

# Exploiting Dynamic Sparsity in Einsum

November 7, 2025

NeurIPS 2025

Christoph Staudt, Mark Blacher, Tim Hoffmann,  
Kaspar Kasche, Olaf Beyersdorff, Joachim Giesen

Faculty of Mathematics and Computer Science



# Overview

Einsum and Tensor Networks

Exponential Separation

Experiments

# Einsum and Tensor Networks

# Einsum and Tensor Networks

## Einsum basics

Operation	Einsum expression
inner vector product	$\text{einsum}(i, i \rightarrow; u, v)$
elementwise vector product	$\text{einsum}(i, i \rightarrow i; u, v)$
matrix-matrix-vector product	$\text{einsum}(ij, jk, k \rightarrow i; A, B, v)$


In practice large expressions are executed in pairwise steps:

$$\text{einsum}(ij, j \rightarrow i; A, \text{einsum}(jk, k \rightarrow j; B, v))$$

# Einsum and Tensor Networks

## Tensor hypernetworks as graphical representation

`einsum` ( $cb, ba, cf, be, ad, fe, ed, fi, eh, dg, ih, hg \rightarrow; A, B, C, D, E, F, G, H, I, J, K, L$ )




# Exponential Separation

# Exponential Separation

## Grid instances



- Each tensor (black node) is a diagonal matrix
- We prove: For such a square grid with side length  $n$ , dense tensors require at least  $2^n$  flops, while sparse tensors only need  $O(n^2)$  flops.



# Exponential Separation

## Grid instances

Synthetic einsum expressions based on the grid formulas used in the proof.






# Experiments

# Experiments

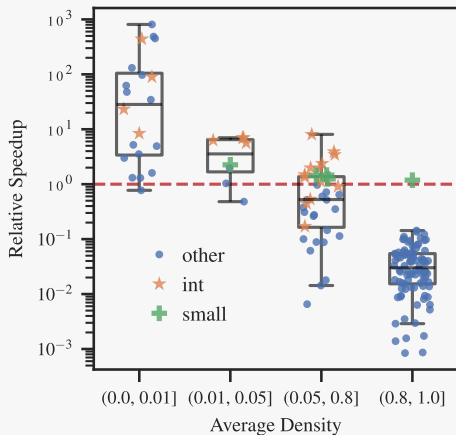
## Dynamic sparsity in the einsum benchmark

Real world instances from the einsum benchmark: <https://benchmark.einsum.org/>.



# Experiments

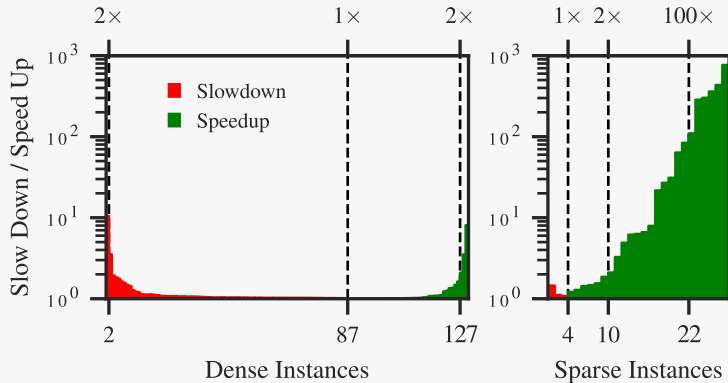
## Sparse Speedup



# Hybrid Algorithm

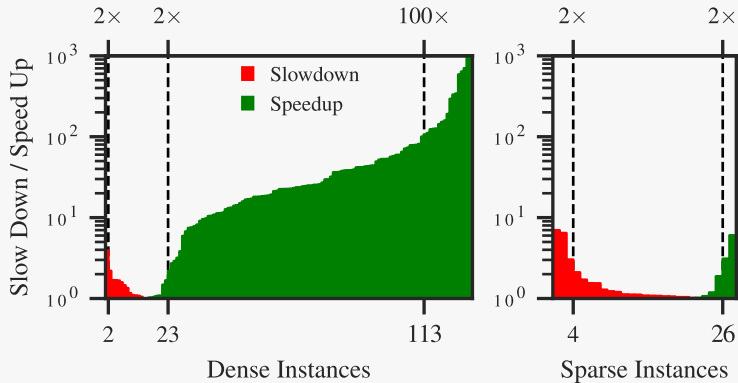
# Hybrid Algorithm

## Hybrid algorithm speedup



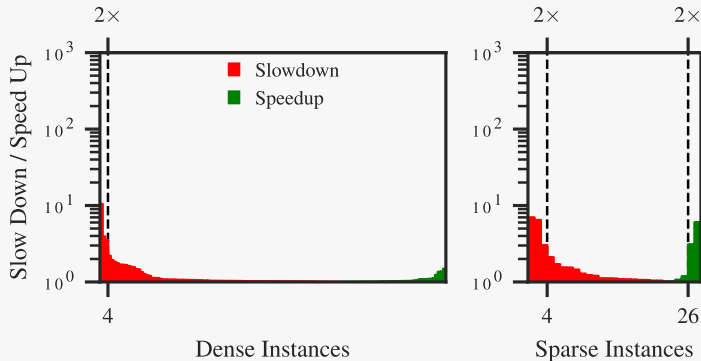
# Hybrid Algorithm

## Hybrid algorithm speedup



# Hybrid Algorithm

## Hybrid algorithm speedup



# Conclusions

1. Large einsum expressions exhibit dynamic sparsity in practice.
2. Significant speed ups are possible by exploiting dynamic sparsity.