

A long time ago, in a room in the basement  
of EE...

# Representation learning of genomic sequence motifs with convolutional neural networks

Peter K. Koo and Sean R. Eddy

*CompBio Faculty  
Candidate, same time*



CompBio Seminar

February 10, 2020

Alyssa LaFleur and Erin Wilson

# Representation learning of genomic sequence motifs with convolutional neural networks

Peter K. Koo and Sean R. Eddy

CompBio Seminar

March 366, 2020

Erin Wilson

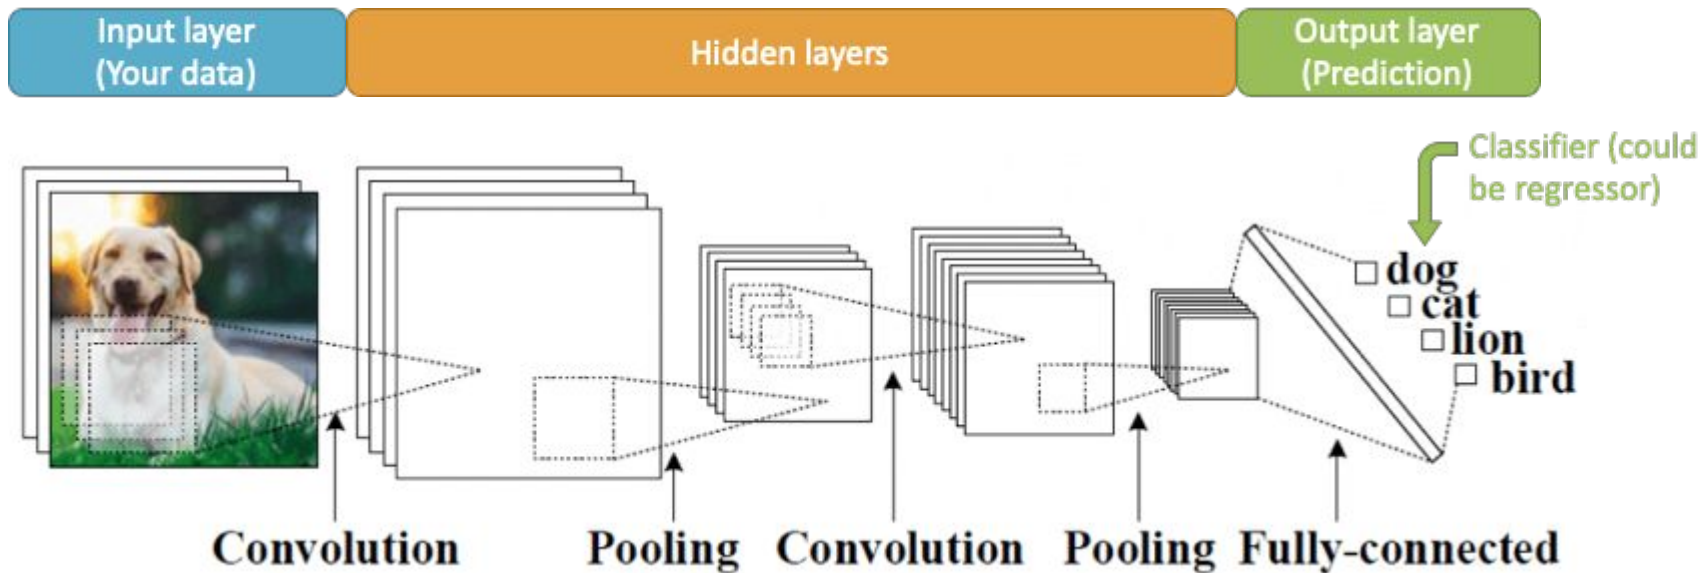
(slides co-created with Alyssa!)

# Overview


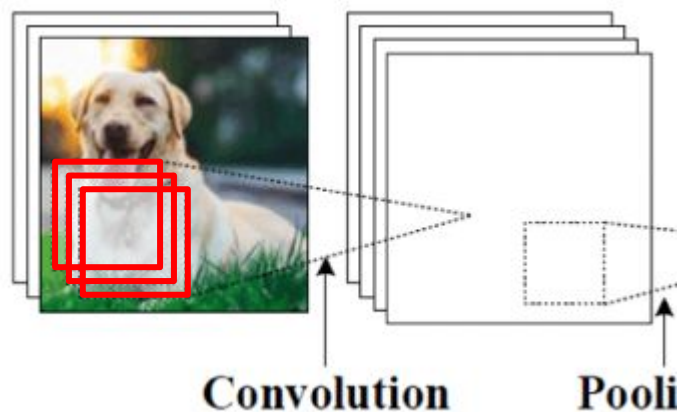
- Background
  - How do CNNs work?
  - Why do people care about finding motifs?
- Methods
  - Synthetic dataset used in this paper
  - Experimental setup for CNN architectures
  - Model vs Motif evaluation metrics
- Results
  - Pulling various CNN architecture levers!
- Main takeaways & discussion



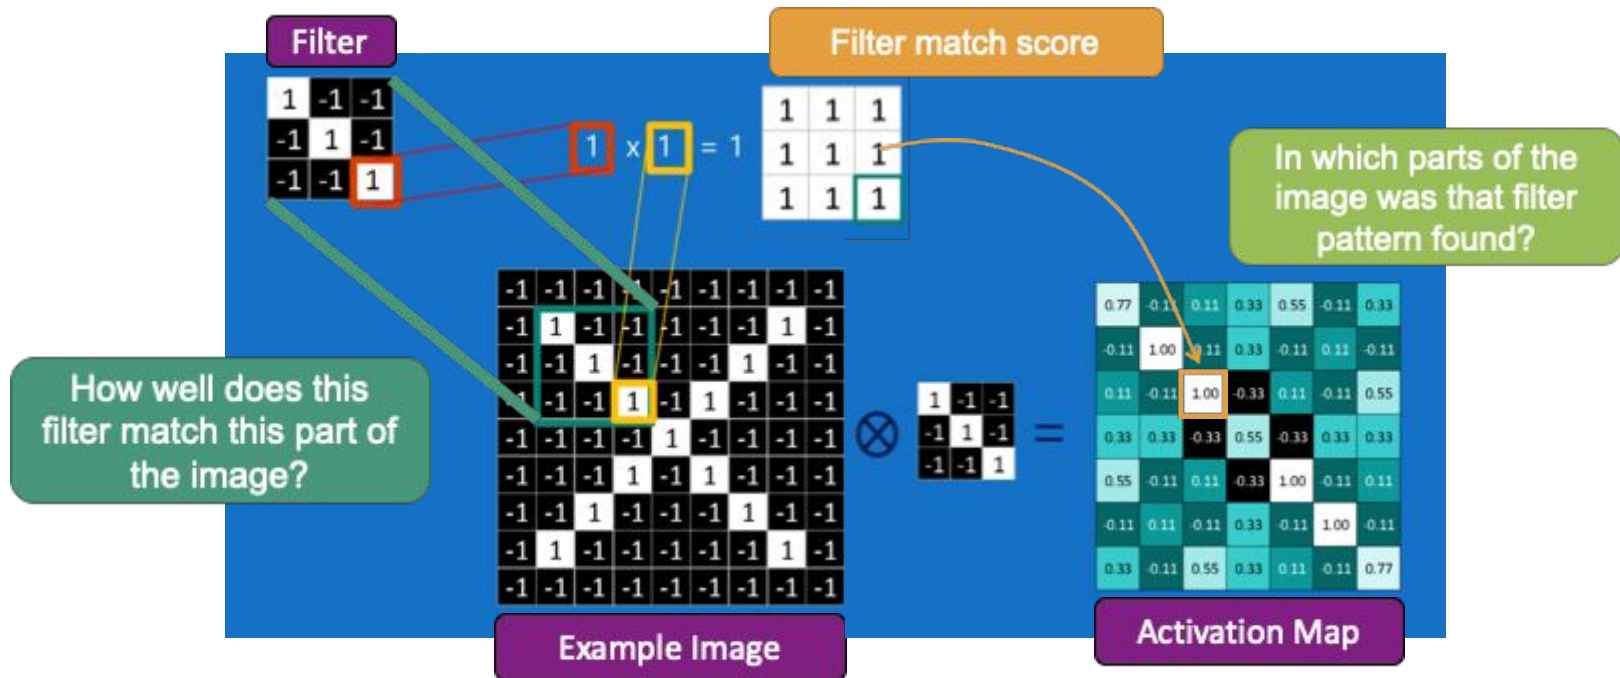
# A shallow dive into deep learning...



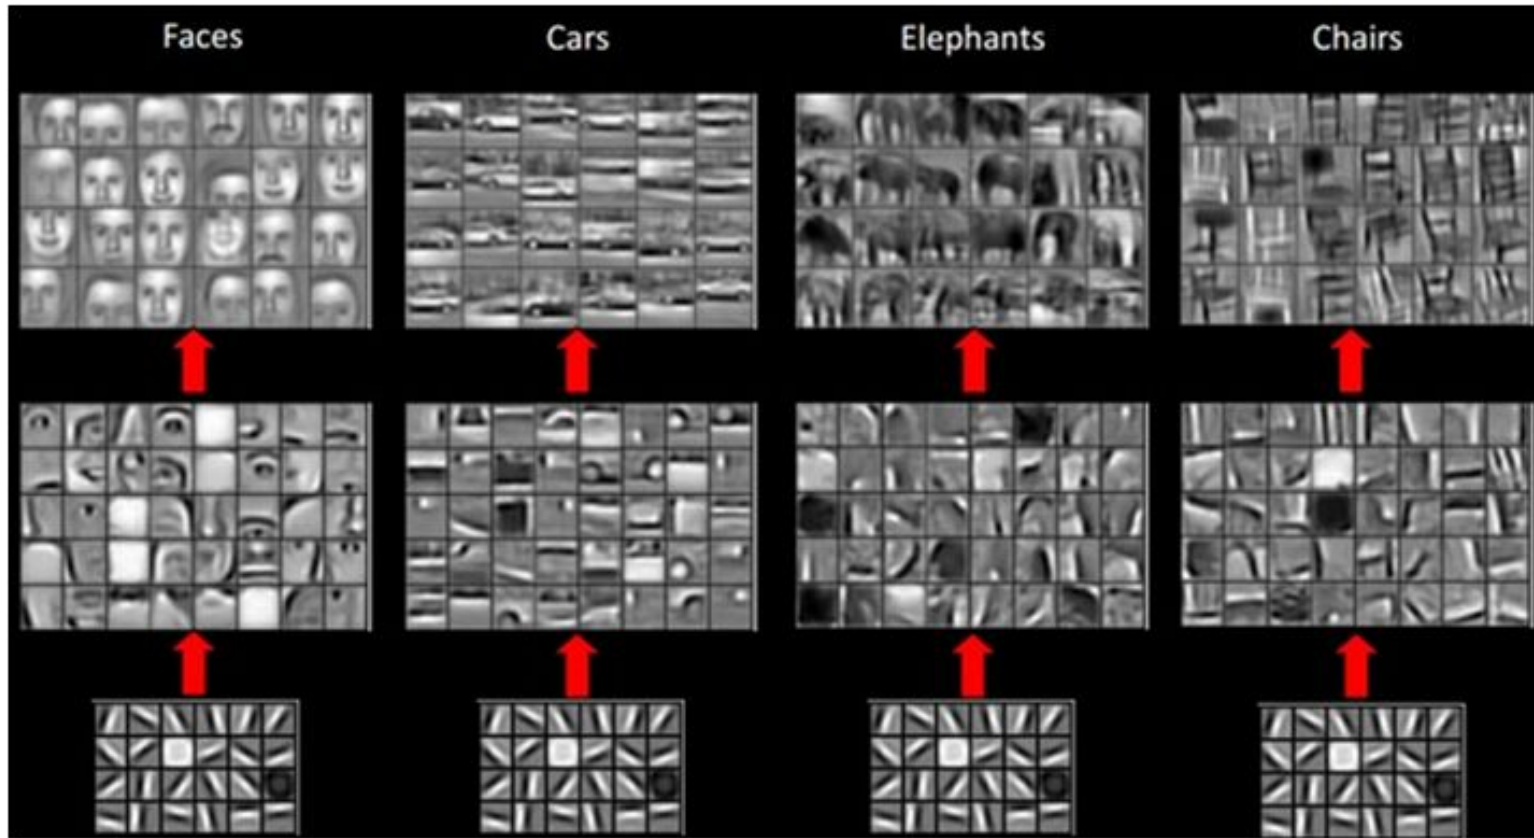
# CNNs capture local spatial information between pixels in an image



# Filters are like small patterns. You can identify areas of the image containing that pattern.

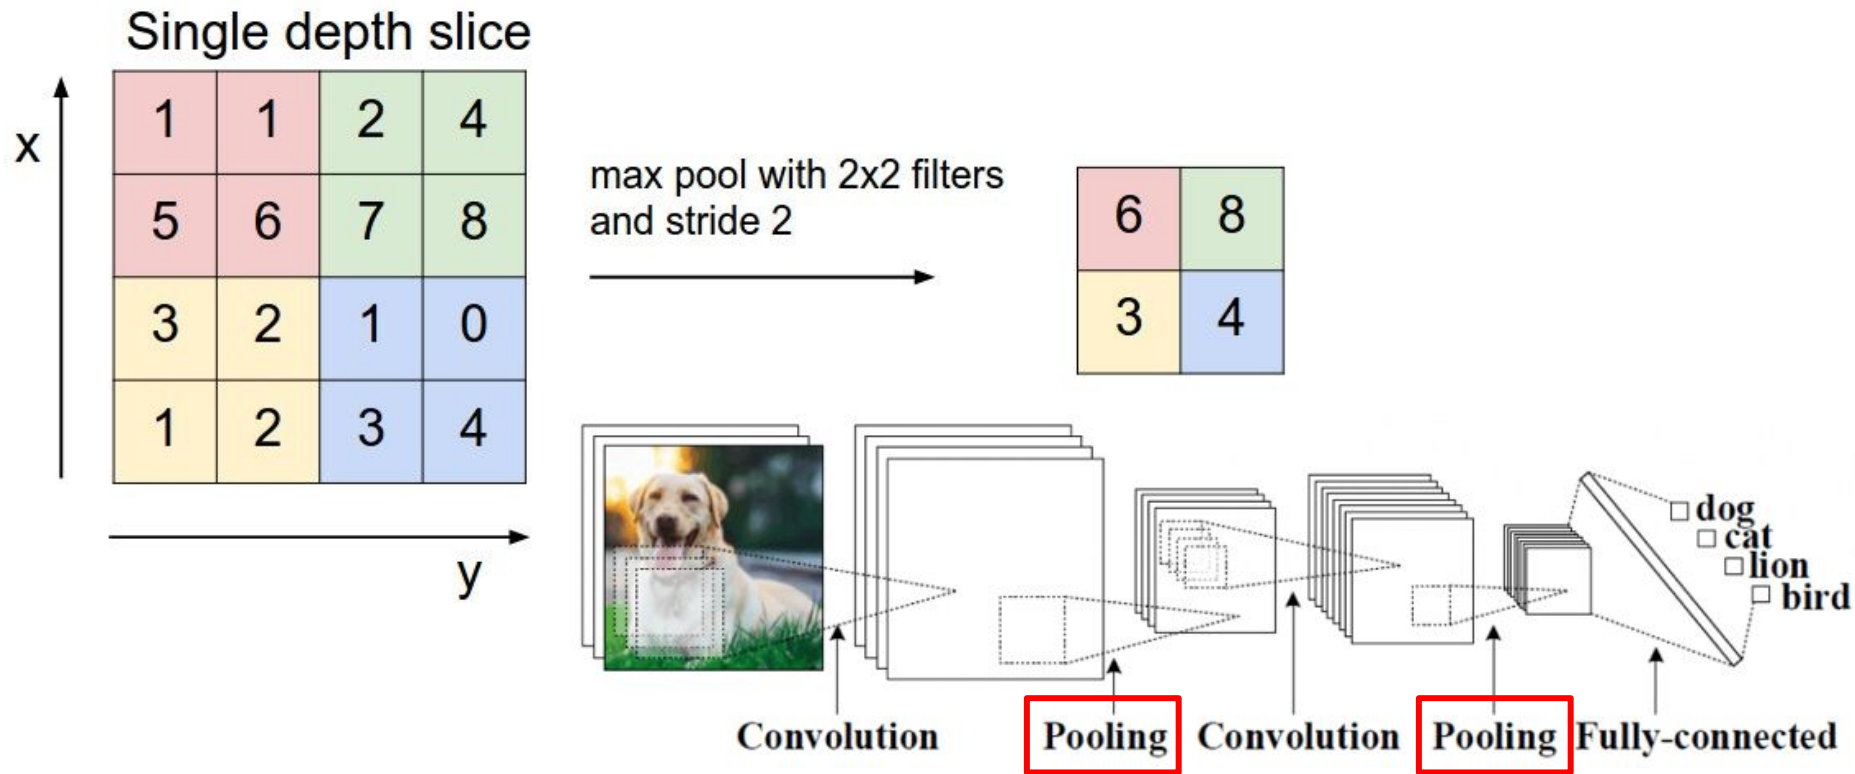


# Filters learn basic patterns that can be composed into more complex features















# Max pooling: reduce image features by taking the max value from a window




# How are CNNs helpful in biology?

Google    

[All](#) [News](#) [Images](#) [Videos](#) [Books](#) [More](#) [Settings](#) [Tools](#)

 kitten  baby  anime  fluffy  white  wallpaper 



International Cat Care | The ultimate ...  
icatcare.org

Cats is being patched with 'improved ...  
theverge.com


Thinking of getting a cat ...  
icatcare.org



The 'Cats' trailer dropped. We have 34 ...  
washingtonpost.com




Cat Excessive Meowing and Yowling: Why ...  
pets.webmd.com



Cats' Bound To Lose At Least \$71M After ...  
deadline.com





Bu  
wir




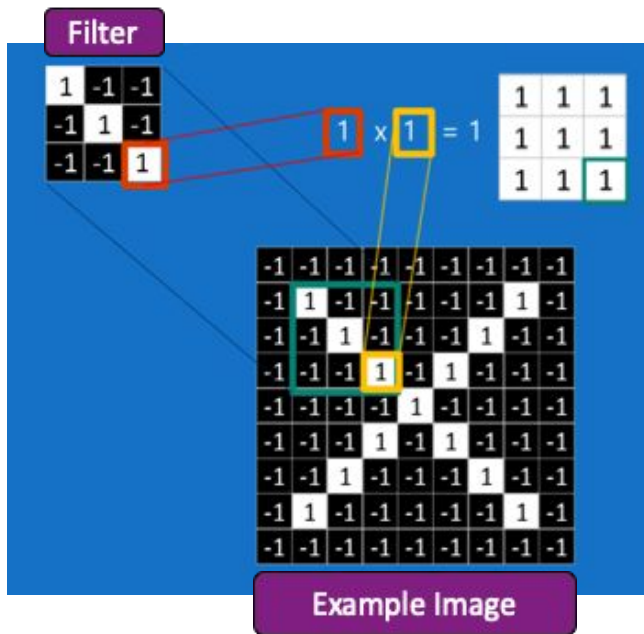
# How to pretend your DNA is a cat.

A	[1,0,0,0]
C	[0,1,0,0]
G	[0,0,1,0]
T	[0,0,0,1]

<u>A</u>	<u>T</u>	<u>G</u>	<u>G</u>	<u>C</u>	<u>T</u>	<u>C</u>	<u>A</u>	<u>T</u>
1	0	0	0	0	0	0	1	0
0	0	0	0	1	0	1	0	0
0	0	1	1	0	0	0	0	0
0	1	0	0	0	1	0	0	1




# What do filters learn?



A	0.0	1.0	0.0	1.0	1.0	0.0
C	0.0	0.0	0.0	0.0	0.0	0.0
G	0.0	0.0	0.0	0.0	0.0	0.0
T	1.0	0.0	1.0	0.0	0.0	1.0
	<b>T</b>	<b>A</b>	<b>T</b>	<b>A</b>	<b>A</b>	<b>T</b>

A	0.1	0.7	0.2	0.4	0.4	0.1
C	0.0	0.0	0.1	0.2	0.2	0.0
G	0.1	0.1	0.1	0.2	0.2	0.0
T	0.8	0.2	0.6	0.2	0.2	0.9
	<b>T</b>	<b>A</b>	<b>T</b>	<b>A</b>	<b>A</b>	<b>T</b>



# CNN filters can learn motifs relevant to the prediction task

How do upstream sequences influence gene expression strength?

Gene

**Learned filters:**




cAAcAUG

Predicts high expression!




cGGcCCGc

Predicts low expression!



Upstream sequence	Expression strength
ATGGCTCATATCTCCG...	204
TATCTCCGCTAATCGA...	50
CTAATCGAACATCGCA...	3
CATCGCATGTCGATTA...	186

# Motifs are landing zones for various DNA binding proteins



Gene

Txn Factor Binding sites  
stability element  
Ribosome binding site

Gene

ACTGCGTATATGGCTCATATCTCGCTAATCGATGATCGCCATGTCGATTACGTATATGCGTCTCTCCTAATAGATCGATGCTAGCTGTACGT

# JASPAR CORE



Total 1964 profiles

<http://jaspar.genereg.net/>

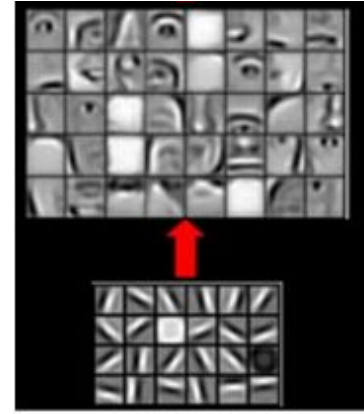
Display  profiles

Filter:

<input type="checkbox"/>	ID	Name	Species	Class	Family	Sequence logo
<input type="checkbox"/>	<b>MA0001.1</b>	AGL3	Arabidopsis thaliana	MADS box factors	MADS	
<input type="checkbox"/>	<b>MA0001.2</b>	AGL3	Arabidopsis thaliana	MADS box factors		
<input type="checkbox"/>	<b>MA0002.1</b>	RUNX1	Homo sapiens	Runt domain factors	Runt-related factors	
<input type="checkbox"/>	<b>MA0002.2</b>	RUNX1	Mus musculus	Runt domain factors	Runt-related factors	
<input type="checkbox"/>	<b>MA0003.1</b>	TFAP2A	Homo sapiens	Basic helix-span-helix factors (bHSH)	AP-2	
<input type="checkbox"/>	<b>MA0003.2</b>	TFAP2A	Homo sapiens	Basic helix-span-helix factors (bHSH)	AP-2	

# Main takeaways:


1.) CNN filters are good at finding small areas of patterns within a bigger pattern that are useful for prediction tasks



2.) For DNA sequence inputs, CNN filters learn DNA motifs





3.) Motifs usually contain some biological relevance for how, when, and where proteins bind to DNA





# Representation learning of genomic sequence motifs with convolutional neural networks

Peter K. Koo <sup>1□\*</sup>, Sean R. Eddy <sup>1,2\*</sup>

Main question:

How does the *architecture* of the CNN influence its ability to *learn whole motifs* in the first convolutional layer?

**A**



Pattern 1: CACGTG




Pattern 2: GTGCAC




Pattern 3: CACNNGTG



Filter 1: GTG



Filter 2: CAC




**A**



Pattern 1: CACGTG





Pattern 2: GTGCAC



Pattern 3: CACNNGTG




Filter 1: GTG



**A**

Pattern 1: CACGTG



Pattern 2




A	0.1	0.7	0.2	0.4	0.4	0.1
C	0.0	0.0	0.1	0.2	0.2	0.0
G	0.1	0.1	0.1	0.2	0.2	0.0
T	0.8	0.2	0.6	0.2	0.2	0.9


T A T A A T




Filter 1: GTG




Filter 2: CAC



Filter 1: GTG




Filter 2: CAC




**A**



Pattern 1: CACGTG




Pattern 2: GTGCAC






Pattern 3: CACNNGTG



Filter 1: GTG



Filter 2: CAC


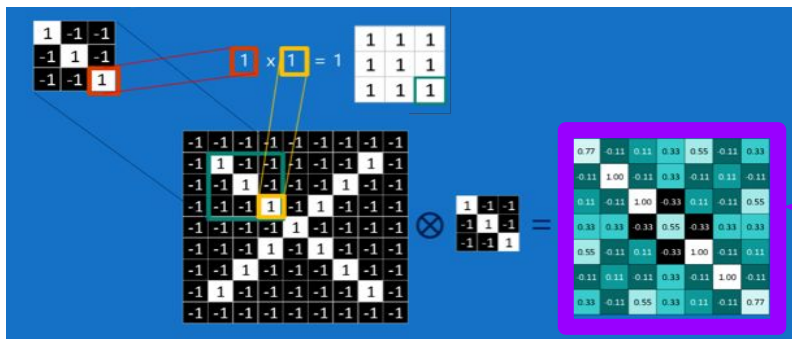



**A**

Pattern 1: CACGTG

Pattern 2: GTGCAC


Pattern 3: CACNNGTG

**Filter 1: GTG****Filter 2: CAC**


↑ Filter activation  
← map

# Second layer convolutional filters

**A**



**B**



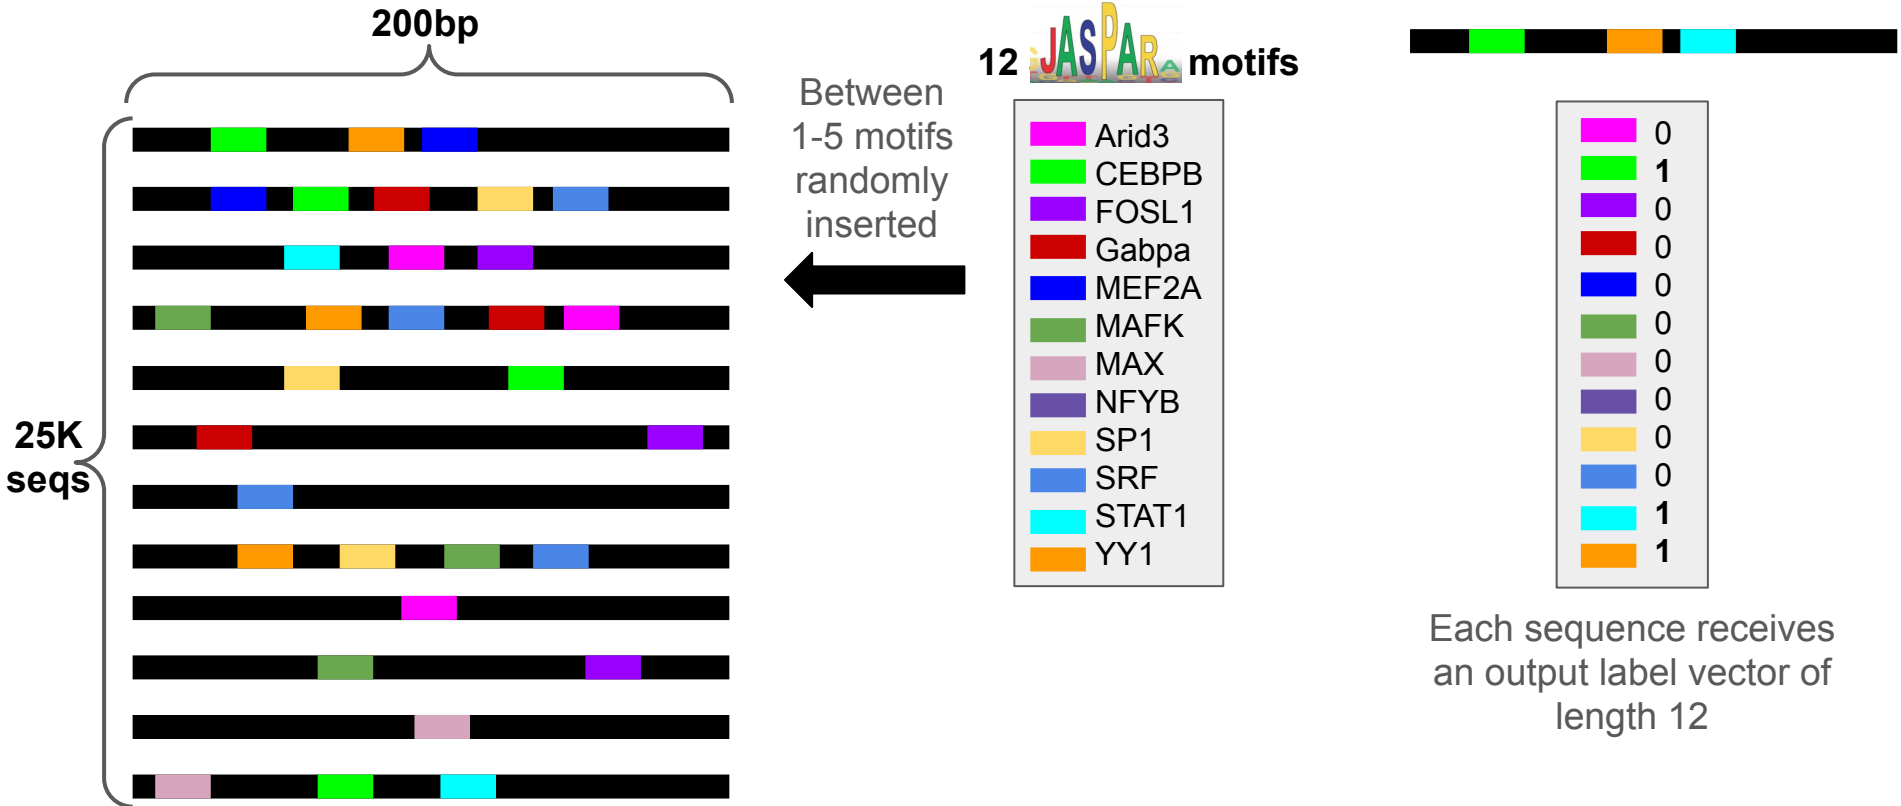
# Overview

- Background
  - How do CNNs work?
  - Why do people care about finding motifs?
- **Methods**
  - **Synthetic dataset used in this paper**
  - **Experimental setup for CNN architectures**
  - **Model vs Motif evaluation metrics**
- Results
  - Pulling various CNN architecture levers!
- Main takeaways & discussion





# Methods: creating a synthetic dataset



# Methods: Network architecture framework

Prediction on 12 motifs  
(sigmoid activation)

Fully connected  
hidden layer  
(512 nodes)

Vary pool size


Max pooling

Vary filter  
size/number

2nd layer convolutions

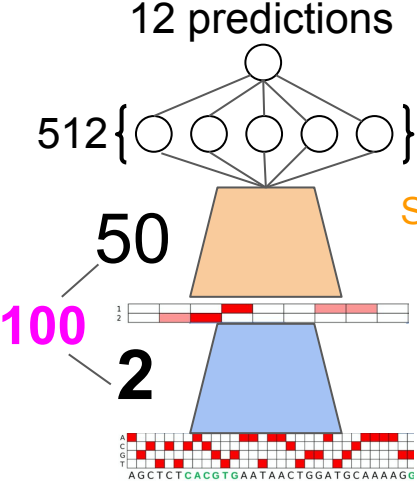
Max pooling

1st layer convolutions

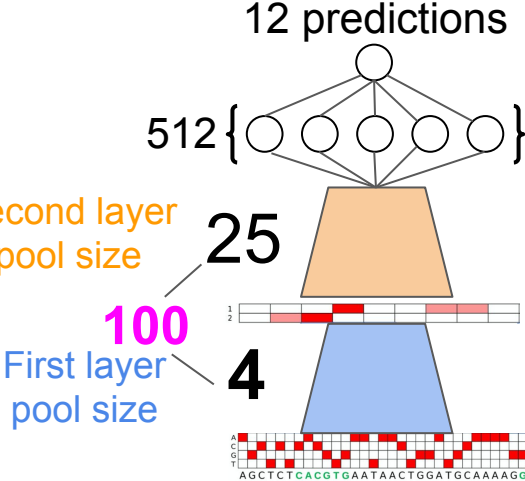


# Methods: Network naming scheme


Max-pooling: **product** of first and second pool sizes is **100**.



CNN-2



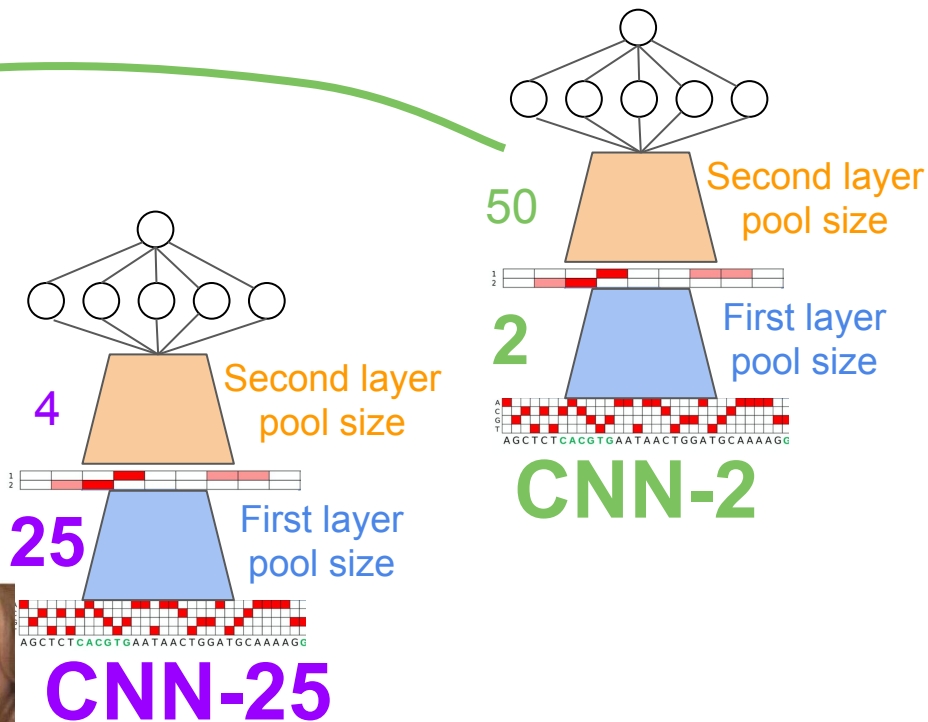

CNN-4



CNN-25

# Models are (mostly) named for their first pool size


Model	Average AU-ROC
CNN-1	0.972±0.001
CNN-2	0.966±0.000
CNN-4	0.955±0.002
CNN-10	0.964±0.007
CNN-25	0.973±0.001
CNN-50	0.961±0.011
CNN-100	0.958±0.012
CNN <sub>9</sub> -4	0.954±0.002
CNN <sub>9</sub> -25	0.958±0.008
CNN <sub>3</sub> -50	
CNN <sub>3</sub> -2	
CNN-50-2	
CNN <sub>19-1</sub> -2	
CNN 25 (60)	
CNN 25 (90)	
CNN-25 (120)	0.963±0.001





# Methods: evaluate models using AU-ROC

True positive rate

$$\frac{\text{TPR / Recall / Sensitivity}}{\text{TP} + \text{FN}}$$



False positive rate =  $\frac{\text{FP}}{\text{TN} + \text{FP}}$



# Evaluate models for consistent AU-ROC, not best!

Model	Average AU-ROC
CNN-1	0.972±0.001
CNN-2	0.966±0.000
CNN-4	0.955±0.002
CNN-10	0.964±0.007
CNN-25	0.973±0.001
CNN-50	0.961±0.011
CNN-100	0.958±0.012
CNN <sub>9</sub> -4	0.954±0.002
CNN <sub>9</sub> -25	0.958±0.008
CNN <sub>3</sub> -50	0.648±0.008
CNN <sub>3</sub> -2	0.968±0.001
CNN-50-2	0.921±0.012
CNN <sub>19-1</sub> -2	0.969 ±0.002
CNN-25 (60)	0.972±0.001
CNN-25 (90)	0.968±0.001
CNN-25 (120)	0.963±0.001

- **Not** concerned with maximizing AU-ROC - want to be **consistent**
- Real question: after change some aspect of network **structure**, and evaluate the **motifs learned** by first layer filters


# Overview

- Background
  - How do CNNs work?
  - Why do people care about finding motifs?
- Methods
  - Synthetic dataset used in this paper
  - Experimental setup for CNN architectures
  - Model vs Motif evaluation metrics
- **Results**
  - **Pulling various CNN architecture levers!**
- Main takeaways




# Results: vary max pooling size


Pool = 2



Pool = 4


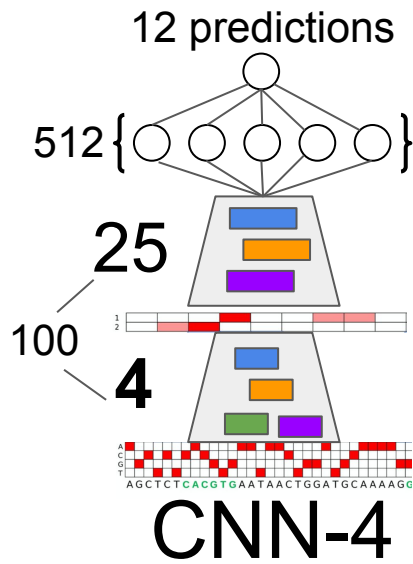
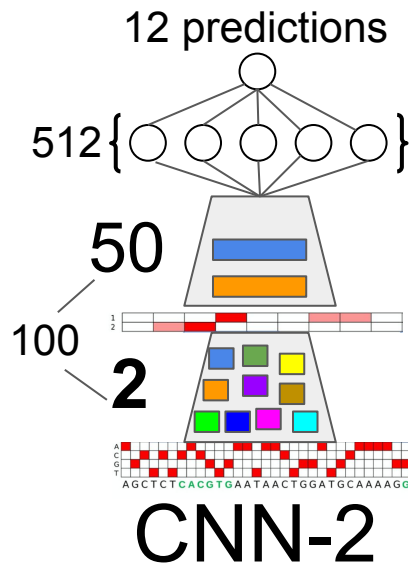


Pool = 25






# Results: vary max pooling size


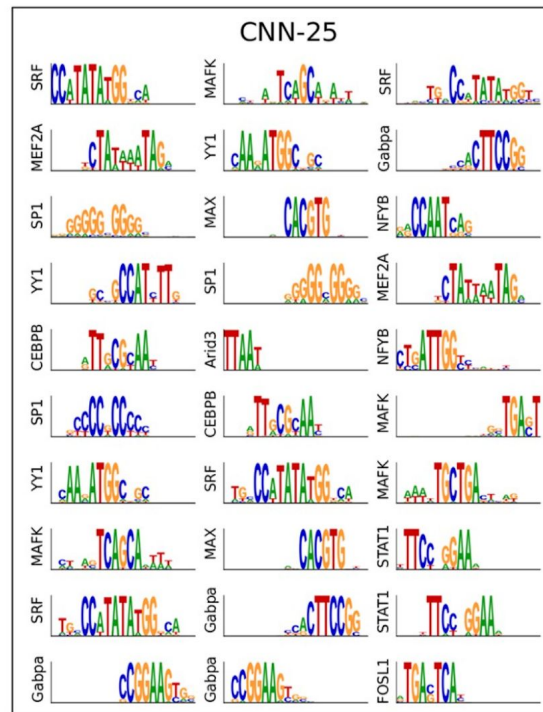
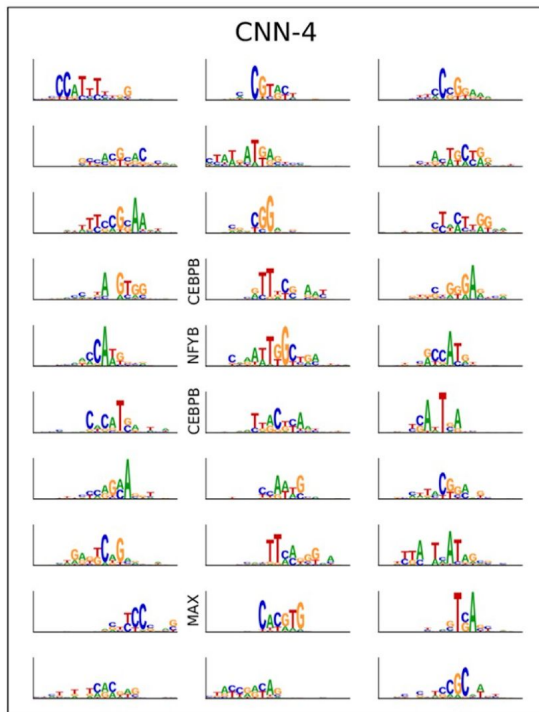
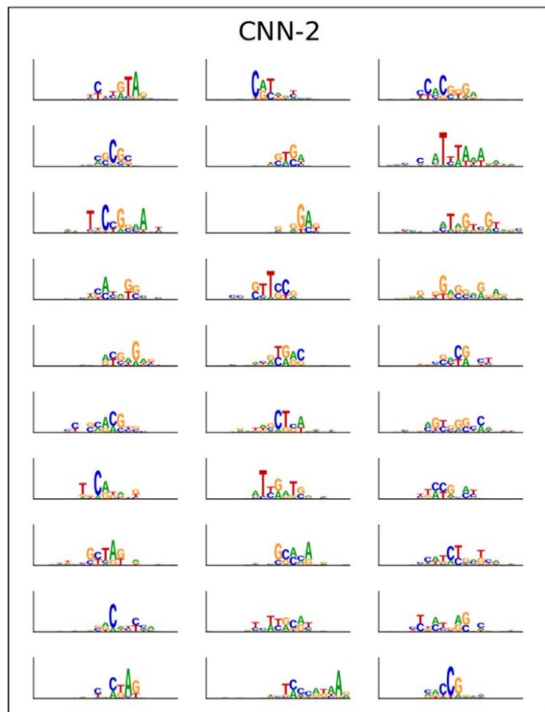


# Results: vary max pooling size


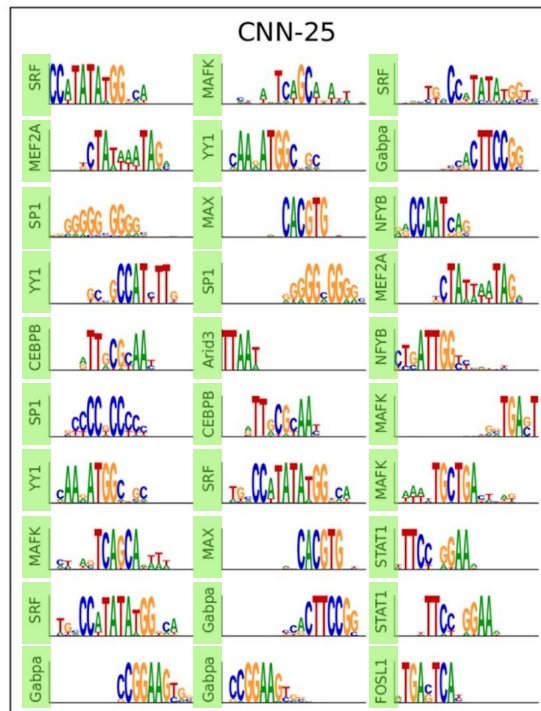
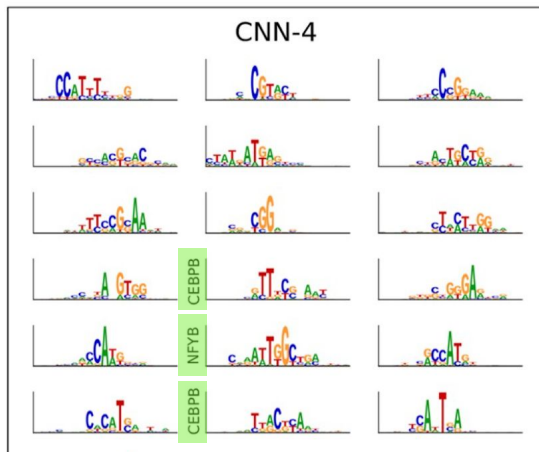
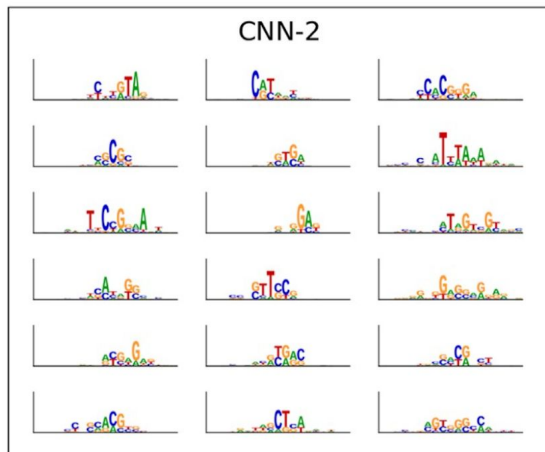
(image versions to use without fixing font size everytime)




# Results: vary max pooling size



# Results: vary max pooling size



## Ground truth

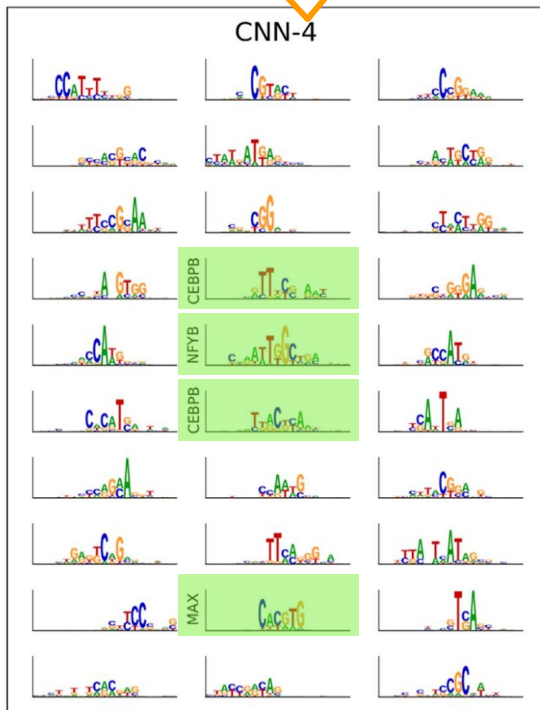
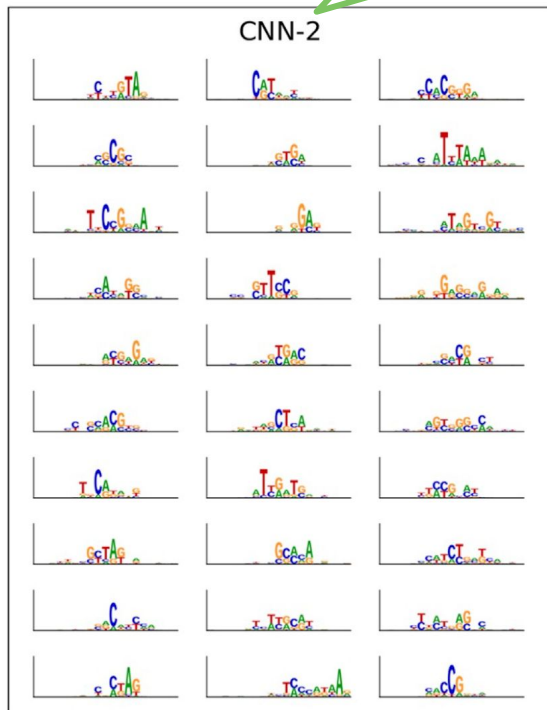


# Results: vary max pooling size

Small pool

Small-ish pool

Large-ish pool



# Results: vary max pooling



12 Ground Truth motifs



Table 1. Performance on the synthetic dataset.

Model	Average AU-ROC	% Motif match (JASPAR)	% Motif match (Relevant)
CNN-1	0.972±0.001	0.240±0.083	0.007±0.013
CNN-2	0.966±0.000	0.240±0.071	0.007±0.013
CNN-4	0.955±0.002	0.453±0.131	0.127±0.080
CNN-10	0.964±0.007	0.987±0.016	0.973±0.033
CNN-25	0.973±0.001	0.987±0.016	0.980±0.027
CNN-50	0.961±0.011	0.933±0.037	0.920±0.045
CNN-100	0.958±0.012	0.887±0.034	0.880±0.034
CNN <sub>9</sub> -4	0.954±0.002	0.260±0.039	0.033±0.030
CNN <sub>9</sub> -25	0.958±0.008	0.993±0.013	0.980±0.016
CNN <sub>3</sub> -50	0.648±0.008	0.160±0.049	0.000±0.000
CNN <sub>3</sub> -2	0.968±0.001	0.233±0.067	0.000±0.000

Small pool

Small-ish pool

Large-ish pool

*Take away: wider pooling size (like CNN-25) forces first row filters to learn WHOLE motifs*

CNN-25 (120)

0.963±0.001

0.933±0.015

0.887±0.025

# Results: vary filter number





Table 1. Performance on the synthetic dataset.

Model	Average AU-ROC	% Motif match (JASPAR)	% Motif match (Relevant)
CNN-1	0.972±0.001	0.240±0.083	0.007±0.013
CNN-2	0.966±0.000	0.240±0.071	0.007±0.013
CNN-4	0.955±0.002	0.453±0.131	0.127±0.080
CNN-10	0.964±0.007	0.987±0.016	0.973±0.033
Default: 30 filters > CNN-25	0.973±0.001	0.987±0.016	0.980±0.027
CNN-50	0.961±0.011	0.933±0.037	0.920±0.045
CNN-100	0.958±0.012	0.887±0.034	0.880±0.034

*Take away: more filters does not improve accuracy and % of filters that learn motifs decreases*

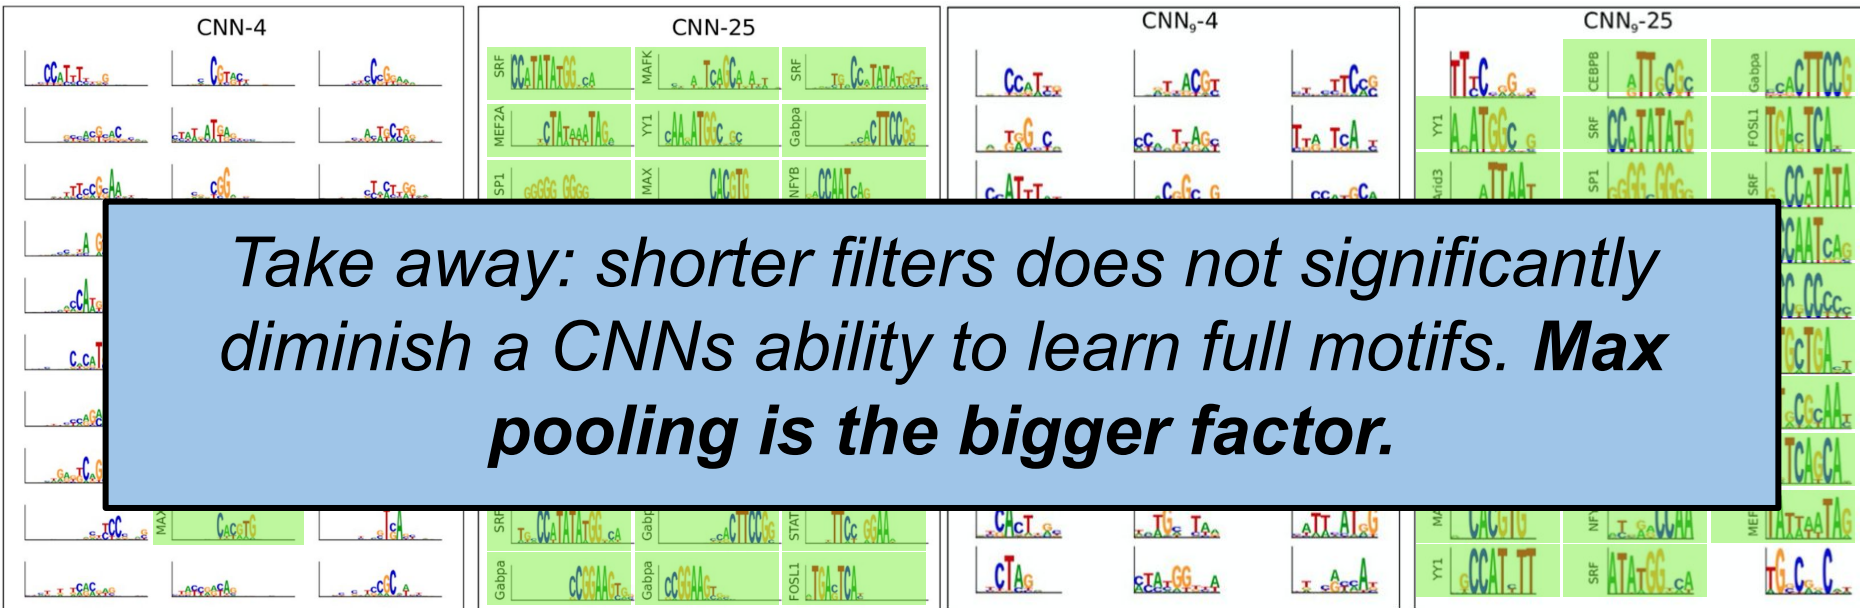
CNN-50-2	0.921±0.012	0.913±0.050	0.893±0.044
CNN <sub>19-1</sub> -2	0.969 ±0.002	0.867±0.056	0.747±0.096
CNN-25 (60)	0.972±0.001	0.973±0.013	0.960±0.023
CNN-25 (90)	0.968±0.001	0.940±0.023	0.909±0.028
CNN-25 (120)	0.963±0.001	0.933±0.015	0.887±0.025

Vary # filters

# Results: vary filter size

Default: filter size = 19

Test filter size = 9



Small-ish  
pool

Large-ish  
pool

Small-ish  
pool

Large-ish  
pool




# Results: Restricting deeper layer assemblies

## CNN-2

2nd layer filter size is usually 5

Still possible to rearrange filters containing partial motifs




# Results: Restricting deeper layer assemblies

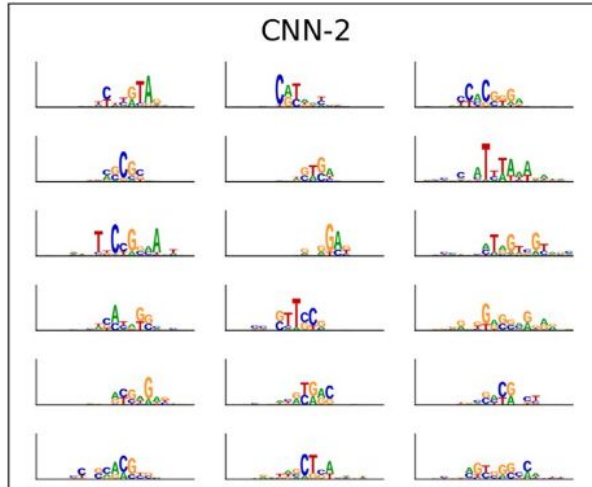
CNN<sub>19-1</sub><sup>-2</sup>

Now, restrict 2nd layer  
filter size to 1

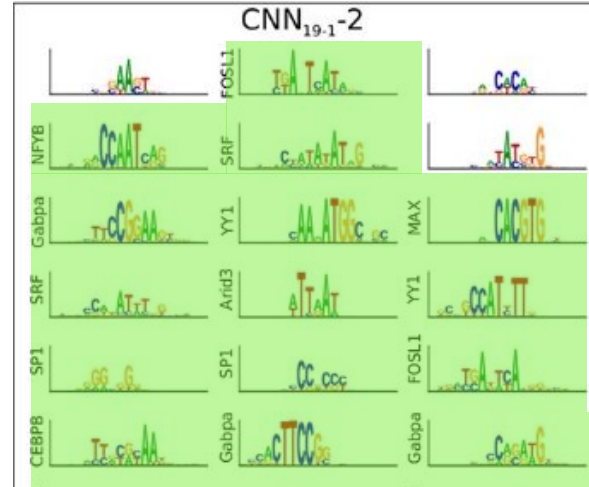
NOW: not possible to  
rearrange filters containing  
partial motifs



# CNN-2



# CNN<sub>19-1</sub>-2



*Take away: learning **WHOLE** motif representations in first layer is affected by the ability of **deeper layers** to hierarchically build motifs.*

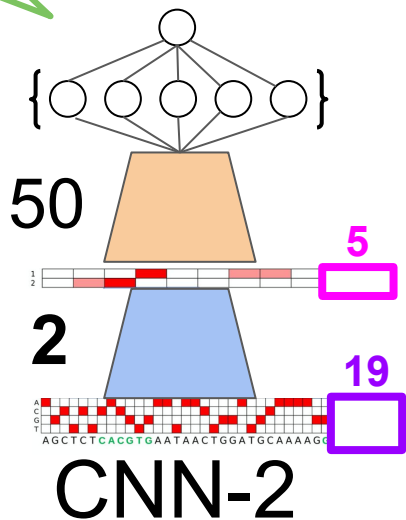
# Results: Restricting deeper layer assemblies

Table 1. Performance on the synthetic dataset.

Model	Average AU-ROC	% Motif match (JASPAR)	% Motif match (Relevant)
CNN-1	0.972±0.001	0.240±0.083	0.007±0.013
CNN-2	0.966±0.000	0.240±0.071	0.007±0.013
CNN-4	0.955±0.002	0.453±0.131	0.127±0.080
CNN-10	0.964±0.007	0.987±0.016	0.973±0.033
CNN-25	0.973±0.001	0.987±0.016	0.980±0.027
CNN-50	0.961±0.011	0.933±0.037	0.920±0.045
CNN-100	0.958±0.012	0.887±0.034	0.880±0.034
CNN <sub>9</sub> -4	0.954±0.002	0.260±0.039	0.033±0.030
CNN <sub>9</sub> -25	0.958±0.008	0.993±0.013	0.980±0.016
CNN <sub>3</sub> -50	0.648±0.008	0.160±0.049	0.000±0.000
CNN <sub>3</sub> -2	0.968±0.001	0.233±0.067	0.000±0.000
CNN-50-2	0.921±0.012	0.913±0.050	0.893±0.044
CNN <sub>19-1</sub> -2	0.969 ±0.002	0.867±0.056	0.747±0.096
CNN-25 (60)	0.972±0.001	0.973±0.013	0.960±0.023
CNN-25 (90)	0.968±0.001	0.940±0.023	0.909±0.028
CNN-25 (120)	0.963±0.001	0.933±0.015	0.887±0.025

# Results: Restricting deeper layer assemblies

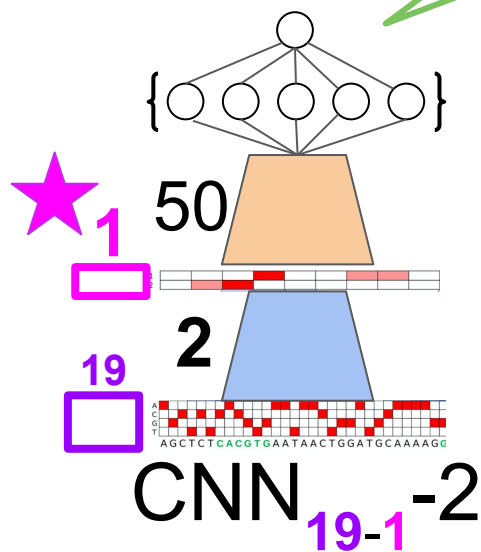
Small pool



2nd layer convolutions

1st layer convolutions

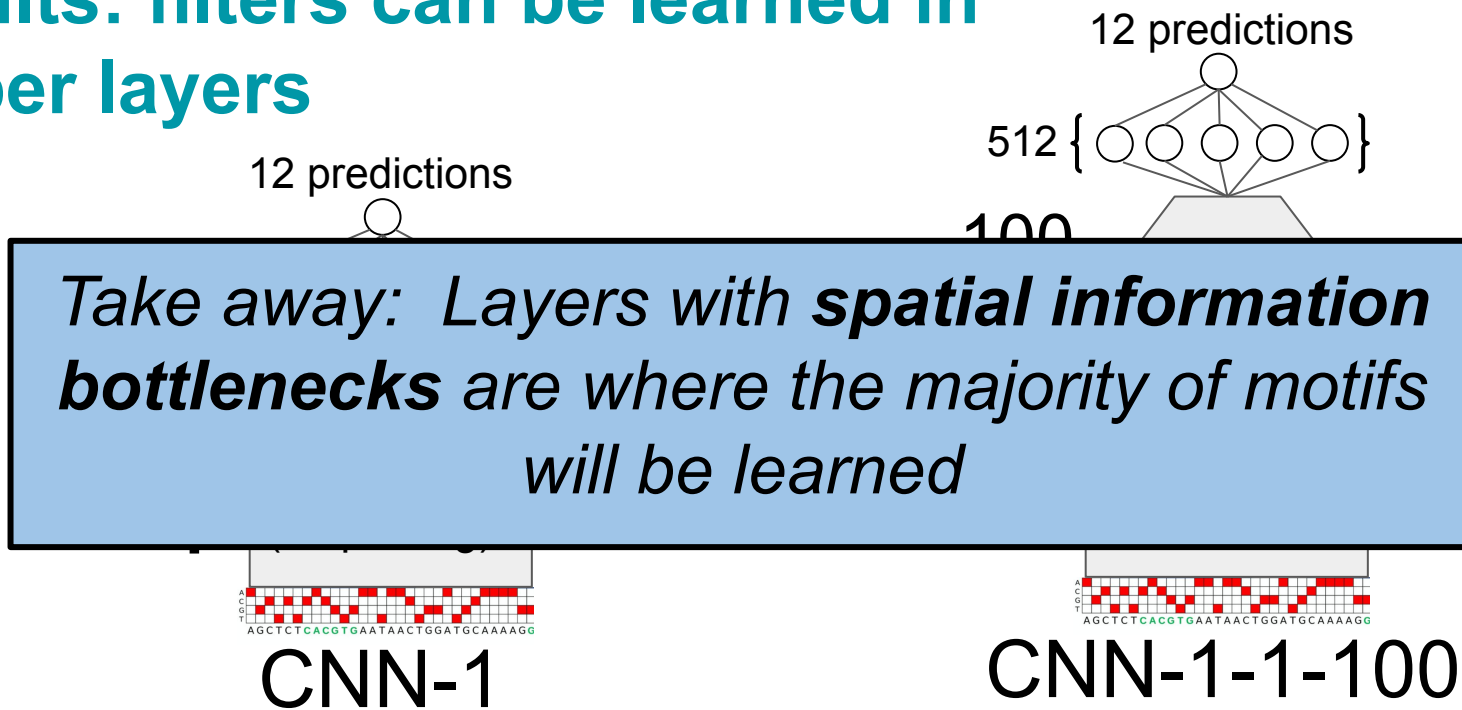
Still Small pool



# Results: Restricting deeper layer assemblies

<b>Model</b>	<b>Average AU-ROC</b>	<b>% Motif match (JASPAR)</b>	<b>% Motif match (Relevant)</b>
CNN-2	0.966±0.000	0.240±0.071	0.007±0.013
CNN-50-2	0.921±0.012	0.913±0.050	0.893±0.044
CNN <sub>19-1</sub> -2	0.969 ±0.002	0.867±0.056	0.747±0.096

# Results: filters can be learned in deeper layers




	<u>First layer filters</u>	<u>Second layer filters</u>
AU-ROC:	0.972	0.972
% JASPAR:	0.240	<b>0.900</b>
% Relevant:	0.007	<b>0.847</b>

	<u>First layer filters</u>	<u>Second layer filters</u>	<u>Third layer filters</u>
AU-ROC:	--	--	--
% JASPAR:	0.147	0.192	<b>0.927</b>
% Relevant:	0.000	0.006	<b>0.891</b>

# Results: motifs are learned at the information bottleneck

Model	Average AU-ROC	% Motif match (JASPAR)	% Motif match (Relevant)
CNN-1	0.972±0.001	0.240±0.083	0.007±0.013
CNN-1 second layer filters →		<b>0.900 +/- 0.024</b>	<b>0.847 +/- 0.021</b>

CNN-1-1-100



CNN-1-1-100

CNN-1-1-100 first layer filters →	0.147 +/- 0.045	0.0 +/- 0.0
CNN-1-1-100 second layer filters →	0.192 +/- 0.022	0.006 +/- 0.006
CNN-1-1-100 second layer filters →	<b>0.927 +/- 0.020</b>	<b>0.891 +/- 0.030</b>

*Take away: Layers with spatial information bottlenecks are where the majority of motifs will be learned*




# Results: *In Vivo* Generalizations

Nat Methods. 2015 Oct;12(10):931-4. doi: 10.1038/nmeth.3547. Epub 2015 Aug 24.


## Predicting effects of noncoding variants with deep learning-based sequence model.

Zhou J<sup>1,2</sup>, Troyanskaya OG<sup>1,3,4</sup>. **Cited ~800 times!**

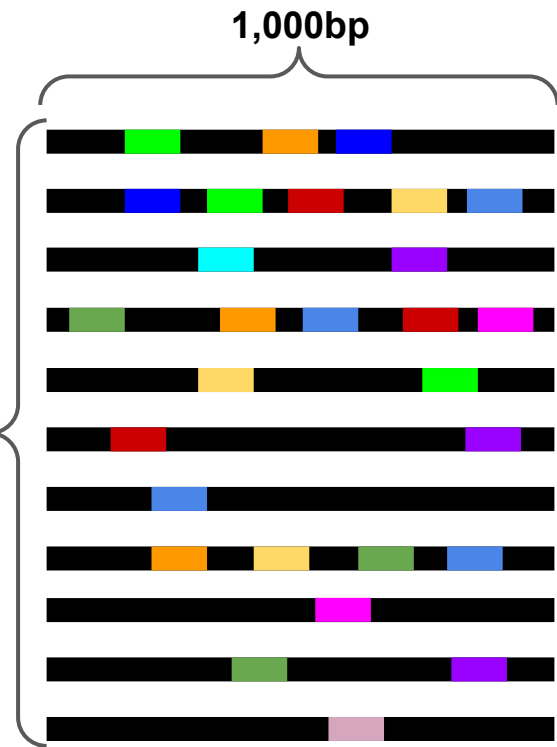
### Human ChIP-seq data



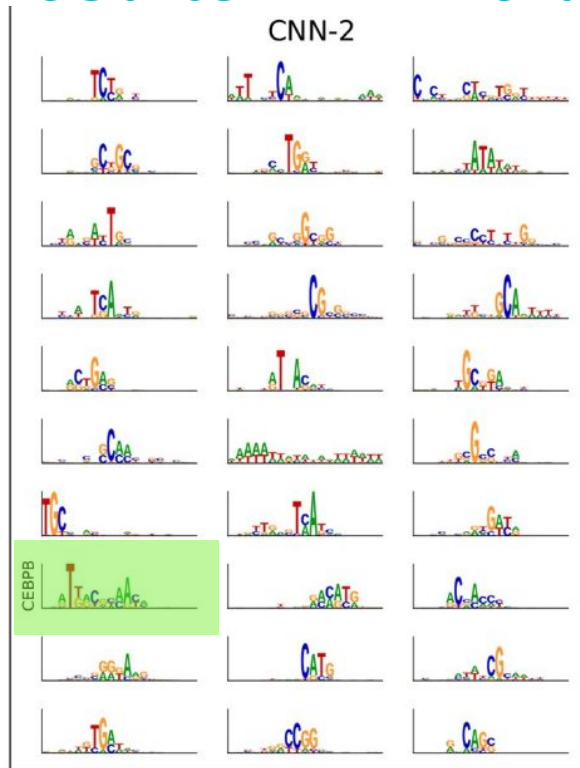
### 12 JASPAR motifs



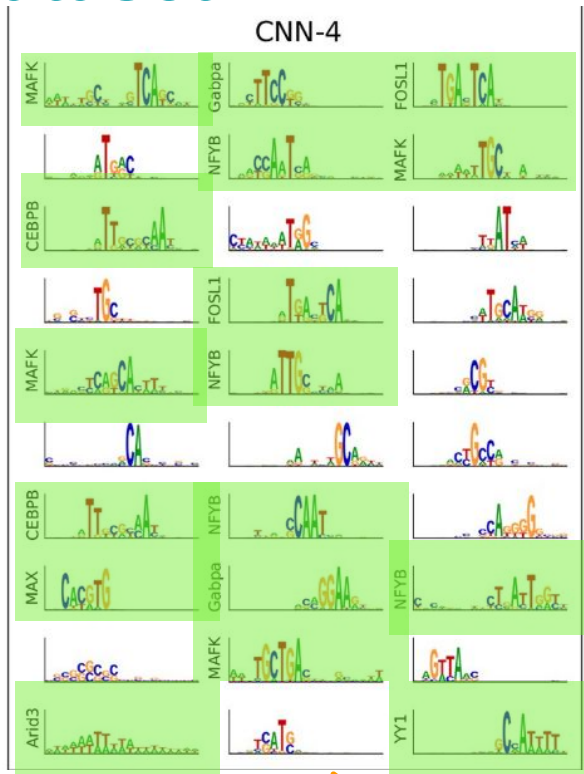
300K  
seqs



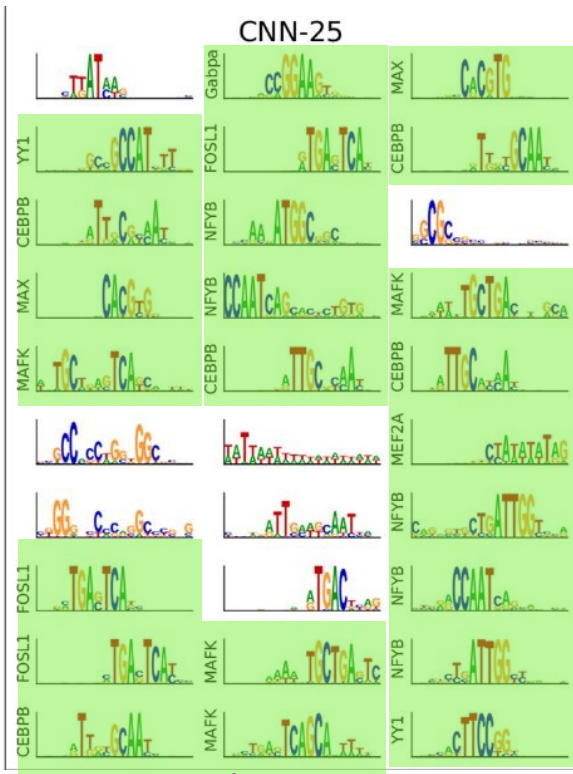
# Results: *in vivo* dataset



Small  
pool



Small-ish  
pool



Large-ish  
pool

Table 2. Performance of deep learning models on the *in vivo* dataset.

Model	Average AU-ROC	Average AU-PR	Motif match (JASPAR)	Motif match (Relevant)
CNN-1	0.918±0.001	0.626±0.000	0.227±0.068	0.020±0.027
CNN-2	0.911±0.003	0.609±0.006	0.333±0.067	0.107±0.057
CNN-4	0.907±0.002	0.601±0.005	0.753±0.045	0.507±0.039
CNN-10	0.903±0.006	0.583±0.020	0.920±0.045	0.753±0.034
CNN-25	0.903±0.003	0.580±0.009	0.933±0.030	0.747±0.040
CNN-50	0.903±0.003	0.582±0.009	0.913±0.034	0.733±0.063
CNN-25 (90)	0.919±0.001	0.628±0.005	0.940±0.023	0.909±0.028
CNN-25 (120)	0.920±0.002	0.637±0.005	0.933±0.015	0.887±0.025

*Take away: Architectures may need more filters to perform better on **in vivo** sequences*

On Synthetic Data:

CNN-25 had the best Relevant match w/ 0.980 +/- 0.027

Now:

CNN-25: 0.747 +/- 0.040 Relevant match

CNN-25 (60): has the best Relevant match performance with 0.960 +/- 0.023

# Results Summary

- CNN architecture choices affect how motifs are learned
  - **Wider pooling** size forces first layer filters to learn whole motifs
  - Filter **number** and filter **size** are less influential
  - **Restricting hierarchical assembly** in deeper layers can increase first layer motif learning
  - Motifs are learned at the **information bottleneck** (can be 1st, 2nd, 3rd layer)
  - With *in vivo* dataset **more filters helped** with distributed representation learning

# Overview

- Background
  - How do CNNs work?
  - Why do people care about finding motifs?
- Methods
  - Synthetic dataset used in this paper
  - Experimental setup for CNN architectures
  - Model vs Motif evaluation metrics
- Results
  - Pulling various CNN architecture levers!
- **Main takeaways & discussion**



# Main Takeaways & Discussion

- Exploration of various CNN architectures to better understand **how** and **where** CNNs learn motifs
  - *Was this a useful aspect to explore?*
- % of 1st layer filters that learn motifs is *not necessarily a useful metric* for assessing biological relevance because CNNs can **assemble partial motifs in deeper layers**
  - *Do you agree? Would you still want this reported?*
- If you want to **enforce** that your CNN learns whole motifs in the **1st layer**, be mindful of your architecture
  - *Would you consider doing this intentionally in your own work?*

Thanks!



Second Beach, La Push, WA