Brain evolution as a machine learning algorithm

Alex Koulakov

Cold Spring Harbor Laboratory

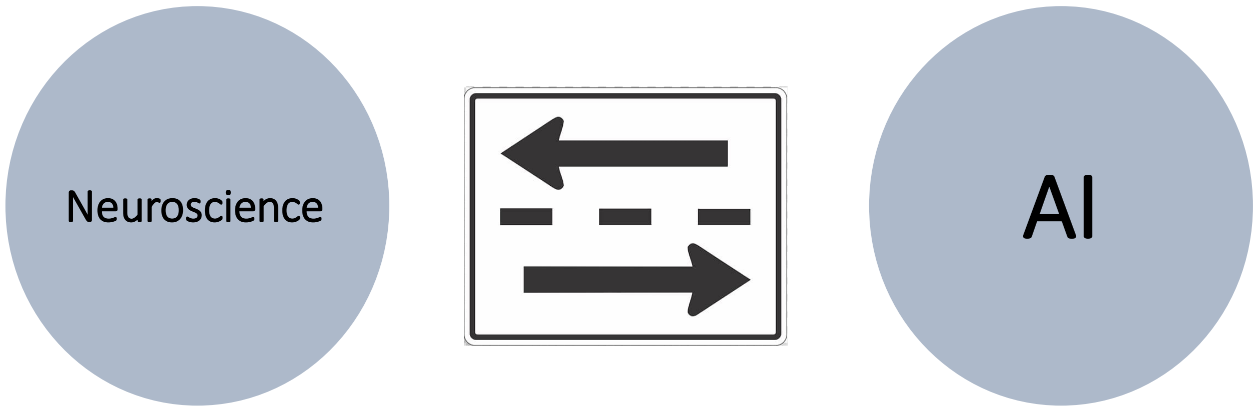
Neuroscience
AI

Neuroscience vs. AI

Neuroscience	AI
Hebbian learning (local LR)	backpropagation (non-local LR)
each neuron has its own set of weights	CNNs (weight pooling)
no translation invariance (complex logarithm maps)	translation invariance
computes with spikes	computes using analog #'s
20W	2000W
connections are instructed by chemical labels	connections are trained
essential behaviors are innate	behaviors are learned

Are these bugs or features?

Examples of innate behaviors

newborn turtles

walking

spider web

burrowing in peromyscus

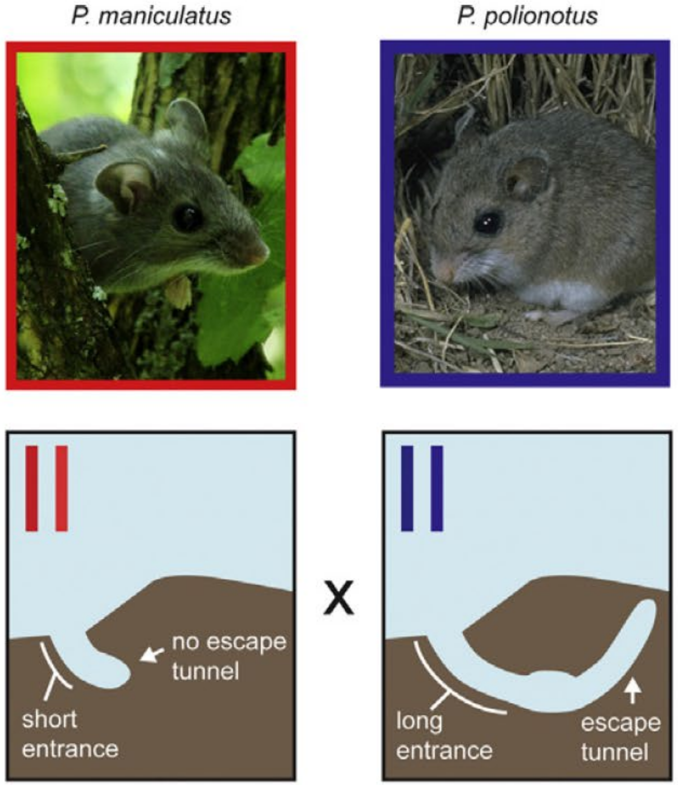

Hopi Hoekstra

flying

capacity for language

Genomic bottleneck approach to learning

Tony Zador

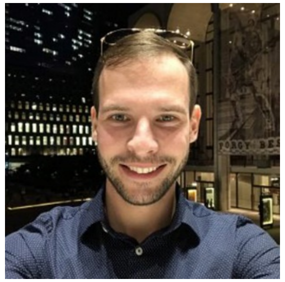

Sergey Shuvaev

This work is supported by Deep Valley Labs, San Jose, CA

Cerebral cortex is a thin sheet of gray matter occupying ~2 sq ft

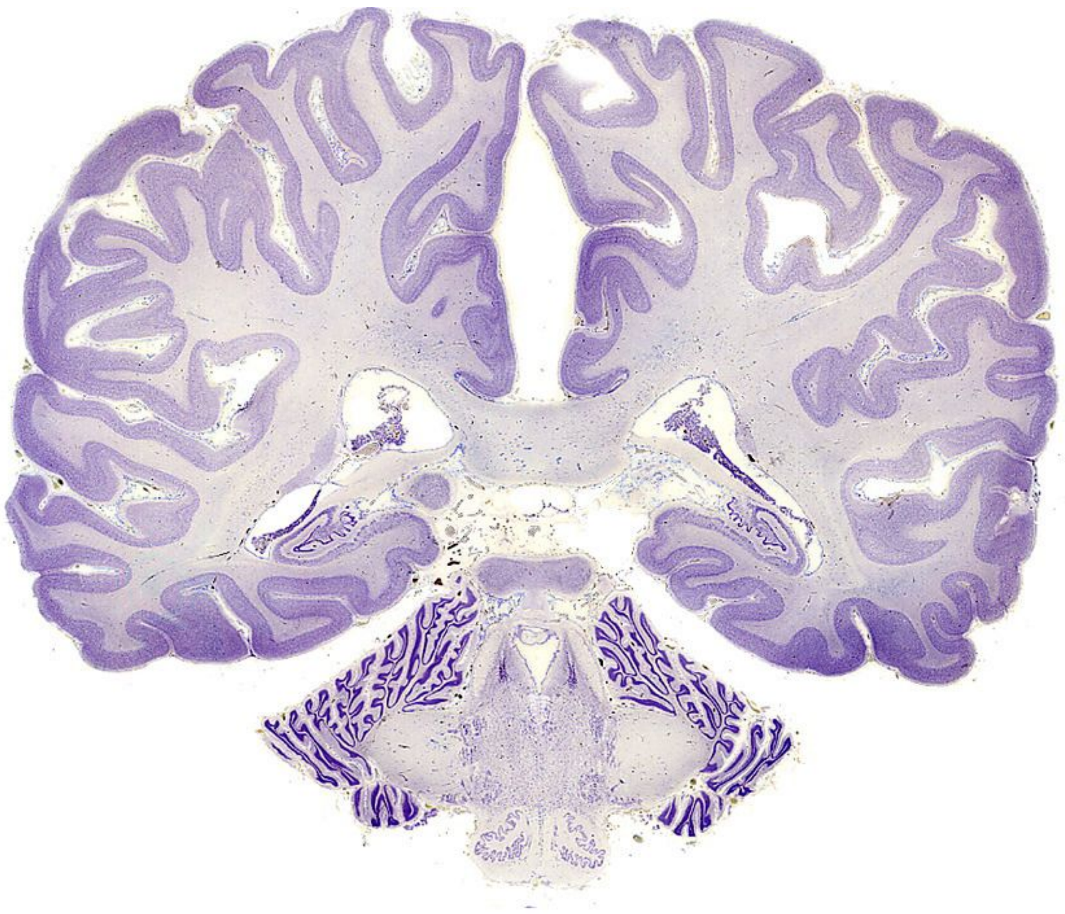

Cortex contains $\sim 10^{10}$ neurons forming about 10^4 synapses each

How much information can be stored in connections?

$N \sim 10^{10}$ neurons in cortex

$s \sim 10^{4}$ synapses per neuron

$$H = Ns \log_2 N \sim 400 \text{ terabytes}$$

$$\sim 45 \text{ years of HD video}$$

Wei, Tsigankov, Koulakov, Annals of NYAS (2013)

Neural development expands data

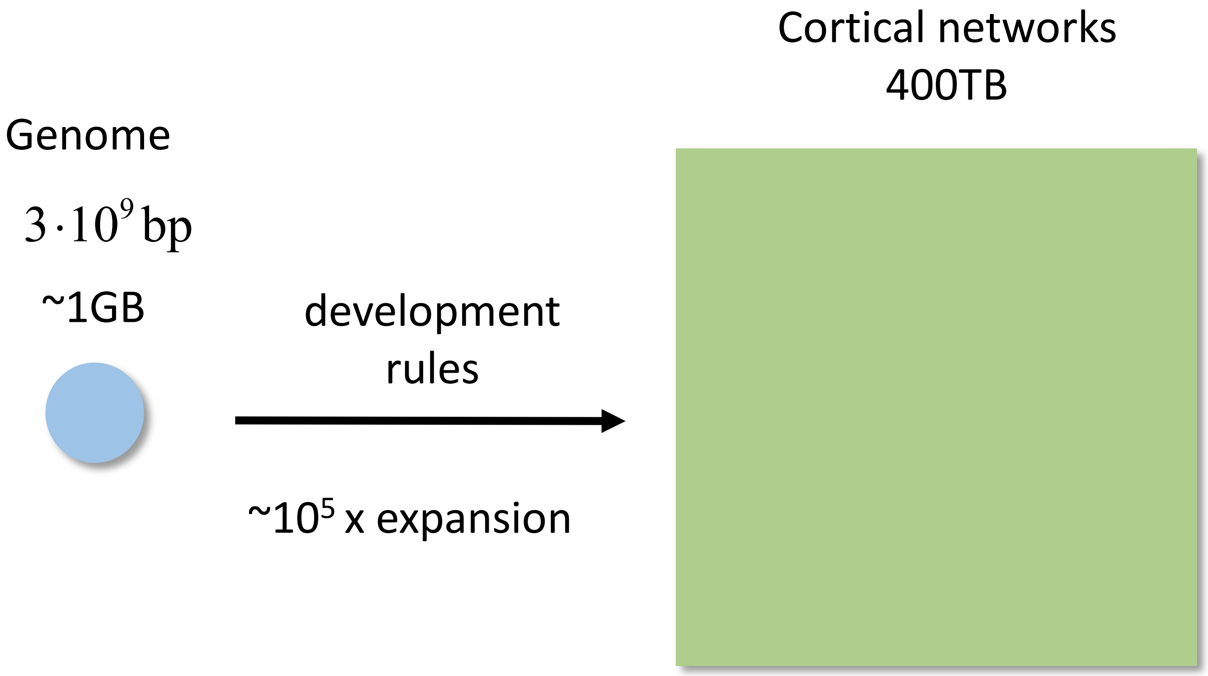

Sperry, PNAS (1963)

Wei, Tsigankov, Koulakov, Annals of NYAS (2013)

Zador, Nature Commun. (2019)

How can 1GB of information set up 400 TB of connections?

Obviously, each synapse cannot be specified in the genome individually

Some simplifying rules are necessary

Developmental mechanisms are forced to find simplifying rules or “organizing principle” for brain networks

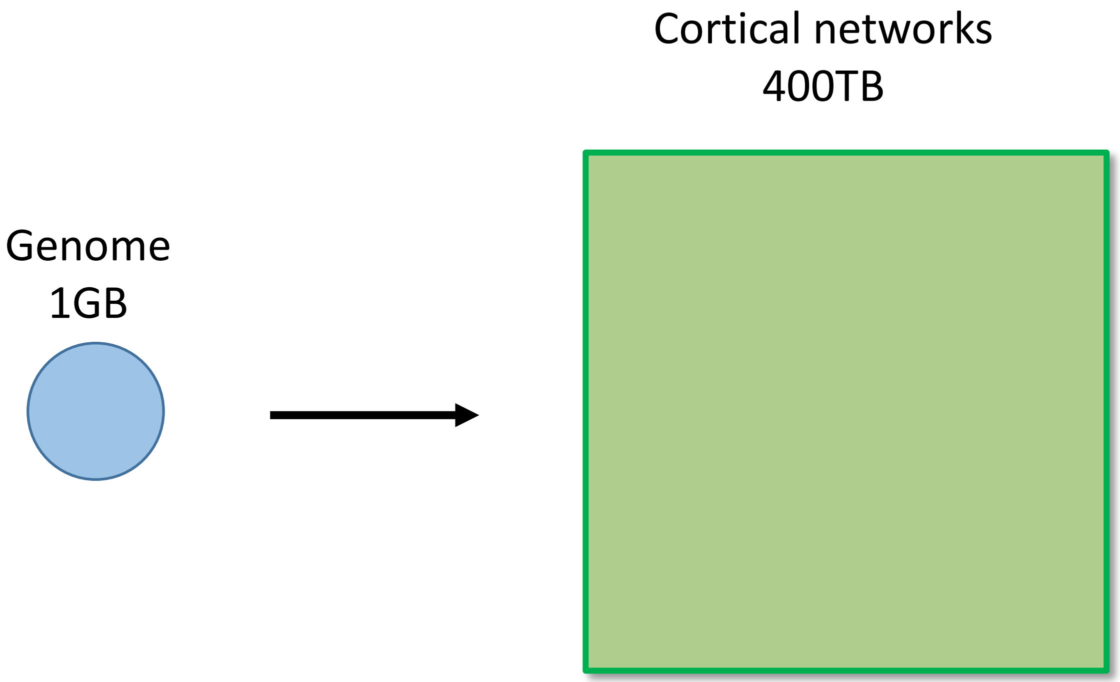

Genomic bottleneck principle:

A: Neurodevelopmental rules contained in the genome (1GB) contain information about the capacity of humans for intelligent behavior

B: The need to compress information about brain architecture into a small volume (<1GB) gives mammalian brain capacity for general intelligence (*Critique of pure learning* by Tony Zador).

Implementation: Pairwise interactions between neurons generate connectivity in the brain

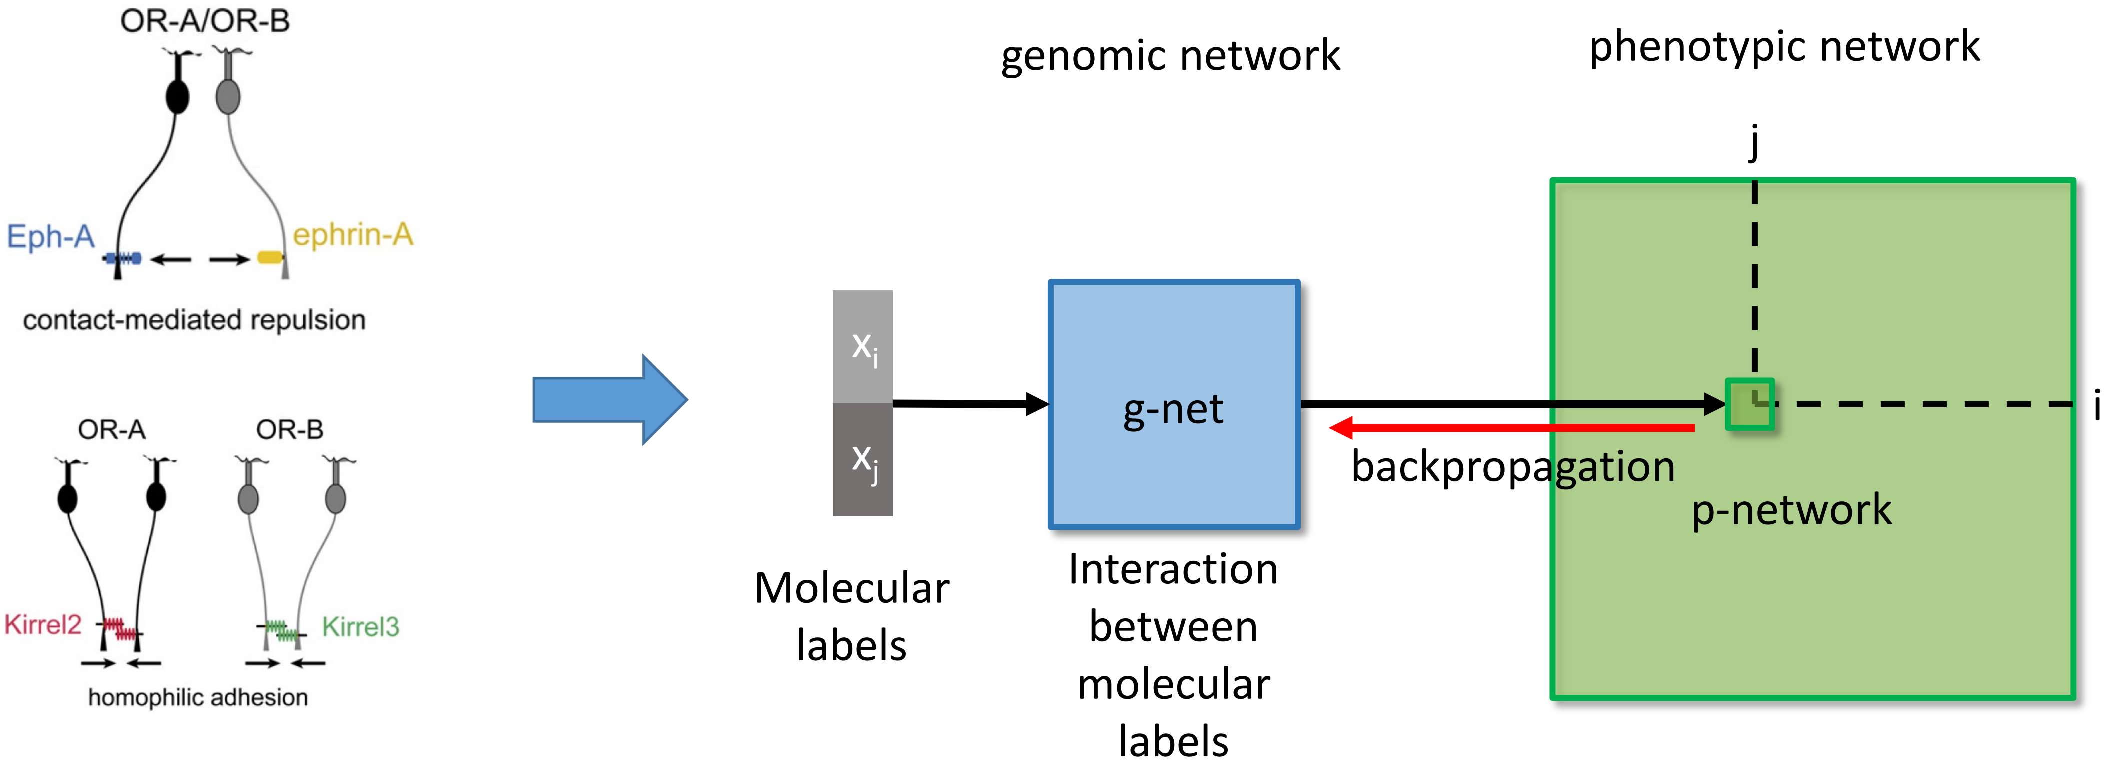

Since molecular labels are not stored in the genome, arbitrary large networks can be encoded

Arrays of molecular labels (x_i and x_j) are represented by binary numbers or locality-based binary Gray codes

Backpropagation ~ evolution

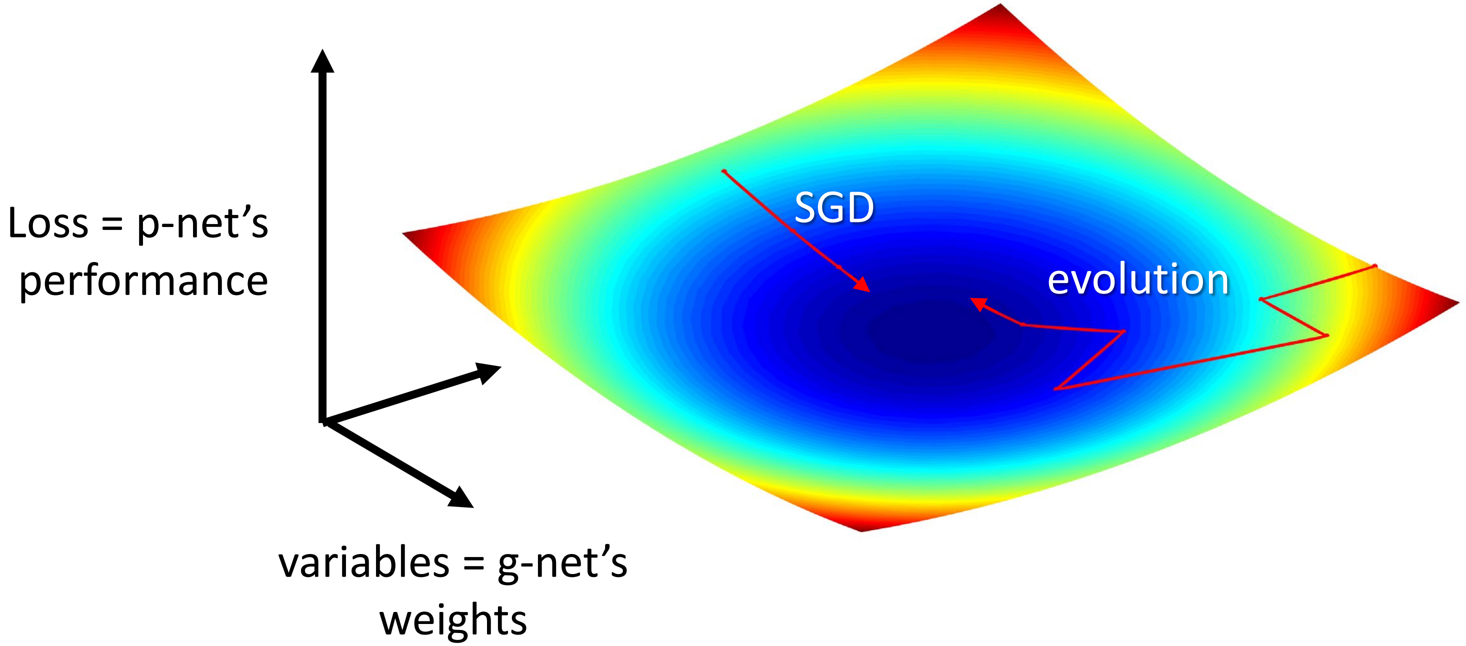

Overall goal is to use g-nets to extract simplifying principles from data

Intermittent co-training of GDN and NN replicates brain evolution

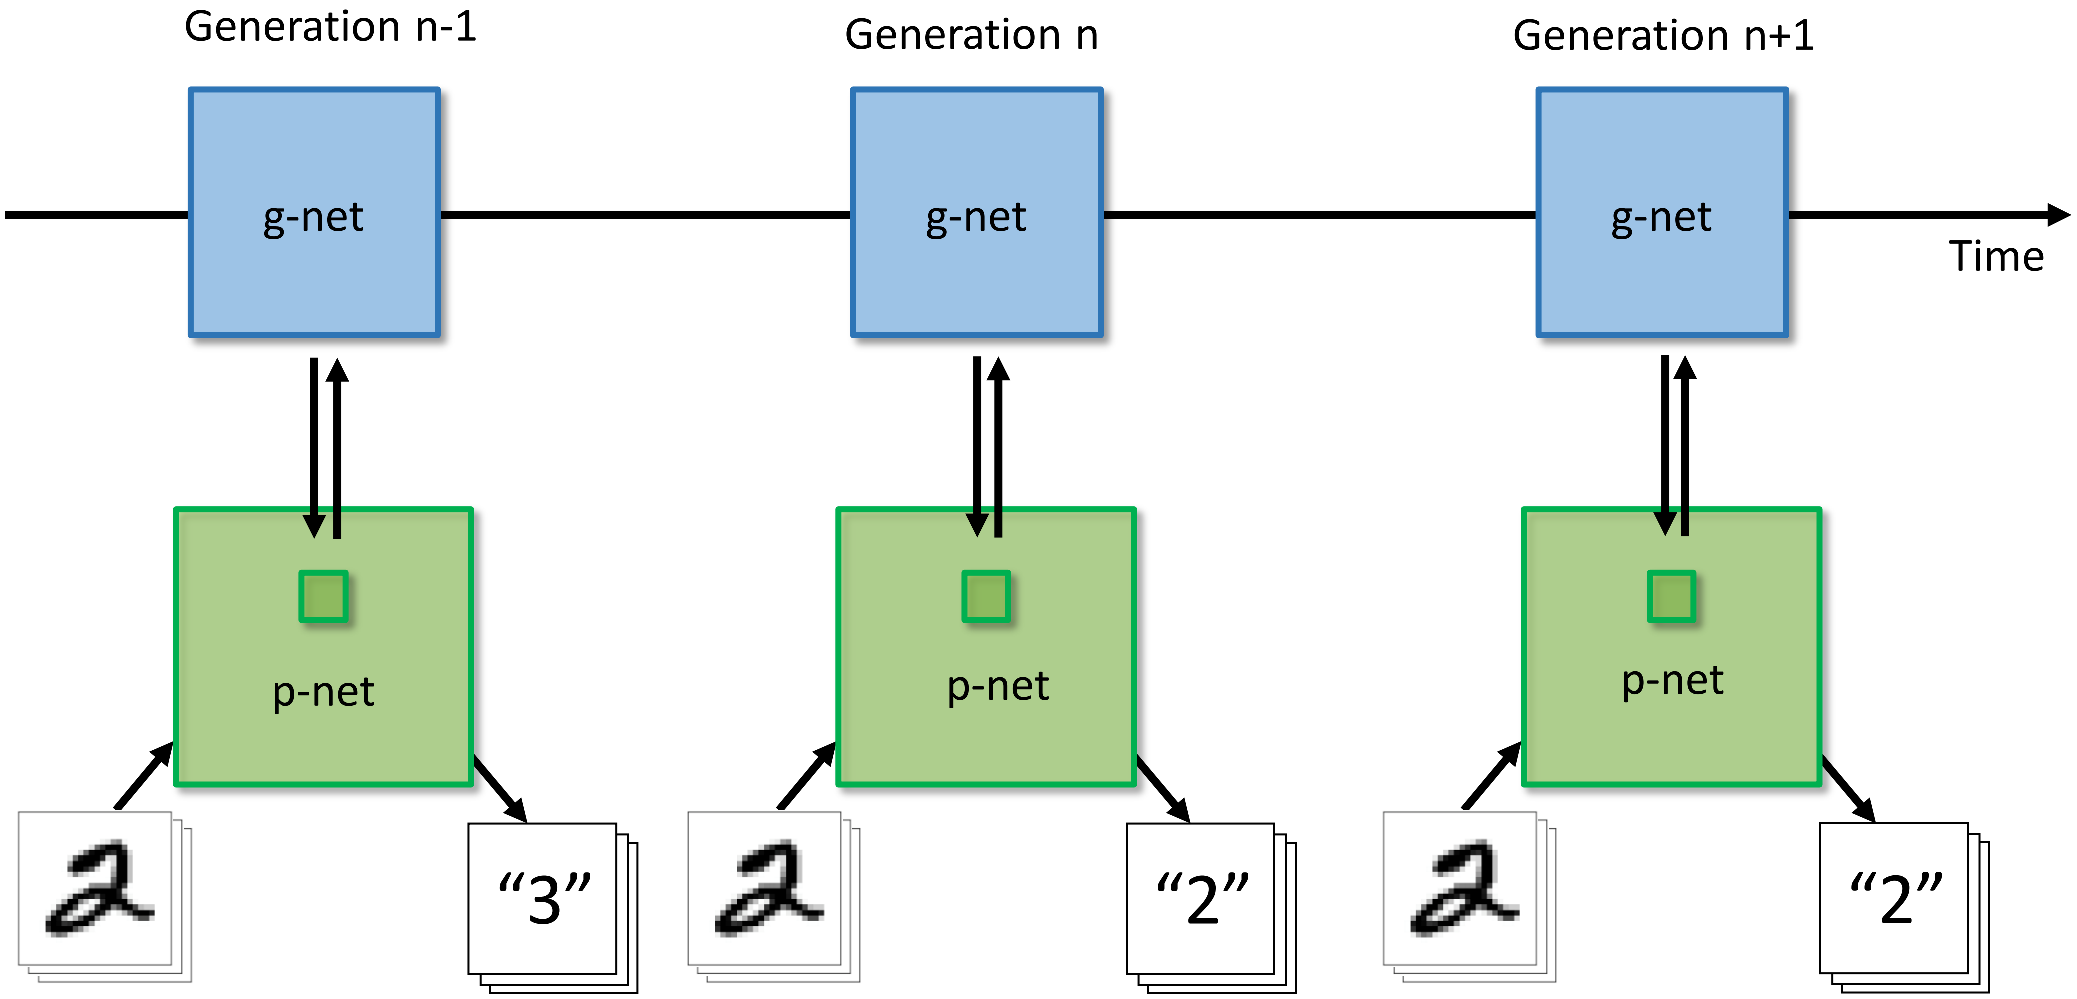

MNIST dataset

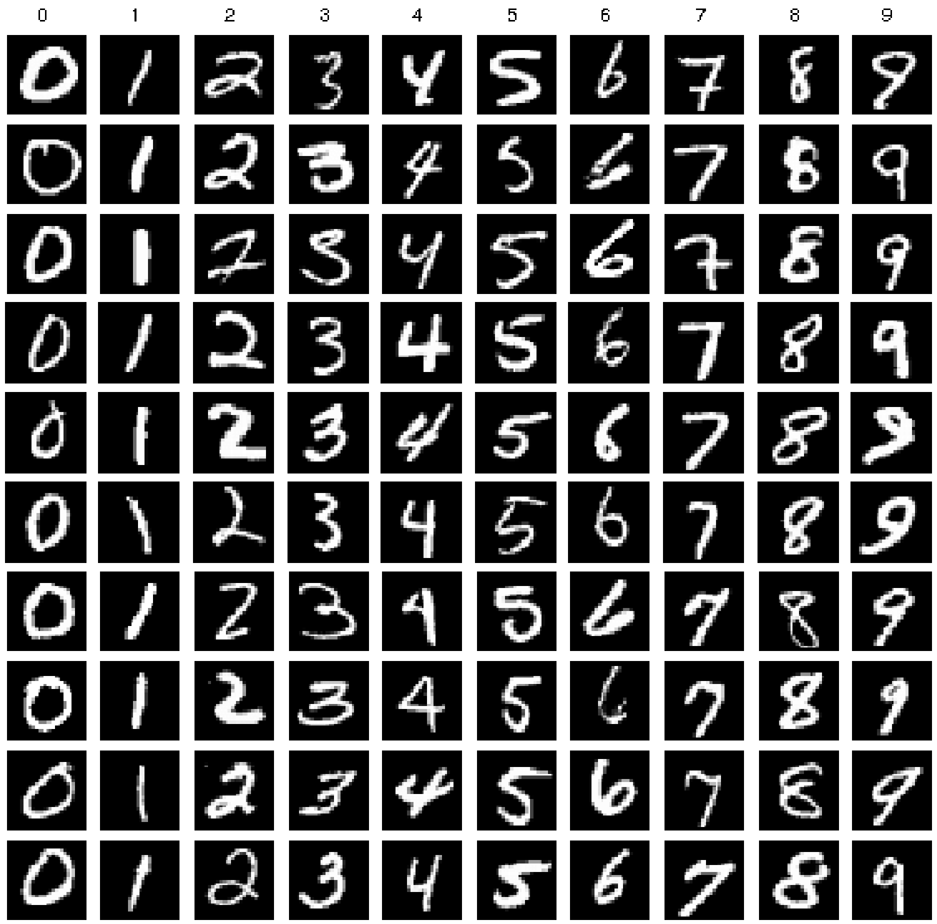

Two-layer MNIST network is encoded by two g-nets

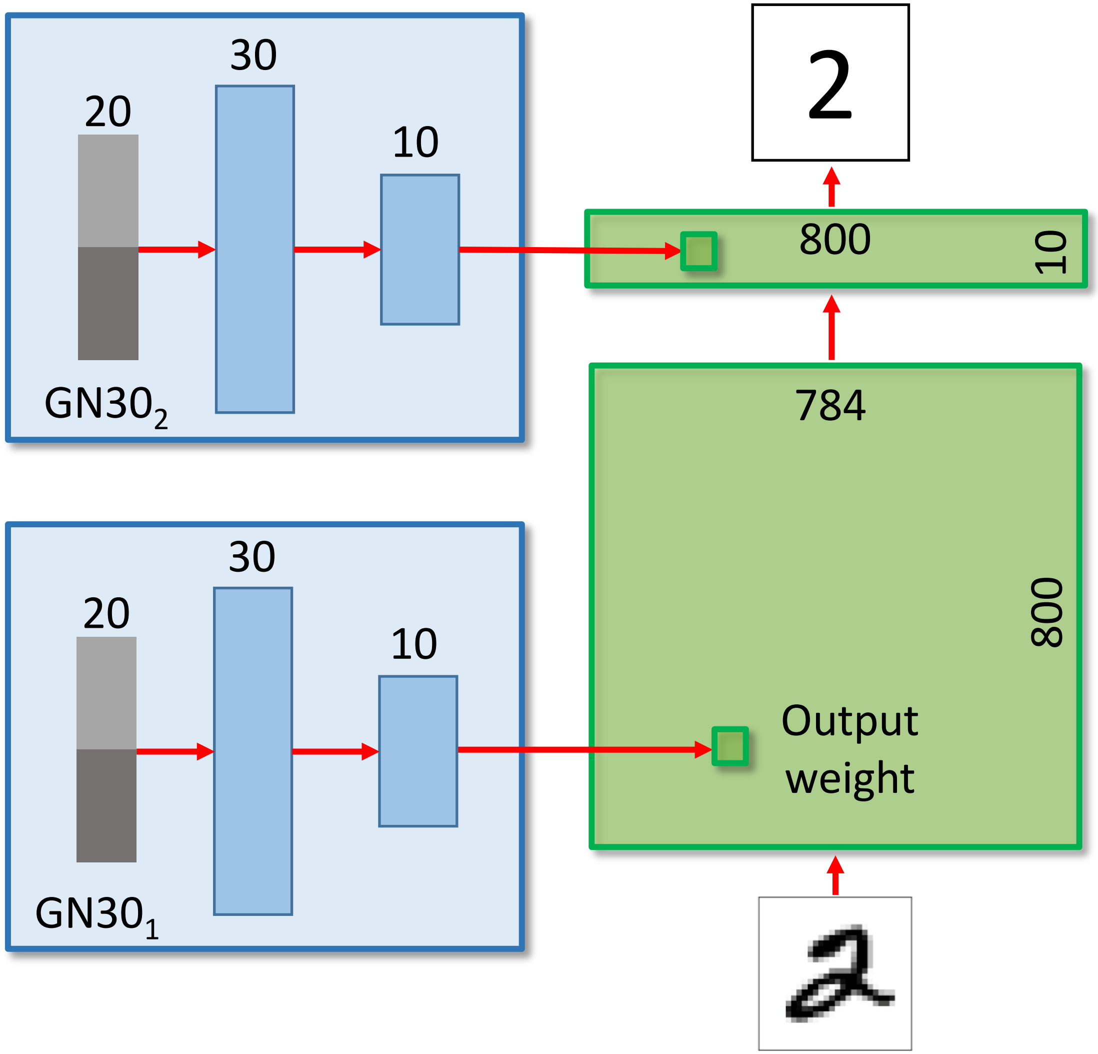

g-nets yield zero-shot learning with 333x compression

Functional network without experience!

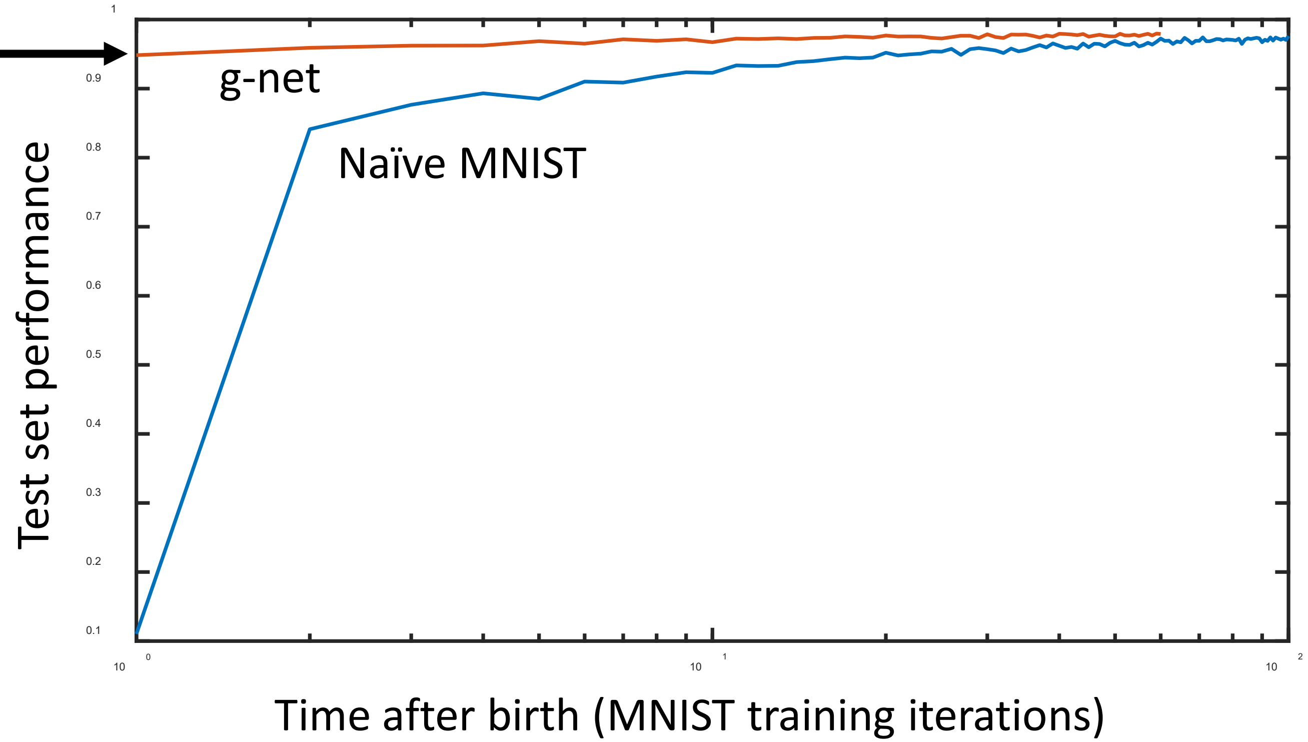

Out of sample weight transfer MNIST-> Fashion MNIST

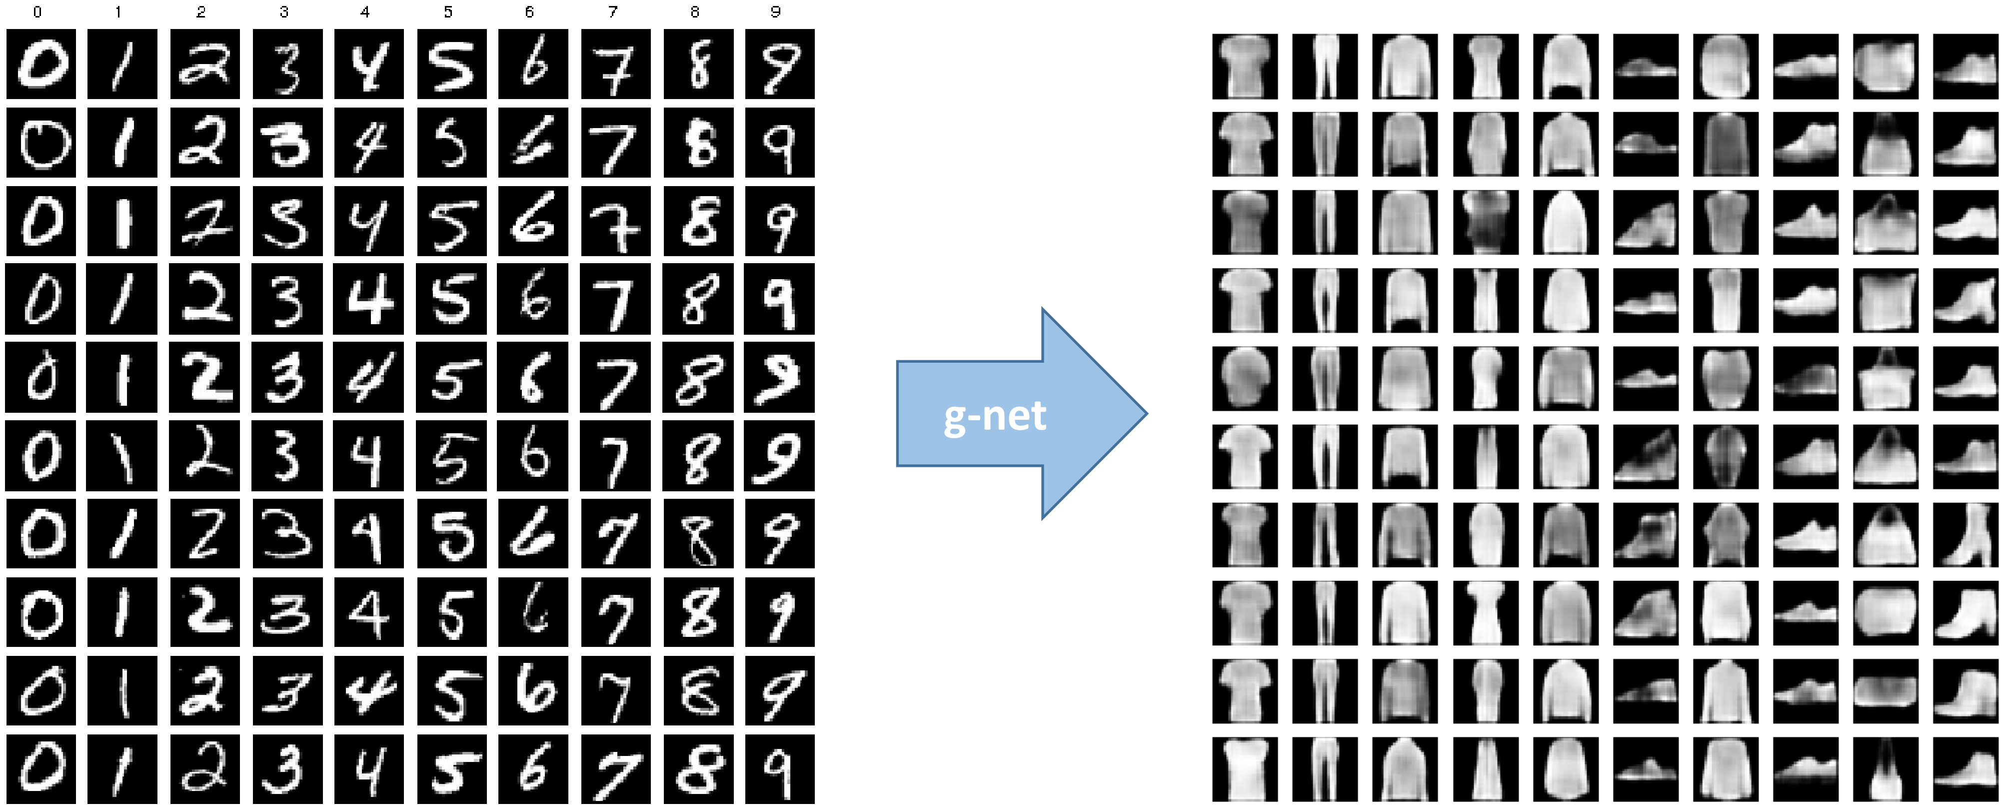

Weight transfer MNIST-> Fashion MNIST

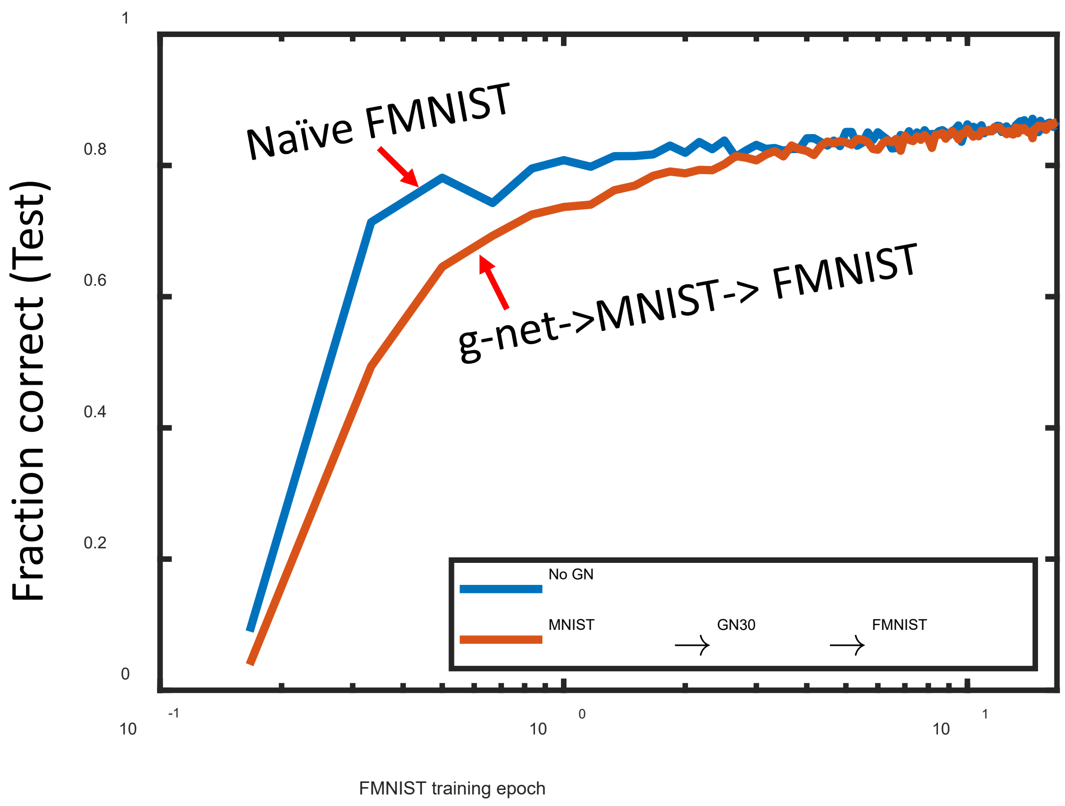

Weight transfer from MNIST makes training on FMNIST dataset slower

g-nets yielded good compression of MNIST network but poor weight transfer performance

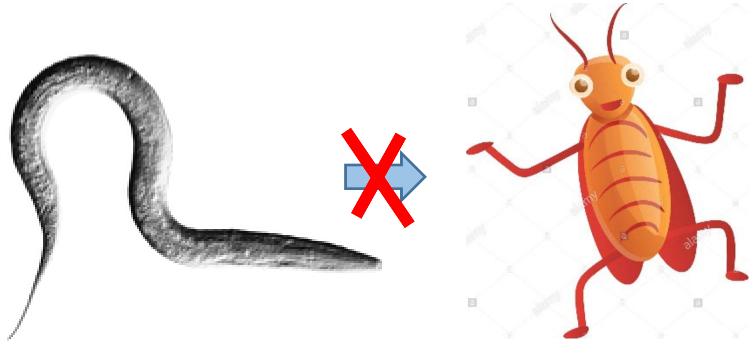

CIFAR10 dataset

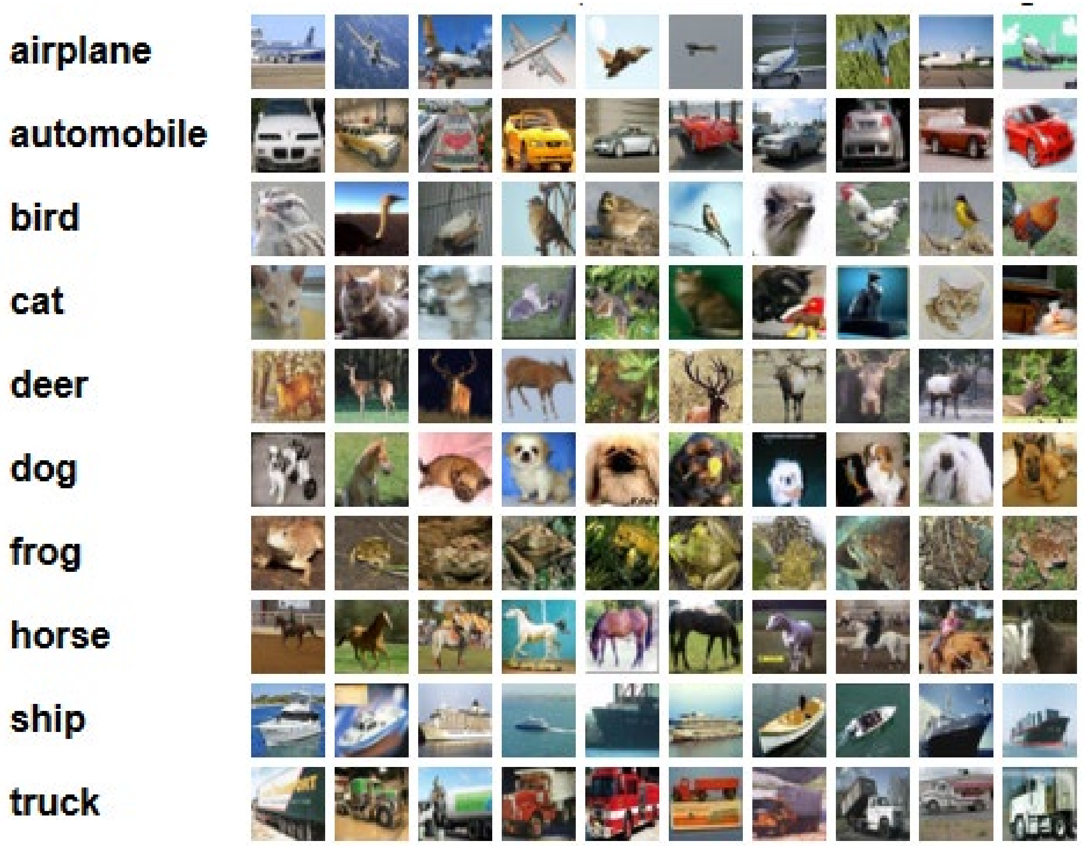

32x32x3 images
(60k/10k)

9-layer all convolutional network:

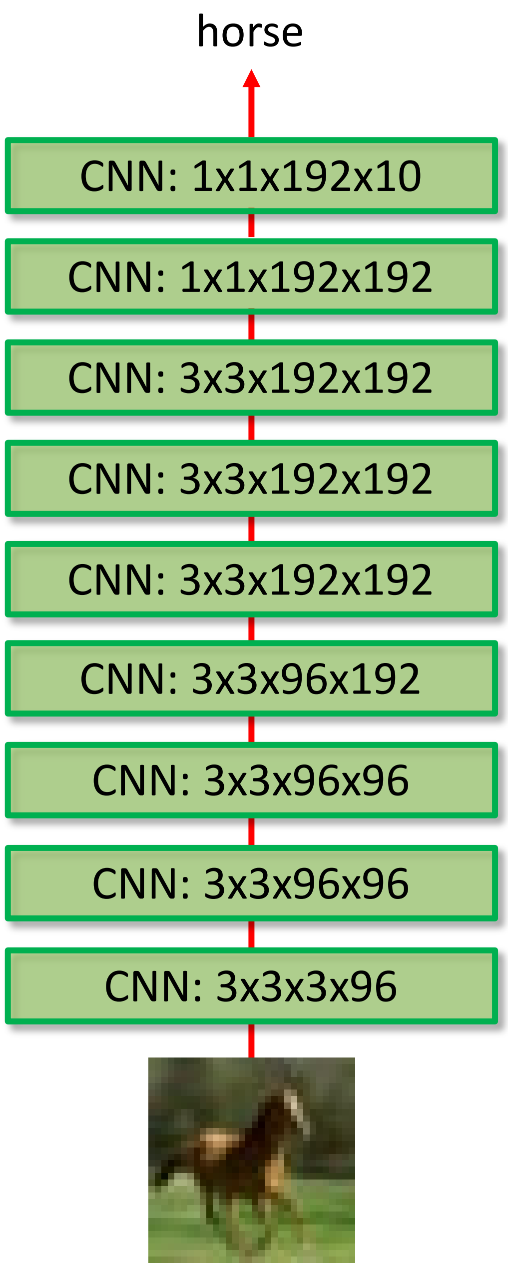

Striving for Simplicity: The All Convolutional Net
arXiv:1412.6806

Co-train g-nets + **nine**-layer CIFAR10

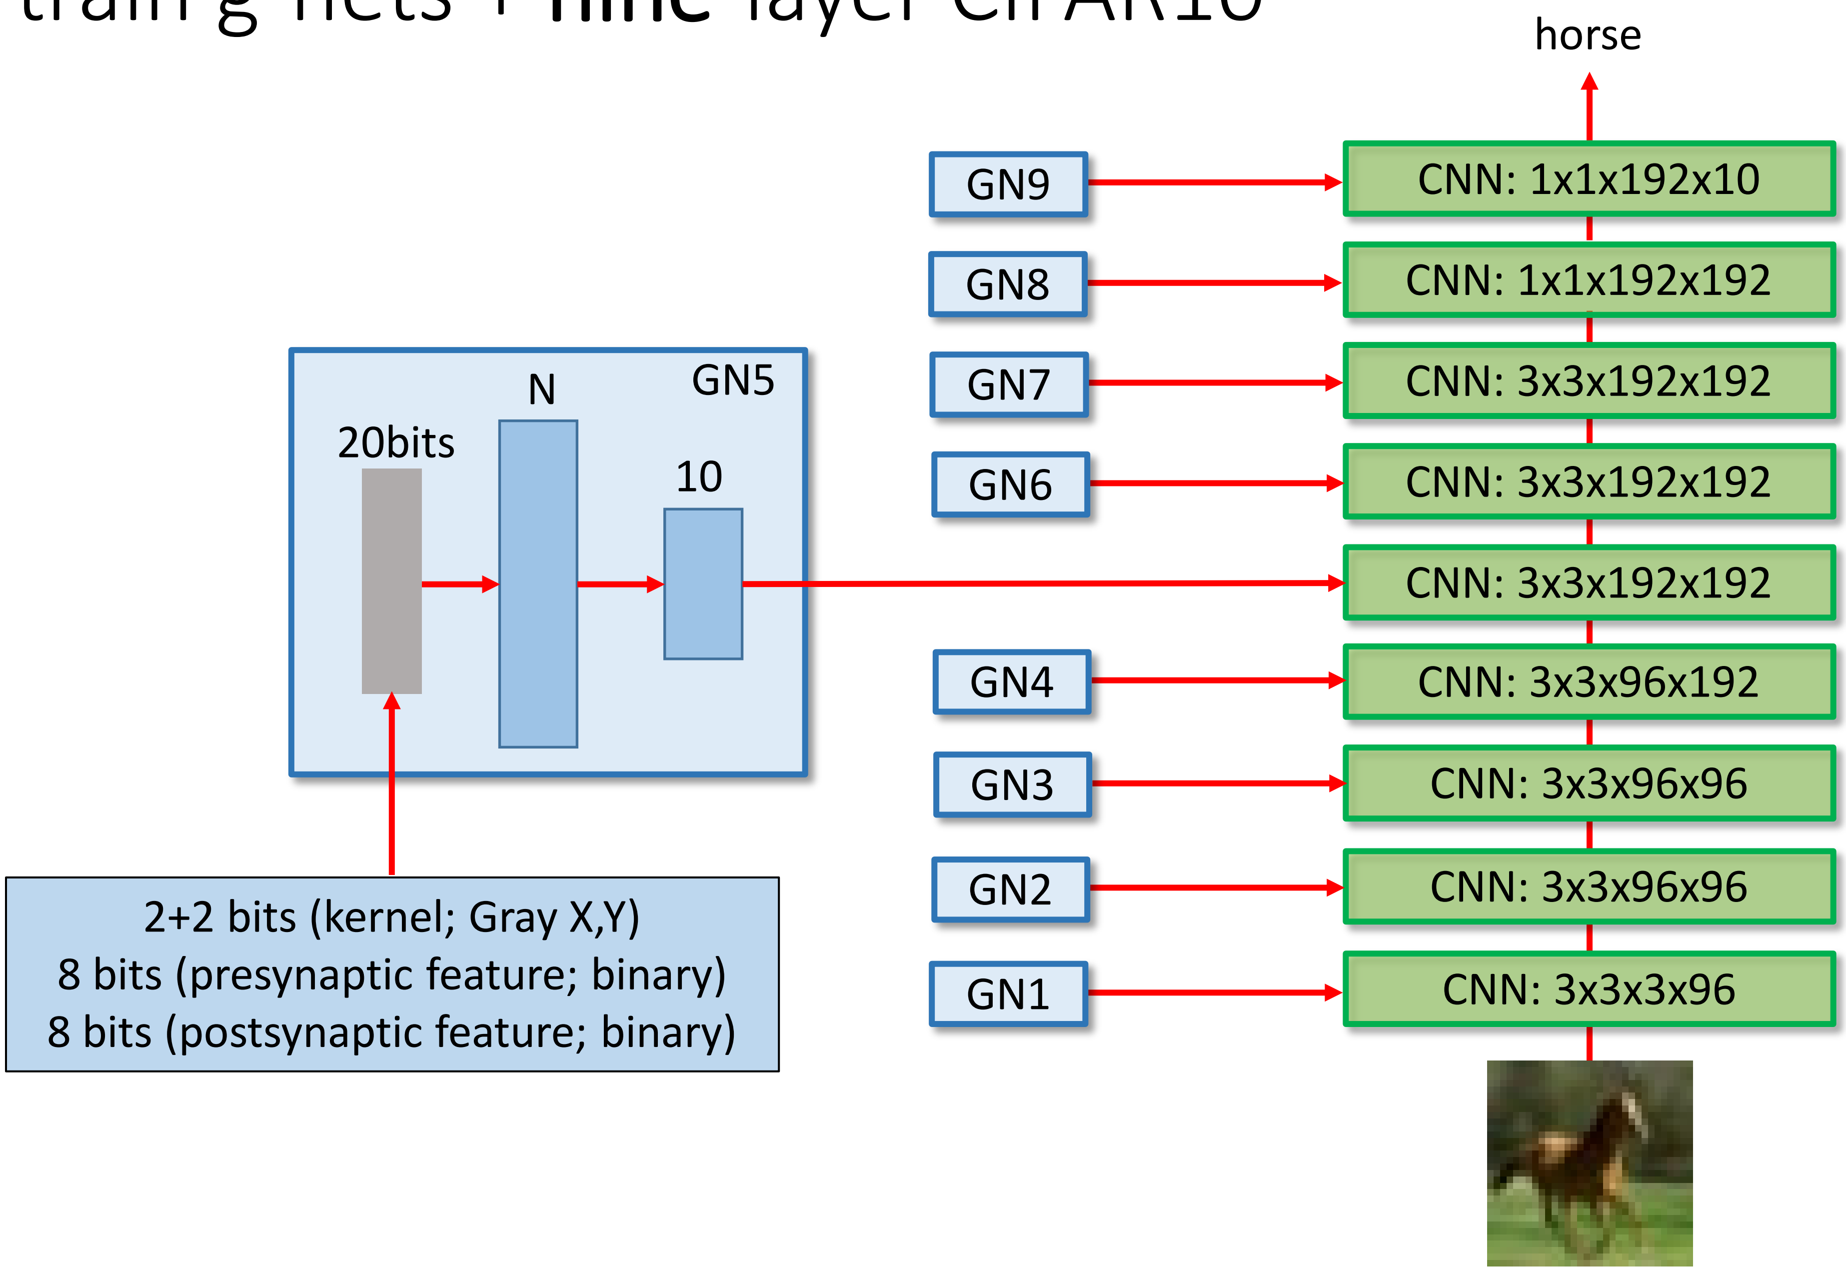

G-nets yield zero-shot performance

Functional network without experience!

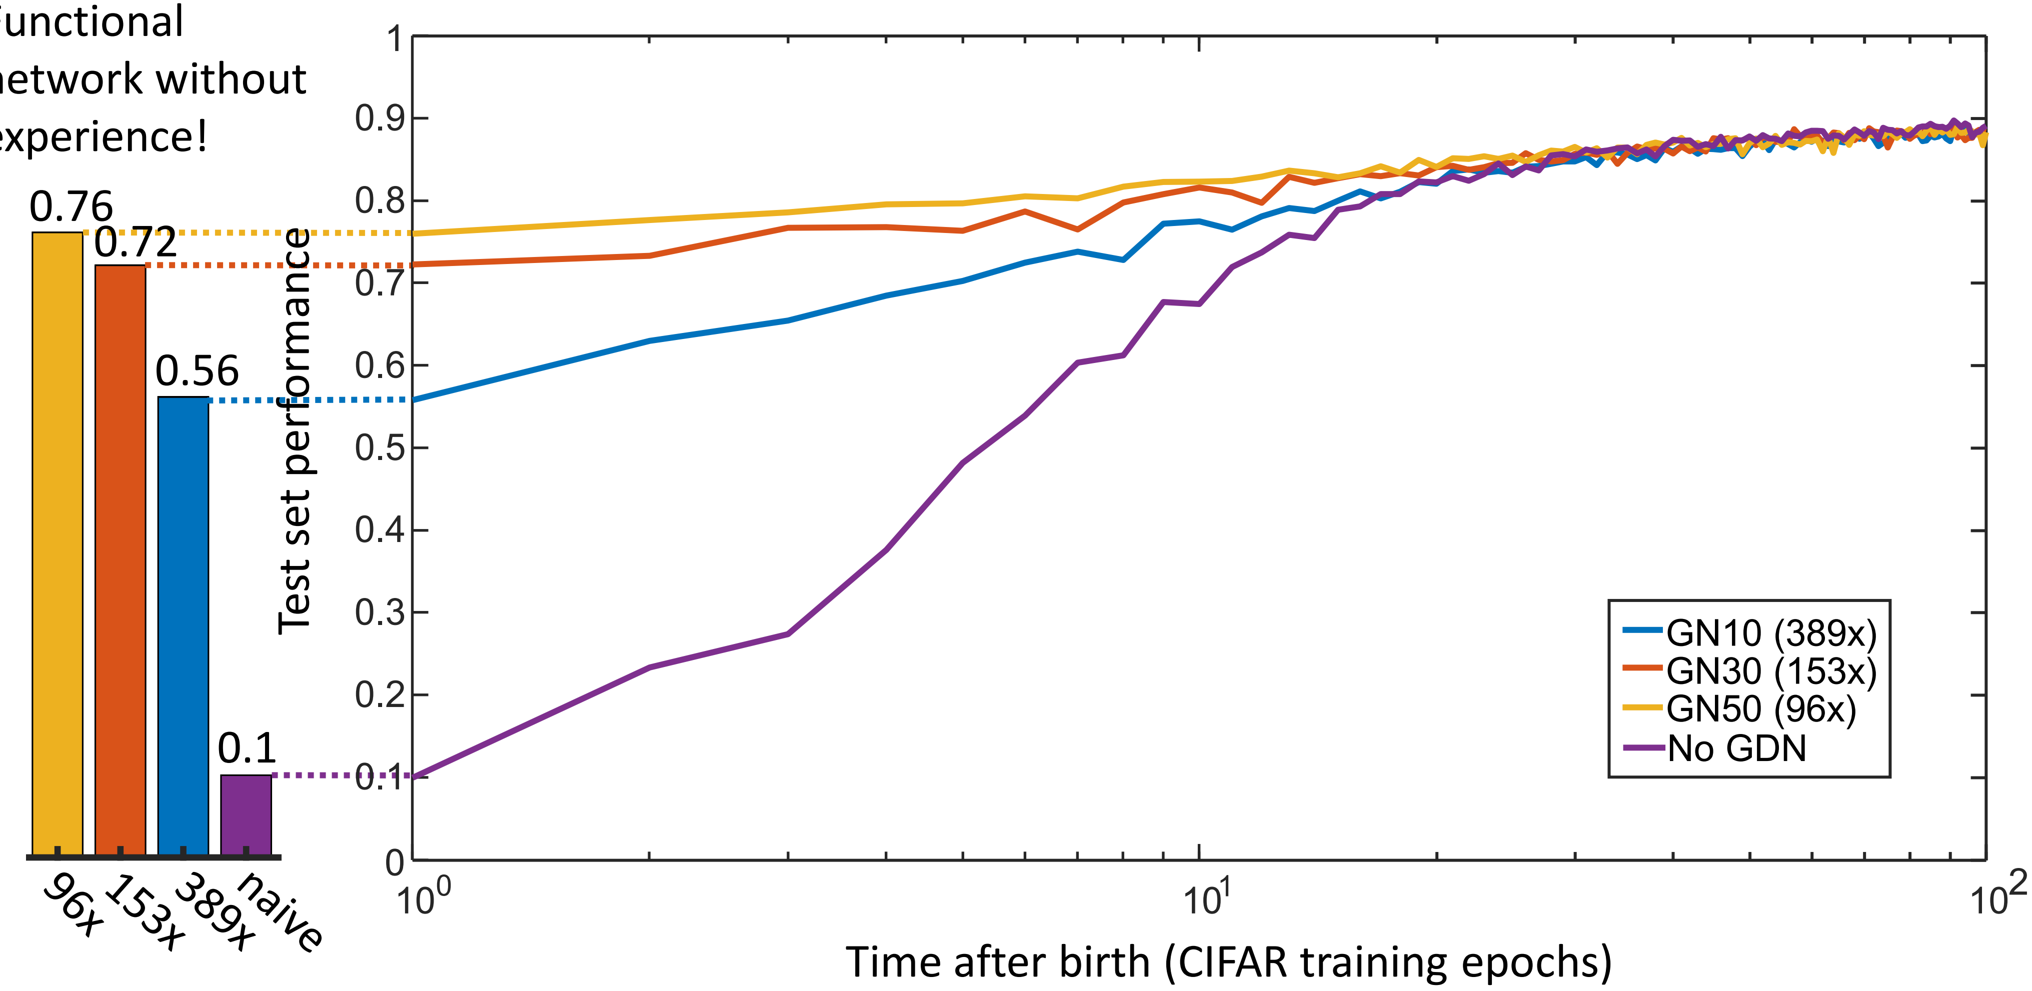

Out of sample weight transfer to the Street View House Numbers (SVHN) dataset

CIFAR10

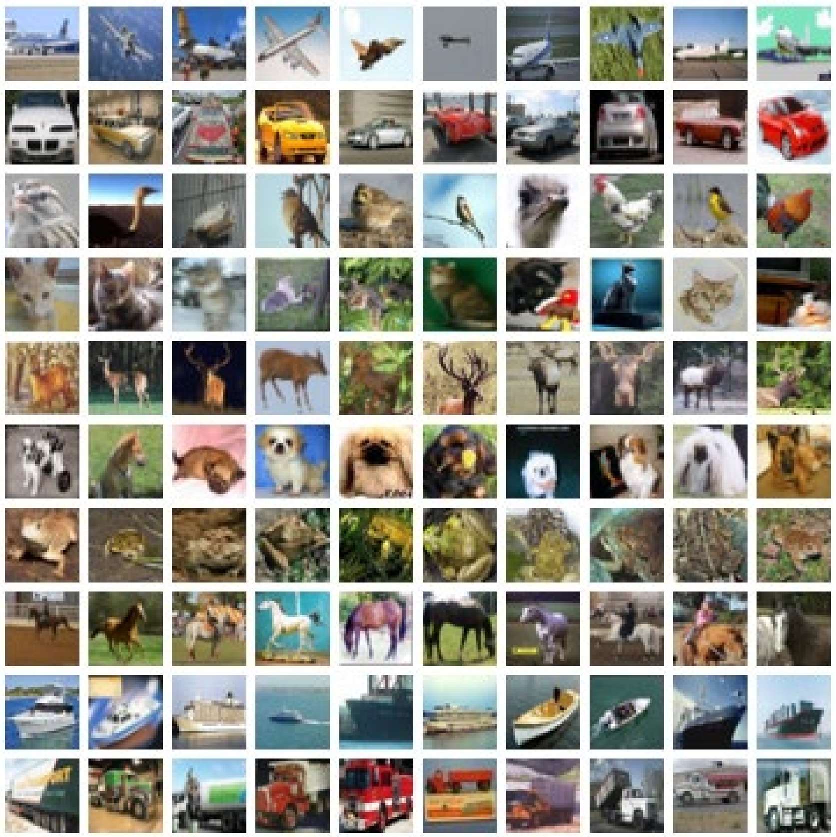

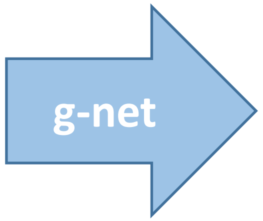

SVHN

GNs outperform basic transfer learning

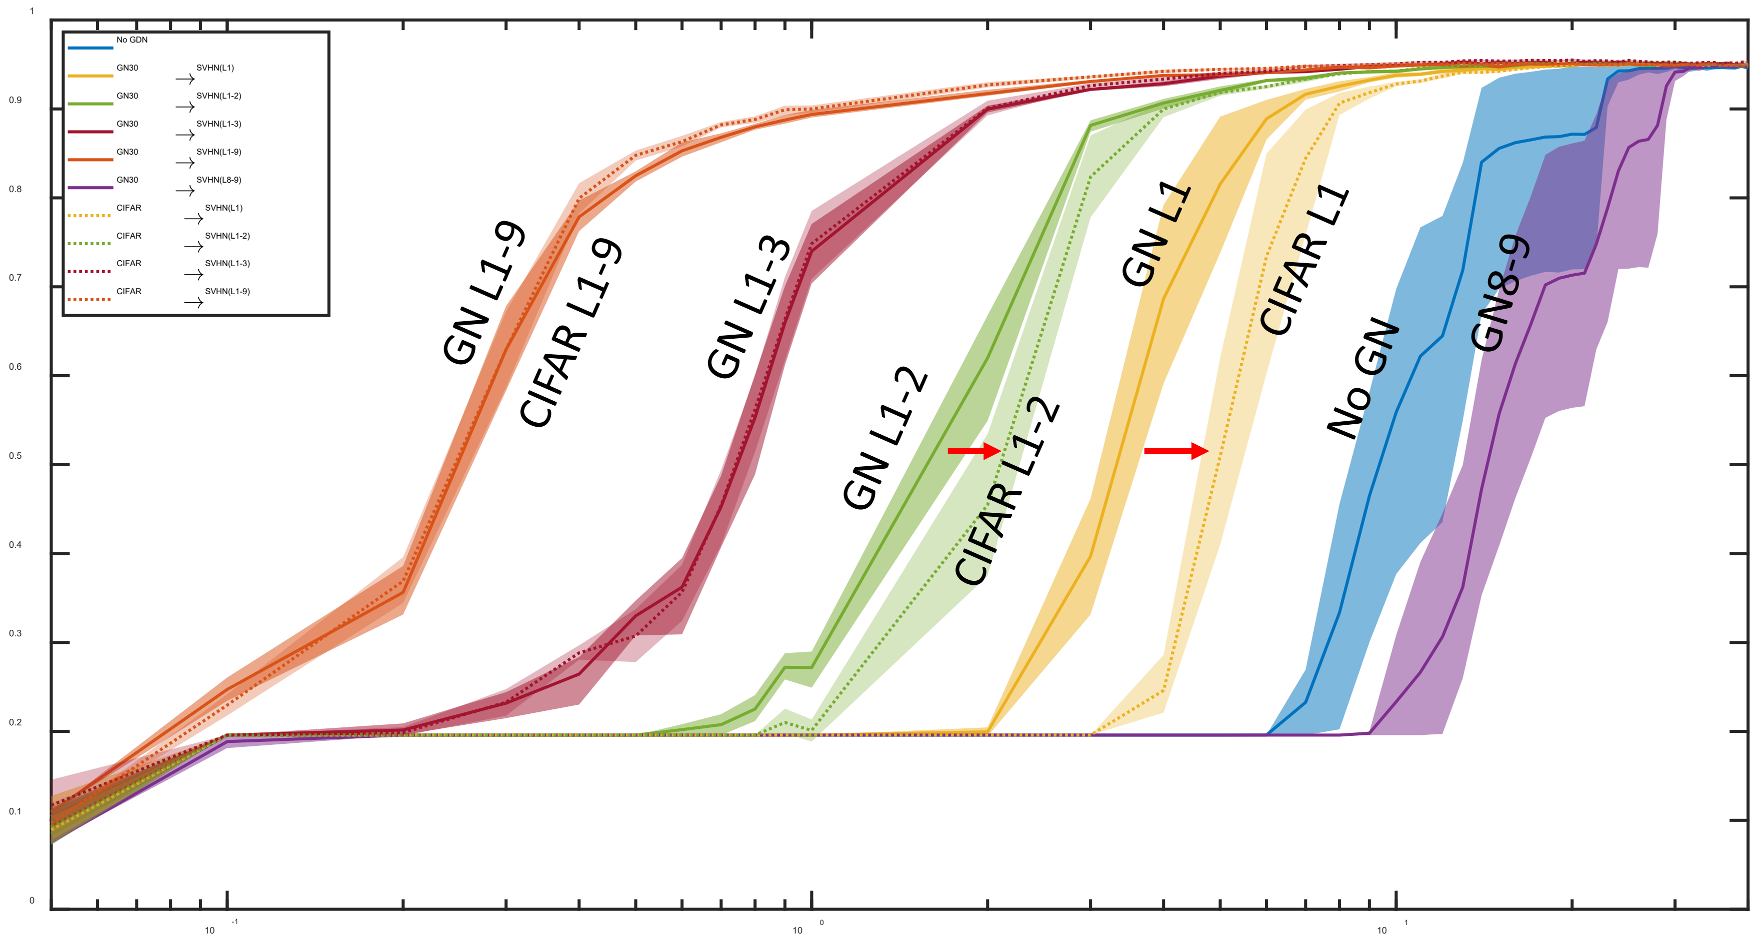

Conclusions: Genomic bottleneck

- g-network manages to find simplifying principle in data that generalizes across datasets
- Genomic bottleneck can force ANN find general principles by imposing an Occam's razor - type constraint

DeepNose: Using artificial neural networks to represent the space of odorants

Tumi Tran

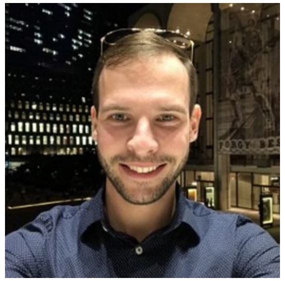

Sergey Shuvaev

Daniel Kepple

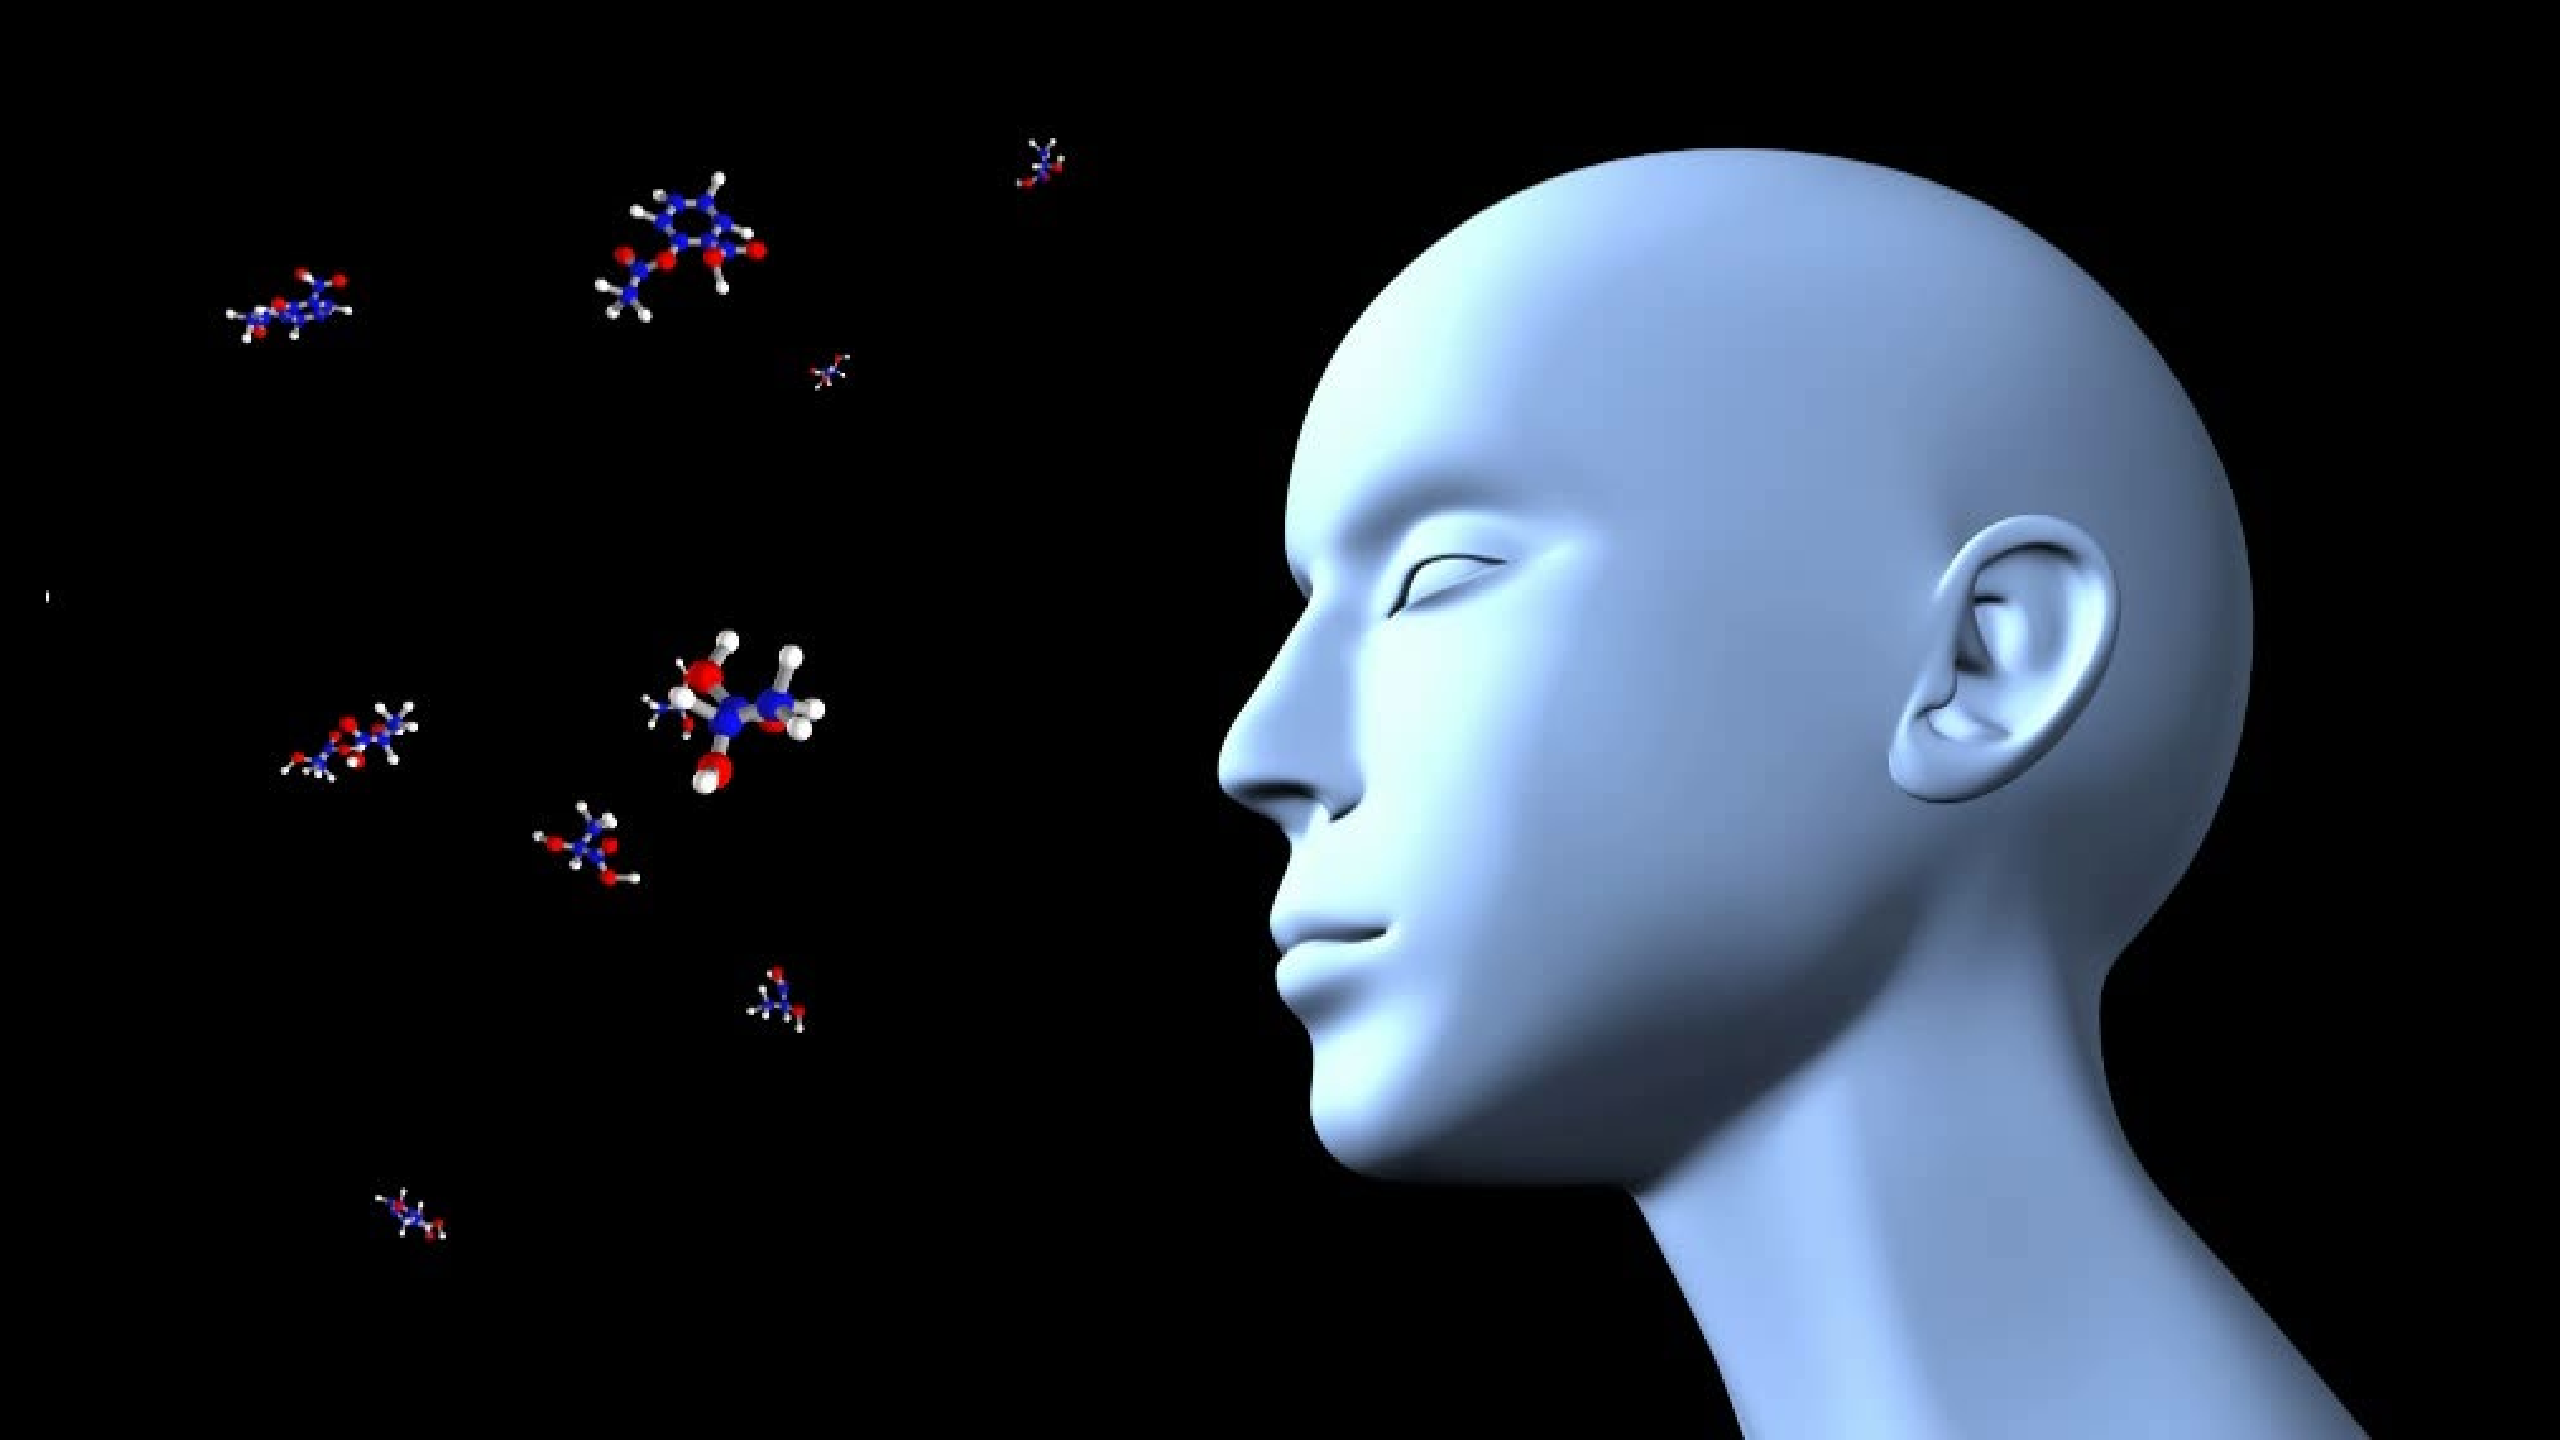

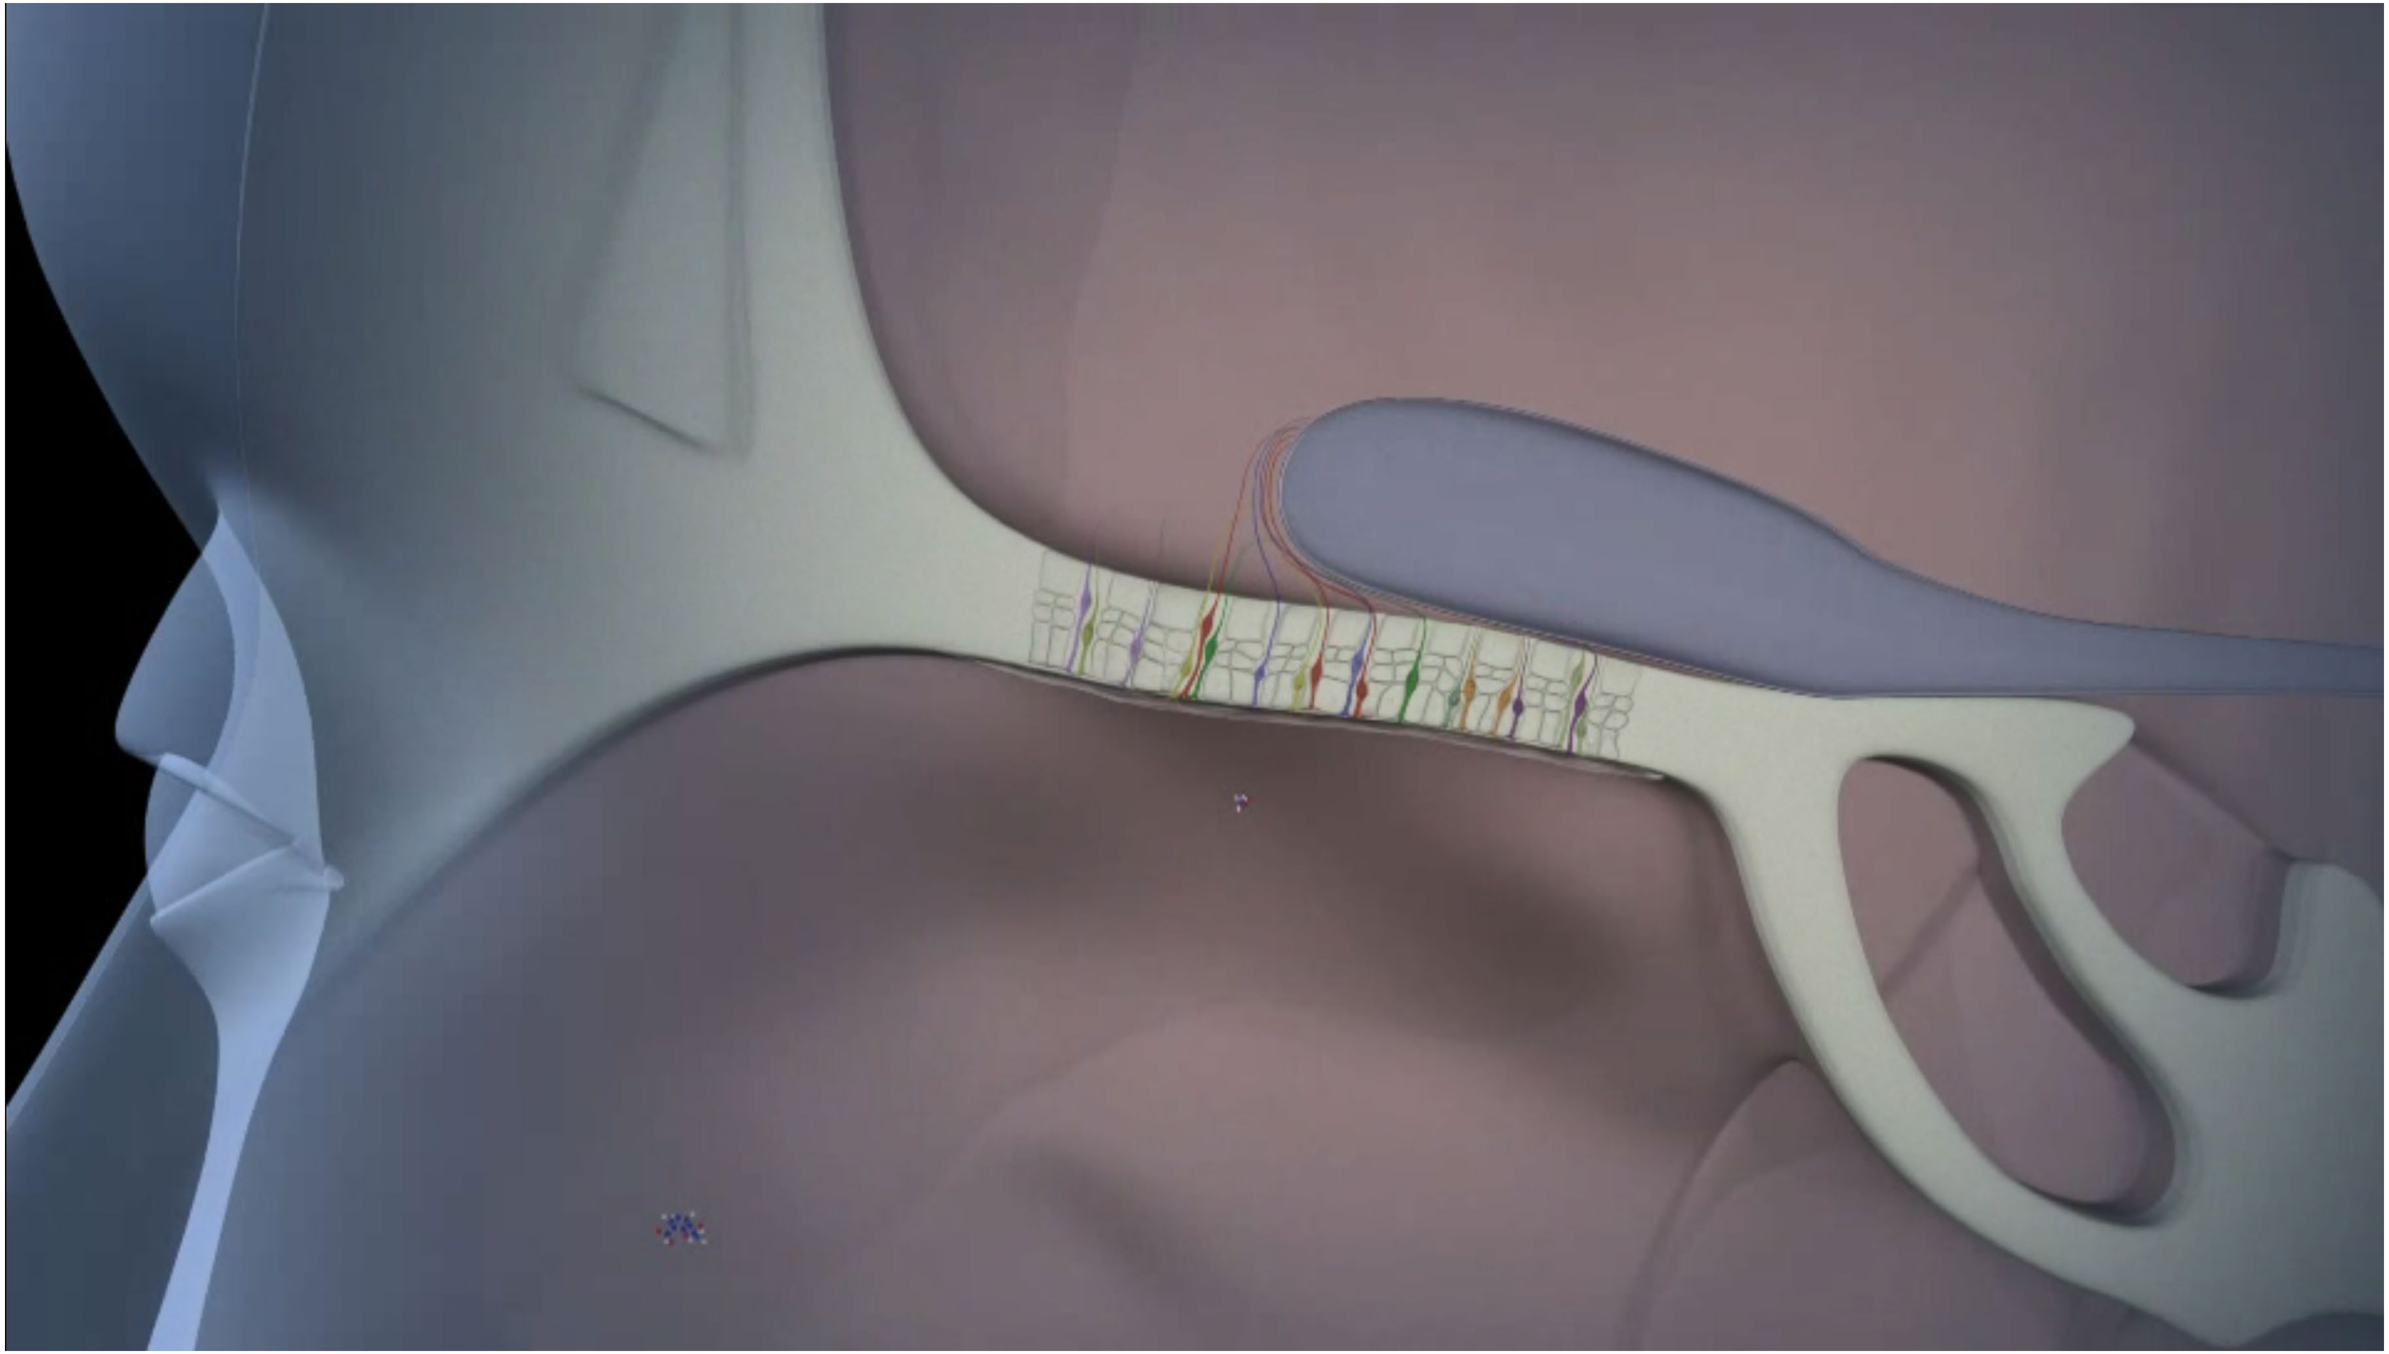

video by dancing-lemon-studio.com

Biological considerations for ML models of olfactory system

- Nose is chiral (mirror images smell differently)
- The relevant stimulus is ambiguous
  - Many possible conformations
  - Many possible orientations

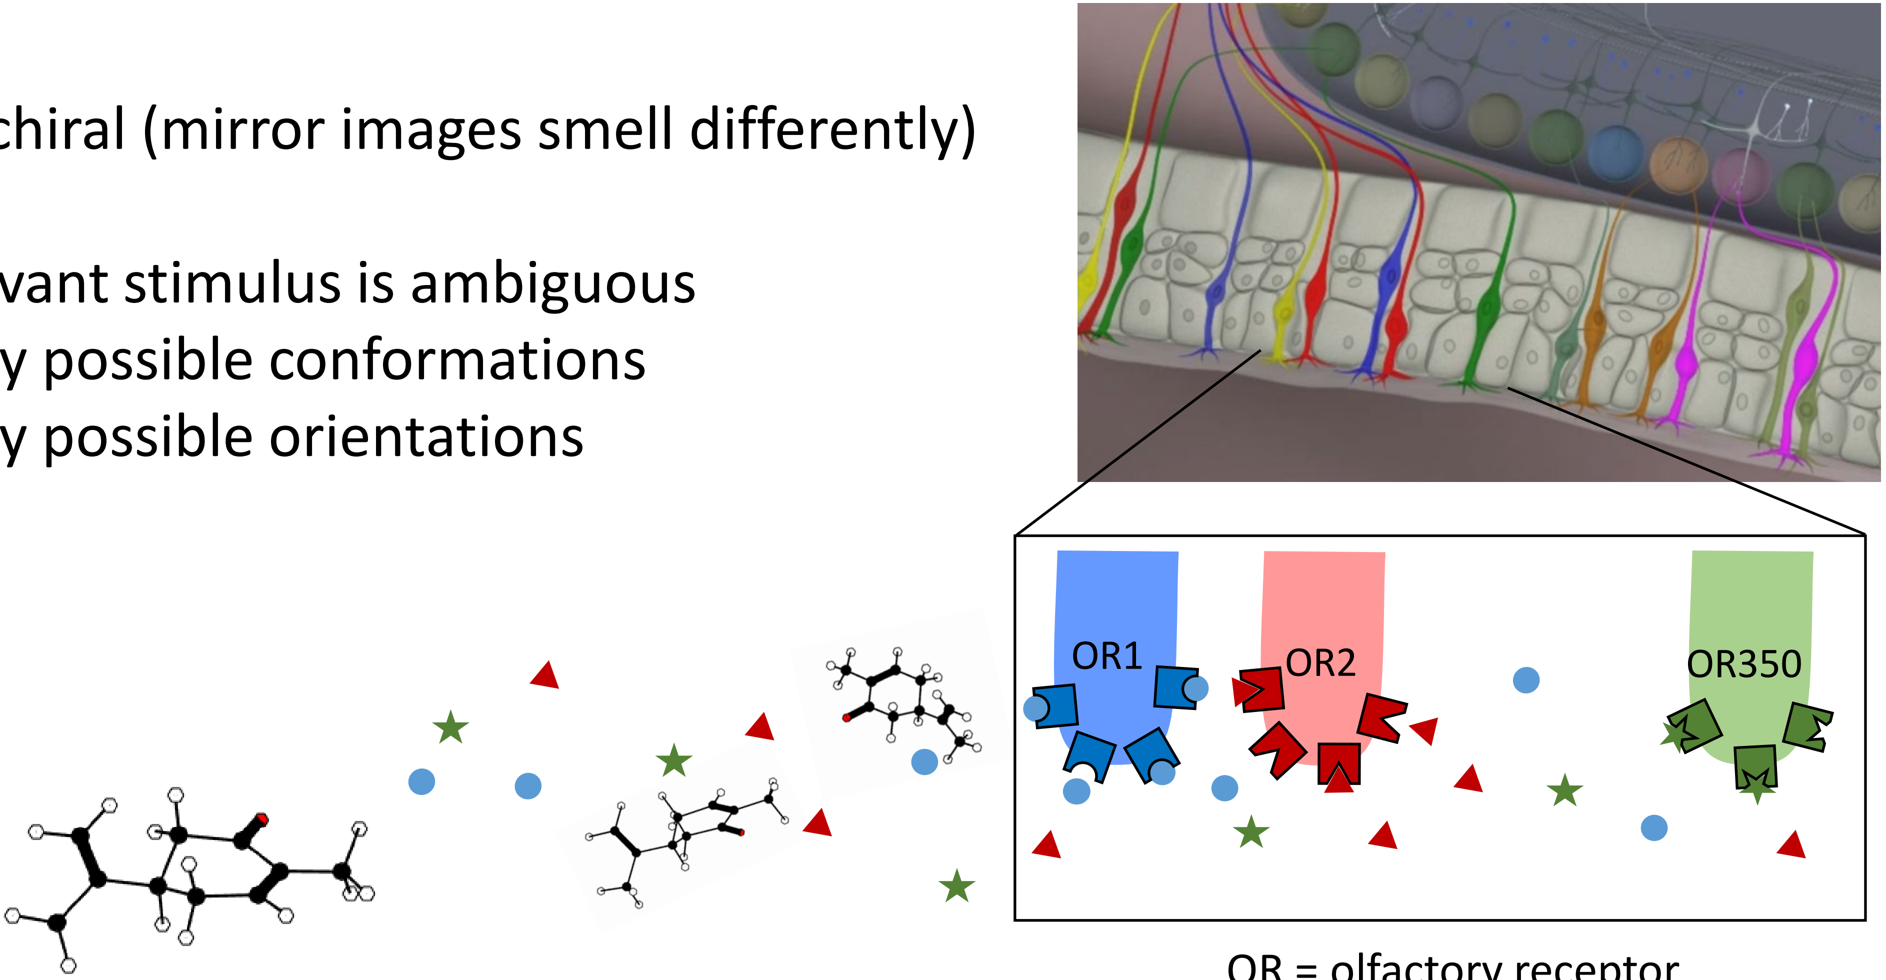

OR = olfactory receptor

Molecules can be represented as 5D tensors

```
    3D space
+   1D element
    1D orientation
----------------------
=   5D tensor
```

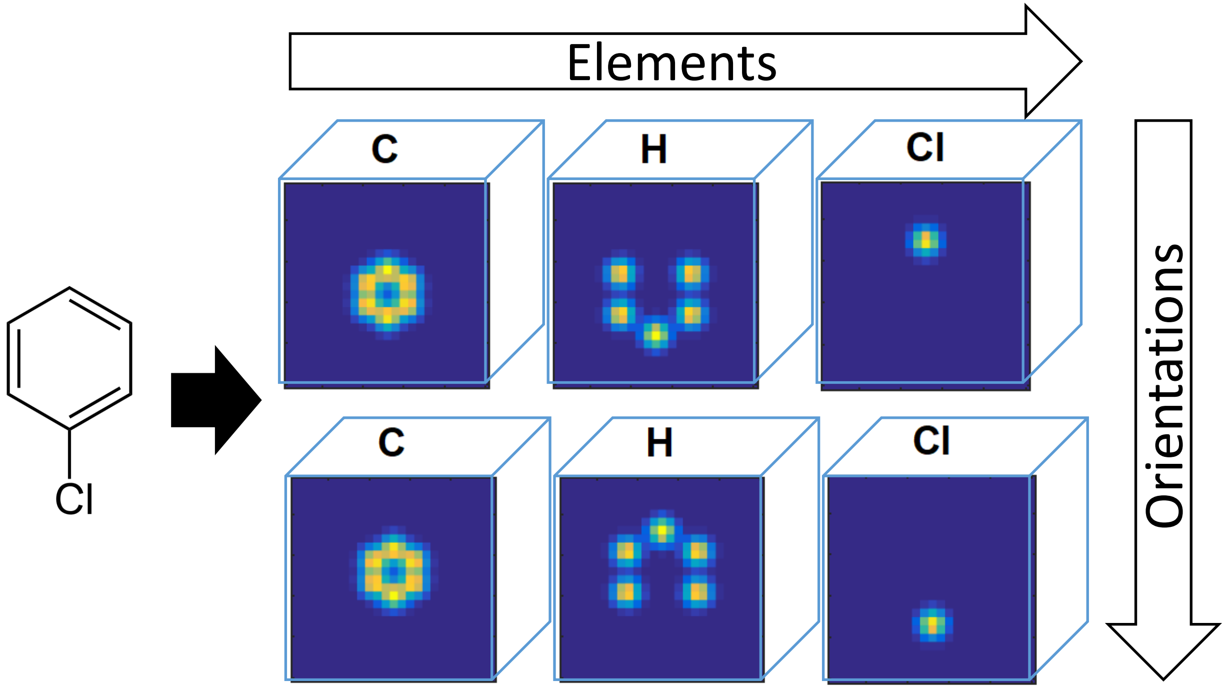

What factors influence the composition of OR ensembles?

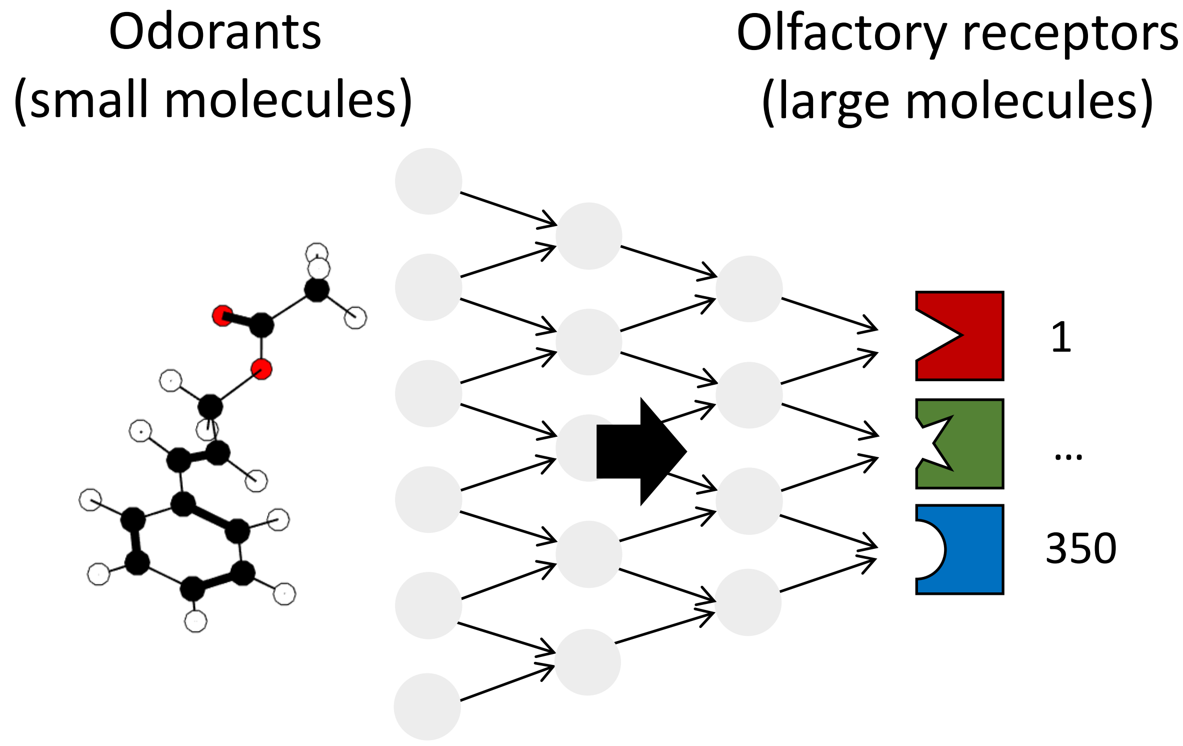

Approach: using CNN to model olfactory receptors.

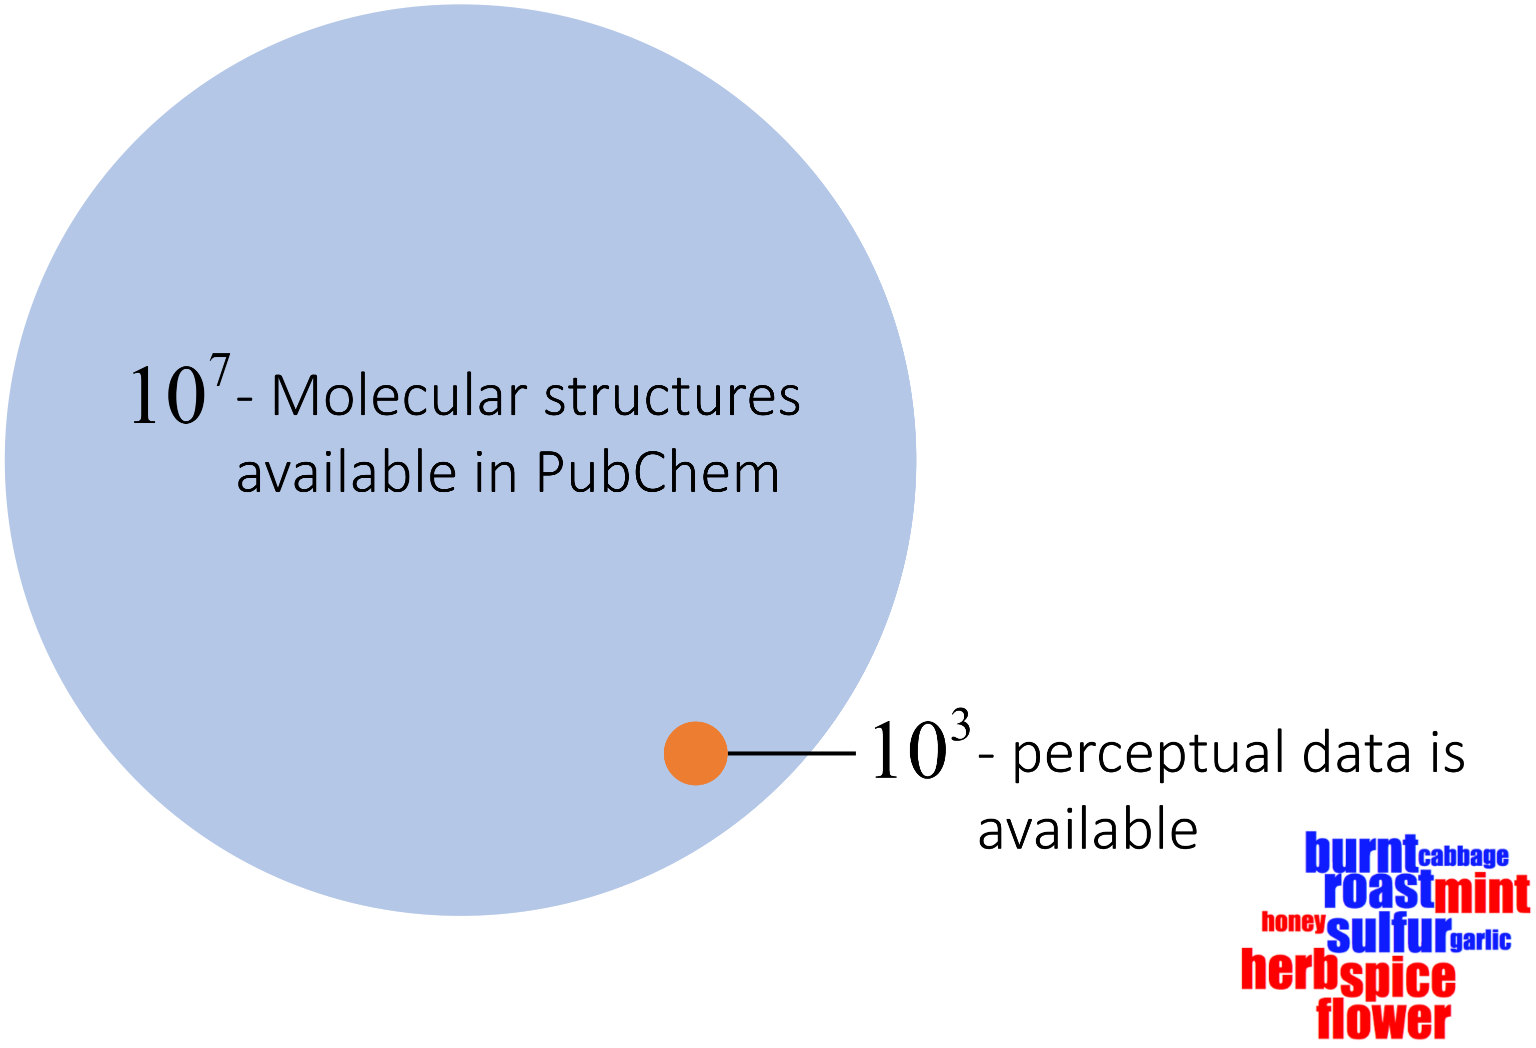
10^7- Molecular structures available in PubChem
10^3- perceptual data is available
burnt
cabbage
roast
mint
honey
sulfur
garlic
herb
spice
flower

DeepNose autoencoder extracts molecular features from molecules

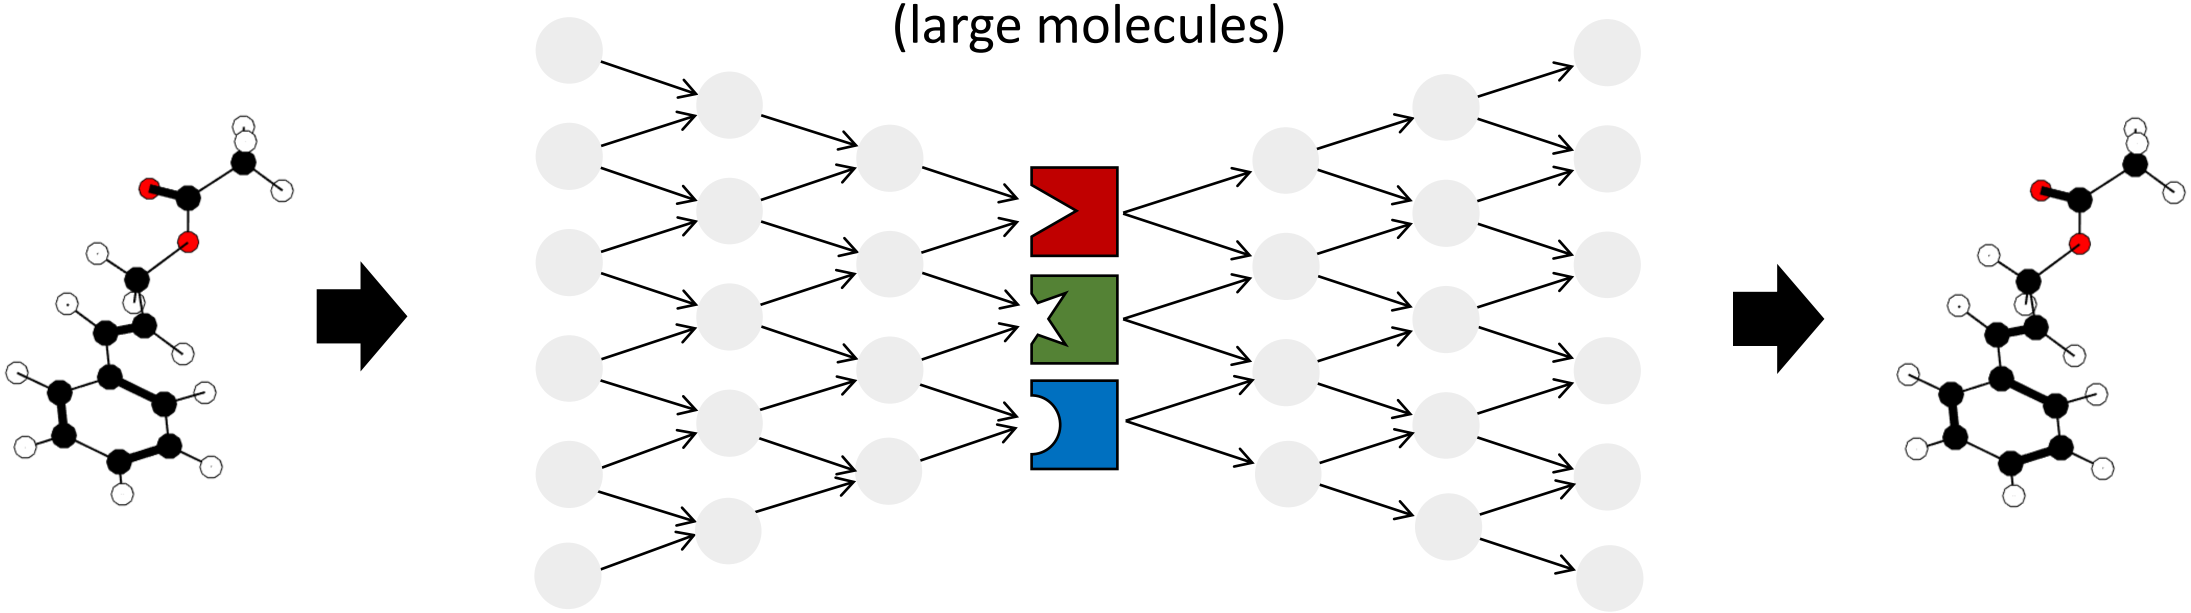

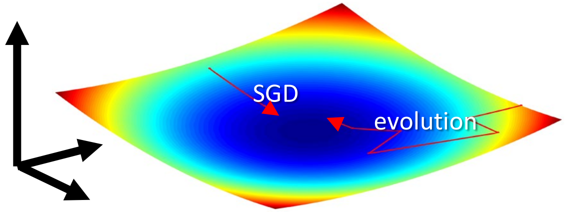

Hinton, G. E. & Salakhutdinov, R. R. *Science* 2006

DeepNose autoencoder reconstructed molecules accurately

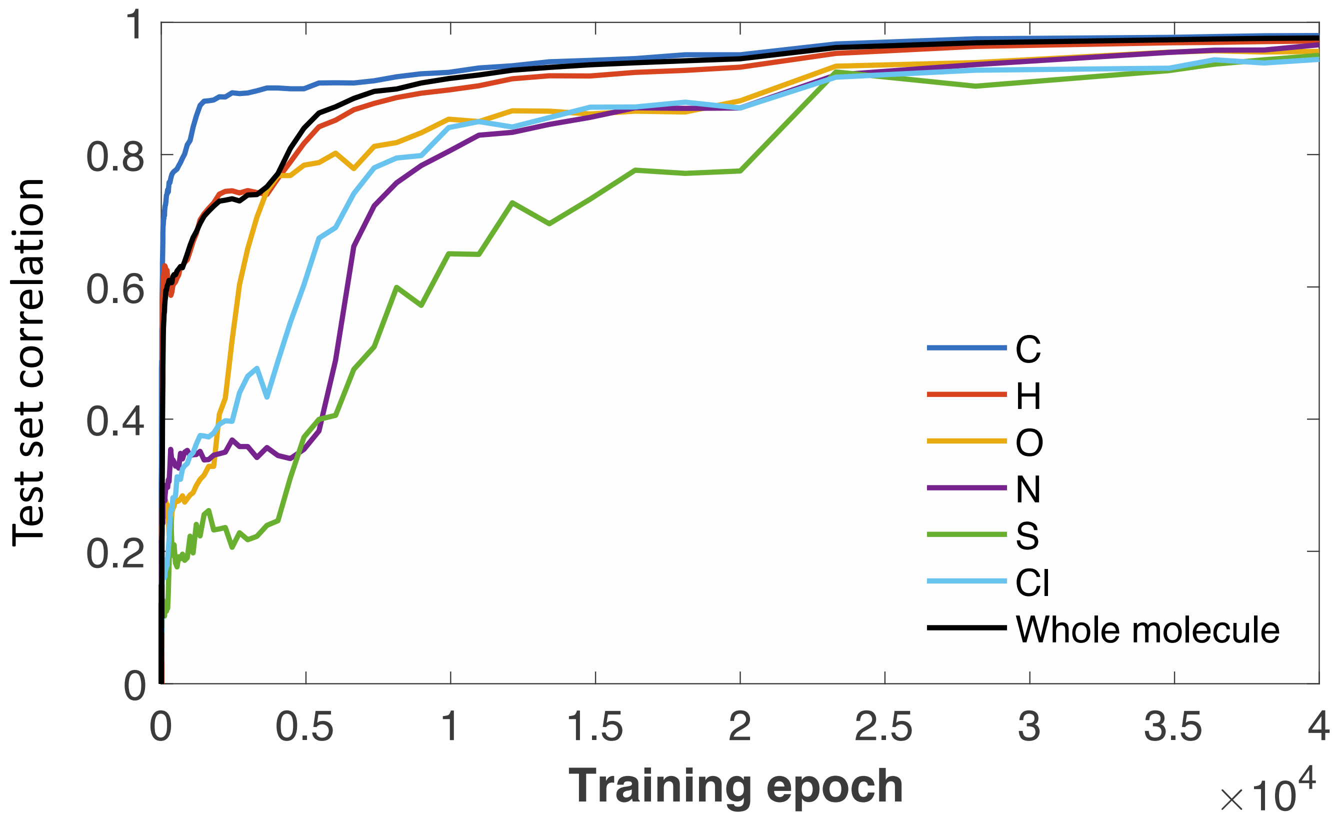

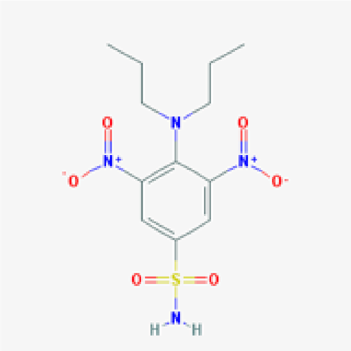

CID 29393

(test set)

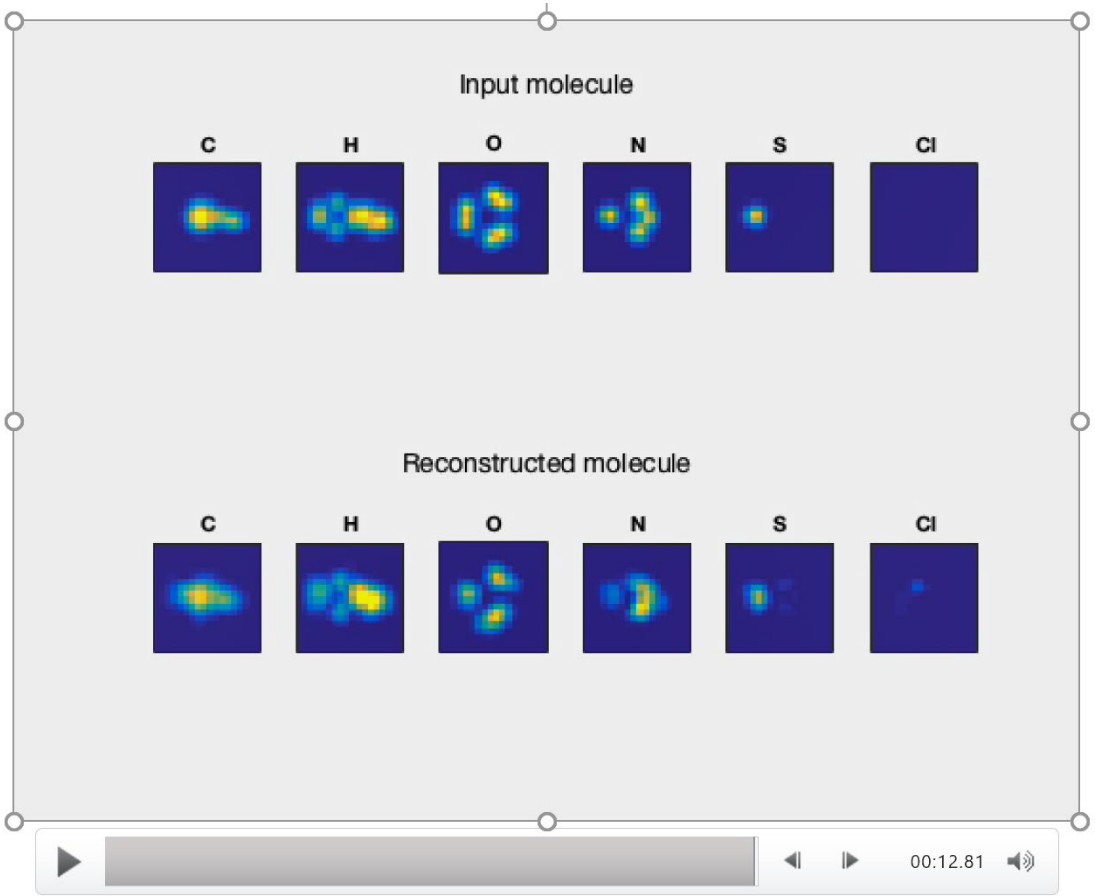

Transferring the encoder’s weights to the classifier

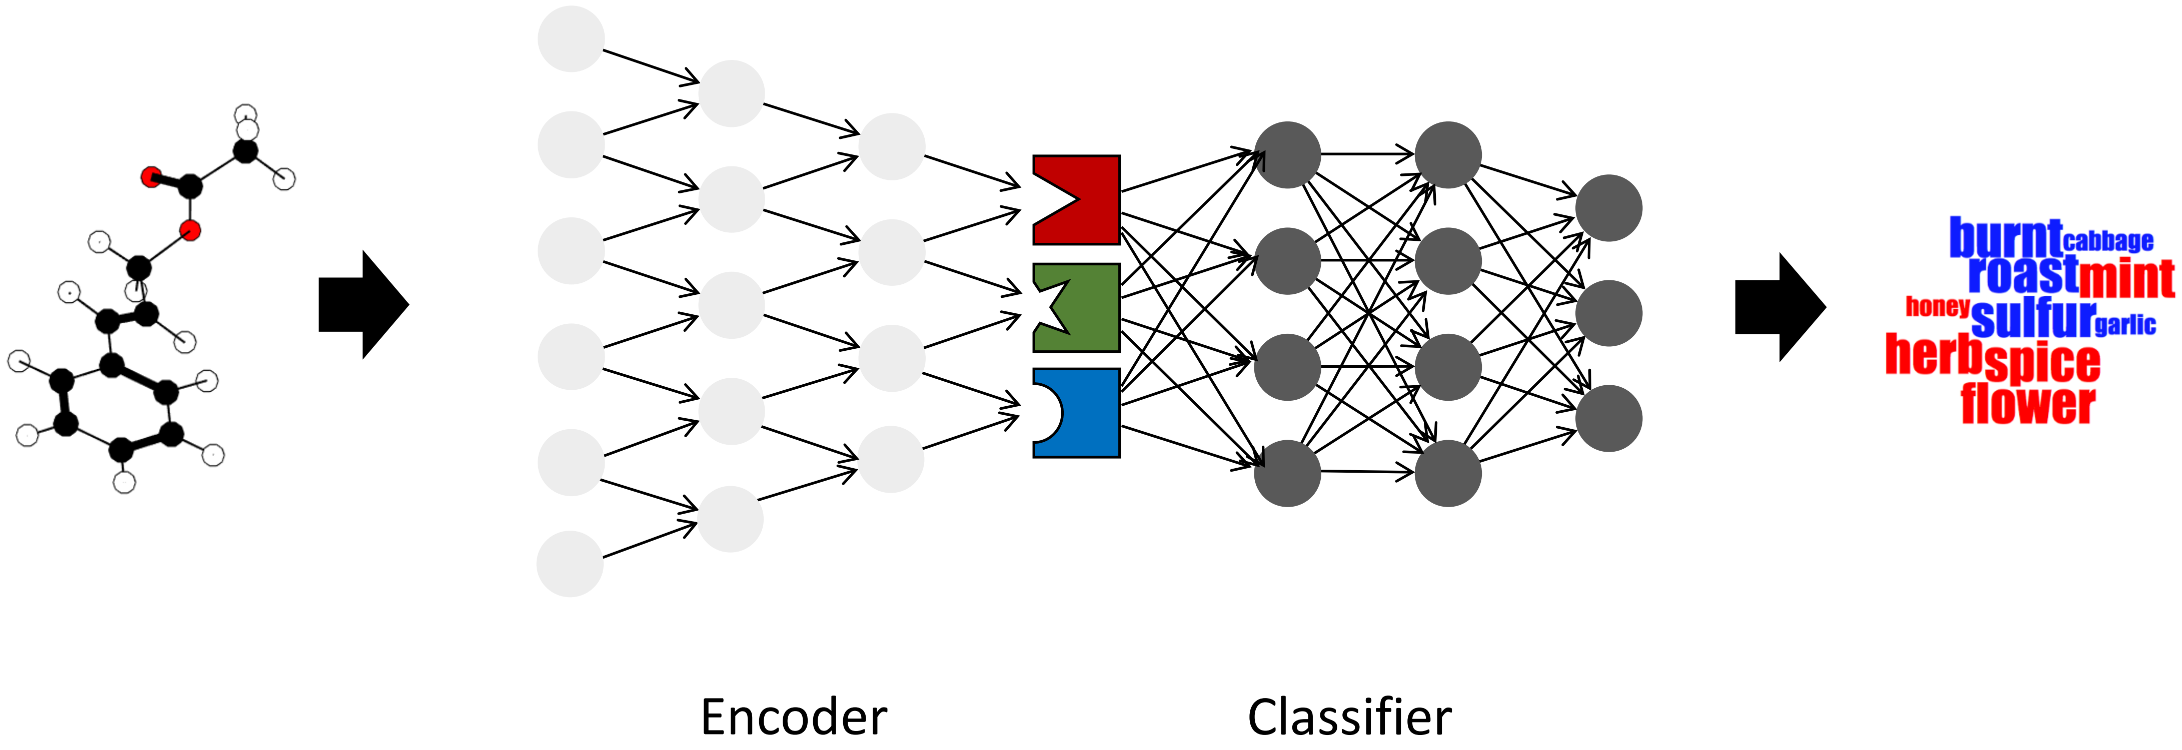

Good Scents dataset contains perceptual descriptors for odorant molecules

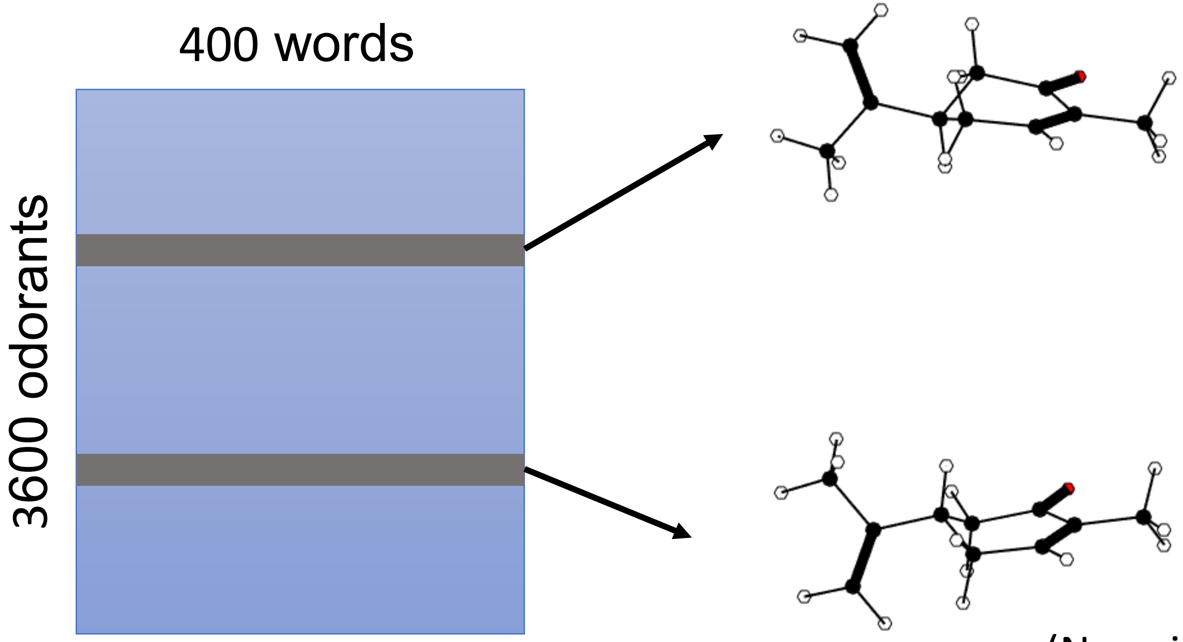

herbal minty spearmint sweet

carraway

(Nose is chiral)

Good Scents dataset

http://thegoodscentscompany.com/

DeepNose can predict olfactory descriptors based on 3D shapes alone

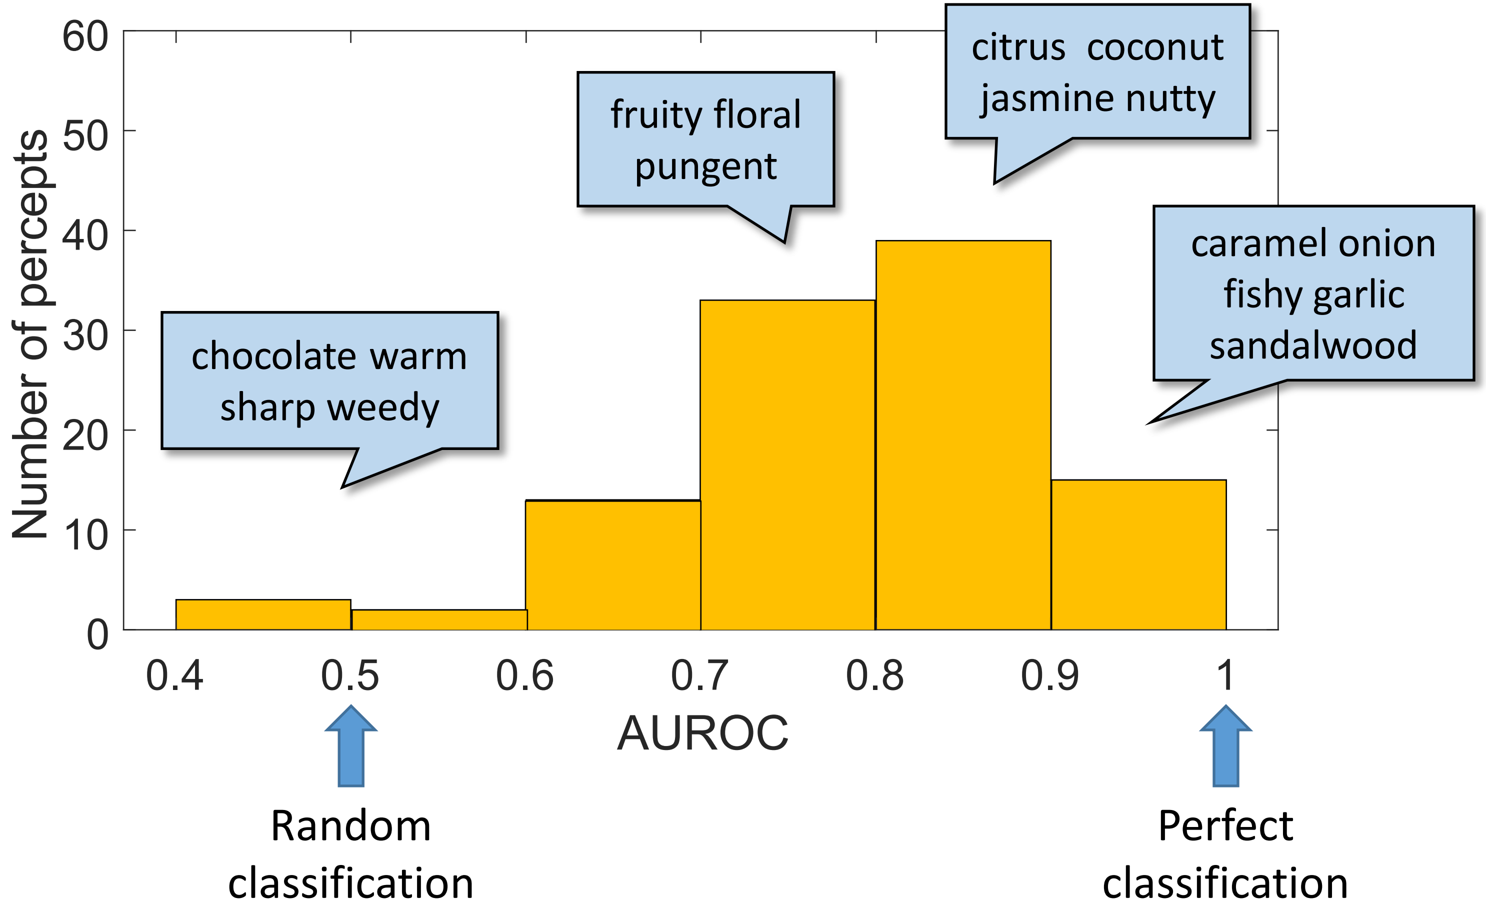

Interpreting DeepNose predictions

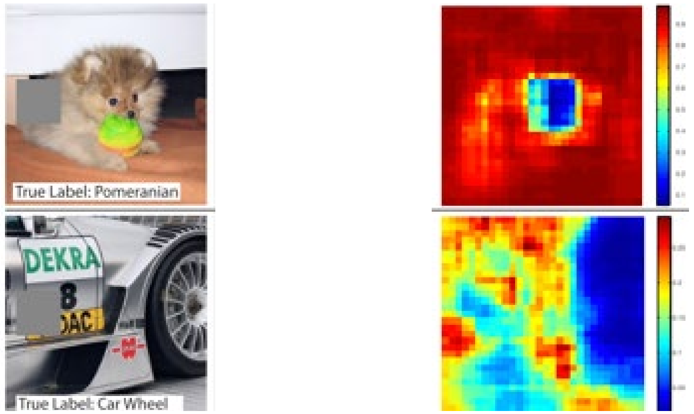

Estimating the "receptive field" of a trained neural networks:

- Occlude a region of the input
- Observe its impact on prediction

Zeiler & Fergus, ECCV 2014

“Fishy” odors – amine groups

trimethylamine

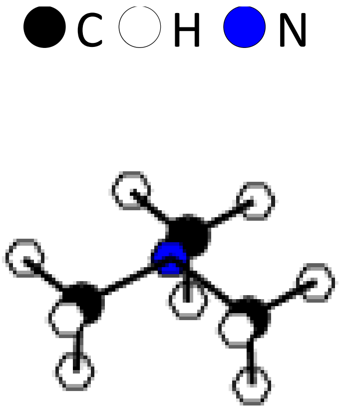

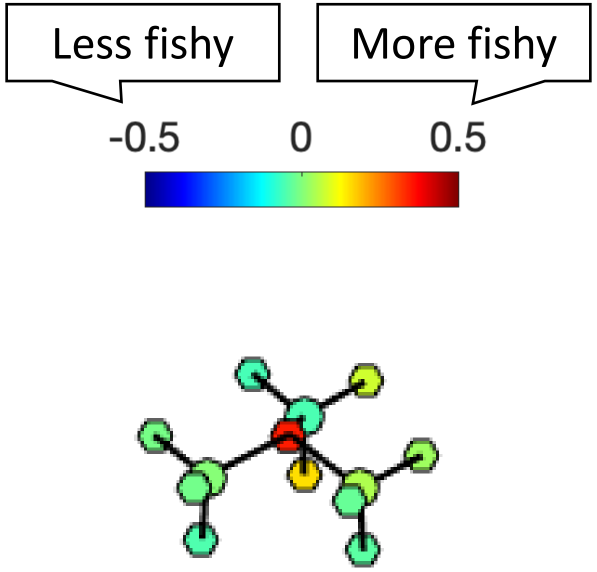

label: fishy

“Fruity” odors – ester groups

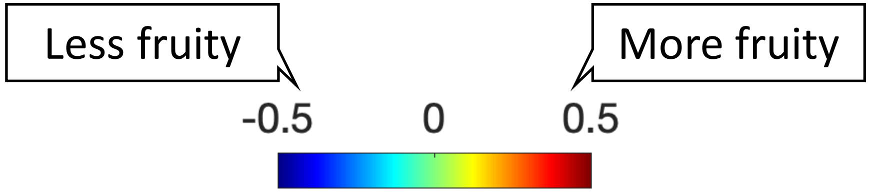

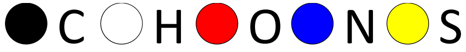

#7758

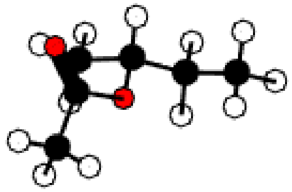

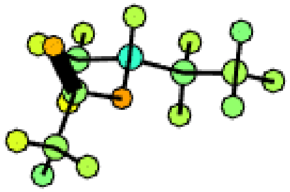

banana <u>fruity</u>

#243783

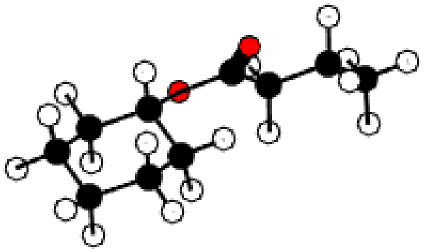

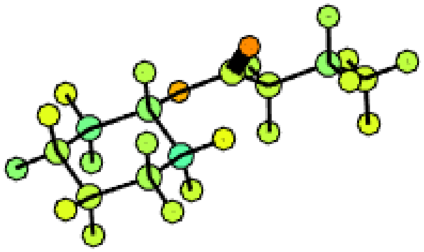

apple floral
<u>fruity</u> pineapple

“Cinnamon” odors – benzene rings

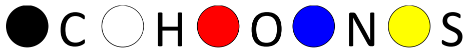

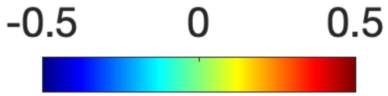

#61010

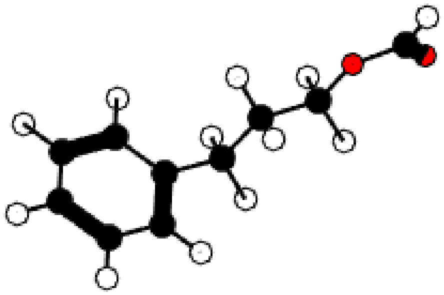

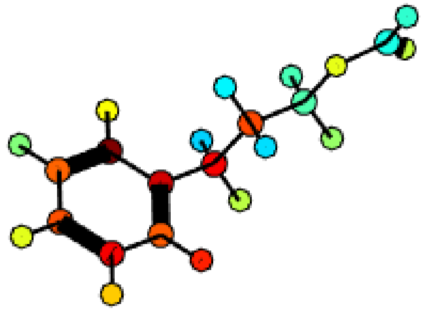

<u>cinnamon</u> floral honey hyacinth sweet

6445555

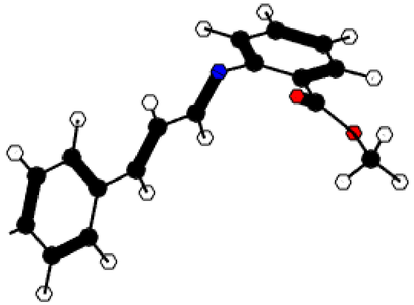

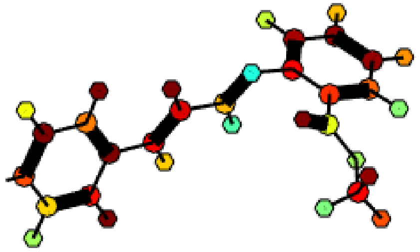

<u>cinnamon</u> sweet

Conclusions

- Olfactory receptors may be viewed as 3D nonlinear convolutional filters that analyzed molecular shapes.

- Deep networks can be used to model the constraint on the evolution of olfactory receptors.

We hire!

Interested in in building new machine learning algorithms inspired by the brain?

Contact me at koulakov@cshl.edu

thanks to:

Lab members

- Hamza Giaffar
- Sergey Shuvaev
- Tumi Tran
- Batu Baserdem
- Daniel Kepple

Collaborators:

olfaction

- Florin Albeanu (CSHL)
- Steve Shea (CSHL)
- Dima Rinberg (NYU)
- Thomas Bozza (Northwestern)
- Alex Fleischman (Brown)
- Bob Datta (Harvard)
- Kevin Franks (Duke)
- Rick Gerkin (ASU)
- Joel Mainland (Monell)
- Tim Holli (WUStL)

RL+theory

- Adam Kepecs (CSHL)
- Sarah Starosta (CSHL)
- Bo Li (CSHL)
- Tony Zador (CSHL)
- Grisha Enikolopov (SUNYSB)
- Pavel Osten (CSHL)

Funding:

This work is supported by Deep Valley Labs, San Jose, CA

NIH DIRECTOR'S TRANSFORMATIVE RESEARCH AWARD

NIDCD National Institute on Deafness and Other Communication Disorders

NIMH National Institute of Mental Health

THE SWARTZ FOUNDATION

NIH National Institute of Neurological Disorders and Stroke

NIH National Institute on Drug Abuse
Advancing Addiction Science

WE HIRE!