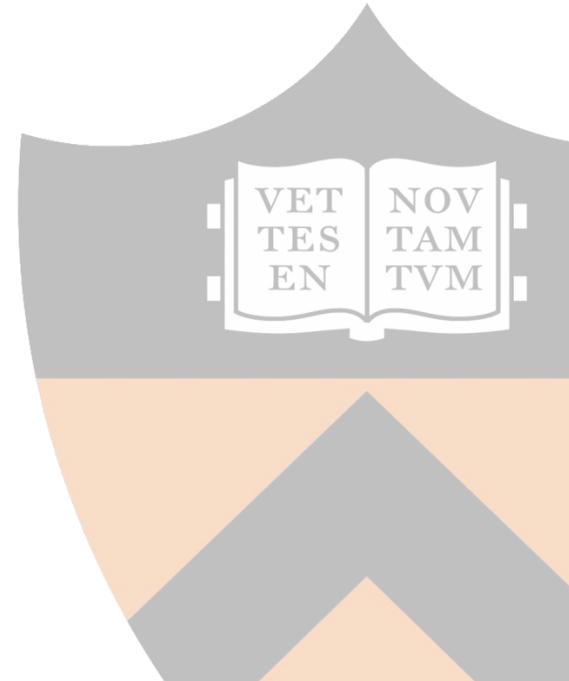


Direct fit

an ecological perspective on biological
and artificial neural networks

Samuel A. Nastase
Princeton Neuroscience Institute
snastase@princeton.edu

 @samnastase



Direct fit

an ecological perspective on biological
and artificial neural networks

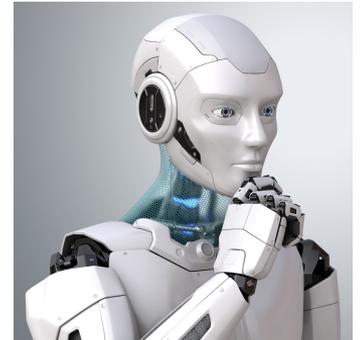
Samuel A. Nastase
Princeton Neuroscience Institute
snastase@princeton.edu

 @samnastase

Uri Hasson



Ariel Goldstein



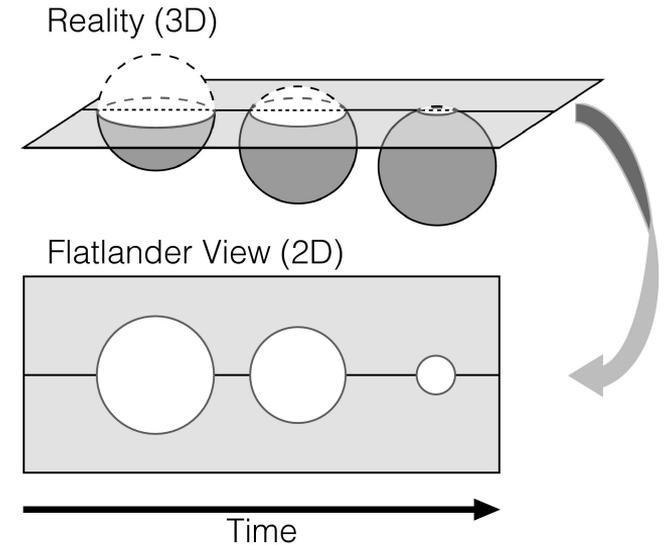
Toward a more ecological neuroscience

Confronted with the complexity of brain and behavior, neuroscientists traditionally design controlled experiments to reduce the dimensionality of the problem to a few intuitive factors.

Meehl, *Psychol Rep*, 1990

Rozenblit & Keil, *Cogn Sci*, 2002

Jolly & Chang, *Topics Cogn Sci*, 2019



Toward a more ecological neuroscience

Models developed under tightly-controlled experimental manipulations are typically tested in more naturalistic, ecological contexts...

...but results have been mixed

Reinagel, *Curr Opin Neurobiol*, 2001
Simoncelli et al., *Annu Rev Neurosci*, 2001
Bartels & Zeki, *Hum Brain Mapp*, 2004
David et al., *Nature*, 2004
Hasson et al., *Science*, 2004
Kayser et al., *Curr Opin Neurobiol*, 2004
Felsen & Dan, *Nat Neurosci*, 2005
Spiers & Maguire, *Trends Cogn Sci*, 2007
Haxby et al., *Neuron*, 2011
Hasson & Honey, *NeuroImage*, 2012
Maguire, *NeuroImage*, 2012
Hamilton & Huth, *Lang Cogn Neurosci*, 2018
Haxby, Nastase et al., *NeuroImage*, 2020

Toward a more ecological neuroscience

Models developed under tightly-controlled experimental manipulations are typically tested in more naturalistic, ecological contexts...

...but results have been mixed

“We can rightfully claim to understand only 10% to 20% of how V1 actually operates under normal conditions.”

Olshausen & Field, *Neural Comput*, 2005

- biased sampling in neural measurement
- biased sampling of stimuli and tasks
- bias toward simple, interpretable theories
- interdependence of neural variables and context
- publication bias

Toward a more ecological neuroscience

Models developed under tightly-controlled experimental manipulations are typically tested in more naturalistic, ecological contexts...

...but results have been mixed

“We can rightfully claim to understand only 10% to 20% of how V1 actually operates under normal conditions.”

Olshausen & Field, *Neural Comput*, 2005

“In order to behave like scientists, [experimental psychologists] must construct situations in which our subjects are totally controlled, manipulated and measured. We must cut our subjects down to size. We construct situations in which they can behave as little like human beings as possible and we do this in order to allow ourselves to make statements about the nature of their humanity.”

Bannister, *Bull Br Psychol Soc*, 1966

A brief history of “ecological validity”

“Ecological validity”—coined by Egon Brunswik in 1947



A brief history of “ecological validity”

“Ecological validity”—coined by Egon Brunswik in 1947

...to mean something else (\neg _ \neg)



A brief history of “ecological validity”

“Ecological validity” — coined by Egon Brunswik in 1947

...to mean something else ($\neg _ \neg$)

Representative design

- psychology maintains a “double standard” in the application of sampling theory
- subjects are sampled with the goal of generalizing to the population, whereas stimuli and tasks generally are not
- ecological generalizability demands a “representative sampling of situations” where “situational instances in an ecology are analogous to individuals in a population”

Brunswik, *Psychol Rev*, 1943

Brunswik, *Psychol Rev*, 1955



A brief history of “ecological validity”

“Ecological validity” — coined by Egon Brunswik in 1947

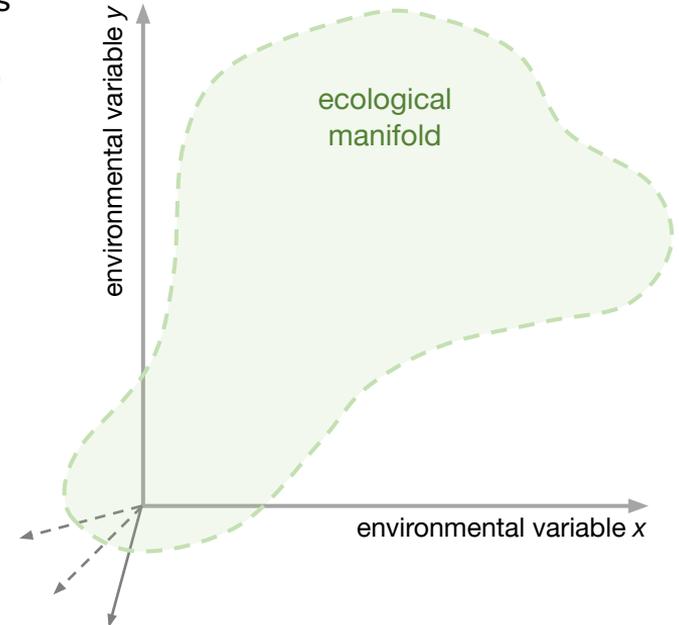
...to mean something else ($\neg _ \neg$)

Representative design

- psychology maintains a “double standard” in the application of sampling theory
- subjects are sampled with the goal of generalizing to the population, whereas stimuli and tasks generally are not
- ecological generalizability demands a “representative sampling of situations” where “situational instances in an ecology are analogous to individuals in a population”

Brunswik, *Psychol Rev*, 1943

Brunswik, *Psychol Rev*, 1955



A brief history of “ecological validity”

“Ecological validity” —coined by Egon Brunswik in 1947

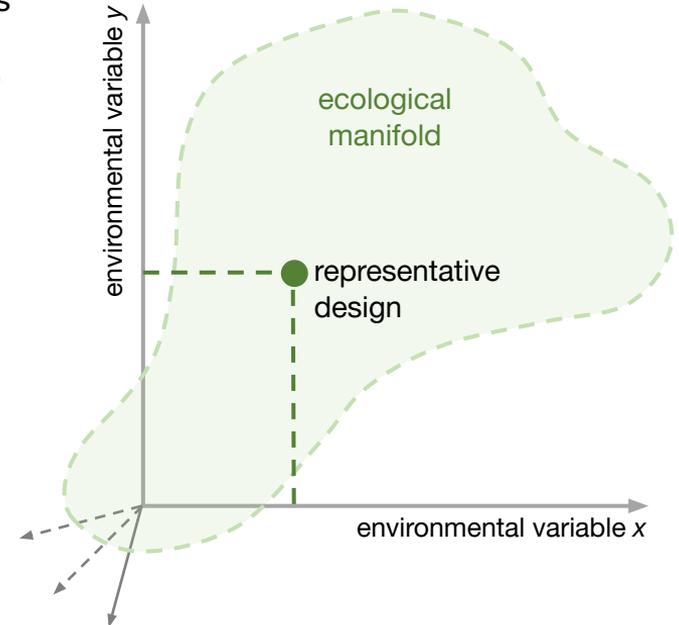
...to mean something else ($\neg _ \neg$)

Representative design

- psychology maintains a “double standard” in the application of sampling theory
- subjects are sampled with the goal of generalizing to the population, whereas stimuli and tasks generally are not
- ecological generalizability demands a “representative sampling of situations” where “situational instances in an ecology are analogous to individuals in a population”

Brunswik, *Psychol Rev*, 1943

Brunswik, *Psychol Rev*, 1955



A brief history of “ecological validity”

“Ecological validity” — coined by Egon Brunswik in 1947

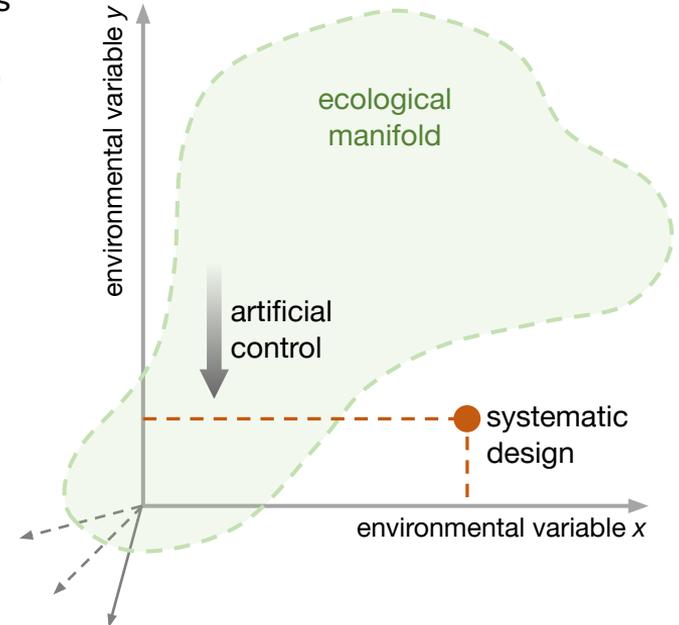
...to mean something else ($\neg _ \neg$)

Representative design

- psychology maintains a “double standard” in the application of sampling theory
- subjects are sampled with the goal of generalizing to the population, whereas stimuli and tasks generally are not
- ecological generalizability demands a “representative sampling of situations” where “situational instances in an ecology are analogous to individuals in a population”

Brunswik, *Psychol Rev*, 1943

Brunswik, *Psychol Rev*, 1955



A brief history of “ecological validity”

“Ecological validity” —coined by Egon Brunswik in 1947

...to mean something else ($\neg _ \neg$)

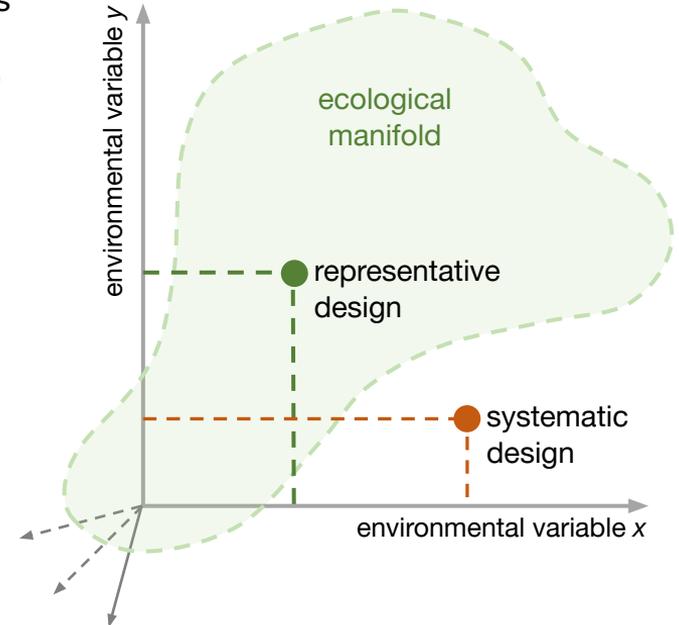
Representative design

- psychology maintains a “double standard” in the application of sampling theory
- subjects are sampled with the goal of generalizing to the population, whereas stimuli and tasks generally are not
- ecological generalizability demands a “representative sampling of situations” where “situational instances in an ecology are analogous to individuals in a population”

Brunswik, *Psychol Rev*, 1943

Brunswik, *Psychol Rev*, 1955

Nastase, S. A., Goldstein, A., & Hasson, U. (2020). Keep it real: rethinking the primacy of experimental control in cognitive neuroscience. *PsyArXiv*.



The “direct fit” model of biological learning

Relinquishing control—machine learning

Modern artificial neural networks leverage millions of noisy, real-world samples to achieve shockingly good performance at a variety of “cognitive” tasks...

—how is this possible?

—**and why do neuroscientists find this so unsatisfying?**

The “direct fit” model of biological learning

Relinquishing control—machine learning

Modern artificial neural networks leverage millions of noisy, real-world samples to achieve shockingly good performance at a variety of “cognitive” tasks...

—how is this possible?

—**and why do neuroscientists find this so unsatisfying?**

Statistical Science
2001, Vol. 16, No. 3, 199–231

Statistical Modeling: The Two Cultures

Leo Breiman

Choosing Prediction Over Explanation in Psychology: Lessons From Machine Learning

Tal Yarkoni and Jacob Westfall

University of Texas at Austin

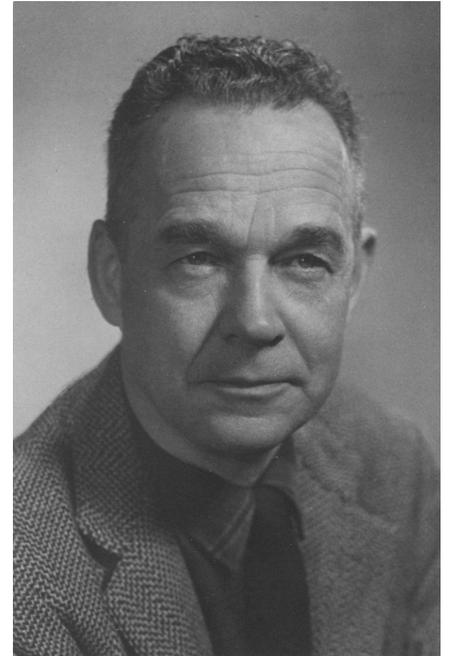
Perspectives on Psychological Science
2017, Vol. 12(6) 1100–1122
© The Author(s) 2017
Reprints and permissions:
sagepub.com/journalsPermissions.nav
DOI: 10.1177/1745691617693393
www.psychologicalscience.org/PPS



The “direct fit” model of biological learning

The *meaningful* environment

“The *affordances* of the environment are what it *offers* the animal, what it *provides* or *furnishes*, either for good or ill. The verb to *afford* is found in the dictionary, but the noun *affordance* is not. I have made it up.” —James J. Gibson (1979)



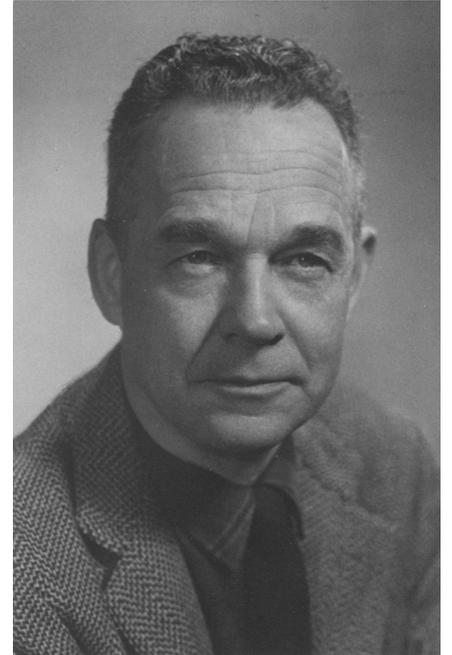
The “direct fit” model of biological learning

The *meaningful* environment

“The *affordances* of the environment are what it *offers* the animal, what it *provides* or *furnishes*, either for good or ill. The verb to *afford* is found in the dictionary, but the noun *affordance* is not. I have made it up.” —James J. Gibson (1979)

Ecological psychology and “direct perception”

- organisms “directly” perceive opportunities for **adaptive behavior** encoded in the environment—or “affordances”
- organism and environment are inextricably coupled
- perception and action are inextricably coupled



The “direct fit” model of biological learning

“In perception, perhaps the nearest anyone came to the level of computational theory was Gibson (1966). However, although some aspects of his thinking were on the right lines, he did not understand properly what information processing was, which led him to seriously underestimate the complexity of the information-processing problems involved in vision and the consequent subtlety that is necessary in approaching them.” —David Marr (1982)



The “direct fit” model of biological learning

“In perception, perhaps the nearest anyone came to the level of computational theory was Gibson (1966). However, although some aspects of his thinking were on the right lines, he did not understand properly what information processing was, which led him to seriously underestimate the complexity of the information-processing problems involved in vision and the consequent subtlety that is necessary in approaching them.” —David Marr (1982)

Direct-fit learning

Neural networks rely on “mindless” overparameterized optimization to map millions of noisy, real-world samples along task-relevant objective functions—providing a mechanistic basis for Gibson’s “direct perception.”

Hasson, U., Nastase, S. A., & Goldstein, A. (2020). Direct fit to nature: an evolutionary perspective on biological and artificial neural networks. *Neuron*, 105(3), 416–434.



What makes a good model?

Scientists favor concise, interpretable models with an emphasis on explanatory power over predictive power.

Sentence

NP + VP

(i)

T + N + VP

(ii)

T + N + Verb + NP

(iii)

the + N + Verb + NP

(iv)

the + man + Verb + NP

(v)

the + man + hit + NP

(vi)

the + man + hit + T + N

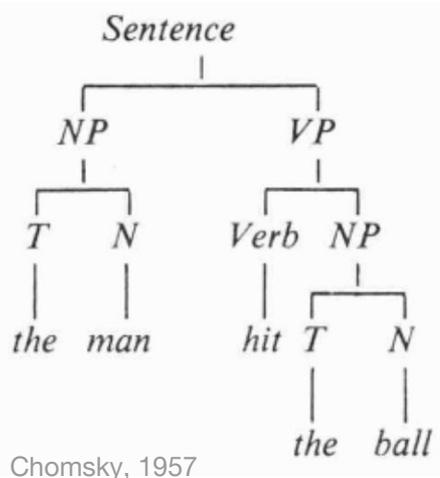
(ii)

the + man + hit + the + N

(iv)

the + man + hit + the + ball

(v)



Chomsky, 1957

What makes a good model?

Scientists favor concise, interpretable models with an emphasis on explanatory power over predictive power.

Sentence

NP + VP

T + N + VP

T + N + Verb + NP

the + N + Verb + NP

the + man + Verb + NP

the + man + hit + NP

the + man + hit + T + N

the + man + hit + the + N

the + man + hit + the + ball

(i)

(ii)

(iii)

(iv)

(v)

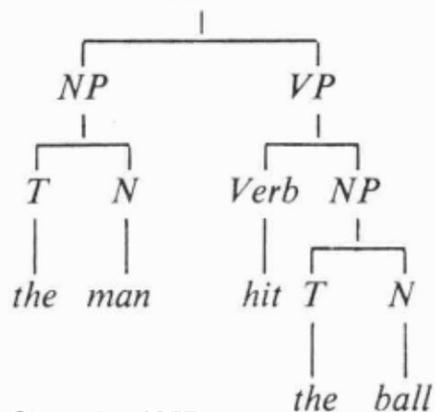
(vi)

(ii)

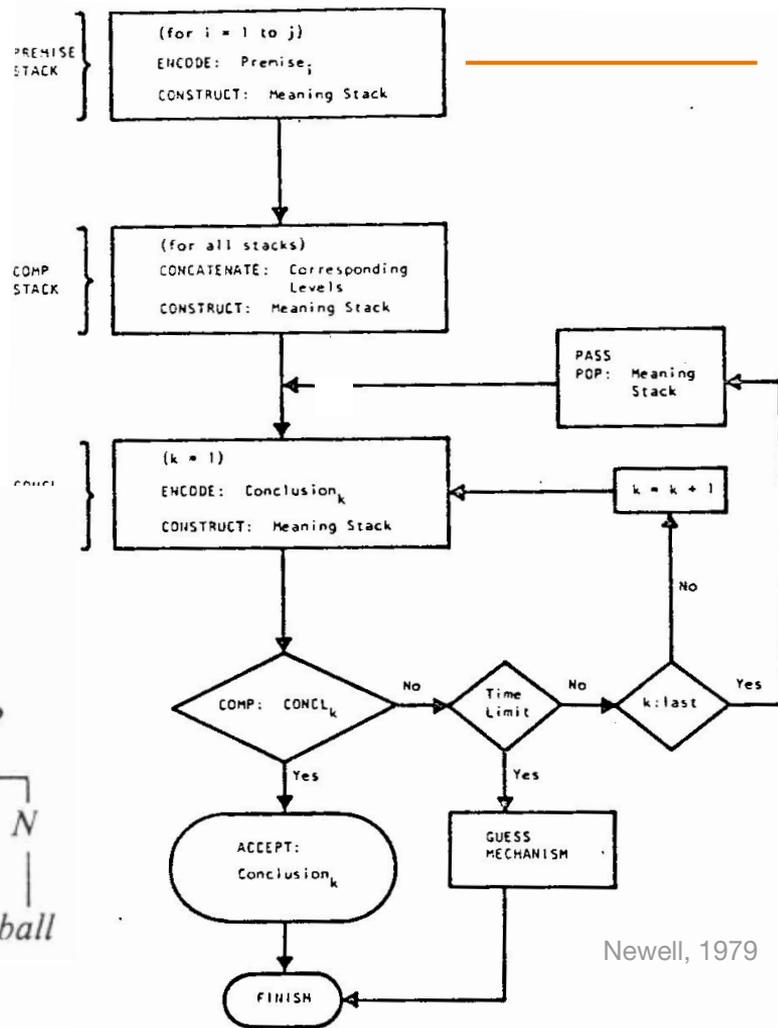
(iv)

(v)

Sentence



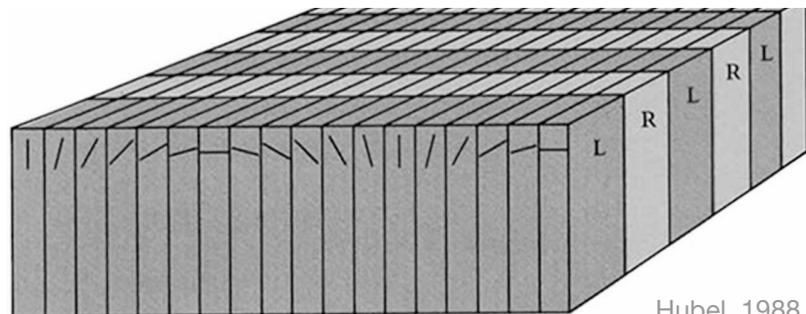
Chomsky, 1957



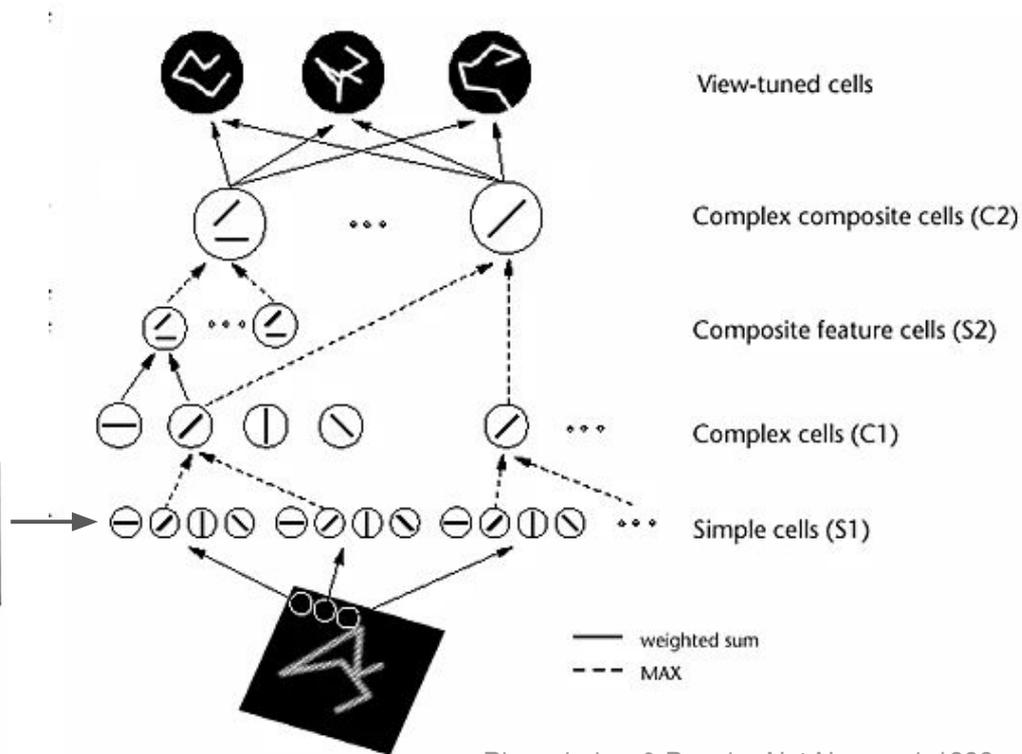
Newell, 1979

What makes a good model?

Scientists favor concise, interpretable models with an emphasis on explanatory power over predictive power.



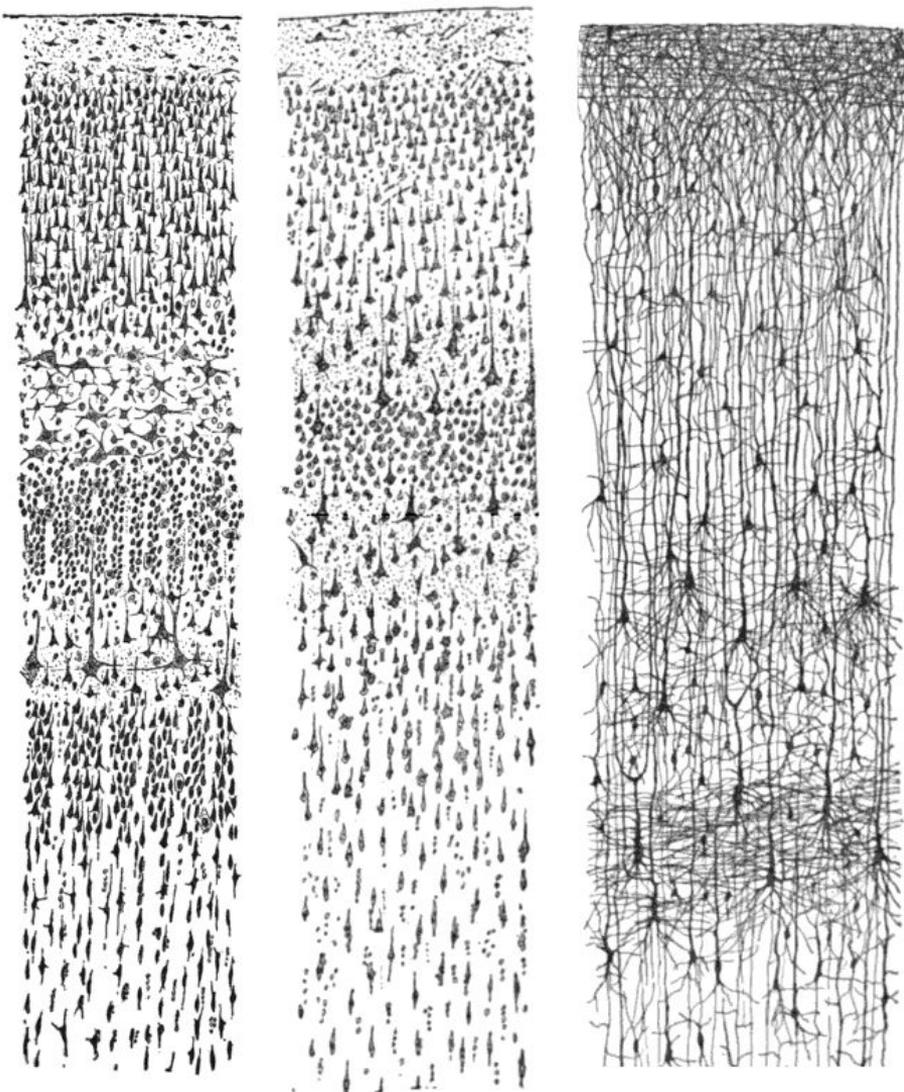
Hubel, 1988



Riesenhuber & Poggio, *Nat Neurosci*, 1999

What makes a good model? _____

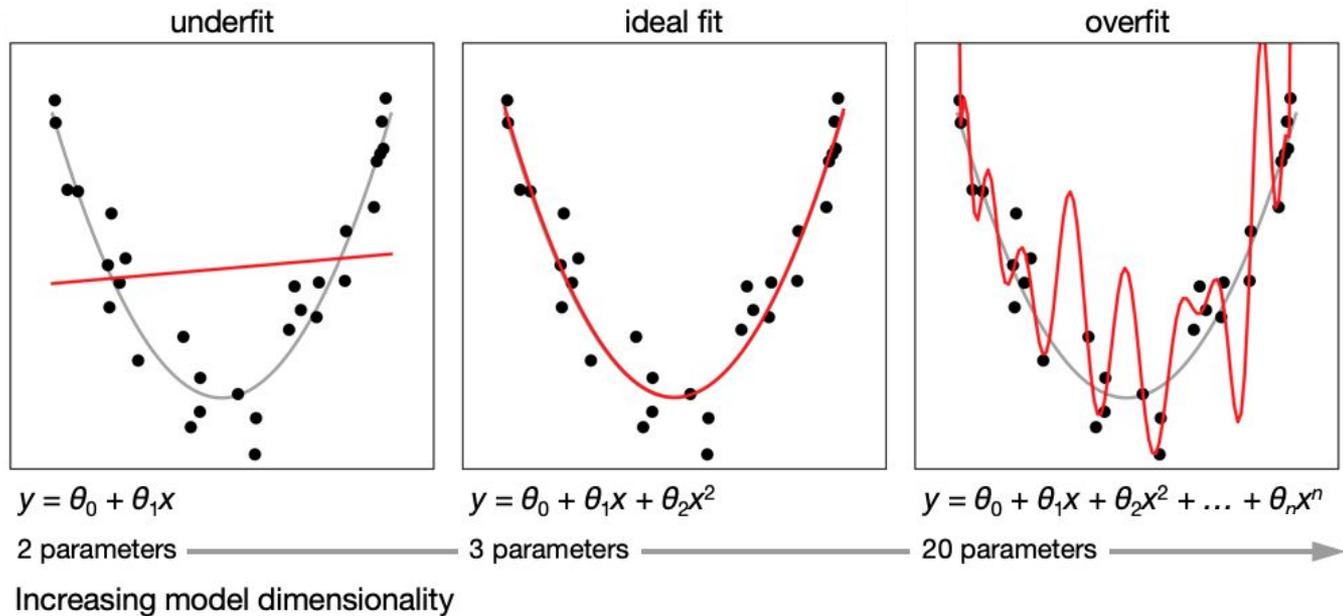
Scientists favor concise, interpretable models with an emphasis on explanatory power over predictive power.



What makes a good model?

The textbook example of underfitting/overfitting

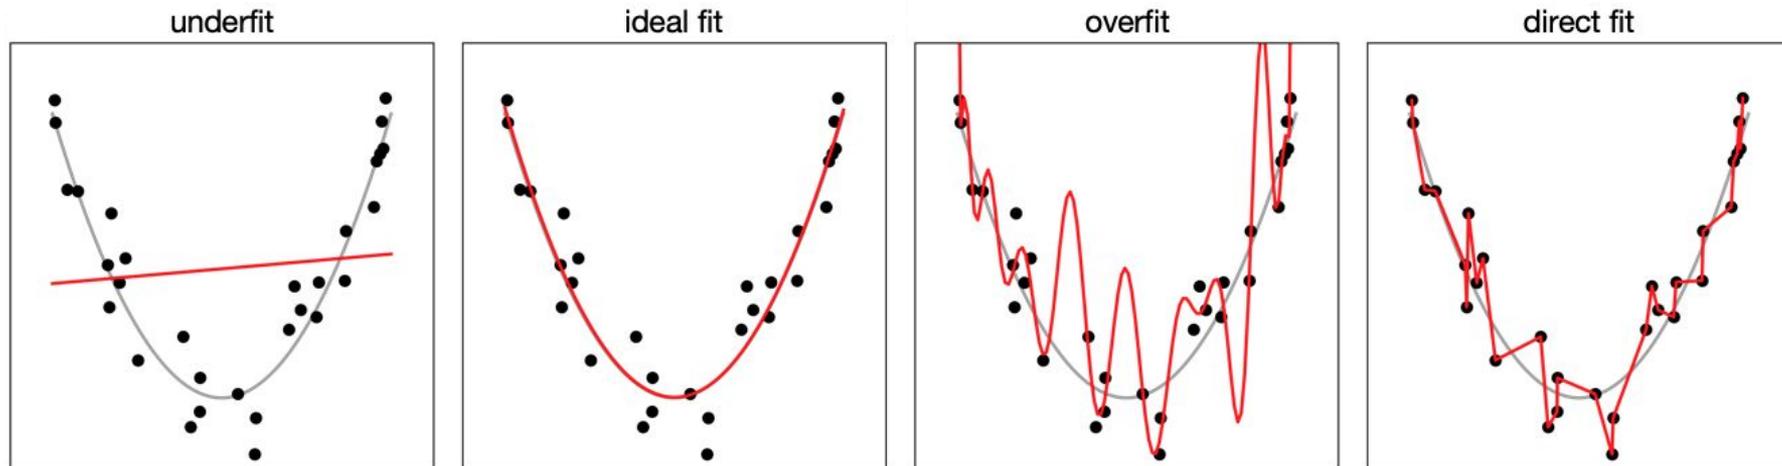
—overly simplistic models will underfit the data, whereas overly complex models will overfit



What makes a good model?

The textbook example of underfitting/overfitting

—overly simplistic models will underfit the data, whereas overly complex models will overfit— **right?**



$$y = \theta_0 + \theta_1 x$$

2 parameters

$$y = \theta_0 + \theta_1 x + \theta_2 x^2$$

3 parameters

$$y = \theta_0 + \theta_1 x + \theta_2 x^2 + \dots + \theta_n x^n$$

20 parameters

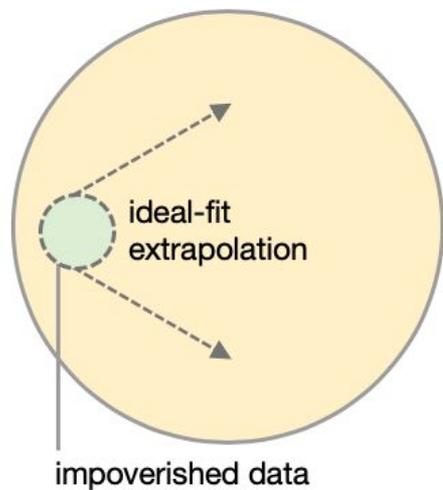
1,000,000 parameters

Increasing model dimensionality

Rethinking overparameterized models

Two types of generalization: interpolation and extrapolation

—experimentalists typically operate in a well-controlled, narrow sampling aperture and seek to discover simple models that will **extrapolate** to novel contexts



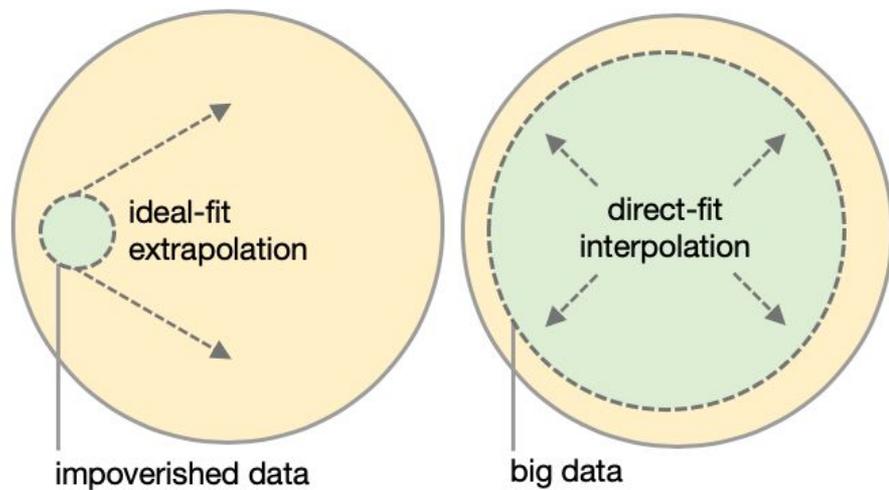
In this constrained sampling regime, ideal-fit **extrapolation** is required for generalization. **Interpolation** only provides weak generalization.

Rethinking overparameterized models

Two types of generalization: interpolation and extrapolation

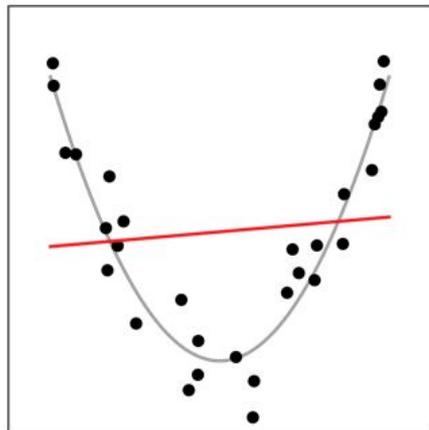
—experimentalists typically operate in a well-controlled, narrow sampling aperture and seek to discover simple models that will **extrapolate** to novel contexts

—machine learning research leverages “big data” and overparameterized models to expand the **interpolation zone**



Rethinking overparameterized models

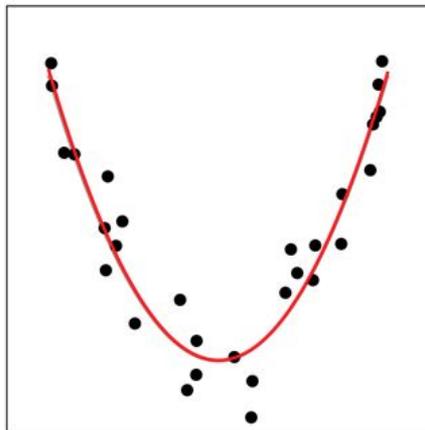
underfit



$$y = \theta_0 + \theta_1 x$$

2 parameters

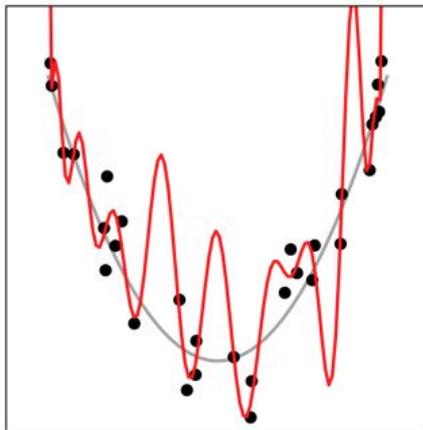
ideal fit



$$y = \theta_0 + \theta_1 x + \theta_2 x^2$$

3 parameters

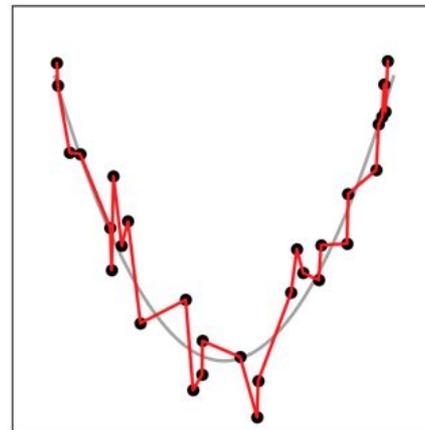
overfit



$$y = \theta_0 + \theta_1 x + \theta_2 x^2 + \dots + \theta_n x^n$$

20 parameters

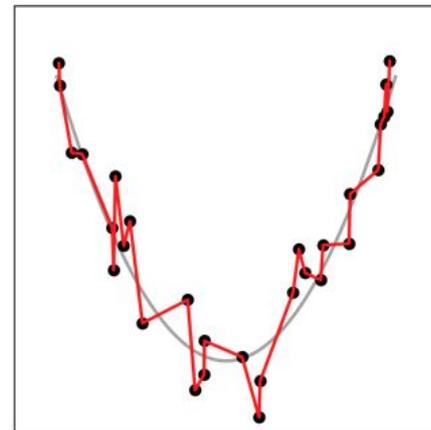
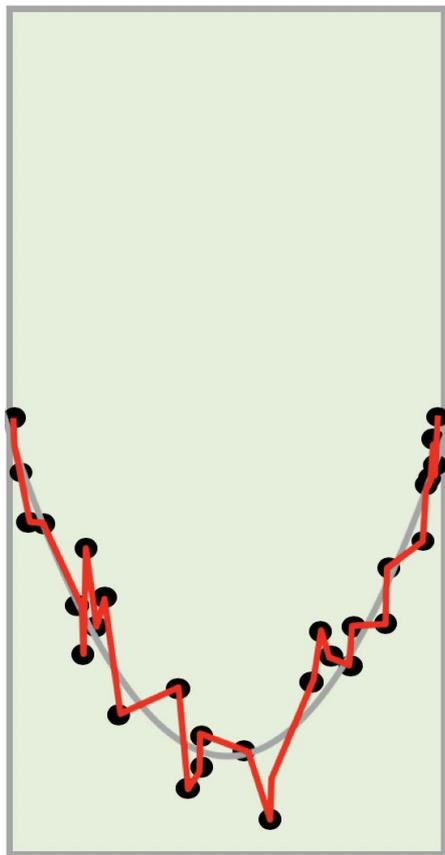
direct fit



1,000,000 parameters

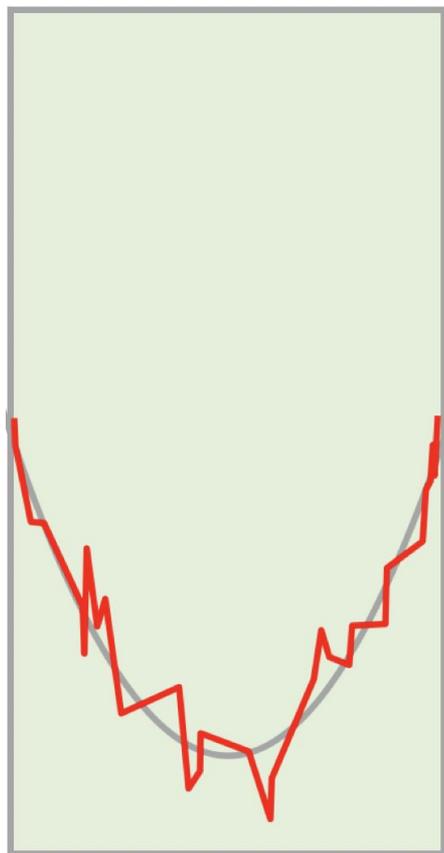
Increasing model dimensionality

Rethinking overparameterized models



● many training samples

Rethinking overparameterized models

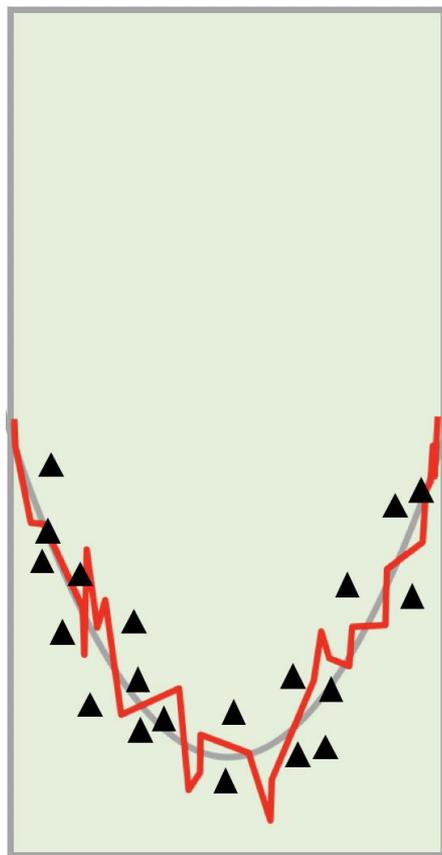


Which is the better model?

— ideal model

— direct-fit model

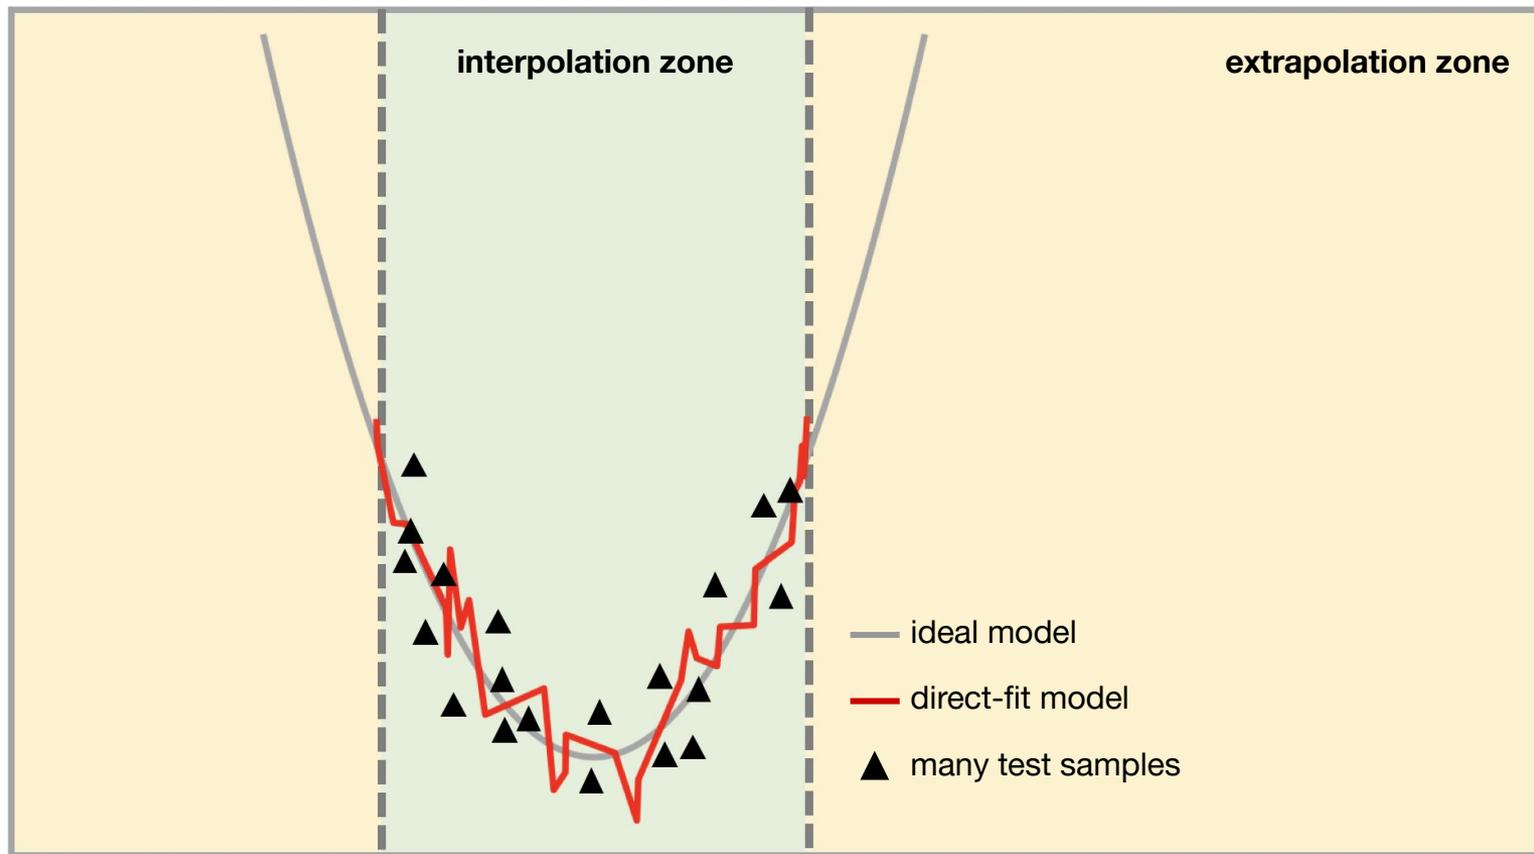
Rethinking overparameterized models



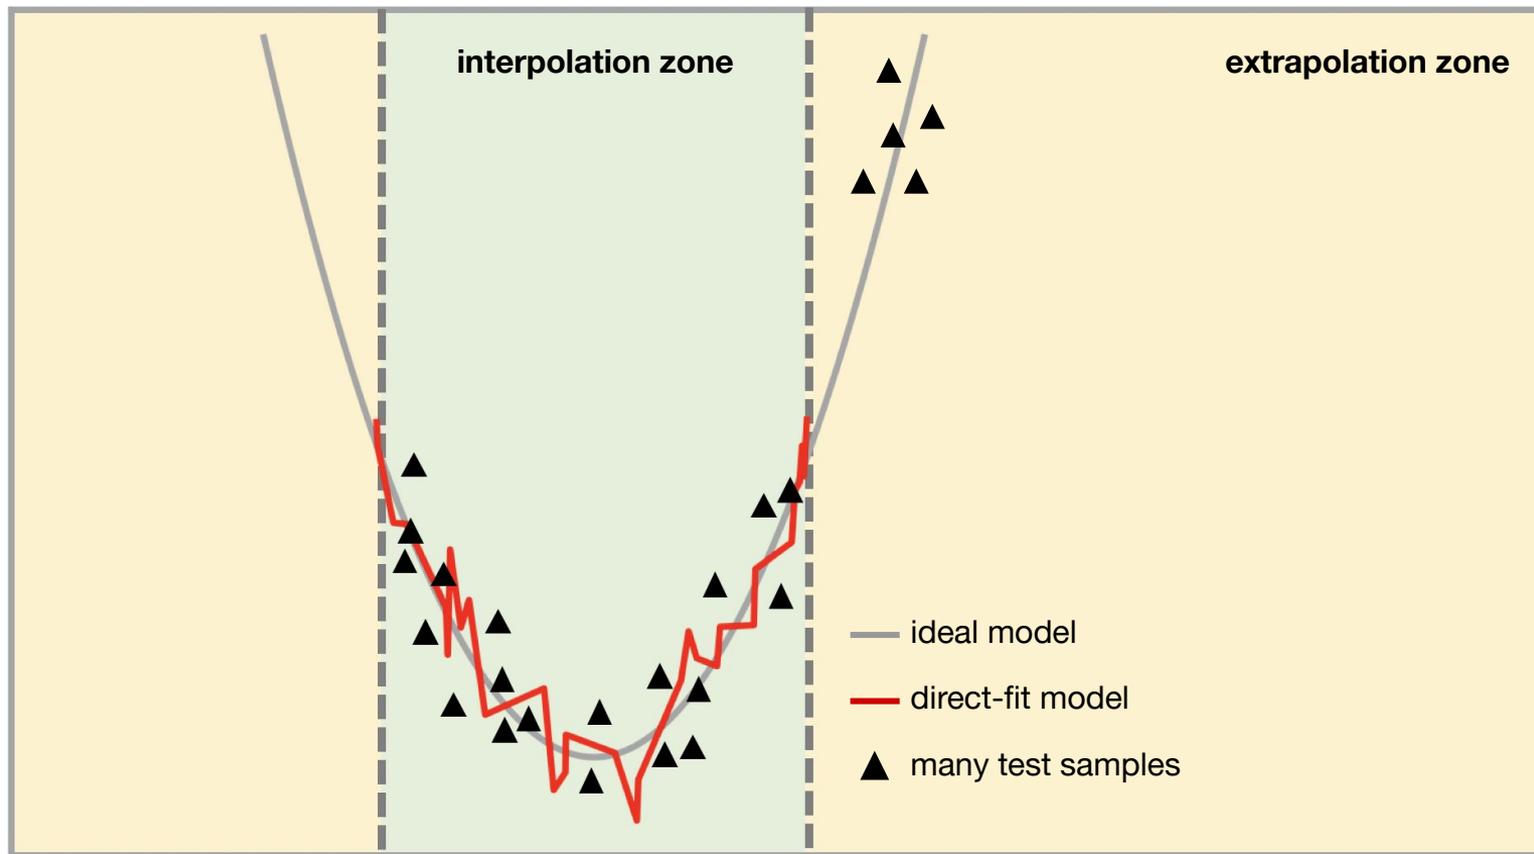
Which is the better model?

- ideal model
- direct-fit model
- ▲ many test samples

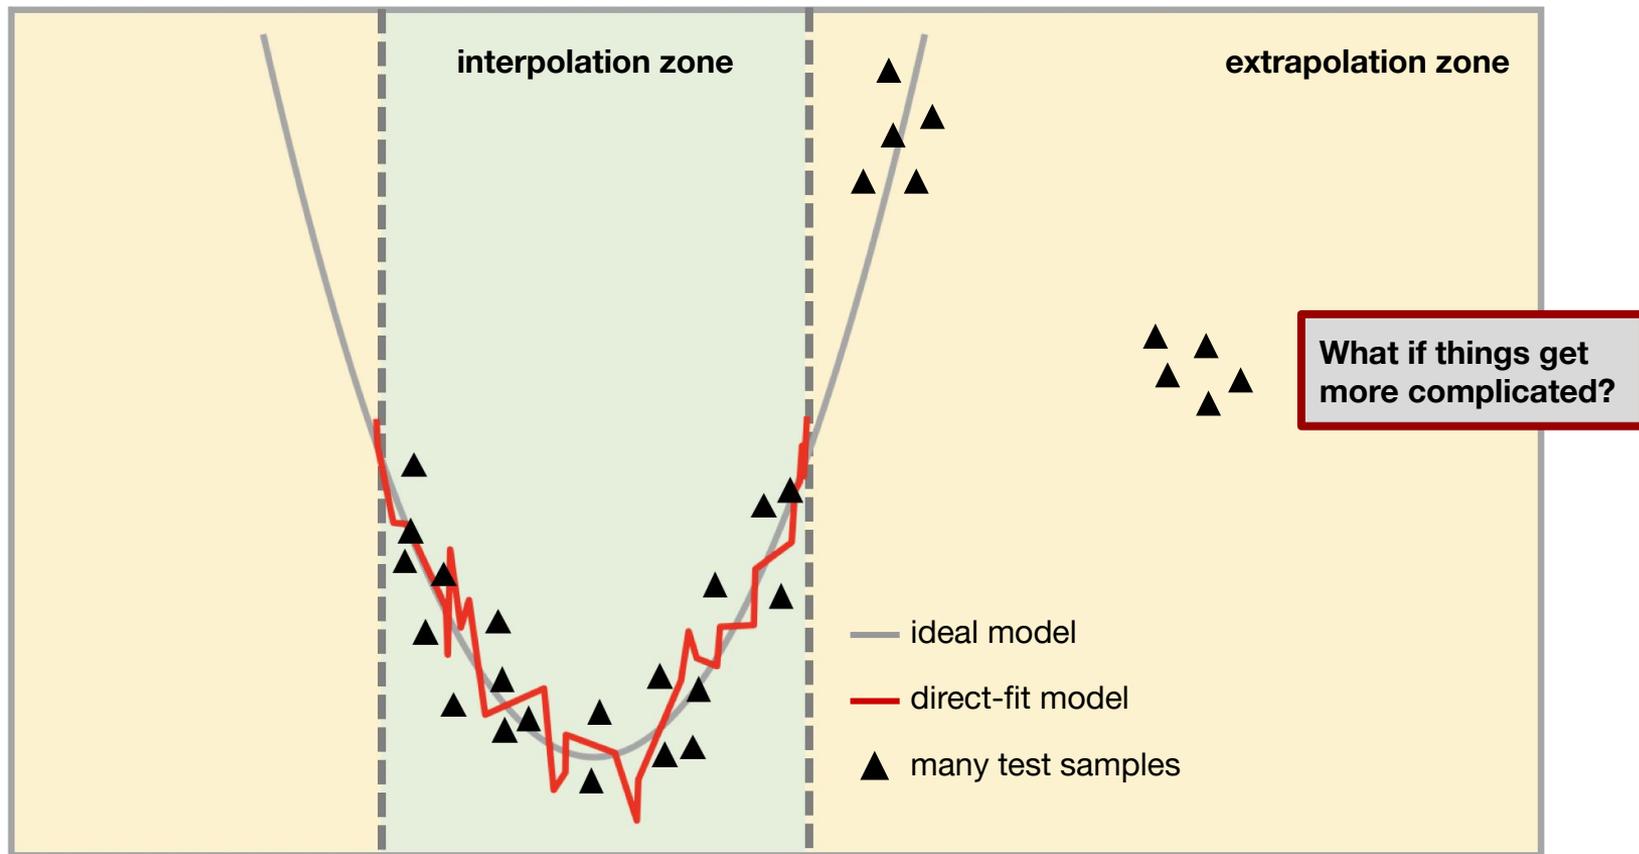
Rethinking overparameterized models



Rethinking overparameterized models



Rethinking overparameterized models

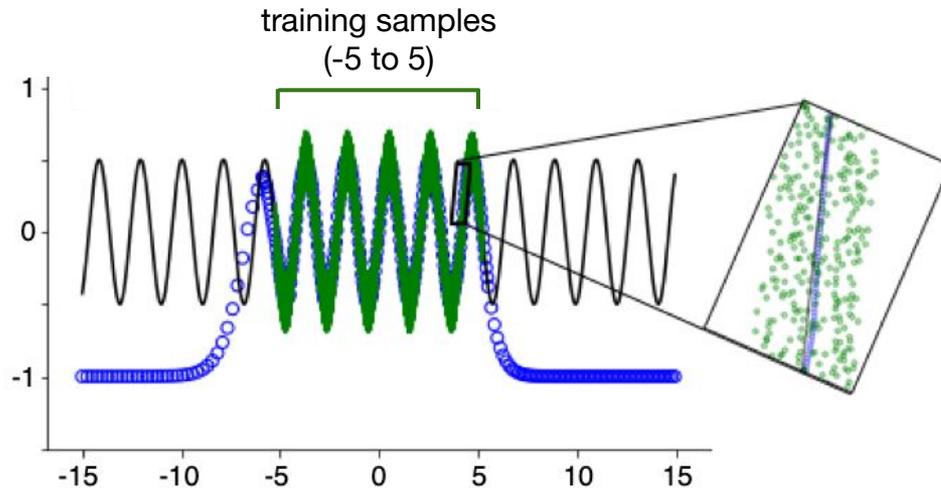


Interpolation and extrapolation

Example—interpolation over space

—simple ANN with 1 input unit, 1 output unit, and three fully connected hidden layers with 300 units each (180,600 parameters)

—trained to predict y values from x values given input–output mappings reflecting by a sine function: **training samples** (even values), **test samples** (odd values)

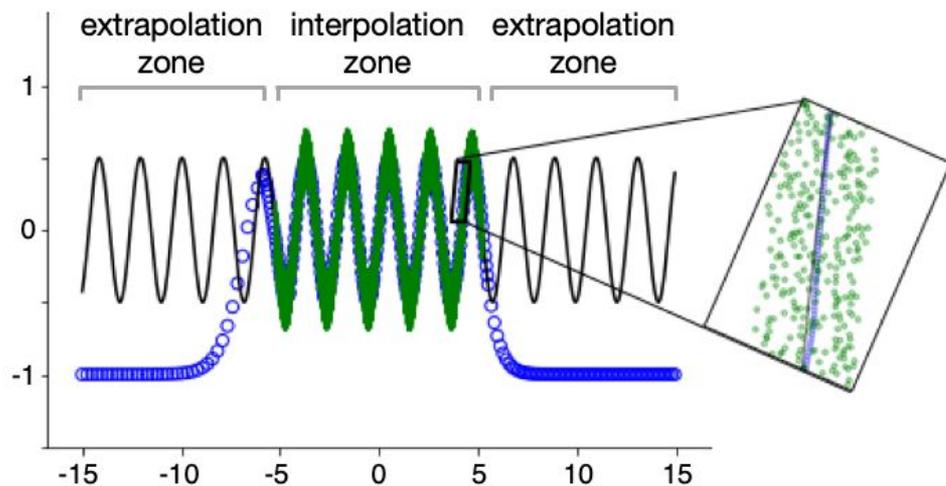


Interpolation and extrapolation

Example—interpolation over space

—simple ANN with 1 input unit, 1 output unit, and three fully connected hidden layers with 300 units each (180,600 parameters)

—trained to predict y values from x values given input–output mappings reflecting by a sine function: **training samples** (even values), **test samples** (odd values)

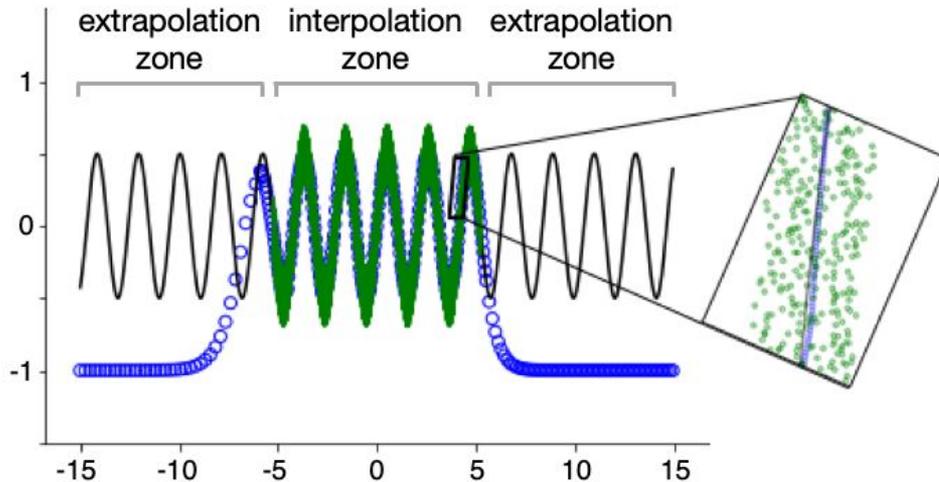


Interpolation and extrapolation

Example—interpolation over space

—simple ANN with 1 input unit, 1 output unit, and three fully connected hidden layers with 300 units each (180,600 parameters)

—trained to predict y values from x values given input–output mappings reflecting by a sine function: **training samples** (even values), **test samples** (odd values)

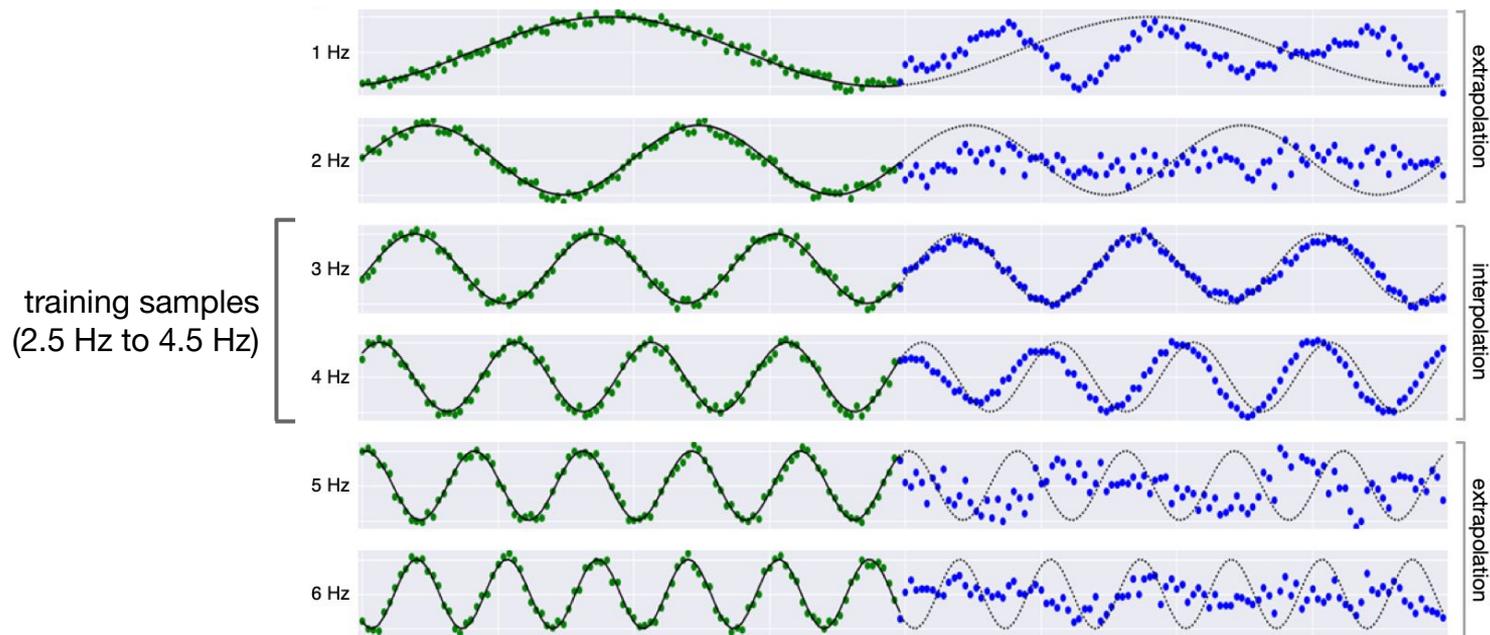


This overparameterized model generalizes well within the scope of the training sample—i.e., the **interpolation zone**. However, the model does not actually learn the ideal sine function necessary for generalization into the **extrapolation zone**.

Interpolation and extrapolation

Example—interpolation over time

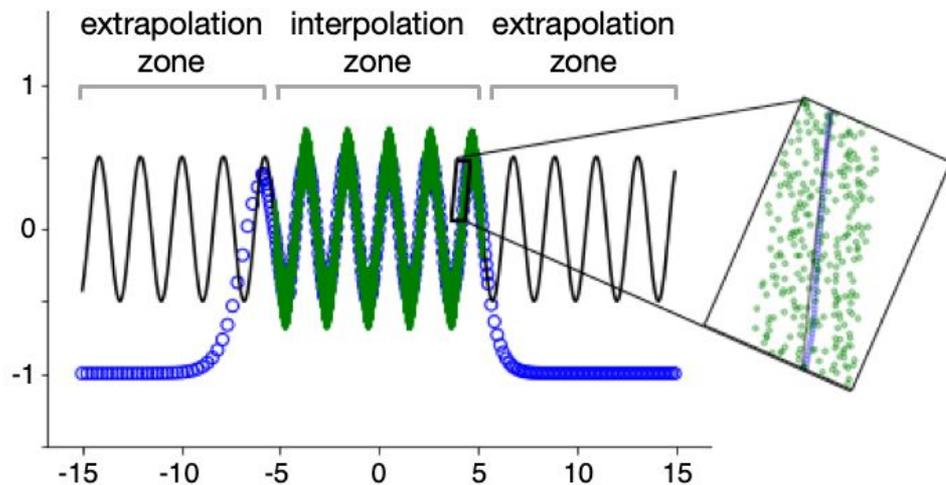
—simple RNN (LSTM) trained to predict the **future sequence** of y values based on the **preceding sequence** of 100 y values sampled within a 1-second input window



The robustness of direct fit

Has this neural network “learned” the sine function?

The model may *appear* to learn the sine function (in the interpolation zone), but this is an incidental byproduct of the input and the fitting procedure.

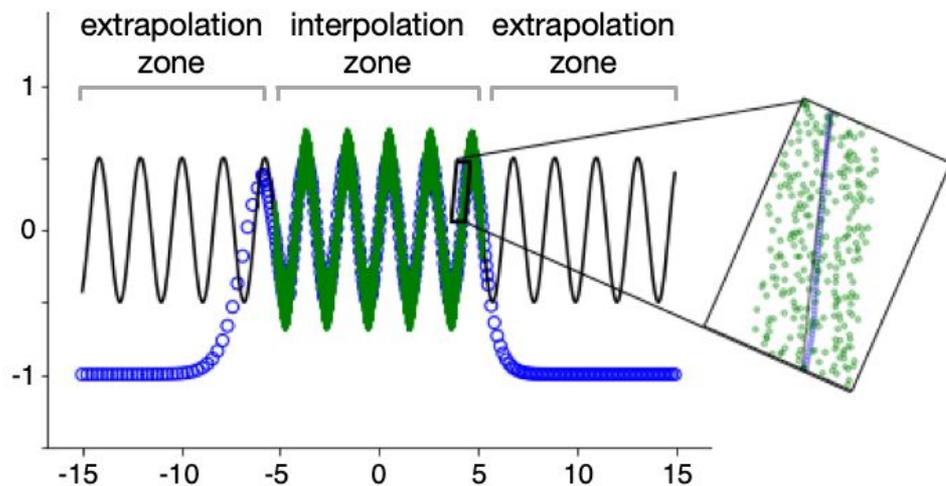


The robustness of direct fit

Has this neural network “learned” the sine function?

The model may *appear* to learn the sine function (in the interpolation zone), but this is an incidental byproduct of the input and the fitting procedure.

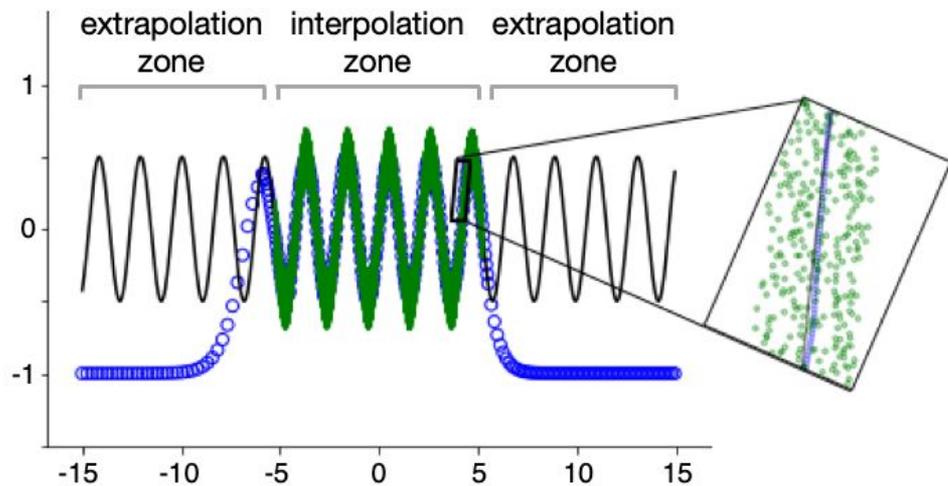
We can interrogate the model for “representations” of the sine function, but these only exist because we injected them into the training data.



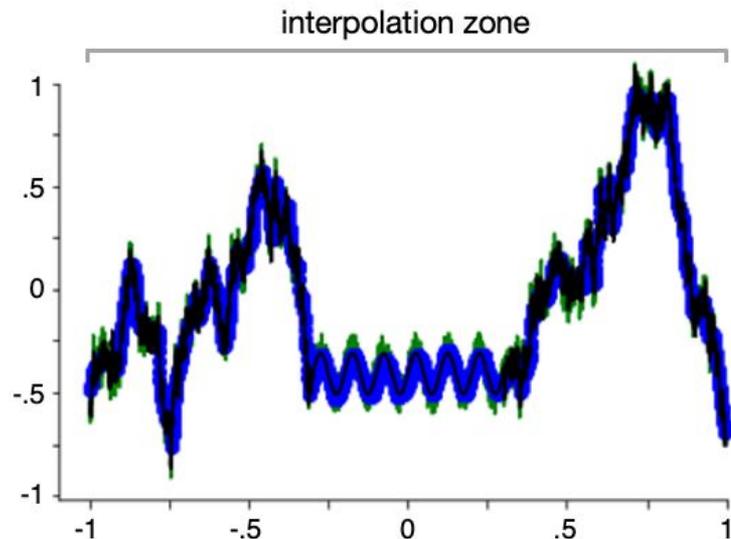
The robustness of direct fit

What's the benefit of interpolation-based direct-fit learning?

—in a more naturalistic “big data” regime, the overparameterized model can learn the structure of a more complex world



Meanwhile—outside the laboratory...



Face recognition in the wild

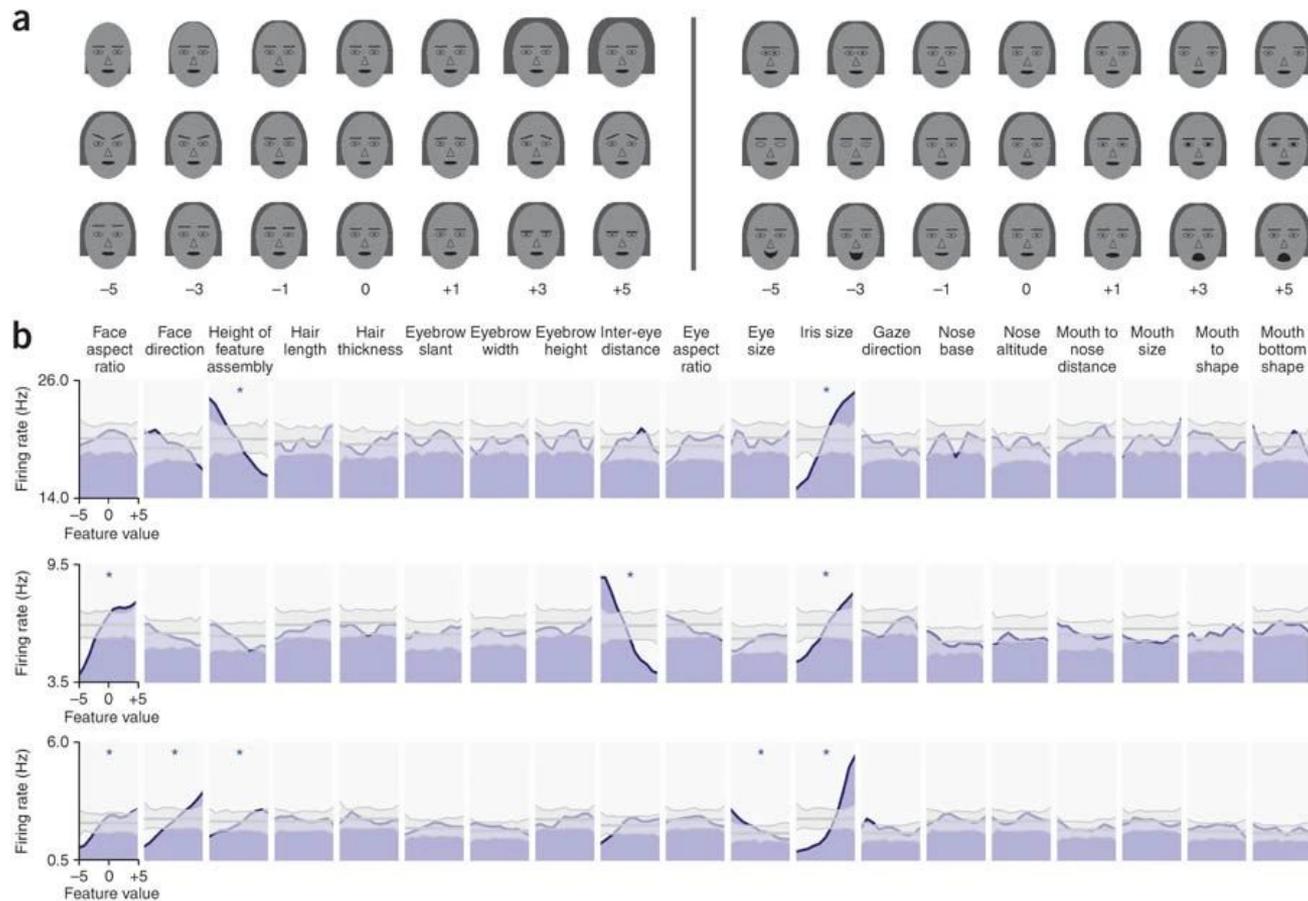
FaceNet—an example at scale

Superhuman performance (95%+ accuracy) based on dense training set of 200 million faces comprising 8 million identities.

Schroff et al, *CVPR*, 2015

layer	size-in	size-out	kernel	param	FLPS
conv1	220×220×3	110×110×64	7×7×3, 2	9K	115M
pool1	110×110×64	55×55×64	3×3×64, 2	0	
rnorm1	55×55×64	55×55×64		0	
conv2a	55×55×64	55×55×64	1×1×64, 1	4K	13M
conv2	55×55×64	55×55×192	3×3×64, 1	111K	335M
rnorm2	55×55×192	55×55×192		0	
pool2	55×55×192	28×28×192	3×3×192, 2	0	
conv3a	28×28×192	28×28×192	1×1×192, 1	37K	29M
conv3	28×28×192	28×28×384	3×3×192, 1	664K	521M
pool3	28×28×384	14×14×384	3×3×384, 2	0	
conv4a	14×14×384	14×14×384	1×1×384, 1	148K	29M
conv4	14×14×384	14×14×256	3×3×384, 1	885K	173M
conv5a	14×14×256	14×14×256	1×1×256, 1	66K	13M
conv5	14×14×256	14×14×256	3×3×256, 1	590K	116M
conv6a	14×14×256	14×14×256	1×1×256, 1	66K	13M
conv6	14×14×256	14×14×256	3×3×256, 1	590K	116M
pool4	14×14×256	7×7×256	3×3×256, 2	0	
concat	7×7×256	7×7×256		0	
fc1	7×7×256	1×32×128	maxout p=2	103M	103M
fc2	1×32×128	1×32×128	maxout p=2	34M	34M
fc7128	1×32×128	1×1×128		524K	0.5M
L2	1×1×128	1×1×128		0	
total				140M	1.6B

Face recognition in the wild



Freiwald et al,
Nat Neurosci, 2009

Face recognition in the wild



How many unique individuals are in this set of images?

Jenkins et al,
Cognition, 2011

Face recognition in the wild



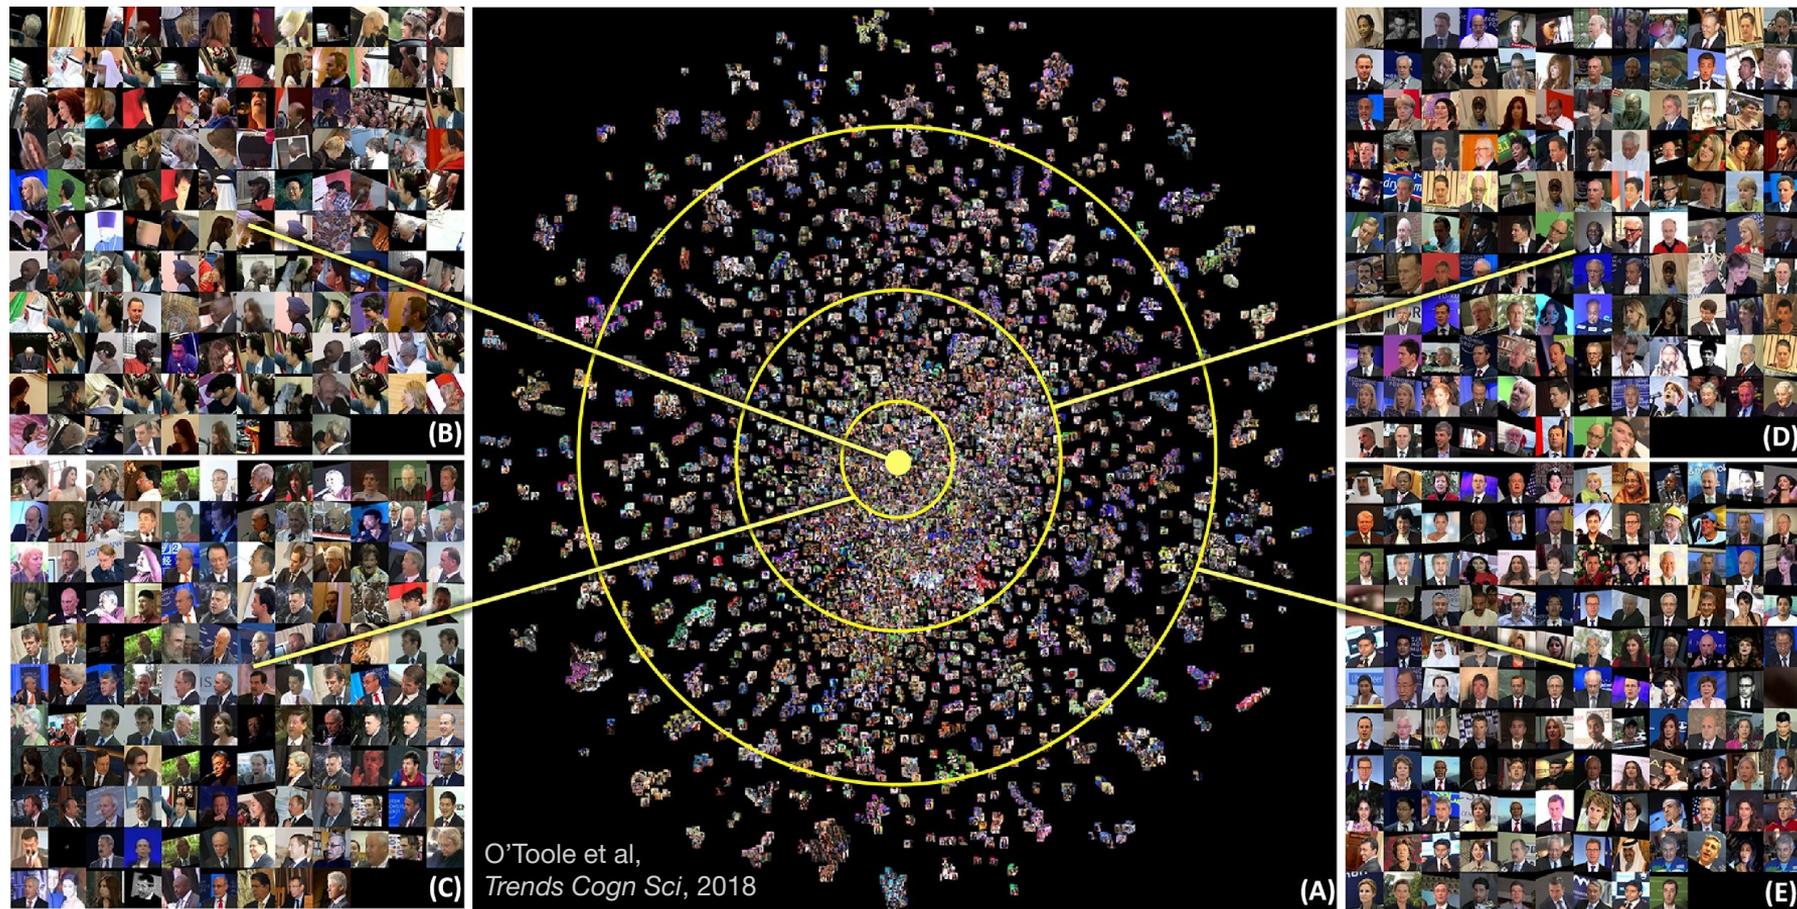
How many unique individuals are in this set of images?

Jenkins et al,
Cognition, 2011

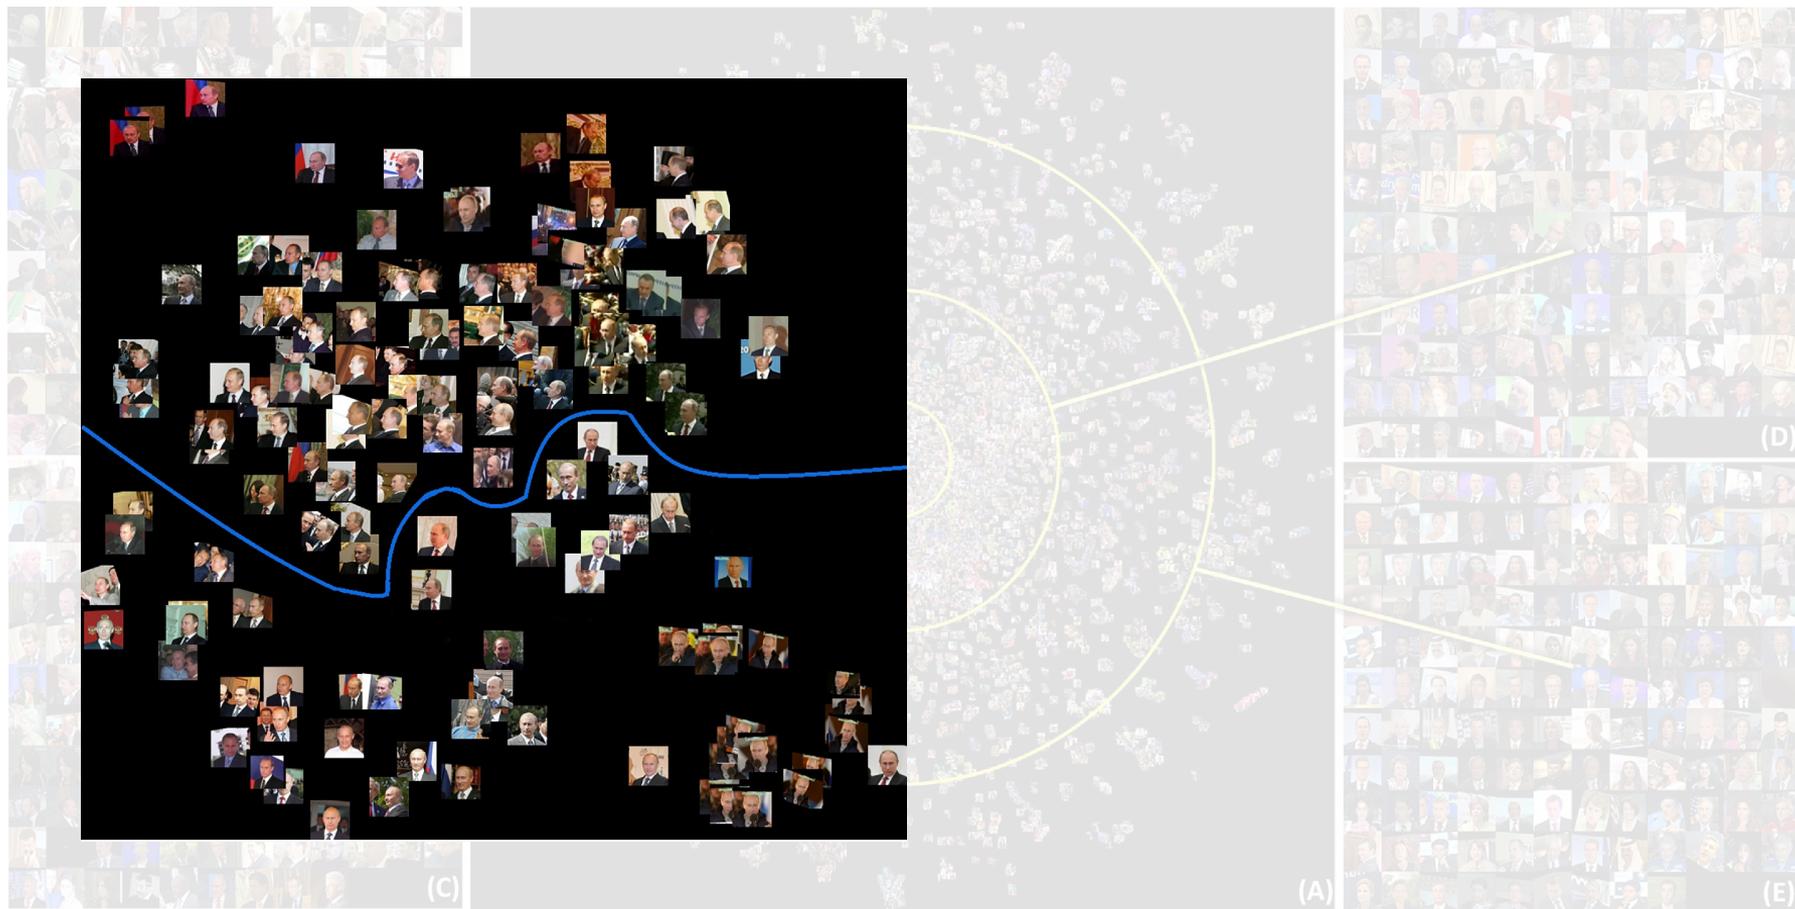
People know roughly 5,000 faces.

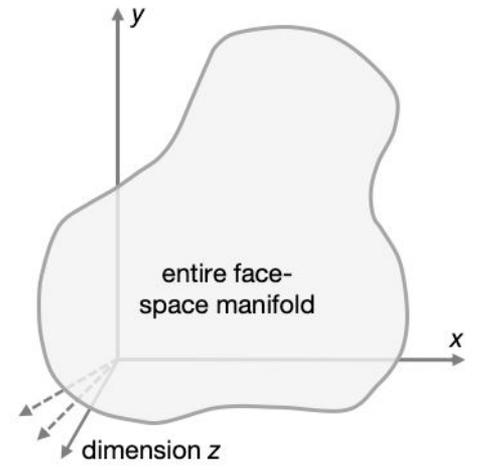
Jenkins et al, *Proc R Soc Lond B Biol Sci*, 2018

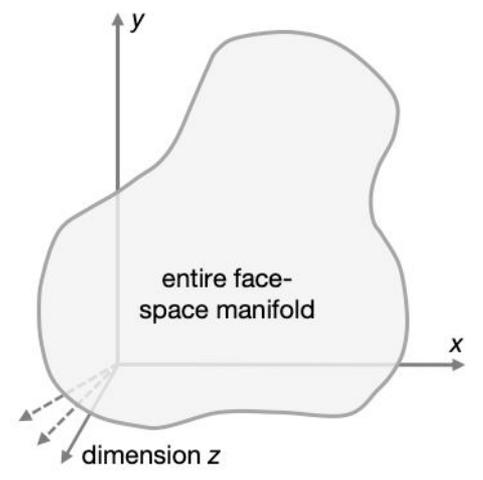
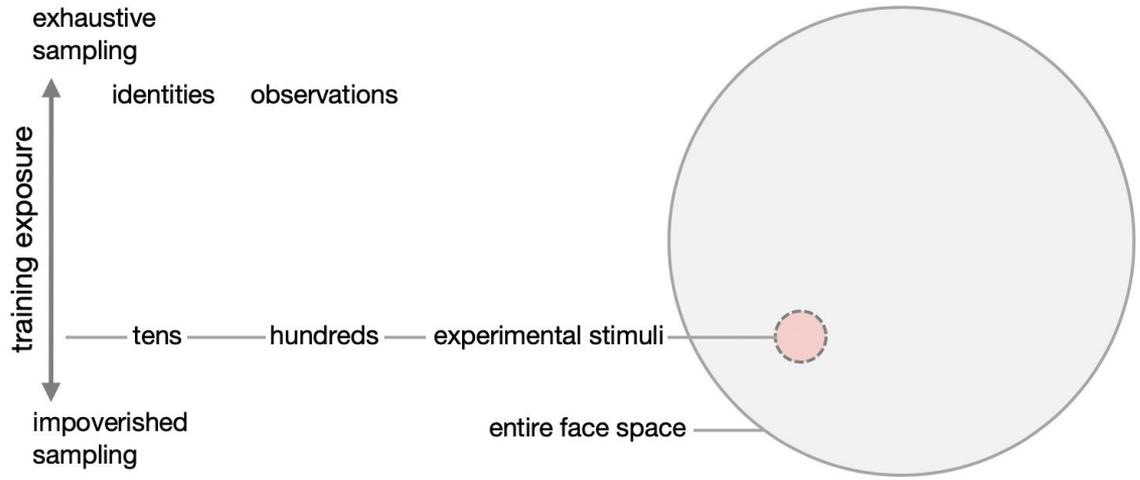
Face recognition in the wild

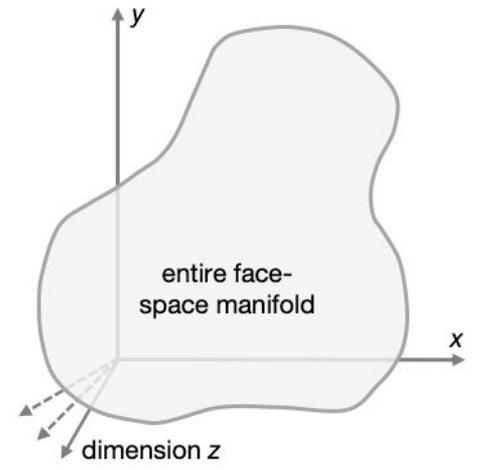
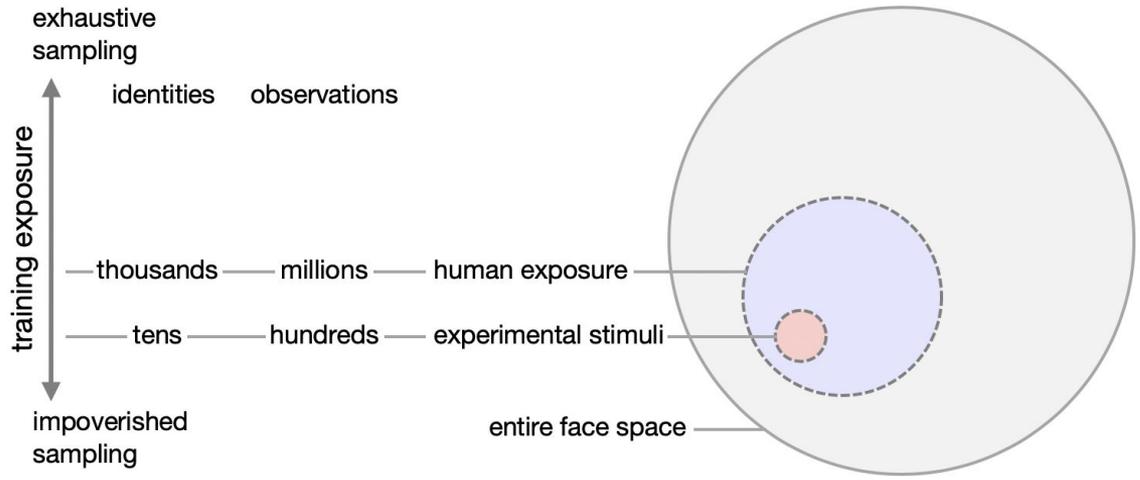


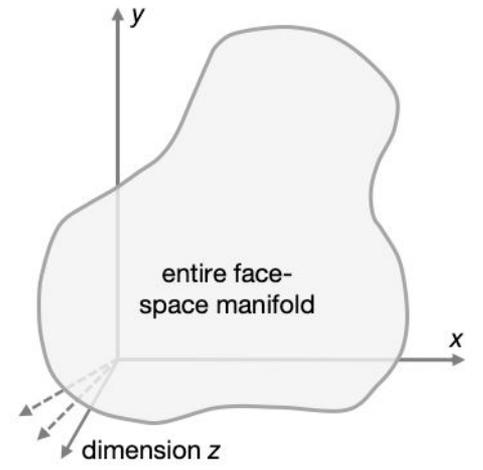
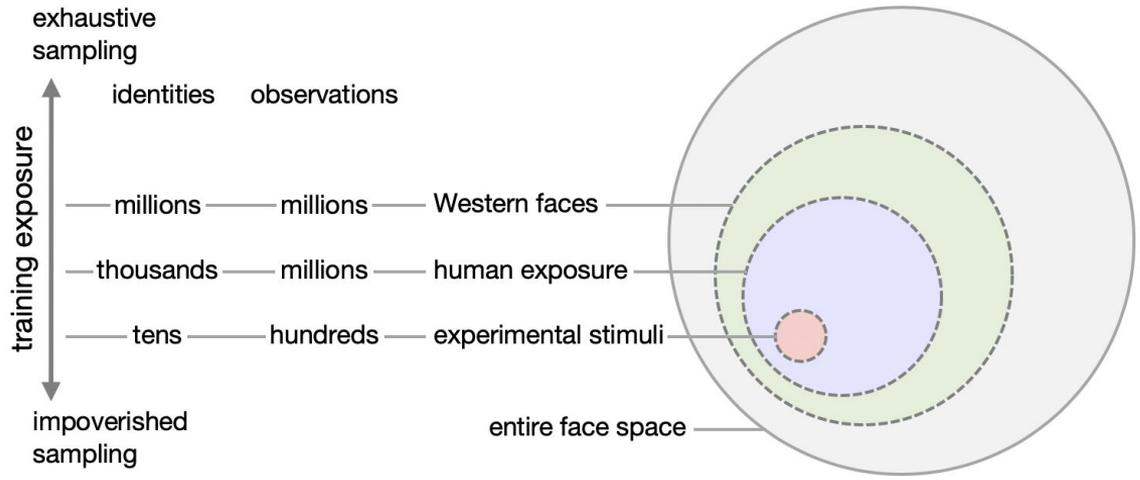
Face recognition in the wild

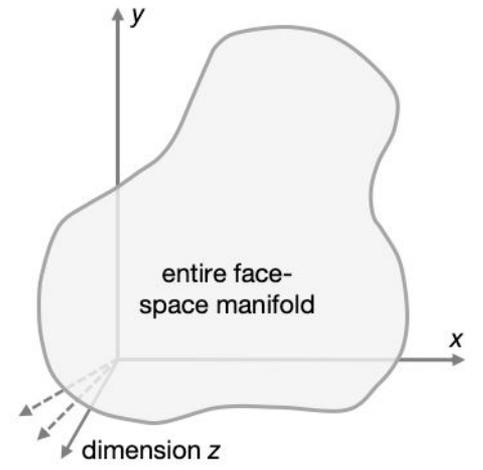
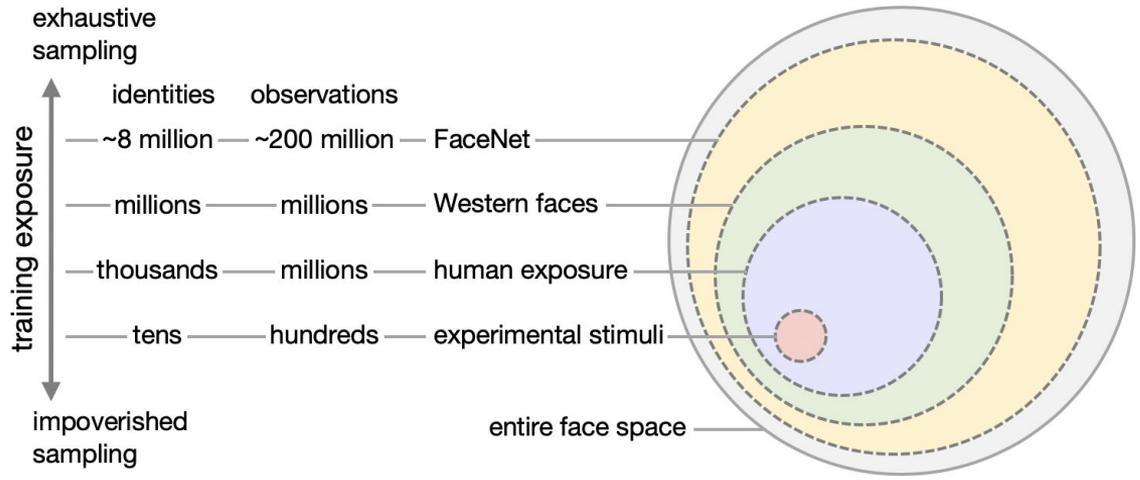


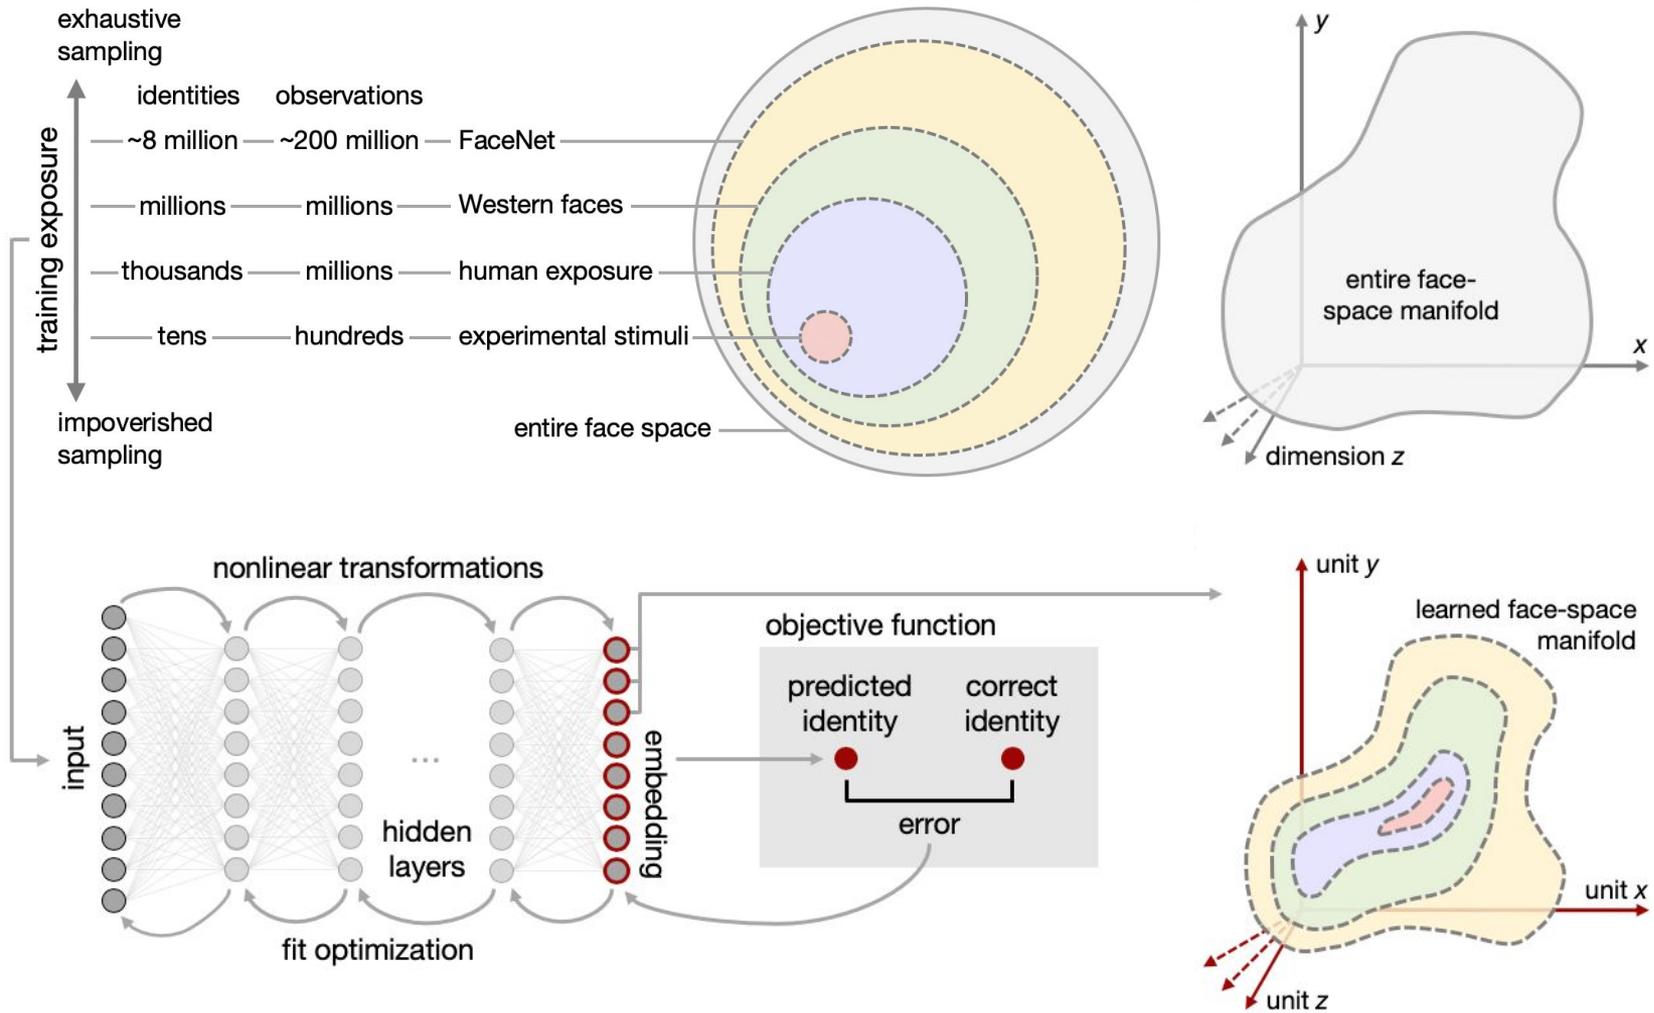












The “direct fit” model of biological learning

Relinquishing control—machine learning

Modern artificial neural networks leverage millions of noisy, real-world samples to achieve shockingly good performance at a variety of “cognitive” tasks...

—how is this possible?

—**and why do neuroscientists find this so unsatisfying?**

The “black box” argument

Artificial neural networks provide an input–output mapping with no explanation of their internal workings. We have simply duplicated the original problem by creating one more black box.

Ashby, 1956

McCloskey, *Psychol Sci*, 1991

The “direct fit” model of biological learning

Relinquishing control—machine learning

Modern artificial neural networks leverage millions of noisy, real-world samples to achieve shockingly good performance at a variety of “cognitive” tasks...

—how is this possible?

—**and why do neuroscientists find this so unsatisfying?**

The “black box” argument

Artificial neural networks provide an input–output mapping with no explanation of their internal workings. We have simply duplicated the original problem by creating one more black box.

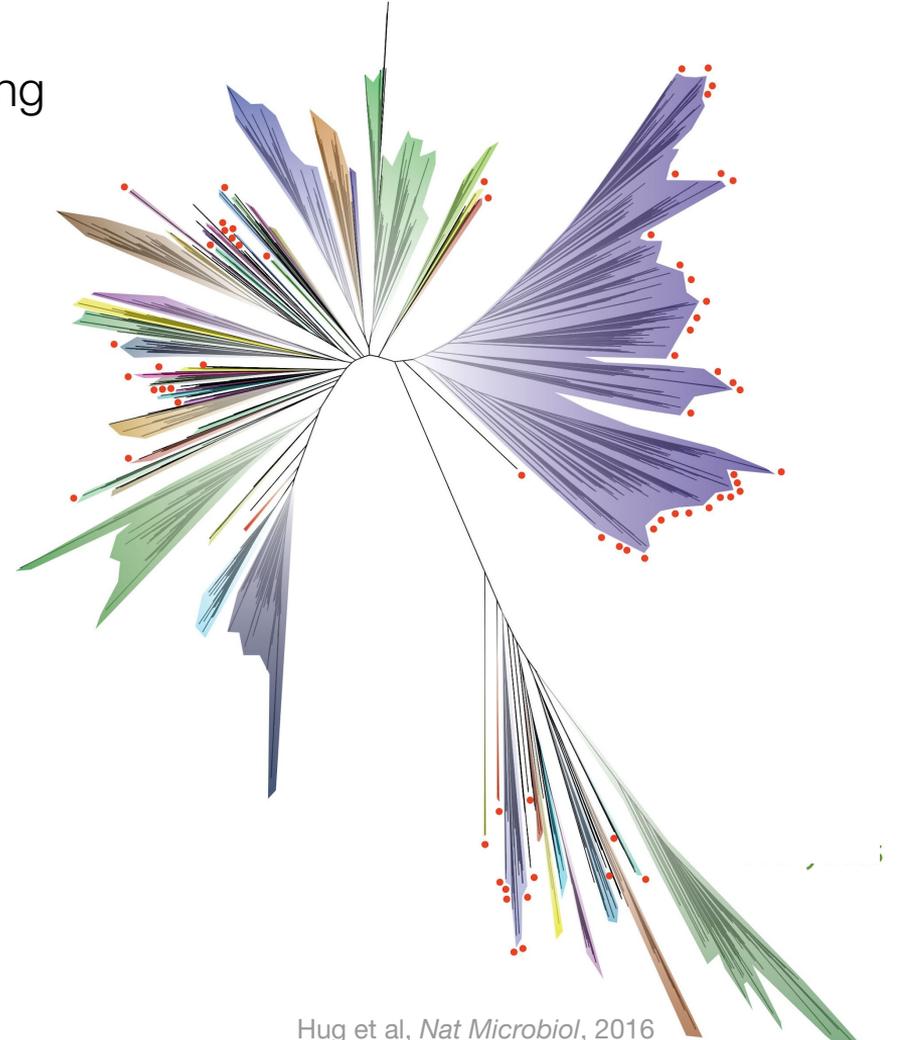
...because we didn’t find what we were looking for

Neuroscientists find this unsatisfying because we **assume** that to perform “cognitive” tasks, a neural network must learn concise, human-interpretable rules about the world.

Parallels between evolution and direct-fit learning

Evolution relies on “blind” optimization to produce complex organisms that over many generations become adapted to an ecological niche.

Evolution does not rely on idealized or interpretable model of the world, but relies on local interpolation to capitalize on “affordances” of the environment.



Parallels between evolution and direct-fit learning

Evolution

Over-production with variation. Over-production of organisms with variation is drives adaptation by sampling the space of feasible solutions for fitting to the environment.

Learning

Over-sampling with variation. Dense and broad sampling of observations expands the interpolation zone and allows the learner to effectively mimic extrapolation.

Parallels between evolution and direct-fit learning

Evolution

Over-production with variation. Over-production of organisms with variation is drives adaptation by sampling the space of feasible solutions for fitting to the environment.

Inheritance. Genes provide a stable substrate for accumulating adaptations (via mutation), thus optimizing the organism to the structure of the world over generations.

Learning

Over-sampling with variation. Dense and broad sampling of observations expands the interpolation zone and allows the learner to effectively mimic extrapolation.

Plasticity. Neural circuits provide a stable substrate for accumulating knowledge over the lifespan of an organism, thus optimizing synaptic weights to support adaptive behavior.

Parallels between evolution and direct-fit learning

Evolution

Over-production with variation. Over-production of organisms with variation is drives adaptation by sampling the space of feasible solutions for fitting to the environment.

Inheritance. Genes provide a stable substrate for accumulating adaptations (via mutation), thus optimizing the organism to the structure of the world over generations.

Combinatorial genetic code. The genetic code shared by all organisms is sufficiently high dimensional to express diverse organisms adapted to diverse ecological niches.

Learning

Over-sampling with variation. Dense and broad sampling of observations expands the interpolation zone and allows the learner to effectively mimic extrapolation.

Plasticity. Neural circuits provide a stable substrate for accumulating knowledge over the lifespan of an organism, thus optimizing synaptic weights to support adaptive behavior.

Combinatorial neural code. High-dimensional neural networks support a broad solution space and are expressive enough to learn complex, multidimensional structure in the world.

Parallels between evolution and direct-fit learning

Evolution

Over-production with variation. Over-production of organisms with variation is drives adaptation by sampling the space of feasible solutions for fitting to the environment.

Inheritance. Genes provide a stable substrate for accumulating adaptations (via mutation), thus optimizing the organism to the structure of the world over generations.

Combinatorial genetic code. The genetic code shared by all organisms is sufficiently high dimensional to express diverse organisms adapted to diverse ecological niches.

Natural selection. External forces and ecological constraints guide evolution by punishing for maladaptive variations.

Learning

Over-sampling with variation. Dense and broad sampling of observations expands the interpolation zone and allows the learner to effectively mimic extrapolation.

Plasticity. Neural circuits provide a stable substrate for accumulating knowledge over the lifespan of an organism, thus optimizing synaptic weights to support adaptive behavior.

Combinatorial neural code. High-dimensional neural networks support a broad solution space and are expressive enough to learn complex, multidimensional structure in the world.

Objective functions. Self-supervised and externally-supervised objective functions guide learning by differentially adjusting synaptic weights to reinforce adaptive behaviors.

Parallels between evolution and direct-fit learning

Evolution

Over-production with variation. Over-production of organisms with variation drives adaptation by sampling the space of feasible solutions for fitting to the environment.

Inheritance. Genes provide a stable substrate for accumulating adaptations (via mutation), thus optimizing the organism to the structure of the world over generations.

Combinatorial genetic code. The genetic code shared by all organisms is sufficiently high dimensional to express diverse organisms adapted to diverse ecological niches.

Natural selection. External forces and ecological constraints guide evolution by punishing for maladaptive variations.

Iteration over generations. Evolution is an iterative optimization process across generations that yields diverse organisms adapted to their ecological niche over billions of years.

Learning

Over-sampling with variation. Dense and broad sampling of observations expands the interpolation zone and allows the learner to effectively mimic extrapolation.

Plasticity. Neural circuits provide a stable substrate for accumulating knowledge over the lifespan of an organism, thus optimizing synaptic weights to support adaptive behavior.

Combinatorial neural code. High-dimensional neural networks support a broad solution space and are expressive enough to learn complex, multidimensional structure in the world.

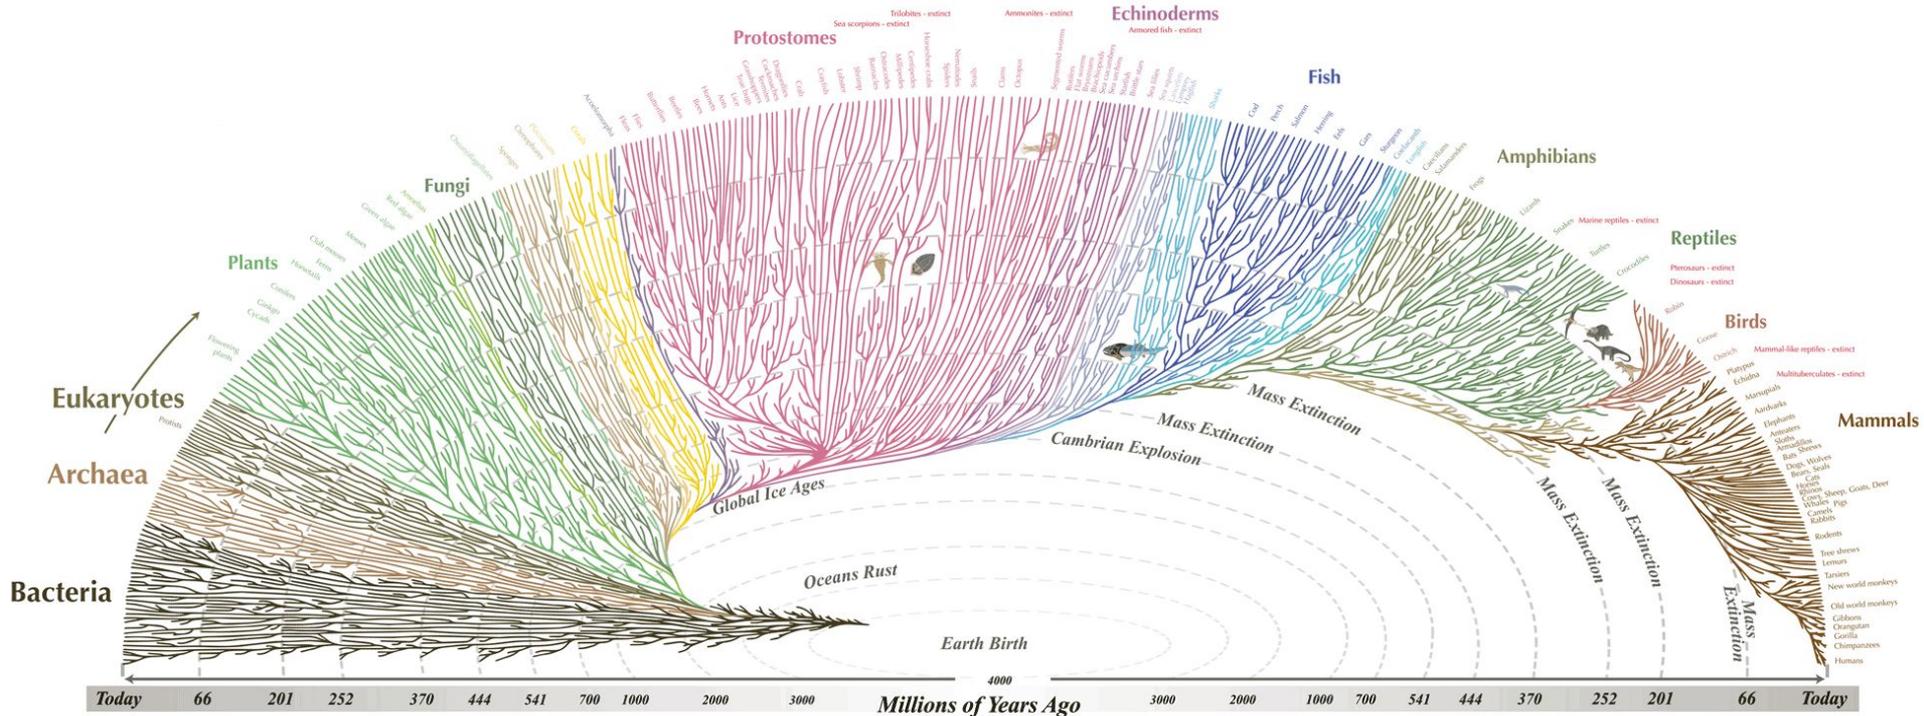
Objective functions. Self-supervised and externally-supervised objective functions guide learning by differentially adjusting synaptic weights to reinforce adaptive behaviors.

Iteration over samples. Direct-fit learning is an iterative optimization process over many percepts and actions over the lifespan that affords an organism flexible, adaptive behavior within an ecological niche.

Parallels between evolution and direct-fit learning

Evolutionary theory—an elegant model?

Evolution does not have the luxury of operating in an idealized, highly controlled parameter space (like the experimenter's laboratory) and neither does biological learning.



Ecological learning and embodiment

Neural networks still have a long way to go

- current models are very limited and brittle
- current models rely on ad hoc domain-specific engineering “tricks”

Ecological learning and embodiment

Neural networks still have a long way to go

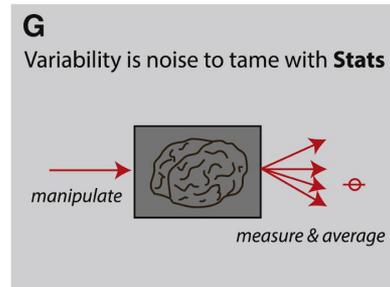
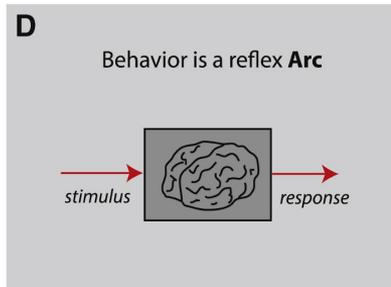
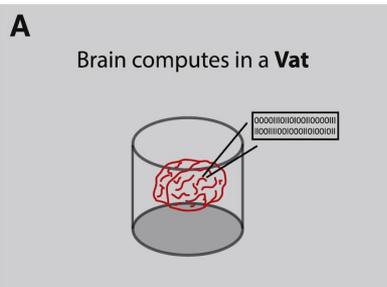
- current models are very limited and brittle
- current models rely on ad hoc domain-specific engineering “tricks”

Evolution, embodiment, and the primacy of action

- all biological learning systems were designed by evolution to produce adaptive behavior
- evolution and learning are continuous organism–environment feedback loops
- we may need to reorient our criteria for “understanding” neural networks

Richards et al, *Nat Neurosci*, 2019

Three principles for living organisms: embodiment, agency and historicity



Inert **INFORMATION**

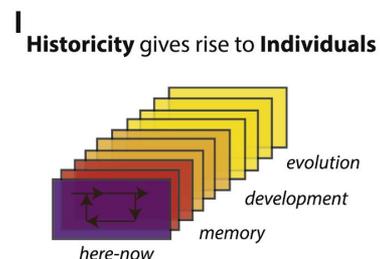
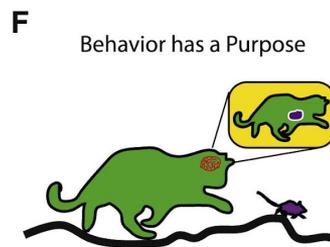
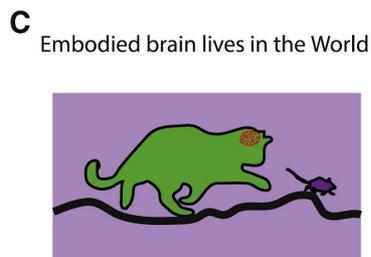
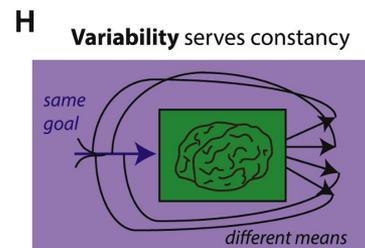
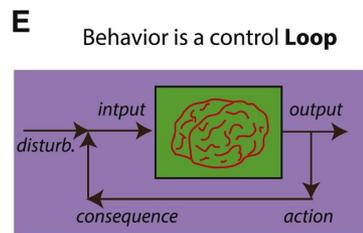
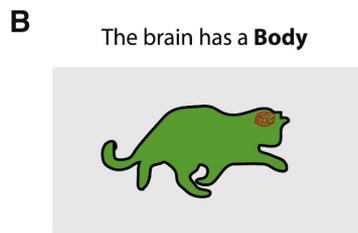
INERTIA

UNIVERSALS

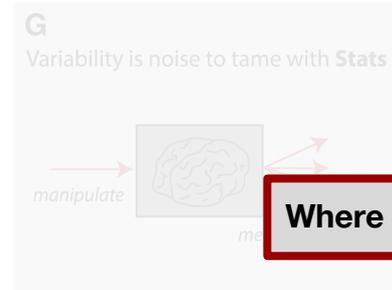
Living **EMBODIMENT / UMWELT**

AGENCY

PARTICULARS



Three principles for living organisms: embodiment, agency and historicity



Where does "System 2" emerge?

Inert

INFORMATION

INERTIA

UNIVERSALS

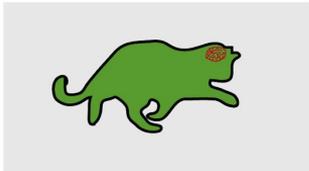
Living

EMBODIMENT / UMWELT

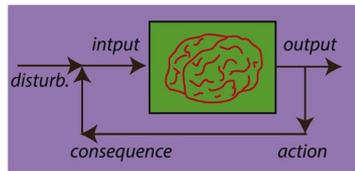
AGENCY

PARTICULARS

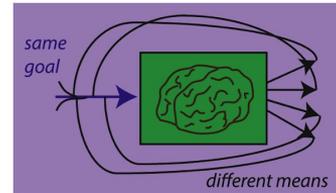
B
The brain has a **Body**



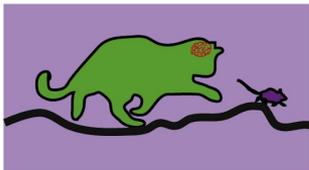
E
Behavior is a control **Loop**



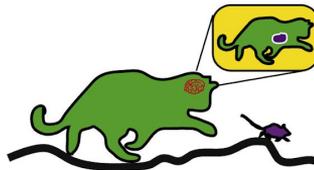
H
Variability serves constancy



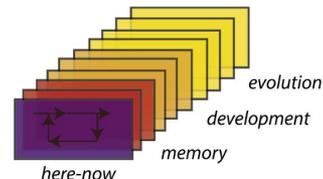
C
Embodied brain lives in the World



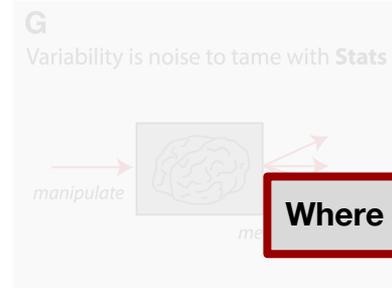
F
Behavior has a Purpose



I
Historicity gives rise to **Individuals**



Three principles for living organisms: embodiment, agency and historicity



Where does "System 2" emerge?

Inert **INFORMATION**

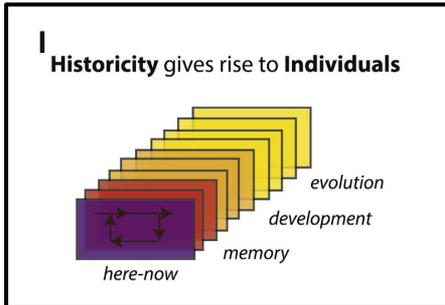
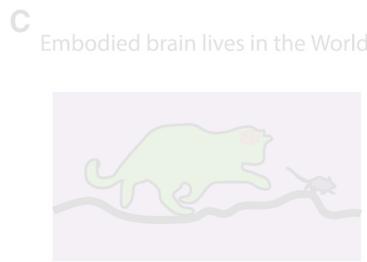
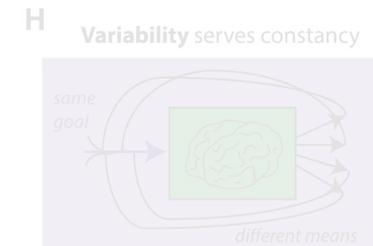
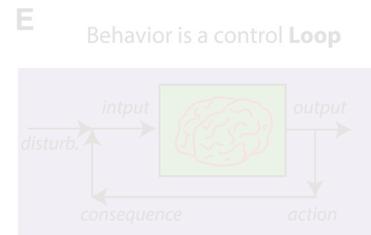
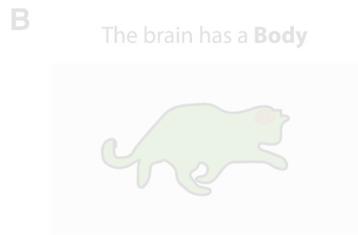
INERTIA

UNIVERSALS

Living **EMBODIMENT / UMWELT**

AGENCY

PARTICULARS



The “direct fit” model of biological learning

Relinquishing control—machine learning

Modern artificial neural networks leverage millions of noisy, real-world samples to achieve shockingly good performance at a variety of “cognitive” tasks...

—how is this possible?

—**and why do neuroscientists find this so unsatisfying?**

The “direct fit” model of biological learning

Relinquishing control—machine learning

Modern artificial neural networks leverage millions of noisy, real-world samples to achieve shockingly good performance at a variety of “cognitive” tasks...

—how is this possible?

—**and why do neuroscientists find this so unsatisfying?**

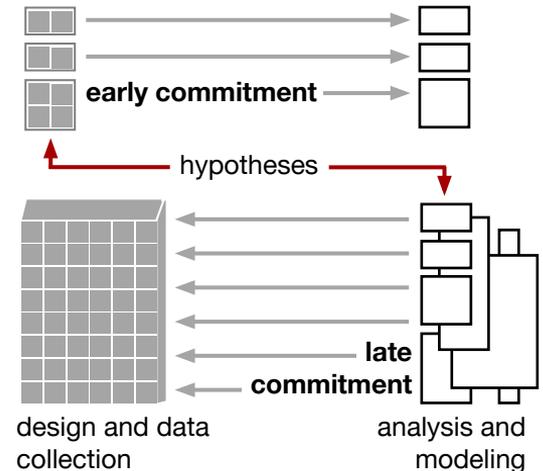
Where now for the neuroscientist?

—ecological considerations must play a central role in both model development and model evaluation to avoid confining ourselves to a “self-created ivory tower ecology”

—in the context of representative design, the “challenge of further [isolating variables] must be met by after-the-fact, mathematical means”

Brunswik, 1955

Nastase, S. A., Goldstein, A., & Hasson, U. (2020). Keep it real: rethinking the primacy of experimental control in cognitive neuroscience. *PsyArXiv*.



The “direct fit” model of biological learning

Exiting the Cartesian Theater

Biological systems are not obligated to learn human-interpretable *rules*, *representations*, or *maps*—there is no “little scientist” inside the organism following the rules or reading the map.

Ryle, 1949

Gibson, 1979

Dennett, 1991

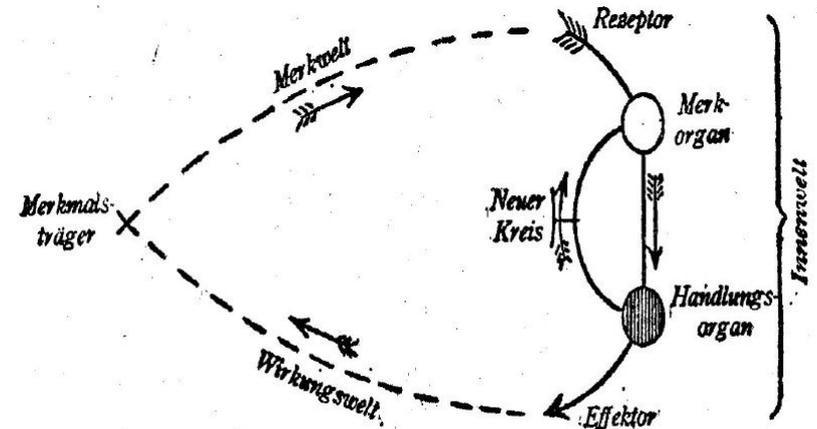


Figure 4.

The “direct fit” model of biological learning

Exiting the Cartesian Theater

Biological systems are not obligated to learn human-interpretable *rules*, *representations*, or *maps*—there is no “little scientist” inside the organism following the rules or reading the map.

The meaningful environment (or *Umwelt*)

The organism cannot be decoupled from its environment—biological (and artificial) neural networks rely on local interpolation to capitalize on “affordances” of the environment.

von Uexküll, 1967

Gibson, 1979

Gomez-Marín & Ghazanfar, *Neuron*, 2020

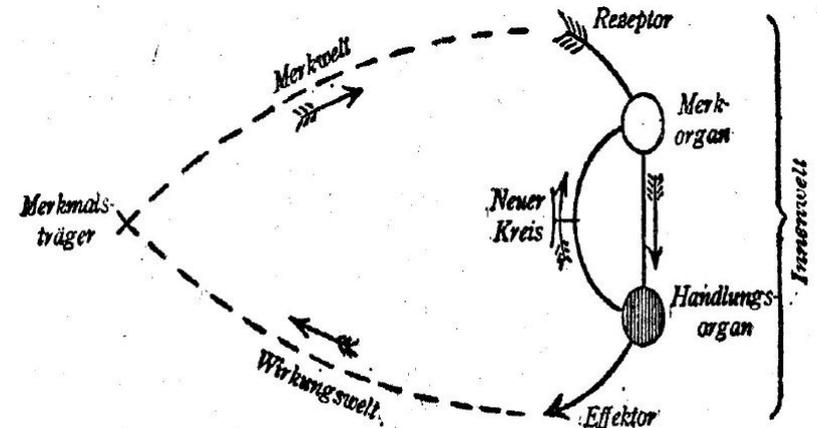


Figure 4.

The “direct fit” model of biological learning

Exiting the Cartesian Theater

Biological systems are not obligated to learn human-interpretable *rules*, *representations*, or *maps*—there is no “little scientist” inside the organism following the rules or reading the map.

The meaningful environment (or *Umwelt*)

The organism cannot be decoupled from its environment—biological (and artificial) neural networks rely on local interpolation to capitalize on “affordances” of the environment.

“There is grandeur in this view of life...”

Both evolution and direct-fit learning models produce solutions that may be mistakenly interpreted in terms of elegant design principles, but in fact reflect the interdigitation of “mindless” optimization processes and the structure of the world.

Darwin, 1859
Dawkins, 1986
Dennett, 1995, 2017

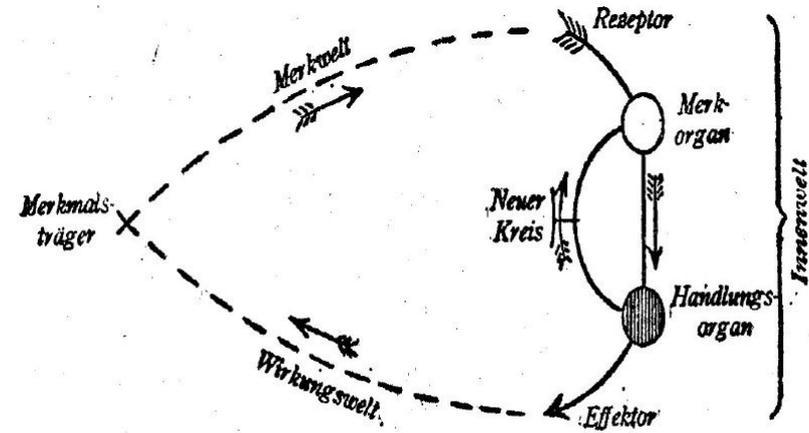


Figure 4.

Thanks!

Uri Hasson

Ariel Goldstein

Kenneth A. Norman

Yun-Fei Liu

Hanna Hillman

Asieh Zadbood

Liat Hasenfratz

Negin Keshavarzian

Janice Chen

Christopher J. Honey

Mai Nguyen

Claire H. C. Chang

Christopher Baldassano

James V. Haxby

M. Ida Gobbini

Yaroslav O. Halchenko

Feilong Ma

J. Swaroop Guntupalli

Andrew C. Connolly



National Institute
of Mental Health

The “direct fit” model of biological learning

Exiting the Cartesian Theater

Biological systems are not obligated to learn human-interpretable *rules*, *representations*, or *maps*—there is no “little scientist” inside the organism following the rules or reading the map.

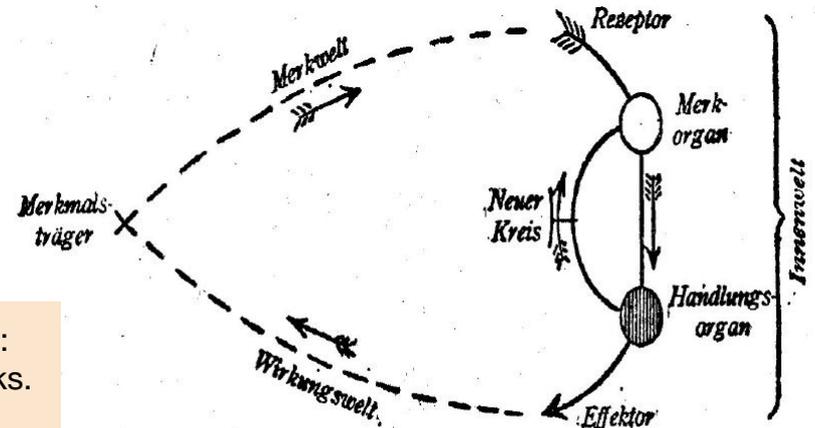
The meaningful environment (or *Umwelt*)

The organism cannot be decoupled from its environment—biological (and artificial) neural networks rely on local interpolation to capitalize on “affordances” of the environment.

“There is grandeur in this view of life...”

Both evolution and direct-fit learning models produce solutions that may be mistakenly interpreted in terms of elegant design principles, but in fact reflect the interdigitation of “mindless” optimization processes and the structure of the world.

Hasson, U., Nastase, S. A., & Goldstein, A. (2020). Direct fit to nature: an evolutionary perspective on biological and artificial neural networks. *Neuron*, 105(3), 416–434.



Figur 4.

von Uexküll, 1920