manish Modani, Sam Stanwyck, and Timothy Costa. 2023. Cuda quantum: The platform for integrated quantum-classical computing. In *2023 60th ACM/IEEE Design Automation Conference (DAC)*. IEEE, 1–4.
- [26] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [27] Fredrik Kjolstad, Willow Ahrens, Shoaib Kamil, and Saman Amarasinghe. 2019. Tensor algebra compilation with workspaces. In *2019 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*. IEEE, 180–192.
- [28] Fredrik Kjolstad, Shoaib Kamil, Stephen Chou, David Lugato, and Saman Amarasinghe. 2017. The tensor algebra compiler. *Proc. ACM Program. Lang.* 1, OOPSLA, Article 77 (Oct. 2017), 29 pages. <https://doi.org/10.1145/3133901>
- [29] Ang Li, Samuel Stein, Sriram Krishnamoorthy, and James Ang. 2023. Qasmbench: A low-level quantum benchmark suite for nisq evaluation and simulation. *ACM Transactions on Quantum Computing* 4, 2 (2023), 1–26.
- [30] A Paszke. 2019. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703* (2019).
- [31] Maximilian Schleich, Amir Shaikhha, and Dan Suciu. 2023. Optimizing tensor programs on flexible storage. *Proceedings of the ACM on Management of Data* 1, 1 (2023), 1–27.
- [32] Maximilian Schüle, Thomas Neumann, and Alfons Kemper. 2024. The Duck’s Brain: Training and Inference of Neural Networks within Database Engines. *Datenbank-Spektrum* 24, 3 (2024), 209–221.
- [33] Hesam Shahrokhi, Amirali Kaboli, Mahdi Ghorbani, and Amir Shaikhha. 2024. Pytond: Efficient python data science on the shoulders of databases. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. IEEE, 423–435.
- [34] Oleksandr Shchur, Maximilian Mumme, Aleksandar Bojchevski, and Stephan Günnemann. 2018. Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868* (2018).
- [35] Johanna Sommer, Matthias Boehm, Alexandre V Evfimievski, Berthold Reinwald, and Peter J Haas. 2019. Mnc: Structure-exploiting sparsity estimation for matrix expressions. In *Proceedings of the 2019 International Conference on Management of Data*. 1607–1623.
- [36] Christoph Staudt, Mark Blacher, Tim Hoffmann, Kaspar Kasche, Olaf Beyersdorff, and Joachim Giesen. 2025. Exploiting Dynamic Sparsity in Einsum. In *The Thirtieth Annual Conference on Neural Information Processing Systems*.
- [37] Yuxin Tang, Zhimin Ding, Dimitrije Jankov, Binhang Yuan, Daniel Bourgeois, and Chris Jermaine. 2023. Auto-differentiation of relational computations for very large scale machine learning. In *International Conference on Machine Learning*. PMLR, 33581–33598.
- [38] Philippe Tillet, Hsiang-Tsung Kung, and David Cox. 2019. Triton: an intermediate language and compiler for tiled neural network computations. In *Proceedings of the 3rd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages*. 10–19.
- [39] Minjie Wang, Da Zheng, Zihao Ye, Quan Gan, Mufei Li, Xiang Song, Jinjing Zhou, Chao Ma, Lingfan Yu, Yu Gai, Tianjun Xiao, Tong He, George Karypis, Jinyang Li, and Zheng Zhang. 2020. Deep Graph Library: A Graph-Centric, Highly-Performant Package for Graph Neural Networks. *arXiv:1909.01315 [cs.LG]* <https://arxiv.org/abs/1909.01315>
- [40] Robert Wille, Rod Van Meter, and Yehuda Naveh. 2019. IBM’s Qiskit tool chain: Working with and developing for real quantum computers. In *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 1234–1240.
- [41] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S Yu. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* 32, 1 (2020), 4–24.
- [42] Lizhi Xiang, Omid Asudeh, Gerald Sabin, Aravind Sukumaran-Rajam, and P Sadayappan. 2025. cuTeSpMM: Accelerating Sparse-Dense Matrix Multiplication using GPU Tensor Cores. *arXiv preprint arXiv:2504.06443* (2025).
- [43] Zhilin Yang, William Cohen, and Ruslan Salakhudinov. 2016. Revisiting semi-supervised learning with graph embeddings. In *International conference on machine learning*. PMLR, 40–48.
- [44] Zihao Ye, Ruihang Lai, Junru Shao, Tianqi Chen, and Luis Ceze. 2023. Sparsetir: Composable abstractions for sparse compilation in deep learning. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3*. 660–678.
- [45] Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. 2021. Do transformers really perform badly for graph representation? *Advances in neural information processing systems* 34 (2021), 28877–28888.
- [46] Binhang Yuan, Dimitrije Jankov, Jia Zou, Yuxin Tang, Daniel Bourgeois, and Chris Jermaine. 2021. Tensor relational algebra for distributed machine learning system design. *Proc. VLDB Endow.* 14, 8 (April 2021), 1338–1350. <https://doi.org/10.14778/3457390.3457399>
- [47] Chen Zhang, Zeyu Song, Haojie Wang, Kaiyuan Rong, and Jidong Zhai. 2021. HyQuas: hybrid partitioner based quantum circuit simulation system on GPU. In *Proceedings of the 35th ACM International Conference on Supercomputing*. 443–454.
- [48] Meng Zhang, Jie Sun, Qinghao Hu, Peng Sun, Zeke Wang, Yonggang Wen, and Tianwei Zhang. 2024. TorchGT: A Holistic System for Large-Scale Graph Transformer Training. In *SC24: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 1–17.

- [49] Lianmin Zheng, Zhuohan Li, Hao Zhang, Yonghao Zhuang, Zhifeng Chen, Yanping Huang, Yida Wang, Yuanzhong Xu, Danyang Zhuo, Eric P Xing, et al. 2022. Alpa: Automating inter-and {Intra-Operator} parallelism for distributed deep learning. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*. 559–578.
- [50] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2020. Graph neural networks: A review of methods and applications. *AI open* 1 (2020), 57–81.
- [51] Rong Zhu, Kun Zhao, Hongxia Yang, Wei Lin, Chang Zhou, Baole Ai, Yong Li, and Jingren Zhou. 2019. AliGraph: A Comprehensive Graph Neural Network Platform. arXiv:1902.08730 [cs.DC] <https://arxiv.org/abs/1902.08730>
- [52] Jia Zou, R Matthew Barnett, Tania Lorido-Botran, Shangyu Luo, Carlos Monroy, Sourav Sikdar, Kia Teymourian, Binhang Yuan, and Chris Jermaine. 2018. PlinyCompute: A platform for high-performance, distributed, data-intensive tool development. In *Proceedings of the 2018 International Conference on Management of Data*. 1189–1204.