# Introduction to Machine Learning (67577) 

## Shai Shalev-Shwartz

School of CS and Engineering, The Hebrew University of Jerusalem

Deep Learning

## Outline

(1) Gradient-Based Learning
(2) Computation Graph and Backpropagation
(3) Expressiveness and Sample Complexity
(4) Computational Complexity
(5) Convolutional Networks
(6) Solving MNIST with LeNet using Tensorflow
(7) Tips and Tricks

## Gradient-Based Learning

- Consider a hypothesis class which is parameterized by a vector $\theta \in \mathbb{R}^{d}$


## Gradient-Based Learning

- Consider a hypothesis class which is parameterized by a vector $\theta \in \mathbb{R}^{d}$
- Loss function of $h_{\theta}$ on example $(x, y)$ is denoted $\ell(\theta ;(x, y))$


## Gradient-Based Learning

- Consider a hypothesis class which is parameterized by a vector $\theta \in \mathbb{R}^{d}$
- Loss function of $h_{\theta}$ on example $(x, y)$ is denoted $\ell(\theta ;(x, y))$
- The true and empirical risks are

$$
L_{\mathcal{D}}(\theta)=\underset{(x, y) \sim \mathcal{D}}{\mathbb{E}}[\ell(\theta ;(x, y))] \quad, \quad L_{S}(\theta)=\frac{1}{m} \sum_{i=1}^{m} \ell\left(\theta ;\left(x_{i}, y_{i}\right)\right)
$$

## Gradient-Based Learning

- Consider a hypothesis class which is parameterized by a vector $\theta \in \mathbb{R}^{d}$
- Loss function of $h_{\theta}$ on example $(x, y)$ is denoted $\ell(\theta ;(x, y))$
- The true and empirical risks are

$$
L_{\mathcal{D}}(\theta)=\underset{(x, y) \sim \mathcal{D}}{\mathbb{E}}[\ell(\theta ;(x, y))] \quad, \quad L_{S}(\theta)=\frac{1}{m} \sum_{i=1}^{m} \ell\left(\theta ;\left(x_{i}, y_{i}\right)\right)
$$

- Assumption: $\ell$ is differentiable w.r.t. $\theta$ and we can calculate $\nabla \ell(\theta ;(x, y))$ efficiently


## Gradient-Based Learning

- Consider a hypothesis class which is parameterized by a vector $\theta \in \mathbb{R}^{d}$
- Loss function of $h_{\theta}$ on example $(x, y)$ is denoted $\ell(\theta ;(x, y))$
- The true and empirical risks are

$$
L_{\mathcal{D}}(\theta)=\underset{(x, y) \sim \mathcal{D}}{\mathbb{E}}[\ell(\theta ;(x, y))] \quad, \quad L_{S}(\theta)=\frac{1}{m} \sum_{i=1}^{m} \ell\left(\theta ;\left(x_{i}, y_{i}\right)\right)
$$

- Assumption: $\ell$ is differentiable w.r.t. $\theta$ and we can calculate $\nabla \ell(\theta ;(x, y))$ efficiently
- Minimize $L_{\mathcal{D}}$ or $L_{S}$ with Stochastic Gradient Descent (SGD): Start with $\theta^{(0)}$ and update $\theta^{(t+1)}=\theta^{(t)}-\eta_{t} \nabla \ell\left(\theta^{(t)} ;(x, y)\right)$


## Gradient-Based Learning

- Consider a hypothesis class which is parameterized by a vector $\theta \in \mathbb{R}^{d}$
- Loss function of $h_{\theta}$ on example $(x, y)$ is denoted $\ell(\theta ;(x, y))$
- The true and empirical risks are

$$
L_{\mathcal{D}}(\theta)=\underset{(x, y) \sim \mathcal{D}}{\mathbb{E}}[\ell(\theta ;(x, y))] \quad, \quad L_{S}(\theta)=\frac{1}{m} \sum_{i=1}^{m} \ell\left(\theta ;\left(x_{i}, y_{i}\right)\right)
$$

- Assumption: $\ell$ is differentiable w.r.t. $\theta$ and we can calculate $\nabla \ell(\theta ;(x, y))$ efficiently
- Minimize $L_{\mathcal{D}}$ or $L_{S}$ with Stochastic Gradient Descent (SGD): Start with $\theta^{(0)}$ and update $\theta^{(t+1)}=\theta^{(t)}-\eta_{t} \nabla \ell\left(\theta^{(t)} ;(x, y)\right)$
- SGD converges for convex problems. It may work for non-convex problems if we initialize "close enough" to a "good minimum"


## Outline

## (1) Gradient-Based Learning

(2) Computation Graph and Backpropagation
(3) Expressiveness and Sample Complexity
(4) Computational Complexity
(5) Convolutional Networks
(6) Solving MNIST with LeNet using Tensorflow
(7) Tips and Tricks

## Computation Graph

## A computation graph for a one dimensional Least Squares

(numbering of nodes corresponds to topological sort):


## Gradient Calculation using the Chain Rule

- Fix $x, y$ and write $\ell$ as a function of $w$ by

$$
\ell(w)=s\left(r_{y}\left(p_{x}(w)\right)\right)=\left(s \circ r_{y} \circ p_{x}\right)(w) .
$$

## Gradient Calculation using the Chain Rule

- Fix $x, y$ and write $\ell$ as a function of $w$ by

$$
\ell(w)=s\left(r_{y}\left(p_{x}(w)\right)\right)=\left(s \circ r_{y} \circ p_{x}\right)(w) .
$$

- Chain rule:

$$
\begin{aligned}
\ell^{\prime}(w) & =\left(s \circ r_{y} \circ p_{x}\right)^{\prime}(w) \\
& =s^{\prime}\left(r_{y}\left(p_{x}(w)\right)\right) \cdot\left(r_{y} \circ p_{x}\right)^{\prime}(w) \\
& =s^{\prime}\left(r_{y}\left(p_{x}(w)\right)\right) \cdot r_{y}^{\prime}\left(p_{x}(w)\right) \cdot p_{x}^{\prime}(w)
\end{aligned}
$$

## Gradient Calculation using the Chain Rule

- Fix $x, y$ and write $\ell$ as a function of $w$ by

$$
\ell(w)=s\left(r_{y}\left(p_{x}(w)\right)\right)=\left(s \circ r_{y} \circ p_{x}\right)(w) .
$$

- Chain rule:

$$
\begin{aligned}
\ell^{\prime}(w) & =\left(s \circ r_{y} \circ p_{x}\right)^{\prime}(w) \\
& =s^{\prime}\left(r_{y}\left(p_{x}(w)\right)\right) \cdot\left(r_{y} \circ p_{x}\right)^{\prime}(w) \\
& =s^{\prime}\left(r_{y}\left(p_{x}(w)\right)\right) \cdot r_{y}^{\prime}\left(p_{x}(w)\right) \cdot p_{x}^{\prime}(w)
\end{aligned}
$$

- Backpropagation: Calculate by a Forward-Backward pass over the graph


## Computation Graph - Forward

- For $t=0,1, \ldots, T-1$
- Layer[t]->output = Layer[t]->function(Layer[t]->inputs)



## Computation Graph — Backward

- Recall: $\ell^{\prime}(w)=s^{\prime}\left(r_{y}\left(p_{x}(w)\right)\right) \cdot r_{y}^{\prime}\left(p_{x}(w)\right) \cdot p_{x}^{\prime}(w)$
- Layer[T-1]->delta = 1
- For $t=T-1, T-2, \ldots, 0$
- For i in Layer [t]->inputs:
- i->delta $=$ Layer[t]->delta * Layer[t]->derivative(i, Layer[t]->inputs)



## Layers

- Nodes in the computation graph are often called layers


## Layers

- Nodes in the computation graph are often called layers
- Each layer is a simple differentiable function


## Layers

- Nodes in the computation graph are often called layers
- Each layer is a simple differentiable function
- Layers can implement multivariate functions


## Layers

- Nodes in the computation graph are often called layers
- Each layer is a simple differentiable function
- Layers can implement multivariate functions
- Example of popular layers:


## Layers

- Nodes in the computation graph are often called layers
- Each layer is a simple differentiable function
- Layers can implement multivariate functions
- Example of popular layers:
- Affine layer: $O=W X+b 1^{\top}$ where $W \in \mathbb{R}^{m, n}, x \in \mathbb{R}^{n, c}, b \in \mathbb{R}^{m}$


## Layers

- Nodes in the computation graph are often called layers
- Each layer is a simple differentiable function
- Layers can implement multivariate functions
- Example of popular layers:
- Affine layer: $O=W X+b 1^{\top}$ where $W \in \mathbb{R}^{m, n}, x \in \mathbb{R}^{n, c}, b \in \mathbb{R}^{m}$
- Unary layer: $\forall i, o_{i}=f\left(x_{i}\right)$ for some $f: \mathbb{R} \rightarrow \mathbb{R}$ e.g.


## Layers

- Nodes in the computation graph are often called layers
- Each layer is a simple differentiable function
- Layers can implement multivariate functions
- Example of popular layers:
- Affine layer: $O=W X+b 1^{\top}$ where $W \in \mathbb{R}^{m, n}, x \in \mathbb{R}^{n, c}, b \in \mathbb{R}^{m}$
- Unary layer: $\forall i, o_{i}=f\left(x_{i}\right)$ for some $f: \mathbb{R} \rightarrow \mathbb{R}$ e.g.
- Sigmoid: $f(x)=(1+\exp (-x))^{-1}$


## Layers

- Nodes in the computation graph are often called layers
- Each layer is a simple differentiable function
- Layers can implement multivariate functions
- Example of popular layers:
- Affine layer: $O=W X+b 1^{\top}$ where $W \in \mathbb{R}^{m, n}, x \in \mathbb{R}^{n, c}, b \in \mathbb{R}^{m}$
- Unary layer: $\forall i, o_{i}=f\left(x_{i}\right)$ for some $f: \mathbb{R} \rightarrow \mathbb{R}$ e.g.
- Sigmoid: $f(x)=(1+\exp (-x))^{-1}$
- Rectified Linear Unit (ReLU): $f(x)=\max \{0, x\}$ (discuss: derivative?)


## Layers

- Nodes in the computation graph are often called layers
- Each layer is a simple differentiable function
- Layers can implement multivariate functions
- Example of popular layers:
- Affine layer: $O=W X+b 1^{\top}$ where $W \in \mathbb{R}^{m, n}, x \in \mathbb{R}^{n, c}, b \in \mathbb{R}^{m}$
- Unary layer: $\forall i, o_{i}=f\left(x_{i}\right)$ for some $f: \mathbb{R} \rightarrow \mathbb{R}$ e.g.
- Sigmoid: $f(x)=(1+\exp (-x))^{-1}$
- Rectified Linear Unit (ReLU): $f(x)=\max \{0, x\}$ (discuss: derivative?)
- Binary layer: $\forall i, o_{i}=f\left(x_{i}, y_{i}\right)$ for some $f: \mathbb{R}^{2} \rightarrow \mathbb{R}$ e.g.


## Layers

- Nodes in the computation graph are often called layers
- Each layer is a simple differentiable function
- Layers can implement multivariate functions
- Example of popular layers:
- Affine layer: $O=W X+b 1^{\top}$ where $W \in \mathbb{R}^{m, n}, x \in \mathbb{R}^{n, c}, b \in \mathbb{R}^{m}$
- Unary layer: $\forall i, o_{i}=f\left(x_{i}\right)$ for some $f: \mathbb{R} \rightarrow \mathbb{R}$ e.g.
- Sigmoid: $f(x)=(1+\exp (-x))^{-1}$
- Rectified Linear Unit (ReLU): $f(x)=\max \{0, x\}$ (discuss: derivative?)
- Binary layer: $\forall i, o_{i}=f\left(x_{i}, y_{i}\right)$ for some $f: \mathbb{R}^{2} \rightarrow \mathbb{R}$ e.g.
- Add layer: $f(x, y)=x+y$


## Layers

- Nodes in the computation graph are often called layers
- Each layer is a simple differentiable function
- Layers can implement multivariate functions
- Example of popular layers:
- Affine layer: $O=W X+b 1^{\top}$ where $W \in \mathbb{R}^{m, n}, x \in \mathbb{R}^{n, c}, b \in \mathbb{R}^{m}$
- Unary layer: $\forall i, o_{i}=f\left(x_{i}\right)$ for some $f: \mathbb{R} \rightarrow \mathbb{R}$ e.g.
- Sigmoid: $f(x)=(1+\exp (-x))^{-1}$
- Rectified Linear Unit (ReLU): $f(x)=\max \{0, x\}$ (discuss: derivative?)
- Binary layer: $\forall i, o_{i}=f\left(x_{i}, y_{i}\right)$ for some $f: \mathbb{R}^{2} \rightarrow \mathbb{R}$ e.g.
- Add layer: $f(x, y)=x+y$
- Hinge loss: $f(x, y)=\left[1-y_{i} x_{i}\right]_{+}$


## Layers

- Nodes in the computation graph are often called layers
- Each layer is a simple differentiable function
- Layers can implement multivariate functions
- Example of popular layers:
- Affine layer: $O=W X+b 1^{\top}$ where $W \in \mathbb{R}^{m, n}, x \in \mathbb{R}^{n, c}, b \in \mathbb{R}^{m}$
- Unary layer: $\forall i, o_{i}=f\left(x_{i}\right)$ for some $f: \mathbb{R} \rightarrow \mathbb{R}$ e.g.
- Sigmoid: $f(x)=(1+\exp (-x))^{-1}$
- Rectified Linear Unit (ReLU): $f(x)=\max \{0, x\}$ (discuss: derivative?)
- Binary layer: $\forall i, o_{i}=f\left(x_{i}, y_{i}\right)$ for some $f: \mathbb{R}^{2} \rightarrow \mathbb{R}$ e.g.
- Add layer: $f(x, y)=x+y$
- Hinge loss: $f(x, y)=\left[1-y_{i} x_{i}\right]_{+}$
- Logistic loss: $f(x, y)=\log \left(1+\exp \left(-y_{i} x_{i}\right)\right)$


## Layers

- Nodes in the computation graph are often called layers
- Each layer is a simple differentiable function
- Layers can implement multivariate functions
- Example of popular layers:
- Affine layer: $O=W X+b 1^{\top}$ where $W \in \mathbb{R}^{m, n}, x \in \mathbb{R}^{n, c}, b \in \mathbb{R}^{m}$
- Unary layer: $\forall i, o_{i}=f\left(x_{i}\right)$ for some $f: \mathbb{R} \rightarrow \mathbb{R}$ e.g.
- Sigmoid: $f(x)=(1+\exp (-x))^{-1}$
- Rectified Linear Unit (ReLU): $f(x)=\max \{0, x\}$ (discuss: derivative?)
- Binary layer: $\forall i, o_{i}=f\left(x_{i}, y_{i}\right)$ for some $f: \mathbb{R}^{2} \rightarrow \mathbb{R}$ e.g.
- Add layer: $f(x, y)=x+y$
- Hinge loss: $f(x, y)=\left[1-y_{i} x_{i}\right]_{+}$
- Logistic loss: $f(x, y)=\log \left(1+\exp \left(-y_{i} x_{i}\right)\right)$


## Layers

- Nodes in the computation graph are often called layers
- Each layer is a simple differentiable function
- Layers can implement multivariate functions
- Example of popular layers:
- Affine layer: $O=W X+b 1^{\top}$ where $W \in \mathbb{R}^{m, n}, x \in \mathbb{R}^{n, c}, b \in \mathbb{R}^{m}$
- Unary layer: $\forall i, o_{i}=f\left(x_{i}\right)$ for some $f: \mathbb{R} \rightarrow \mathbb{R}$ e.g.
- Sigmoid: $f(x)=(1+\exp (-x))^{-1}$
- Rectified Linear Unit (ReLU): $f(x)=\max \{0, x\}$ (discuss: derivative?)
- Binary layer: $\forall i, o_{i}=f\left(x_{i}, y_{i}\right)$ for some $f: \mathbb{R}^{2} \rightarrow \mathbb{R}$ e.g.
- Add layer: $f(x, y)=x+y$
- Hinge loss: $f(x, y)=\left[1-y_{i} x_{i}\right]_{+}$
- Logistic loss: $f(x, y)=\log \left(1+\exp \left(-y_{i} x_{i}\right)\right)$


## Main message

Computation graph enables us to construct very complicated functions from simple building blocks

## Backpropgation for multivariate layers

- Recall the backpropagation rule:
- For i in Layer[t]->inputs:
- i->delta $=$ Layer[t]->delta *

Layer[t]->derivative(i, Layer[t]->inputs)

## Backpropgation for multivariate layers

- Recall the backpropagation rule:
- For i in Layer[t]->inputs:
- i->delta $=$ Layer[t]->delta *

Layer[t]->derivative(i, Layer [t]->inputs)

- "delta" is now a vector (same dimension as the output of the layer)


## Backpropgation for multivariate layers

- Recall the backpropagation rule:
- For i in Layer[t]->inputs:
- i->delta $=$ Layer[t]->delta *

Layer[t]->derivative(i, Layer[t]->inputs)

- "delta" is now a vector (same dimension as the output of the layer)
- "derivative" is the Jacobian matrix:

The Jacobian of $\mathbf{f}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}$ at $\mathbf{x} \in \mathbb{R}^{n}$, denoted $J_{\mathbf{x}}(\mathbf{f})$, is the $m \times n$ matrix whose $i, j$ element is the partial derivative of $f_{i}: \mathbb{R}^{n} \rightarrow \mathbb{R}$ w.r.t. its $j$ 'th variable at $\mathbf{x}$.

## Backpropgation for multivariate layers

- Recall the backpropagation rule:
- For i in Layer[t]->inputs:
- i->delta $=$ Layer[t]->delta *

Layer[t]->derivative(i, Layer[t]->inputs)

- "delta" is now a vector (same dimension as the output of the layer)
- "derivative" is the Jacobian matrix:

The Jacobian of $\mathbf{f}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}$ at $\mathbf{x} \in \mathbb{R}^{n}$, denoted $J_{\mathbf{x}}(\mathbf{f})$, is the $m \times n$ matrix whose $i, j$ element is the partial derivative of $f_{i}: \mathbb{R}^{n} \rightarrow \mathbb{R}$ w.r.t. its $j$ 'th variable at $\mathbf{x}$.

- The multiplication is matrix multiplication


## Backpropgation for multivariate layers

- Recall the backpropagation rule:
- For i in Layer[t]->inputs:

$$
\begin{aligned}
& \text { - i->delta = Layer[t]->delta } * \\
& \text { Layer[t]->derivative(i, Layer[t]->inputs) }
\end{aligned}
$$

- "delta" is now a vector (same dimension as the output of the layer)
- "derivative" is the Jacobian matrix:

The Jacobian of $\mathbf{f}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}$ at $\mathbf{x} \in \mathbb{R}^{n}$, denoted $J_{\mathbf{x}}(\mathbf{f})$, is the $m \times n$ matrix whose $i, j$ element is the partial derivative of $f_{i}: \mathbb{R}^{n} \rightarrow \mathbb{R}$ w.r.t. its $j$ 'th variable at $\mathbf{x}$.

- The multiplication is matrix multiplication
- The correctness of the algorithm follows from the multivariate chain rule

$$
J_{\mathbf{w}}(\mathbf{f} \circ \mathbf{g})=J_{g(\mathbf{w})}(\mathbf{f}) J_{\mathbf{w}}(\mathbf{g})
$$

## Jacobian - Examples

- If $\mathbf{f}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{n}$ is element-wise application of $\sigma: \mathbb{R} \rightarrow \mathbb{R}$ then $J_{\mathbf{x}}(\mathbf{f})=\operatorname{diag}\left(\left(\sigma^{\prime}\left(x_{1}\right), \ldots, \sigma^{\prime}\left(x_{n}\right)\right)\right)$.


## Jacobian - Examples

- If $\mathbf{f}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{n}$ is element-wise application of $\sigma: \mathbb{R} \rightarrow \mathbb{R}$ then $J_{\mathbf{x}}(\mathbf{f})=\operatorname{diag}\left(\left(\sigma^{\prime}\left(x_{1}\right), \ldots, \sigma^{\prime}\left(x_{n}\right)\right)\right)$.
- Let $\mathbf{f}(\mathbf{x}, \mathbf{w}, b)=\mathbf{w}^{\top} \mathbf{x}+b$ for $\mathbf{w}, \mathbf{x} \in \mathbb{R}^{n}, b \in \mathbb{R}^{1}$. Then:

$$
J_{\mathbf{x}}(\mathbf{f})=\mathbf{w}^{\top} \quad, \quad J_{\mathbf{w}}(\mathbf{f})=\mathbf{x}^{\top} \quad, \quad J_{b}(\mathbf{f})=1
$$

## Jacobian - Examples

- If $\mathbf{f}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{n}$ is element-wise application of $\sigma: \mathbb{R} \rightarrow \mathbb{R}$ then $J_{\mathbf{x}}(\mathbf{f})=\operatorname{diag}\left(\left(\sigma^{\prime}\left(x_{1}\right), \ldots, \sigma^{\prime}\left(x_{n}\right)\right)\right)$.
- Let $\mathbf{f}(\mathbf{x}, \mathbf{w}, b)=\mathbf{w}^{\top} \mathbf{x}+b$ for $\mathbf{w}, \mathbf{x} \in \mathbb{R}^{n}, b \in \mathbb{R}^{1}$. Then:

$$
J_{\mathbf{x}}(\mathbf{f})=\mathbf{w}^{\top} \quad, \quad J_{\mathbf{w}}(\mathbf{f})=\mathbf{x}^{\top} \quad, \quad J_{b}(\mathbf{f})=1
$$

- Let $\mathbf{f}(W, \mathbf{x})=W \mathbf{x}$. Then:

$$
J_{\mathbf{x}}(\mathbf{f})=W \quad, \quad J_{W}(\mathbf{f})=\left(\begin{array}{cccc}
\mathbf{x}^{\top} & 0 & \cdots & 0 \\
0 & \mathbf{x}^{\top} & \cdots & 0 \\
\vdots & \vdots & \ddots & \vdots \\
0 & 0 & \cdots & \mathbf{x}^{\top}
\end{array}\right)
$$

## Outline

## (1) Gradient-Based Learning

(C) Computation Graph and Backpropagation
(3) Expressiveness and Sample Complexity
(4) Computational Complexity
(5) Convolutional Networks
© Solving MNIST with LeNet using Tensorflow
(7) Tips and Tricks

## Sample Complexity

- If we learn $d$ parameters, and each one is stored in, say, 32 bits, then the number of hypotheses in our class is at most $2^{32 d}$. It follows that the sample complexity is order of $d$.


## Sample Complexity

- If we learn $d$ parameters, and each one is stored in, say, 32 bits, then the number of hypotheses in our class is at most $2^{32 d}$. It follows that the sample complexity is order of $d$.
- Other ways to improve generalization is all sort of regularization


## Expressiveness

- So far in the course we considered hypotheses of the form $x \mapsto w^{\top} x+b$
- Now, consider the following computation graph, known as "one hidden layer network":



## Expressiveness of "One Hidden Layer Network"

- Claim: Every Boolean function $f:\{ \pm 1\}^{n} \rightarrow\{ \pm 1\}$ can be expressed by a one hidden layer network.


## Expressiveness of "One Hidden Layer Network"

- Claim: Every Boolean function $f:\{ \pm 1\}^{n} \rightarrow\{ \pm 1\}$ can be expressed by a one hidden layer network.
- Proof:


## Expressiveness of "One Hidden Layer Network"

- Claim: Every Boolean function $f:\{ \pm 1\}^{n} \rightarrow\{ \pm 1\}$ can be expressed by a one hidden layer network.
- Proof:
- Show that for integer $x$ we have $\operatorname{sign}(x)=2\left([x+1]_{+}-[x]_{+}\right)-1$


## Expressiveness of "One Hidden Layer Network"

- Claim: Every Boolean function $f:\{ \pm 1\}^{n} \rightarrow\{ \pm 1\}$ can be expressed by a one hidden layer network.
- Proof:
- Show that for integer $x$ we have $\operatorname{sign}(x)=2\left([x+1]_{+}-[x]_{+}\right)-1$
- Show that any $f$ can be written as $f(x)=\vee_{i}\left(x==u_{i}\right)$ for some vectors $u_{1}, \ldots, u_{k}$


## Expressiveness of "One Hidden Layer Network"

- Claim: Every Boolean function $f:\{ \pm 1\}^{n} \rightarrow\{ \pm 1\}$ can be expressed by a one hidden layer network.
- Proof:
- Show that for integer $x$ we have $\operatorname{sign}(x)=2\left([x+1]_{+}-[x]_{+}\right)-1$
- Show that any $f$ can be written as $f(x)=\vee_{i}\left(x==u_{i}\right)$ for some vectors $u_{1}, \ldots, u_{k}$
- Show that $\operatorname{sign}\left(x^{\top} u_{i}-(n-1)\right)$ is an indicator to $\left(x==u_{i}\right)$


## Expressiveness of "One Hidden Layer Network"

- Claim: Every Boolean function $f:\{ \pm 1\}^{n} \rightarrow\{ \pm 1\}$ can be expressed by a one hidden layer network.
- Proof:
- Show that for integer $x$ we have $\operatorname{sign}(x)=2\left([x+1]_{+}-[x]_{+}\right)-1$
- Show that any $f$ can be written as $f(x)=\vee_{i}\left(x==u_{i}\right)$ for some vectors $u_{1}, \ldots, u_{k}$
- Show that $\operatorname{sign}\left(x^{\top} u_{i}-(n-1)\right)$ is an indicator to $\left(x==u_{i}\right)$
- Conclude that we can adjust the weights so that $y p(x) \geq 1$ for all examples ( $x, y$ )


## Expressiveness of "One Hidden Layer Network"

- Claim: Every Boolean function $f:\{ \pm 1\}^{n} \rightarrow\{ \pm 1\}$ can be expressed by a one hidden layer network.
- Proof:
- Show that for integer $x$ we have $\operatorname{sign}(x)=2\left([x+1]_{+}-[x]_{+}\right)-1$
- Show that any $f$ can be written as $f(x)=\vee_{i}\left(x==u_{i}\right)$ for some vectors $u_{1}, \ldots, u_{k}$
- Show that $\operatorname{sign}\left(x^{\top} u_{i}-(n-1)\right)$ is an indicator to $\left(x==u_{i}\right)$
- Conclude that we can adjust the weights so that $y p(x) \geq 1$ for all examples ( $x, y$ )
- Theorem: For every $n$, let $s(n)$ be the minimal integer such that there exists a one hidden layer network with $s(n)$ hidden neurons that implements all functions from $\{0,1\}^{n}$ to $\{0,1\}$. Then, $s(n)$ is exponential in $n$.


## Expressiveness of "One Hidden Layer Network"

- Claim: Every Boolean function $f:\{ \pm 1\}^{n} \rightarrow\{ \pm 1\}$ can be expressed by a one hidden layer network.
- Proof:
- Show that for integer $x$ we have $\operatorname{sign}(x)=2\left([x+1]_{+}-[x]_{+}\right)-1$
- Show that any $f$ can be written as $f(x)=\vee_{i}\left(x==u_{i}\right)$ for some vectors $u_{1}, \ldots, u_{k}$
- Show that $\operatorname{sign}\left(x^{\top} u_{i}-(n-1)\right)$ is an indicator to $\left(x==u_{i}\right)$
- Conclude that we can adjust the weights so that $y p(x) \geq 1$ for all examples ( $x, y$ )
- Theorem: For every $n$, let $s(n)$ be the minimal integer such that there exists a one hidden layer network with $s(n)$ hidden neurons that implements all functions from $\{0,1\}^{n}$ to $\{0,1\}$. Then, $s(n)$ is exponential in $n$.
- Proof: Think on the VC dimension ...


## Expressiveness of "One Hidden Layer Network"

- Claim: Every Boolean function $f:\{ \pm 1\}^{n} \rightarrow\{ \pm 1\}$ can be expressed by a one hidden layer network.
- Proof:
- Show that for integer $x$ we have $\operatorname{sign}(x)=2\left([x+1]_{+}-[x]_{+}\right)-1$
- Show that any $f$ can be written as $f(x)=\vee_{i}\left(x==u_{i}\right)$ for some vectors $u_{1}, \ldots, u_{k}$
- Show that $\operatorname{sign}\left(x^{\top} u_{i}-(n-1)\right)$ is an indicator to $\left(x==u_{i}\right)$
- Conclude that we can adjust the weights so that $y p(x) \geq 1$ for all examples ( $x, y$ )
- Theorem: For every $n$, let $s(n)$ be the minimal integer such that there exists a one hidden layer network with $s(n)$ hidden neurons that implements all functions from $\{0,1\}^{n}$ to $\{0,1\}$. Then, $s(n)$ is exponential in $n$.
- Proof: Think on the VC dimension ...
- What type of functions can be implemented by small size networks?


## Geometric Intuition

- One hidden layer networks can express intersection of halfspaces



## Geometric Intuition

- Two hidden layer networks can express unions of intersection of halfspaces



## What can we express with $T$-depth networks?

- Theorem: Let $T: \mathbb{N} \rightarrow \mathbb{N}$ and for every $n$, let $\mathcal{F}_{n}$ be the set of functions that can be implemented using a Turing machine using runtime of at most $T(n)$. Then, there exist constants $b, c \in \mathbb{R}_{+}$such that for every $n$, there is a network of depth at most $T$ and size at most $c T(n)^{2}+b$ such that it implements all functions in $\mathcal{F}_{n}$.


## What can we express with $T$-depth networks?

- Theorem: Let $T: \mathbb{N} \rightarrow \mathbb{N}$ and for every $n$, let $\mathcal{F}_{n}$ be the set of functions that can be implemented using a Turing machine using runtime of at most $T(n)$. Then, there exist constants $b, c \in \mathbb{R}_{+}$such that for every $n$, there is a network of depth at most $T$ and size at most $c T(n)^{2}+b$ such that it implements all functions in $\mathcal{F}_{n}$.
- Sample complexity is order of number of variables (in our case polynomial in $T$ )


## What can we express with $T$-depth networks?

- Theorem: Let $T: \mathbb{N} \rightarrow \mathbb{N}$ and for every $n$, let $\mathcal{F}_{n}$ be the set of functions that can be implemented using a Turing machine using runtime of at most $T(n)$. Then, there exist constants $b, c \in \mathbb{R}_{+}$such that for every $n$, there is a network of depth at most $T$ and size at most $c T(n)^{2}+b$ such that it implements all functions in $\mathcal{F}_{n}$.
- Sample complexity is order of number of variables (in our case polynomial in $T$ )
- Conclusion: A very weak notion of prior knowledge suffices - if we only care about functions that can be implemented in time $T(n)$, we can use neural networks of depth $T$ and size $O\left(T(n)^{2}\right)$, and the sample complexity is also bounded by polynomial in $T(n)$ !


## The ultimate hypothesis class



## Outline

## (1) Gradient-Based Learning

(2) Computation Graph and Backpropagation
(3) Expressiveness and Sample Complexity

4 Computational Complexity
(5) Convolutional Networks
(6) Solving MNIST with LeNet using Tensorflow
(7) Tips and Tricks

## Runtime of learning neural networks

- Theorem: It is NP hard to implement the ERM rule even for one hidden layer networks with just 4 neurons in the hidden layer.


## Runtime of learning neural networks

- Theorem: It is NP hard to implement the ERM rule even for one hidden layer networks with just 4 neurons in the hidden layer.
- But, maybe ERM is hard but some improper algorithm works ?


## Runtime of learning neural networks

- Theorem: It is NP hard to implement the ERM rule even for one hidden layer networks with just 4 neurons in the hidden layer.
- But, maybe ERM is hard but some improper algorithm works ?
- Theorem: Under some average case complexity assumption, it is hard to learn one hidden layer networks with $\omega(\log (d))$ hidden neurons even improperly



## How to train neural network ?

- So, neural networks can form an excellent hypothesis class, but it is intractable to train it.


## How to train neural network ?

- So, neural networks can form an excellent hypothesis class, but it is intractable to train it.
- How is this different than the class of all Python programs that can be implemented in code length of $b$ bits?


## How to train neural network ?

- So, neural networks can form an excellent hypothesis class, but it is intractable to train it.
- How is this different than the class of all Python programs that can be implemented in code length of $b$ bits?
- Main technique: Gradient-based learning (using SGD)


## How to train neural network ?

- So, neural networks can form an excellent hypothesis class, but it is intractable to train it.
- How is this different than the class of all Python programs that can be implemented in code length of $b$ bits ?
- Main technique: Gradient-based learning (using SGD)
- Not convex, no guarantees, can take a long time, but:


## How to train neural network ?

- So, neural networks can form an excellent hypothesis class, but it is intractable to train it.
- How is this different than the class of all Python programs that can be implemented in code length of $b$ bits ?
- Main technique: Gradient-based learning (using SGD)
- Not convex, no guarantees, can take a long time, but:
- Often (but not always) still works fine, finds a good solution


## How to train neural network ?

- So, neural networks can form an excellent hypothesis class, but it is intractable to train it.
- How is this different than the class of all Python programs that can be implemented in code length of $b$ bits?
- Main technique: Gradient-based learning (using SGD)
- Not convex, no guarantees, can take a long time, but:
- Often (but not always) still works fine, finds a good solution
- Easier than optimizing over Python programs ...


## How to train neural network ?

- So, neural networks can form an excellent hypothesis class, but it is intractable to train it.
- How is this different than the class of all Python programs that can be implemented in code length of $b$ bits ?
- Main technique: Gradient-based learning (using SGD)
- Not convex, no guarantees, can take a long time, but:
- Often (but not always) still works fine, finds a good solution
- Easier than optimizing over Python programs ...
- Need to apply some tricks (initialization, learning rate, mini-batching, architecture), and need some luck


## Outline

## (1) Gradient-Based Learning

(2) Computation Graph and Backpropagation
(3) Expressiveness and Sample Complexity

4 Computational Complexity
(5) Convolutional Networks
(6) Solving MNIST with LeNet using Tensorflow
(7) Tips and Tricks

## Deep Learning Golden age in Vision

- 20I2-2014 Imagenet results:

|  | CNN |
| :--- | :---: |
|  | non-CNN |
| 2014 Teams | \%error |
| GoogLeNet | 6.6 |
| VGG (Oxford) | 7.3 |
| MSRA | 8.0 |
| A. Howard | 8.1 |
| DeeperVision | 9.5 |
| NUS-BST | 9.7 |
| TTIC-ECP | 10.2 |
| XYZ | 11.2 |
| UvA | 12.1 |

- 2015 results: MSRA under $\mathbf{3 . 5 \%}$ error. (using a CNN with 150 layers!)
figures from Yann LeCun's CVPR'I 5 plenary


## Convolution Layer

- Input: $C$ images


## Convolution Layer

- Input: $C$ images
- Output: $C^{\prime}$ images


## Convolution Layer

- Input: $C$ images
- Output: $C^{\prime}$ images
- Calculation:

$$
O\left[c^{\prime}, h^{\prime}, w^{\prime}\right]=b^{\left(c^{\prime}\right)}+\sum_{c=0}^{C-1} \sum_{h=0}^{k-1} \sum_{w=0}^{k-1} W^{\left(c^{\prime}\right)}[c, h, w] X\left[c, h+h^{\prime}, w+w^{\prime}\right]
$$

## Convolution Layer

- Input: $C$ images
- Output: $C^{\prime}$ images
- Calculation:

$$
O\left[c^{\prime}, h^{\prime}, w^{\prime}\right]=b^{\left(c^{\prime}\right)}+\sum_{c=0}^{C-1} \sum_{h=0}^{k-1} \sum_{w=0}^{k-1} W^{\left(c^{\prime}\right)}[c, h, w] X\left[c, h+h^{\prime}, w+w^{\prime}\right]
$$

- Observe: equivalent to an Affine layer with weight sharing


## Convolution Layer

- Input: $C$ images
- Output: $C^{\prime}$ images
- Calculation:

$$
O\left[c^{\prime}, h^{\prime}, w^{\prime}\right]=b^{\left(c^{\prime}\right)}+\sum_{c=0}^{C-1} \sum_{h=0}^{k-1} \sum_{w=0}^{k-1} W^{\left(c^{\prime}\right)}[c, h, w] X\left[c, h+h^{\prime}, w+w^{\prime}\right]
$$

- Observe: equivalent to an Affine layer with weight sharing
- Observe: can be implemented as a combination of Im2Col layer and Affine layer


## Im2Col Layer

- Im2Col for $3 \times 3$ convolution

| 0 | 1 | 2 | 3 | 4 |
| :---: | :---: | :---: | :---: | :---: |
| 5 | 6 | 7 | 8 | 9 |
| 10 | 11 | 12 | 13 | 14 |
| 15 | 16 | 17 | 18 | 19 |
| 20 | 21 | 22 | 23 | 24 |


|  |  |  |  |  | 0 | 0 | 1 |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
|  |  |  |  | 0 | 1 | 2 |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  | 1 | 2 | 3 |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  | 0 | 5 | 6 |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  | 5 | 6 | 7 |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  | 6 | 7 | 8 |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  | 0 | 10 | 11 |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  | 10 | 11 | 12 |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  | 11 | 12 | 13 |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |

## Im2Col Layer

- Im2Col for $3 \times 3$ convolution with 2 input channels



## Parameters of Convolutions layer

- Kernel height and kernel width
- Stride height and stride width
- zero padding (True or False)
- Number of output channels


## Pooling Layer

- Input: Image of size $H \times W$


## Pooling Layer

- Input: Image of size $H \times W$
- Output: Image of size $(H / k) \times(W / k)$


## Pooling Layer

- Input: Image of size $H \times W$
- Output: Image of size $(H / k) \times(W / k)$
- Calculation: Divide input image to $k \times k$ windows and for each such window output the maximal value (or average value)


## Pooling Layer

- Input: Image of size $H \times W$
- Output: Image of size $(H / k) \times(W / k)$
- Calculation: Divide input image to $k \times k$ windows and for each such window output the maximal value (or average value)
- Observe: equivalent to $\operatorname{Im} 2 \mathrm{Col}+$ reduce operation


## Pooling Layer

- Input: Image of size $H \times W$
- Output: Image of size $(H / k) \times(W / k)$
- Calculation: Divide input image to $k \times k$ windows and for each such window output the maximal value (or average value)
- Observe: equivalent to $\operatorname{Im} 2 \mathrm{Col}+$ reduce operation
- Discuss: how to calculate derivative ?


## Outline

## (1) Gradient-Based Learning

(2) Computation Graph and Backpropagation
(3) Expressiveness and Sample Complexity

4 Computational Complexity
(5) Convolutional Networks
(6) Solving MNIST with LeNet using Tensorflow
(7) Tips and Tricks

## Example: LeNet for MNIST

- The task: Handwritten digits recognition


## Example: LeNet for MNIST

- The task: Handwritten digits recognition
- Input space: $\mathcal{X}=\{0,1, \ldots, 255\}^{28 \times 28}$


## Example: LeNet for MNIST

- The task: Handwritten digits recognition
- Input space: $\mathcal{X}=\{0,1, \ldots, 255\}^{28 \times 28}$
- Output space: $\mathcal{Y}=\{0,1, \ldots, 9\}$


## Example: LeNet for MNIST

- The task: Handwritten digits recognition
- Input space: $\mathcal{X}=\{0,1, \ldots, 255\}^{28 \times 28}$
- Output space: $\mathcal{Y}=\{0,1, \ldots, 9\}$
- Multiclass categorization:


## Example: LeNet for MNIST

- The task: Handwritten digits recognition
- Input space: $\mathcal{X}=\{0,1, \ldots, 255\}^{28 \times 28}$
- Output space: $\mathcal{Y}=\{0,1, \ldots, 9\}$
- Multiclass categorization:
- We take hypotheses of the form $h: \mathcal{X} \rightarrow \mathbb{R}^{|\mathcal{Y}|}$


## Example: LeNet for MNIST

- The task: Handwritten digits recognition
- Input space: $\mathcal{X}=\{0,1, \ldots, 255\}^{28 \times 28}$
- Output space: $\mathcal{Y}=\{0,1, \ldots, 9\}$
- Multiclass categorization:
- We take hypotheses of the form $h: \mathcal{X} \rightarrow \mathbb{R}^{|\mathcal{Y}|}$
- We interpret $h(x)$ as a vector that gives scores for all the labels


## Example: LeNet for MNIST

- The task: Handwritten digits recognition
- Input space: $\mathcal{X}=\{0,1, \ldots, 255\}^{28 \times 28}$
- Output space: $\mathcal{Y}=\{0,1, \ldots, 9\}$
- Multiclass categorization:
- We take hypotheses of the form $h: \mathcal{X} \rightarrow \mathbb{R}^{|\mathcal{Y}|}$
- We interpret $h(x)$ as a vector that gives scores for all the labels
- The actual prediction is the label with the highest score: $\operatorname{argmax}_{i} h_{i}(x)$


## Example: LeNet for MNIST

- The task: Handwritten digits recognition
- Input space: $\mathcal{X}=\{0,1, \ldots, 255\}^{28 \times 28}$
- Output space: $\mathcal{Y}=\{0,1, \ldots, 9\}$
- Multiclass categorization:
- We take hypotheses of the form $h: \mathcal{X} \rightarrow \mathbb{R}^{|\mathcal{Y}|}$
- We interpret $h(x)$ as a vector that gives scores for all the labels
- The actual prediction is the label with the highest score: $\operatorname{argmax}_{i} h_{i}(x)$
- Network architecture: $x \rightarrow \operatorname{Conv}(5 \times 5,1 \times 1$, no-pad,20) $\rightarrow$ Pool( $2 \times 2$ 2) $\rightarrow$ Conv(5×5, $1 \times 1$, no-pad,50) $\rightarrow \operatorname{Pool}(2 \times 2) \rightarrow$ Affine(500) $\rightarrow$ ReLU $\rightarrow$ Affine(10).


## Example: LeNet for MNIST

- The task: Handwritten digits recognition
- Input space: $\mathcal{X}=\{0,1, \ldots, 255\}^{28 \times 28}$
- Output space: $\mathcal{Y}=\{0,1, \ldots, 9\}$
- Multiclass categorization:
- We take hypotheses of the form $h: \mathcal{X} \rightarrow \mathbb{R}^{|\mathcal{Y}|}$
- We interpret $h(x)$ as a vector that gives scores for all the labels
- The actual prediction is the label with the highest score: $\operatorname{argmax}_{i} h_{i}(x)$
- Network architecture: $x \rightarrow \operatorname{Conv}(5 \times 5,1 \times 1$, no-pad,20) $\rightarrow \operatorname{Pool}(2 \times 2) \rightarrow$ Conv(5×5, $1 \times 1$, no-pad,50) $\rightarrow \operatorname{Pool}(2 \times 2) \rightarrow$ Affine(500) $\rightarrow$ ReLU $\rightarrow$ Affine(10).
- Logistic loss for multiclass categorization:


## Example: LeNet for MNIST

- The task: Handwritten digits recognition
- Input space: $\mathcal{X}=\{0,1, \ldots, 255\}^{28 \times 28}$
- Output space: $\mathcal{Y}=\{0,1, \ldots, 9\}$
- Multiclass categorization:
- We take hypotheses of the form $h: \mathcal{X} \rightarrow \mathbb{R}^{|\mathcal{Y}|}$
- We interpret $h(x)$ as a vector that gives scores for all the labels
- The actual prediction is the label with the highest score: $\operatorname{argmax}_{i} h_{i}(x)$
- Network architecture: $x \rightarrow \operatorname{Conv}(5 \times 5,1 \times 1$, no-pad,20) $\rightarrow$ Pool( $2 \times 2$ 2) $\rightarrow$ Conv(5×5, $1 \times 1$, no-pad,50) $\rightarrow \operatorname{Pool}(2 \times 2) \rightarrow$ Affine(500) $\rightarrow$ ReLU $\rightarrow$ Affine(10).
- Logistic loss for multiclass categorization:
- SoftMax: $\forall i, \quad p_{i}=\frac{\exp \left(h_{i}(x)\right)}{\sum_{j} \exp \left(h_{j}(x)\right)}$


## Example: LeNet for MNIST

- The task: Handwritten digits recognition
- Input space: $\mathcal{X}=\{0,1, \ldots, 255\}^{28 \times 28}$
- Output space: $\mathcal{Y}=\{0,1, \ldots, 9\}$
- Multiclass categorization:
- We take hypotheses of the form $h: \mathcal{X} \rightarrow \mathbb{R}^{|\mathcal{Y}|}$
- We interpret $h(x)$ as a vector that gives scores for all the labels
- The actual prediction is the label with the highest score: $\operatorname{argmax}_{i} h_{i}(x)$
- Network architecture: $x \rightarrow \operatorname{Conv}(5 \times 5,1 \times 1$,no-pad,20) $\rightarrow \operatorname{Pool}(2 \times 2) \rightarrow$ Conv(5×5, $1 \times 1$,no-pad,50) $\rightarrow \operatorname{Pool}(2 \times 2) \rightarrow$ Affine(500) $\rightarrow$ ReLU $\rightarrow$ Affine(10).
- Logistic loss for multiclass categorization:
- SoftMax: $\forall i, \quad p_{i}=\frac{\exp \left(h_{i}(x)\right)}{\sum_{j} \exp \left(h_{j}(x)\right)}$
- LogLoss: If the correct label is $y$ then the loss is

$$
-\log \left(p_{y}\right)=\log \left(\sum_{j} \exp \left(h_{j}(x)-h_{i}(x)\right)\right)
$$

## Outline

## (1) Gradient-Based Learning

(2) Computation Graph and Backpropagation
(3) Expressiveness and Sample Complexity
4. Computational Complexity
(5) Convolutional Networks
(6) Solving MNIST with LeNet using Tensorflow
(7) Tips and Tricks

## Reduction Layers

- The complexity of convolutional layers is $C_{\text {in }} \times C_{\text {out }} \times H \times W$
- A "reduction layer" is a $1 \times 1$ convolution aiming at reducing $C_{\text {in }}$
- It can greatly reduce the computational complexity (less time) and sample complexity (fewer parameters)


## Inception modules



Figure 2: Inception module

- Szegedy et al (Google)
- Won the ImageNet 2014 challenge ( $6.67 \%$ error)


## Residual Networks



Figure 2. Residual learning: a building block.

- He, Zhang, Ren, Sun (Microsoft)
- Won the ImageNet 2015 challenge with a 152 layers network (3.57\% error)


## Some Training Tricks

- Input normalization: divide each element of $x$ by 255 to make sure it is in $[0,1]$


## Some Training Tricks

- Input normalization: divide each element of $x$ by 255 to make sure it is in $[0,1]$
- Initialization is important: One trick that works well in practice is to initialize the bias to be zero and initialize the rows of $W$ to be random in $[-1 / \sqrt{n}, 1 / \sqrt{n}]$


## Some Training Tricks

- Input normalization: divide each element of $x$ by 255 to make sure it is in $[0,1]$
- Initialization is important: One trick that works well in practice is to initialize the bias to be zero and initialize the rows of $W$ to be random in $[-1 / \sqrt{n}, 1 / \sqrt{n}]$
- Mini-batches: At each iteration of SGD we calculate the average loss on $k$ random examples for $k>1$. Advantages:


## Some Training Tricks

- Input normalization: divide each element of $x$ by 255 to make sure it is in $[0,1]$
- Initialization is important: One trick that works well in practice is to initialize the bias to be zero and initialize the rows of $W$ to be random in $[-1 / \sqrt{n}, 1 / \sqrt{n}]$
- Mini-batches: At each iteration of SGD we calculate the average loss on $k$ random examples for $k>1$. Advantages:
- Reduces the variance of the update direction (w.r.t. the full gradient), hence converges faster


## Some Training Tricks

- Input normalization: divide each element of $x$ by 255 to make sure it is in $[0,1]$
- Initialization is important: One trick that works well in practice is to initialize the bias to be zero and initialize the rows of $W$ to be random in $[-1 / \sqrt{n}, 1 / \sqrt{n}]$
- Mini-batches: At each iteration of SGD we calculate the average loss on $k$ random examples for $k>1$. Advantages:
- Reduces the variance of the update direction (w.r.t. the full gradient), hence converges faster
- We don't pay a lot in time because of parallel implementation


## Some Training Tricks

- Input normalization: divide each element of $x$ by 255 to make sure it is in $[0,1]$
- Initialization is important: One trick that works well in practice is to initialize the bias to be zero and initialize the rows of $W$ to be random in $[-1 / \sqrt{n}, 1 / \sqrt{n}]$
- Mini-batches: At each iteration of SGD we calculate the average loss on $k$ random examples for $k>1$. Advantages:
- Reduces the variance of the update direction (w.r.t. the full gradient), hence converges faster
- We don't pay a lot in time because of parallel implementation
- Learning rate: Choice of learning rate is important. One way is to start with some fixed $\eta$ and decrease it by $1 / 2$ whenever the training stops making progress.


## Some Training Tricks

- Input normalization: divide each element of $x$ by 255 to make sure it is in $[0,1]$
- Initialization is important: One trick that works well in practice is to initialize the bias to be zero and initialize the rows of $W$ to be random in $[-1 / \sqrt{n}, 1 / \sqrt{n}]$
- Mini-batches: At each iteration of SGD we calculate the average loss on $k$ random examples for $k>1$. Advantages:
- Reduces the variance of the update direction (w.r.t. the full gradient), hence converges faster
- We don't pay a lot in time because of parallel implementation
- Learning rate: Choice of learning rate is important. One way is to start with some fixed $\eta$ and decrease it by $1 / 2$ whenever the training stops making progress.
- Variants of SGD: There are plenty of variants that work better than vanilla SGD.


## Failures of Deep Learning

- Parity of more than 30 bits
- Multiplication of large numbers
- Matrix inversion
- ...


## Summary

- Deep Learning can be used to construct the ultimate hypothesis class
- Worst-case complexity is exponential
- ... but, empirically, it works reasonably well and leads to state-of-the-art on many real world problems

