

einsum optimizer

is all you need


Einsum

`einsum('ij,jk->ik', f1, f2)`

$$\sum_j f1(i, j) f2(j, k) \rightarrow g(i, k)$$


$f1, f2$ ← Factors

i, j, k ← Indices



Beyond original einstein summation

`einsum('ij,jk,kl->il', f1, f2, f3)`



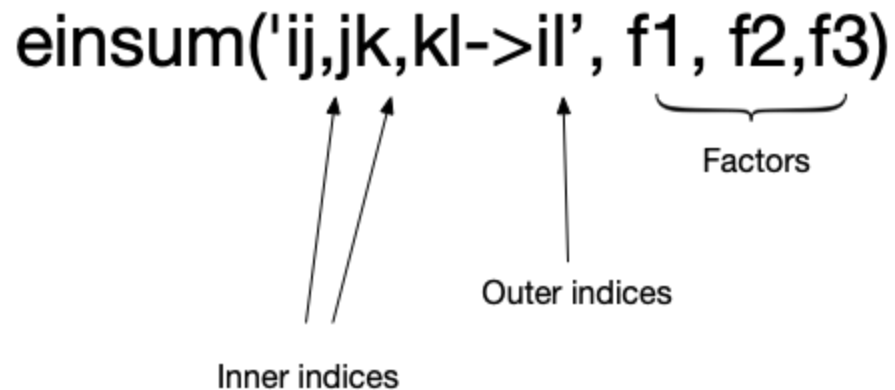
- outer indices can be empty

$$\sum_i f(i)$$

`einsum('i->', f)`

Beyond original einstein summation

`einsum('ij,jk,kl->il', f1, f2, f3)`



The diagram shows the `einsum` function call `einsum('ij,jk,kl->il', f1, f2, f3)`. Below the string, two parallel arrows point from the text 'Inner indices' to the 'ij' and 'jk' parts of the string. A single arrow points from the text 'Outer indices' to the 'il' part. A curly brace under the 'f1, f2, f3' arguments is labeled 'Factors'.


- inner indices can be empty

`einsum('ij->ij', f)`

Beyond original einstein summation

`einsum('ij->ji', f)`

`einsum('ij,jk,kl->il', f1, f2,f3)`



- order matters

`einsum('ij->ji', f)`

Beyond original einstein summation

`einsum('ij->ji', f)`

`einsum('ij,jk,kl->il', f1, f2, f3)`

Inner indices

Outer indices

Factors

- outer indices can repeat on left side: `einsum('ii->i', f)`

Matmul as einsum

$$\begin{pmatrix} a & b & c & d \\ e & f & g & h \\ i & j & k & l \end{pmatrix} \begin{pmatrix} 1 \\ 2 \\ 3 \\ 4 \end{pmatrix}$$

`einsum('ij,j->i', A, x)`

Convolution as einsum

a	b
---	---

1	2	3	4
---	---	---	---

$$\begin{pmatrix} a & b & 0 & 0 \\ 0 & a & b & 0 \\ 0 & 0 & a & b \end{pmatrix} \begin{pmatrix} 1 \\ 2 \\ 3 \\ 4 \end{pmatrix}$$


better memory

$$\begin{pmatrix} 1 & 2 \\ 2 & 3 \\ 3 & 4 \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix}$$

better time

- need a linearly indexed view
- ie, "fold" or "im2col" operation

Factored convolutions




$$\begin{pmatrix} 1 & 2 \\ 2 & 3 \\ 3 & 4 \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} \begin{pmatrix} c \\ d \end{pmatrix}$$

Factored convolutions

EINCONV: EXPLORING UNEXPLORED TENSOR NETWORK DECOMPOSITIONS FOR CONVOLUTIONAL NEURAL NETWORKS

- factored convolution
 - depthwise convolution
 - spatially separable convolution
 - bottleneck layer
 - 1x1 convolution
- 492 3D convolutions

Linear Neural Network





$$y' = x'ABC$$

```
einsum('i,ij,jk,kl->l', x, A, B, C)
```

=sum over weighted walks

Derivative



JacA

JacB

JacC

`einsum('ij,jk,kl->i', jacA, jacB, jacC)`

Hessian

$$\sum_{a,b,c,d} (f1(a,b)f2(b,c)f3(c,d))'$$

$$\sum_{a,b,c,d} f1'(a,b)f2(b,c)f3(c,d)$$

$$\sum_{a,b,c,d} f1(a,b)f2'(b,c)f3(c,d)$$




$$\sum_{a,b,c,d} f1(a,b)f2(b,c)f3'(c,d)$$

$h1(a,b,i)=f1'(a,b)$ if $i==1$ else $f1(a,b)$

$$\sum_{a,b,c,d,i} h1(a,b,i)h2(b,c,i)h3(c,d,i)$$

Efficient computation: factoring

$$\sum_{a,b,c,d} f1(a,b)f2(b,c)f3(c,d) = \sum_a \sum_b f(a,b) \left(\sum_c f(b,c) \left(\sum_d f(c,d) \right) \right)$$





Efficient computation: factoring

$$\sum_{ijkl} f1(i)f2(j)f3(k)f4(l) = \sum_i f1(i) \sum_j f2(j) \sum_k f3(k) \sum_l f4(l)$$

<https://colab.research.google.com/drive/1ItfFMp6WGdZLSrFtl-ppnVcD2ahLxIHg#scrollTo=ICV7-1SrEDA1>

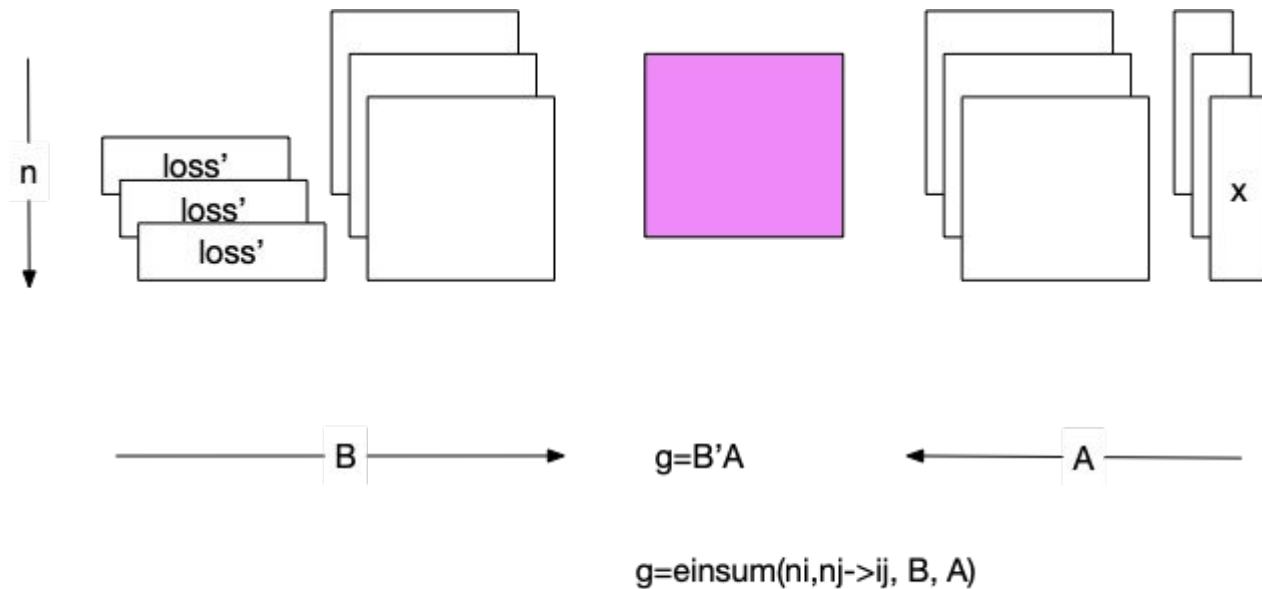
More gradients




$$g = b'a$$

$$g = \text{einsum}(i, j \rightarrow ij, b, a)$$

More gradients



Per example gradients



$$\text{B} \longrightarrow g = B \otimes^{KH} A \longleftarrow \text{A}$$

$$g = \text{einsum}(ni, nj \rightarrow nij, B, A)$$

Gradient norms squared

per-example gradient norms squared

```
g=einsum(i,i,j,j->, b, b, a, a)
```

```
g=einsum(ni,ni,nj,nj->n, B, B, A, A)
```

stick into einsum optimizer => discover the trick from Goodfellow **“Efficient Per-Example Gradient Computations”**

```
(B*B).sum(dim=1) * (A*A).sum(dim=1)
```


Gradient norms squared: conv layers?

- hessian trace
- gradient norms
- Hessian-vector products
- batch of per-example Hessian traces

https://colab.research.google.com/drive/16nKr_LmiiH8pgGkF1gNNahK83WVCpqk4#scrollTo=E8yzltiWaUS

Graphical view


$\text{einsum}(\underbrace{abcd}_{f1}, \underbrace{def-}_{f2} \rightarrow)$



factor vertices

needs hyper-edges

Graphical view




terms are factor nodes
indices are vertex nodes

$$\text{einsum}(\underbrace{abcd}_{f_1}, \underbrace{def}_{f_2} - >)$$

factor graph

Graphical view



vertices connected if they share a factor

index vertices

$$\text{einsum}(\underbrace{abcd}_{f_1}, \underbrace{def}_{f_2} \rightarrow)$$