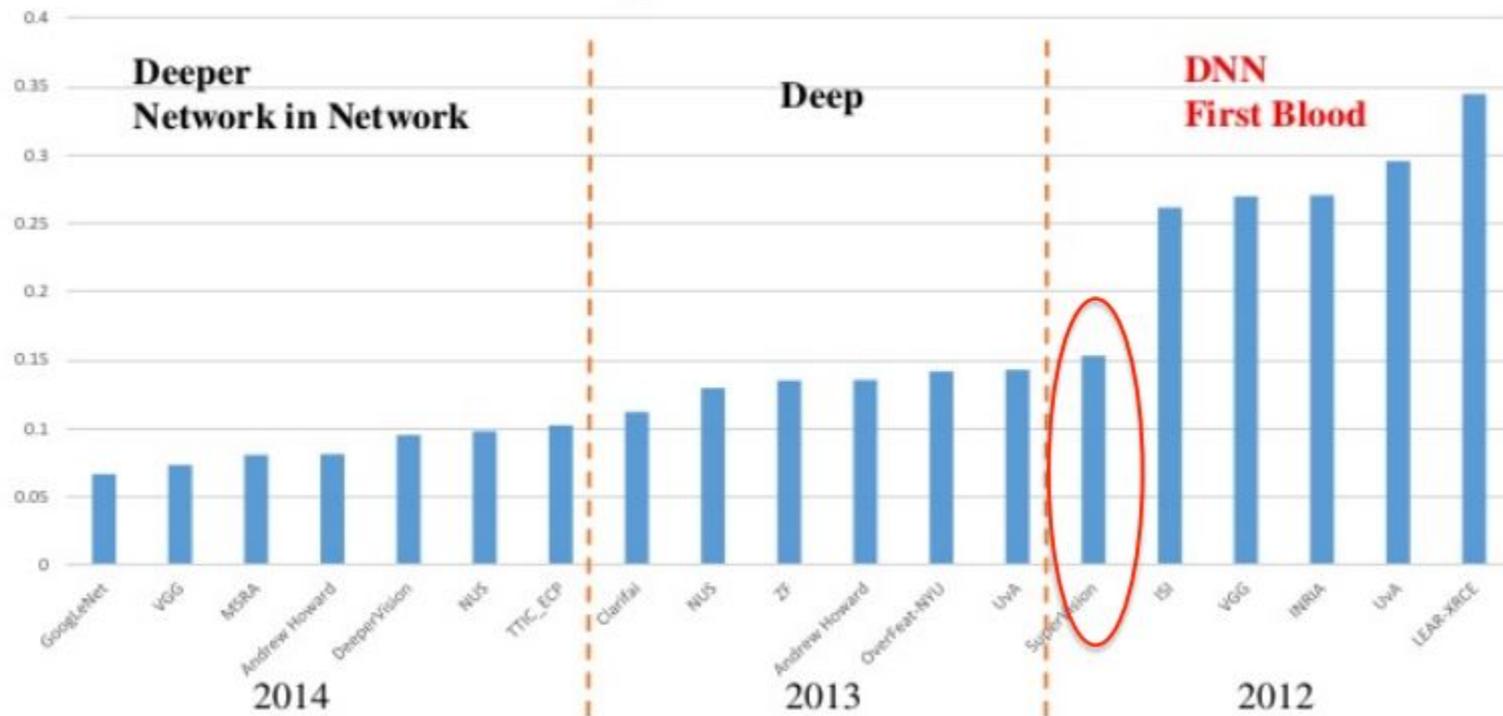

Automated Deep Learning by Neural Architecture Search

—

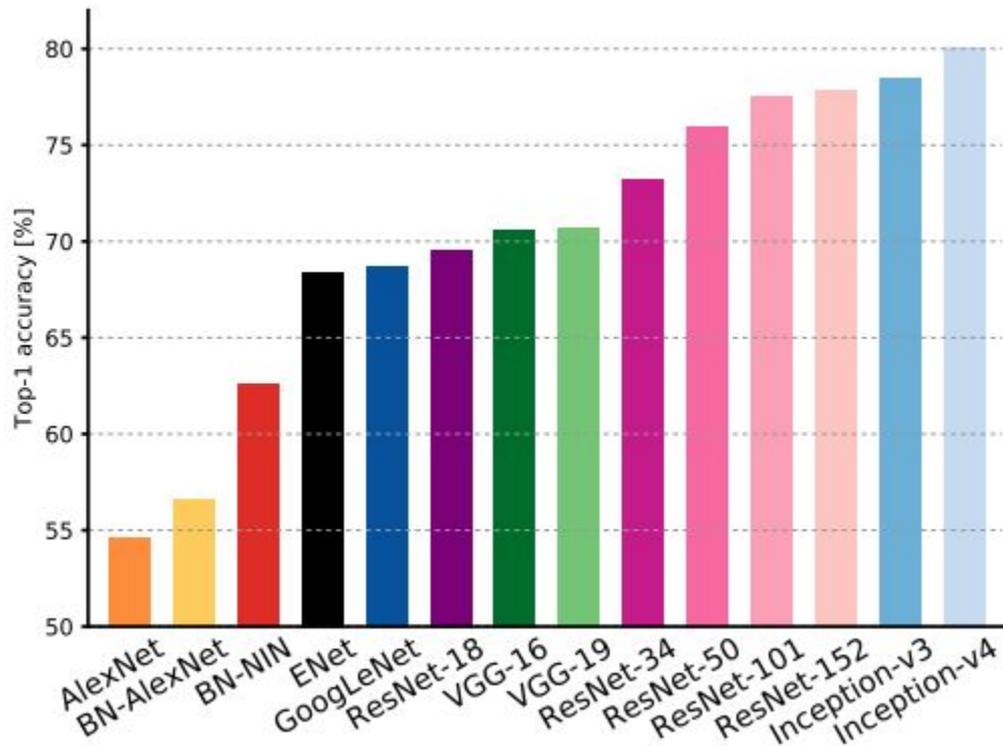
Yuu Jinnai
Brown University

Deep Learning in Computer Vision



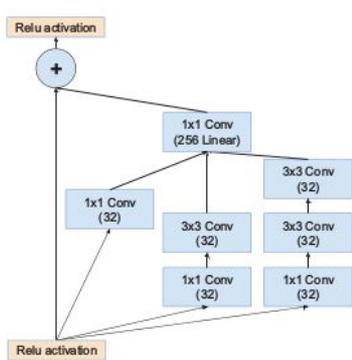
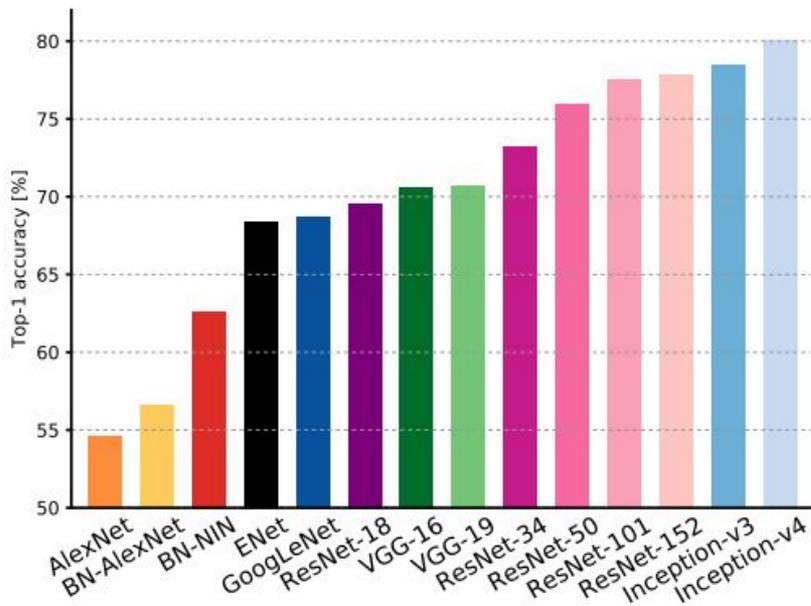
(FeiFei Li 2014)

Deep Learning in Computer Vision

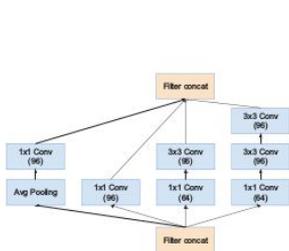


(Canziani et al. 2016)

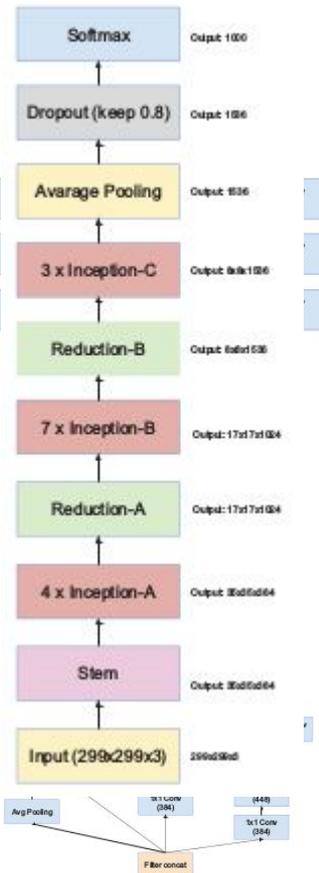
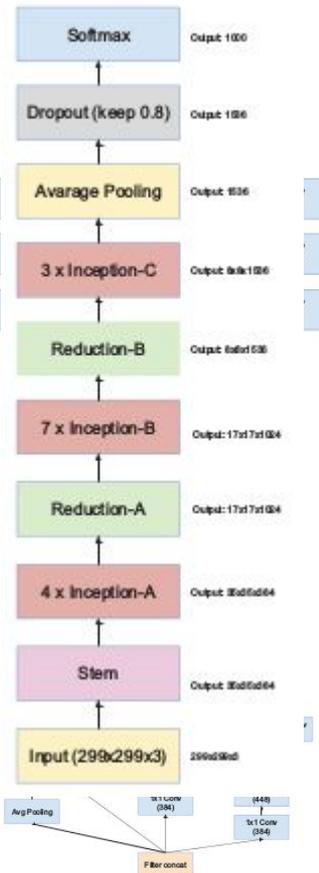
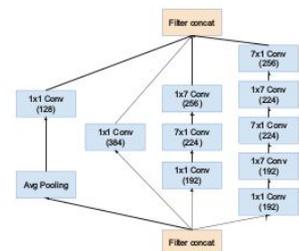
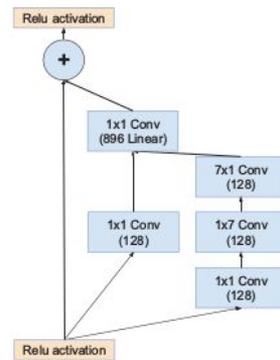
No More Engineering...?



Inception-v1



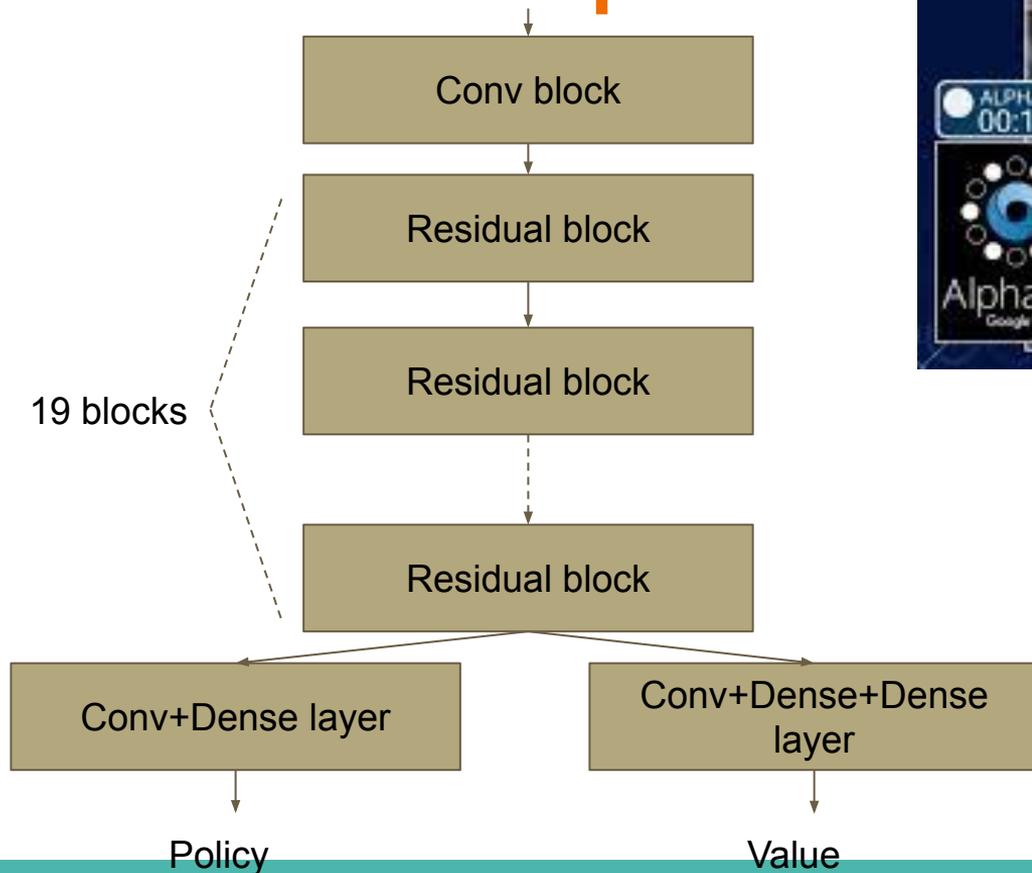
Inception-v4



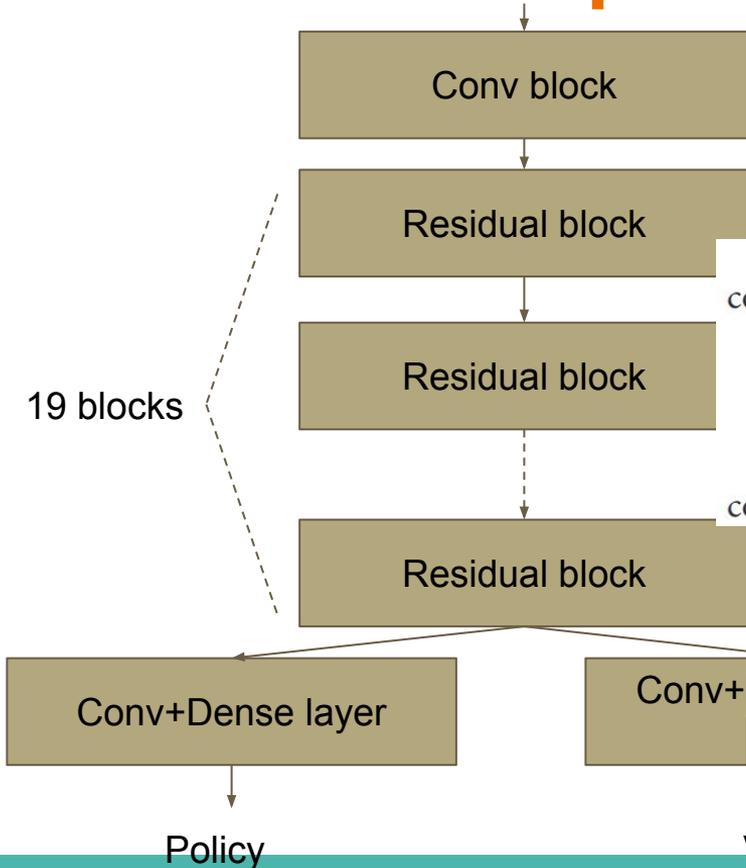
Architecture of AlphaGo



Architecture of AlphaGo



Architecture of AlphaZero



The convolutional block applies the following modules:

- (1) A convolution of 256 filters of kernel size 3×3 with stride 1
- (2) Batch normalization¹⁸
- (3) A rectifier nonlinearity

Each residual block applies the following modules sequentially to its input:

- (1) A convolution of 256 filters of kernel size 3×3 with stride 1
- (2) Batch normalization
- (3) A rectifier nonlinearity
- (4) A convolution of 256 filters of kernel size 3×3 with stride 1
- (5) Batch normalization
- (6) A skip connection that adds the input to the block
- (7) A rectifier nonlinearity

The output of the residual tower is passed into two separate 'heads' for computing the policy and value. The policy head applies the following modules:

- (1) A convolution of 2 filters of kernel size 1×1 with stride 1
- (2) Batch normalization
- (3) A rectifier nonlinearity

(4) A fully connected linear layer that outputs a vector of size $19^2 + 1 = 362$, corresponding to logit probabilities for all intersections and the pass move

The value head applies the following modules:

- (1) A convolution of 1 filter of kernel size 1×1 with stride 1
- (2) Batch normalization
- (3) A rectifier nonlinearity

(4) A fully connected linear layer to a hidden layer of size 25

(5) A rectifier nonlinearity

(6) A fully connected linear layer to a scalar

“Machine Learning have become an Alchemy”



Claim: We Need to Automate

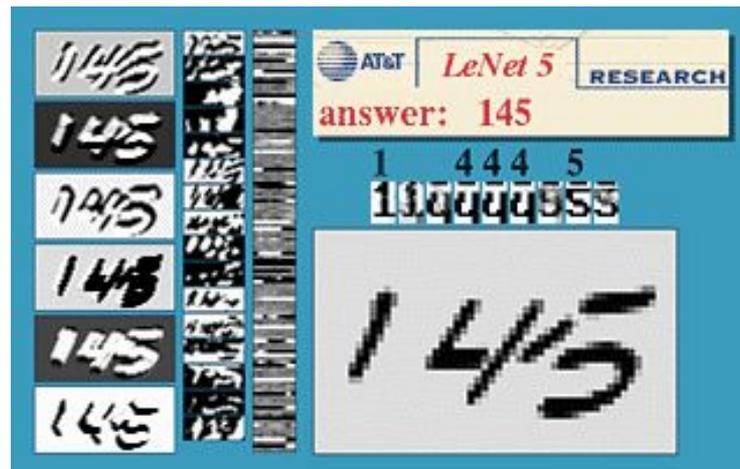
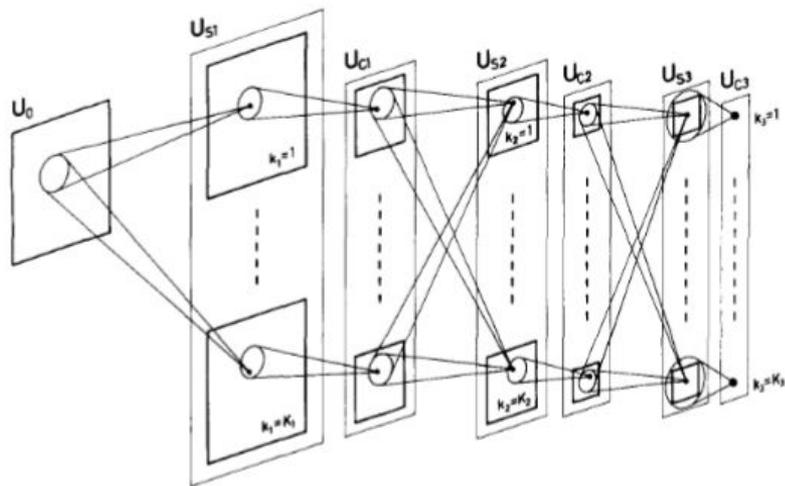
Claim: Structure of efficient neural network is so complicated that human expert cannot design by heuristic

-> **We need AutoML!**

Deep Learning Techniques

1. **Convolution and Sparsification**
2. Vanishing Gradient Problem: Skip Connection, LSTM
3. Hierarchical Architecture

Convolution: Sparsification



(Fukushima 1980; LeCun et al. 1998)

Deep Deep Architectures

Will smaller number of parameters, CNN can be deeper than Densenet

Deeper networks often perform better (Simonyan&Zisserman 2014)

Deeper networks can learn higher level features (Nguyen&Hein 2017)

Overparameterization as a momentum (Arora et al. 2018)

Deep Deep Architectures

Methods	LeNet-5[48]	AlexNet [7]	OverFeat (fast)[8]	VGG-16[9]	GoogLeNet [10]	ResNet-50(v1)[11]
Top-5 errors	n/a	16.4	14.2	7.4	6.7	5.3
Input size	28x28	227x227	231x231	224x224	224x224	224x224
Number of Conv Layers	2	5	5	16	21	50
Filter Size	5	3,5,11	3,7	3	1,3,5,7	1,3,7
Number of Feature Maps	1,6	3-256	3-1024	3-512	3-1024	3-1024
Stride	1	1,4	1,4	1	1,2	1,2
Number of Weights	26k	2.3M	16M	14.7M	6.0M	23.5M
Number of MACs	1.9M	666M	2.67G	15.3G	1.43G	3.86G
Number of FC layers	2	3	3	3	1	1
Number of Weights	406k	58.6M	130M	124M	1M	1M
Number of MACs	405k	58.6M	130M	124M	1M	1M
Total Weights	431k	61M	146M	138M	7M	25.5M
Total MACs	2.3M	724M	2.8G	15.5G	1.43G	3.9G

Deep Learning Techniques

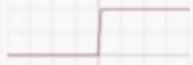
1. Convolution and Sparsification
2. **Vanishing Gradient Problem: Skip Connection, LSTM**
3. Hierarchical Architecture

Vanishing Gradient Problem (Hochreiter 1991)

Gradient gets smaller in deep architectures (or RNN)

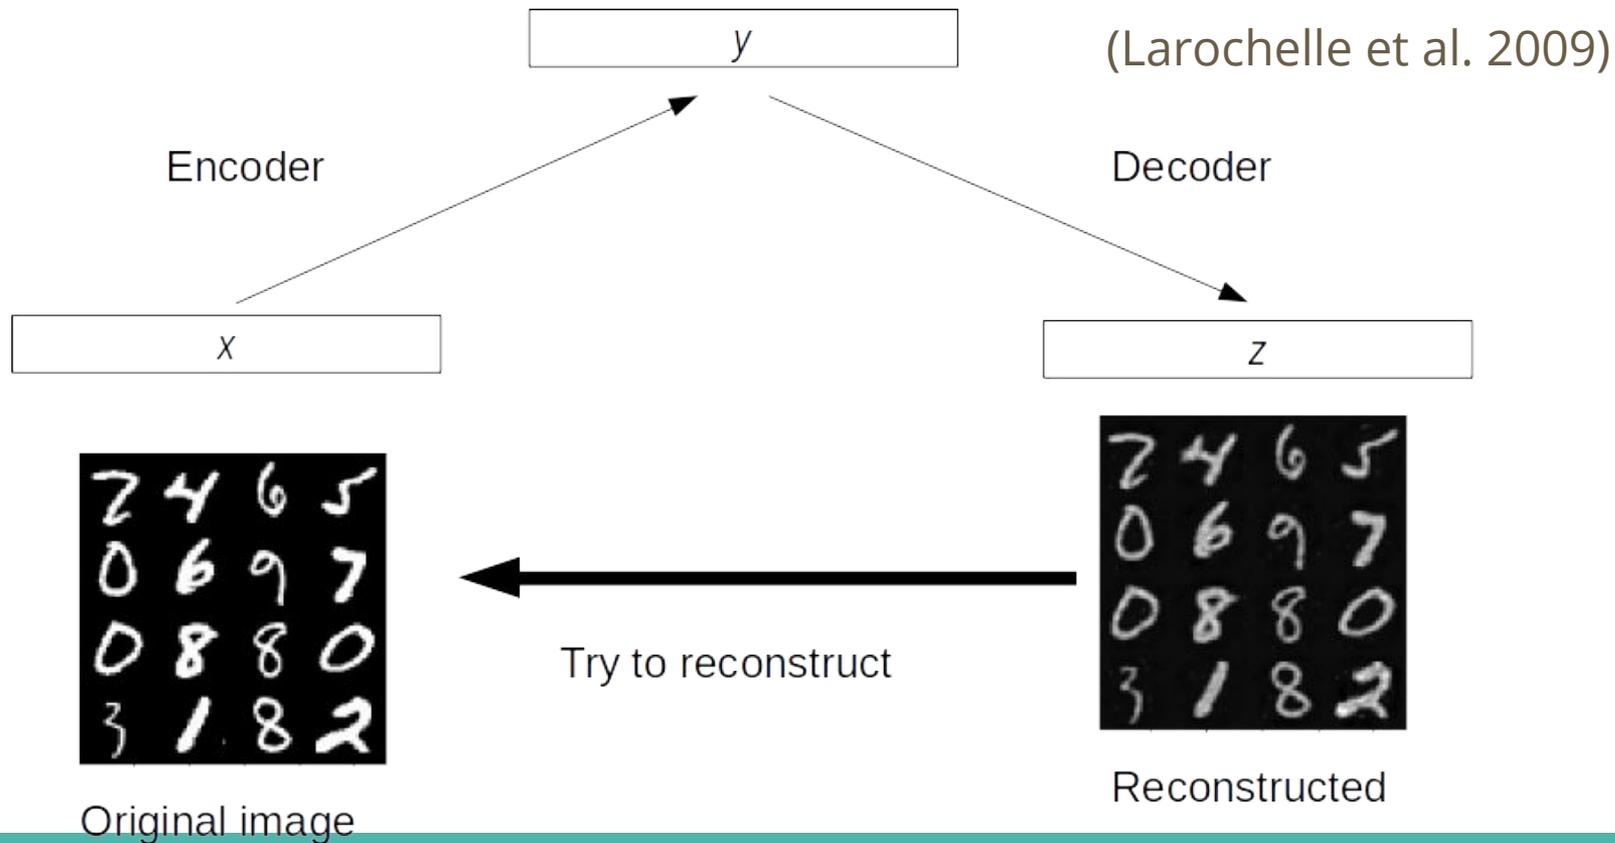
Example: Hyperbolic tangent activation function have gradient in (0,1) range

Solution 1. Activation Functions

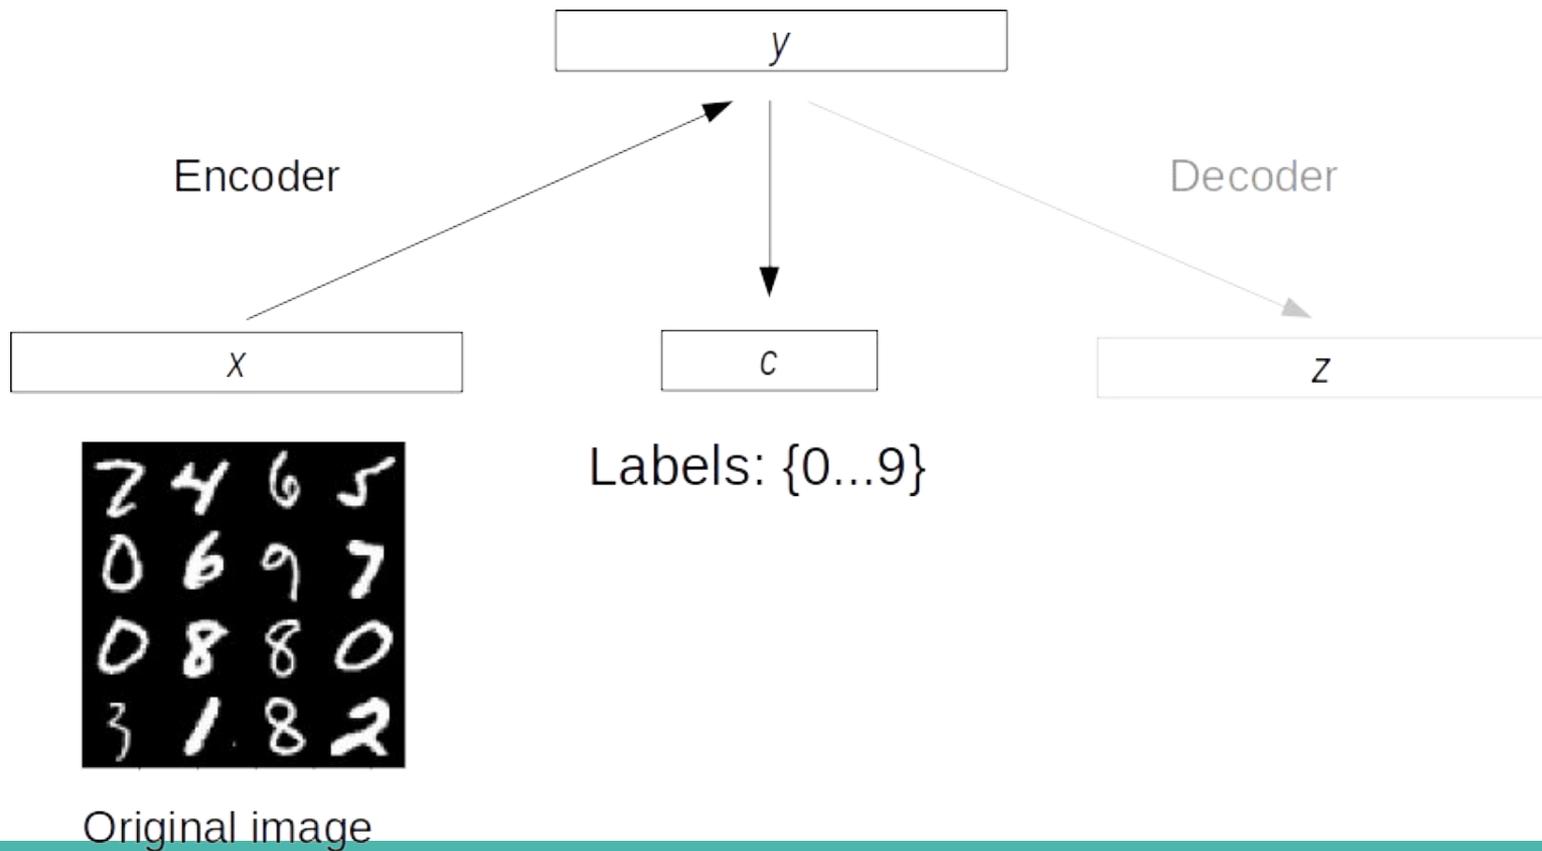
Name	Plot	Equation	Derivative
Identity		$f(x) = x$	$f'(x) = 1$
Binary step		$f(x) = \begin{cases} 0 & \text{for } x < 0 \\ 1 & \text{for } x \geq 0 \end{cases}$	$f'(x) = \begin{cases} 0 & \text{for } x \neq 0 \\ ? & \text{for } x = 0 \end{cases}$
Logistic (a.k.a Soft step)		$f(x) = \frac{1}{1 + e^{-x}}$	$f'(x) = f(x)(1 - f(x))$
Tanh		$f(x) = \tanh(x) = \frac{2}{1 + e^{-2x}} - 1$	$f'(x) = 1 - f(x)^2$
ArcTan		$f(x) = \tan^{-1}(x)$	$f'(x) = \frac{1}{x^2 + 1}$
Rectified Linear Unit (ReLU)		$f(x) = \begin{cases} 0 & \text{for } x < 0 \\ x & \text{for } x \geq 0 \end{cases}$	$f'(x) = \begin{cases} 0 & \text{for } x < 0 \\ 1 & \text{for } x \geq 0 \end{cases}$

(Figure from <https://towardsdatascience.com/activation-functions-neural-networks-1cbd9f8d91d6>)

Solution 2. Pretraining with RBM/Autoencoder

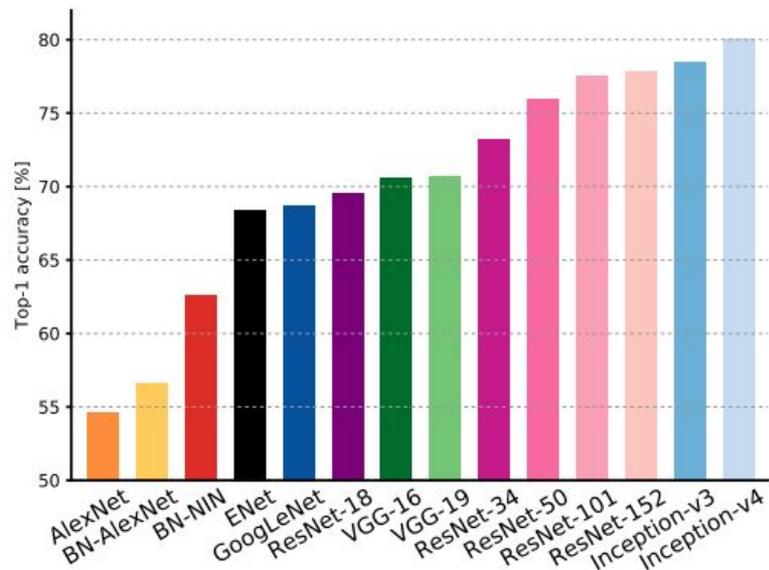
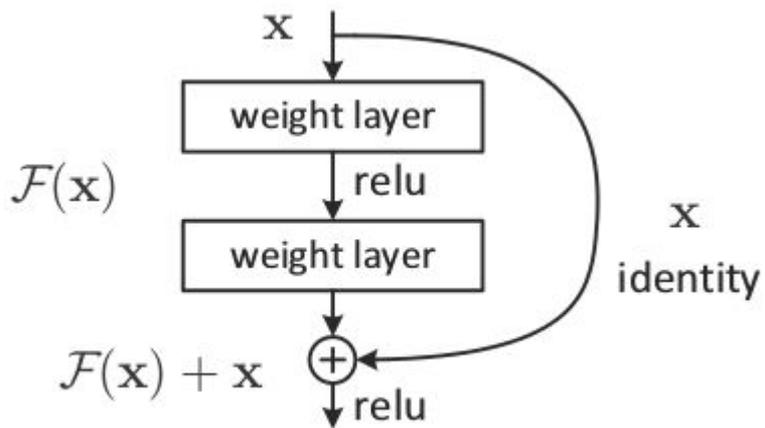


Solution 2. Pretraining with RBM/Autoencoder



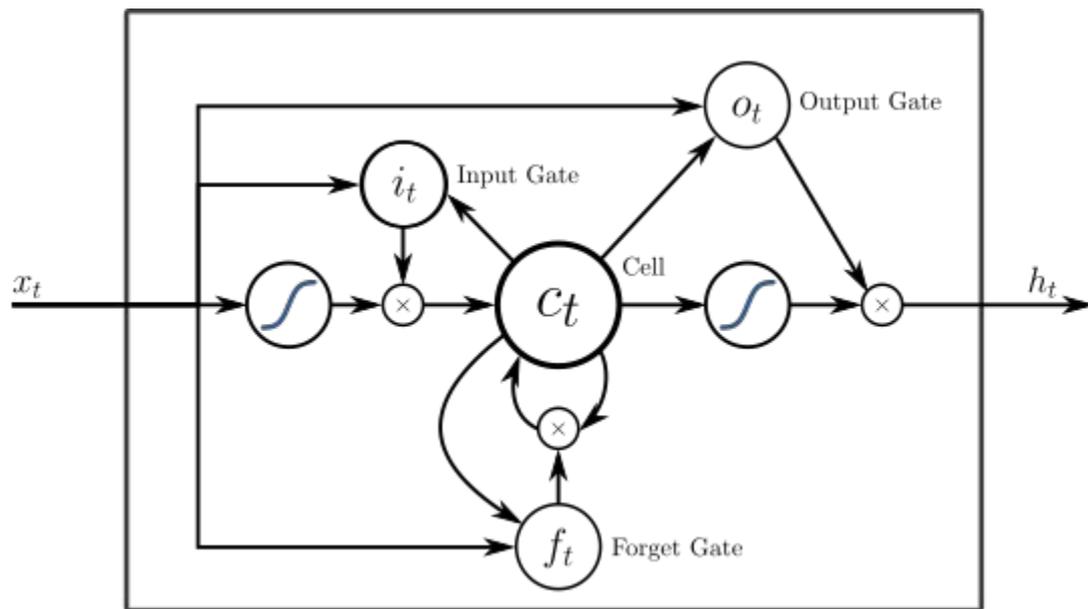
Solution 3. Skip Connections

ResNet (Kaiming et al. 2015)



Solution 3. Skip Connections

LSTM/GRU (Hochreiter&Schmidhuber1997; Chung et al. 2014)



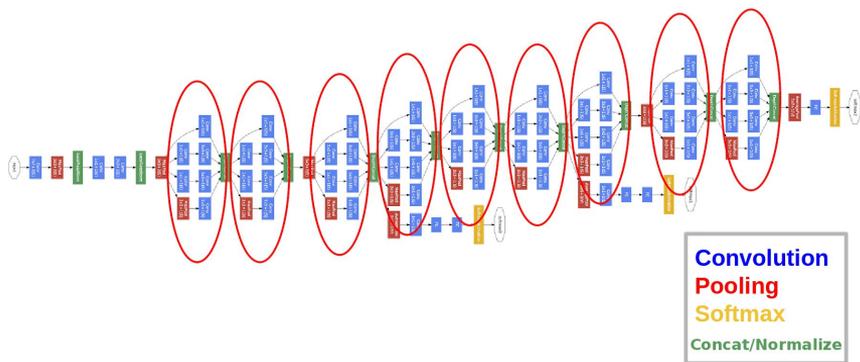
Deep Learning Techniques

1. Convolution and Sparsification
2. Vanishing Gradient Problem: Skip Connection, LSTM
3. **Hierarchical Architecture**

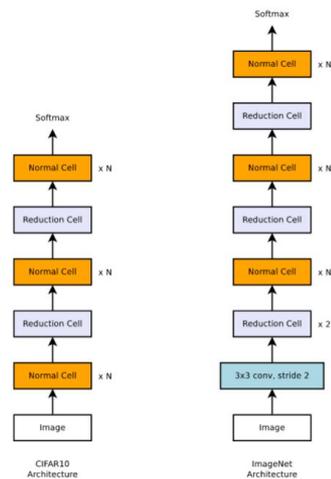
Hierarchical Architecture

Almost all sota NNs have hierarchical structures

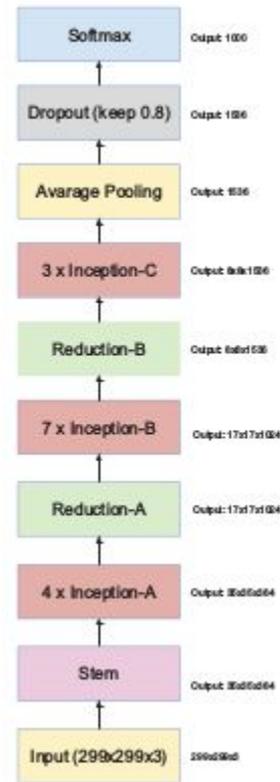
Repetition of some structures is effective



GoogLeNet



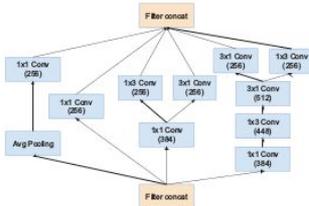
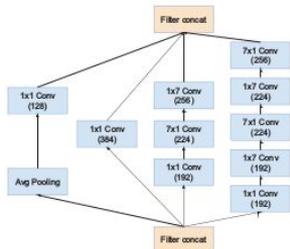
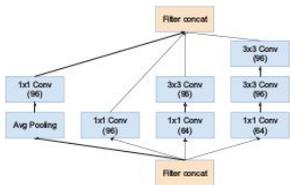
NasNet



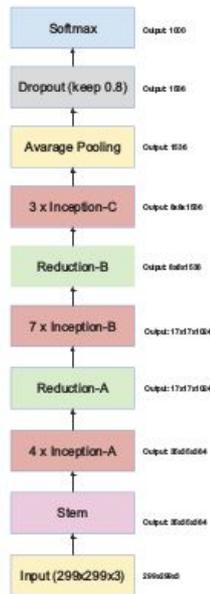
Inception-v4

What We Know Today

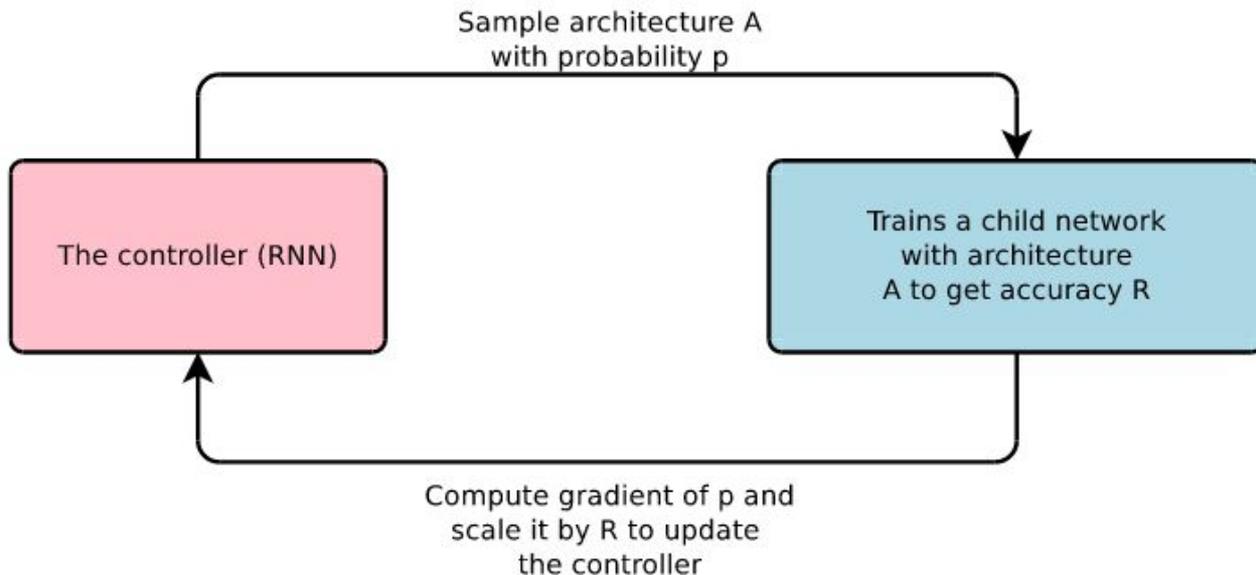
1. Hierarchical architecture works great in many domains
2. Efficient substructures are often very complex and domain dependent



Inception-v4



Neural Architecture Search (Zoph&Le 2016)



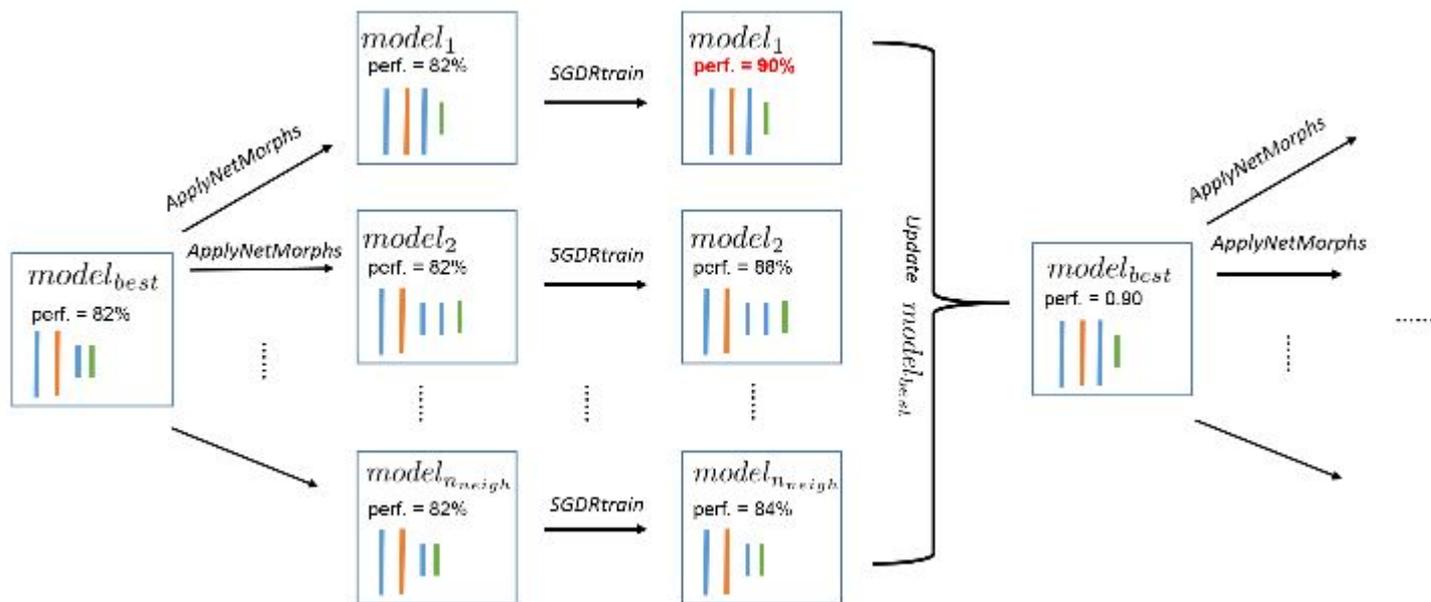
Challenges for Architecture Search

Exponentially Grown Search Spaces:

- Layer choices: *convolution, pooling, batch normalization, etc. (8 types of layer)*
- Parameters in a layer: *e.g. filters, kernel size, strides in conv layer. (at least two design choices in a layer)*
- Each layer has ~16 choices
- Total architectures exponentially grow with depth (*e.g. 2 layers = 16^2 , 3 layers = 16^3 , and 10 layers, 16^{10}*)
- **We need 1.5 million years to fully train all the possible architectures on a TITAN-XP**

Exhaustive Search is infeasible; and a smart choice must be made

Hill Climbing (Liu et al. 2017)



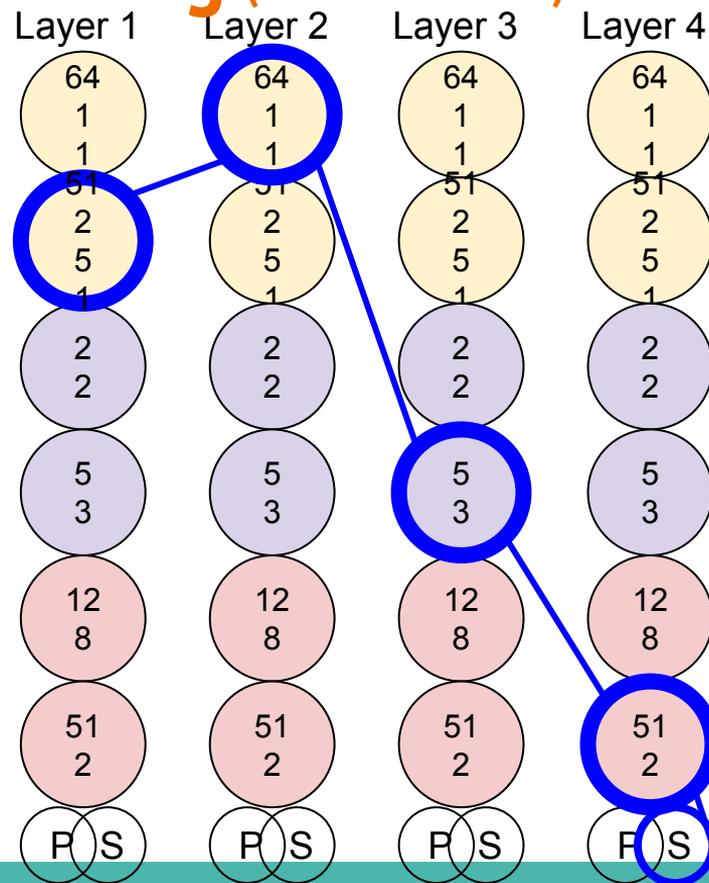
Hill Climbing (Liu et al. 2017)

model	resources spent	# params (mil.)	error (%)
Shake-Shake (Gastaldi, 2017)	2 days, 2 GPUs	26	2.9
WRN 28-10 (Loshchilov & Hutter, 2017)	1 day, 1 GPU	36.5	3.86
Baker et al. (2016)	8-10 days, 10 GPUs	11	6.9
Cai et al. (2017)	3 days, 5 GPUs	19.7	5.7
Zoph & Le (2017)	800 GPUs, ? days	37.5	3.65
Real et al. (2017)	250 GPUs, 10 days	5.4	5.4
Saxena & Verbeek (2016)	?	21	7.4
Brock et al. (2017)	3 days, 1 GPU	16.0	4.0
Ours (random networks, $n_{steps} = 5, n_{neigh} = 1$)	4.5 hours	4.4	6.5
Ours ($n_{steps} = 5, n_{neigh} = 8, 10$ runs)	0.5 days, 1 GPU	5.7	5.7
Ours ($n_{steps} = 8, n_{neigh} = 8, 4$ runs)	1 day, 1 GPU	19.7	5.2
Ours (snapshot ensemble, 4 runs)	2 days, 1 GPU	57.8	4.7
Ours (ensemble across runs)	1 day, 4 GPUs	88	4.4

Reinforcement Learning: Q-Learning (Baker et al. 2016)

$$Q(s, a) := Q(s, a) + \alpha \left[r + \gamma \max_{a'} Q(s', a') - Q(s, a) \right]$$

(r = accuracy)



Markov Decision Processes (MDP)

$$M = (S, A, T, R, \gamma)$$

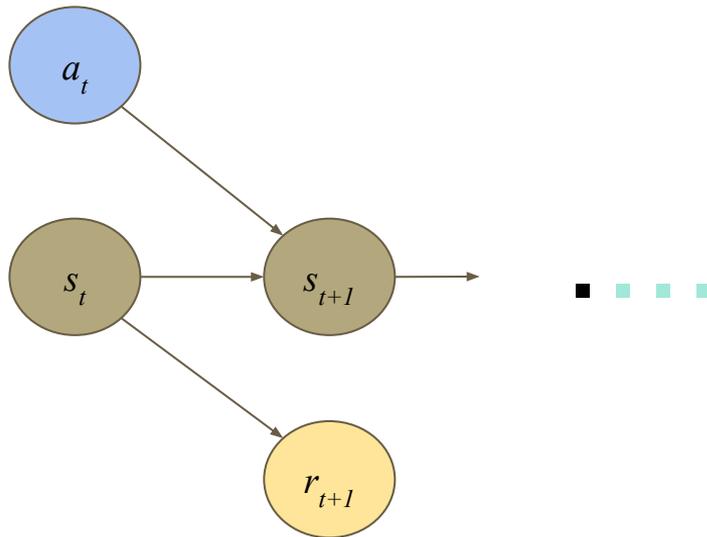
S : set of states

A : set of actions

T : transitions

R : reward

γ : discount factor



Markov Decision Processes (MDP)

$$M = (S, A, T, R, \gamma)$$

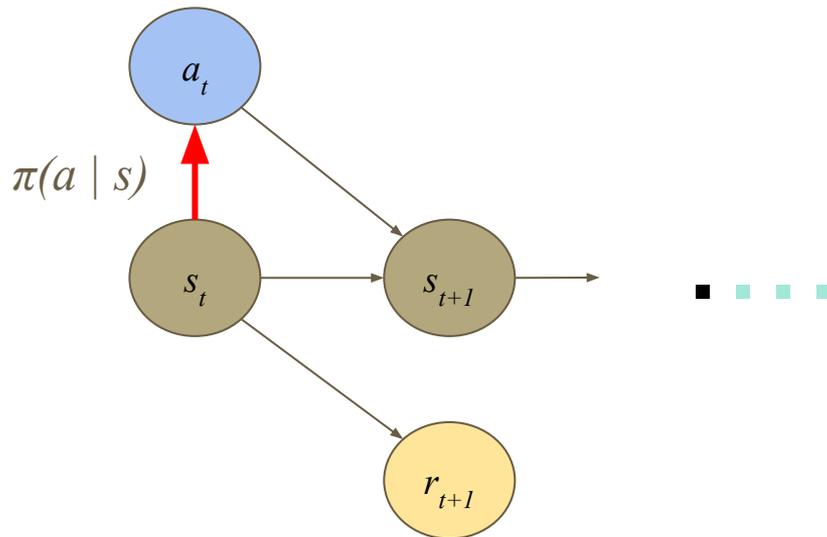
S : set of states

A : set of actions

T : transitions

R : reward

γ : discount factor

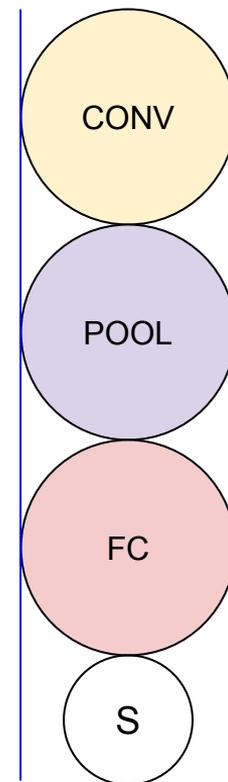


Objective: Find a policy $\pi(a | s)$ which maximizes total discounted reward

State: Network architecture

Set of network layers and their parameters

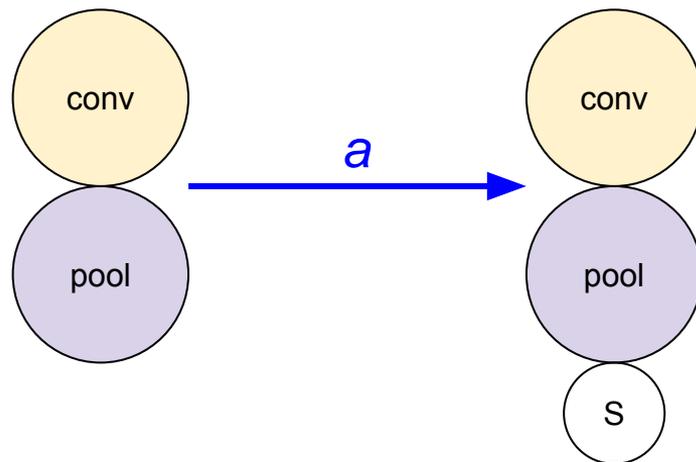
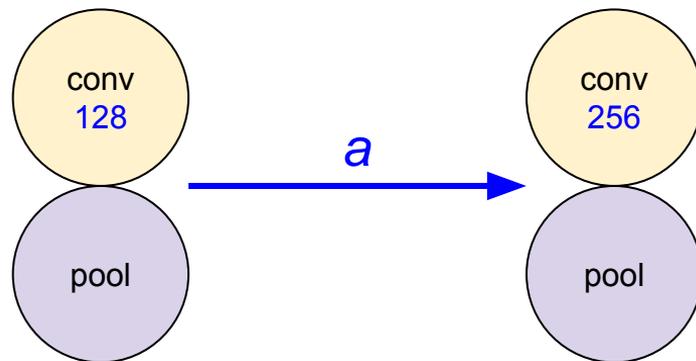
Layer Type	Layer Parameters	Parameter Values
Convolution (C)	$i \sim$ Layer depth $f \sim$ Receptive field size $\ell \sim$ Stride $d \sim$ # receptive fields $n \sim$ Representation size	< 12 Square. $\in \{1, 3, 5\}$ Square. Always equal to 1 $\in \{64, 128, 256, 512\}$ $\in \{(\infty, 8), (8, 4), (4, 1)\}$
Pooling (P)	$i \sim$ Layer depth $(f, \ell) \sim$ (Receptive field size, Strides) $n \sim$ Representation size	< 12 Square. $\in \{(5, 3), (3, 2), (2, 2)\}$ $\in \{(\infty, 8), (8, 4) \text{ and } (4, 1)\}$
Fully Connected (FC)	$i \sim$ Layer depth $n \sim$ # consecutive FC layers $d \sim$ # neurons	< 12 < 3 $\in \{512, 256, 128\}$
Termination State	$s \sim$ Previous State $t \sim$ Type	Global Avg. Pooling/Softmax



Action: Network Modification

Modify existing layers, add new layers, or terminate

- ϵ -Greedy choice rule
- Only trains terminating networks
- Transition function limits inappropriate networks



Reward: Classification Accuracy

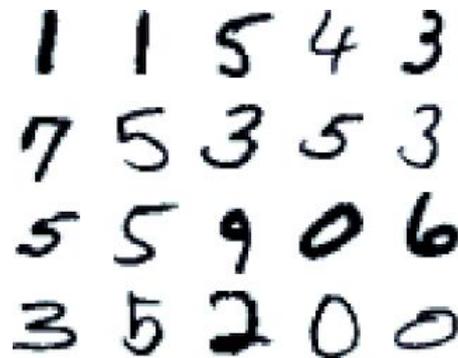
Performance at standard classification tasks



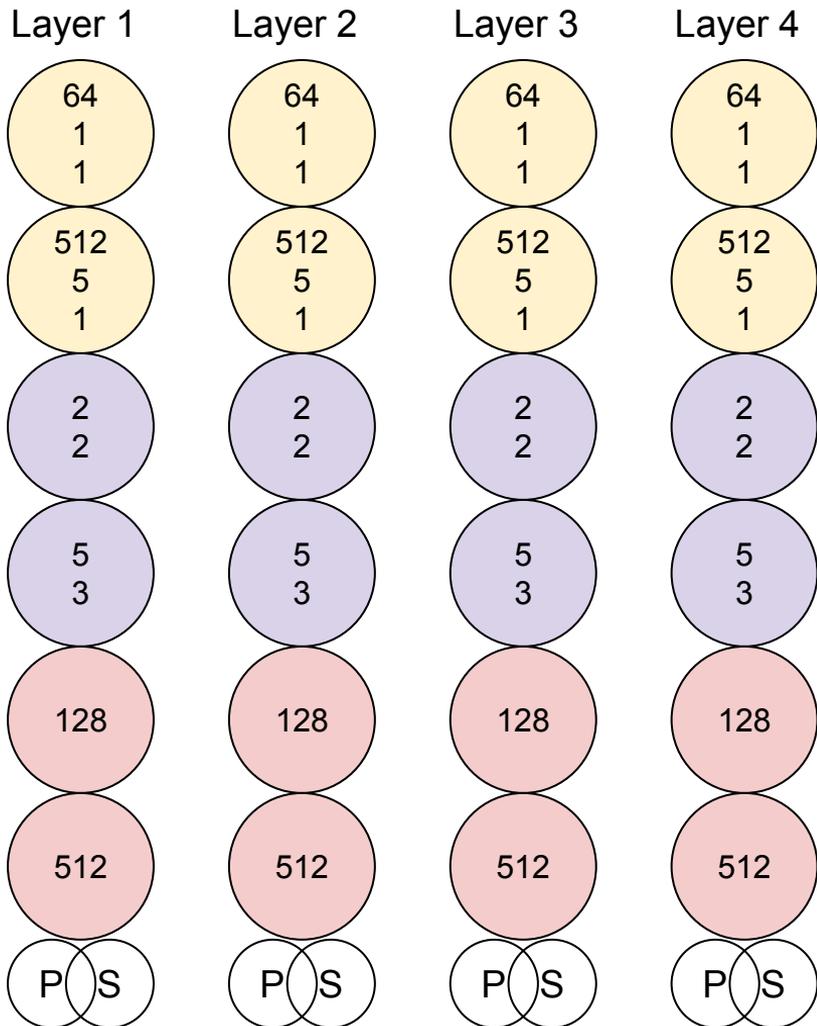
SVHN

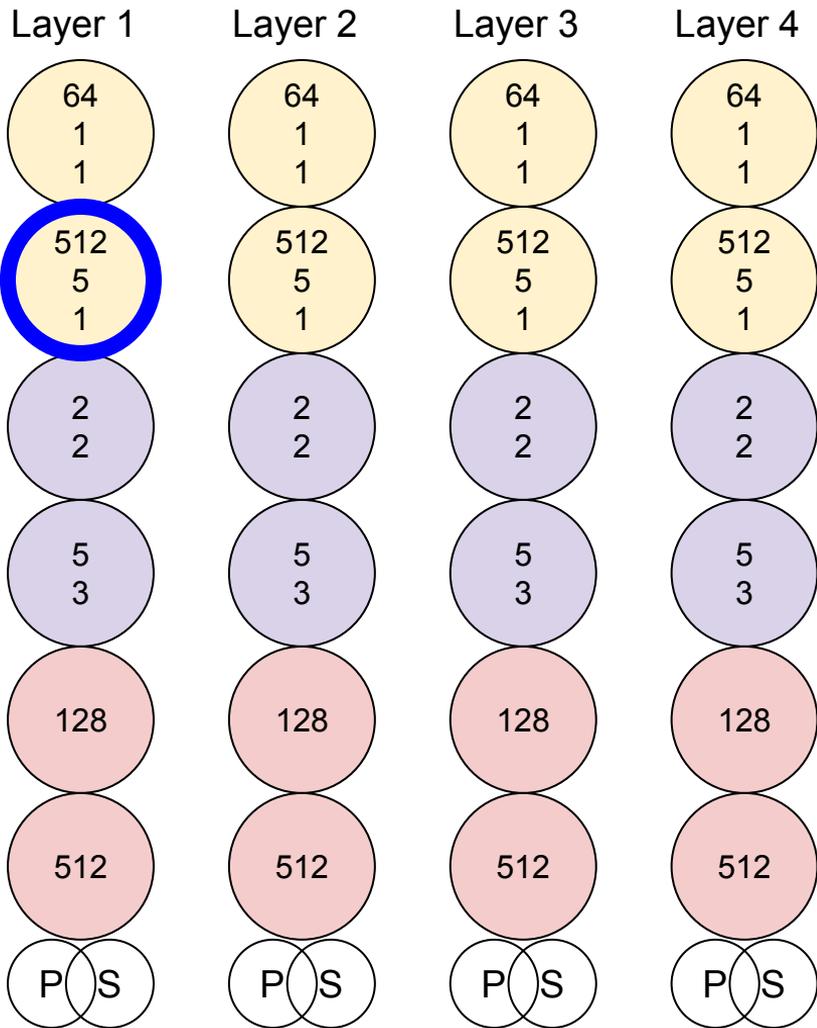


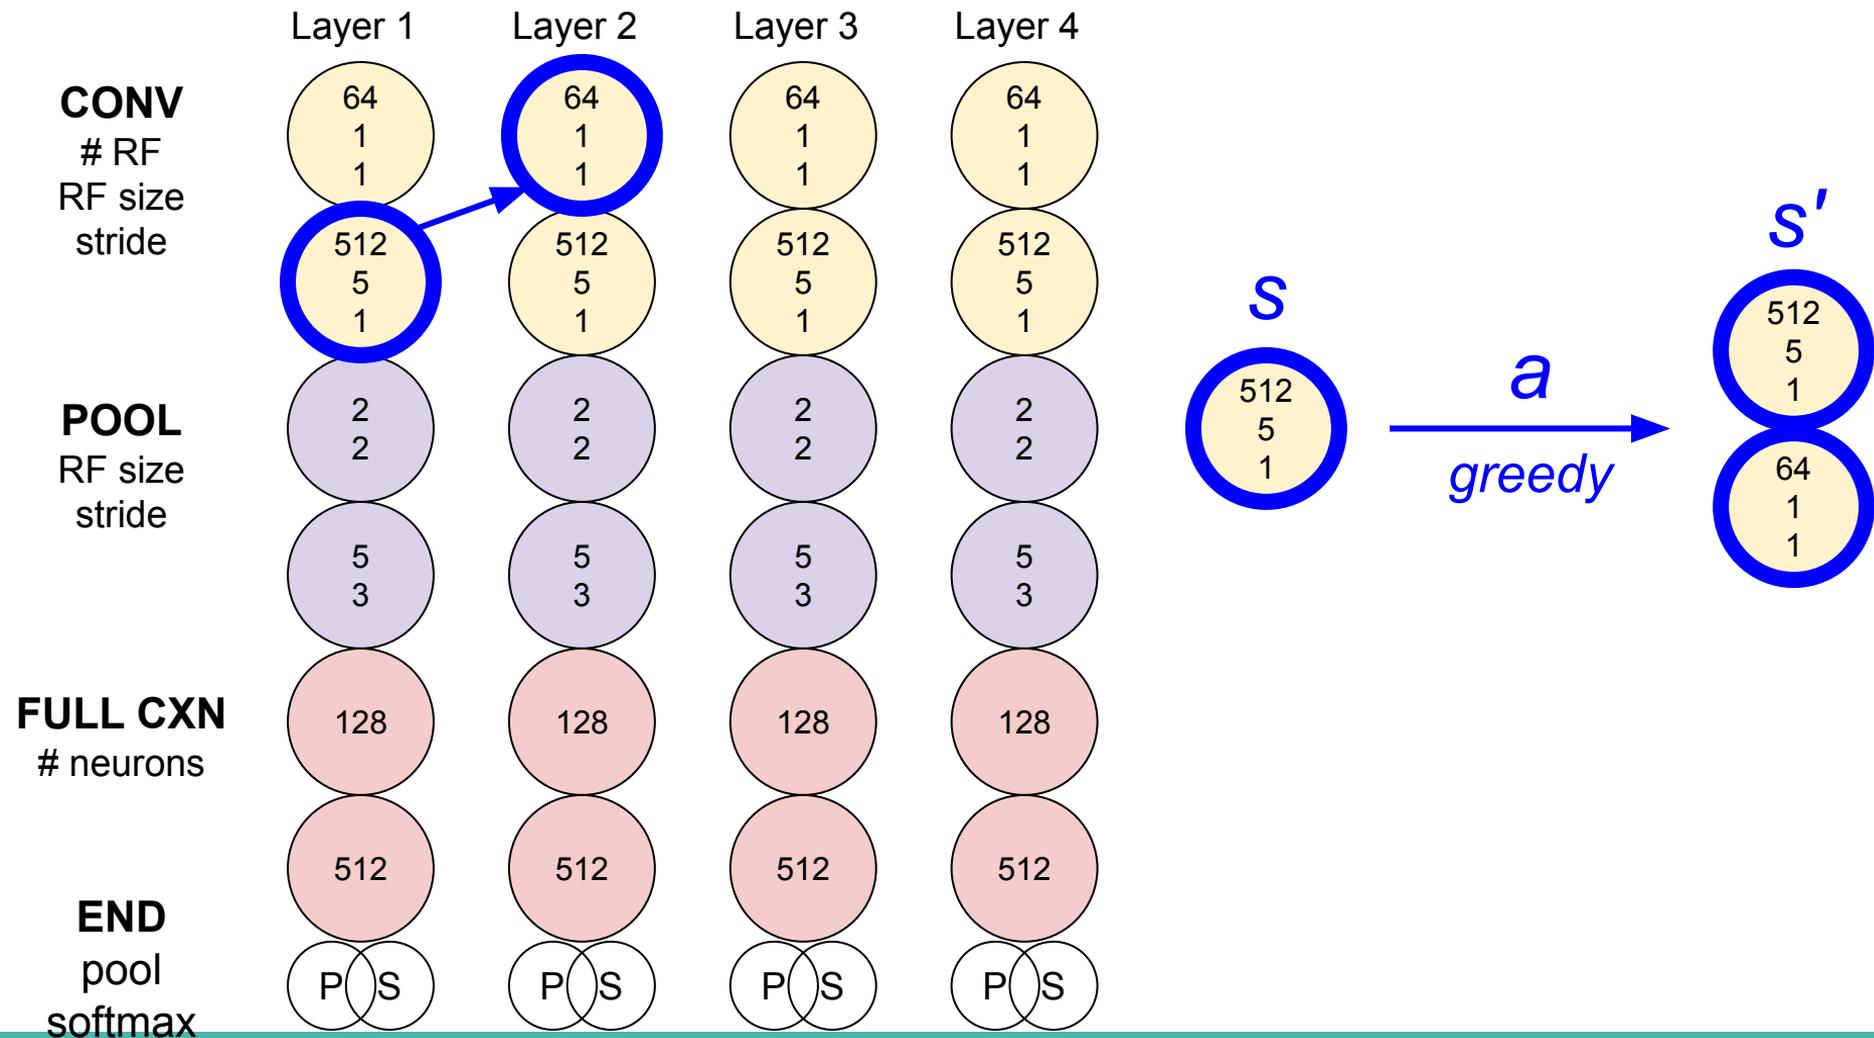
CIFAR-10

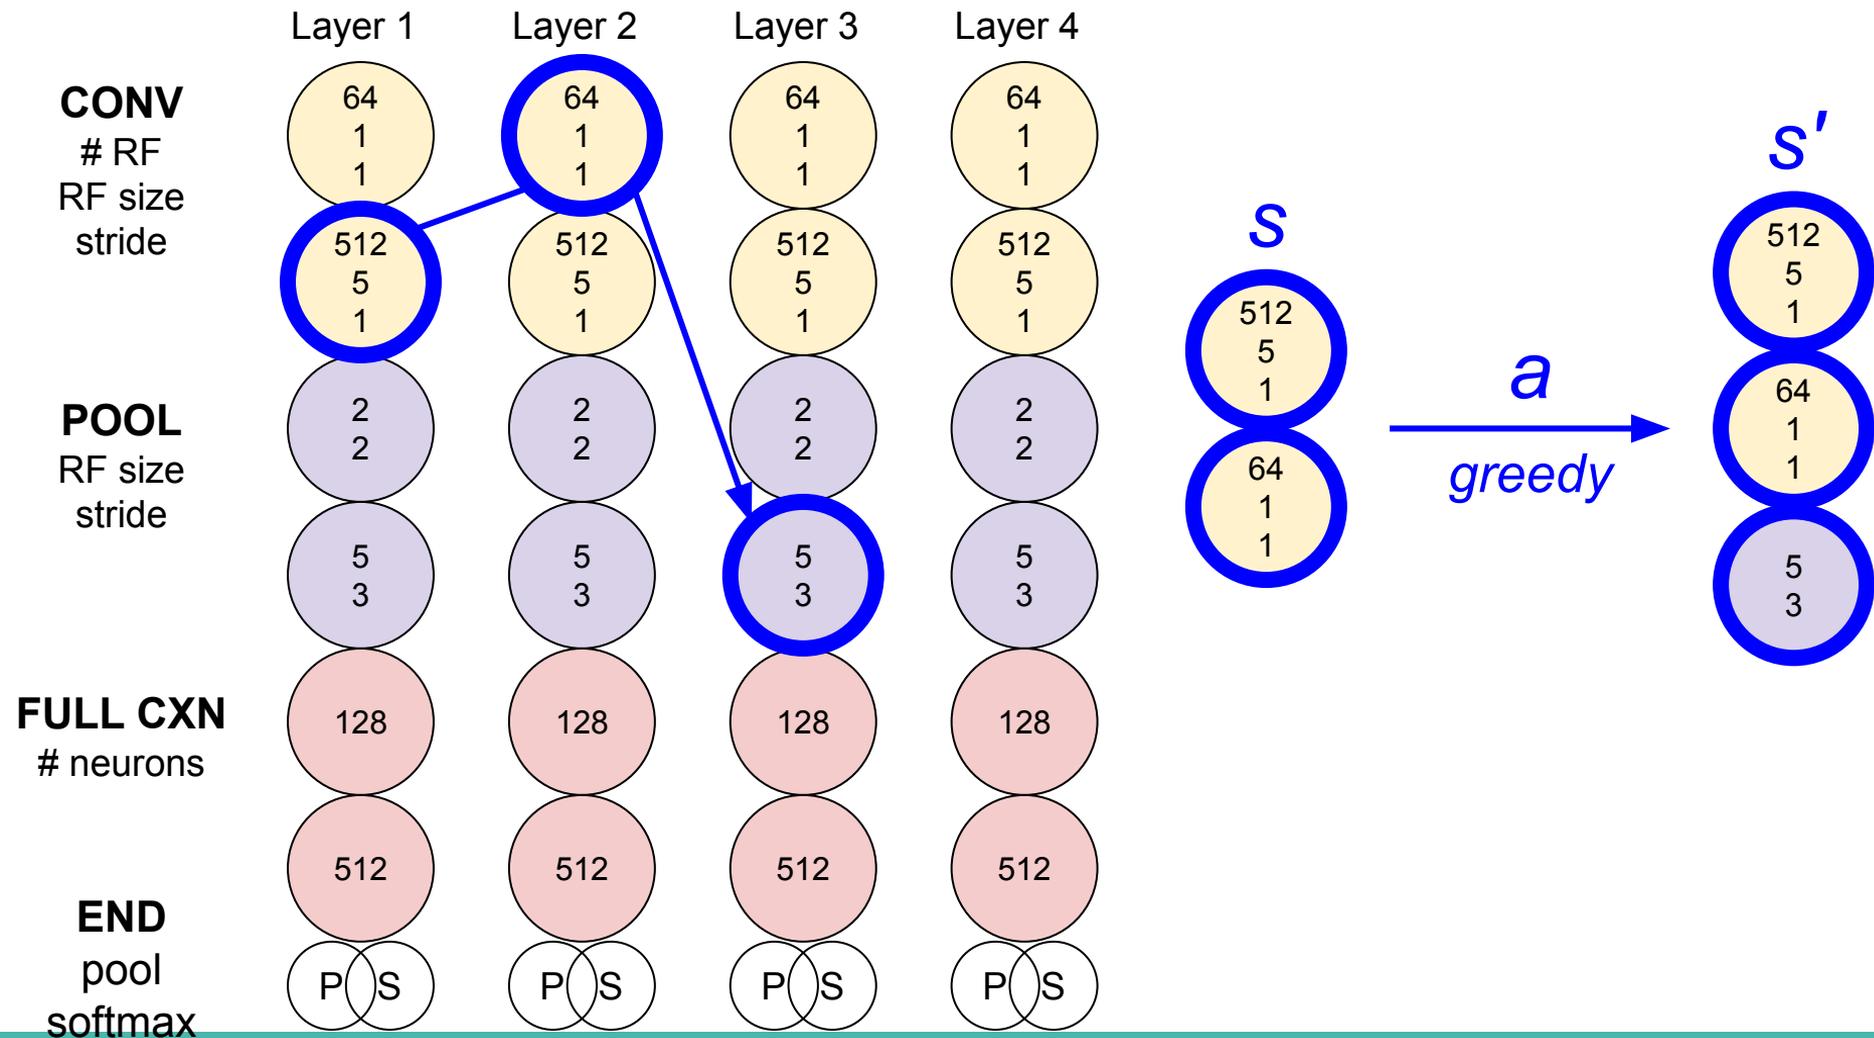


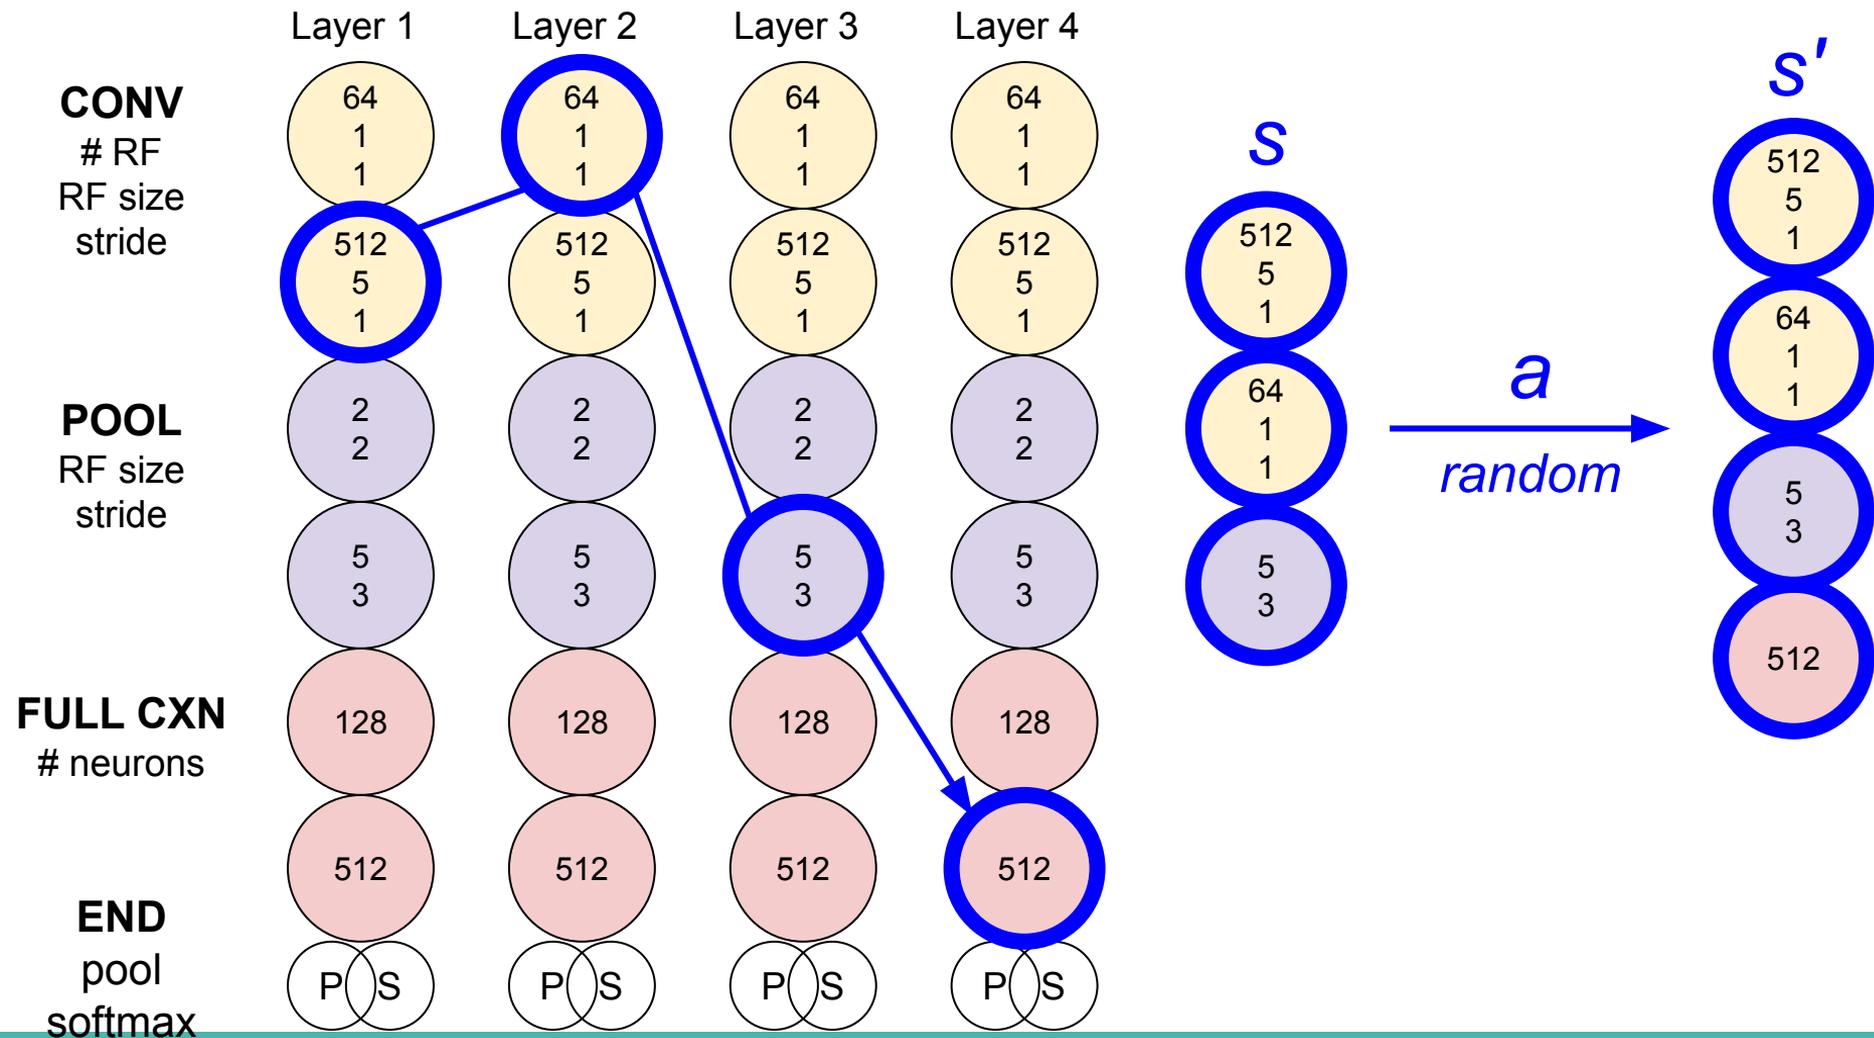
MNIST

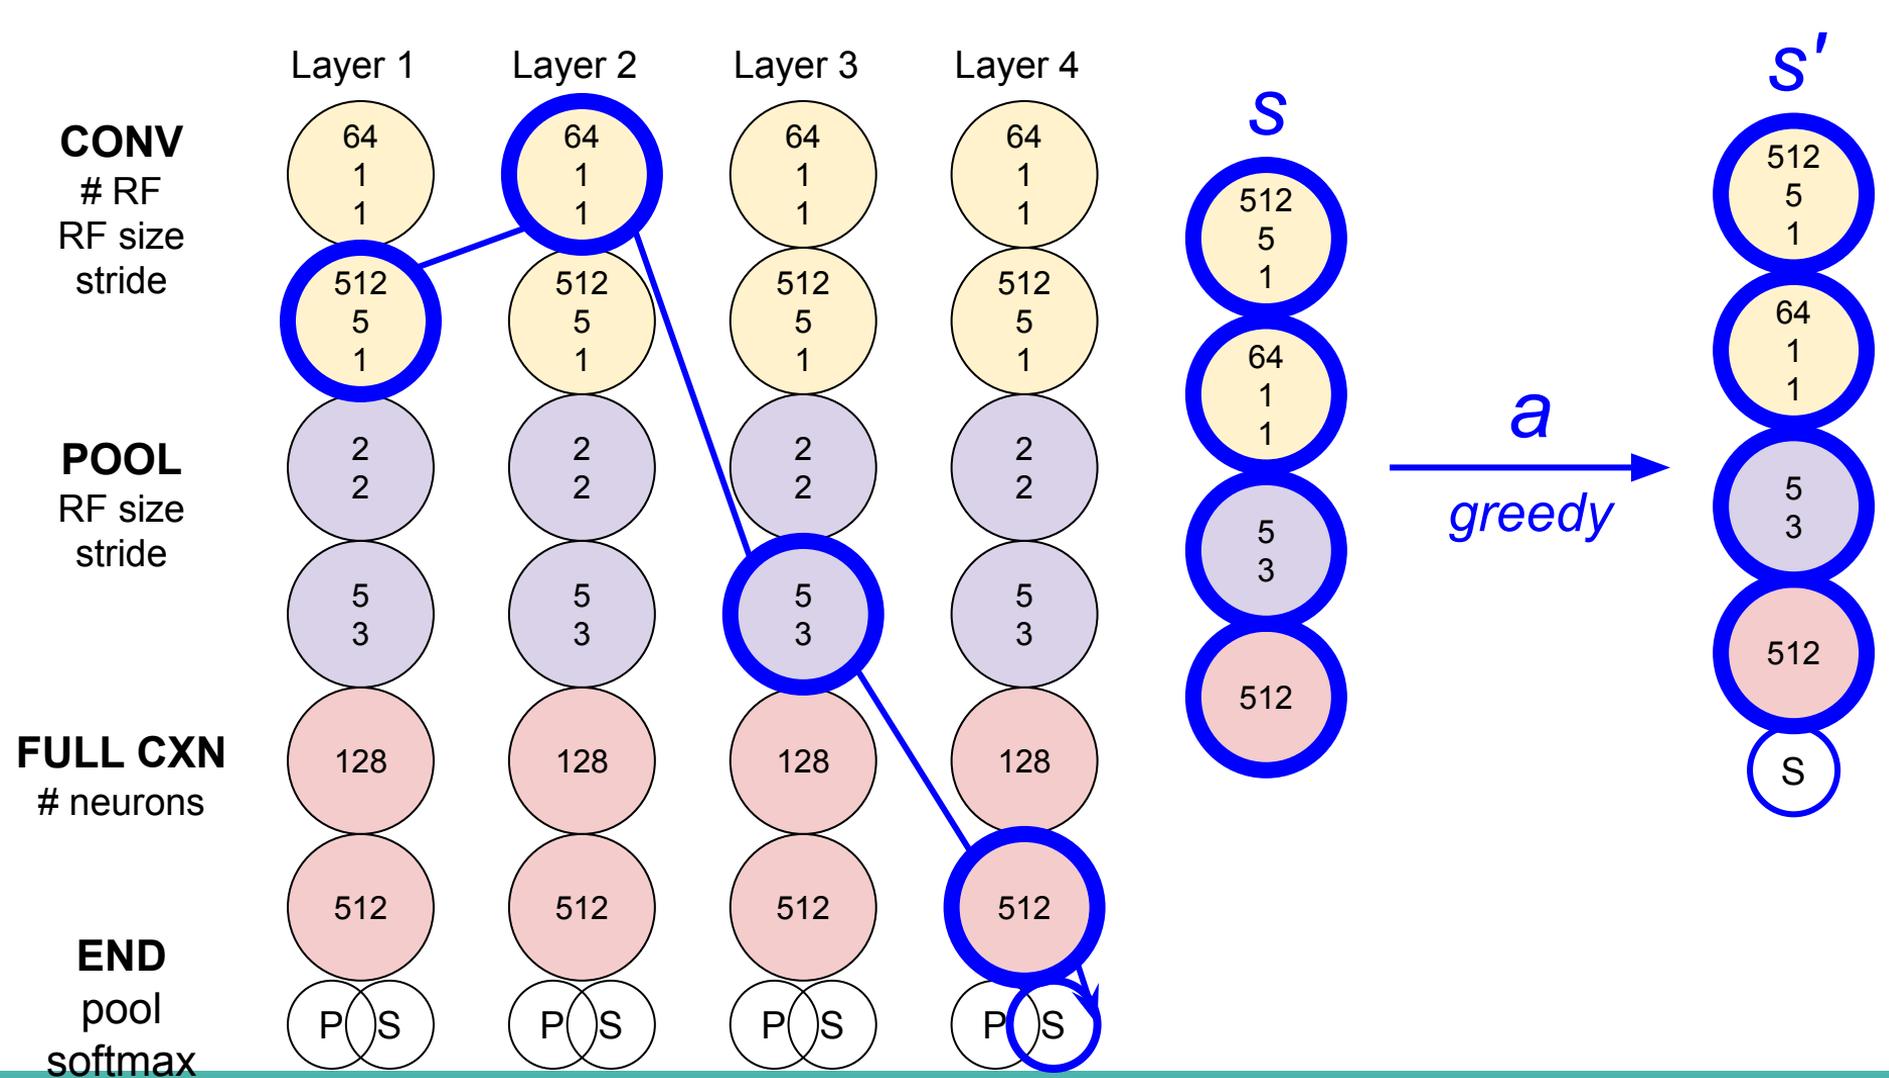




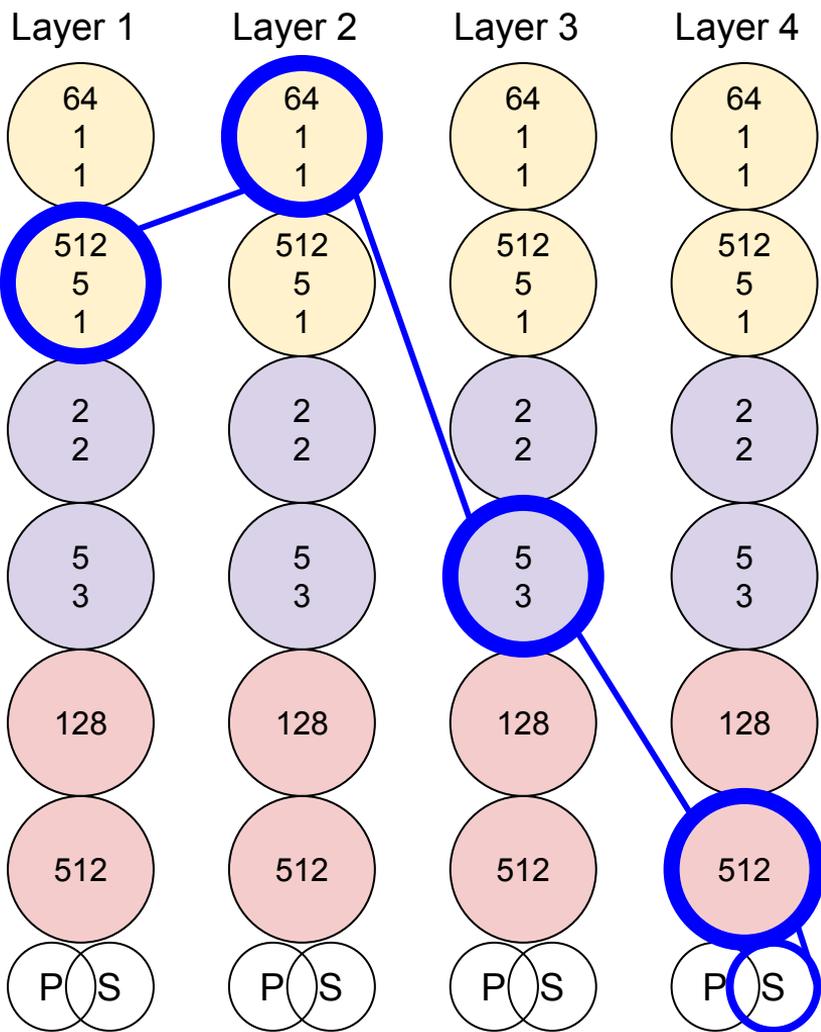








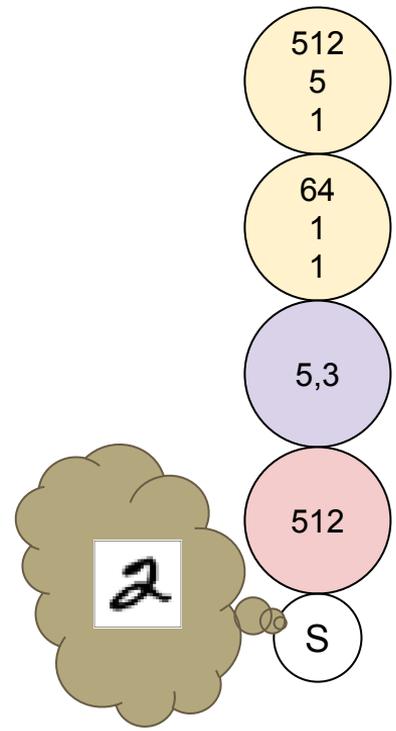
CONV
RF
RF size
stride



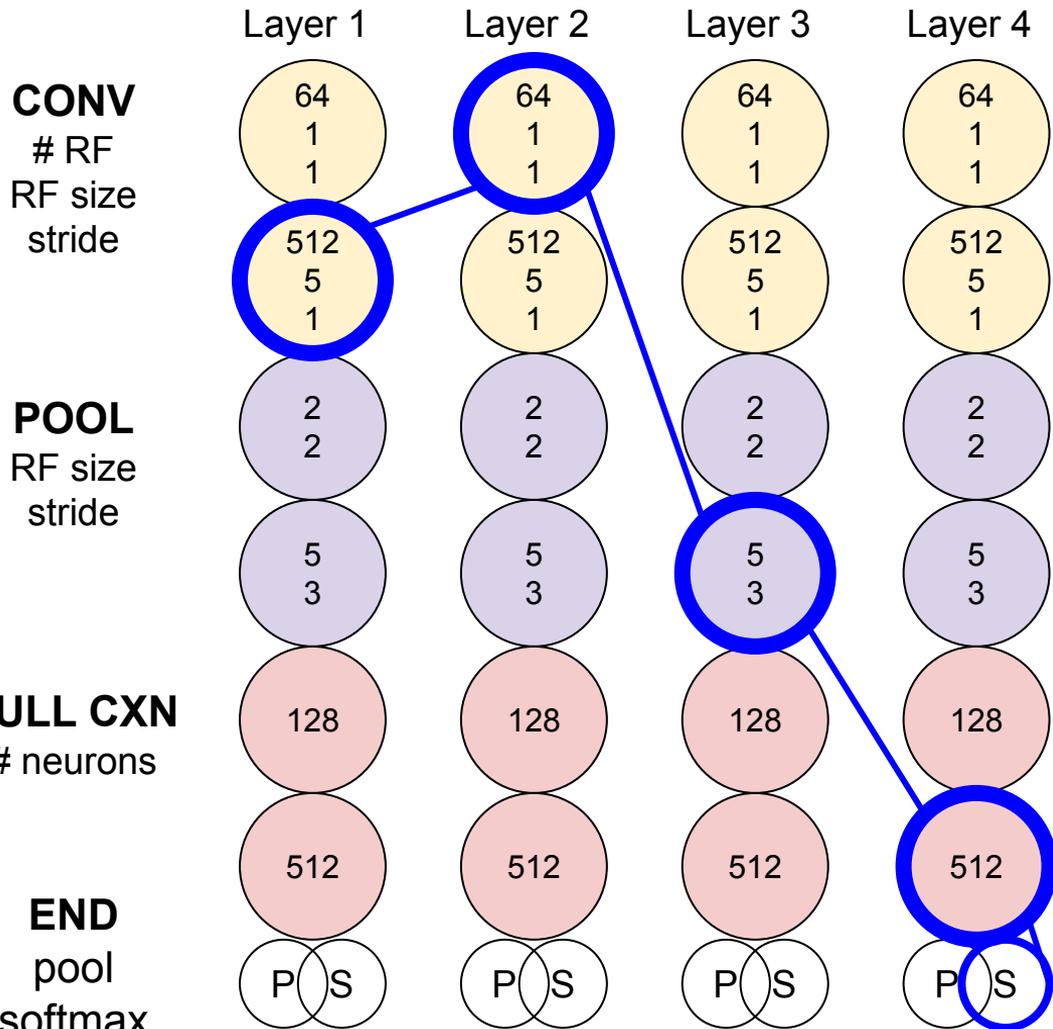
POOL
RF size
stride

FULL CXN
neurons

END
pool
softmax

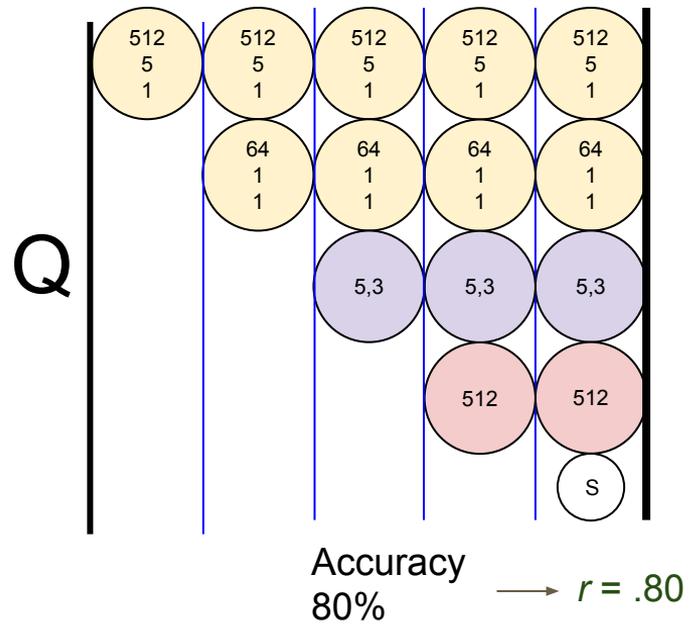


Accuracy
80%



$$Q(s, a) := Q(s, a) + \alpha \left[r + \gamma \max_{a'} Q(s', a') - Q(s, a) \right]$$

(r = accuracy)



Reinforcement Learning: Q-Learning (Baker et al. 2016)

Method	CIFAR-10	SVHN	MNIST	CIFAR-100
DropConnect (Wan et al., 2013)	9.32	1.94	0.57	-
DSN (Lee et al., 2015)	8.22	1.92	0.39	34.57
R-CNN (Liang & Hu, 2015)	7.72	1.77	0.31	31.75
MetaQNN (ensemble)	7.32	2.06	0.32	-
MetaQNN (top model)	6.92	2.28	0.44	27.14*
Resnet(110) (He et al., 2015)	6.61	-	-	-
Resnet(1001) (He et al., 2016)	4.62	-	-	22.71
ELU (Clevert et al., 2015)	6.55	-	-	24.28
Tree+Max-Avg (Lee et al., 2016)	6.05	1.69	0.31	32.37

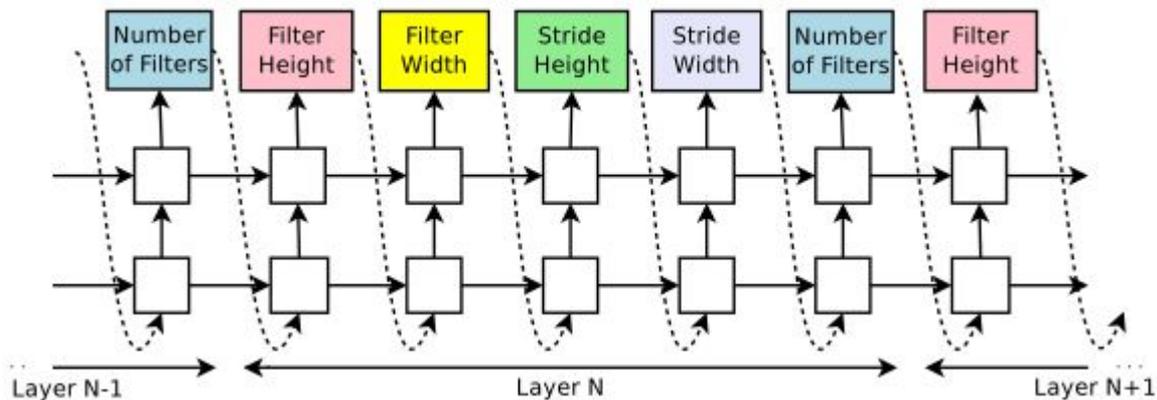
Reinforcement Learning: Policy Gradient (Zoph&Le 2016)

REINFORCE: Learn a policy to optimize the reward (e.g. accuracy)

$$J(\theta_c) = E_{P(a_{1:T}; \theta_c)}[R]$$
$$\nabla_{\theta_c} J(\theta_c) = \sum_{t=1}^T E_{P(a_{1:T}; \theta_c)} [\nabla_{\theta_c} \log P(a_t | a_{(t-1):1}; \theta_c) R]$$

Reinforcement Learning: Policy Gradient (Zoph&Le 2016)

REINFORCE: Learn a policy to optimize the reward (e.g. accuracy)



Reinforcement Learning: Policy Gradient (Zoph&Le 2016)

REINFORCE: Learn a policy to optimize the reward (e.g. accuracy)

Model	Depth	Parameters	Error rate (%)
Network in Network (Lin et al., 2013)	-	-	8.81
All-CNN (Springenberg et al., 2014)	-	-	7.25
Deeply Supervised Net (Lee et al., 2015)	-	-	7.97
Highway Network (Srivastava et al., 2015)	-	-	7.72
Scalable Bayesian Optimization (Snoek et al., 2015)	-	-	6.37
FractalNet (Larsson et al., 2016)	21	38.6M	5.22
with Dropout/Drop-path	21	38.6M	4.60
ResNet (He et al., 2016a)	110	1.7M	6.61
ResNet (reported by Huang et al., (2016c))	110	1.7M	6.41
ResNet with Stochastic Depth (Huang et al., 2016c)	110 1202	1.7M 10.2M	5.23 4.91
Wide ResNet (Zagoruyko & Komodakis, 2016)	16 28	11.0M 36.5M	4.81 4.17
ResNet (pre-activation) (He et al., 2016b)	164 1001	1.7M 10.2M	5.46 4.62
DenseNet ($L = 40, k = 12$) (Huang et al., 2016a)	40	1.0M	5.24
DenseNet ($L = 100, k = 12$) (Huang et al., 2016a)	100	7.0M	4.10
DenseNet ($L = 100, k = 24$) (Huang et al., 2016a)	100	27.2M	3.74
DenseNet-BC ($L = 100, k = 40$) (Huang et al., 2016b)	190	25.6M	3.46
Neural Architecture Search v1 no stride or pooling	15	4.2M	5.50
Neural Architecture Search v2 predicting strides	20	2.5M	6.01
Neural Architecture Search v3 max pooling	39	7.1M	4.47
Neural Architecture Search v3 max pooling + more filters	39	37.4M	3.65

Exploration-Exploitation Problem

Reinforcement learning optimizes policy based on its previous experience

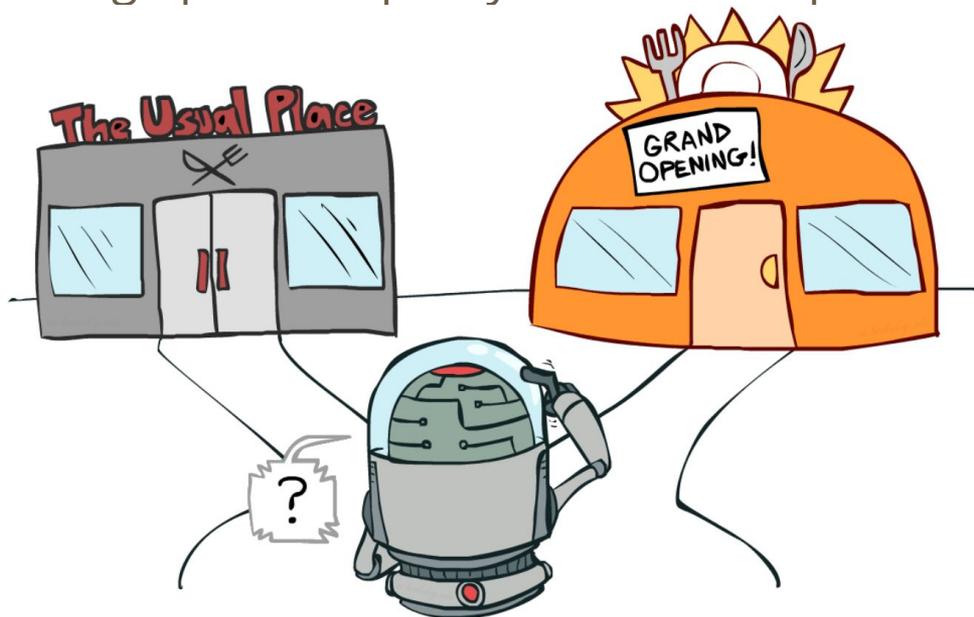


Figure from http://ai.berkeley.edu/lecture_slides.html

Problem 1: Epsilon scheduling

Hyperparameters of Q-learning

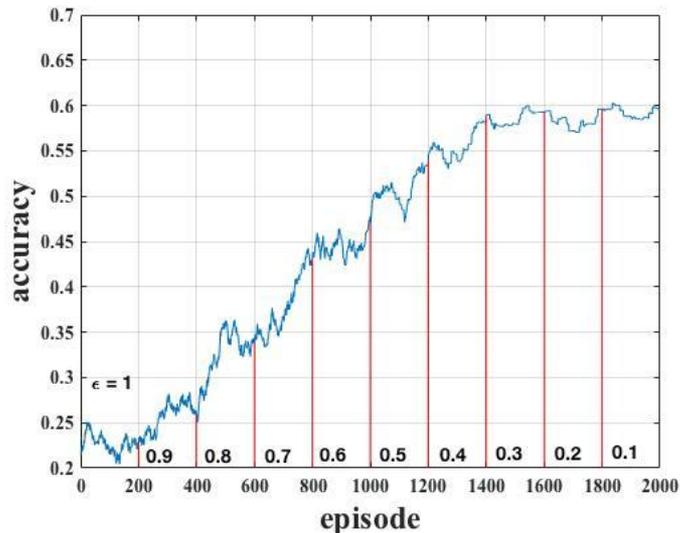
1. ϵ_t : Exploration probability for every episode

ϵ	1.0	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1
# Models Trained	1500	100	100	100	150	150	150	150	150	150

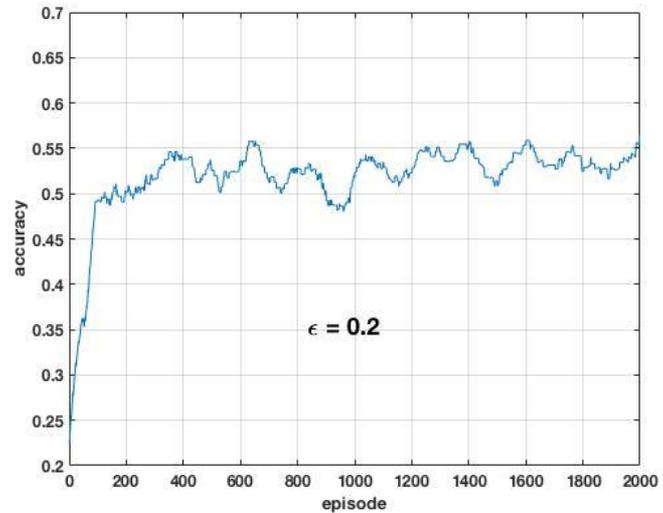
2. α : Learning rate alpha 0.01

The whole goal was to get rid of hyperparameters!

Results: Q-Learning with/without epsilon scheduling



epsilon scheduling

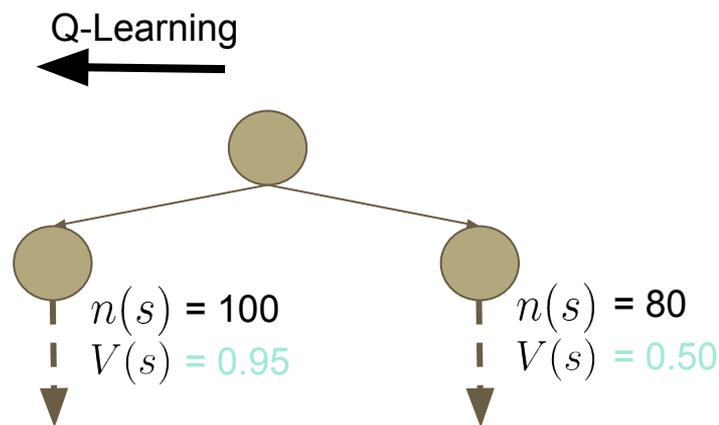
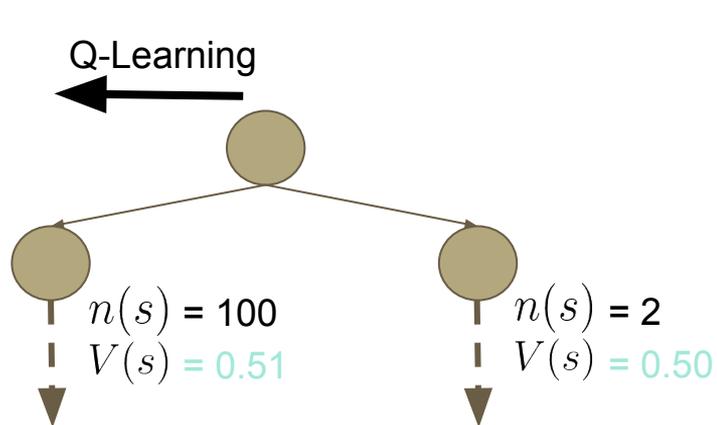


epsilon = 0.2

Problem 2: Constant exploration chance for all states

$V(s)$: value estimate of state s

$n(s)$: number of visits to a state s



Exploration chance should be set for each state rather than for each episode

1. No Epsilon Scheduling Required

$$\pi(s) = \max_{a \in A} \underbrace{\bar{X}(s')}_{\text{Exploitation}} + C_p \underbrace{\sqrt{\frac{2 \ln n(s)}{n(s')}}}_{\text{Exploration}} \quad \text{where } s' = T(s, a)$$

$\bar{X}(s)$: average accuracy
 $n(s)$: number of visits to a state s
 C_p : hyperparameter

Exploitation-Exploration dilemma is automatically tuned by the number of visits

2. Exploration Factor for Each State

$$\pi(s) = \underbrace{\max_{a \in A} \bar{X}(s')}_{\text{Exploitation}} + C_p \underbrace{\sqrt{\frac{2 \ln n(s)}{n(s)}}}_{\text{Exploration}}$$

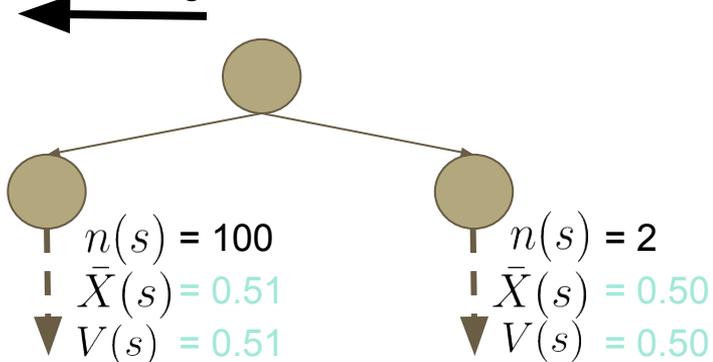
where $s' = T(s, a)$

$\bar{X}(s)$: average accuracy

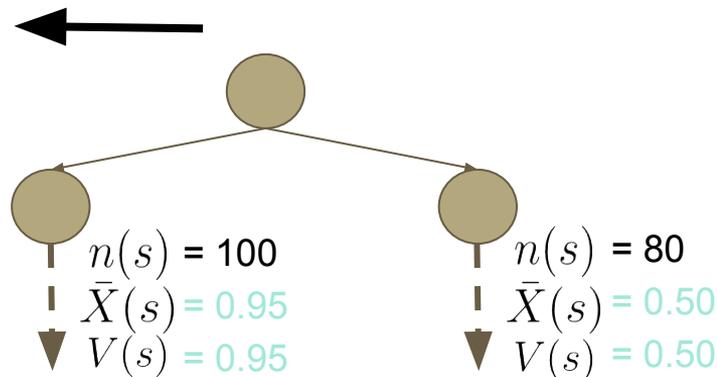
$n(s)$: number of visits to a state s

C_p : hyperparameter

Q-Learning



Q-Learning



2. Exploration Factor for Each State

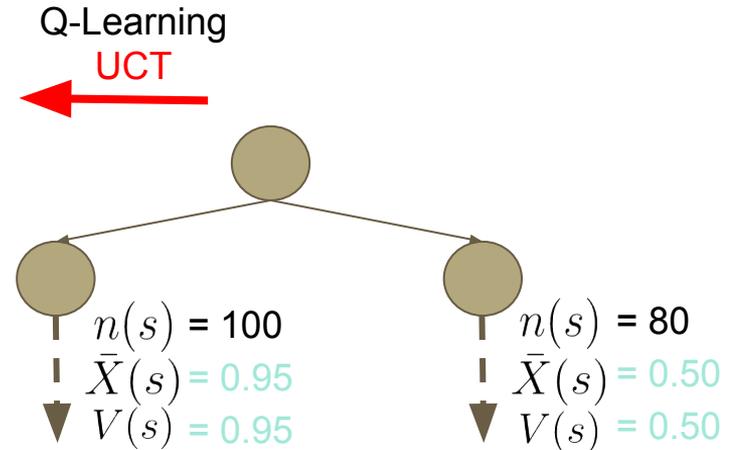
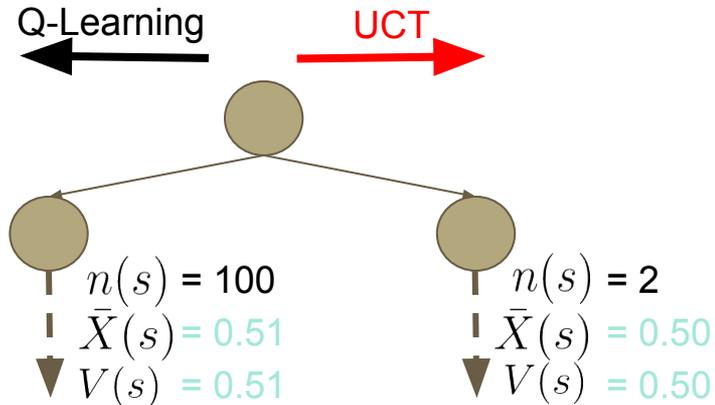
$$\pi(s) = \max_{a \in A} \underbrace{\bar{X}(s')}_{\text{Exploitation}} + C_p \underbrace{\sqrt{\frac{2 \ln n(s)}{n(s)}}}_{\text{Exploration}}$$

where $s' = T(s, a)$

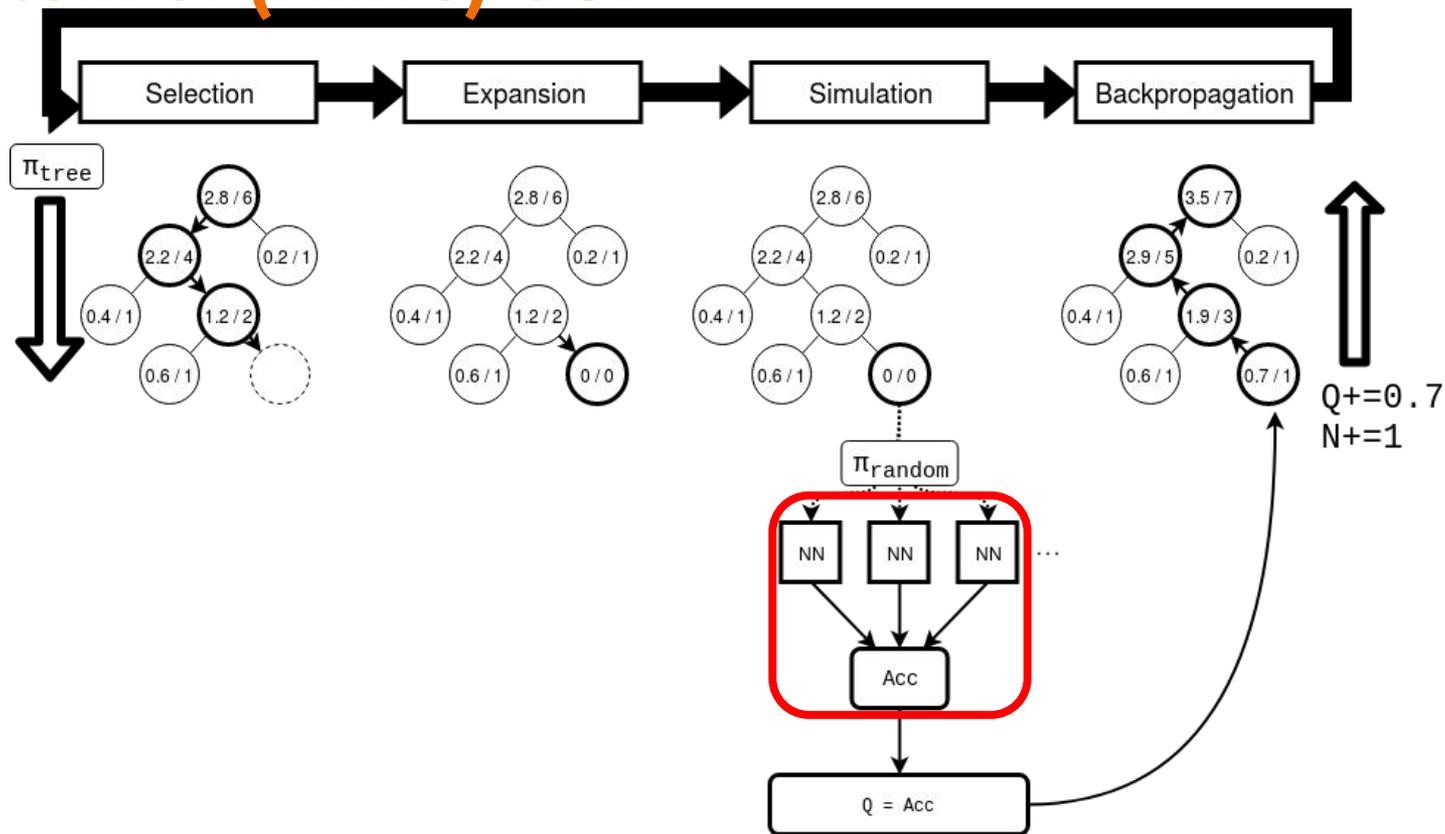
$\bar{X}(s)$: average accuracy

$n(s)$: number of visits to a state s

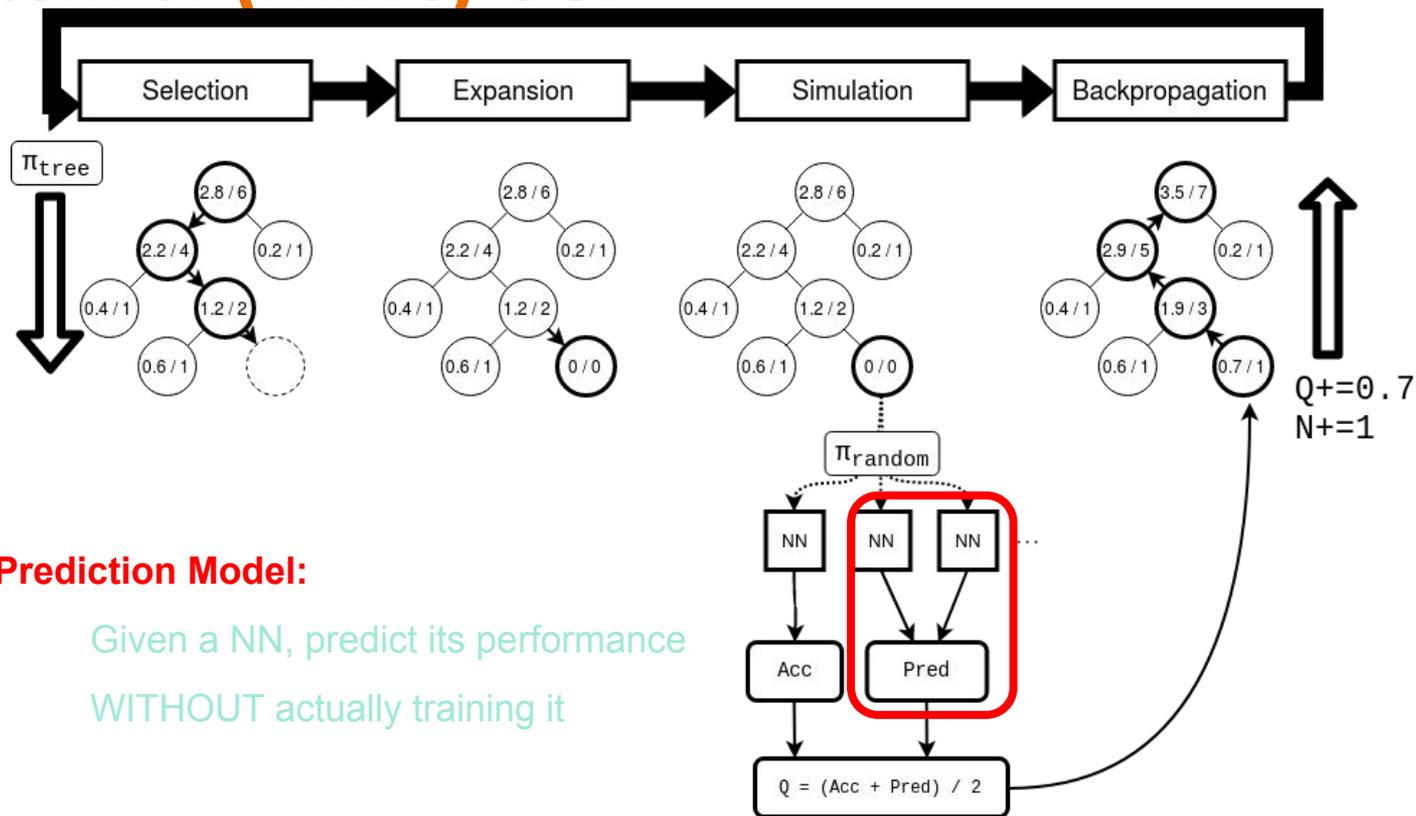
C_p : hyperparameter



Problem of (Naive) UCT



Problem of (Naive) UCT

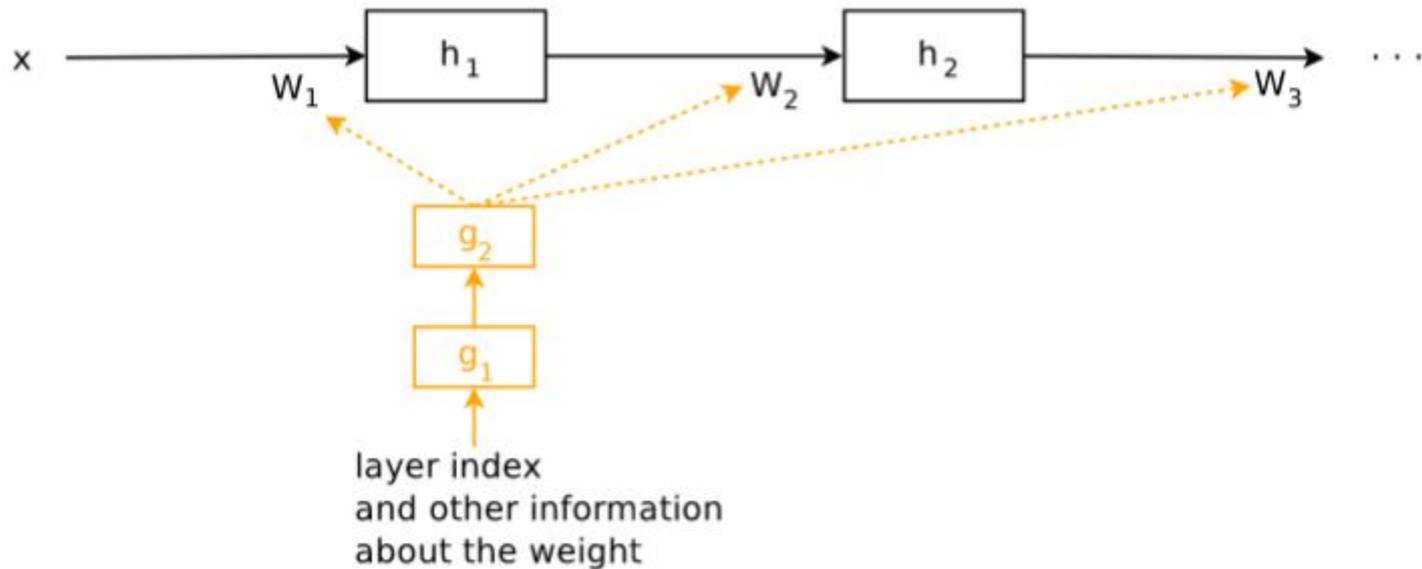


What is the State-of-the-Art of 2018?

Transfer learning

Reuse a weight of previously trained NN for a novel NN

HyperNetworks (Ha et al. 2016)



SMASH (Brock et al. 2017)

Transfer weights using HyperNetworks

Algorithm 1 SMASH

input Space of all candidate architectures, \mathbb{R}_c

Initialize HyperNet weights H

loop

Sample input minibatch x_i , random architecture c and architecture weights $W = H(c)$

Get training error $E_t = f_c(W, x_i) = f_c(H(c), x_i)$, backprop and update H

end loop

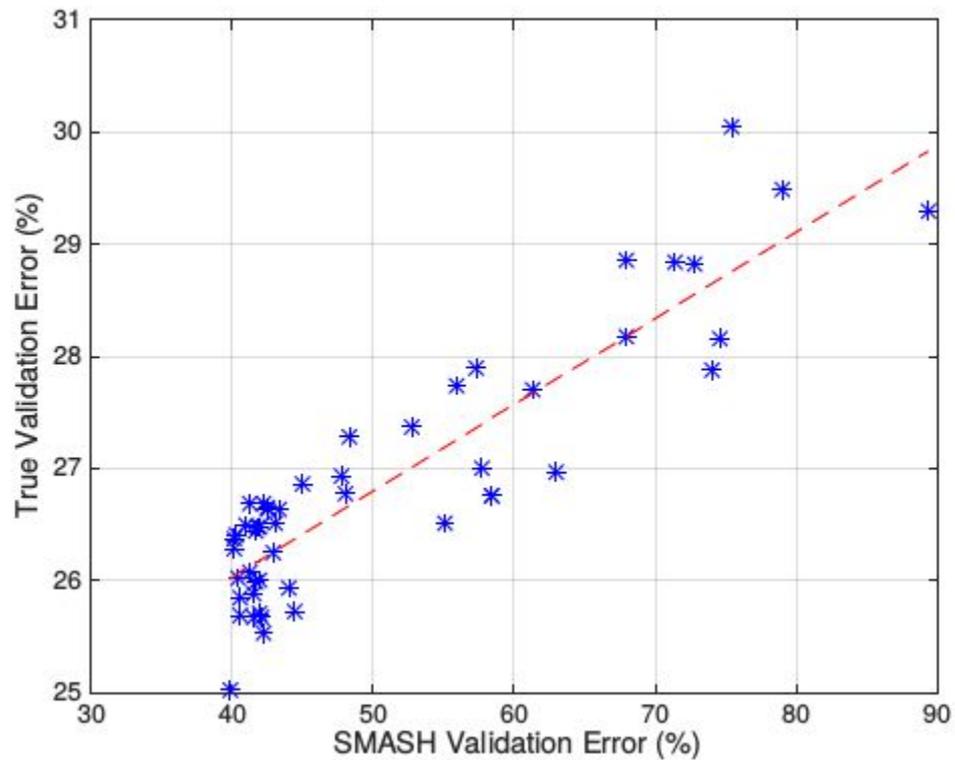
loop

Sample random c and evaluate error on validation set $E_v = f_c(H(c), x_v)$

end loop

Fix architecture and train normally with freely-varying weights W

SMASH (Brock et al. 2017)



SMASH (Brock et al. 2017)

Method	Depth	Params	C10+	C100+
FractalNet [20]	21	38.6M	5.22	23.30
with Dropout/Drop-path	21	38.6M	4.60	23.73
Wide ResNet [43]	16	11.0M	4.81	22.07
	28	36.5M	4.17	20.50
DenseNet-BC ($k = 24$) [15]	250	15.3M	3.62	17.60
DenseNet-BC ($k = 40$)	190	25.6M	3.46	17.18
Shake-Shake [11]	26	26.2M	2.86	15.85
Neural Architecture Search w/ RL [44]	39	32.0M	3.84	-
MetaQNN [3]	9	11.18M	6.92	27.14
Large-Scale Evolution [26]	-	5.4M	5.40	-
	-	40.4 M	-	23.7
CGP-CNN [38]	-	1.68M	5.98	-
SMASHv1	116	4.6M	5.53	22.07
SMASHv2	211	16M	4.03	20.60

Net2Net (Chen et al. 2016)

Traditional Workflow

Initial Design



Training



Rebuild the Model

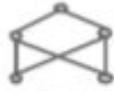


Training



Net2Net Workflow

Initial Design



Training



Reuse the Model



Net2Net Operator

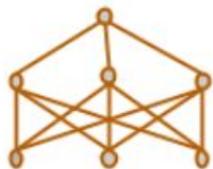


Training



Net2Net (Chen et al. 2016)

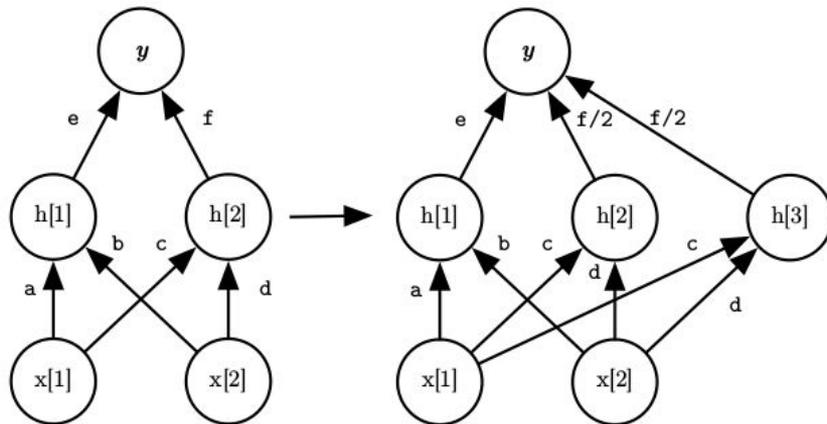
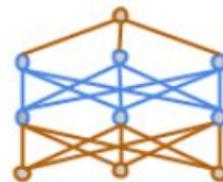
Original Model



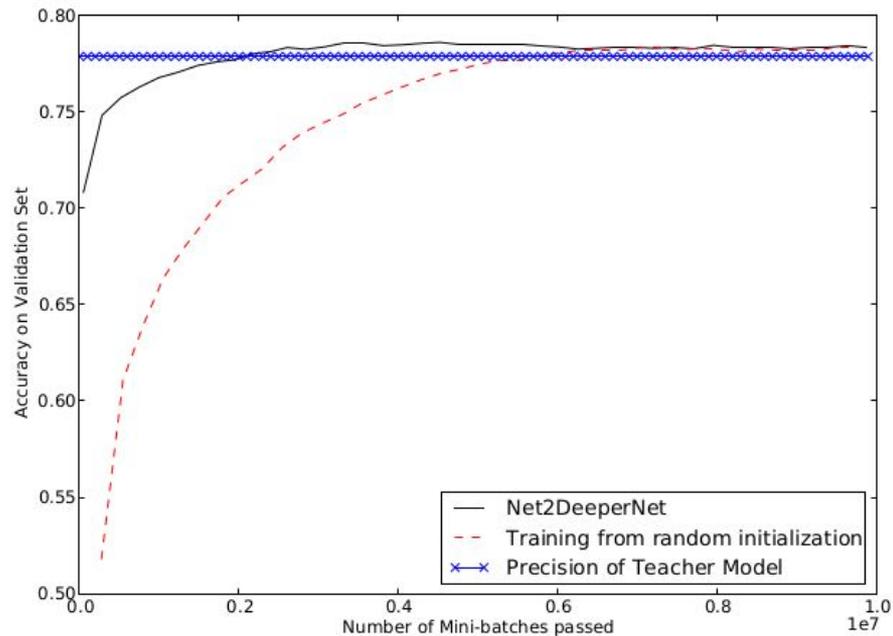
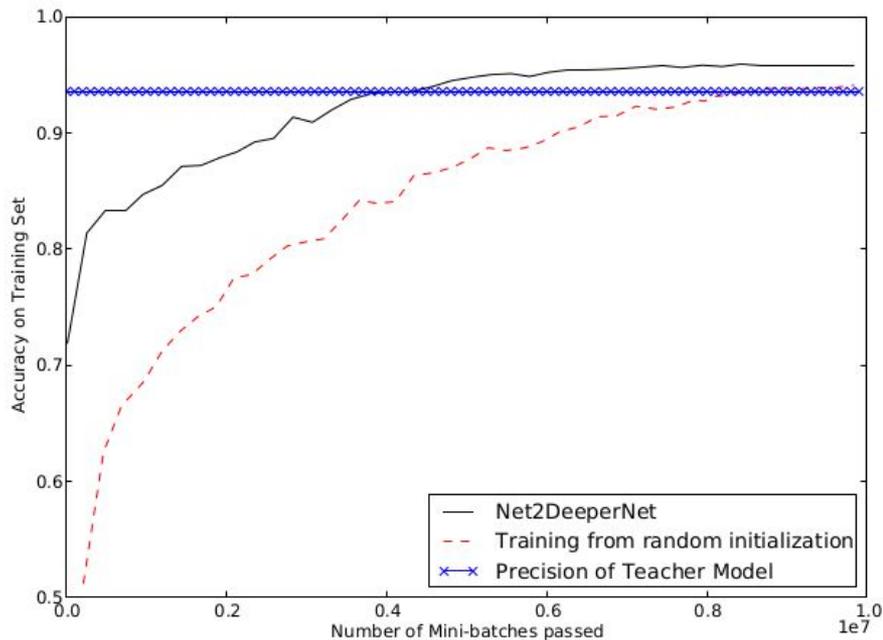
Layers that Initialized as Identity Mapping



A Deeper Model Contains Identity Mapping Initialized Layers



Net2Net (Chen et al. 2016)



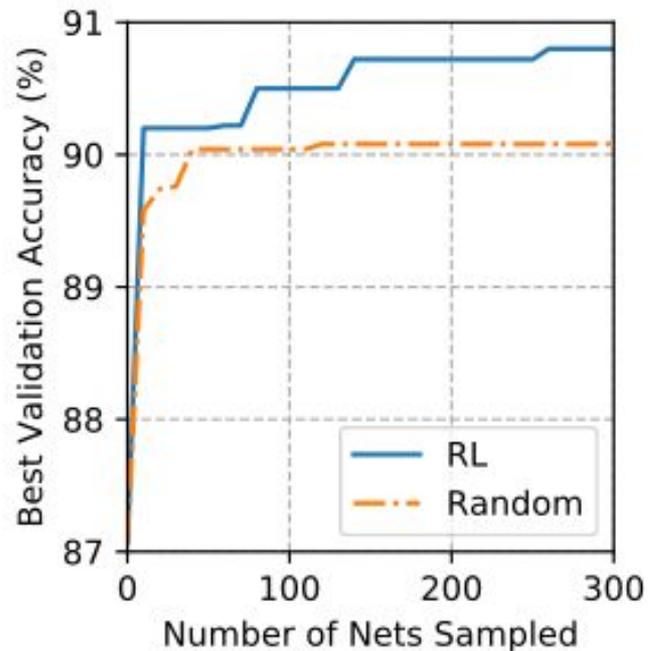
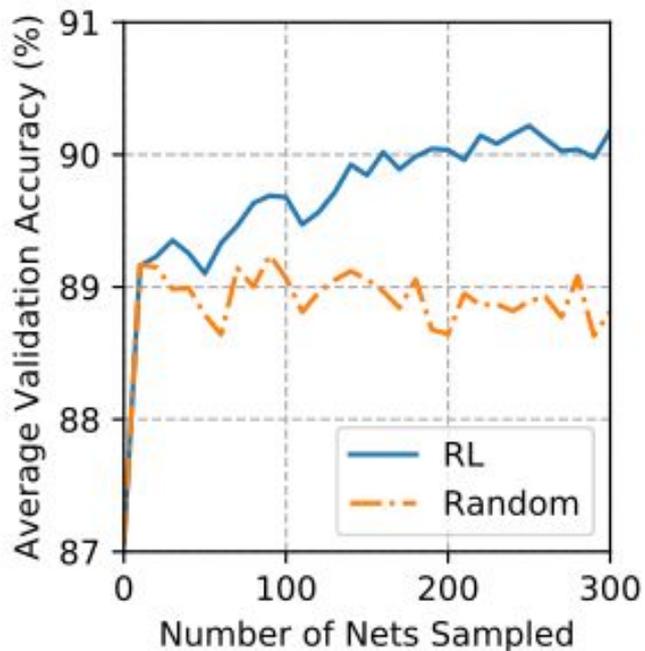
Network Transformation (Cai et al. 2017)

Net2Net applied to RL-based search

Table 3: Test error rate (%) comparison with state-of-the-art architectures.

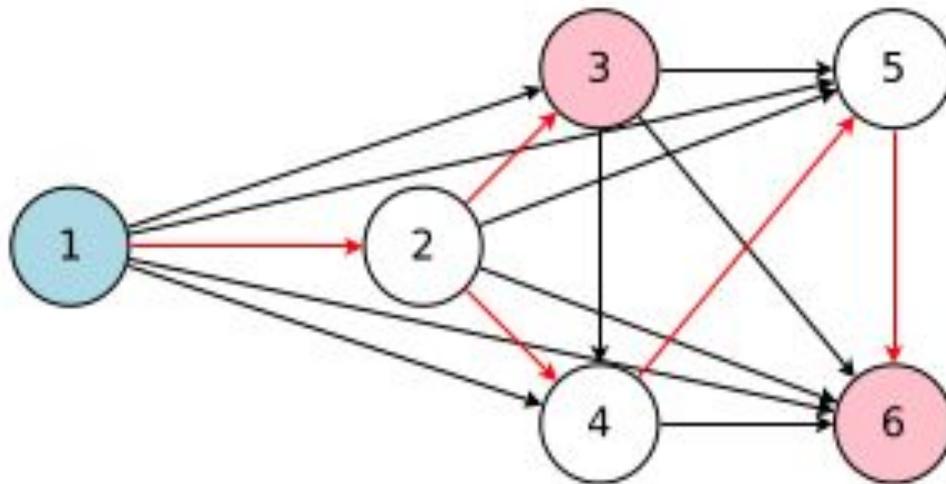
	Model	Depth	Params	C10+
human designed	ResNet (He et al. 2016a)	110	1.7M	6.61
	ResNet (stochastic depth) (Huang et al. 2017)	1202	10.2M	4.91
	Wide ResNet (Zagoruyko and Komodakis 2016)	16	11.0M	4.81
	Wide ResNet (Zagoruyko and Komodakis 2016)	28	36.5M	4.17
	ResNet (pre-activation) (He et al. 2016b)	1001	10.2M	4.62
	DenseNet ($L = 40, k = 12$) (Huang et al. 2017)	40	1.0M	5.24
	DenseNet-BC ($L = 100, k = 12$) (Huang et al. 2017)	100	0.8M	4.51
	DenseNet-BC ($L = 190, k = 40$) (Huang et al. 2017)	190	25.6M	3.46
auto designed	Large-Scale Evolution (250 GPUs)(Real et al. 2017)	-	5.4M	5.40
	NAS (predicting strides, 800 GPUs) (Zoph and Le 2017)	20	2.5M	6.01
	NAS (max pooling, 800 GPUs) (Zoph and Le 2017)	39	7.1M	4.47
	NAS (post-processing, 800 GPUs) (Zoph and Le 2017)	39	37.4M	3.65
	EAS (plain CNN, 5 GPUs)	20	23.4M	4.23

Network Transformation (Cai et al. 2017)



ENAS (Pham et al. 2018)

Sharing weights among child networks

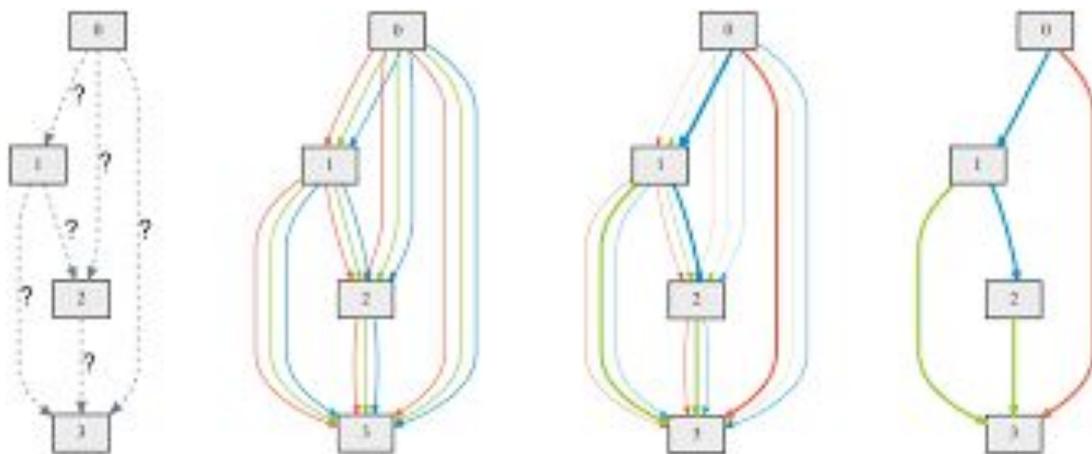


ENAS (Pham et al. 2018)

Sharing weights among child networks

Architecture	Additional Techniques	Params (million)	Test PPL
LSTM (Zaremba et al., 2014)	Vanilla Dropout	66	78.4
LSTM (Gal & Ghahramani, 2016)	VD	66	75.2
LSTM (Inan et al., 2017)	VD, WT	51	68.5
LSTM (Melis et al., 2017)	Hyper-parameters Search	24	59.5
LSTM (Yang et al., 2018)	VD, WT, ℓ_2 , AWD, MoC	22	57.6
LSTM (Merity et al., 2017)	VD, WT, ℓ_2 , AWD	24	57.3
LSTM (Yang et al., 2018)	VD, WT, ℓ_2 , AWD, MoS	22	56.0
RHN (Zilly et al., 2017)	VD, WT	24	66.0
NAS (Zoph & Le, 2017)	VD, WT	54	62.4
ENAS	VD, WT, ℓ_2	24	55.8

DARTS: Differentiable NN Search (Liu et al. 2018)



$$\bar{o}^{(i,j)}(x) = \sum_{o \in \mathcal{O}} \frac{\exp(\alpha_o^{(i,j)})}{\sum_{o' \in \mathcal{O}} \exp(\alpha_{o'}^{(i,j)})} o(x)$$

DARTS: Differentiable NN Search (Liu et al. 2018)

Architecture	Test Error (%)	Params (M)	Search Cost (GPU days)	Search Method
DenseNet-BC (Huang et al., 2017)	3.46	25.6	–	manual
NASNet-A + cutout (Zoph et al., 2017)	2.65	3.3	1800	RL
NASNet-A + cutout (Zoph et al., 2017) [†]	2.83	3.1	3150	RL
AmoebaNet-A + cutout (Real et al., 2018)	3.34 ± 0.06	3.2	3150	evolution
AmoebaNet-A + cutout (Real et al., 2018) [†]	3.12	3.1	3150	evolution
AmoebaNet-B + cutout (Real et al., 2018)	2.55 ± 0.05	2.8	3150	evolution
Hierarchical Evo (Liu et al., 2017b)	3.75 ± 0.12	15.7	300	evolution
PNAS (Liu et al., 2017a)	3.41 ± 0.09	3.2	225	SMBO
ENAS + cutout (Pham et al., 2018b)	2.89	4.6	0.5	RL
Random + cutout	3.49	3.1	–	–
DARTS (first order) + cutout	2.94	2.9	1.5	gradient-based
DARTS (second order) + cutout	2.83 ± 0.06	3.4	4	gradient-based

Gitgraph: Subgraph Mining (Bennani-Smires et al. 2018)

Search for Problem

Mine Frequent
Subgraphs

Define Search Space

Query

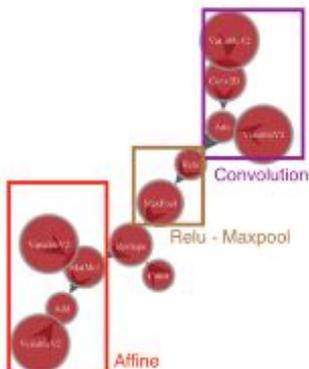
Machine +
Translation

Reinforcement +
Learning

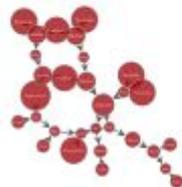
Imagenet

Knowledge Base

GitGraph



Probable
Starting
Points

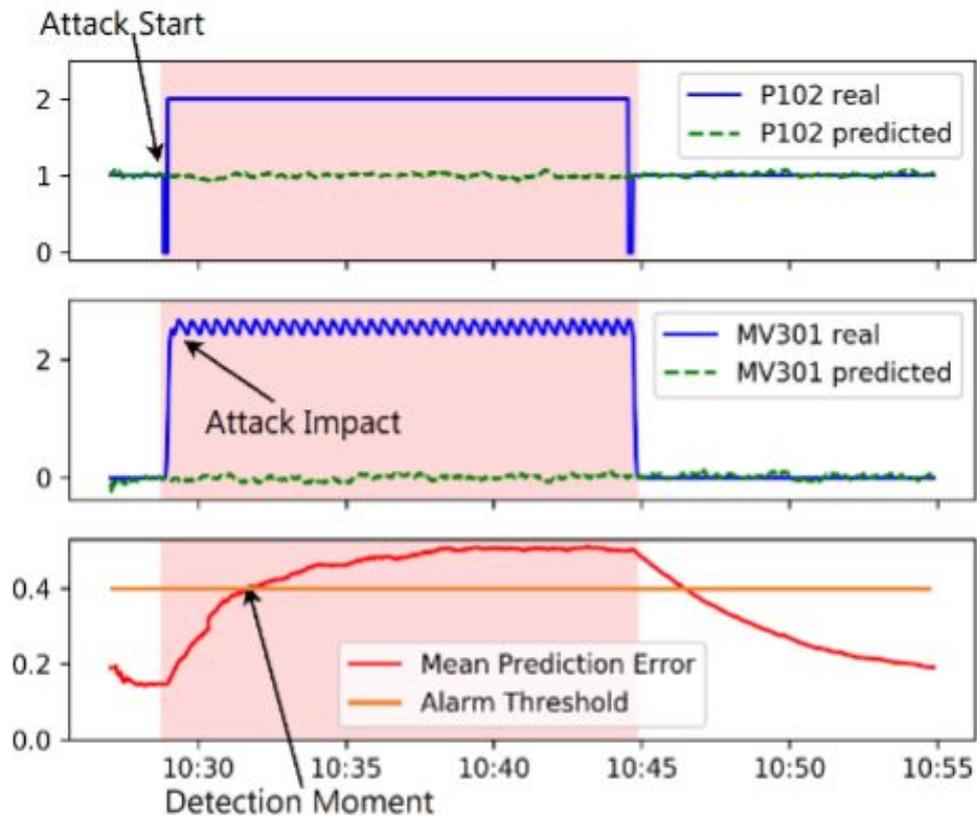


Problem
Specific
Toolbox



Learn Policy

Application to Anomaly Detection (Shalyga et al. 2018)



Model	NAB score	F_1	Precision	Recall
MLP	69.612	0.812	0.967	0.696
CNN	34.225	0.808	0.952	0.702
RNN	36.924	0.796	0.936	0.692
SVM*	-	0.796	0.925	0.699
DNN*	-	0.802	0.982	0.678

Architecture Search is...

