From Architecting Networks to Architecting Network Generators

Ross Girshick



facebook AI Research

Many slides adapted from Saining Xie

Before 2012 (... but also after 1990)

• Engineering features (SIFT, ...), learning weights (SVM, ...)



Engineered



Learned

2012 – now-ish

- Engineering features (SIFT, ...), learning weights (SVM, ...)
- Engineering networks (AlexNet, VGG, ...), learning features and weights



Millions of parameters

- Internal filters
- Classifier weights

Engineered

Learned

Now-ish - ???

- Engineering features (SIFT, ...), learning weights (SVM, ...)
- Engineering networks (AlexNet, VGG, ...), learning features and weights
- Engineering network generators, learning networks, features &



Engineered



Learned

• Think of a Python function:

def vgg_net(input_res, depth, conv1_out_channels):
... # code to make the model
return model

• Think of a Python function:

```
def vgg_net(input_res, depth, conv1_out_channels):
... # code to make the model
return model
```

• Input: some arguments (resolution, widths, depths, strides, ...)

• Think of a Python function:

def vgg_net(input_res, depth, conv1_out_channels):
... # code to make the model
return model

- Input: some arguments (resolution, widths, depths, strides, ...)
- Output: instantiation of a network

• Think of a Python function:

def vgg_net(input_res, depth, conv1_out_channels):
... # code to make the model
return model

- Input: some arguments (resolution, widths, depths, strides, ...)
- Output: instantiation of a network
- All valid input args \rightarrow generates a population of related networks

From Single Networks to Populations of Networks

- Comparing individual models ("point estimates")
 - E.g., AlexNet vs. VGG-16



- No controls
- Why is one better than another?

Poster #61 in session 1.2 tomorrow (Tuesday)

From Single Networks to Populations of Networks

• Comparing a few hand-designed variants ("curve estimates")

facebook AI Research

• E.g., ResNet-50, ResNet-101, ... vs. ResNeXt-50-32x4d, ResNeXt-101-32x4d, ...



From Single Networks to Populations of Networks

- Comparing generated populations of models ("distribution estimates")
 - Define network generators, evaluate population

facebook AI Research



- Discover general principles
 - E.g., grouped conv (extreme: depthwise conv) is good

- Discover general principles
 - E.g., grouped conv (extreme: depthwise conv) is good
- Understand the landscape
 - Are all models basically good? (e.g., insight into NAS)

	#ops	#nodes	output	#cells (B)
NASNet [41]	13	5	L	71,465,842
Amoeba [30]	8	5	L	556,628
PNAS [20]	8	5	А	556,628
ENAS [29]	5	5	L	5,063
DARTS [21]	8	4	А	242

Changes to NAS design spaces in a sequence of NAS papers

- Discover general principles
 - E.g., grouped conv (extreme: depthwise conv) is good
- Understand the landscape
 - Are all models basically good? (e.g., insight into NAS)

facebook AI Research

- Discover general principles
 - E.g., grouped conv (extreme: depthwise conv) is good
- Understand the landscape
 - Are all models basically good? (e.g., insight into NAS)

facebook AI Research

- Discover general principles
 - E.g., grouped conv (extreme: depthwise conv) is good
- Understand the landscape
 - Are all models basically good? (Insight into NAS)
- Explore scientific questions

Explore Scientific Questions

How does random wiring compare to engineered wiring?

Historical Motivation

Minsky (1951) randomly wired "SNARC"

Turing (1948) *"Unorganized machines"*

facebook AI Research

Rosenblatt (1959)

first Perceptron machine

Neuroscience Motivation

facebook AI Research

Bullmore and Sporns. "The economy of brain network organization." Nature Reviews Neuroscience, 2012 Bullmore and Sporns. "Complex brain networks: graph theoretical analysis of structural and functional systems." Nature reviews neuroscienc

Practical Motivation

Practical Motivation

Prior Work: Random Micro Wiring

- Macro vs. micro wiring
 - Macro: wires between "op level" units (e.g., conv) [our focus]
 - Micro: wires inside "op level" units (e.g., sparse conv filters)

E.g., micro wiring: random expander graphs

Deep Expander Networks: Efficient Deep Networks from Graph Theory. Ameya Prabhu, Girish Varma, Anoop Namboodiri. ECCV 2018.

Prior Work: Connectivity within ResNe(X)t

- Optimize for weights and connectivity masks with gradient descent
- Focus: learned wiring

Karim Ahmed and Lorenzo Torresani, "MaskConnect: Connectivity Learning by Gradient Descent", ECCV 2018

Prior Work: Learned Wiring with NAS

Searching for (1) operations (2) wiring patterns

Network Generators (More Formally)

$$g: \Theta \mapsto \mathcal{N}$$

g is a mapping from a parameter space Θ , to a network architecture space \mathcal{N}

Network Generators (More Formally)

$$g: \Theta \mapsto \mathcal{N}$$

g is a mapping from a parameter space Θ , to a network architecture space \mathcal{N}

• The parameters $\theta \in \Theta$ specify the instantiated network

Network Generators (More Formally)

$$g: \Theta \mapsto \mathcal{N}$$

g is a mapping from a parameter space Θ , to a network architecture space \mathcal{N}

- The parameters $\theta \in \Theta$ specify the instantiated network
- E.g., ResNet generator: θ = {#stages, #blocks per stage, depth, width, filter sizes, ...}

$$G(\theta) = \frac{34 \cdot 134 \cdot 1}{33 \cdot 10^{10} \cdot 13} = \frac{$$

Stochastic Network Generators

• The network generator $g(\theta)$ performs a deterministic mapping

Stochastic Network Generators

- The network generator $g(\theta)$ performs a deterministic mapping
- s: seed of a pseudo-random number generator used by g

Stochastic Network Generators

- The network generator $g(\theta)$ performs a deterministic mapping
- s: seed of a pseudo-random number generator used by g
- $g(\theta, s)$ is a stochastic network generator

Network Generator Perspective

- Encapsulates the *entire* network generation process
 - Empirical priors (what we already know works well)
 - Learnable parameters θ
 - Randomization s
 - Misc. heuristic rules are revealed

E.g., NAS from the Network Generator Perspective

• LSTM controller weights: θ

E.g., NAS from the Network Generator Perspective

- LSTM controller weights: θ
- Sample actions (insert an conv, connect two nodes, ...) determined by θ ,

S

E.g., NAS from the Network Generator Perspective

- LSTM controller weights: θ
- Sample actions (insert an conv, connect two nodes, ...) determined by θ , s
- Allowed wiring patterns is a small subset of all possible graphs
 - E.g., 5 nodes in a cell always have input degree 2 and output degree 1

The generator design involves a strong prior

"Misc. heuristic rules"

Network Generators for Randomly Wired Networks

How do we design neural networks like these?

Step 1. Generate a general graph G

Step 2. Map graph G to a valid neural network

Random Graph Models

Erdös-Rényi (ER), 1959

Barabási–Albert (BA), 1998

"Scale-free" graphs

m = 3

Watts–Strogatz (WS), 1998

"Small-world" graphs

Example: Small-World Ggraphs 100 nodes, average degree = 4

Network Generator

facebook AI Research

Ő

Network Generator

Network Generator

Notes on Nodes and Edges

- Overall computation \propto # nodes
- # nodes is fixed
- Tiny extra computation on each edge ($\times w_i$)

• Differences in accuracy are about wiring, not FLOPs or parameters

random (fixed seed)

deterministic

Small Model Regime – Engineered Baseline

facebook AI Research

Results – Small Computation Regime (< 600M FLOPs)

	Network	Top-1 Acc	FLOPs
esigned -	MobileNet	70.6	569
	MobileNet v2	74.7	585
	ShuffleNet	73.7	524
	ShuffleNet v2	74.9	591

Hand de

Results – Small Computation Regime (< 600M FLOPs)

	Network	Тор-1 Асс	FLOPs
and designed -	MobileNet	70.6	569
	MobileNet v2	74.7	585
	ShuffleNet	73.7	524
	ShuffleNet v2	74.9	591
Γ	NASNet-A	74.0	564
NAS –	NASNet-C	72.5	558
	AmoebaNet-A	74.5	555
	AmoebaNet-C	75.7	570
	PNAS	74.2	588
	DARTS	73.1	595

Η

Results – Small Computation Regime (< 600M FLOPs)

	Network	Top-1 Acc	FLOPs
Γ	MobileNet	70.6	569
Hand designed	MobileNet v2	74.7	585
	ShuffleNet	73.7	524
	ShuffleNet v2	74.9	591
ſ	NASNet-A	74.0	564
	NASNet-C	72.5	558
NAC	AmoebaNet-A	74.5	555
NAS -	AmoebaNet-C	75.7	570
	PNAS	74.2	588
	DARTS	73.1	595
Randomly Wired -	RandWire-WS	74.7 ±0.25	583 ±6.2

ResNet-50

RandWire (same FLOPs)

Conclusions

- The next frontier: Architecting network generators?
- Study *populations* of networks
- We explored randomly wired nets *they can work surprisingly well!*
- Heuristic rules are pervasive even in "random" and "automatic" generators