Google Research


# Automatic differentiation 

## Mathieu Blondel

February 12, 2021

## Gradient-based learning

- Gradient-based training algorithms are the workhorse of modern machine learning.
- Deriving gradients by hand is tedious and error prone.
- This becomes quickly infeasible for complex models.
- Changes to the model require rederiving the gradient.

■ Deep learning = GPU + data + autodiff

## Automatic differentiation

- Evaluates the derivatives of a function at a given point.
- Not the same as numerical differentiation.
- Not the same as symbolic differentiation, which returns a "human-readable" expression.

■ In a neural network context, reverse autodiff is often known as backpropagation.

## Automatic differentiation

$\square$ A program is defined as the composition of primitive operations that we know how to derive.

- The user can focus on the forward computation / model.

```
import jax.numpy as jnp
from jax import grad, jit
def predict(params, inputs):
        for W, b in params:
            outputs = jnp.dot(inputs, W) + b
            inputs = jnp.tanh(outputs)
        return outputs
def loss_fun(params, inputs, targets):
        preds = predict(params, inputs)
        return jnp.sum((preds - targets)**2)
grad_fun = jit(grad(loss_fun))
```


## Automatic differentiation

■ Modern frameworks support higher-order derivatives

```
def tanh(x):
    y = jnp.exp(-2.0 * x)
    return (1.0 - y) / (1.0 + y)
```

$\mathrm{fp}=\operatorname{grad}(\tanh )$
$\mathrm{fpp}=\operatorname{grad}(\operatorname{grad}(\tanh ))$

## Outline

## 1 Numerical differentiation

## 2 Chain compositions

3 Computational graphs

4 Implementation

## 5 Advanced topics

6 Conclusion

## Derivatives

$\square$ Definition of derivative of $g: \mathbb{R} \rightarrow \mathbb{R}$

$$
g^{\prime}(a)=\frac{\partial g(a)}{\partial a}=\lim _{h \rightarrow 0} \frac{g(a+h)-g(a)}{h}
$$

- $g^{\prime}(a)$ is called Lagrange notation.
$\square \frac{\partial g(a)}{\partial a}$ is called Leibniz notation.
■ Interpretations: instantaneous rate of change of $g$, slope of the tangent of $g$ at $a$.


## Gradient

$\square$ The gradient of $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is

$$
\nabla f(\mathbf{x})=\left[\begin{array}{c}
\frac{\partial f}{\partial x_{1}}(\mathbf{x}) \\
\vdots \\
\frac{\partial f}{\partial x_{n}}(\mathbf{x})
\end{array}\right] \in \mathbb{R}^{n}
$$

i.e., a vector that gathers the partial derivatives of $f$.

- Applying the definition of derivative coordinate-wise:

$$
[\nabla f(\mathbf{x})]_{j}=\frac{\partial f}{\partial x_{j}}(\mathbf{x})=\lim _{h \rightarrow 0} \frac{f\left(\mathbf{x}+h \mathbf{e}_{j}\right)-f(\mathbf{x})}{h} \quad j \in\{1, \ldots, n\}
$$

where $\mathbf{e}_{j}=[0,0, \ldots, 0, \underbrace{1}_{j}, 0, \ldots, 0]^{\top} \in\{0,1\}^{n}$ is the $j^{\text {th }}$ standard
basis vector.

## Numerical gradient

■ Finite difference:

$$
[\nabla f(\mathbf{x})]_{j}=\frac{\partial f}{\partial x_{j}}(\mathbf{x}) \approx \frac{f\left(\mathbf{x}+\varepsilon \mathbf{e}_{j}\right)-f(\mathbf{x})}{\varepsilon} \quad j \in\{1, \ldots, n\}
$$

where $\varepsilon$ is a small value (e.g., $10^{-6}$ ).

- Central finite difference:

$$
[\nabla f(\mathbf{x})]_{j}=\frac{\partial f}{\partial x_{j}}(\mathbf{x}) \approx \frac{f\left(\mathbf{x}+\varepsilon \mathbf{e}_{j}\right)-f\left(\mathbf{x}-\varepsilon \mathbf{e}_{j}\right)}{2 \varepsilon} \quad j \in\{1, \ldots, n\}
$$

■ Computing $\nabla f(\mathbf{x})$ approximately by (central) finite difference is $n+1$ times ( $2 n$ times) as costly as evaluating $f$.

## Directional derivative

- Derivative of $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ in the direction of $\mathbf{v} \in \mathbb{R}^{n}$

$$
D_{\mathbf{v}} f(\mathbf{x})=\lim _{h \rightarrow 0} \frac{f(\mathbf{x}+h \mathbf{v})-f(\mathbf{x})}{h} \in \mathbb{R}
$$

$\square$ Interpretation: rate of change of $f$ in the direction of $\mathbf{v}$, when moving away from $\mathbf{x}$.
$\square[\nabla f(\mathbf{x})]_{i}$ is the derivative in the direction of $\mathbf{e}_{i}$.
■ Finite difference (and similarly for the central finite difference):

$$
D_{\mathbf{v}} f(\mathbf{x}) \approx \frac{f(\mathbf{x}+\varepsilon \mathbf{v})-f(\mathbf{x})}{\varepsilon}
$$

Only 2 calls to $f$ are needed, i.e., independent of $n$.

## Directional derivative

- Fact. The directional derivative is equal to the scalar product between the gradient and $\mathbf{v}$, i.e.,

$$
D_{\mathbf{v}} f(\mathbf{x})=\nabla f(\mathbf{x}) \cdot \mathbf{v}
$$

$\square$ Proof. Let $g(t)=f(\mathbf{x}+t \mathbf{v})$. We have

$$
g^{\prime}(t)=\lim _{h \rightarrow 0} \frac{f(\mathbf{x}+(t+h) \mathbf{v})-f(\mathbf{x}+t \mathbf{v})}{h}
$$

and therefore $g^{\prime}(0)=D_{\mathbf{v}}(\mathbf{x})$. By the chain rule, we also have

$$
g^{\prime}(t)=\nabla f(\mathbf{x}+t \mathbf{v}) \cdot \mathbf{v}
$$

Hence, $g^{\prime}(0)=D_{\mathbf{v}}(\mathbf{x})=\nabla f(\mathbf{x}) \cdot \mathbf{v}$.

## Jacobian

$■$ The Jacobian of $\mathbf{f}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}$

$$
\begin{aligned}
J_{f}(\mathbf{x})=\frac{\partial \mathbf{f}(\mathbf{x})}{\partial \mathbf{x}} & =\left[\begin{array}{ccc}
\frac{\partial f_{1}}{\partial x_{1}} & \cdots & \frac{\partial f_{1}}{\partial x_{n}} \\
\vdots & \ddots & \vdots \\
\frac{\partial f_{m}}{\partial x_{1}} & \cdots & \frac{\partial f_{m}}{\partial x_{n}}
\end{array}\right] \\
& =\left[\frac{\partial \mathbf{f}}{\partial x_{1}}, \ldots, \frac{\partial \mathbf{f}}{\partial x_{n}}\right] \\
& =\left[\begin{array}{c}
\nabla f_{1}(\mathbf{x})^{\top} \\
\vdots \\
\nabla f_{m}(\mathbf{x})^{\top}
\end{array}\right]
\end{aligned}
$$

■ The size of the Jacobian matrix is $m \times n$.

- The gradient's transpose is thus a "wide" Jacobian $(m=1)$.


## Jacobian vector product ("JVP")

■ Right-multiply the Jacobian with a vector $\mathbf{v} \in \mathbb{R}^{n}$

$$
\begin{aligned}
J_{\mathbf{f}}(\mathbf{x}) \mathbf{v} & =\left[\begin{array}{c}
\nabla f_{1}(\mathbf{x})^{\top} \\
\vdots \\
\nabla f_{m}(\mathbf{x})^{\top}
\end{array}\right] \mathbf{v} \\
& =\left[\begin{array}{c}
\nabla f_{1}(\mathbf{x}) \cdot \mathbf{v} \\
\vdots \\
\nabla f_{m}(\mathbf{x}) \cdot \mathbf{v}
\end{array}\right] \\
& =\lim _{h \rightarrow 0} \frac{\mathbf{f}(\mathbf{x}+h \mathbf{v})-\mathbf{f}(\mathbf{x})}{h}
\end{aligned}
$$

■ Finite difference (and similarly for the central finite difference):

$$
J_{\mathbf{f}}(\mathbf{x}) \mathbf{v} \approx \frac{\mathbf{f}(\mathbf{x}+\varepsilon \mathbf{v})-\mathbf{f}(\mathbf{x})}{\varepsilon}
$$

■ Computing the JVP approximately by (central) finite difference requires only 2 calls to $f$.

## Vector Jacobian Product ("VJP")

■ Left-multiply the Jacobian with a vector $\mathbf{u} \in \mathbb{R}^{m}$

$$
\mathbf{u}^{\top} J_{\mathbf{f}}(\mathbf{x})=\mathbf{u}^{\top}\left[\frac{\partial \mathbf{f}}{\partial x_{1}}, \ldots, \frac{\partial \mathbf{f}}{\partial x_{n}}\right]=\left[\mathbf{u} \cdot \frac{\partial \mathbf{f}}{\partial x_{1}}, \ldots, \mathbf{u} \cdot \frac{\partial \mathbf{f}}{\partial x_{n}}\right]
$$

- Finite difference (and similarly for the central finite difference):

$$
\frac{\partial \mathbf{f}}{\partial x_{i}} \approx \frac{\mathbf{f}\left(\mathbf{x}+\varepsilon \mathbf{e}_{i}\right)-\mathbf{f}(\mathbf{x})}{\varepsilon}
$$

■ Computing the VJP approximately by (central) finite difference requires $n+1$ calls ( $2 n$ calls) to f .

## Outline

## 1 Numerical differentiation

## 2 Chain compositions

3 Computational graphs

4 Implementation

5 Advanced topics

6 Conclusion

## Chain rule

$\square$ Let $F(x)=f(g(x))=f \circ g(x)$, where $f, g: \mathbb{R} \rightarrow \mathbb{R}$. Then,

$$
F^{\prime}(x)=f^{\prime}(g(x)) g^{\prime}(x)
$$

- Alternatively, let $y=g(x)$ and $z=f(y)$, then

$$
\frac{\partial z}{\partial x}=\frac{\partial z}{\partial y} \frac{\partial y}{\partial x}=\left.\left.\frac{\partial z}{\partial y}\right|_{y=g(x)} \frac{\partial y}{\partial x}\right|_{x=x}
$$

$\square$ Let $f(\mathbf{x})=h(\mathbf{g}(\mathbf{x}))$, where $\mathbf{g}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{d}$ and $h: \mathbb{R}^{d} \rightarrow \mathbb{R}$. Then,

$$
\underbrace{\nabla f(\mathbf{x})}_{n \times 1}=(\underbrace{\nabla h(\mathbf{g}(\mathbf{x}))^{\top}}_{1 \times d} \underbrace{J_{\mathbf{g}}(\mathbf{x})}_{d \times n})^{\top}=\underbrace{J_{\mathbf{g}}(\mathbf{x})^{\top}}_{n \times d} \underbrace{\nabla h(\mathbf{g}(\mathbf{x}))}_{d \times 1}
$$

$\square$ and similarly using Leibniz notation

## Chain compositions



■ Assume $f: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}$ decomposes as follows:

$$
\begin{aligned}
\mathbf{0} & =\mathbf{f}(\mathbf{x}) \\
& =\mathbf{f}_{4} \circ \mathbf{f}_{3} \circ \mathbf{f}_{2} \circ \mathbf{f}_{1}(\mathbf{x}) \\
& =\mathbf{f}_{4}\left(\mathbf{f}_{3}\left(\mathbf{f}_{2}\left(\mathbf{f}_{1}(\mathbf{x})\right)\right)\right)
\end{aligned}
$$

where $\mathbf{f}_{1}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m_{1}}, \mathbf{f}_{2}: \mathbb{R}^{m_{1}} \rightarrow \mathbb{R}^{m_{2}}, \ldots, \mathbf{f}_{4}: \mathbb{R}^{m_{3}} \rightarrow \mathbb{R}^{m}$.

- How to compute the Jacobian $J_{f}(x)=\frac{\partial \circ}{\partial \mathrm{x}} \in \mathbb{R}^{m \times n}$ efficiently?


## Chain rule

■ Sequence of operations

$$
\begin{aligned}
\mathbf{x}_{1} & =\mathbf{x} \\
\mathbf{x}_{2} & =\mathbf{f}_{1}\left(\mathbf{x}_{1}\right) \\
\mathbf{x}_{3} & =\mathbf{f}_{2}\left(\mathbf{x}_{2}\right) \\
\mathbf{x}_{4} & =\mathbf{f}_{3}\left(\mathbf{x}_{3}\right) \\
\mathbf{0} & =\mathbf{f}_{4}\left(\mathbf{x}_{4}\right)
\end{aligned}
$$

- By the chain rule, we have

$$
\begin{aligned}
\frac{\partial \mathbf{o}}{\partial \mathbf{x}} & =\frac{\partial o}{\partial \mathbf{x}_{4}} \frac{\partial \mathbf{x}_{4}}{\partial \mathbf{x}_{3}} \frac{\partial \mathbf{x}_{3}}{\partial \mathbf{x}_{2}} \frac{\partial \mathbf{x}_{2}}{\partial \mathbf{x}} \\
& =\frac{\partial \mathbf{f}_{4}\left(\mathbf{x}_{4}\right)}{\partial \mathbf{x}_{4}} \frac{\partial \mathbf{f}_{3}\left(\mathbf{x}_{3}\right)}{\partial \mathbf{x}_{3}} \frac{\partial \mathbf{f}_{2}\left(\mathbf{x}_{2}\right)}{\partial \mathbf{x}_{2}} \frac{\partial \mathbf{f}_{1}(\mathbf{x})}{\partial \mathbf{x}} \\
& =J_{\mathbf{f}_{4}}\left(\mathbf{x}_{4}\right) J_{\mathbf{f}_{3}}\left(\mathbf{x}_{3}\right) J_{\mathbf{f}_{2}}\left(\mathbf{x}_{2}\right) J_{\mathbf{f}_{1}}(\mathbf{x})
\end{aligned}
$$

## Forward differentiation

$\square$ Recall that $\frac{\partial \mathbf{f}}{\partial x_{j}} \in \mathbb{R}^{m}$ is the $j^{\text {th }}$ column of $J_{\mathbf{f}}(\mathbf{x})$.
■ Jacobian vector product (JVP) with $\mathbf{e}_{j} \in \mathbb{R}^{n}$ extracts the $j^{\text {th }}$ column

$$
\begin{aligned}
J_{\mathbf{f}}(\mathbf{x}) \mathbf{e}_{1} & =\frac{\partial \mathbf{f}}{\partial x_{1}} \\
J_{\mathbf{f}}(\mathbf{x}) \mathbf{e}_{2} & =\frac{\partial \mathbf{f}}{\partial x_{2}} \\
\vdots & \\
J_{\mathbf{f}}(\mathbf{x}) \mathbf{e}_{n} & =\frac{\partial \mathbf{f}}{\partial x_{n}}
\end{aligned}
$$

■ Computing a gradient ( $m=1$ ) requires $n$ JVPs with $\mathbf{e}_{1}, \ldots, \mathbf{e}_{n}$.

## Forward differentiation

■ Jacobian-vector product with $v \in \mathbb{R}^{n}$

$$
J_{\mathbf{f}}(\mathbf{x}) \vee=\underbrace{J_{\mathbf{f}_{4}}\left(\mathbf{x}_{4}\right)}_{m \times m_{3}} \underbrace{J_{\mathbf{t}_{3}}\left(\mathbf{x}_{3}\right)}_{m_{3} \times m_{2}} \underbrace{J_{\mathbf{f}_{2}}\left(\mathbf{x}_{2}\right)}_{m_{2} \times m_{1}} \underbrace{J_{\mathbf{f}_{1}}(\mathbf{x})}_{m_{1} \times n}
$$

Multiplication from right to left is more efficient.
■ Cost of computing $n$ JVPs:

$$
n\left(m m_{3}+m_{3} m_{2}+m_{2} m_{1}+m_{1} n\right)
$$

$\square$ Cost of computing a gradient $\left(m=1, m_{3}=m_{2}=m_{1}=n\right)$ :

$$
O\left(n^{3}\right)
$$

## Forward differentiation

$\square 0=\mathbf{f}(\mathbf{x})=\mathbf{f}_{K} \circ \cdots \circ \mathbf{f}_{2} \circ \mathbf{f}_{1}(\mathbf{x})$
$\square\left[\mathcal{J}_{\mathbf{f}}(\mathbf{x})\right]_{:, j}=J_{\mathbf{f}_{K}}\left(\mathbf{x}_{K}\right) \ldots J_{\mathbf{f}_{2}}\left(\mathbf{x}_{2}\right) \mathcal{J}_{\mathbf{f}_{1}}(\mathbf{x}) \mathbf{e}_{j} \quad j \in\{1, \ldots, n\}$
Algorithm 1 Compute $0=\mathbf{f}(\mathbf{x})$ and $J_{f}(\mathbf{x})$ alongside
1: Input: $x \in \mathbb{R}^{n}$
2: $\mathbf{x}_{1} \leftarrow \mathbf{x}$
3: $\mathrm{v}_{j} \leftarrow \mathbf{e}_{j} \in \mathbb{R}^{n} \quad j \in\{1, \ldots, n\}$
4: for $k=1$ to $K$ do
5: $\quad \mathbf{x}_{k+1} \leftarrow \mathbf{f}_{k}\left(\mathbf{x}_{k}\right)$
6: $\quad \mathrm{v}_{j} \leftarrow \mathrm{~J}_{\mathrm{f}_{k}}\left(\mathbf{x}_{k}\right) \mathrm{v}_{j} j \in\{1, \ldots, n\}$
7: end for
8: Returns: $\mathbf{o}=\mathbf{x}_{K+1},\left[J_{\mathfrak{f}}(\mathbf{x})\right]_{:, j}=\mathrm{v}_{j} j \in\{1, \ldots, n\}$

## Backward differentiation

$\square$ Recall that $\nabla f_{i}(\mathbf{x})^{\top} \in \mathbb{R}^{n}$ is the $i^{\text {th }}$ row of $J_{\mathbf{f}}(\mathbf{x})$.
■ Vector Jacobian product (VJP) with $\mathbf{e}_{i} \in \mathbb{R}^{m}$ extracts the $i^{\text {th }}$ row

$$
\begin{gathered}
\mathbf{e}_{1}^{\top} J_{\mathbf{f}}(\mathbf{x})=\nabla f_{1}(\mathbf{x})^{\top} \\
\mathbf{e}_{2}^{\top} J_{\mathbf{f}}(\mathbf{x})=\nabla f_{2}(\mathbf{x})^{\top} \\
\vdots \\
\mathbf{e}_{m}^{\top} J_{\mathbf{f}}(\mathbf{x})=\nabla f_{m}(\mathbf{x})^{\top}
\end{gathered}
$$

■ Computing a gradient $(m=1)$ requires only 1 VJP with $\mathbf{e}_{1} \in \mathbb{R}^{1}$.

## Backward differentiation

- Vector Jacobian product with $u \in \mathbb{R}^{m}$

$$
\mathbf{u}^{\top} \underbrace{J_{\mathbf{f}_{4}}\left(\mathbf{X}_{4}\right)}_{m \times m_{3}} \underbrace{J_{\mathbf{f}_{3}}\left(\mathbf{X}_{3}\right)}_{m_{3} \times m_{2}} \underbrace{J_{\mathbf{f}_{2}}\left(\mathbf{X}_{2}\right)}_{m_{2} \times m_{1}} \underbrace{J_{\mathbf{f}_{1}}(\mathbf{X})}_{m_{1} \times n}
$$

Multiplication from left to right is more efficient.

- Cost of computing $m$ VJPs:

$$
m\left(m m_{3}+m_{3} m_{2}+m_{2} m_{1}+m_{1} n\right)
$$

$\square$ Cost of computing a gradient $\left(m=1, m_{3}=m_{2}=m_{1}=n\right)$ :

$$
O\left(n^{2}\right)
$$

## Backward differentiation

$$
\begin{aligned}
& \square 0=\mathbf{f}(\mathbf{x})=\mathbf{f}_{K} \circ \cdots \circ \mathbf{f}_{2} \circ \mathbf{f}_{1}(\mathbf{x}) \\
& \square\left[J_{\mathbf{f}}(\mathbf{x})\right]_{i,:}=\mathbf{e}_{i}^{\top} J_{\mathbf{f}_{K}}\left(\mathbf{x}_{K}\right) \ldots J_{\mathbf{f}_{2}}\left(\mathbf{x}_{2}\right) J_{\mathbf{f}_{1}}(\mathbf{x}) \quad i \in\{1, \ldots, m\}
\end{aligned}
$$

Algorithm 2 Compute $0=\mathbf{f}(\mathbf{x})$ and $J_{\mathbf{f}}(\mathbf{x})$
1: Input: $x \in \mathbb{R}^{n}$
2: $\mathbf{x}_{1} \leftarrow \mathbf{x}, \mathrm{u}_{i} \leftarrow \mathbf{e}_{i} \in \mathbb{R}^{m} \quad i \in\{1, \ldots, m\}$
3: for $k=1$ to $K$ do
4: $\quad \mathbf{x}_{k+1} \leftarrow \mathbf{f}_{k}\left(\mathbf{x}_{k}\right)$
5: end for
6: for $k=K$ to 1 do
7: $\quad u_{i}^{\top} \leftarrow u_{i}^{\top} J_{f_{k}}\left(\mathbf{x}_{k}\right) \quad i \in\{1, \ldots, m\}$
8: end for
9: Returns: $0=\mathbf{x}_{K+1},\left[J_{\mathbf{f}}(\mathrm{x})\right]_{i,:}=u_{i}^{\top} \quad i \in\{1, \ldots, m\}$

## Feedforward networks


$\square$ Each function can now have two arguments: $\mathbf{f}_{k}\left(\mathbf{x}_{k}, \theta_{k}\right)$, where $\mathbf{x}_{k}$ is the previous output and $\theta_{k}$ are learnable parameters.

■ Example one hidden layer, one output layer, squared loss

$$
\begin{array}{rlrl}
\mathbf{f} & =\mathbf{f}_{4} \circ \cdots \circ \mathbf{f}_{1} & \\
\mathbf{x}_{2} & =\mathbf{f}_{1}\left(\mathbf{x}, W_{1}\right)=W_{1} \mathbf{x} & W_{1} \in \mathbb{R}^{m_{1} \times n} \\
\mathbf{x}_{3} & =\mathbf{f}_{2}\left(\mathbf{x}_{2}, \emptyset\right)=\operatorname{relu}\left(\mathbf{x}_{2}\right) & & \\
\mathbf{x}_{4} & =\mathbf{f}_{3}\left(\mathbf{x}_{3}, W_{3}\right)=W_{3} \mathbf{x}_{3} & W_{3} \in \mathbb{R}^{1 \times m_{3}} \\
0 & =\mathbf{f}_{4}\left(\mathbf{x}_{4}, \mathbf{y}\right)=\frac{1}{2}\left\|\mathbf{x}_{4}-\mathbf{y}\right\|^{2} & &
\end{array}
$$

## Feedforward network example



- Applying the chain rule once again we have

$$
\begin{aligned}
& \frac{\partial \mathrm{o}}{\partial \theta_{4}} \\
& \frac{\partial \mathrm{o}}{\partial \theta_{3}}=\frac{\partial \mathrm{o}}{\partial \mathbf{x}_{4}} \frac{\partial \mathbf{x}_{4}}{\partial \theta_{3}} \\
& \frac{\partial \mathrm{o}}{\partial \theta_{2}}=\frac{\partial \mathrm{o}}{\partial \mathbf{x}_{4}} \frac{\partial \mathbf{x}_{4}}{\partial \mathbf{x}_{3}} \frac{\partial \mathbf{x}_{3}}{\partial \theta_{2}}
\end{aligned}
$$

$\square$ Apart from the last multiplication, the Jacobians $\frac{\partial \circ}{\partial \mathbf{x}_{k}}$ and $\frac{\partial \circ}{\partial \theta_{k}}$ share the same computations!

## Backprop for feedforward networks

Algorithm 3 Compute $0=\mathbf{f}\left(\mathbf{x}, \theta_{1}, \ldots, \theta_{K}\right)$ and its Jacobians.
1: Input: $x \in \mathbb{R}^{n}, \theta_{1}, \ldots, \theta_{K}$
2: $\mathbf{x}_{1} \leftarrow \mathbf{x}$
3: $u_{i} \leftarrow \mathbf{e}_{i} \in \mathbb{R}^{m} \quad i \in\{1, \ldots, m\}$
4: for $k=1$ to $K$ do
5: $\quad \mathbf{x}_{k+1} \leftarrow \mathbf{f}_{k}\left(\mathbf{x}_{k}, \theta_{k}\right)$
6: end for
7: for $k=K$ to 1 do
8: $\quad \mathrm{j}_{i, k} \leftarrow \mathrm{u}_{i}^{\top} \frac{\partial \mathbf{f}_{k}\left(\mathbf{x}_{k}, \theta_{k}\right)}{\partial \theta_{k}} \quad i \in\{1, \ldots, m\}$
9: $\quad \mathrm{u}_{i}^{\top} \leftarrow \mathrm{u}_{i}^{\top} \frac{\partial \mathbf{f}_{k}\left(\mathbf{x}_{k}, \theta_{k}\right)}{\partial \mathbf{x}_{k}} \quad i \in\{1, \ldots, m\}$
10: end for
11: Returns: $o=\mathbf{x}_{K+1},\left[\frac{\partial \circ}{\partial \mathrm{x}}\right]_{i,:}=u_{i}^{\top},\left[\frac{\partial \circ}{\partial \theta_{k}}\right]_{i,:}=\mathrm{j}_{i, k} \quad i \in\{1, \ldots, m\}, k \in\{1, \ldots, K\}$

## Examples of VJPs

Let $W \in \mathbb{R}^{a \times b}, u \in \mathbb{R}^{a}, x \in \mathbb{R}^{b}$.
$\square \mathbf{f}(x)=g(x)$ (element-wise)
■ $\mathbf{f}$ maps $\mathbb{R}^{b}$ to $\mathbb{R}^{b}$
$\square J_{\mathbf{f}}(x)=J_{\mathbf{f}}(x)^{\top}=\operatorname{diag}\left(g^{\prime}(x)\right)$ maps $\mathbb{R}^{b}$ to $\mathbb{R}^{b}$, i.e., $b \times b$ matrix
■ $u^{\top} J_{\mathfrak{f}}(x)=J_{\mathfrak{f}}(x)^{\top} u=u * g^{\prime}(x) \in \mathbb{R}^{b}$, where $*$ means element-wise multiplication
$\square \mathbf{f}(x)=W x$

- $\mathbf{f} \operatorname{maps} \mathbb{R}^{b}$ to $\mathbb{R}^{a}$

■ $J_{\mathbf{f}}(x)=W$ maps $\mathbb{R}^{b}$ to $\mathbb{R}^{a}$, i.e., $a \times b$ matrix
$\square J_{\mathrm{f}}(x)^{\top}=W^{\top}$ maps $\mathbb{R}^{a}$ to $\mathbb{R}^{b}$, i.e., $b \times$ a matrix

- $u^{\top} J_{\mathfrak{f}}(x)=J_{\mathfrak{f}}(x)^{\top} u=W^{\top} u \in \mathbb{R}^{b}$


## Examples of VJPs

$\mathbf{f}(W)=W x$
$■ \mathbf{f}$ maps $\mathbb{R}^{a \times b}$ to $\mathbb{R}^{a}$
■ $J_{\mathrm{f}}(W)$ maps $\mathbb{R}^{a \times b}$ to $\mathbb{R}^{a}$, i.e., $a \times(a \times b)$ matrix
$\square J_{\mathbf{f}}(W)^{\top}$ maps $\mathbb{R}^{a}$ to $\mathbb{R}^{a \times b}$, i.e., $(a \times b) \times$ a matrix

- $J_{\mathfrak{f}}(W)^{\top} u=u x^{\top}$

VJPs make things easier when dealing with matrix or tensor inputs.

## Summary: Forward vs. Backward

- Forward

■ Uses Jacobian vector products (JVPs)
■ Each JVP call builds one column of the Jacobian

- Efficient for tall Jacobians ( $m \geq n$ )

■ Need not store intermediate computations
■ Backward
■ Uses vector Jacobian products (VJPs)
■ Each VJP call builds one row of the Jacobian

- Efficient for wide matrices ( $m \leq n$ )

■ Needs to store intermediate computations

## Machine learning use case

- Most objectives in machine learning can be written in the form

$$
\min _{\mathbf{x} \in \mathbb{R}^{n}} f(\mathbf{x})=\sum_{i=1}^{N} \ell_{i}\left(f_{i}(\mathbf{x})\right)
$$

where $f: \mathbb{R}^{n} \rightarrow \mathbb{R}^{M}$ and $\ell_{i}: \mathbb{R}^{M} \rightarrow \mathbb{R}$.

- The minimization needs to be w.r.t. a scalar valued loss.
- This corresponds to the $m=1$ setting, for which backward differentiation is more efficient.
- This explains the immense success of reverse autodiff in machine learning.


## Outline

## 1 Numerical differentiation

## 2 Chain compositions

## 3 Computational graphs

## 4 Implementation

5 Advanced topics

6 Conclusion

## Computational graph

$$
f\left(x_{1}, x_{2}\right)=x_{2} e^{x_{1}} \sqrt{x_{1}+x_{2} e^{x_{1}}}
$$

■ Operations in topological order

$$
\begin{aligned}
& x_{3}=f_{3}\left(x_{1}\right)=e^{x_{1}} \\
& x_{4}=f_{4}\left(x_{2}, x_{3}\right)=x_{2} x_{3} \\
& x_{5}=f_{5}\left(x_{1}, x_{4}\right)=x_{1}+x_{4} \\
& x_{6}=f_{6}\left(x_{5}\right)=\sqrt{x_{5}} \\
& x_{7}=f_{7}\left(x_{4}, x_{6}\right)=x_{4} x_{6}
\end{aligned}
$$

- Directed acyclic graph traversal



## Forward differentiation example


$\square \mathbf{x}_{4}$ is influenced by $\mathbf{x}_{3}$ and $\mathbf{x}_{2}$, therefore

$$
\frac{\partial \mathbf{x}_{4}}{\partial \mathbf{x}_{1}}=\frac{\partial \mathbf{x}_{4}}{\partial \mathbf{x}_{3}} \frac{\partial \mathbf{x}_{3}}{\partial \mathbf{x}_{1}}+\frac{\partial \mathbf{x}_{4}}{\partial \mathbf{x}_{2}} \frac{\partial \mathbf{x}_{2}}{\partial \mathbf{x}_{1}}
$$

$\square \mathbf{x}_{7}$ is influenced by $\mathbf{x}_{4}$ and $\mathbf{x}_{6}$, therefore

$$
\frac{\partial \mathbf{x}_{7}}{\partial \mathbf{x}_{1}}=\frac{\partial \mathbf{x}_{7}}{\partial \mathbf{x}_{4}} \frac{\partial \mathbf{x}_{4}}{\partial \mathbf{x}_{1}}+\frac{\partial \mathbf{x}_{7}}{\partial \mathbf{x}_{6}} \frac{\partial \mathbf{x}_{6}}{\partial \mathbf{x}_{1}}
$$

## Forward differentiation example



- Recurse in topological order

$$
\begin{aligned}
& \frac{\partial \mathbf{x}_{1}}{\partial \mathbf{x}_{1}}=\operatorname{ld} \mathrm{d}_{n} \\
& \frac{\partial \mathbf{x}_{2}}{\partial \mathbf{x}_{2}}=\operatorname{ld} \mathbf{d}_{n} \\
& \frac{\partial \mathbf{x}_{3}}{\partial \mathbf{x}_{1}}=\frac{\partial \mathbf{x}_{3}}{\partial \mathbf{x}_{1}} \frac{\partial \mathbf{x}_{1}}{\partial \mathbf{x}_{1}} \\
& \frac{\partial \mathbf{x}_{4}}{\partial \mathbf{x}_{1}}=\frac{\partial \mathbf{x}_{4}}{\partial \mathbf{x}_{3}} \frac{\partial \mathbf{x}_{3}}{\partial \mathbf{x}_{1}}+\frac{\partial \mathbf{x}_{4}}{\partial \mathbf{x}_{2}} \frac{\partial \mathbf{x}_{2}}{\partial \mathbf{x}_{1}}
\end{aligned}
$$

■ Everything can be computed in terms of JVPs

## Forward differentiation



■ In the general case, we have

$$
\frac{\partial x_{j}}{\partial \mathbf{x}_{1}}=\sum_{i \in \operatorname{Parents}(j)} \frac{\partial x_{j}}{\partial \mathbf{x}_{i}} \frac{\partial \mathbf{x}_{i}}{\partial \mathbf{x}_{1}}
$$

$\square \frac{\partial x_{j}}{\partial \mathbf{x}_{i}}$ is easy to compute as $f_{j}$ is a direct function of $\mathbf{x}_{i}$.
$\square \frac{\partial \mathbf{x}_{i}}{\partial \mathrm{x}_{1}}$ is obtained from the previous iterations in topological order.

## Backward differentiation example



■ $\mathbf{x}_{5}$ influences only $\mathbf{x}_{6}$, therefore

$$
\frac{\partial \mathrm{o}}{\partial \mathbf{x}_{5}}=\frac{\partial \mathrm{o}}{\partial \mathbf{x}_{6}} \frac{\partial \mathbf{x}_{6}}{\partial \mathbf{x}_{5}}
$$

■ $\mathbf{x}_{4}$ influences $\mathbf{x}_{5}$ and $\mathbf{x}_{7}$, therefore

$$
\frac{\partial \mathbf{o}}{\partial \mathbf{x}_{4}}=\frac{\partial \mathbf{o}}{\partial \mathbf{x}_{5}} \frac{\partial \mathbf{x}_{5}}{\partial \mathbf{x}_{4}}+\frac{\partial \mathbf{o}}{\partial \mathbf{x}_{7}} \frac{\partial \mathbf{x}_{7}}{\partial \mathbf{x}_{4}}
$$

## Backward differentiation example



■ Recurse in reverse topological order

$$
\begin{aligned}
& \frac{\partial \mathrm{o}}{\partial \mathbf{x}_{7}}=\frac{\partial \mathbf{x}_{7}}{\partial \mathbf{x}_{7}}=\mathrm{Id} \\
& m \\
& \frac{\partial \mathrm{o}}{\partial \mathbf{x}_{6}}=\frac{\partial \mathrm{o}}{\partial \mathbf{x}_{7}} \frac{\partial \mathbf{x}_{7}}{\partial \mathbf{x}_{6}} \\
& \frac{\partial \mathrm{o}}{\partial \mathbf{x}_{5}}=\frac{\partial \mathrm{o}}{\partial \mathbf{x}_{6}} \frac{\partial \mathbf{x}_{6}}{\partial \mathbf{x}_{5}} \\
& \frac{\partial \mathrm{o}}{\partial \mathbf{x}_{4}}=\frac{\partial \mathrm{o}}{\partial \mathbf{x}_{5}} \frac{\partial \mathbf{x}_{5}}{\partial \mathbf{x}_{4}}+\frac{\partial \mathrm{o}}{\partial \mathbf{x}_{7}} \frac{\partial \mathbf{x}_{7}}{\partial \mathbf{x}_{4}}
\end{aligned}
$$

■ Everything can be computed in terms of VJPs

## Backward differentiation



■ In the general case, we have

$$
\frac{\partial o}{\partial x_{j}}=\sum_{k \in \text { Children }(j)} \frac{\partial o}{\partial x_{k}} \frac{\partial x_{k}}{\partial x_{j}}
$$

- $\frac{\partial 0}{\partial \mathrm{x}_{k}}$ is obtained from previous iterations (reverse topological order) and is known as "adjoint".
$\square \frac{\partial x_{k}}{\partial x_{j}}$ is easy to compute as $f_{k}$ is a direct function of $x_{j}$.


## Outline

## 1 Numerical differentiation

## 2 Chain compositions

3 Computational graphs

## 4 Implementation

## 5 Advanced topics

6 Conclusion

## Obtaining the computational graph

- Ahead of time

■ Read from source or abstract syntax tree (AST). Ex: Tangent.

- API for composing primitive operations (the graph is fully built before the program is evaluated). Ex: Tensorflow.
- Just in time
- Tracing: monitor the program execution (the graph is built while the program is being executed). Ex: Tensorflow Eager, JAX, PyTorch.

```
import jax.numpy as jnp
from jax import grad
def add(a, b):
    return a + b
a = jnp.array([1, 2, 3])
b = jnp.array([4, 5, 6])
print(grad(add)(a, b))
```


## Key components of an implementation

- VJP for all primitive operations
- Node class
- Topological sort
- Forward pass
- Backward pass

We will now briefly review each component using a rudimentary implementation (link to code).

## VJPs for primitive operations

```
def dot(x, W):
    return np.dot(W, x)
def dot_make_vjp(x, W):
    def vjp(u):
        return W.T.dot(u), np.outer(u, x)
    return vjp
dot.make_vjp = dot_make_vjp
def add(a, b):
    return a + b
def add_make_vjp(a, b):
    gprime = np.ones(len(a))
    def vjp(u):
        return u * gprime, u * gprime
    return vjp
add.make_vjp = add_make_vjp
```


## Node class

```
class Node(object):
```

```
def __init__(self, value=None, func=None, parents=None, name="")
```

def __init__(self, value=None, func=None, parents=None, name="")
\# Value stored in the node.
\# Value stored in the node.
self.value = value
self.value = value
\# Function producing the node.
\# Function producing the node.
self.func = func
self.func = func
\# Inputs to the function.
\# Inputs to the function.
self.parents = [] if parents is None else parents
self.parents = [] if parents is None else parents
\# Unique name of the node (for debugging and hashing).
\# Unique name of the node (for debugging and hashing).
self.name = name
self.name = name
\# Gradient / Jacobian.
\# Gradient / Jacobian.
self.grad = 0
self.grad = 0
if not name:
if not name:
raise ValueError("Each node must have a unique name.")
raise ValueError("Each node must have a unique name.")
def __hash__(self):
return hash(self.name)
def __repr__(self):
return "Node(%s)" % self.name

```

\section*{DAG}

```

def create_dag(x):
x1 = Node(value=np.array([x[0]]), name="x1")
x2 = Node(value=np.array([x[1]]), name="x2")
x3 = Node(func=exp, parents=[x1], name="x3")
x4 = Node(func=mul, parents=[x2, x3], name="x4")
x5 = Node(func=add, parents=[x1, x4], name="x5")
x6 = Node(func=sqrt, parents=[x5], name="x6")
x7 = Node(func=mul, parents=[x4, x6], name="x7")
return x7

```

A good implementation would support tracing, instead of building the DAG manually.

\section*{Topological sort}
```

def dfs(node, visited):
visited.add(node)
for parent in node.parents:
if not parent in visited:
\# Yield parent nodes first.
yield from dfs(parent, visited)
\# And current node later.
yield node
def topological_sort(end_node):
visited = set()
sorted_nodes = []

```
    \# All non-visited nodes reachable from end_node.
    for node in dfs(end_node, visited):
        sorted_nodes.append (node)
    return sorted_nodes

\section*{Forward pass}
def evaluate_dag(sorted_nodes):
for node in sorted_nodes:
if node.value is None:
values = [p.value for \(p\) in node.parents] node.value \(=\) node.func(*values)
return sorted_nodes [-1].value

\section*{Backward pass}
```

def backward_diff_dag(sorted_nodes):
value = evaluate_dag(sorted_nodes)
m = value.shape[0] \# Output size
\# Initialize recursion.
sorted_nodes[-1].grad = np.eye(m)
for node_k in reversed(sorted_nodes):
if not node_k.parents:
\# We reached a node without parents.
continue

```

    \# Values of the parent nodes.
    values = [p.value for \(p\) in node_k.parents]
    \# Iterate over outputs.
    for i in range(m):
            \# A list of size len(values) containing the vjps.
            vjps = node_k.func.make_vjp(*values)(node_k.grad[i])
            for node_j, vjp in zip(node_k.parents, vjps):
            node_j.grad += vjp
    return sorted_nodes

\section*{Checkoointing (best seen in presentation mode)}
- During the forward pass, save computations at intermediate locations only (checkpoints).
- During the backward pass, recompute other locations on the fly, starting from the checkpoints.
- Tradeoff between memory and computation time.


\section*{Checkoointing (best seen in presentation mode)}
- During the forward pass, save computations at intermediate locations only (checkpoints).
- During the backward pass, recompute other locations on the fly, starting from the checkpoints.
- Tradeoff between memory and computation time.


\section*{Checkoointing (best seen in presentation mode)}
- During the forward pass, save computations at intermediate locations only (checkpoints).
- During the backward pass, recompute other locations on the fly, starting from the checkpoints.
- Tradeoff between memory and computation time.


\section*{Checkoointing (best seen in presentation mode)}
- During the forward pass, save computations at intermediate locations only (checkpoints).
- During the backward pass, recompute other locations on the fly, starting from the checkpoints.
- Tradeoff between memory and computation time.


\section*{Checkoointing (best seen in presentation mode)}
- During the forward pass, save computations at intermediate locations only (checkpoints).
- During the backward pass, recompute other locations on the fly, starting from the checkpoints.
- Tradeoff between memory and computation time.


\section*{Checkoointing (best seen in presentation mode)}
- During the forward pass, save computations at intermediate locations only (checkpoints).
- During the backward pass, recompute other locations on the fly, starting from the checkpoints.
- Tradeoff between memory and computation time.


\section*{Checkoointing (best seen in presentation mode)}
- During the forward pass, save computations at intermediate locations only (checkpoints).
- During the backward pass, recompute other locations on the fly, starting from the checkpoints.
- Tradeoff between memory and computation time.


\section*{Checkoointing (best seen in presentation mode)}
- During the forward pass, save computations at intermediate locations only (checkpoints).
- During the backward pass, recompute other locations on the fly, starting from the checkpoints.
- Tradeoff between memory and computation time.


\section*{Checkoointing (best seen in presentation mode)}

■ During the forward pass, save computations at intermediate locations only (checkpoints).
- During the backward pass, recompute other locations on the fly, starting from the checkpoints.
- Tradeoff between memory and computation time.


■ NumPy and SciPy compatible
- Automatic differentiation (grad)
- Just-in-time compilation (jit)

- Automatic vectorization (vmap)
- Code transformations are composable
- Actively developed by Google

■ Gaining a lot of popularity among ML and science researchers

\section*{Outline}

\section*{1 Numerical differentiation}

\section*{2 Chain compositions}

3 Computational graphs

4 Implementation

\section*{5 Advanced topics}

\section*{6 Conclusion}

\section*{Hessian}
- The matrix gathering second-order derivatives
\[
\nabla^{2} f=\left[\begin{array}{cccc}
\frac{\partial^{2} f}{\partial x_{1}^{2}} & \frac{\partial^{2} f}{\partial x_{1} \partial x_{2}} & \cdots & \frac{\partial^{2} f}{\partial x_{1} \partial x_{n}} \\
\frac{\partial^{2} f}{\partial x_{2} \partial x_{1}} & \frac{\partial^{2} f}{\partial x_{2}^{2}} & \cdots & \frac{\partial^{2} f}{\partial x_{2} \partial x_{n}} \\
\vdots & \vdots & \ddots & \vdots \\
\frac{\partial^{2} f}{\partial x_{n} \partial x_{1}} & \frac{\partial^{2} f}{\partial x_{n} \partial x_{2}} & \cdots & \frac{\partial^{2} f}{\partial x_{n}^{2}}
\end{array}\right]
\]

■ Hessian vector product = gradient of directional derivative
\[
\nabla^{2} f(\mathbf{x}) \mathbf{v}=\nabla(\nabla f(\mathbf{x}) \cdot \mathbf{v})
\]

■ JAX supports fully closed tracing: we can "trace through tracing"

\section*{Recovering JVPs from VJPs}

■ Suppose we already have a VJP routine for computing \(u^{\top} J_{f}(x)\)
■ By linearity we have
\[
\frac{\partial \mathrm{u}^{\top} J_{\mathbf{f}}(\mathbf{x})}{\partial \mathrm{u}}=J_{\mathbf{f}}(\mathbf{x})^{\top}
\]
- and therefore
\[
\mathbf{v}^{\top} \frac{\partial \mathrm{u}^{\top} J_{\mathbf{f}}(\mathbf{x})}{\partial \mathrm{u}}=\mathbf{v}^{\top} J_{\mathbf{f}}(\mathbf{x})^{\top}=\left(\mathcal{J}_{\mathbf{f}}(\mathbf{x}) \mathbf{v}\right)^{\top}
\]
\(■\) The VJP w.r.t. \(u\) of the VJP w.r.t. x is equal to the transopose of the JVP w.r.t. x.

■ The trick does not work in the other direction!

\section*{Differentiating min problems}
- Consider the function
\[
f(\theta)=\min _{x} E(x, \theta)=E\left(x^{\star}(\theta), \theta\right)
\]

■ From Danskin's theorem (a.k.a. envelope theorem)
\[
\nabla f(\theta)=\nabla_{2} E\left(x^{\star}(\theta), \theta\right)
\]
where \(\nabla_{2}\) indicates the gradient w.r.t. the second argument.
- Informally, the theorem says that we can treat \(x^{\star}(\theta)\) as if it did not depend on \(\theta\).

\section*{Differentiating argmin problems}

■ Now, consider the function
\[
\begin{aligned}
x^{\star}(\theta) & =\underset{x}{\operatorname{argmin}} E(x, \theta) \\
f(\theta) & =L\left(x^{\star}(\theta), \theta\right)
\end{aligned}
\]
- By the chain rule, we have
\[
\nabla f(\theta)=\left(J x^{\star}(\theta)\right)^{\top} \nabla_{1} L\left(x^{\star}(\theta), \theta\right)+\nabla_{2} L\left(x^{\star}(\theta), \theta\right)
\]
- How to compute \(J x^{\star}(\theta)=\frac{\partial x^{\star}(\theta)}{\partial \theta}\) ?

\section*{Fixed points}
- Consider the following fixed point iteration
\[
x^{\star}(\theta)=g\left(x^{\star}(\theta), \theta\right) \Leftrightarrow h\left(x^{\star}(\theta), \theta\right)=0
\]
where \(h(x, \theta)=x-g(x, \theta)\)
- By the implicit function theorem
\[
J x^{\star}(\theta)=-\left(J_{1} h\left(x^{\star}(\theta), \theta\right)\right)^{-1} J_{2} h\left(x^{\star}(\theta), \theta\right)
\]
where \(J_{1}\) and \(J_{2}\) are the Jacobians w.r.t. the 1st and 2nd variables

\section*{Differentiating argmin problems}
- Recall that
\[
x^{\star}(\theta)=\underset{x}{\operatorname{argmin}} E(x, \theta)
\]
- We have the fixed point iteration (gradient descent)
\[
x^{\star}(\theta)=x^{\star}(\theta)-\nabla_{1} E\left(x^{\star}(\theta), \theta\right)
\]

■ Choosing \(h(x, \theta)=\nabla_{1} E(x, \theta)\), we get
\[
\begin{aligned}
J x^{\star}(\theta) & =-\left(J_{1} \nabla_{1} E\left(x^{\star}(\theta), \theta\right)\right)^{-1} J_{2} \nabla_{1} E\left(x^{\star}(\theta), \theta\right) \\
& =-\left(\nabla_{1}^{2} E\left(x^{\star}(\theta), \theta\right)\right)^{-1} J_{2} \nabla_{1} E\left(x^{\star}(\theta), \theta\right)
\end{aligned}
\]

■ In practice, we need to replace \(x^{\star}(\theta)\) by an approximate solution.

\section*{Differentiating argmin problems}

■ Example: hyper-parameter optimization for ridge regression
\[
\begin{aligned}
E(x, \theta) & =\frac{1}{2}\|A x-b\|^{2}+\frac{\theta}{2}\|x\|^{2} \in \mathbb{R} \\
\nabla_{1} E(x, \theta) & =A^{\top}(A x-b)+\theta x \in \mathbb{R}^{d} \\
\nabla_{1}^{2} E(x, \theta) & =A^{\top} A+\theta I \in \mathbb{R}^{d \times d} \\
J_{2} \nabla_{1} E(x, \theta) & =x \in \mathbb{R}^{d \times 1} \\
x^{\star}(\theta) & =\left(A^{\top} A+\theta I\right)^{-1} A^{\top} b
\end{aligned}
\]

■ \(J x^{\star}(\theta)\) is therefore obtained by solving the following linear system
\[
\left(A^{\top} A+\theta I\right)\left[J x^{\star}(\theta)\right]=-x^{\star}(\theta)
\]

\section*{Differentiating argmin problems}
- An alternative idea to obtain \(J x^{\star}(\theta)\) is to to backpropagate through gradient descent:
\[
x^{t+1}(\theta)=x^{t}(\theta)-\eta_{t} \nabla_{1} E\left(x^{t}(\theta), \theta\right)
\]
- No longer needs to solve a linear system...

■ ...but needs to store intermediate iterates \(x^{t}(\theta)\) or checkpoints
- Possibility to use truncated backpropagation

■ Possibility to use reversible dynamics in some cases

\section*{Inference in graphical models}

■ Gibbs distribution
\[
\mathbb{P}(Y=y ; \theta) \propto \exp (y \cdot \theta)
\]
where \(y \in \mathcal{Y} \subset\{0,1\}^{n}\)
- Log-partition function
\[
f(\theta)=\log \sum_{y \in \mathcal{Y}} \exp (y \cdot \theta)
\]
- Fact.
\[
\left(\mathbb{P}\left(Y_{i}=1 ; \theta\right)\right)_{i=1}^{n}=\mathbb{E}[Y]=\nabla f(\theta)
\]
- If we know how to compute \(f(\theta)\), we can get expectations / marginal probabilities by autodiff! Recovers forward-backward algorithms as special case. For a proof, see e.g. this paper.

\section*{Outline}

\section*{1 Numerical differentiation}

\section*{2 Chain compositions}

3 Computational graphs

4 Implementation

\section*{5 Advanced topics}

6 Conclusion
- Automatic differentiation is one of the keys that enabled the deep learnnig "revolution".

■ Backward / reverse differentiation is more efficient when the function has more inputs than outputs.
- Which is the de-facto setting in machine learning!

■ Even if you use Tensorflow / JAX / PyTorch, implementing a rudimentary autodiff library is a very good exercise.

\section*{References}

The following tutorials have been a great inspiration:
- Automatic Differentiation, Matthew Johnson, Deep Learning Summer School Montreal, 2017.

■ Differential programming, Gabriel Peyré, Mathematical Coffees, 2018.

Two minimalist implementations of autodiff:
- Autodidact, by Matthew Johnson.

■ Micrograd, by Andrej Karpathy.```

