

Department of Computer Science
University of Bristol

COMSM0045 – Applied Deep Learning
comsm0045-applied-deep-learning.github.io

2020/21

Lecture 05


COST FUNCTIONS, REGULARISATION, DEPTH

Tilo Burghardt | tilo@cs.bris.ac.uk

29 Slides

Agenda Lecture 5

- Recap: Stochastic Gradient Descent
- Key Loss Functions
- L1 and L2 Weight Decay
- Dropout and Noise
- Data Augmentation
- Why deep is advantageous...
- Scalability Considerations...



RECAP: SGD



(Online) Backpropagation so far: Notational Compaction

initialise all weights w_{ij}^l randomly

for $t=0, 1, 2, \dots$ **do**

pick next training sample $([f_1^0, f_2^0, \dots], [f_1^*, f_2^*, \dots])$

FORWARD PASS: compute all $s_j^l = \sum_{i=1}^{d(l-1)} w_{ij}^l f_i^{l-1}$ and $f_j^l = g_j^l(s_j^l)$


compute top deltas $\delta_j^N = g_j^{\prime N}(s_j^N) \cdot \partial J / \partial f_j^N$

BACKWARD PASS: compute all $\delta_i^{l-1} = g_i^{\prime(l-1)}(s_i^{l-1}) \sum_{j=1}^{d(l)} w_{ij}^l \delta_j^l$

update weights $w_{ij}^l \leftarrow w_{ij}^l - \eta f_i^{l-1} \delta_j^l$

check if stopping criteria are met to break loop

return final weights w_{ij}^l



initialise all weights W randomly

for $t=0, 1, 2, \dots$ **do**

pick next training sample (x, f^*)

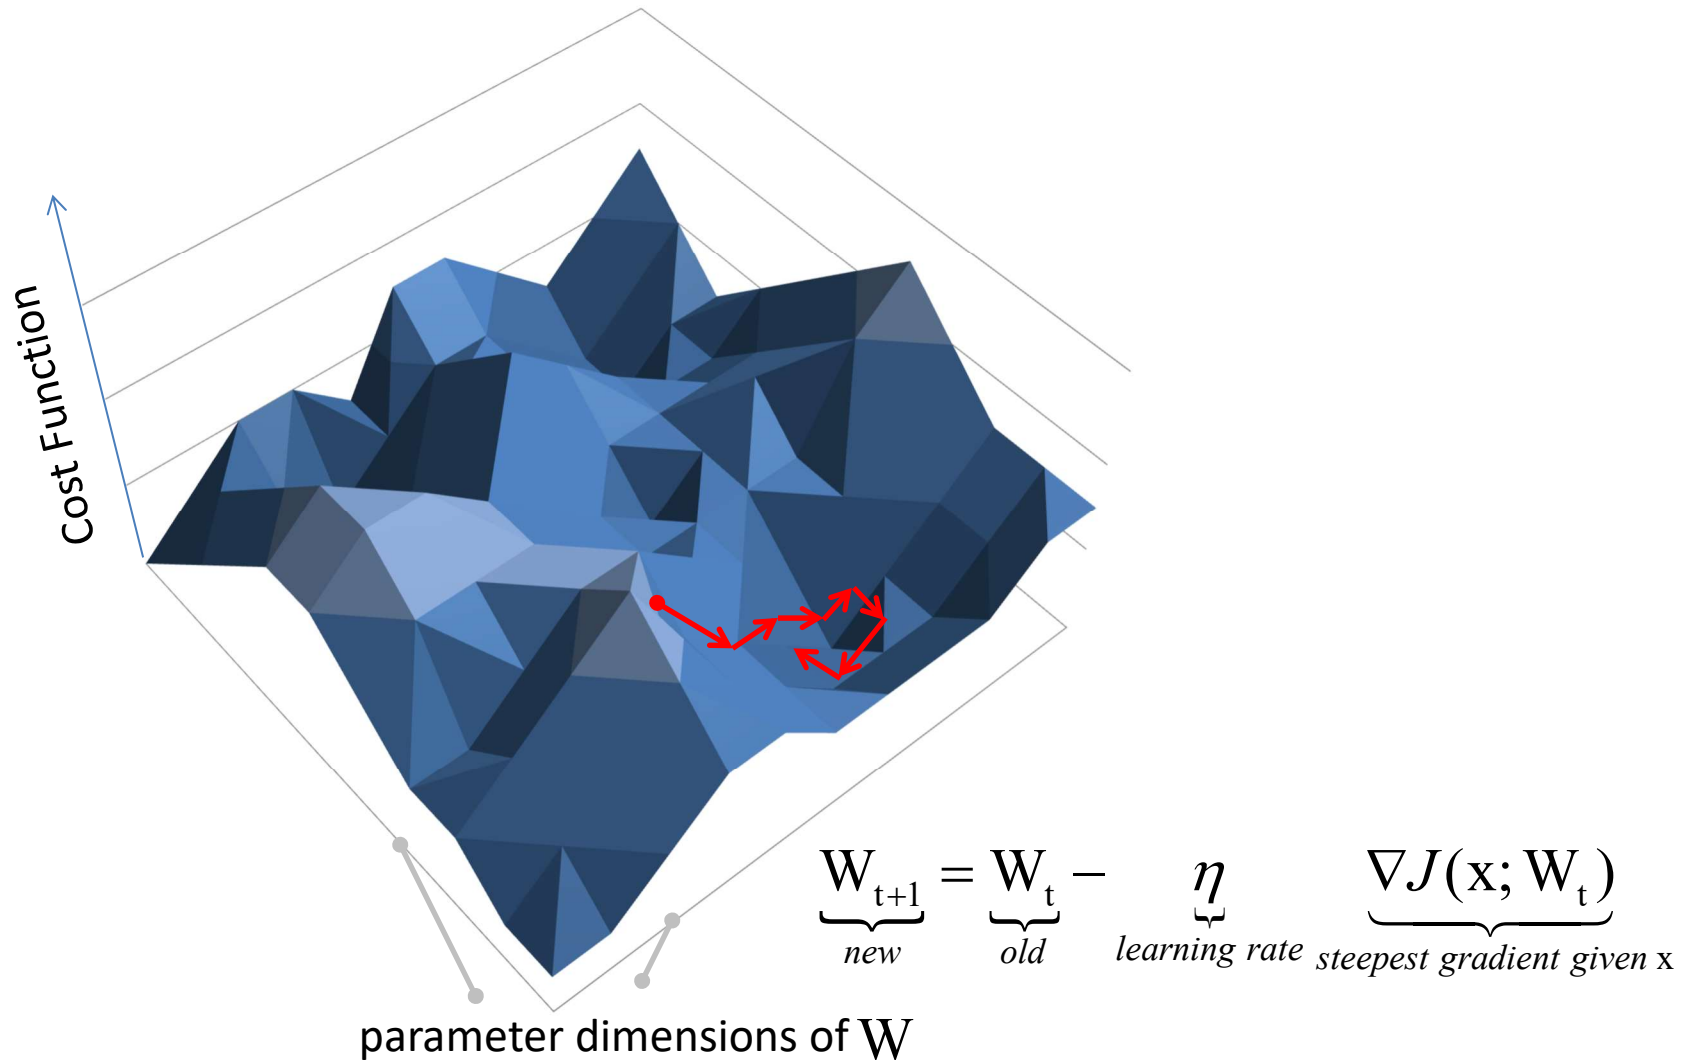
FORWARD-BACKWARD PASS: compute ∇J

update weights $W \leftarrow W - \eta \nabla J$

check if stopping criteria are met to break loop

return final weights W

Noisy Gradient Descent due to Online Sampling



SOME KEY COST FUNCTIONS



Adjusting the Output Layer for Classification

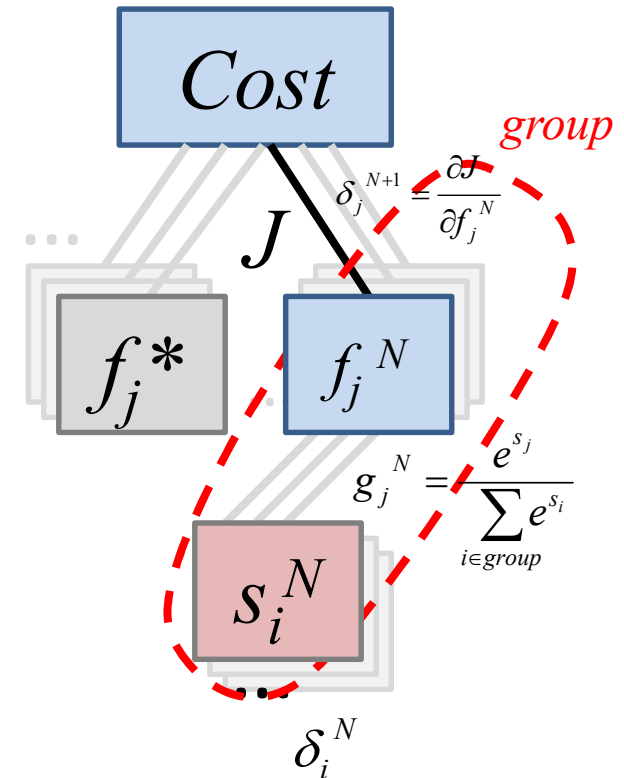
Practical Step 1: We want outputs to represent probabilities of class labels, thus outputs should be within $\{0, 1\}$ and sum to 1 ! – Can we force outputs to reflect such a distribution by creating a layer-wide, normalising non-linearity?

→ yes, introduce a **softmax neuron group** in the last layer with:

$$g_j^N(s_j^N) = \frac{e^{s_j^N}}{\sum_{i \in \text{group}} e^{s_i^N}};$$

$$g'_j{}^N(s_j^N) = f_j^N(1 - f_j^N); g'_j{}^N(s_{i \neq j}^N) = -f_j^N f_i^N$$

- now all outputs f_i^N range between 0 and 1 , while the **group** output sums to 1



Adjusting Loss for Classification


Now, so far we have changed the layout of the last layer to represent a classification setting.

What is an appropriate cost function?

→ use cross-entropy
as the *group's* cost function:

$$J = - \sum_{j \in \text{group}} f_j^* \log f_j^N$$


$$\delta_i^N = \sum_{j \in \text{group}} \frac{\partial J}{\partial f_j^N} \frac{\partial f_j^N}{\partial s_i^N} = f_i^N - f_i^*$$



- steepness of the **cost function derivative** now cancels the shallowness of the **softmax derivative** exactly, leading to an **MSE-style delta** propagated backwards from layer N – can we show why this works?

Counter-Action: Softmax exactly balances Cross-Entropy Loss

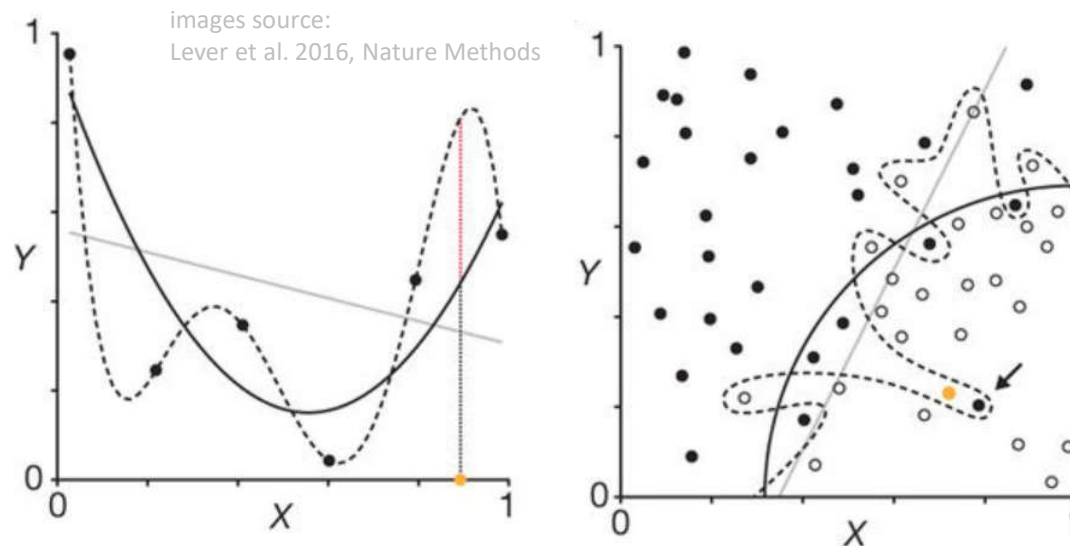
$$\begin{aligned}
 \delta_i^N &= \frac{\partial J}{\partial s_i^N} = - \sum_{j \in \text{group}} f_j^* \underbrace{\frac{\partial \log f_j^N}{\partial s_i^N}}_{\text{apply chain rule}} \\
 &= - \sum_{j \in \text{group}} f_j^* \frac{1}{f_j^N} \frac{\partial f_j^N}{\partial s_i^N} \quad \text{where} \quad f_j^N = \frac{e^{s_j^N}}{\sum_{i \in \text{group}} e^{s_i^N}} \\
 &= - \sum_{j=i} f_j^* \frac{1}{f_j^N} \overbrace{\frac{\partial f_j^N}{\partial s_i^N}}_{\text{for } i=j} - \sum_{j \neq i} f_j^* \frac{1}{f_j^N} \underbrace{\left(- f_j^N f_i^N \right)}_{\frac{\partial f_j^N}{\partial s_i^N} \text{ for } i \neq j} \\
 &= - f_i^* (1 - f_i^N) + \sum_{j \neq i} f_j^* \frac{f_j^N f_i^N}{f_j^N} \\
 &= - f_i^* + f_i^* f_i^N + \underbrace{\sum_{j \neq i} f_j^* f_i^N}_{= f_i^N \sum_{j \in \text{group}} f_j^*} \\
 &= f_i^N \left(\underbrace{\sum_{j \in \text{group}} f_j^*}_{=1} \right) - f_i^* \\
 &= f_i^N - f_i^*
 \end{aligned}$$



this is propagated backwards as the delta - basically, just the discrepancy for a particular output as one would intuitively expect...

Problem with High Parameter Spaces: Overfitting

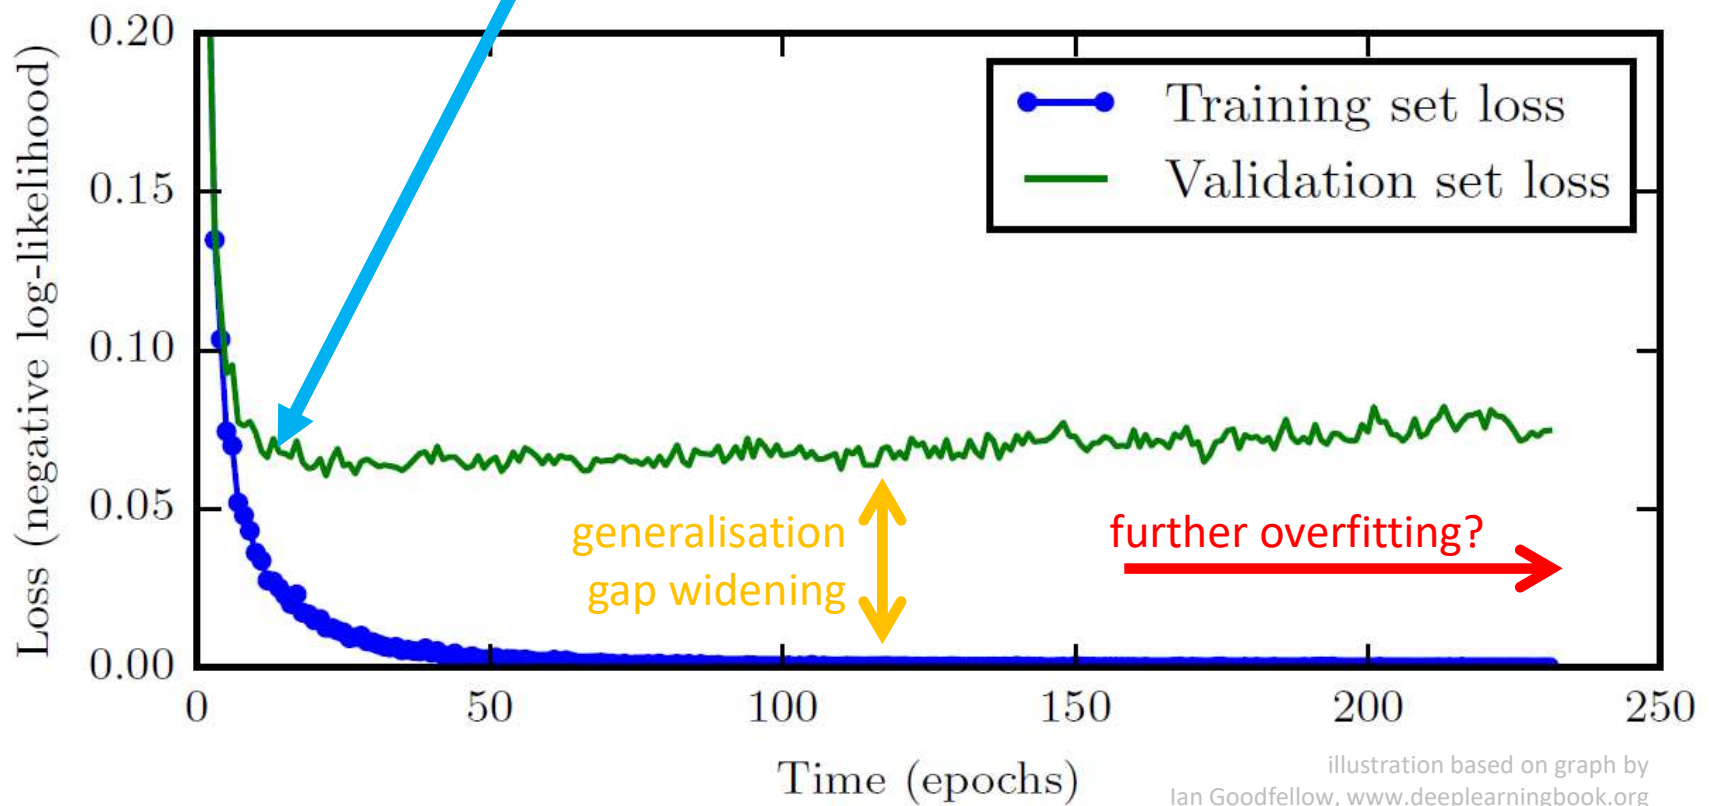
- consider ImageNet (*Krizhevsky et al. 2012*), which is a highly successful, by now classic network with (only) 8 layers; yet, it has already approx. 650,000 neurons, 60,000,000 parameters and 630 million connections!
→ models with such a high degrees of freedom are particularly prone to overfitting



overfitting (dotted) and underfitting (grey) examples for regression (left) and classification (right) task

Loss across Training: 'Pocket' Algorithm and Early Stopping

Keep best N results at any stage ('in your pocket') based on validation set and decide based on further performance whether to revert back to it. One may also decide to terminate learning altogether (early stopping) while the validation performance is better.



REGULARISATION



Regularisation

“Regularization is any modification we make to a learning algorithm that is intended to reduce its generalization error, but not its training error.”

From: Ian Goodfellow 2015

→ One may view many recent developments in deep learning as attempts to reduce the complexity of neural network training models – including **L-regularization** (constraining weight space), **dropout** (stochastic network thinning), **depth** (composition over concatenation), as well as parameter (CNNs) and network (RNNs) **sharing**.

L2-Regularisation

- **Assumption:** we target a local minimum with *small-magnitude weights* to combat overfitting
- **Idea:** introduce a penalty for every weight based on its squared value directly in the cost function J :

$$J(\mathbf{X}; \mathbf{W}) = \frac{1}{|\mathbf{X}|} \sum_{\mathbf{x} \in \mathbf{X}} L(f(\mathbf{x}; \mathbf{W}), f^*(\mathbf{x})) + \frac{1}{2} \sum_{w \in \mathbf{W}} \lambda w^2$$

- practically, this is easiest implemented as scaled weight decay towards zero during training:

$$\mathbf{W}_{t+1} = \mathbf{W}_t - \eta (\nabla J_{plain}(\mathbf{X}; \mathbf{W}_t) + \lambda \mathbf{W}_t)$$

L1-Regularisation


- **Assumption:** we target a local minimum with *sparse weights* against overfitting (Occam's Razor)
- **Idea:** introduce a penalty for every weight based on its magnitude directly in the cost function J :

$$J(\mathbf{X}; \mathbf{W}) = \frac{1}{|\mathbf{X}|} \sum_{\mathbf{x} \in \mathbf{X}} L(f(\mathbf{x}; \mathbf{W}), f^*(\mathbf{x})) + \sum_{w \in \mathbf{W}} \mu |w|$$


- practically, this is easiest implemented as absolute weight decay towards zero during training:

$$\mathbf{W}_{t+1} = \mathbf{W}_t - \eta (\nabla J_{plain}(\mathbf{X}; \mathbf{W}_t) + \mu \text{sgn}(\mathbf{W}_t))$$

DROPOUT



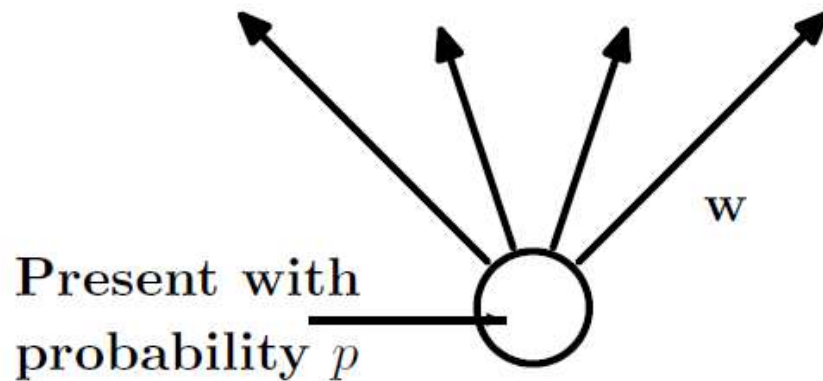
Neural Networks as Ensemble of Sub-networks



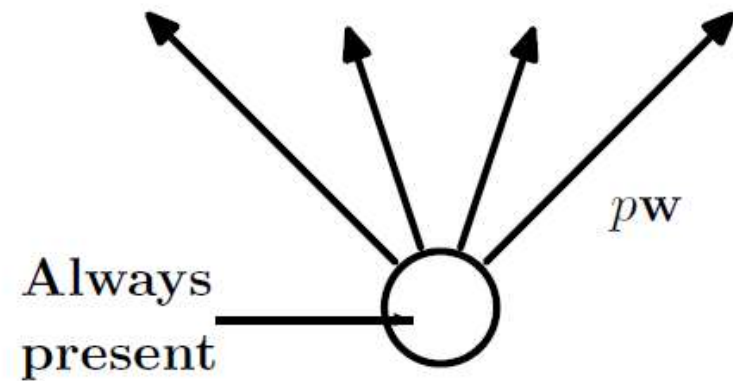
adjusted based on graph by
Ian Goodfellow, www.deeplearningbook.org

Dropout for Regularisation (*Srivastava et al. 2014*)

- **Idea:** during training, for each loop of weight updates 'immobilise' a random subset of neurons with probability $(1-p)$ by forcing their output to zero (particularly effective when applied to final fully-connected layers of networks)
→ these neurons now do not contribute to the network output nor can they be used for adjusting weights or passing gradients through them; p is often set to a value around 0.5 before tuning on validation data




Training Time



Testing Time

image source:
Srivastava et al. 2014

MNIST Test Performance Impact of Dropout on Test Data



DropConnect (*Wan et al. 2013*)




image source:
Wan Li, <http://cs.nyu.edu/~wanli/dropc>

- **Idea:** instead of dropping neurons, DropConnect sets weights to zero with a probability $(1-p)$ providing a more fine-grained mechanism for regularisation based on the same underlying concept.

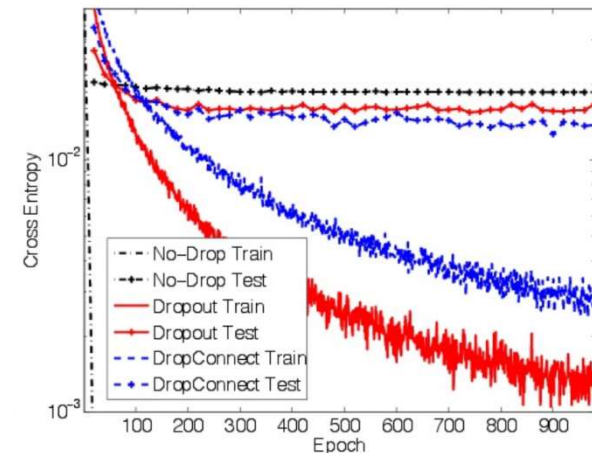


image source:
Wan Li, <http://cs.nyu.edu/~wanli/dropc>

Data against Overfitting

- so far, we have addressed overfitting by limiting the representational capacity of our networks by introduction of model constraints and by interpretation of the networks as ensembles
- however, overfitting can always be addressed by more data, more representative data and/or strategically sampled data





image source: Andrew Ng

- given the high dimensionality of deep net parameter spaces, taking advantage of more data is often possible without reaching the representational limits of the network

DATA AUGMENTATION



Realistic Variance Matters: Augmentation for Images

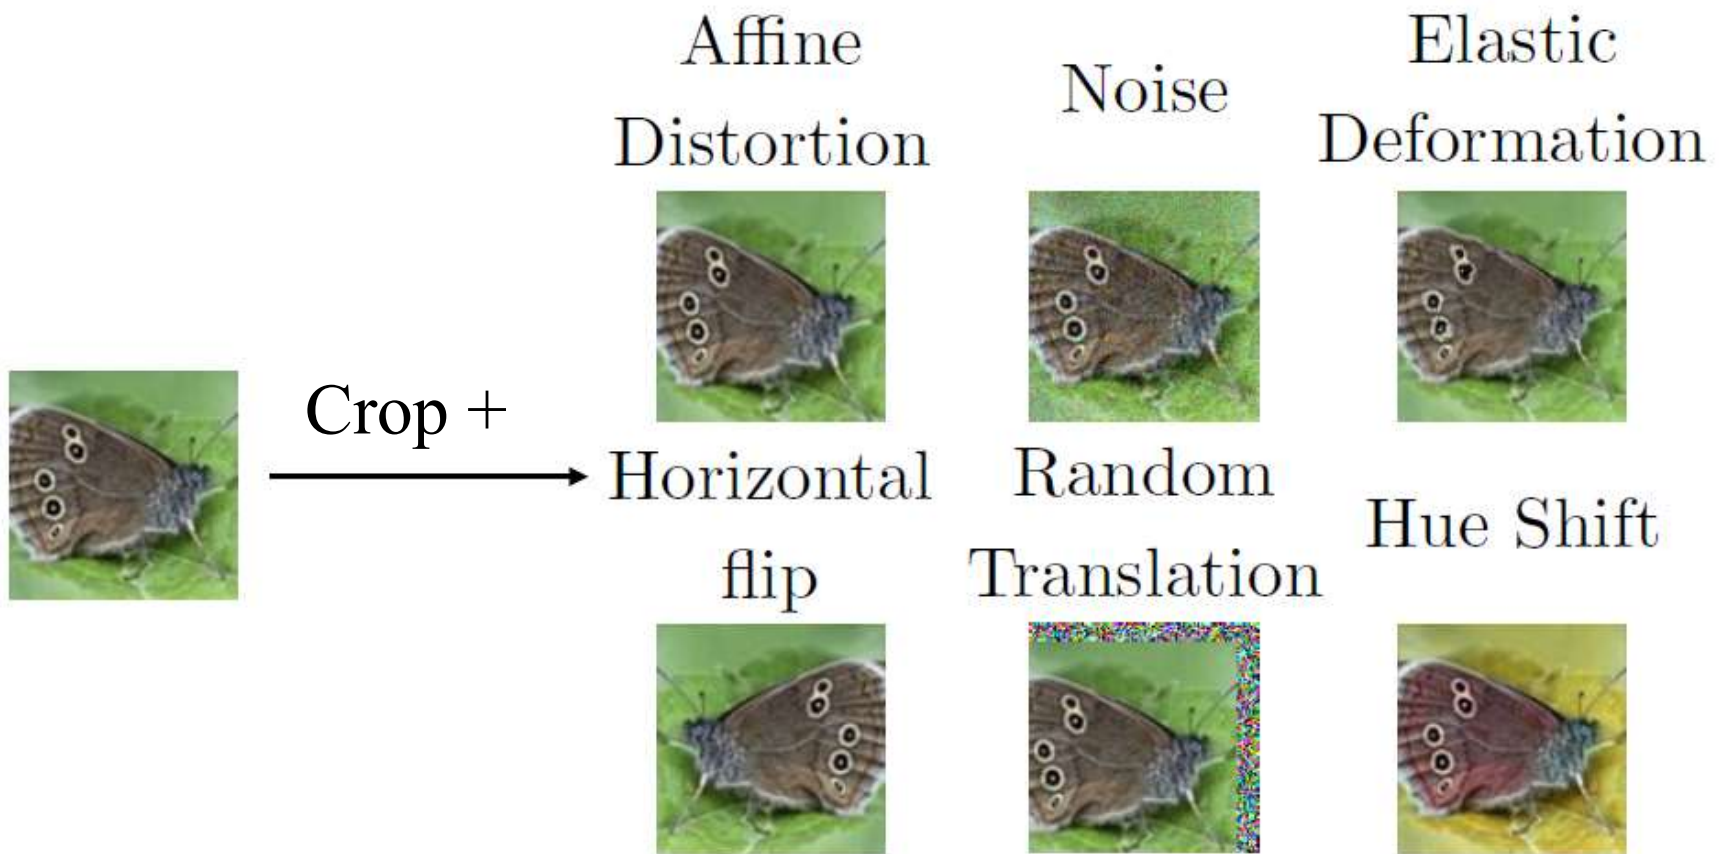


image source: Ian Goodfellow


Example: Augmentation for Images



image source: V Yadav, <https://github.com/vxy10>

WHY DEEP AND SCALABILITY

(OPEN TOPIC)



Point 1: Hierarchical Automatic Modularisation

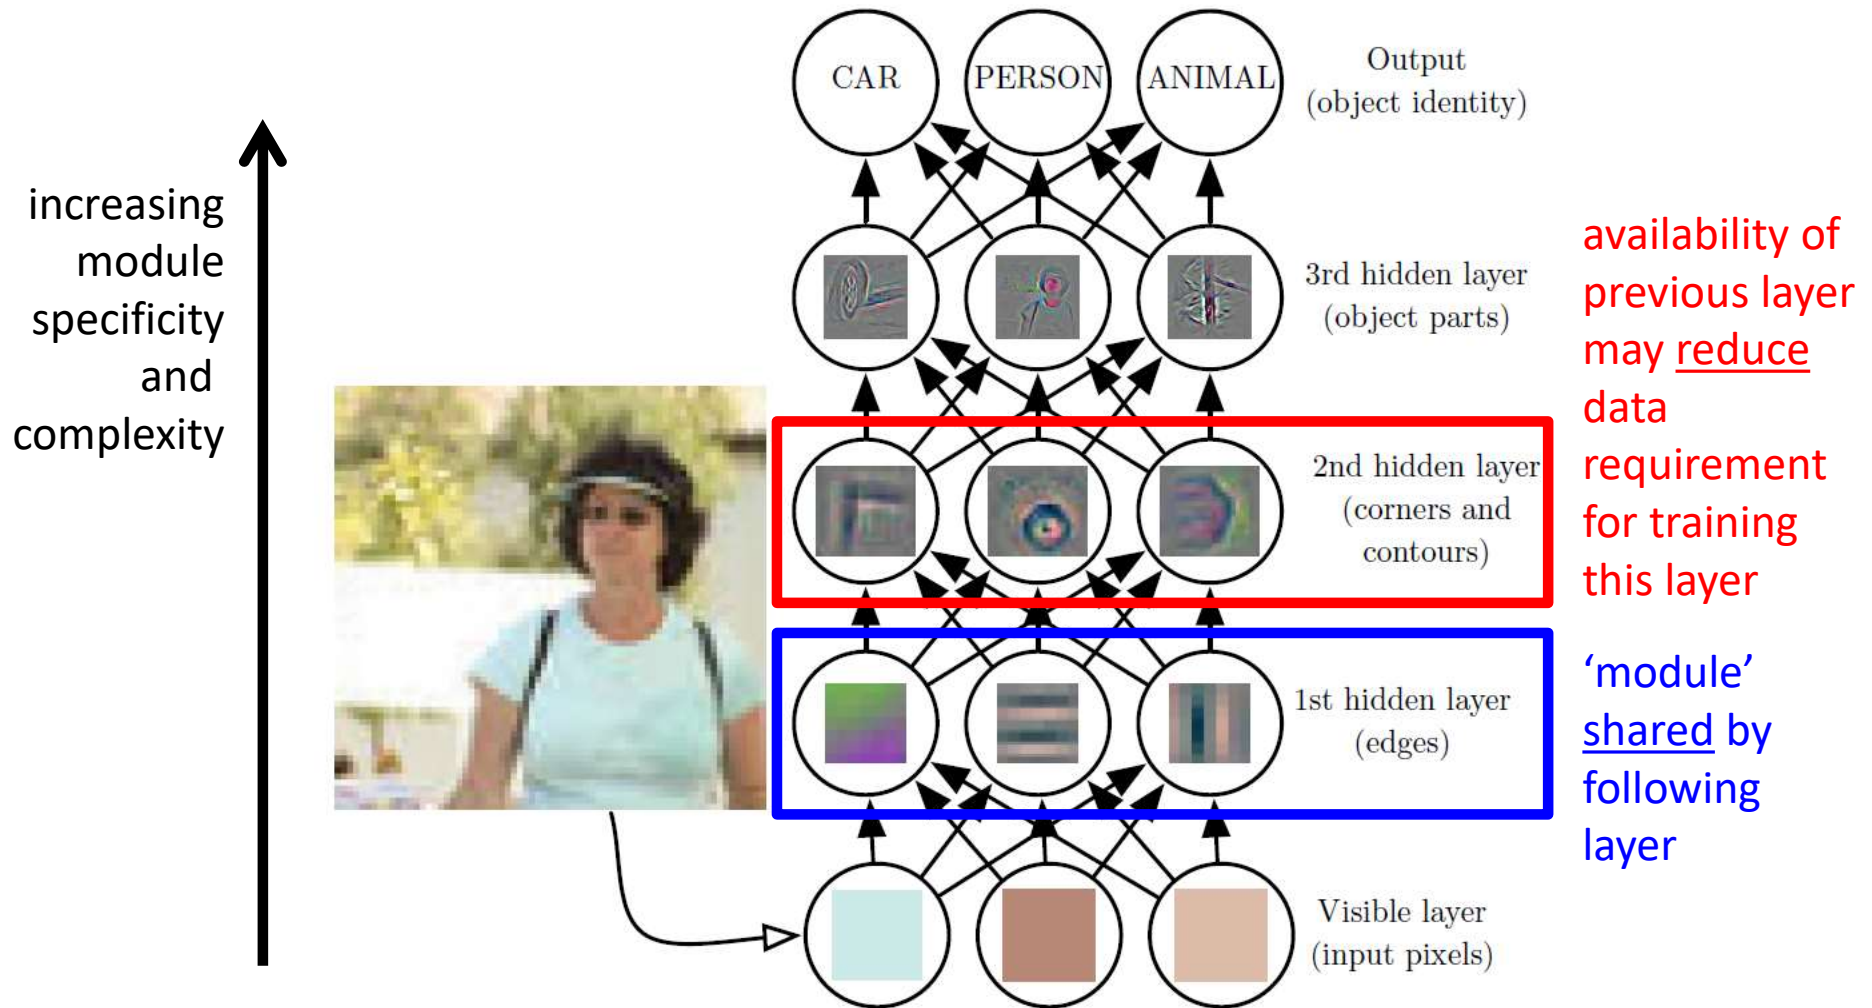


image adapted from Ian Goodfellow, www.deeplearningbook.org

Point 2: Practical Performance (Example)

- **Example:** Conversational Speech Transcription Using Context-Dependent Deep Neural Networks (*Seide et al. 2011*)
- direct comparison of different layer depth shows performance advantage for deeper structures across many domains
- however, required training times increase rapidly given large datasets and deep/large network architectures

Hidden Layers x Size	Error Rate (%) (word-error)
$1 \times 2k$	24.2
$2 \times 2k$	20.4
$3 \times 2k$	18.4
$4 \times 2k$	17.8
$5 \times 2k$	17.2
$7 \times 2k$	17.1
$1 \times 3.8k$	22.5
$1 \times 4.6k$	22.6
$1 \times 16k$	22.1

data source:
Seide et al 2011

Point 3: Sketch of Oscillation Argument (Telgarsky 2016)

- **Key Claim:** There are functions f that can be represented by a deep ReLU net with a polynomial number of neurons, where a shallow network would require exponentially many units.
- Telgarsky shows that k ReLU-like layers h can generate in the order of 2^{k-1} oscillations to approximate functions, while shallow networks have exponentially less peaks.
- **Open Question:** Are these functions f the type of 'natural' functions of interest to practical problems?


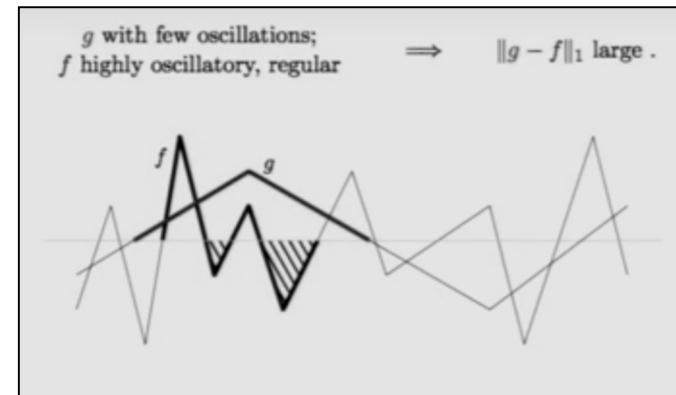


image source:
Telgarsky 2016, watch
<https://www.youtube.com/watch?v=ssaXJqG9Dz4>

Some Current Ideas related to Scale and Depth

- **Idea of Mixtures of Experts:** extend network capacity and yet keep parameters per training run at bay via learning to switch on/off parts of networks dependent on their appropriateness per example (*Shazeer et al. 2017*)
- **Idea of Network Distillation:** learn deep ensembles, but then 'compress' deep networks into shallow or more efficient compactions (*Ba and Caruana 2014, Hinton 2017*)

