Reinforcement Learning for Generating (hopefully) Useful Combinatorial Data

Karan Srivastava | Department of Mathematics at the University of Wisconsin-Madison

Research supported in part by NSF Award DMS-2023239

Under supervision of Jordan Ellenberg | University of Wisconsin-Madison

Collaboration with Adam Z. Wagner | Tel Aviv University

Some Vocabulary

Machine Learning:

- Computational process of (gradually) learning a task from given information

Some Vocabulary

Machine Learning:

- Computational process of (gradually) learning a task from given information

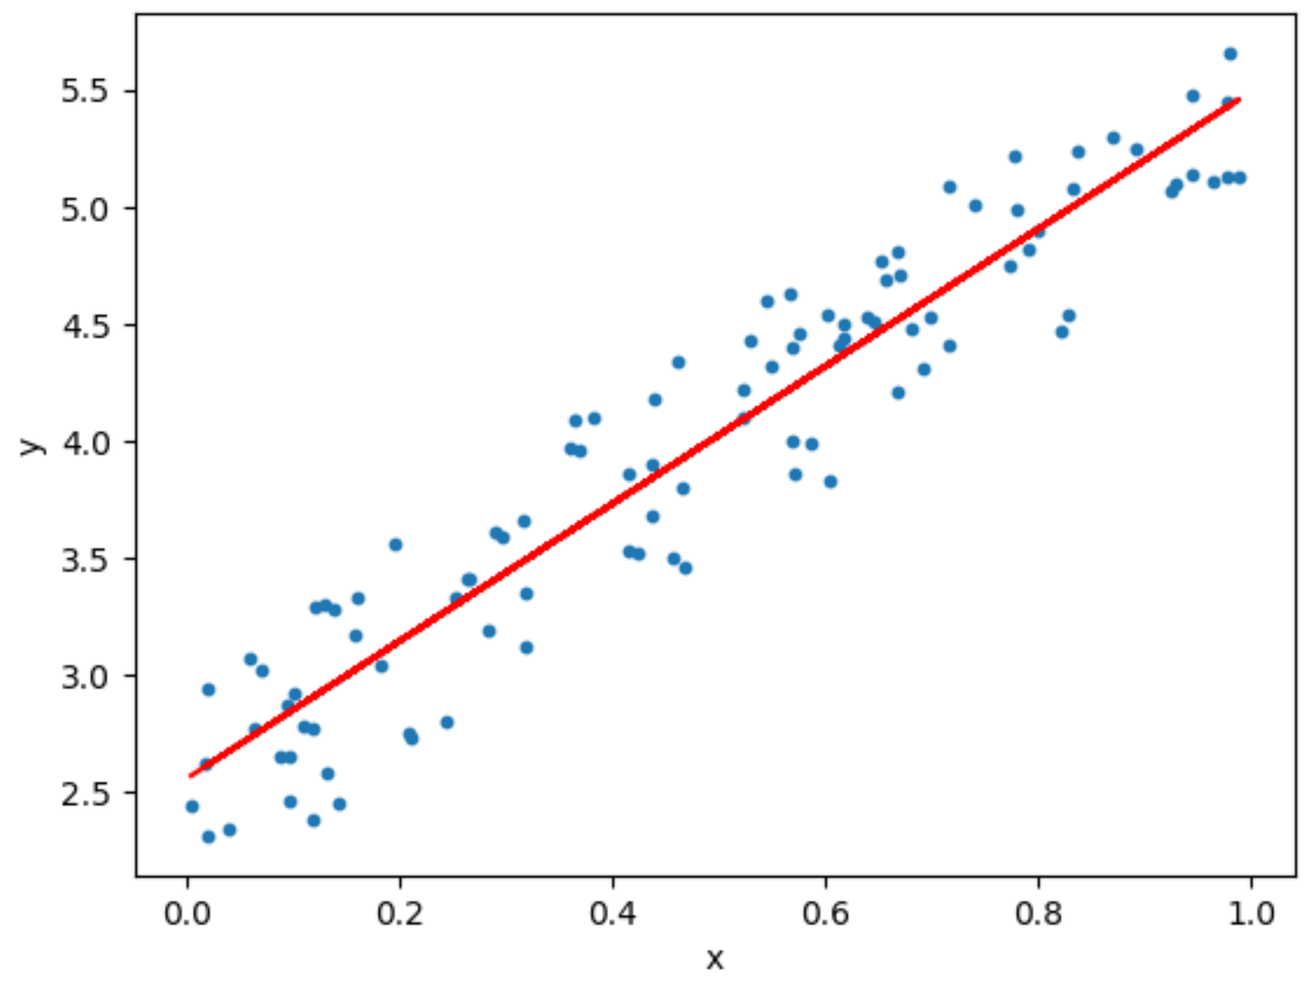

Some Vocabulary

Machine Learning:

- Computational process of (gradually) learning a task from given information

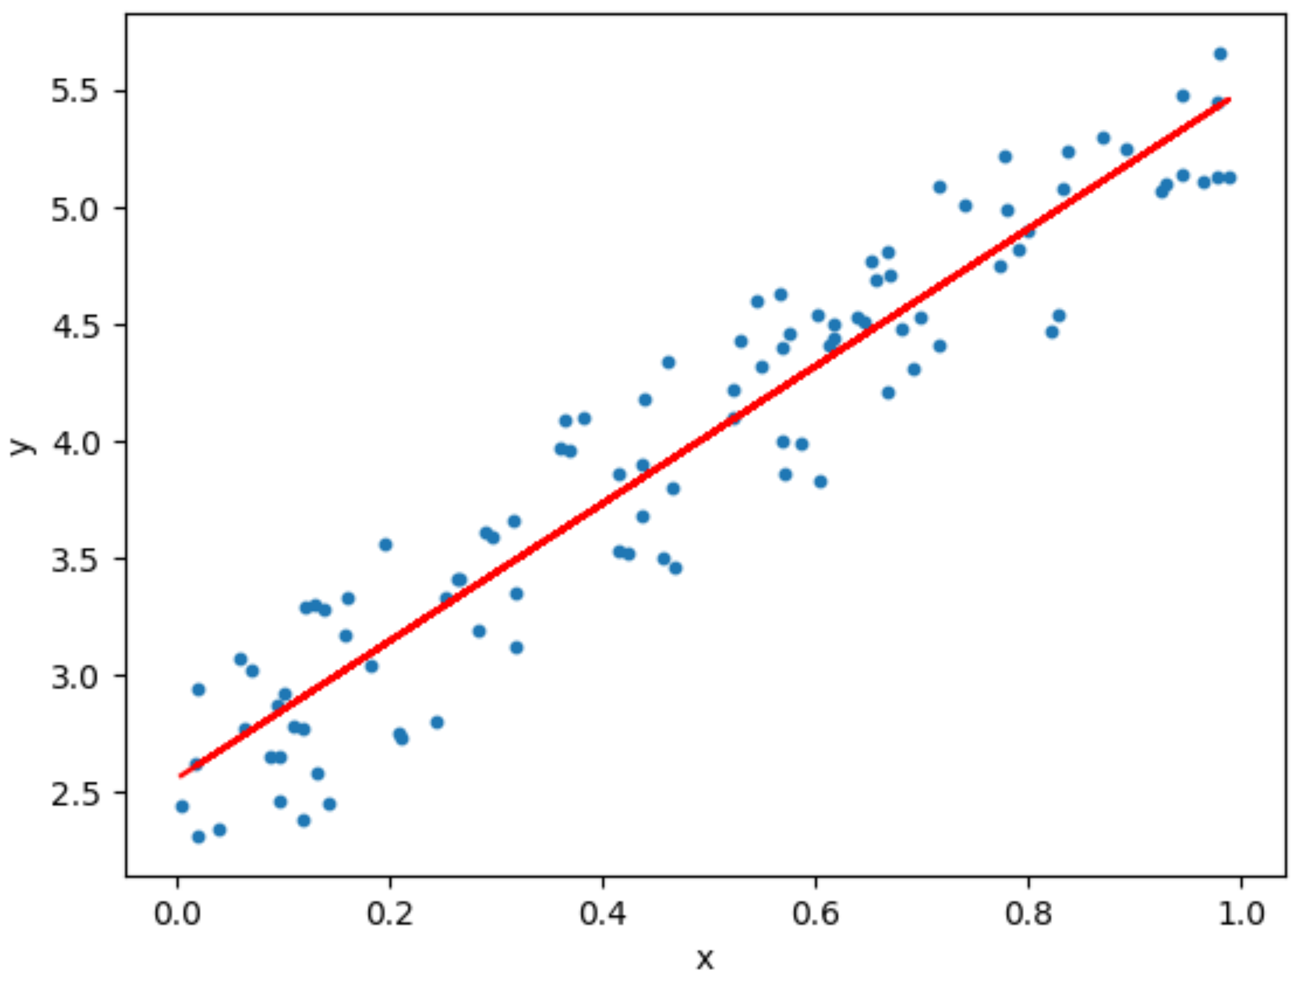

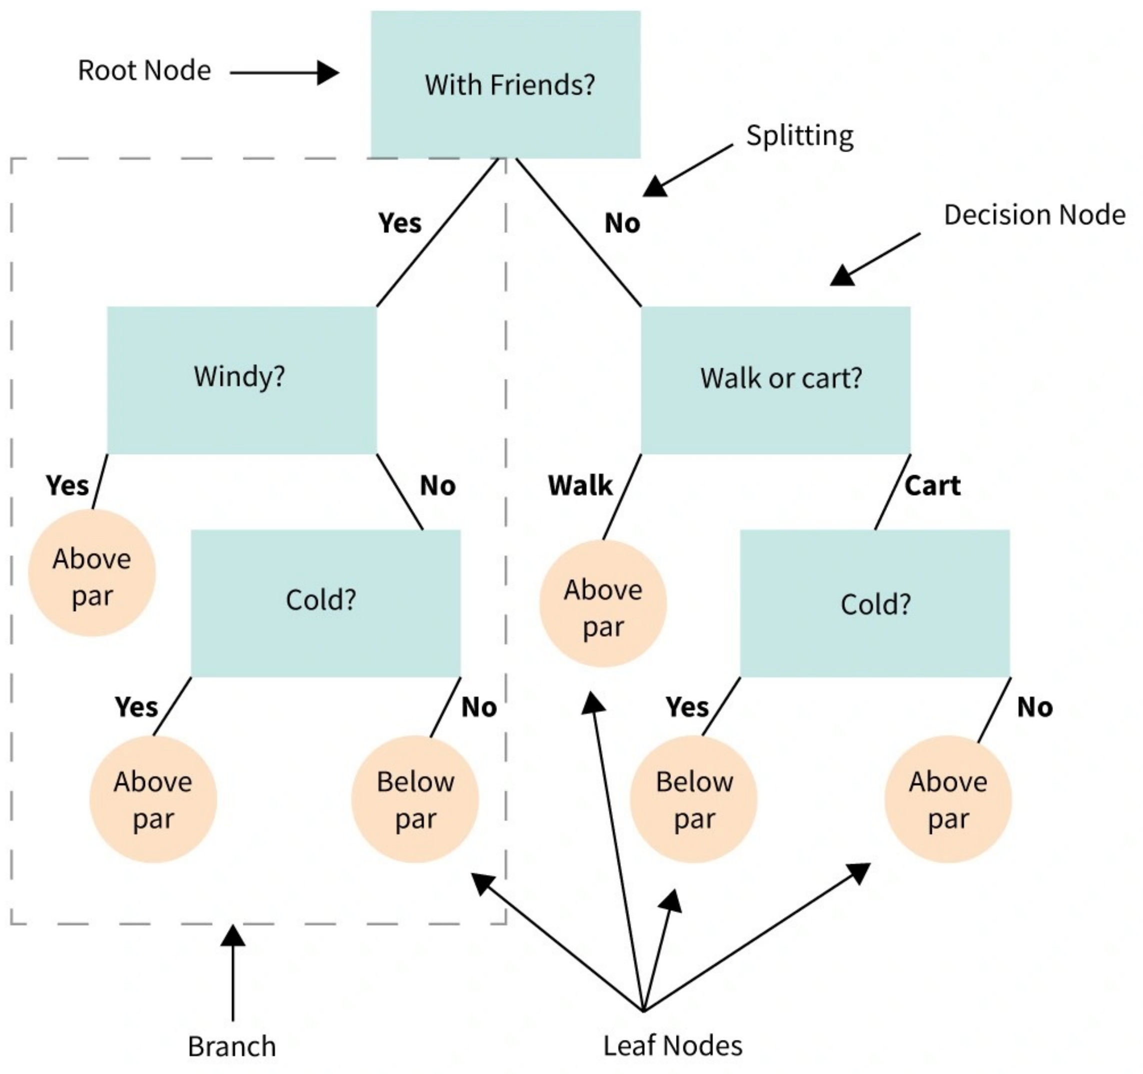

Some Vocabulary

Machine Learning:

- Computational process of (gradually) learning a task from given information

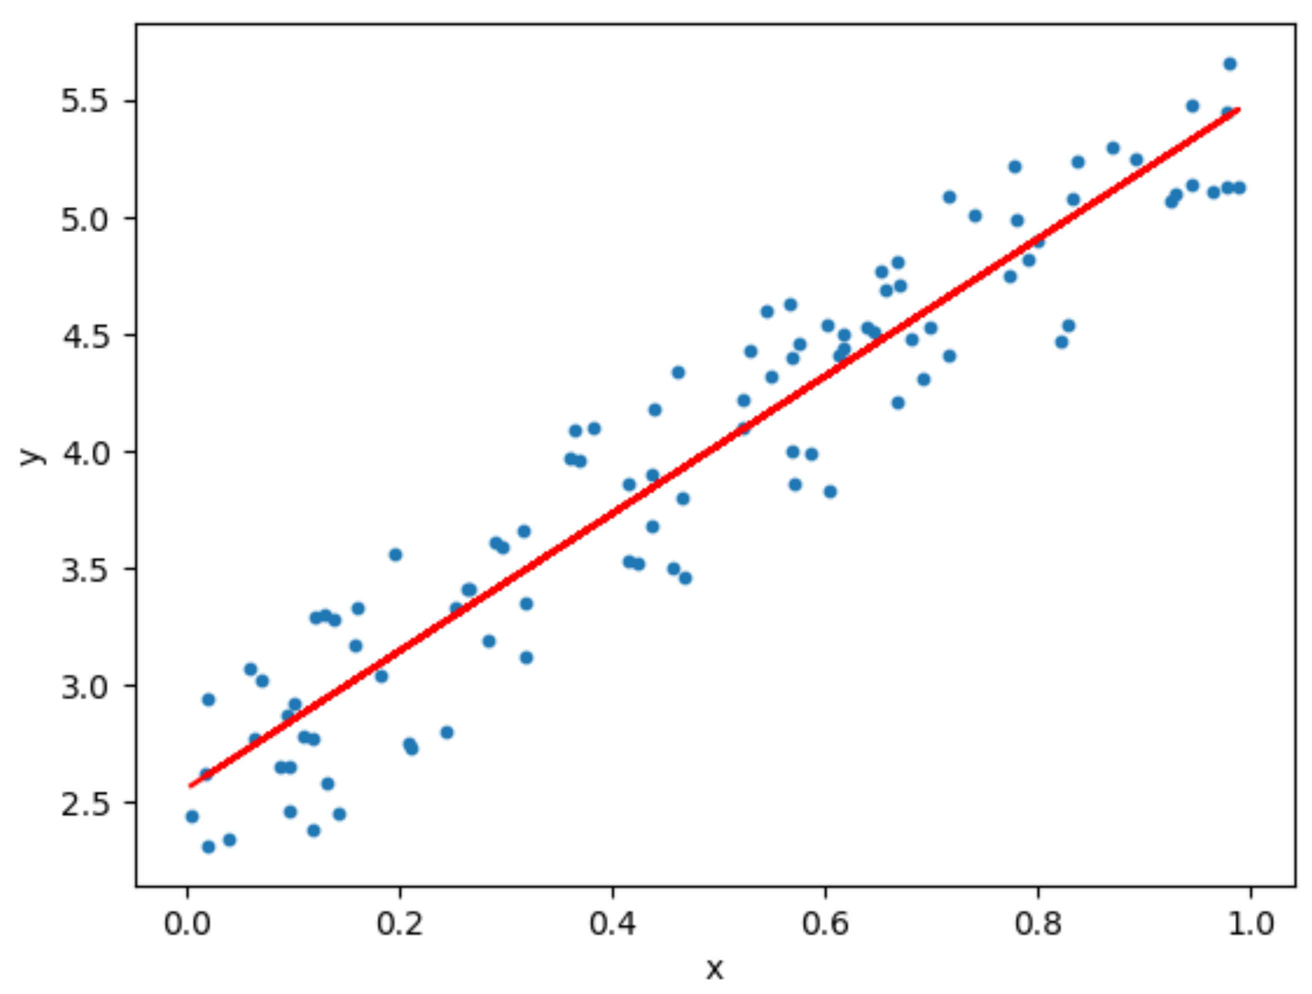

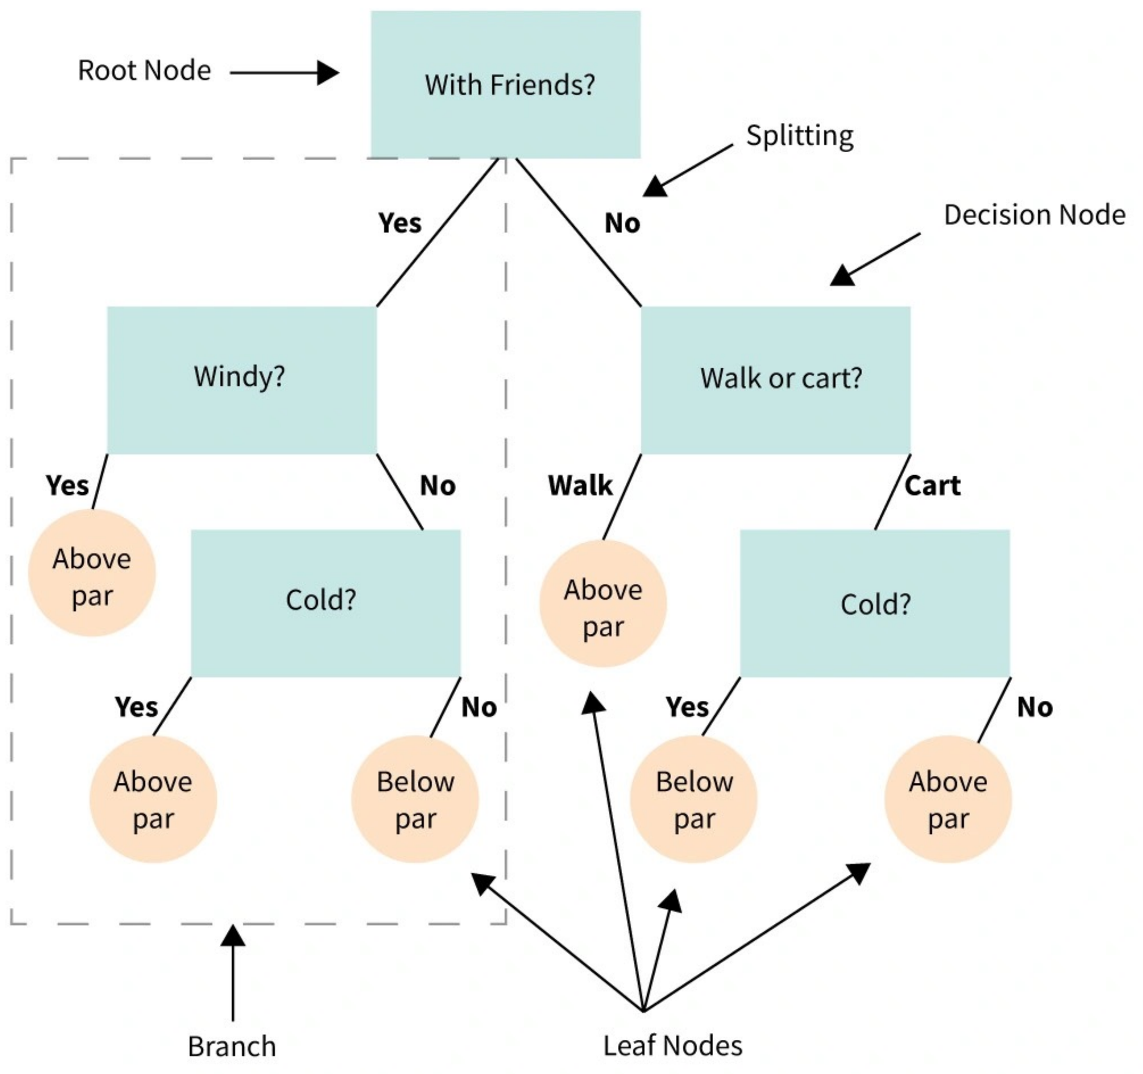

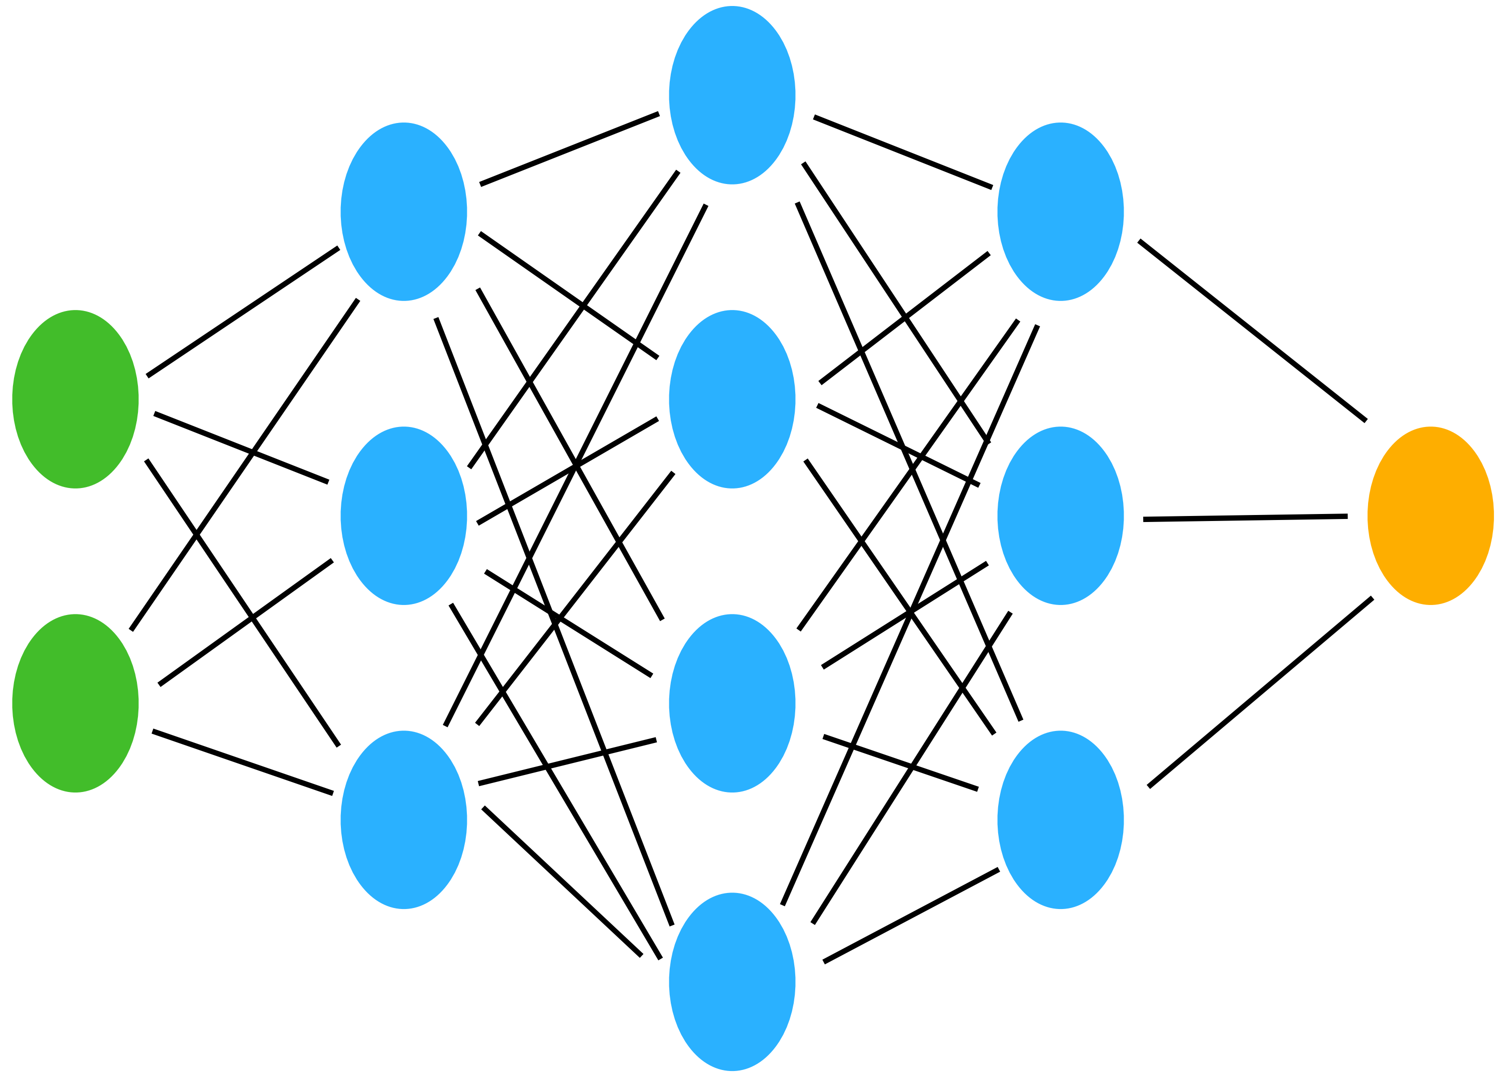

Some Vocabulary

Machine Learning:

- Computational process of (gradually) learning a task from given information

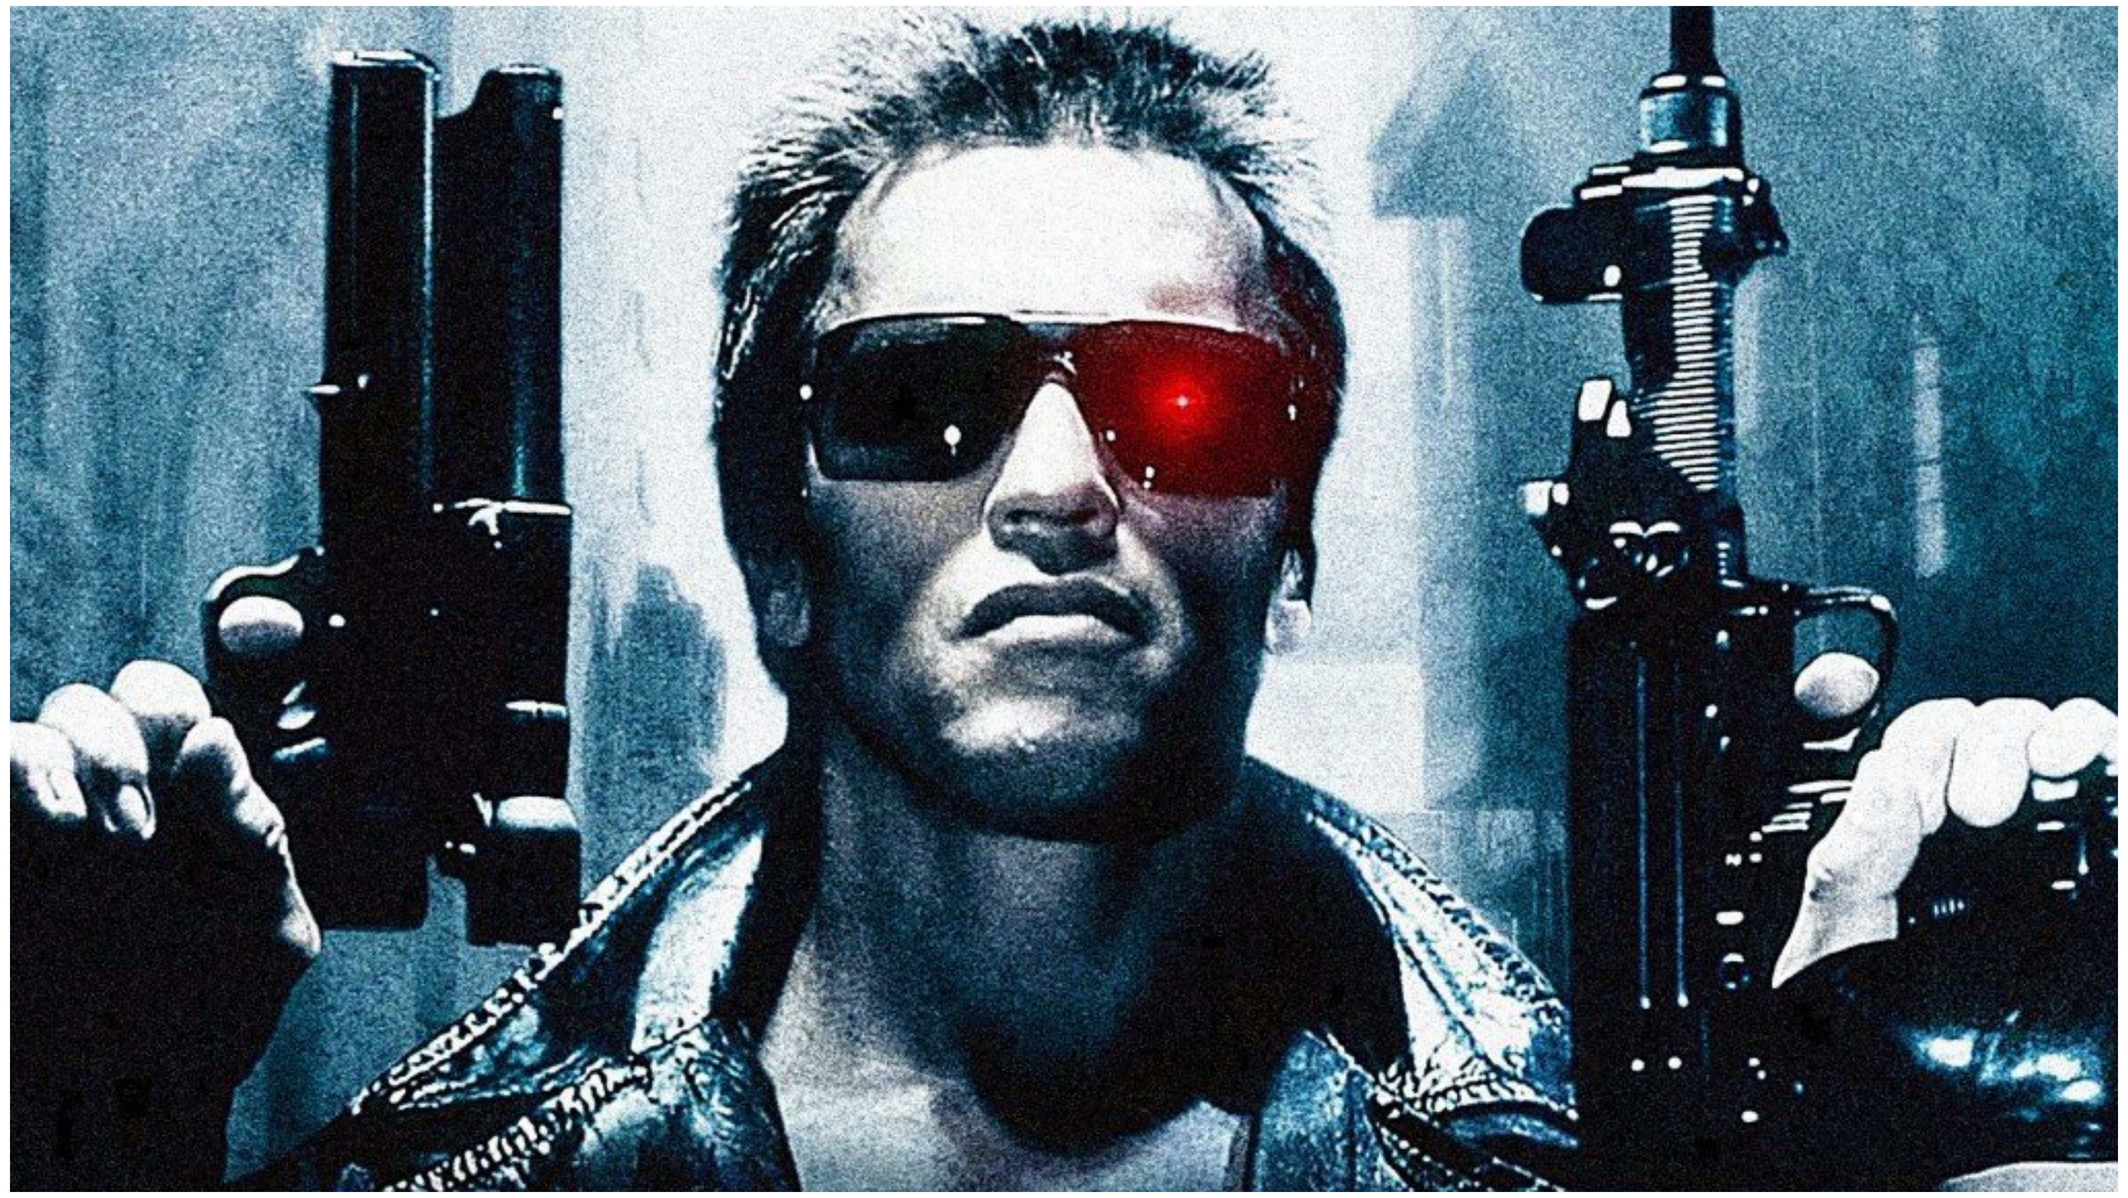

Some Vocabulary

Machine Learning:

- Computational process of (gradually) learning a task from given information

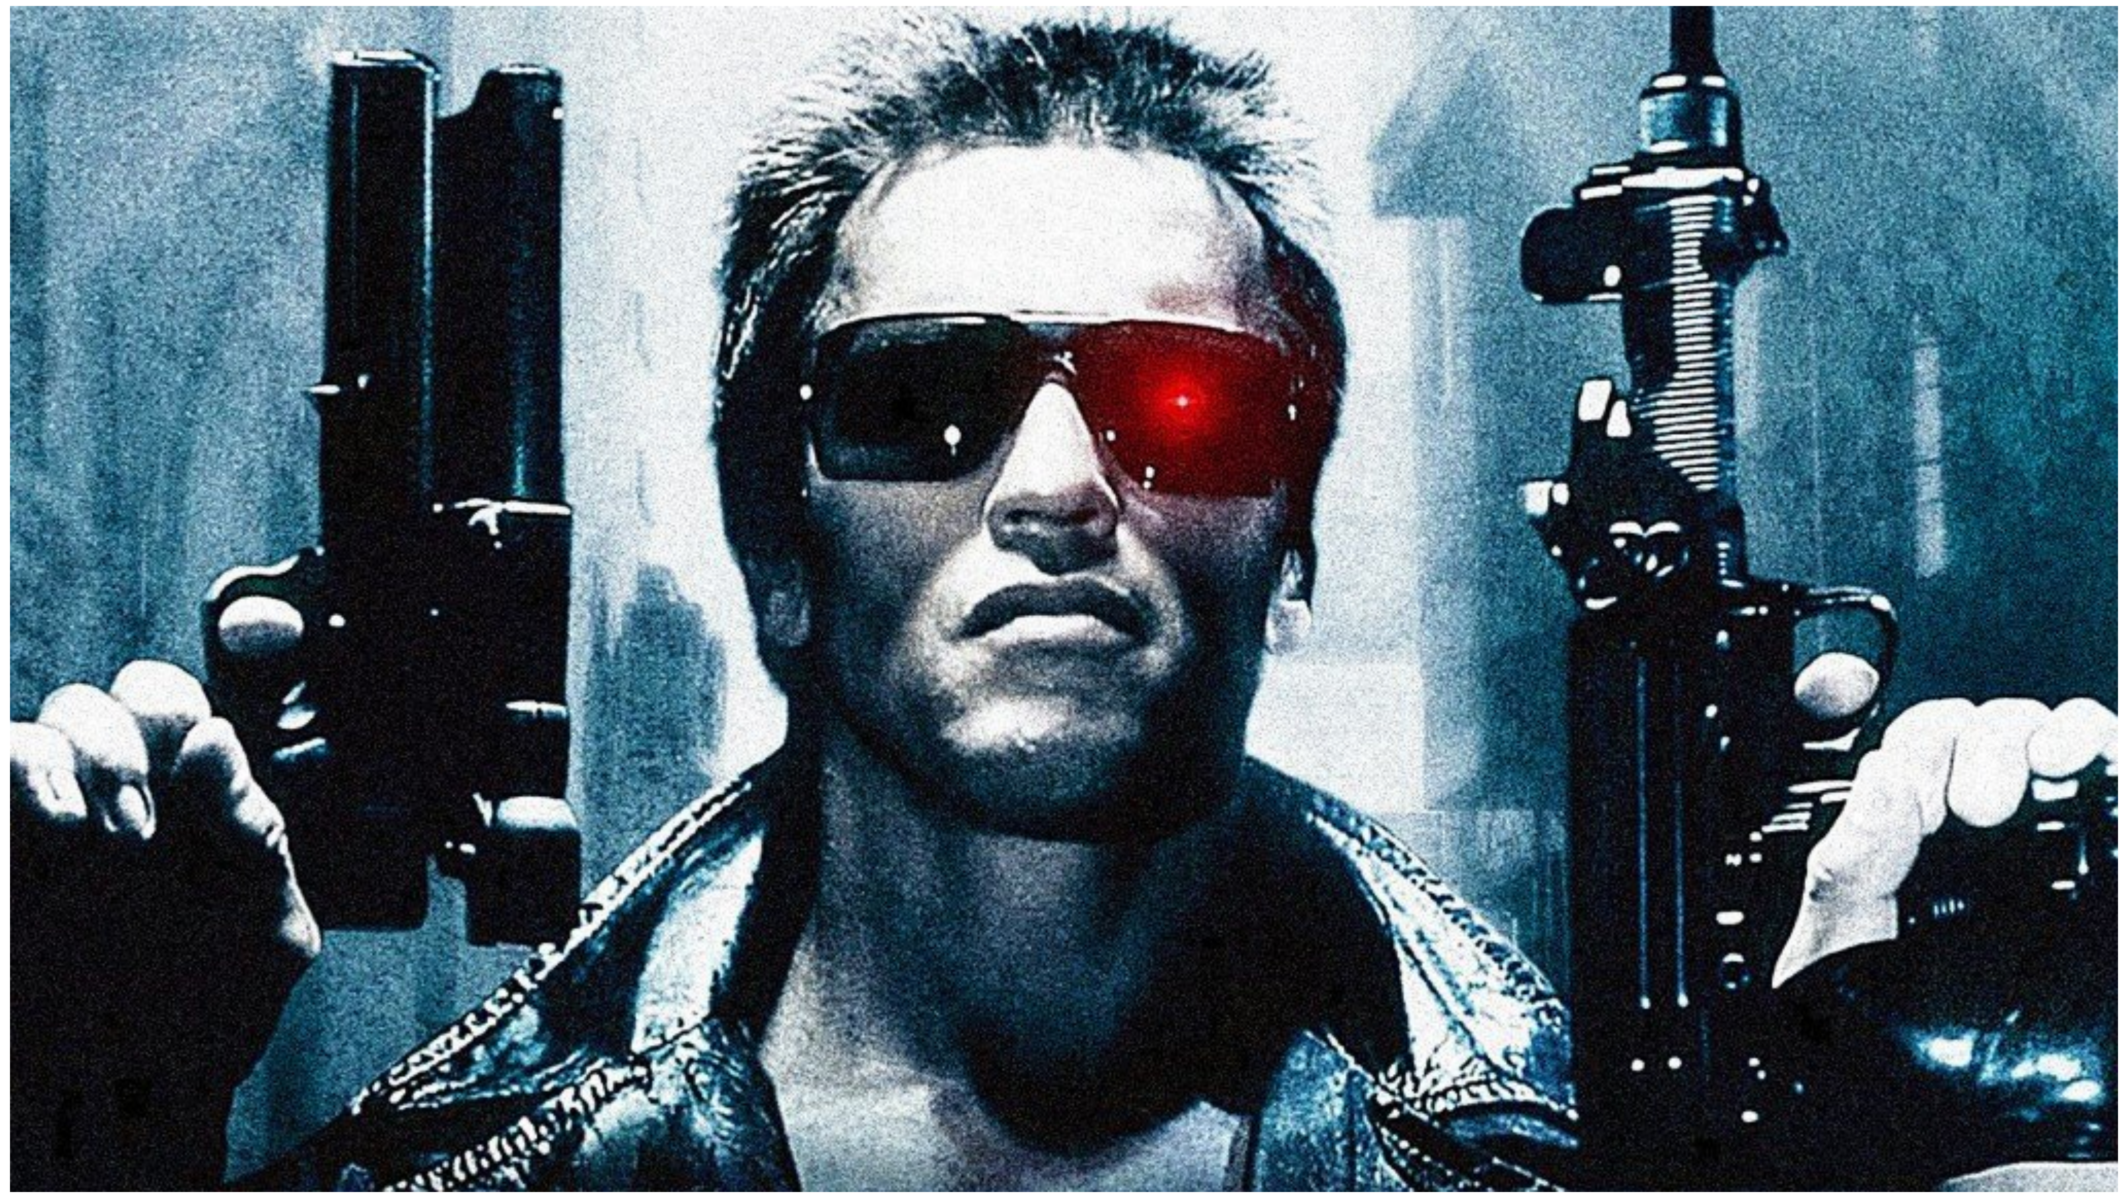

≈

Some Vocabulary

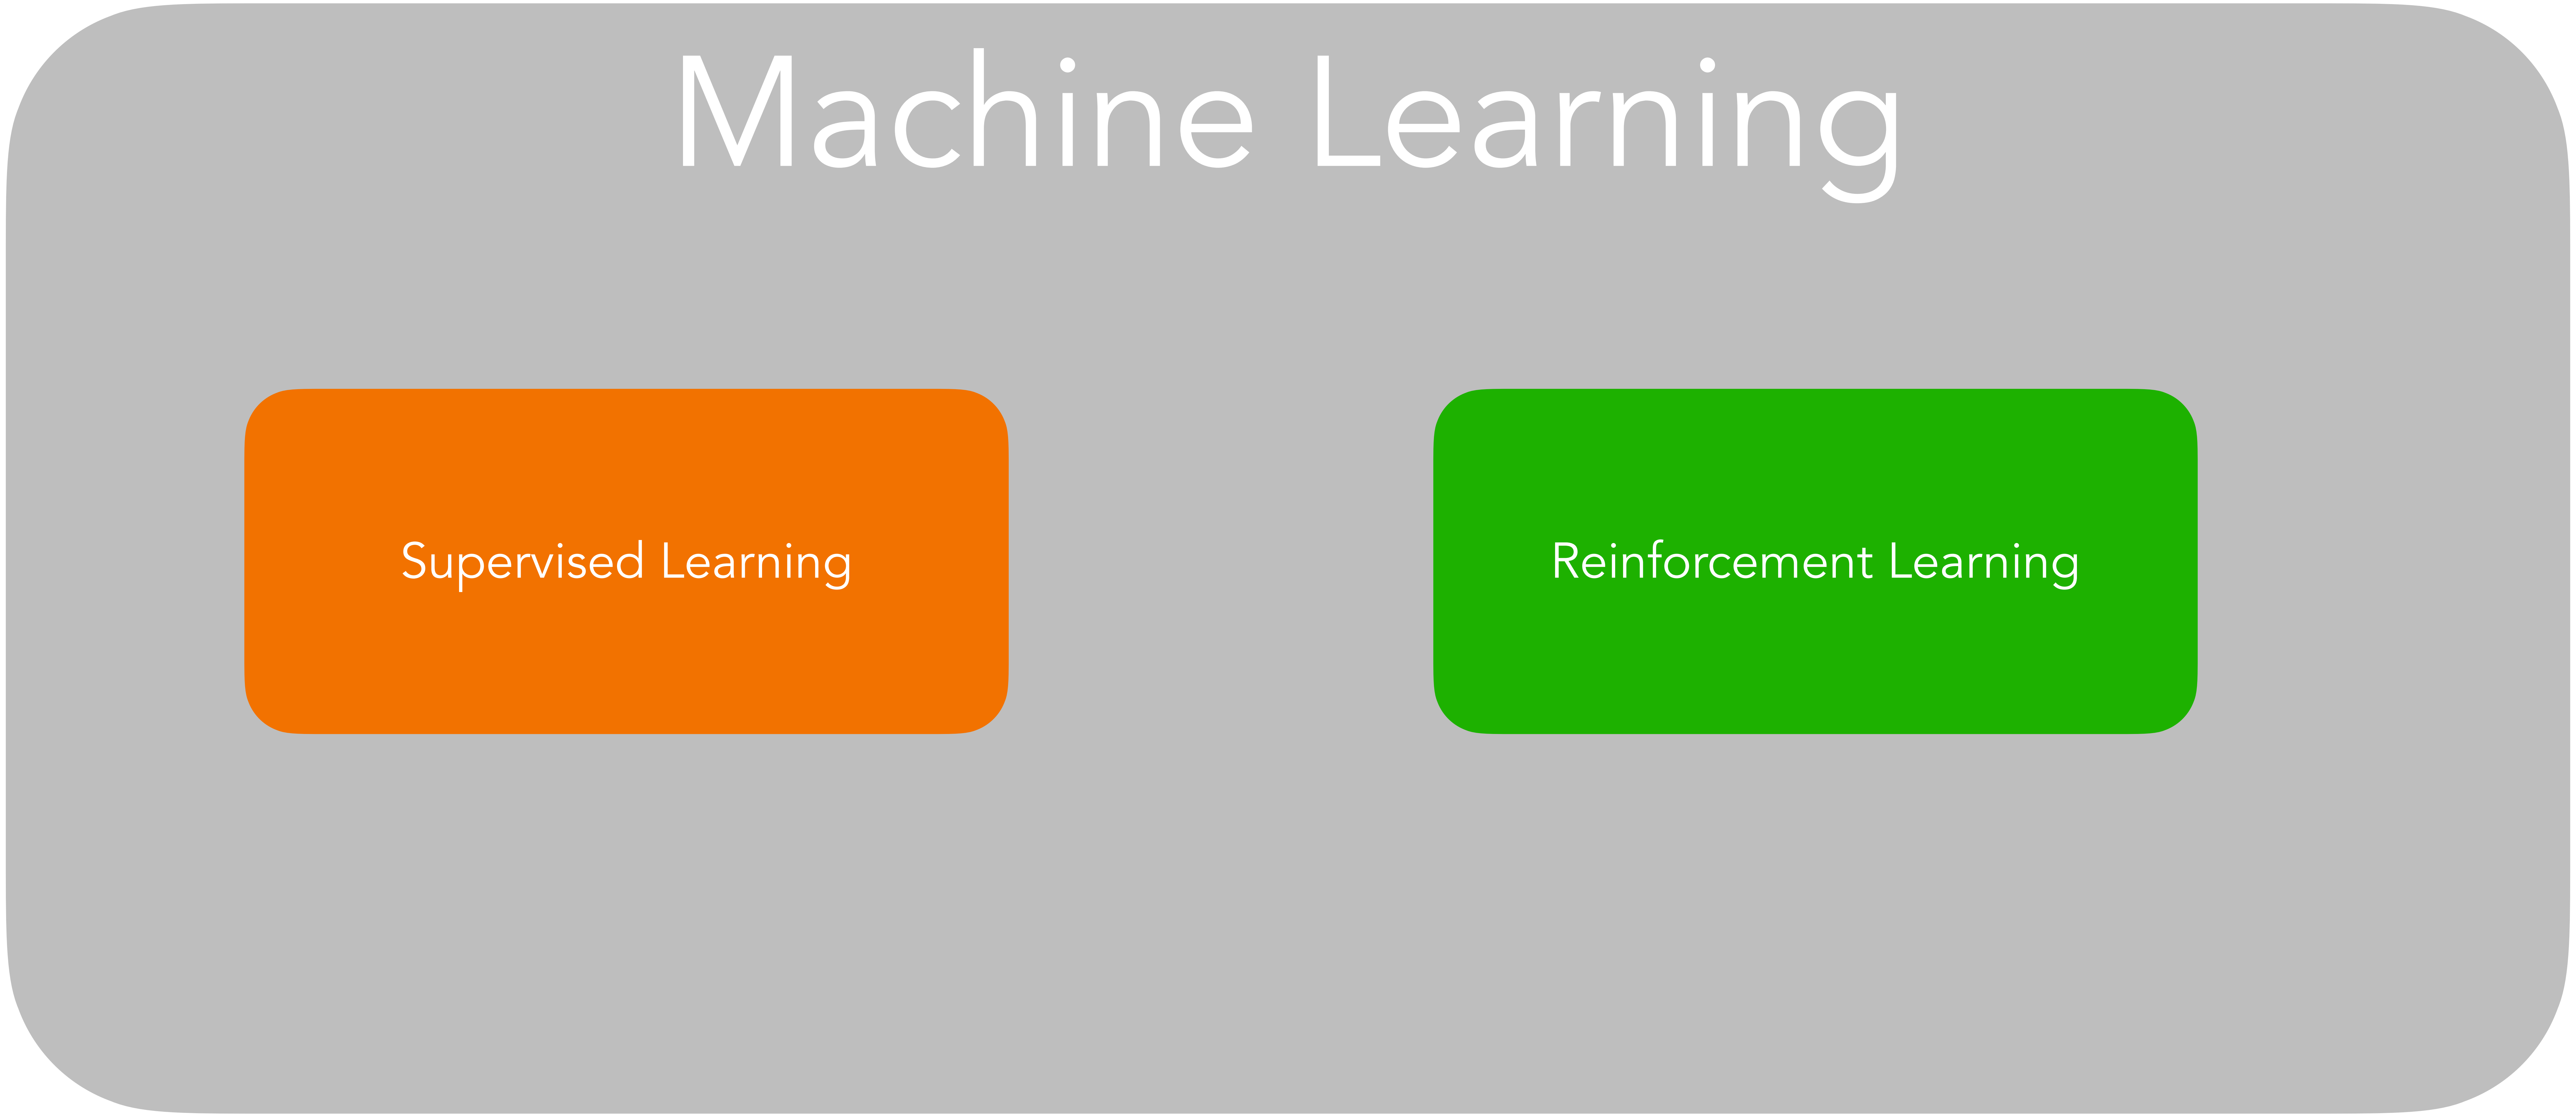

Some Vocabulary

Supervised Learning: Learning with Sampled Data

Some Vocabulary

Supervised Learning: Learning with Sampled Data

Known Labeled Data
$(\mathcal{X}, Y)$

Some Vocabulary

Supervised Learning: Learning with Sampled Data

Known Labeled Data

$(\mathscr{X}, Y)$

Some Vocabulary

Supervised Learning: Learning with Sampled Data

Known Labeled Data

$(\mathscr{X}, Y)$

Some Vocabulary

Supervised Learning: Learning with Sampled Data

Known Labeled Data
$(\mathscr{X}, Y)$

Unseen Data
$\hat{f}(\mathscr{X}) \approx Y$

Some Vocabulary

Reinforcement Learning: Learning without Sampled Data

Some Vocabulary

Reinforcement Learning: Learning without Sampled Data

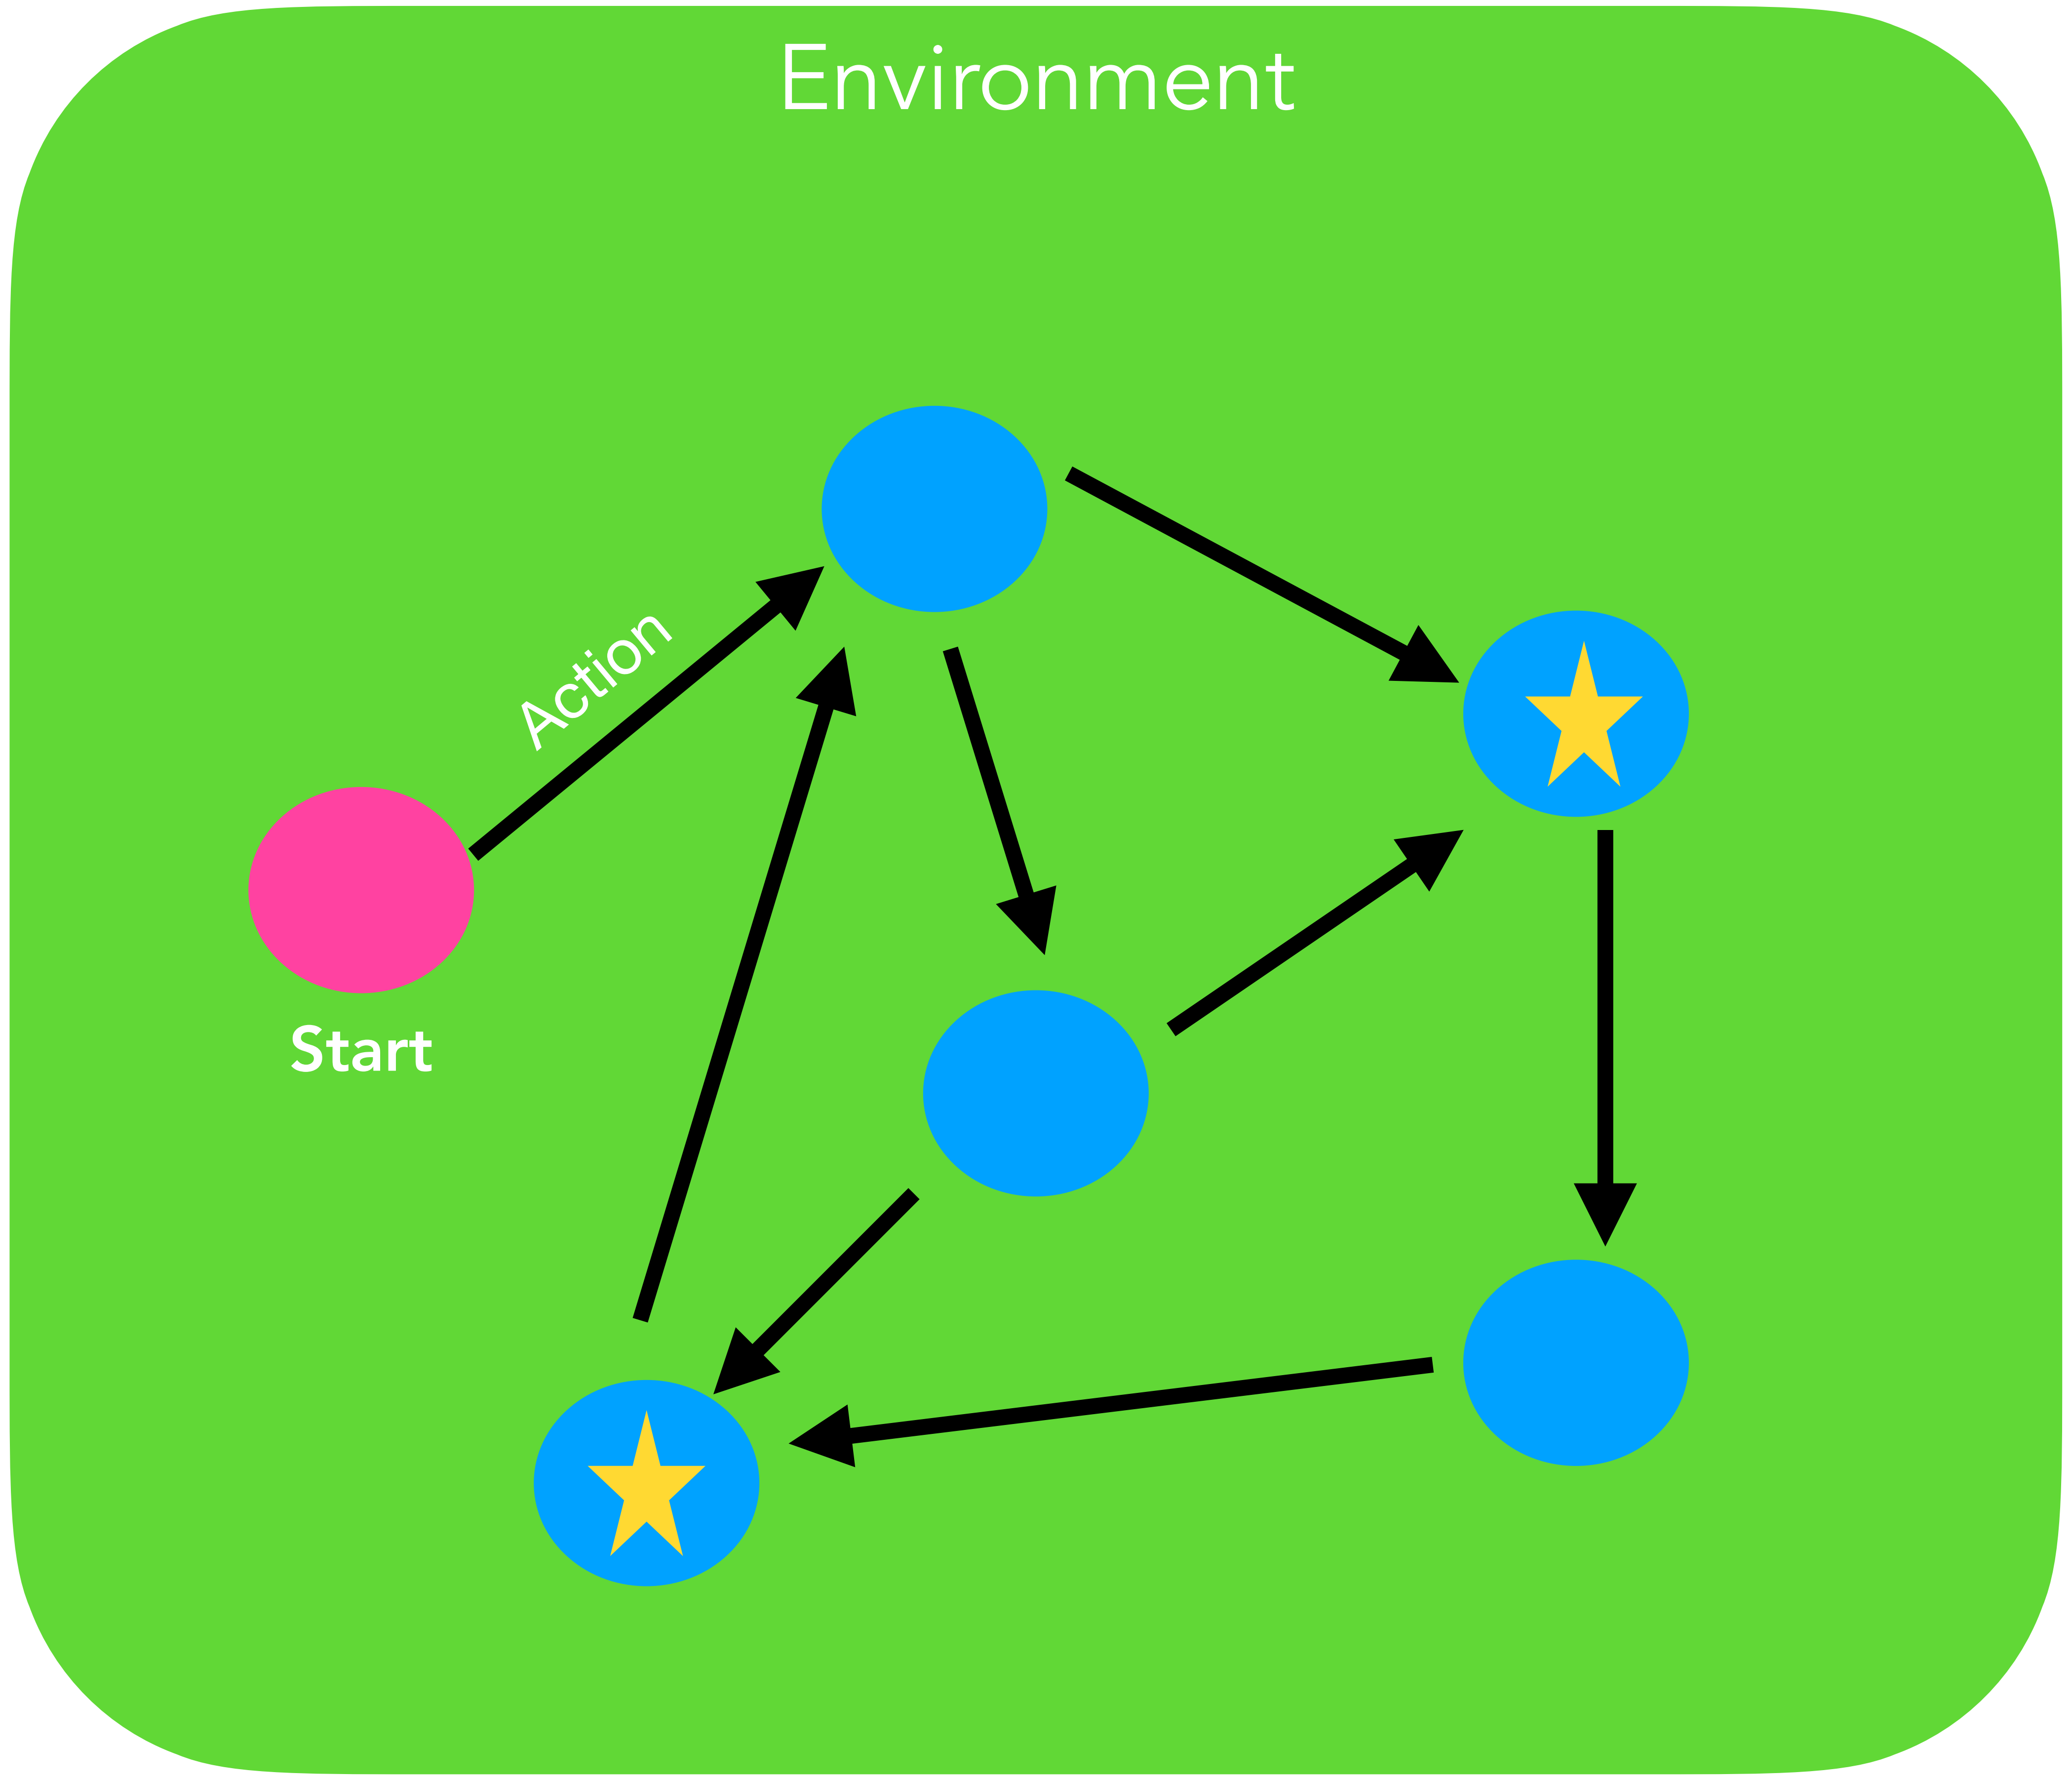

Some Vocabulary

Reinforcement Learning: Learning without Sampled Data

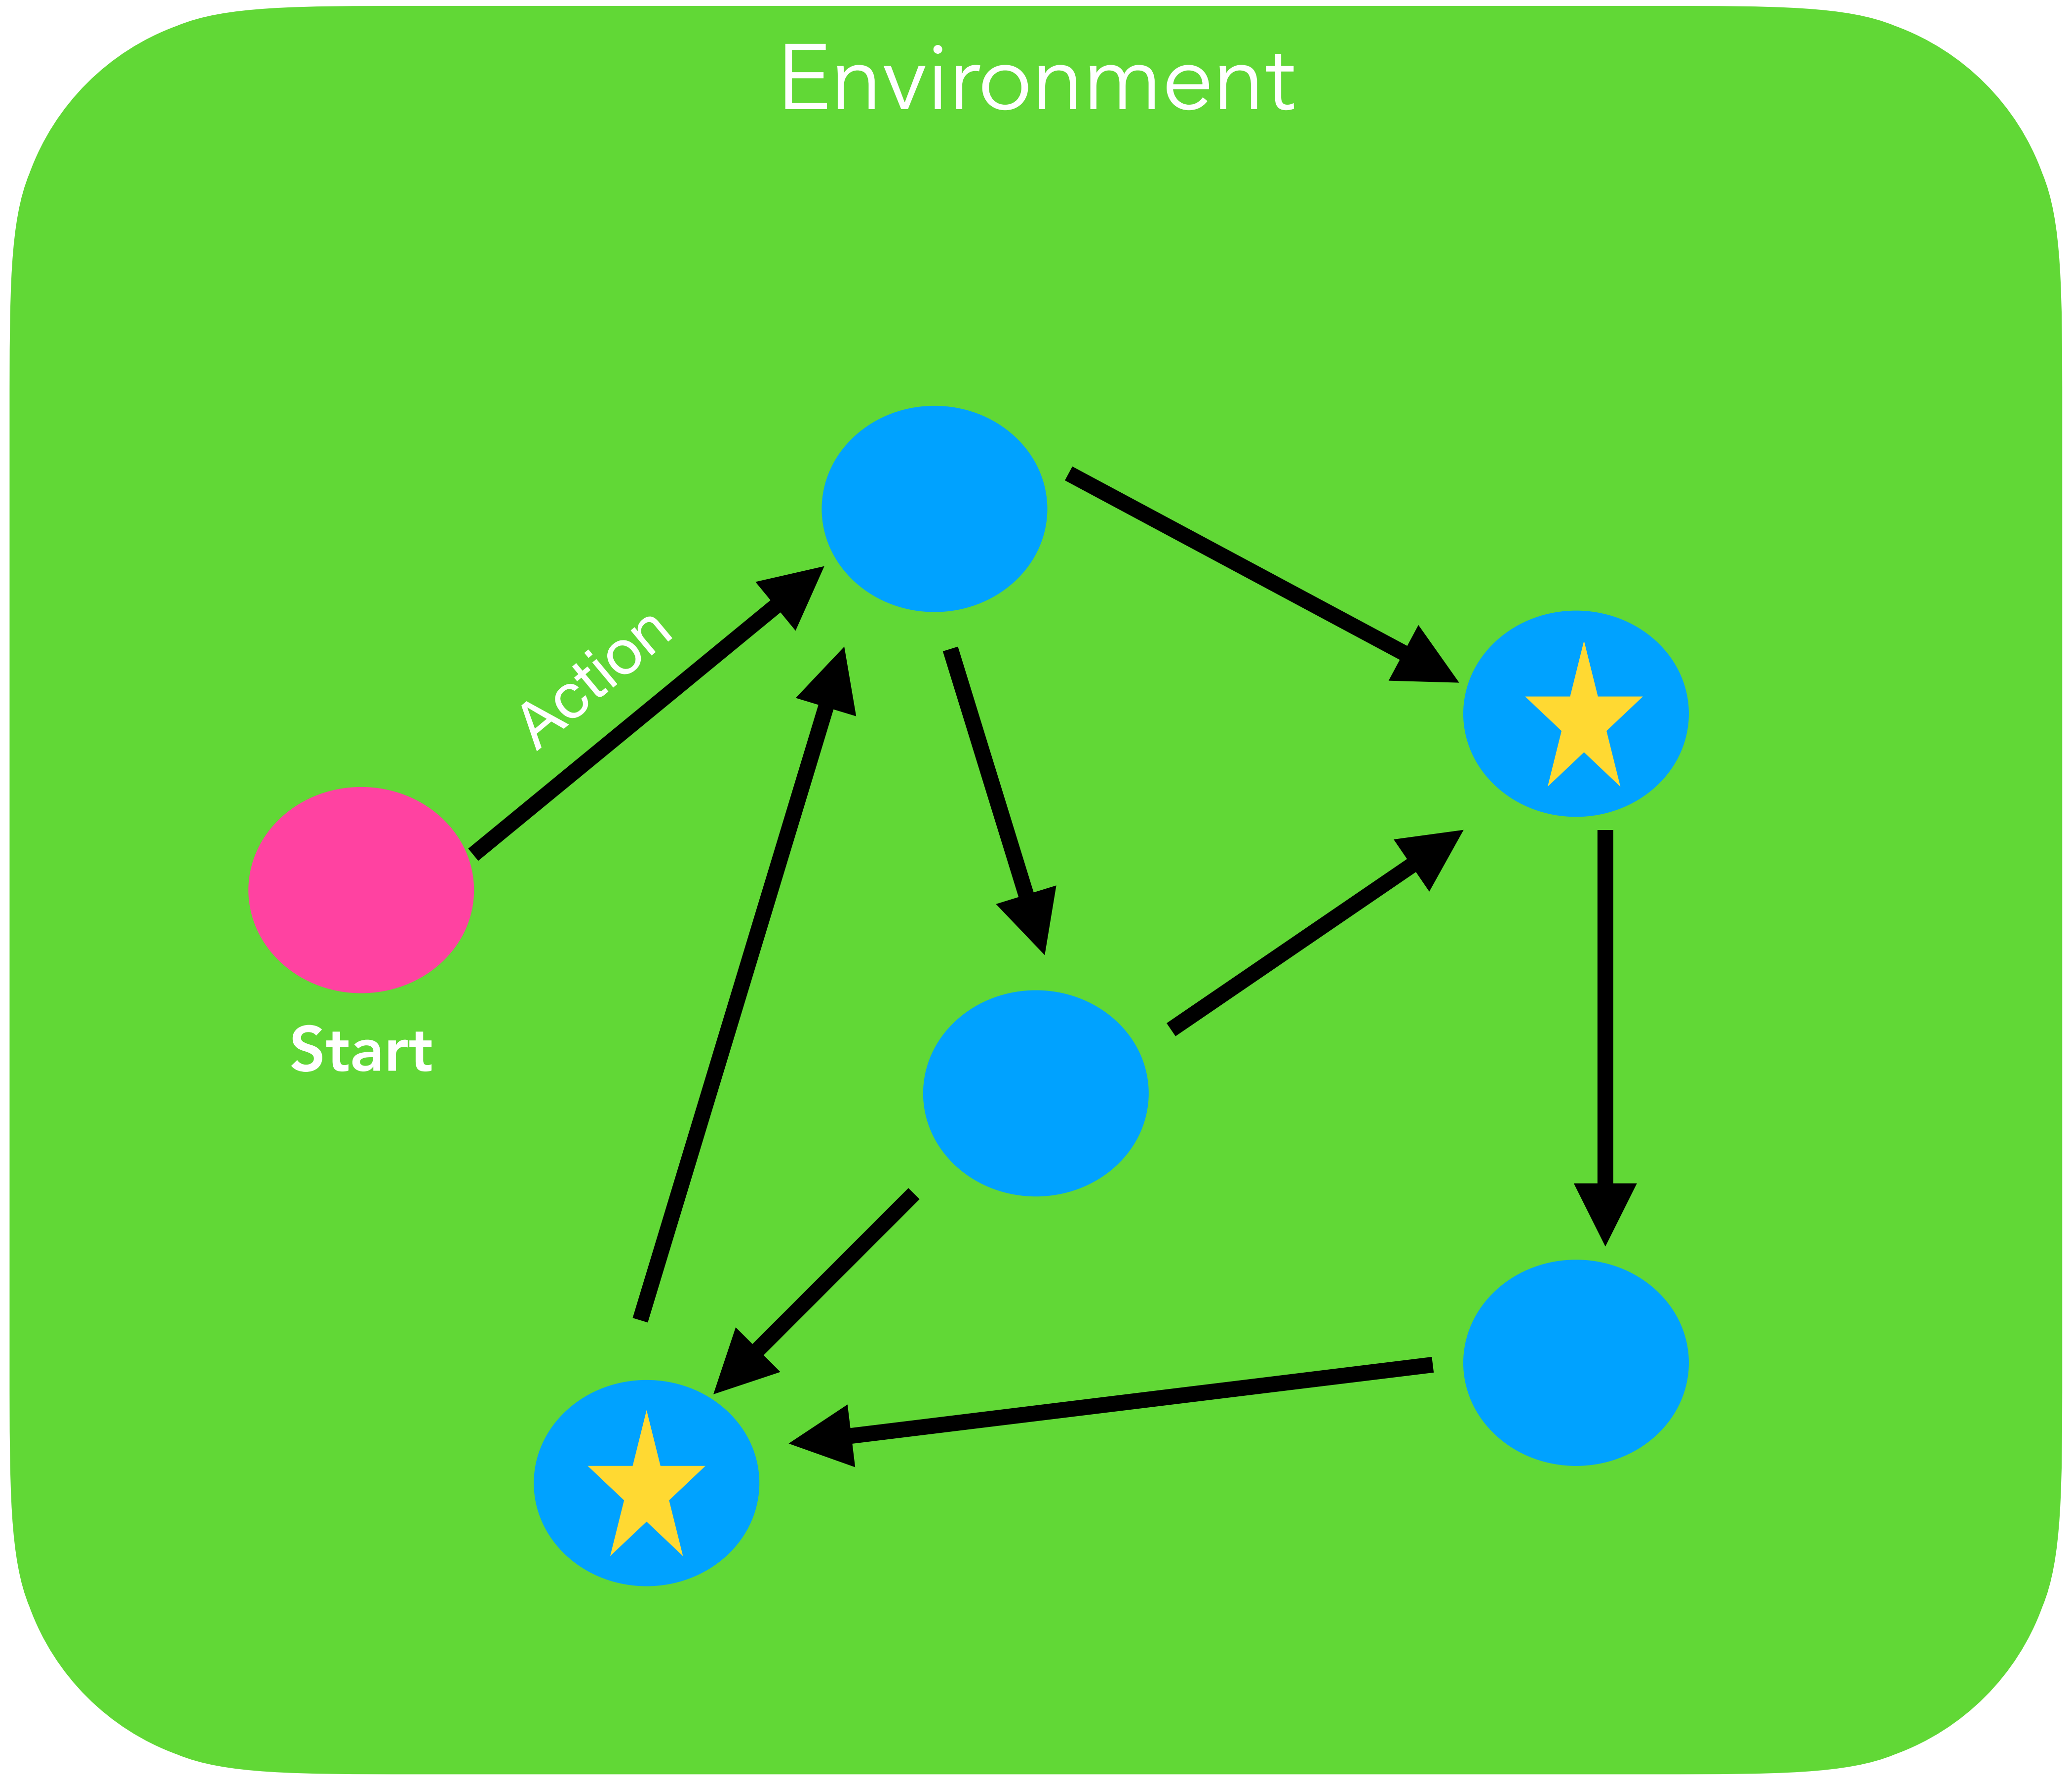

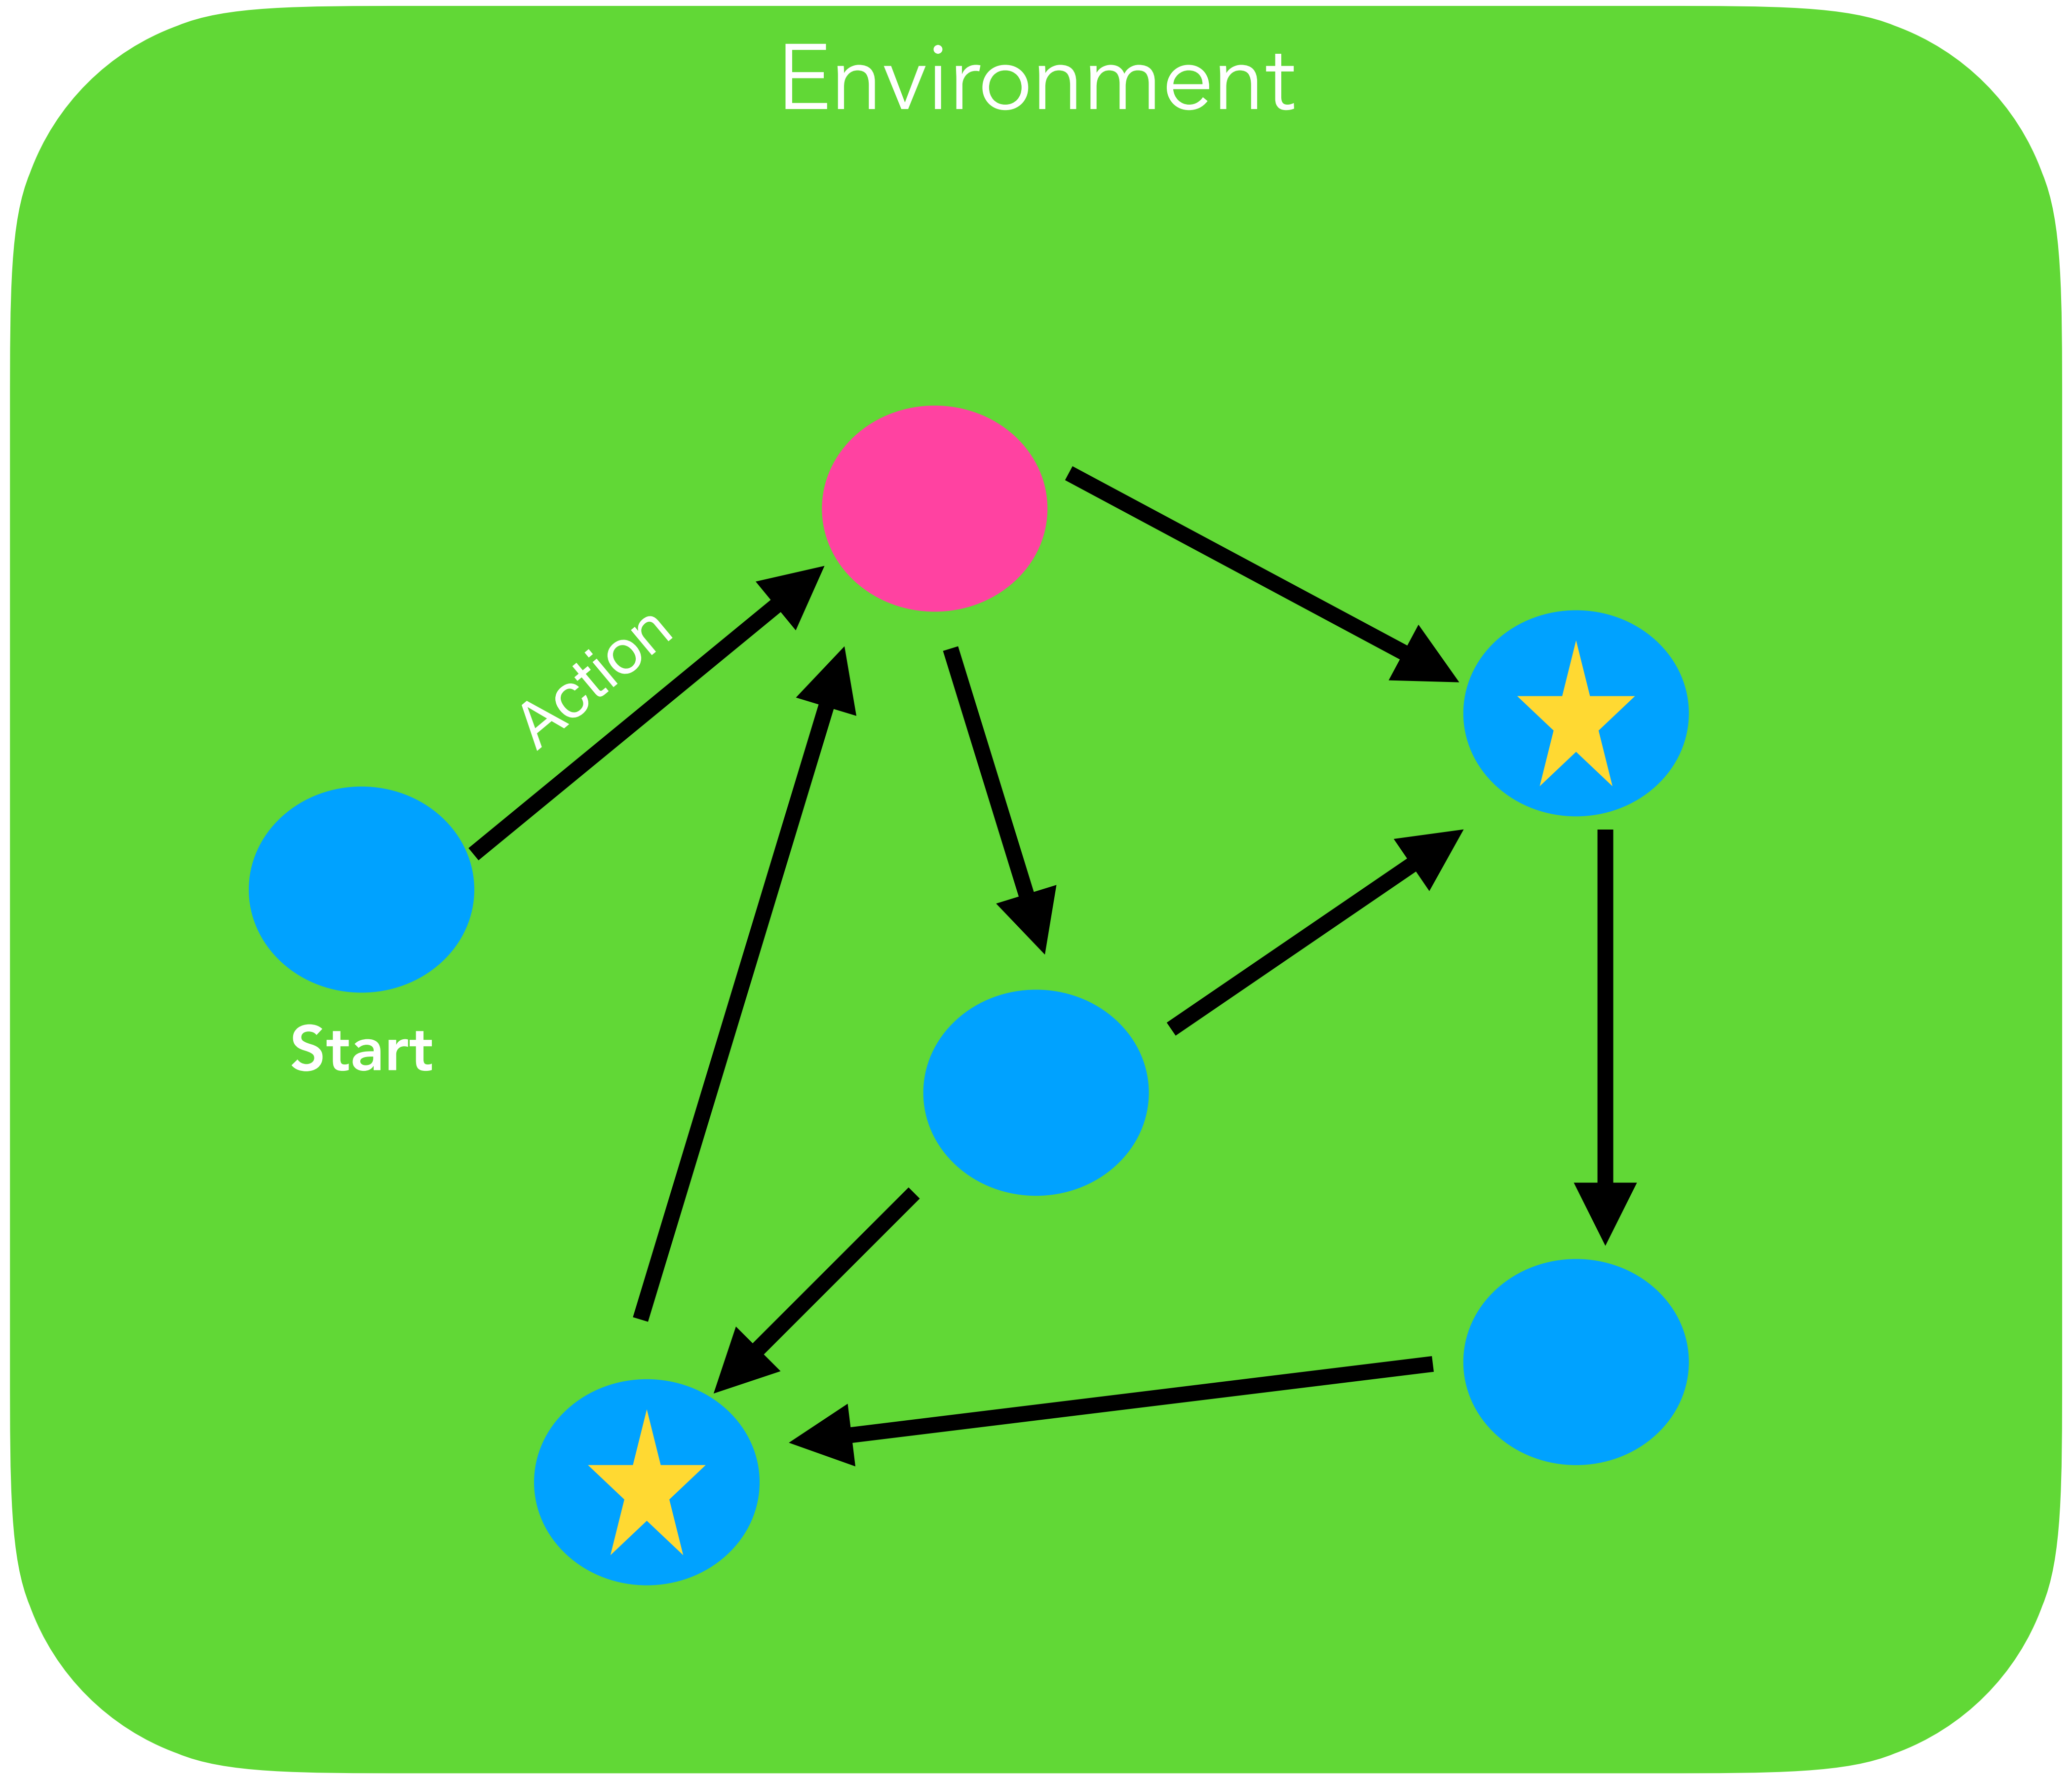

Some Vocabulary

Neural Networks:

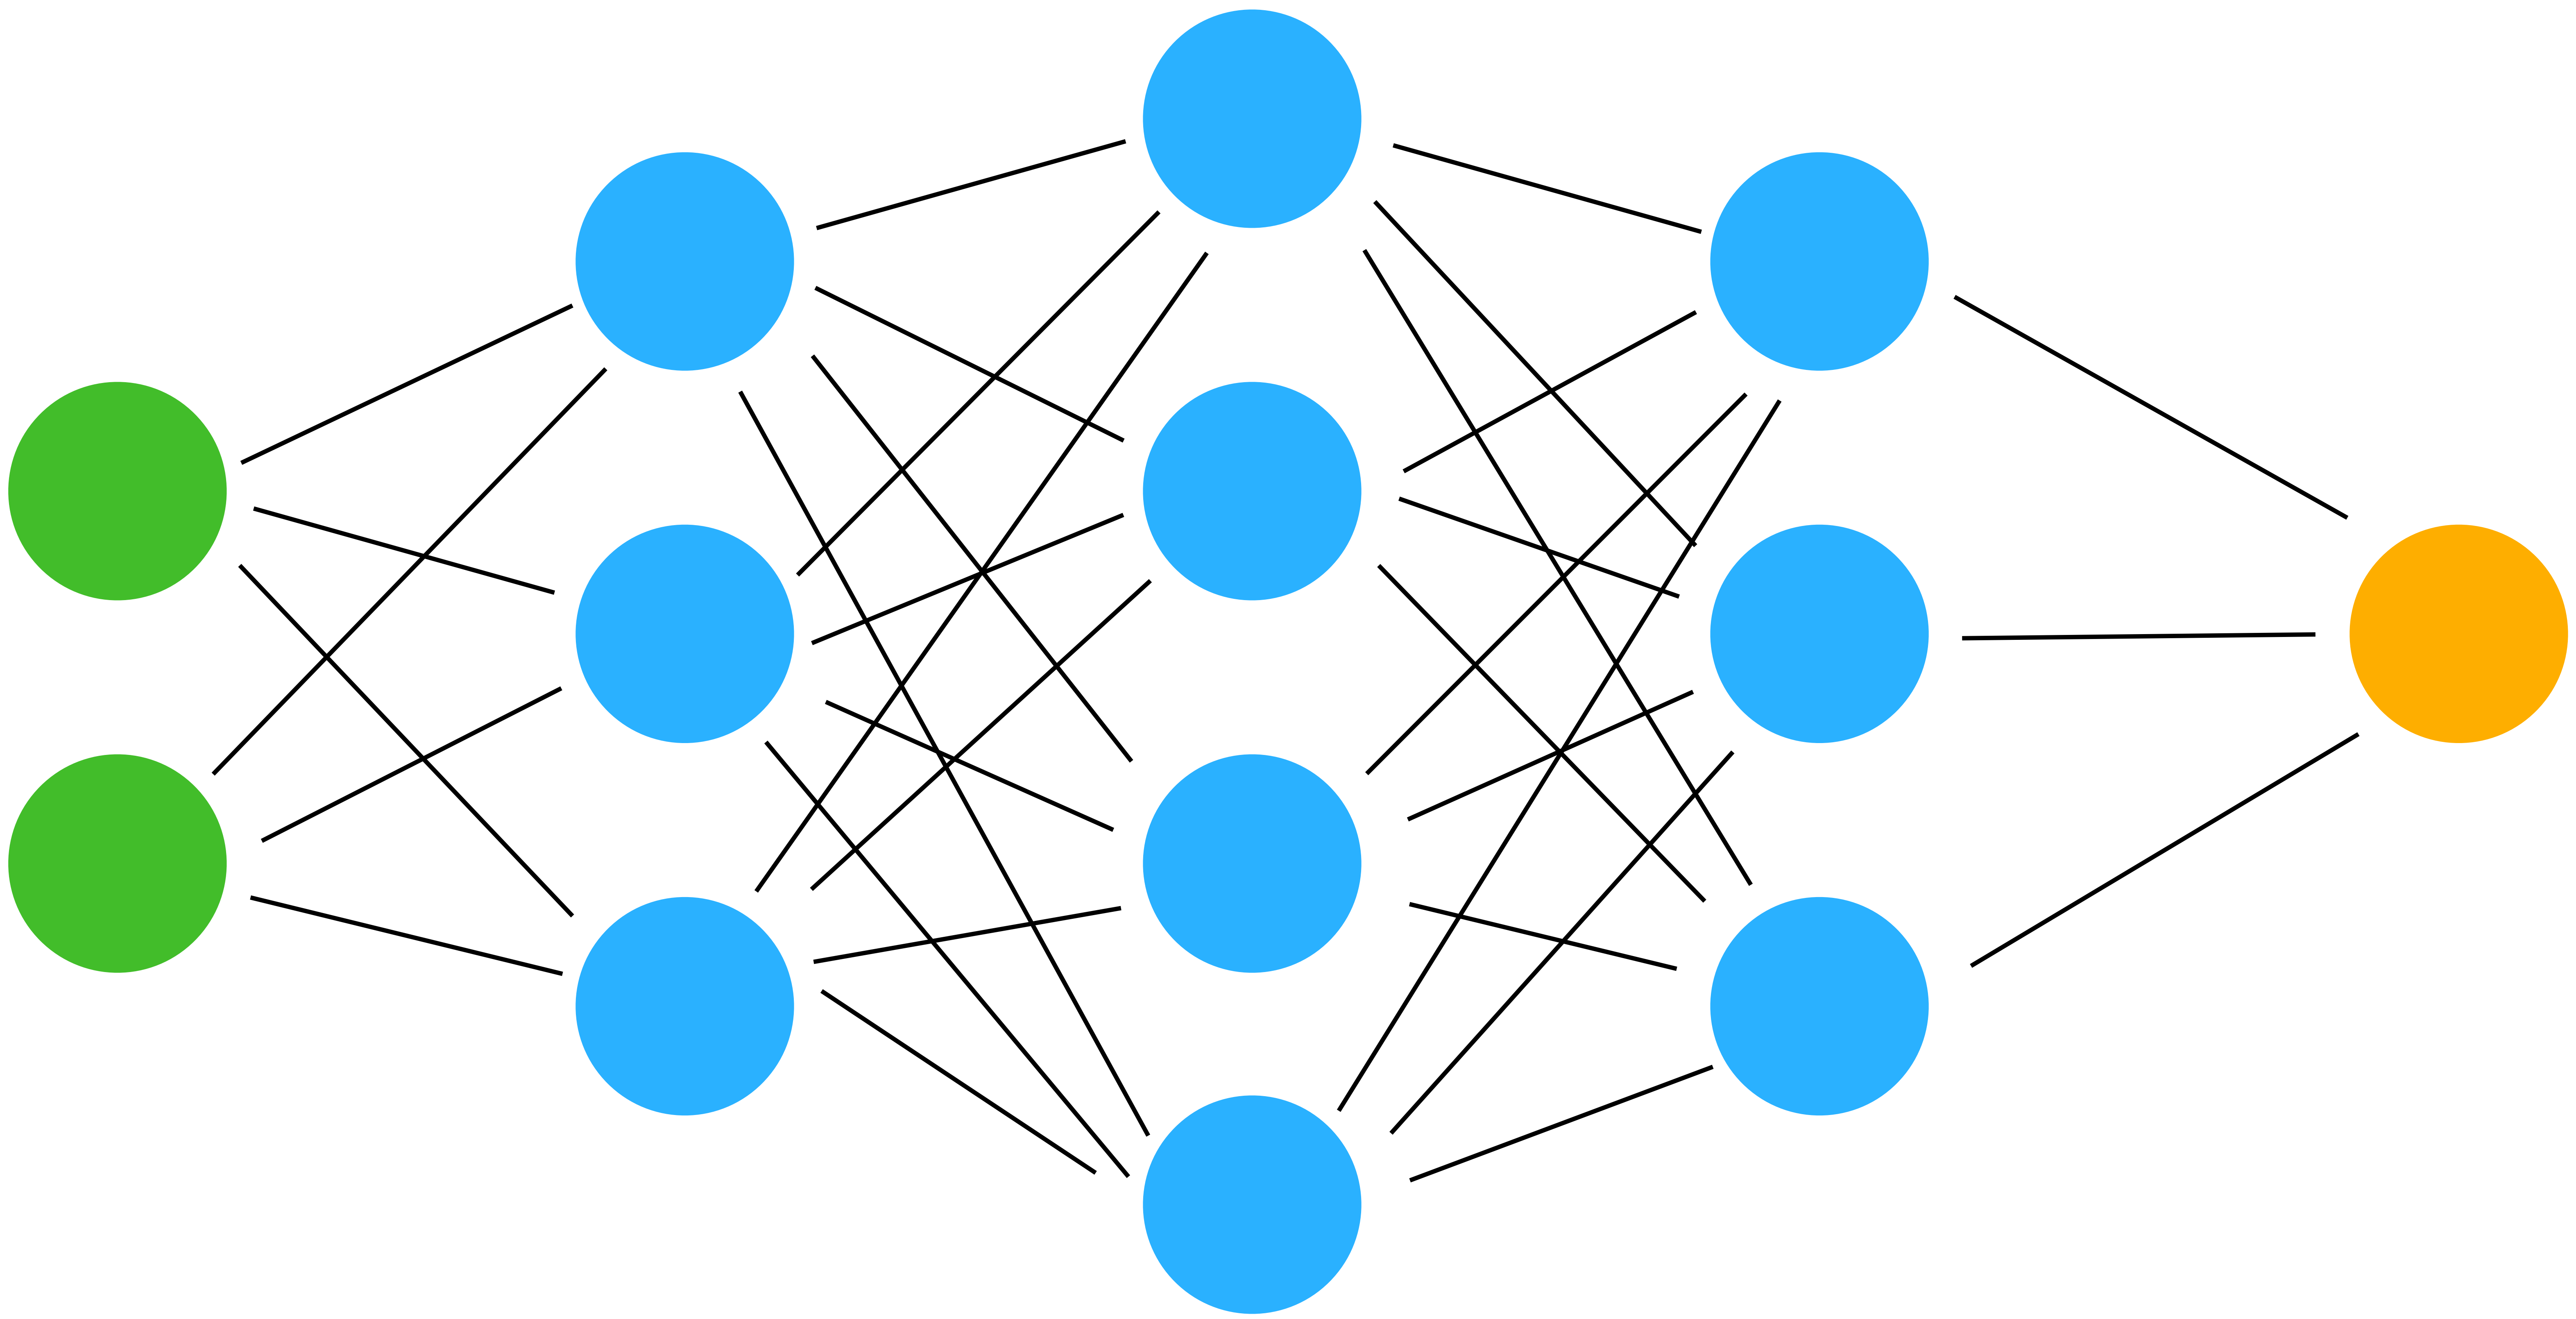

Some Vocabulary

Neural Networks:

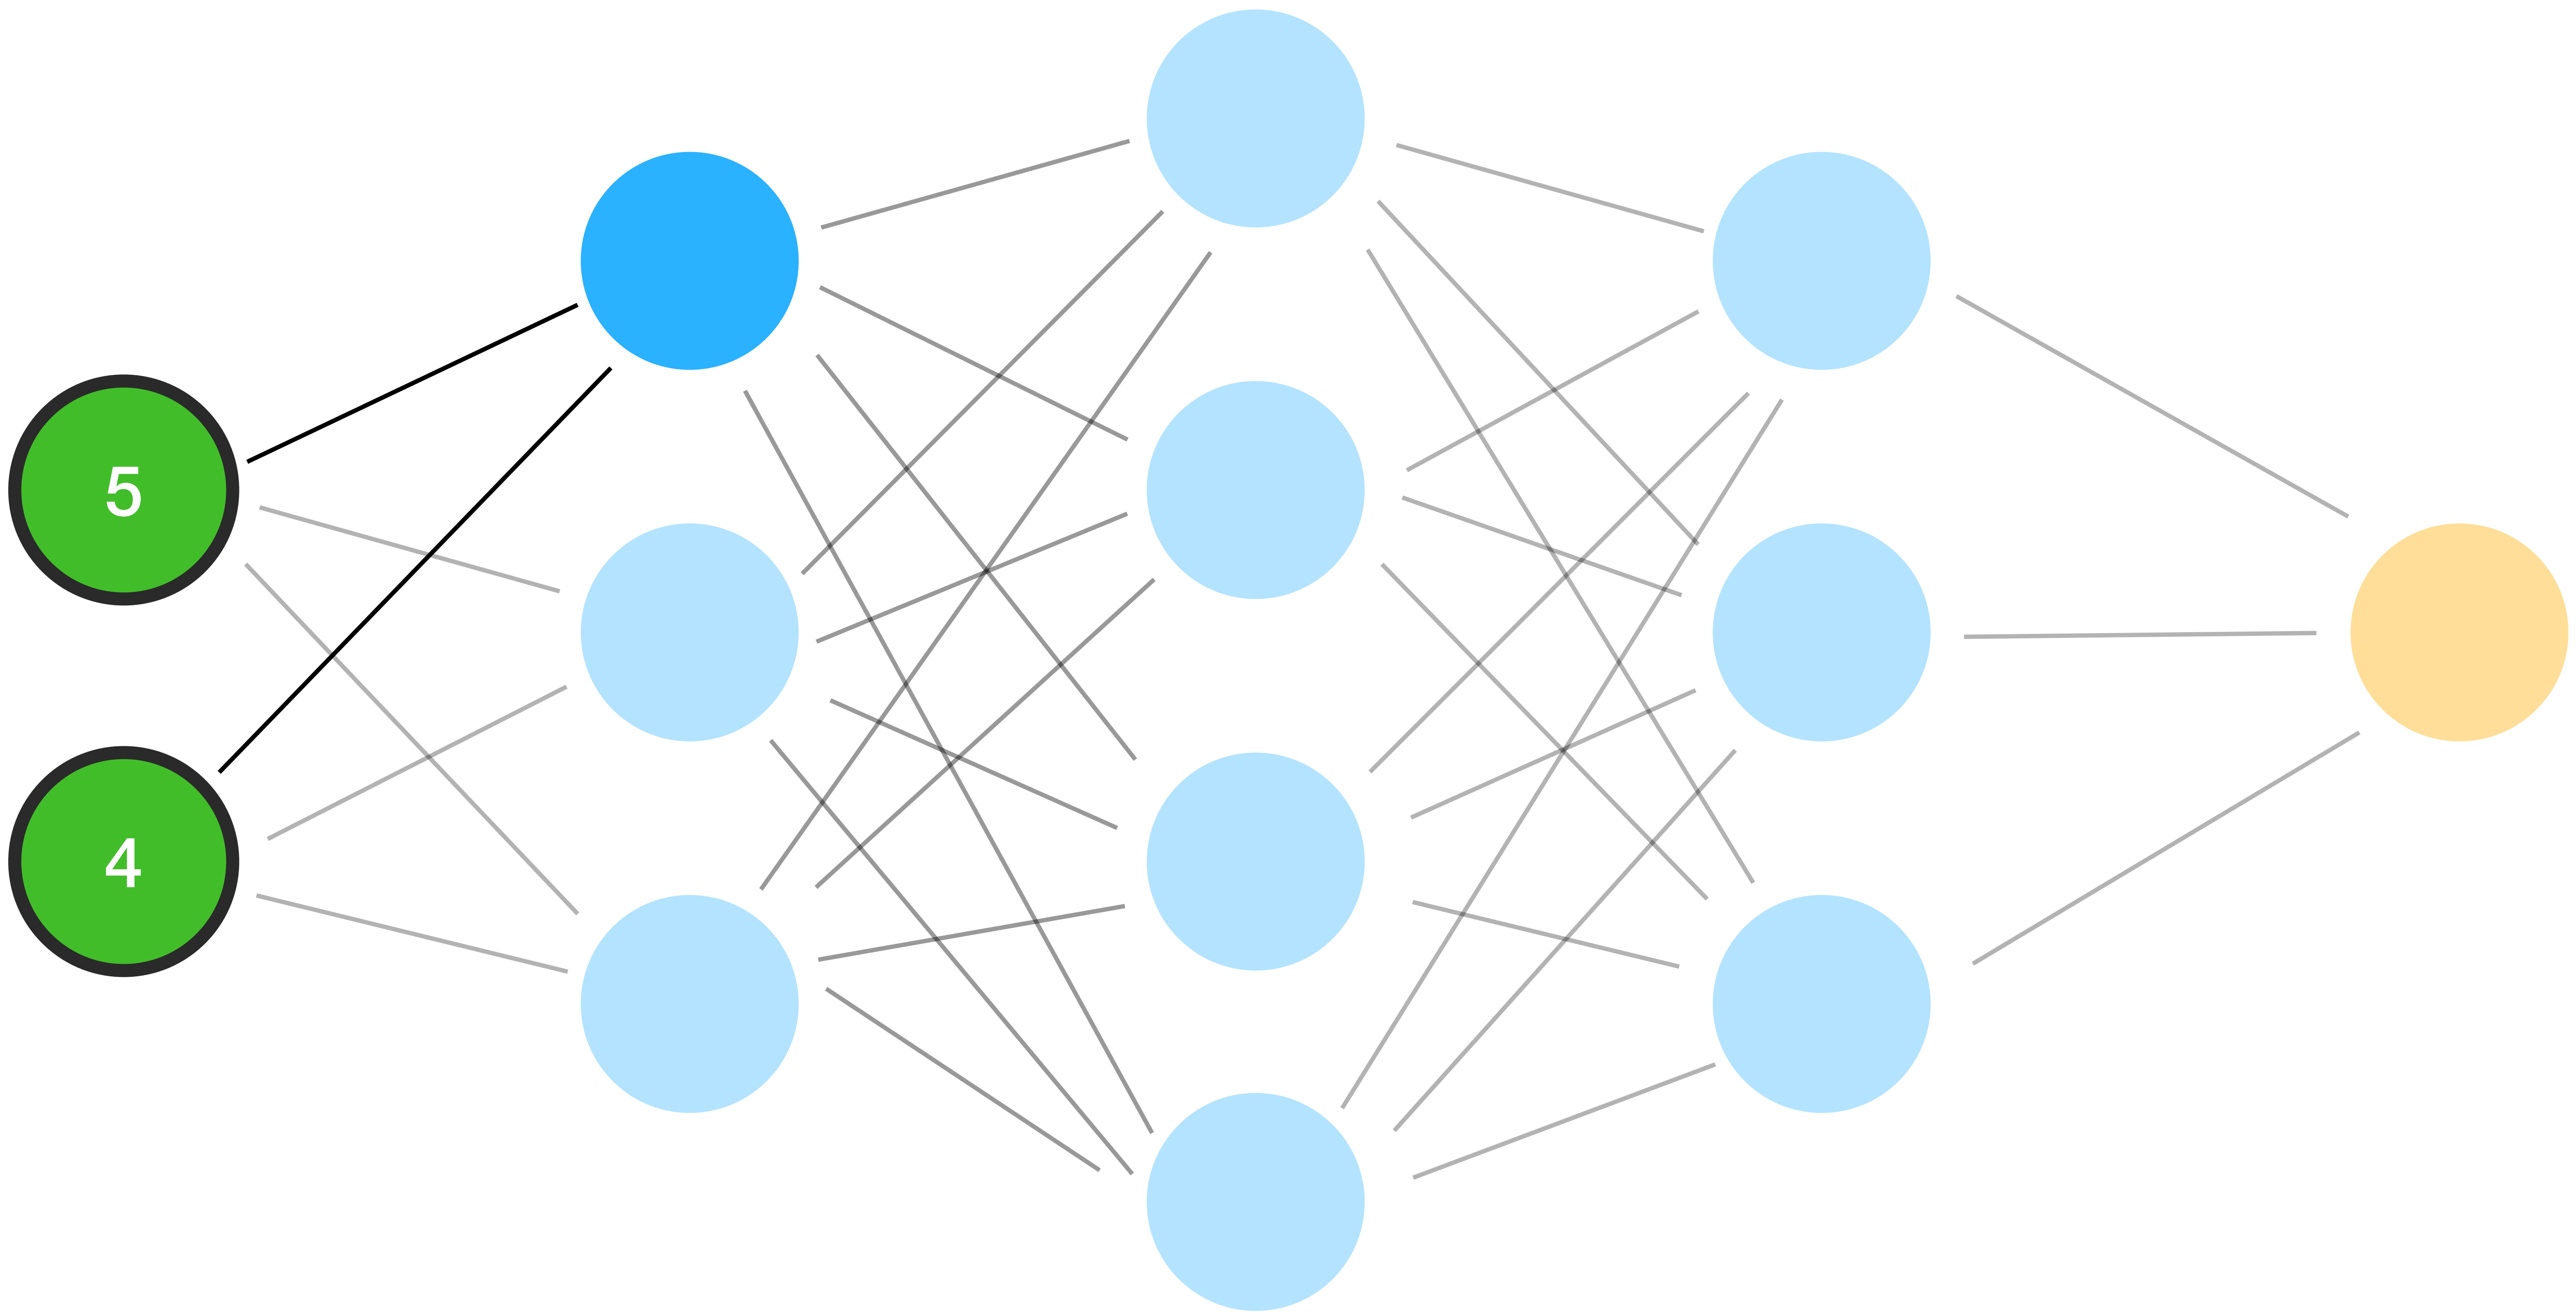

Some Vocabulary

Neural Networks:

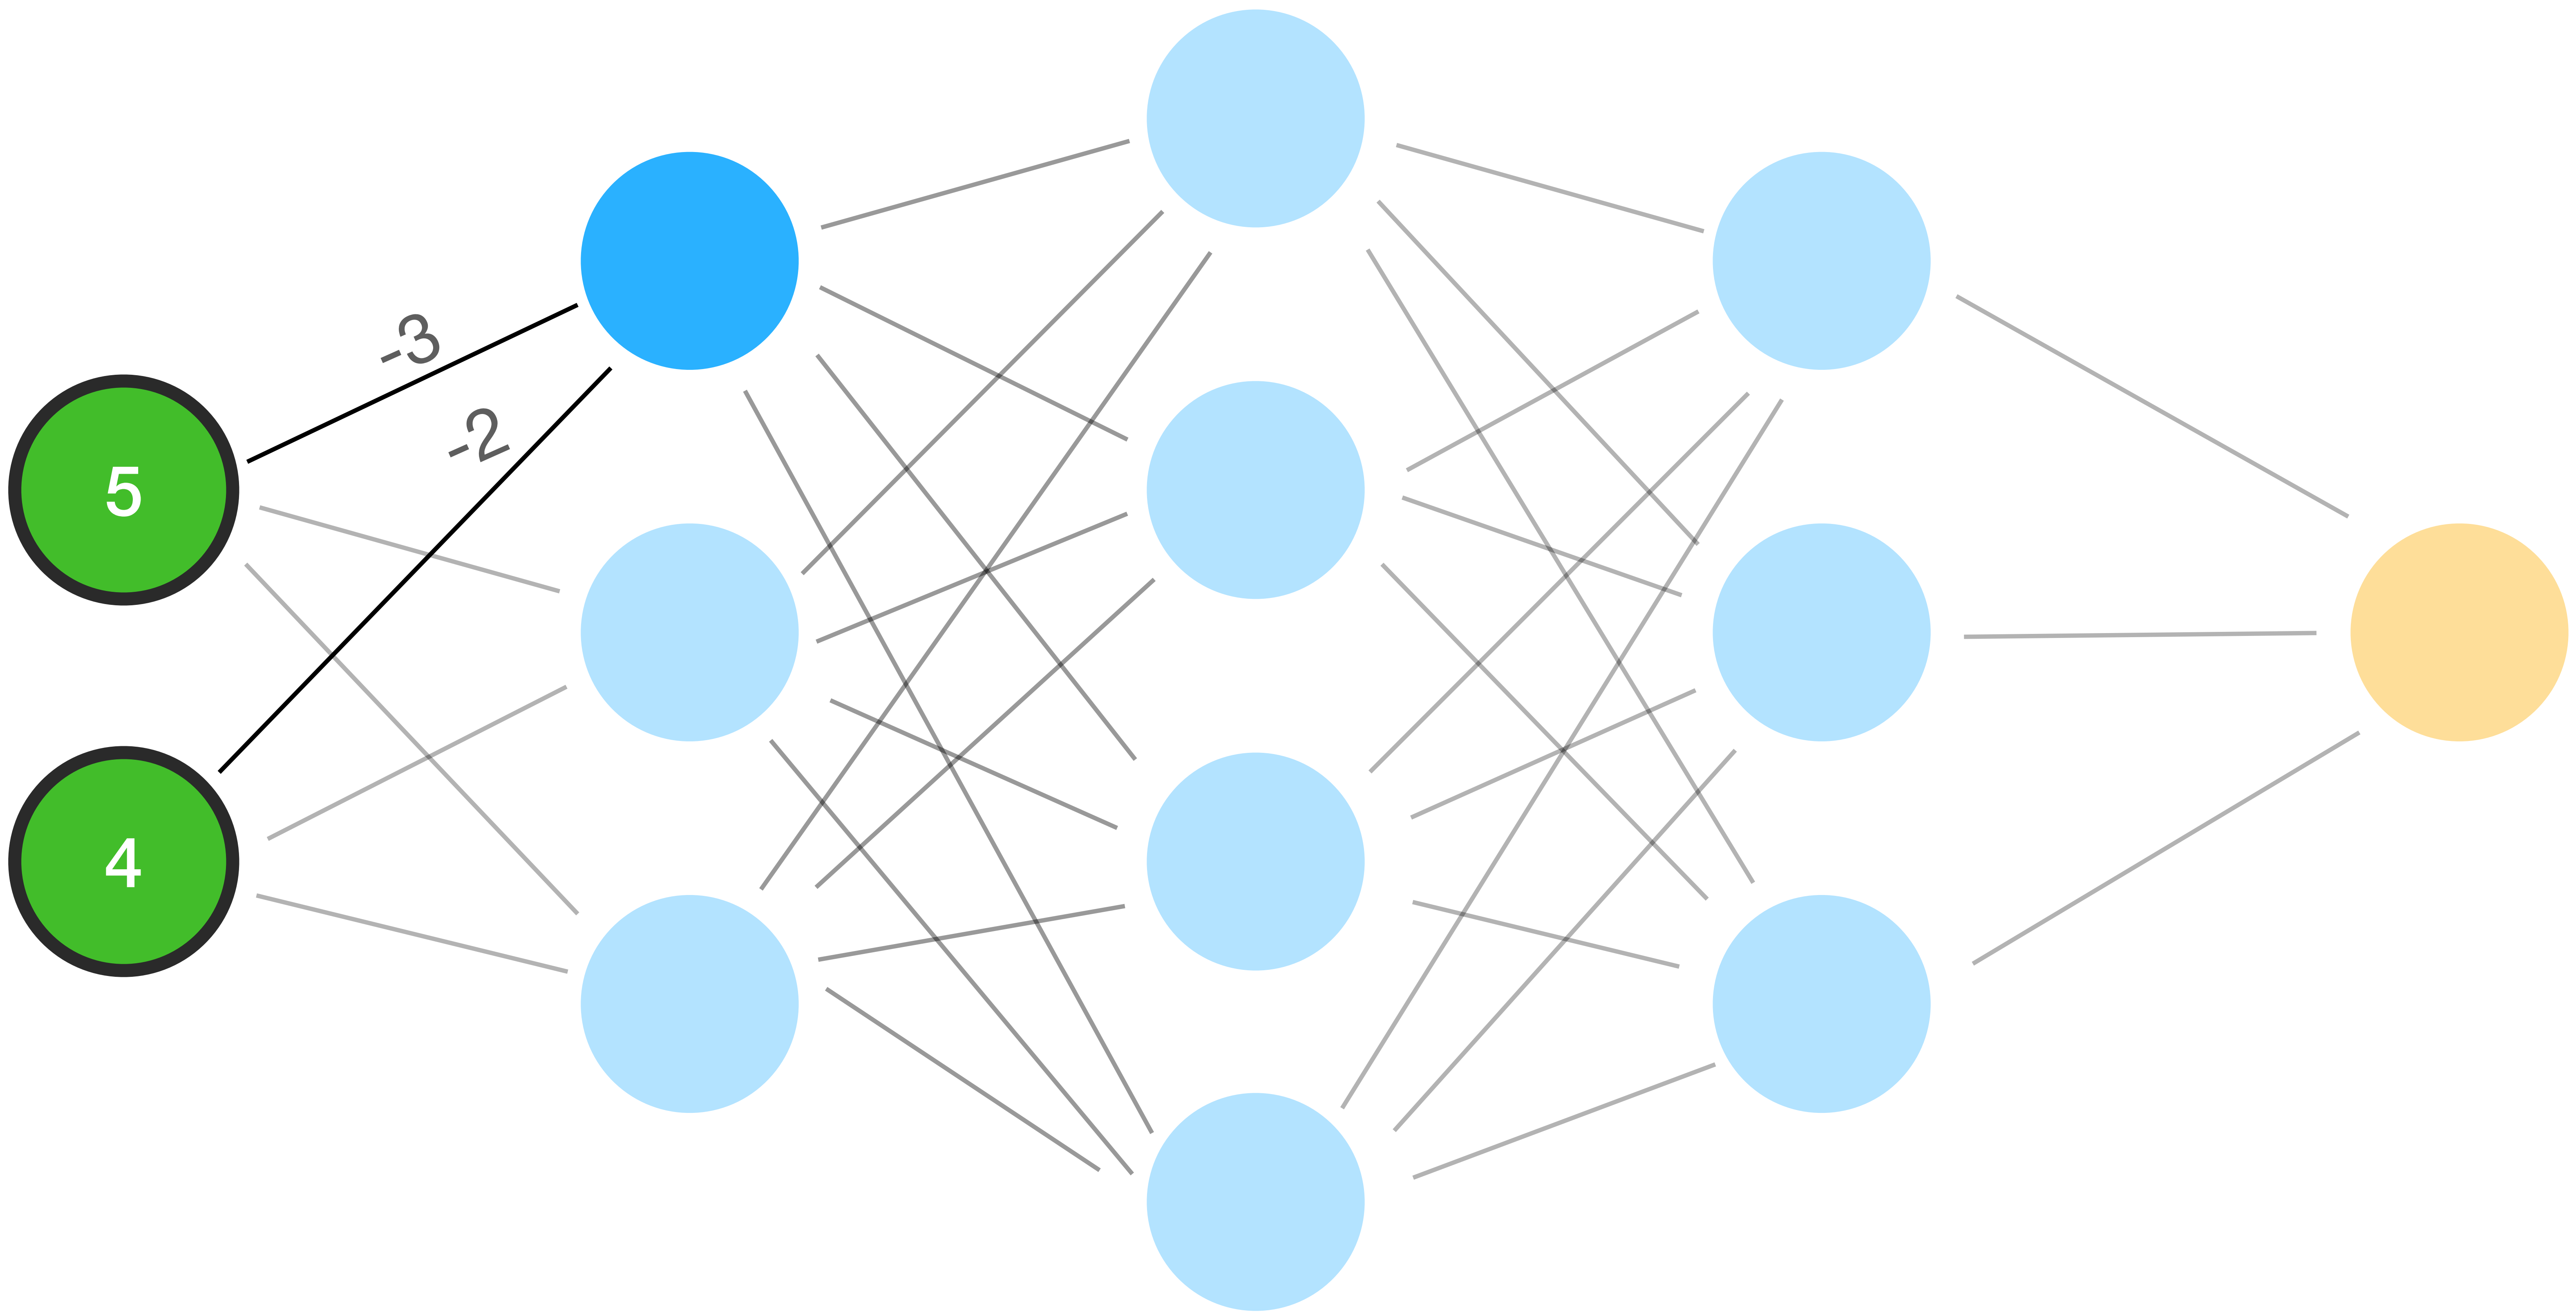

Some Vocabulary

Neural Networks:

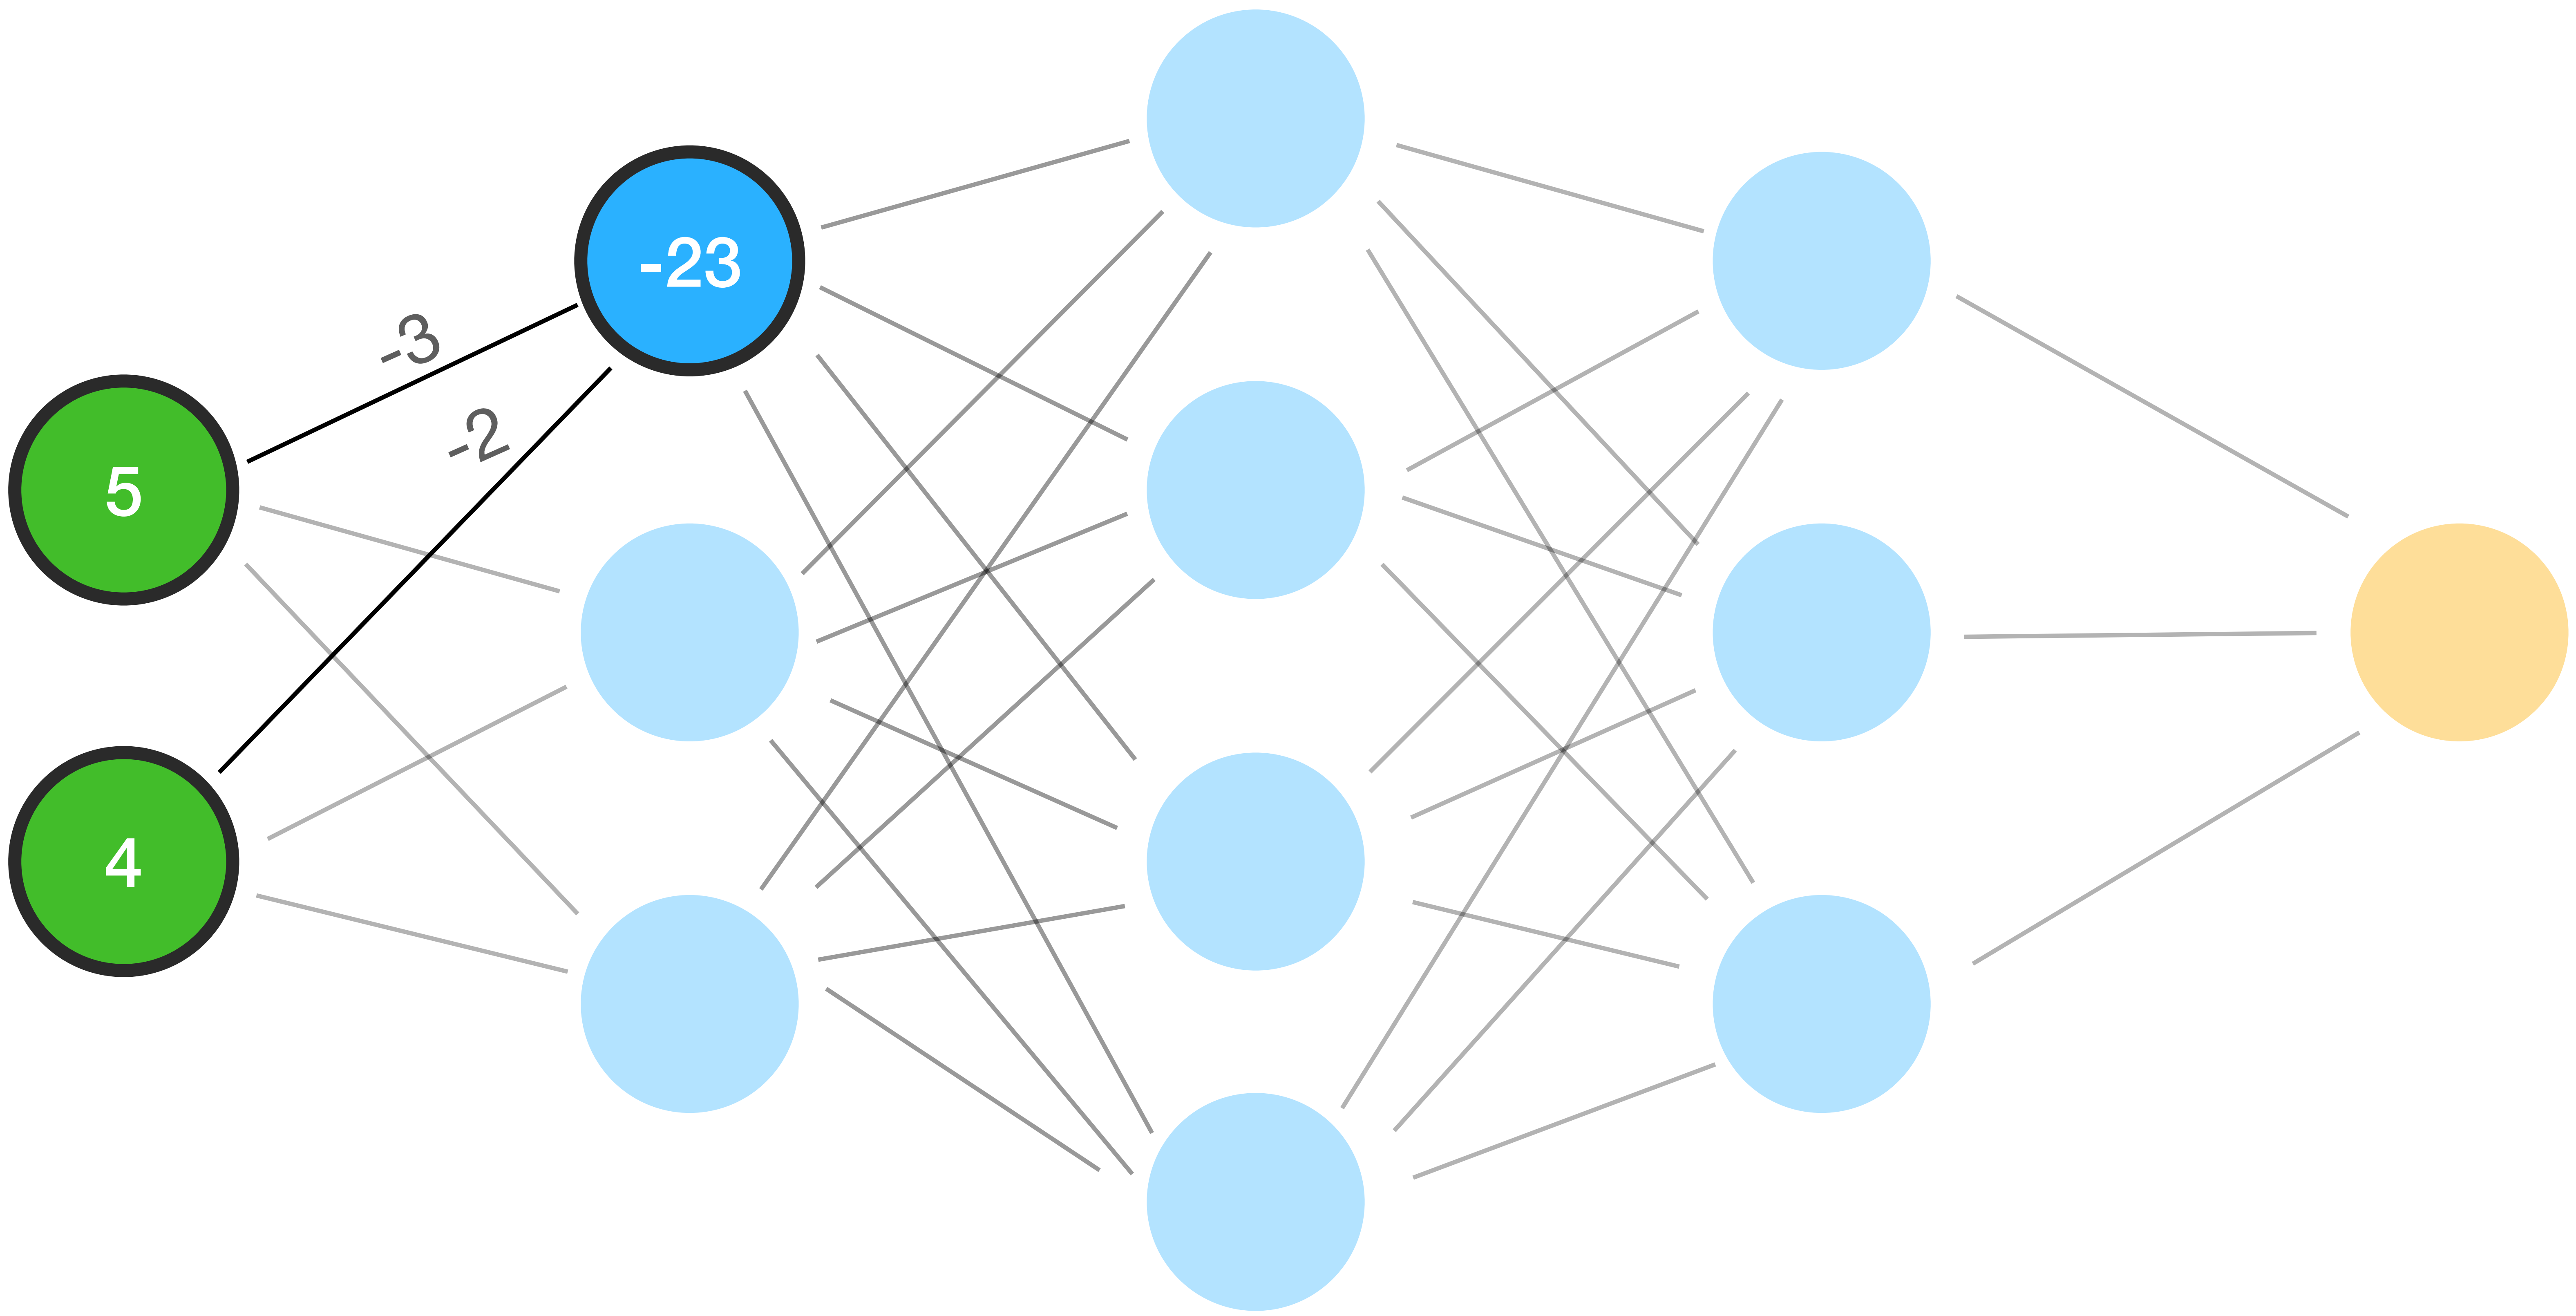

Some Vocabulary

Neural Networks:

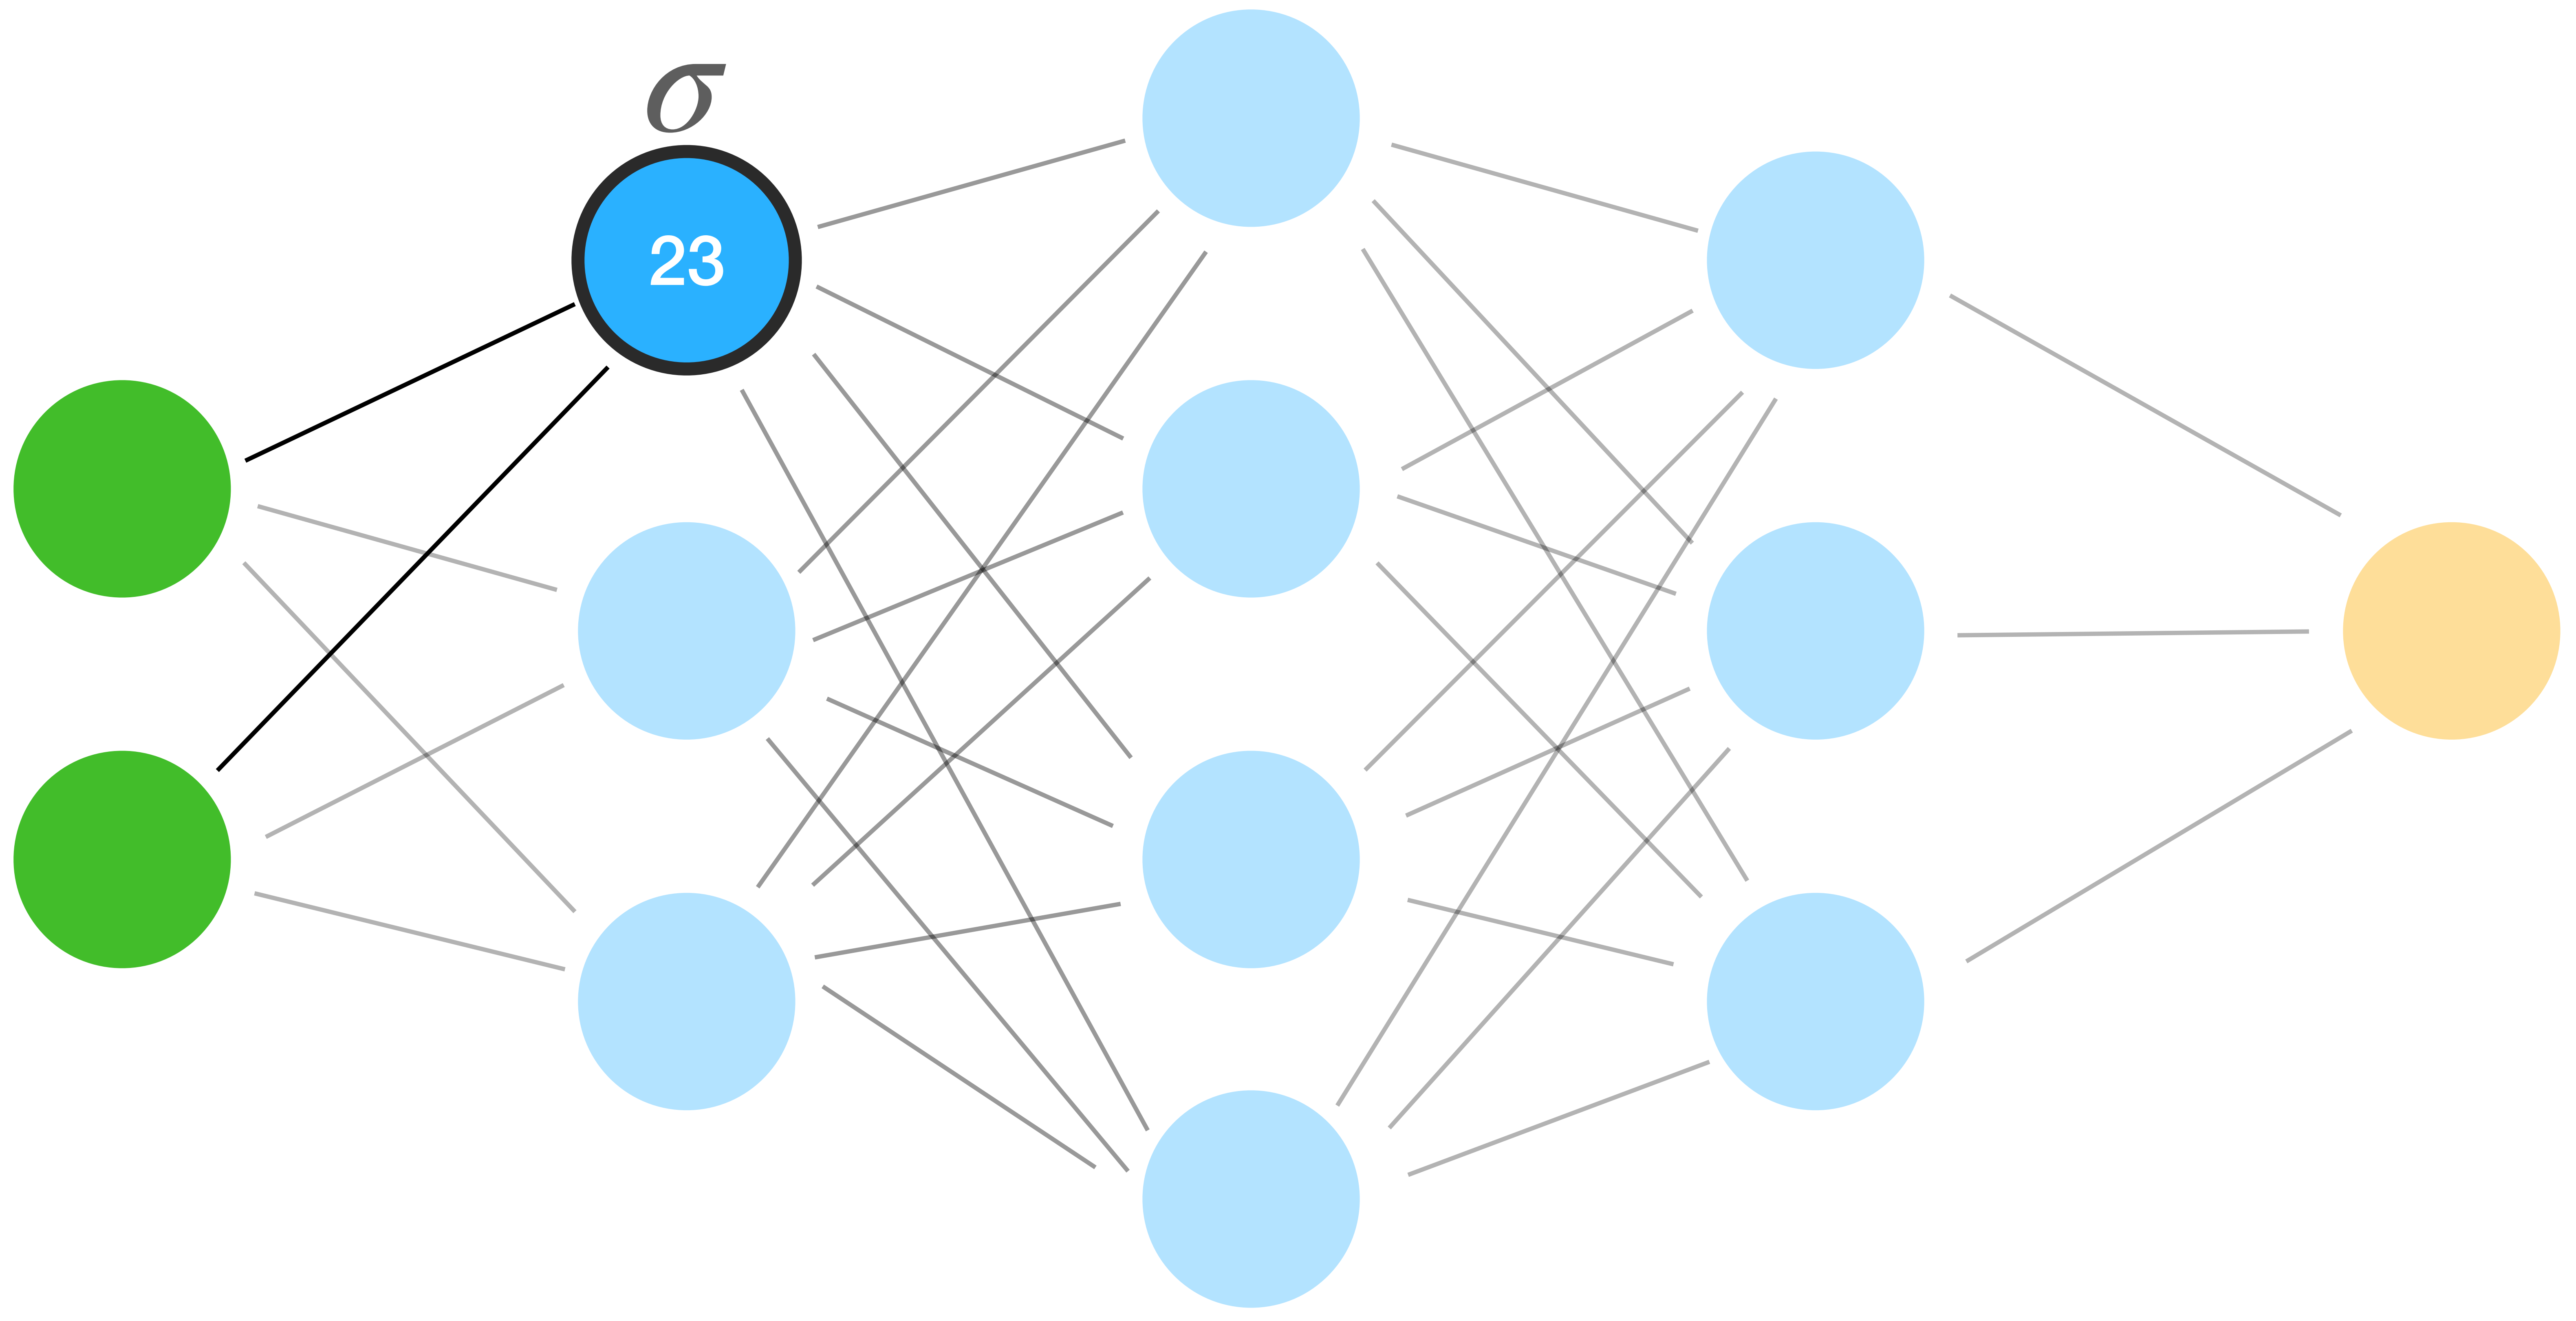

Some Vocabulary

Neural Networks:

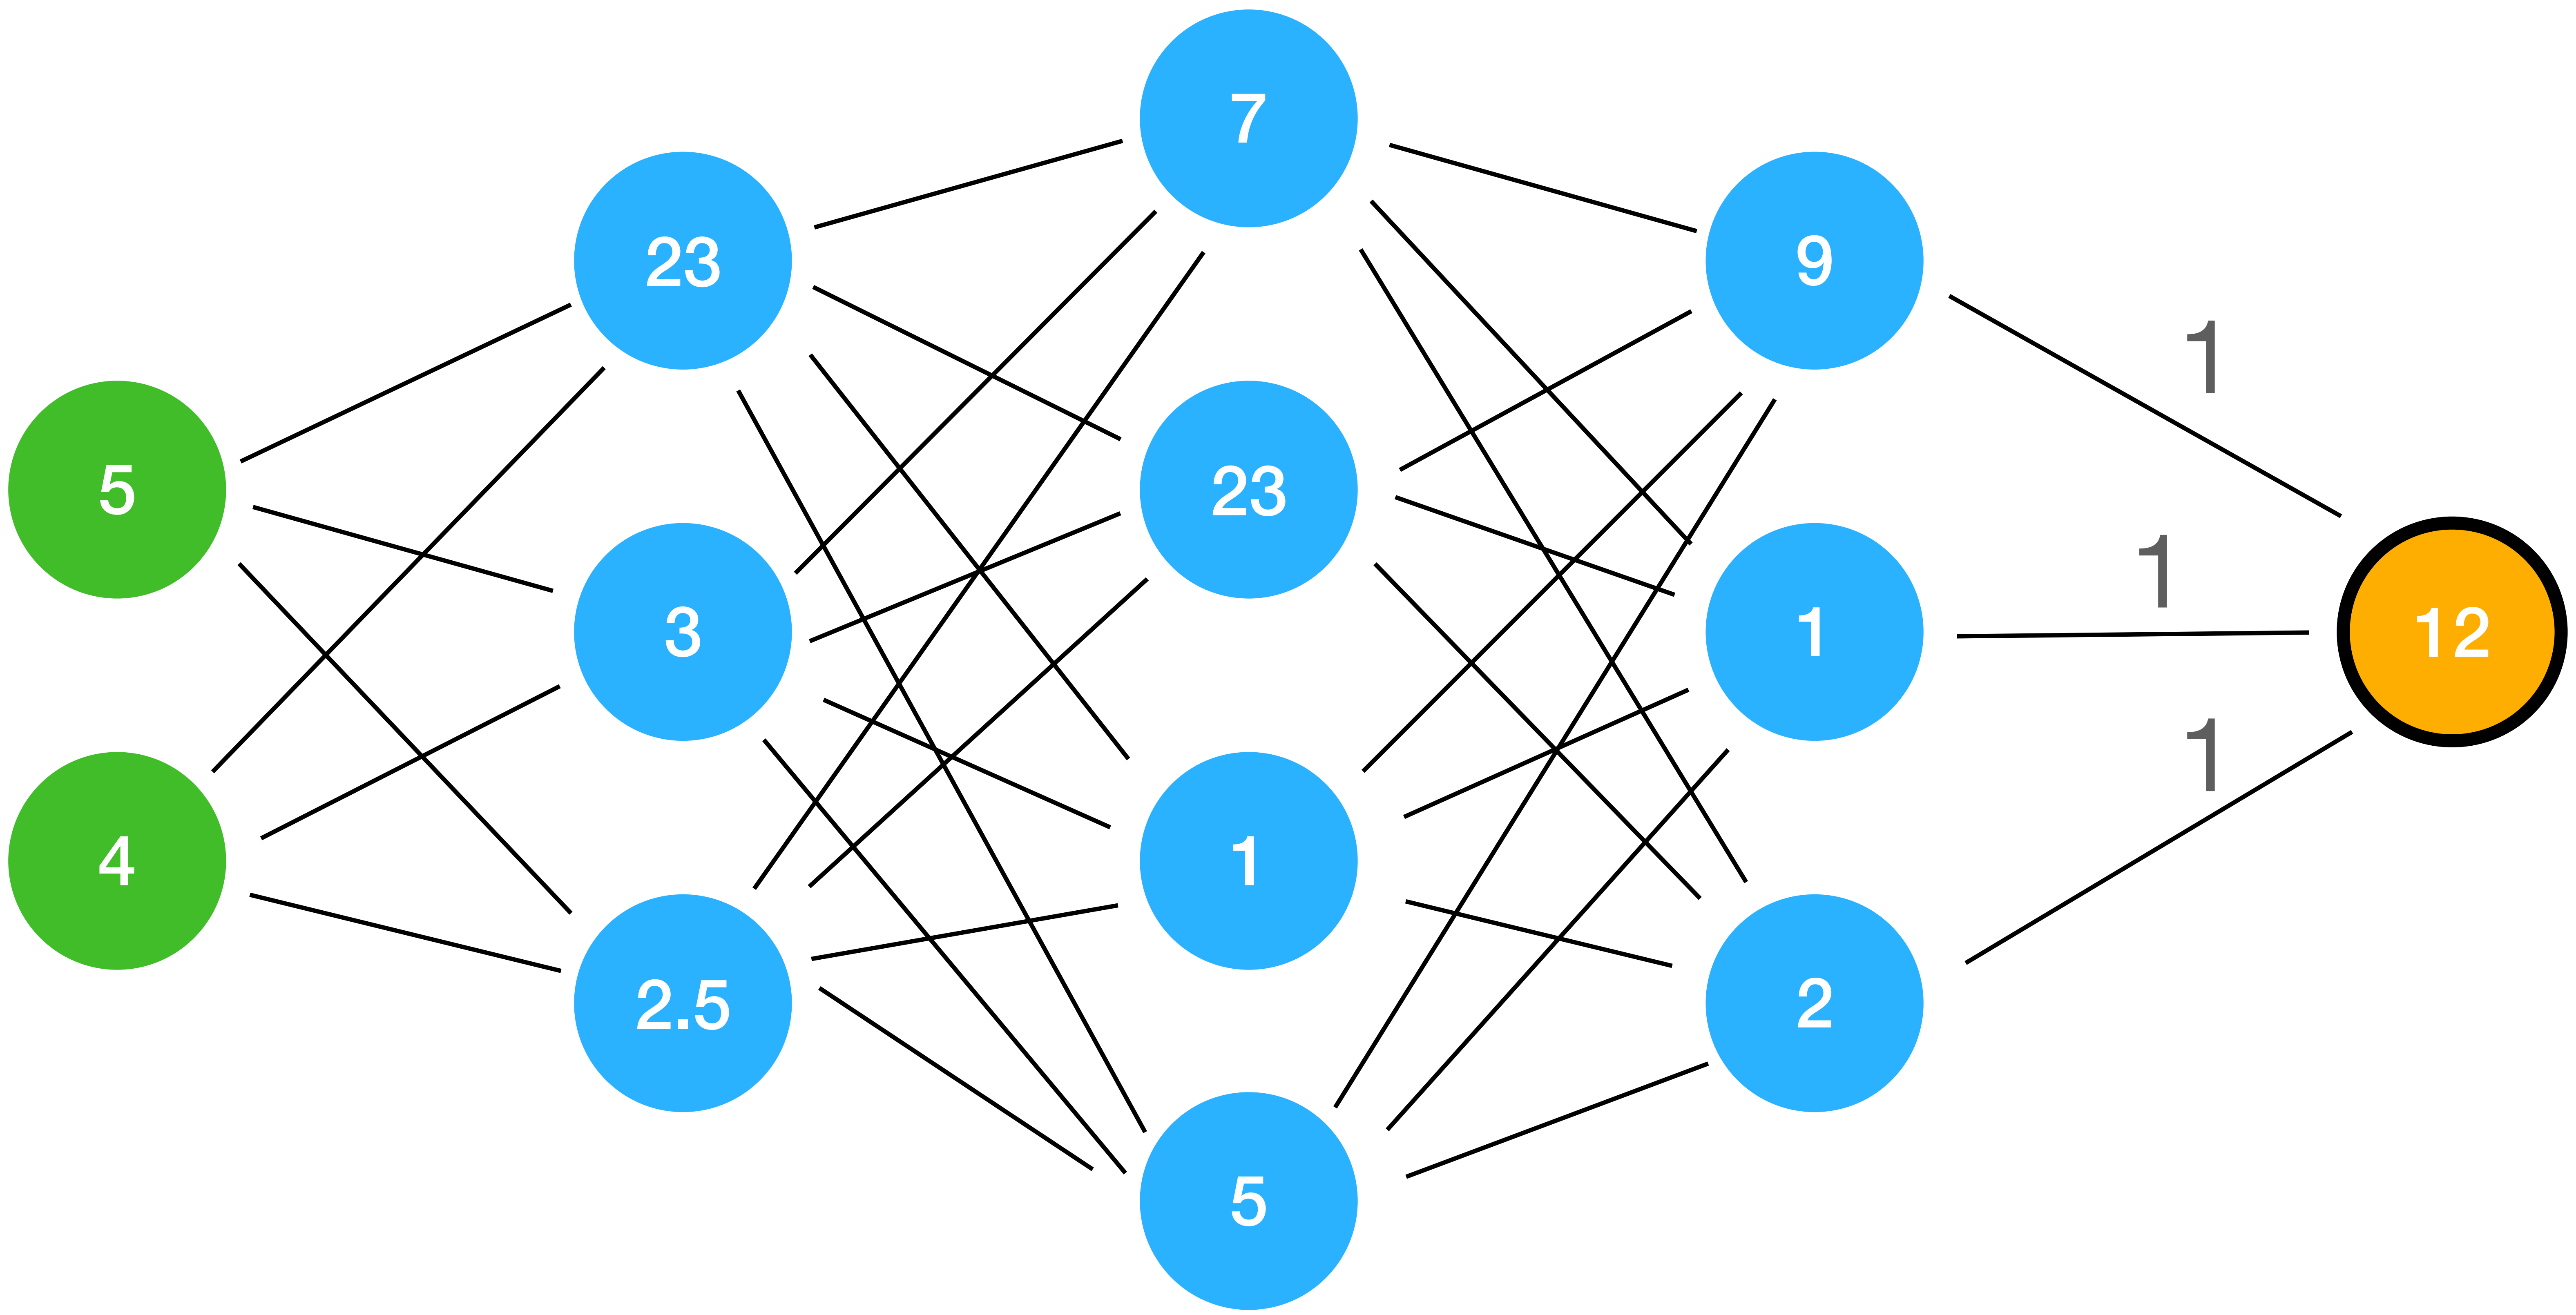

Some Vocabulary

Neural Networks:

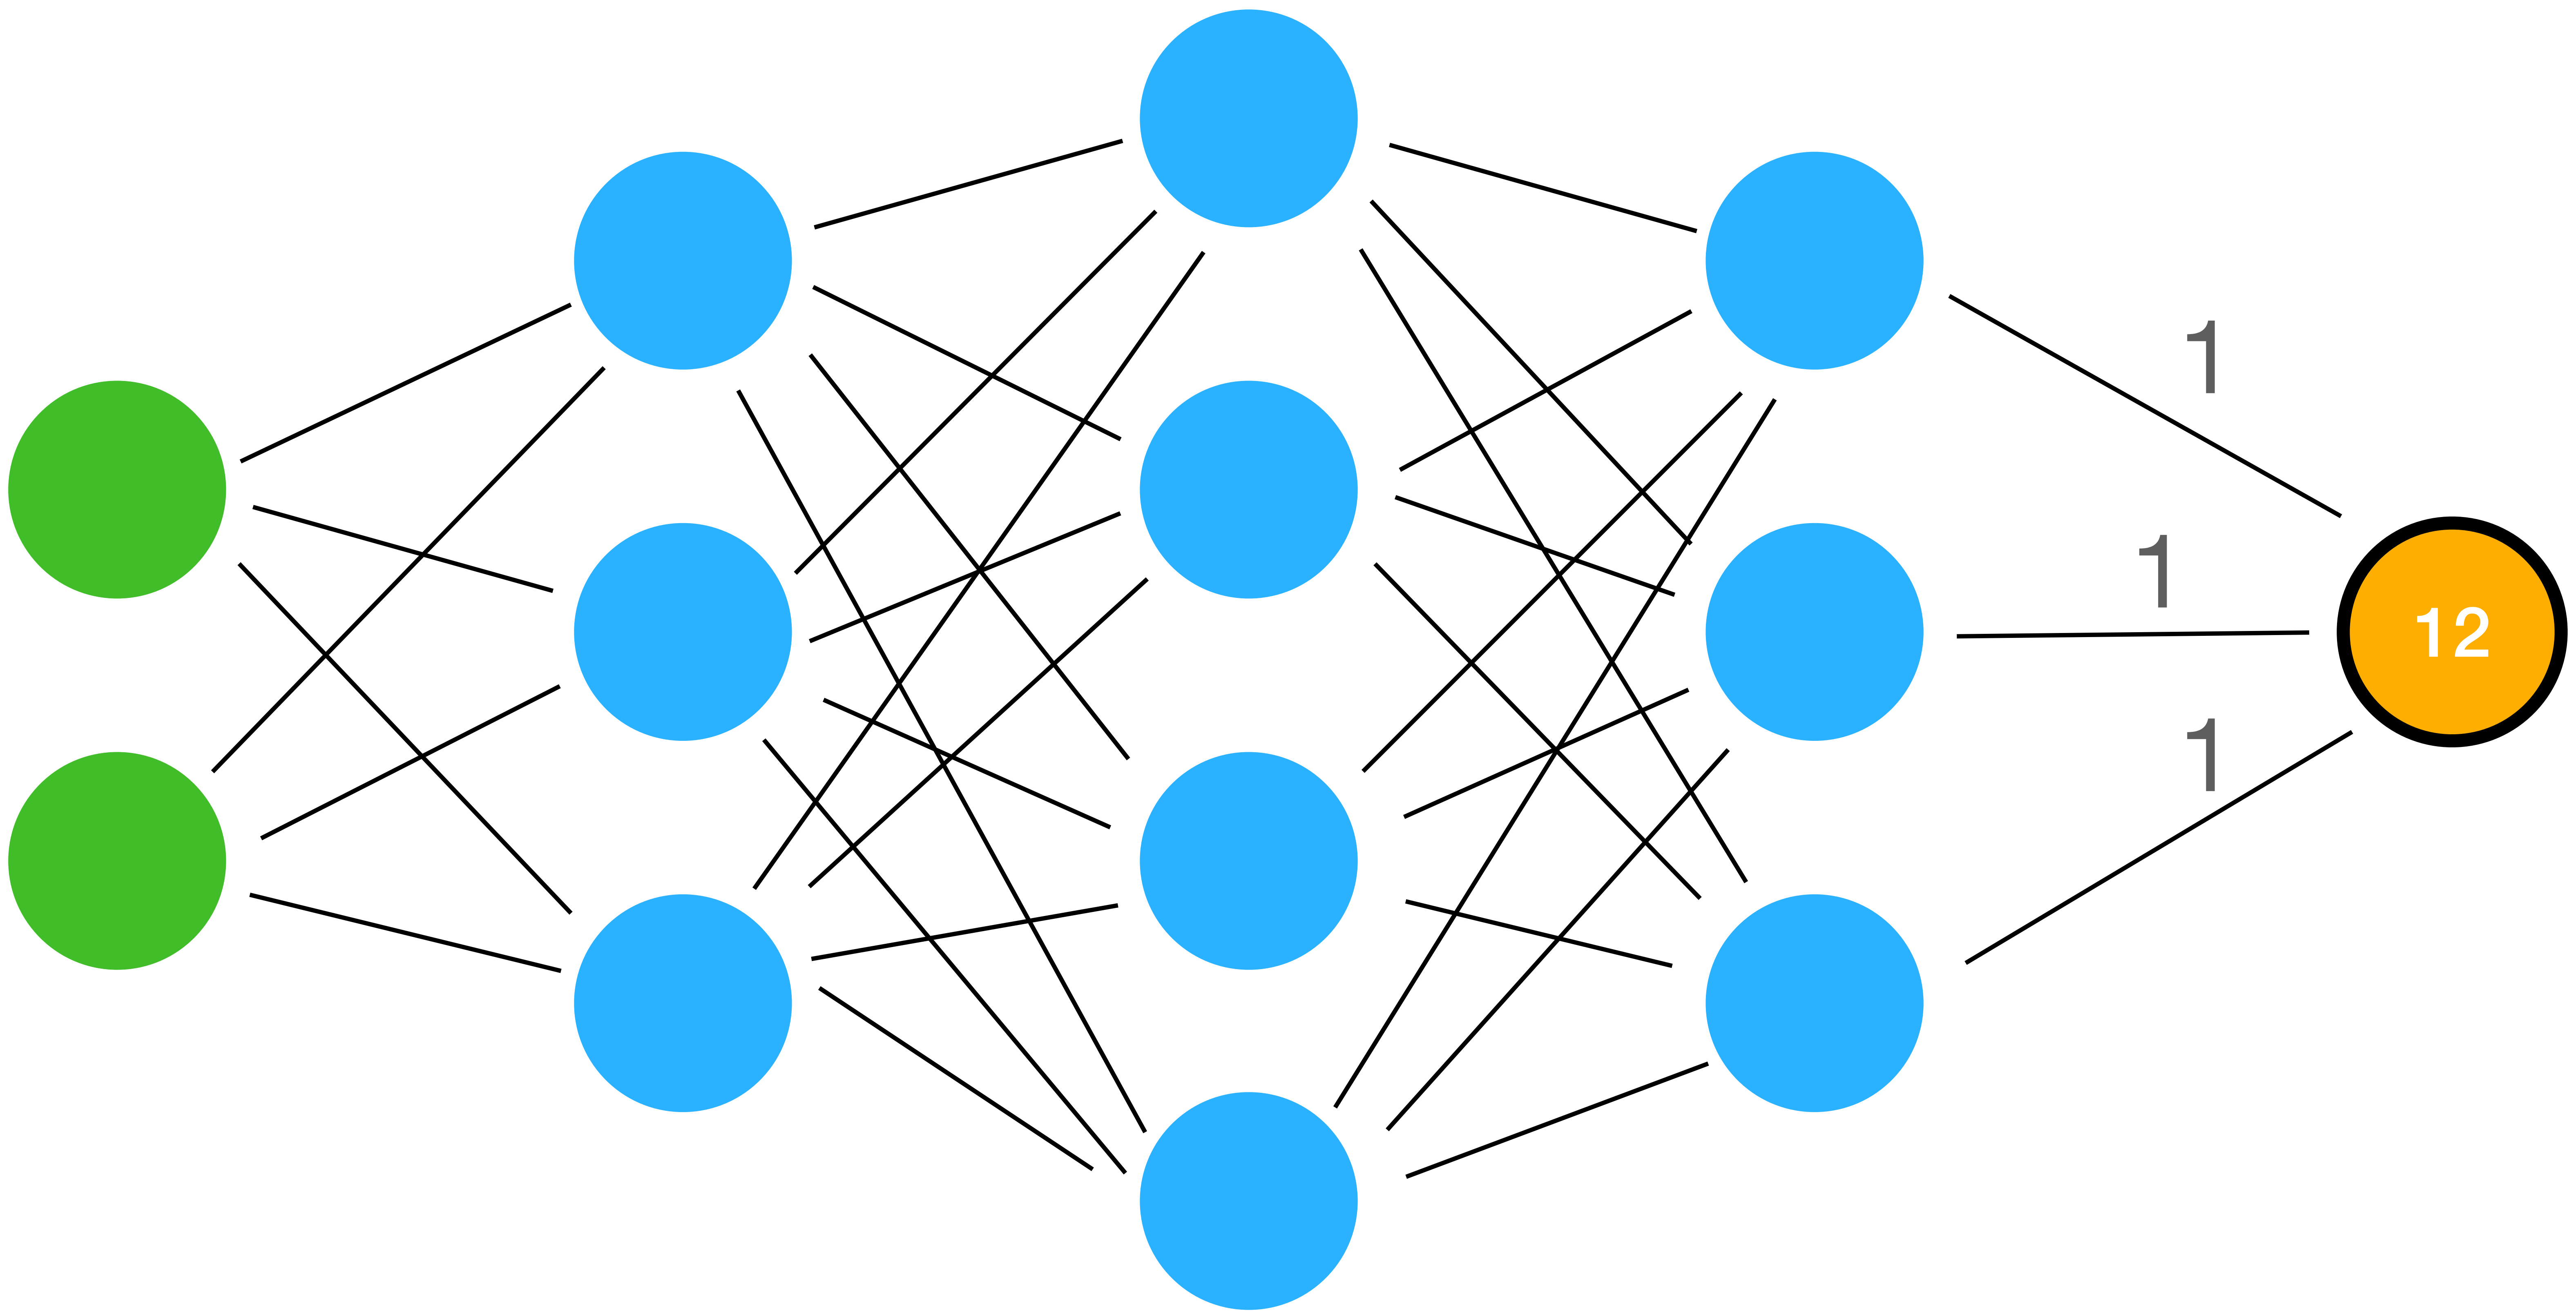

$l(\hat{y}, y)$?

Some Vocabulary

Neural Networks:

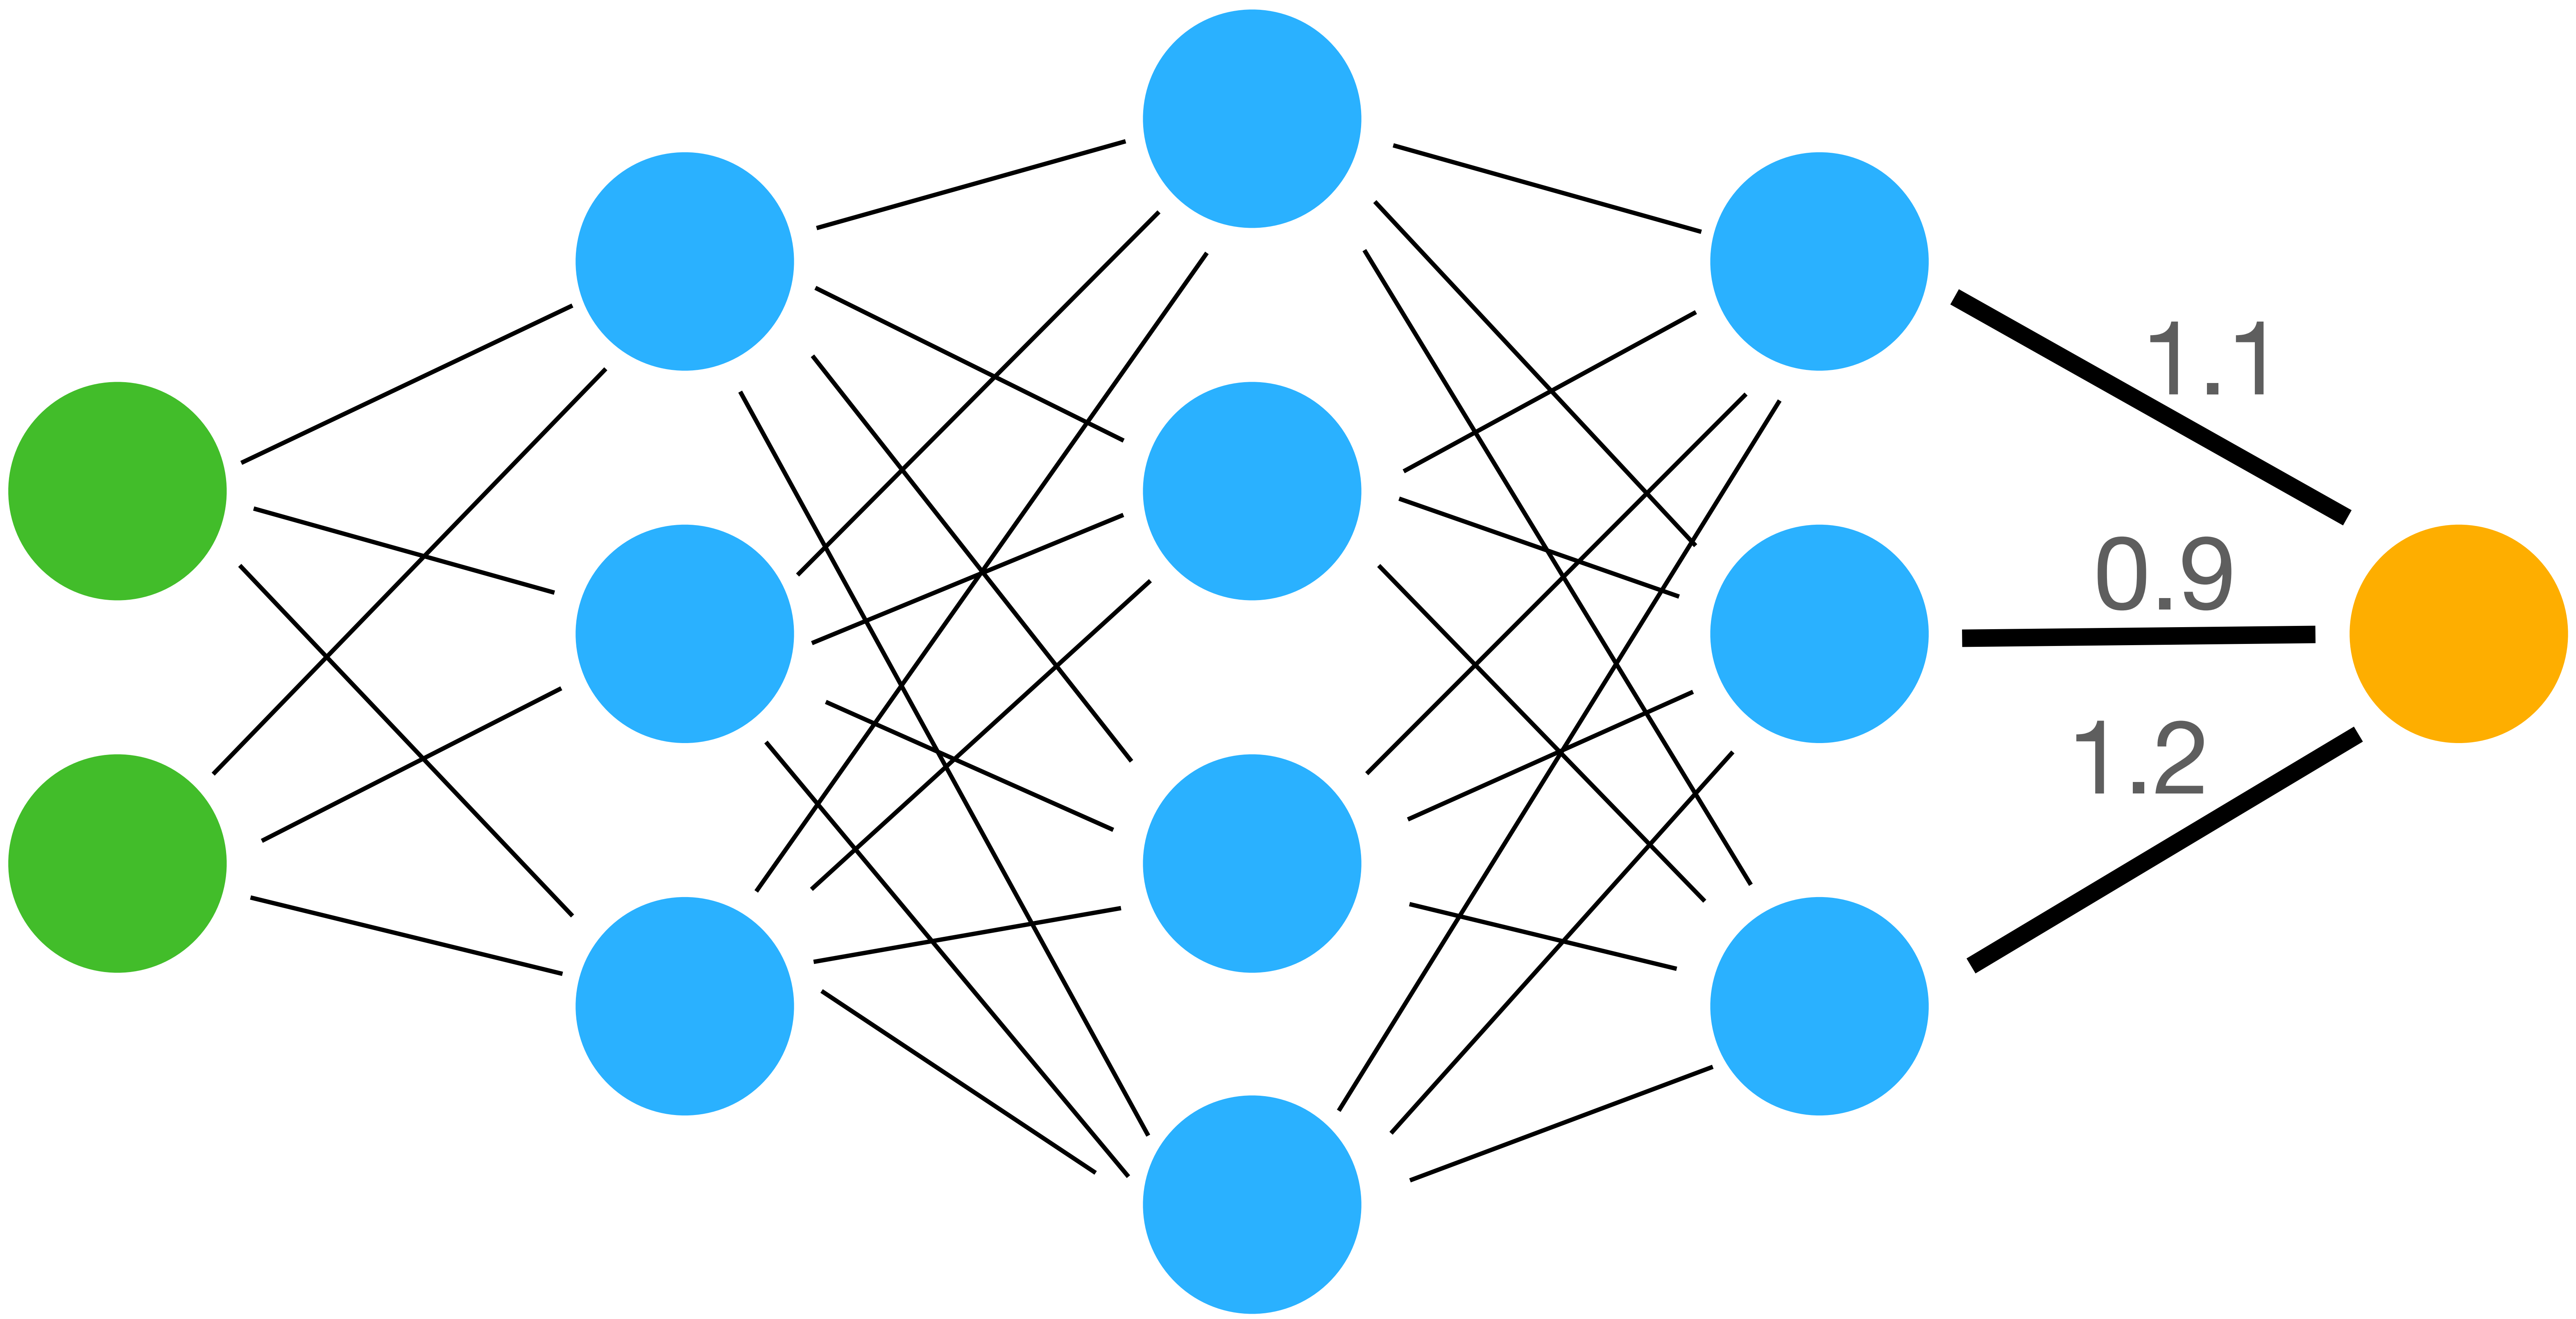

Some Vocabulary

Neural Networks:

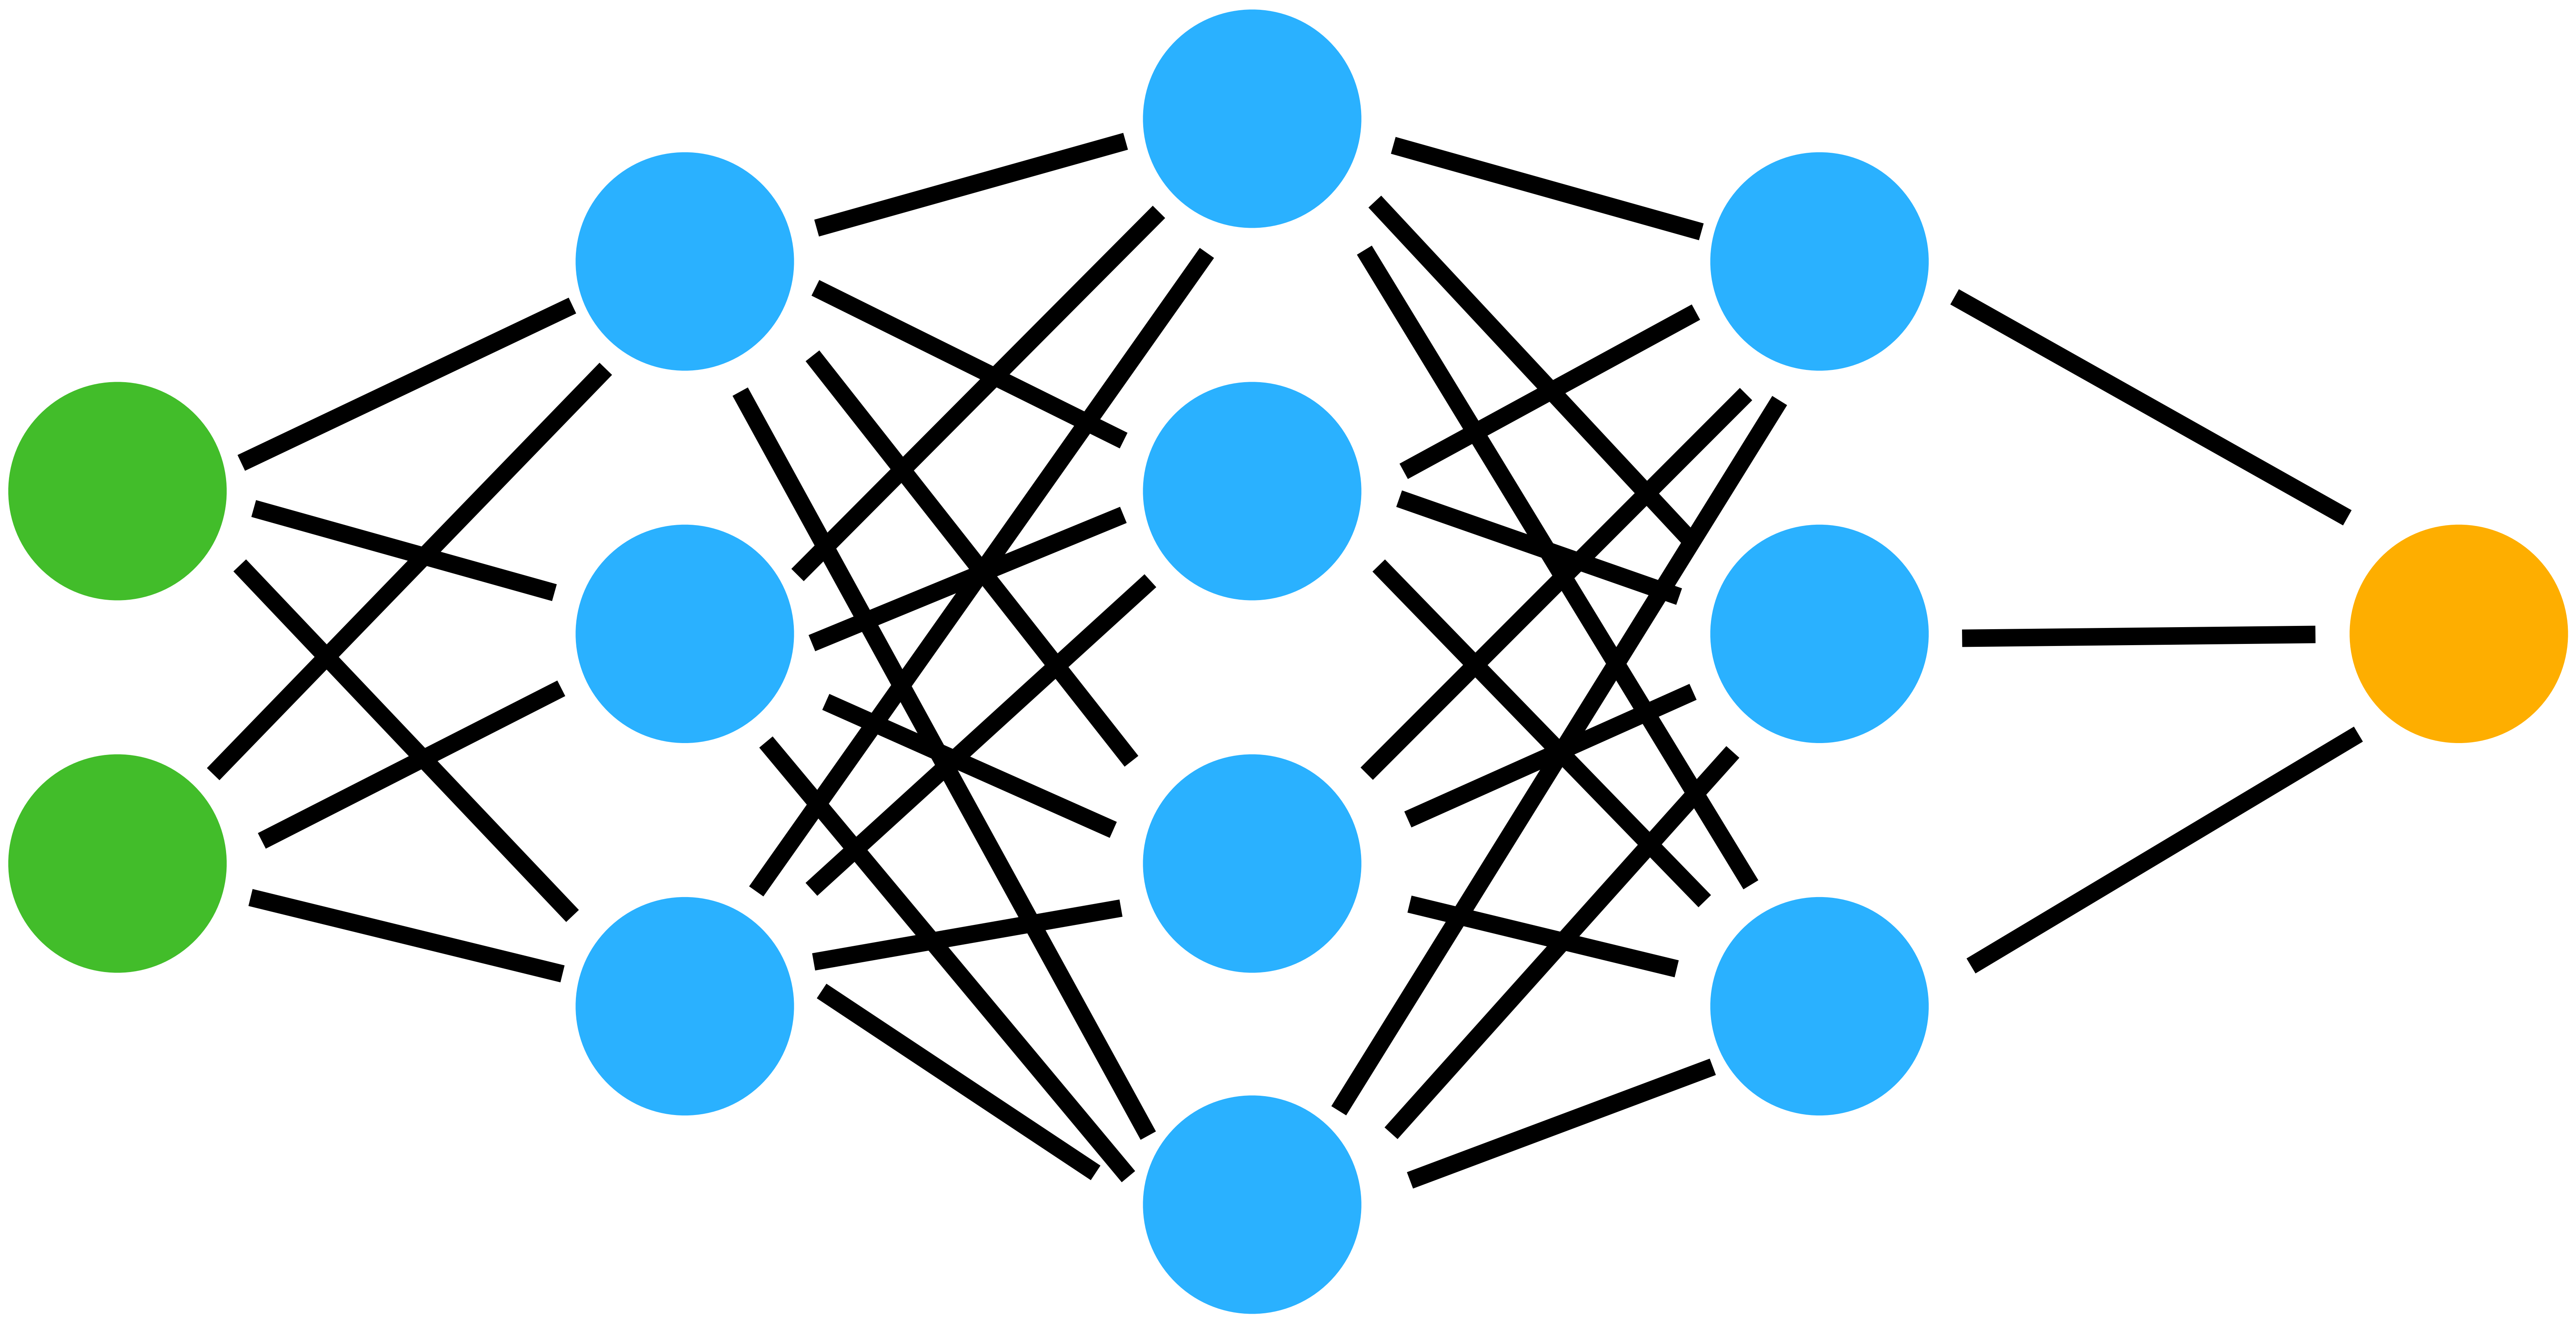

Some Vocabulary

Neural Networks:

Inspiration

nature

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature › articles › article

Article | Open Access | Published: 01 December 2021

Advancing mathematics by guiding human intuition with AI

Alex Davies, Petar Veličković, Lars Buesing, Sam Blackwell, Daniel Zheng, Nenad Tomašev, Richard Tanburn, Peter Battaglia, Charles Blundell, András Juhász, Marc Lackenby, Geordie Williamson, Demis Hassabis & Pushmeet Kohli

Nature **600**, 70–74 (2021) | Cite this article

Abstract

The practice of mathematics involves discovering patterns and using these to formulate and prove conjectures, resulting in theorems. Since the 1960s, mathematicians have used computers to assist in the discovery of patterns and formulation of conjectures[1], most famously in the Birch and Swinnerton-Dyer conjecture[2], a Millennium Prize Problem[3]. Here we provide examples of new fundamental results in pure mathematics that have been discovered with the assistance of machine learning—demonstrating a method by which machine learning can aid mathematicians in discovering new conjectures and theorems. We propose a process of using machine learning to discover potential patterns and relations between mathematical objects, understanding them with attribution techniques and using these observations to guide intuition and propose conjectures. We outline this machine-

[1]

nature communications

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature › nature communications › articles › article

Article | Open Access | Published: 12 April 2023

Combining data and theory for derivable scientific discovery with AI-Descartes

Cristina Cornelio, Sanjeeb Dash, Vernon Austel, Tyler R. Josephson, Joao Goncalves, Kenneth L. Clarkson, Nimrod Megiddo, Bachir El Khadir & Lior Horesh

Nature Communications **14**, Article number: 1777 (2023) | Cite this article

Abstract

Scientists aim to discover meaningful formulae that accurately describe experimental data. Mathematical models of natural phenomena can be manually created from domain knowledge and fitted to data, or, in contrast, created automatically from large datasets with machine-learning algorithms. The problem of incorporating prior knowledge expressed as constraints on the functional form of a learned model has been studied before, while finding models that are consistent with prior knowledge expressed via general logical axioms is an open problem. We develop a method to enable principled derivations of models of natural phenomena from axiomatic knowledge and experimental data by combining logical reasoning with symbolic regression. We demonstrate these concepts for Kepler's third law of planetary motion, Einstein's relativistic time-dilation law, and Langmuir's theory of adsorption. We show we can discover governing laws from few data points when logical reasoning is used to distinguish between candidate formulae having similar error on the data.

[2]

nature

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature › articles › article

Article | Open Access | Published: 05 October 2022

Discovering faster matrix multiplication algorithms with reinforcement learning

Alhussein Fawzi, Matej Balog, Aja Huang, Thomas Hubert, Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Francisco J. R. Ruiz, Julian Schrittwieser, Grzegorz Swirszcz, David Silver, Demis Hassabis & Pushmeet Kohli

Nature **610**, 47–53 (2022) | Cite this article

Abstract

Improving the efficiency of algorithms for fundamental computations can have a widespread impact, as it can affect the overall speed of a large amount of computations. Matrix multiplication is one such primitive task, occurring in many systems—from neural networks to scientific computing routines. The automatic discovery of algorithms using machine learning offers the prospect of reaching beyond human intuition and outperforming the current best human-designed algorithms. However, automating the algorithm discovery procedure is intricate, as the space of possible algorithms is enormous. Here we report a deep reinforcement learning approach based on AlphaZero[1] for discovering efficient and provably correct algorithms for the multiplication of arbitrary matrices. Our agent, AlphaTensor, is trained to play a single-player game where the objective is finding tensor decompositions within a finite factor space. AlphaTensor discovered algorithms that outperform the state-of-the-art complexity for many matrix sizes. Particularly relevant is the case of 4×4 matrices in a finite field, where AlphaTensor's algorithm improves on Strassen's two-level algorithm for the

[3]

Inspiration

nature

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature › articles › article

Article | Open Access | Published: 01 December 2021

Advancing mathematics by guiding human intuition with AI

Alex Davies ✉, Petar Veličković, Lars Buesing, Sam Blackwell, Daniel Zheng, Nenad Tomašev, Richard Tanburn, Peter Battaglia, Charles Blundell, András Juhász, Marc Lackenby, Geordie Williamson, Demis Hassabis & Pushmeet Kohli ✉

Nature **600**, 70–74 (2021) | Cite this article

Abstract

The practice of mathematics involves discovering patterns and using these to formulate and prove conjectures, resulting in theorems. Since the 1960s, mathematicians have used computers to assist in the discovery of patterns and formulation of conjectures[1], most famously in the Birch and Swinnerton-Dyer conjecture[2], a Millennium Prize Problem[3]. Here we provide examples of new fundamental results in pure mathematics that have been discovered with the assistance of machine learning—demonstrating a method by which machine learning can aid mathematicians in discovering new conjectures and theorems. We propose a process of using machine learning to discover potential patterns and relations between mathematical objects, understanding them with attribution techniques and using these observations to guide intuition and propose conjectures. We outline this machine-

[1]

Idea from [1]

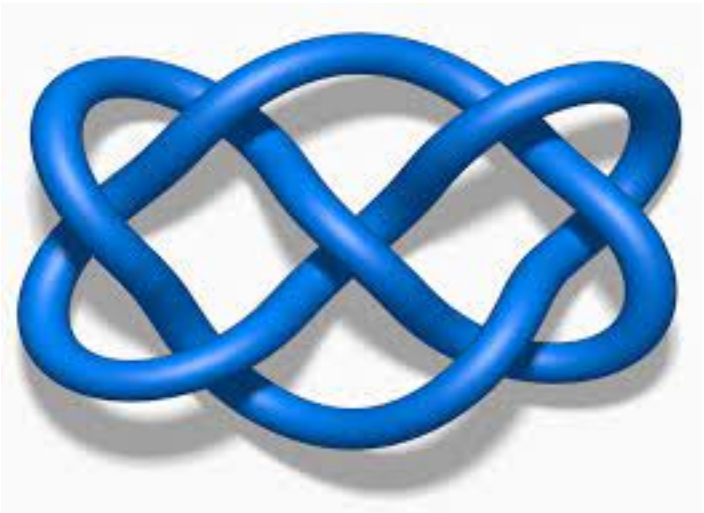

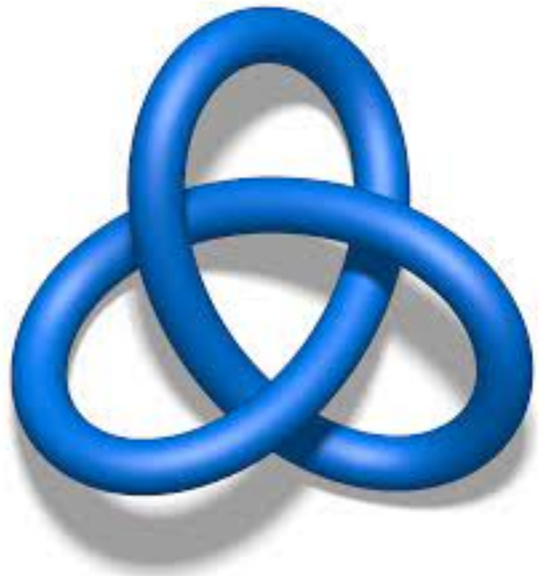

Idea from [1]

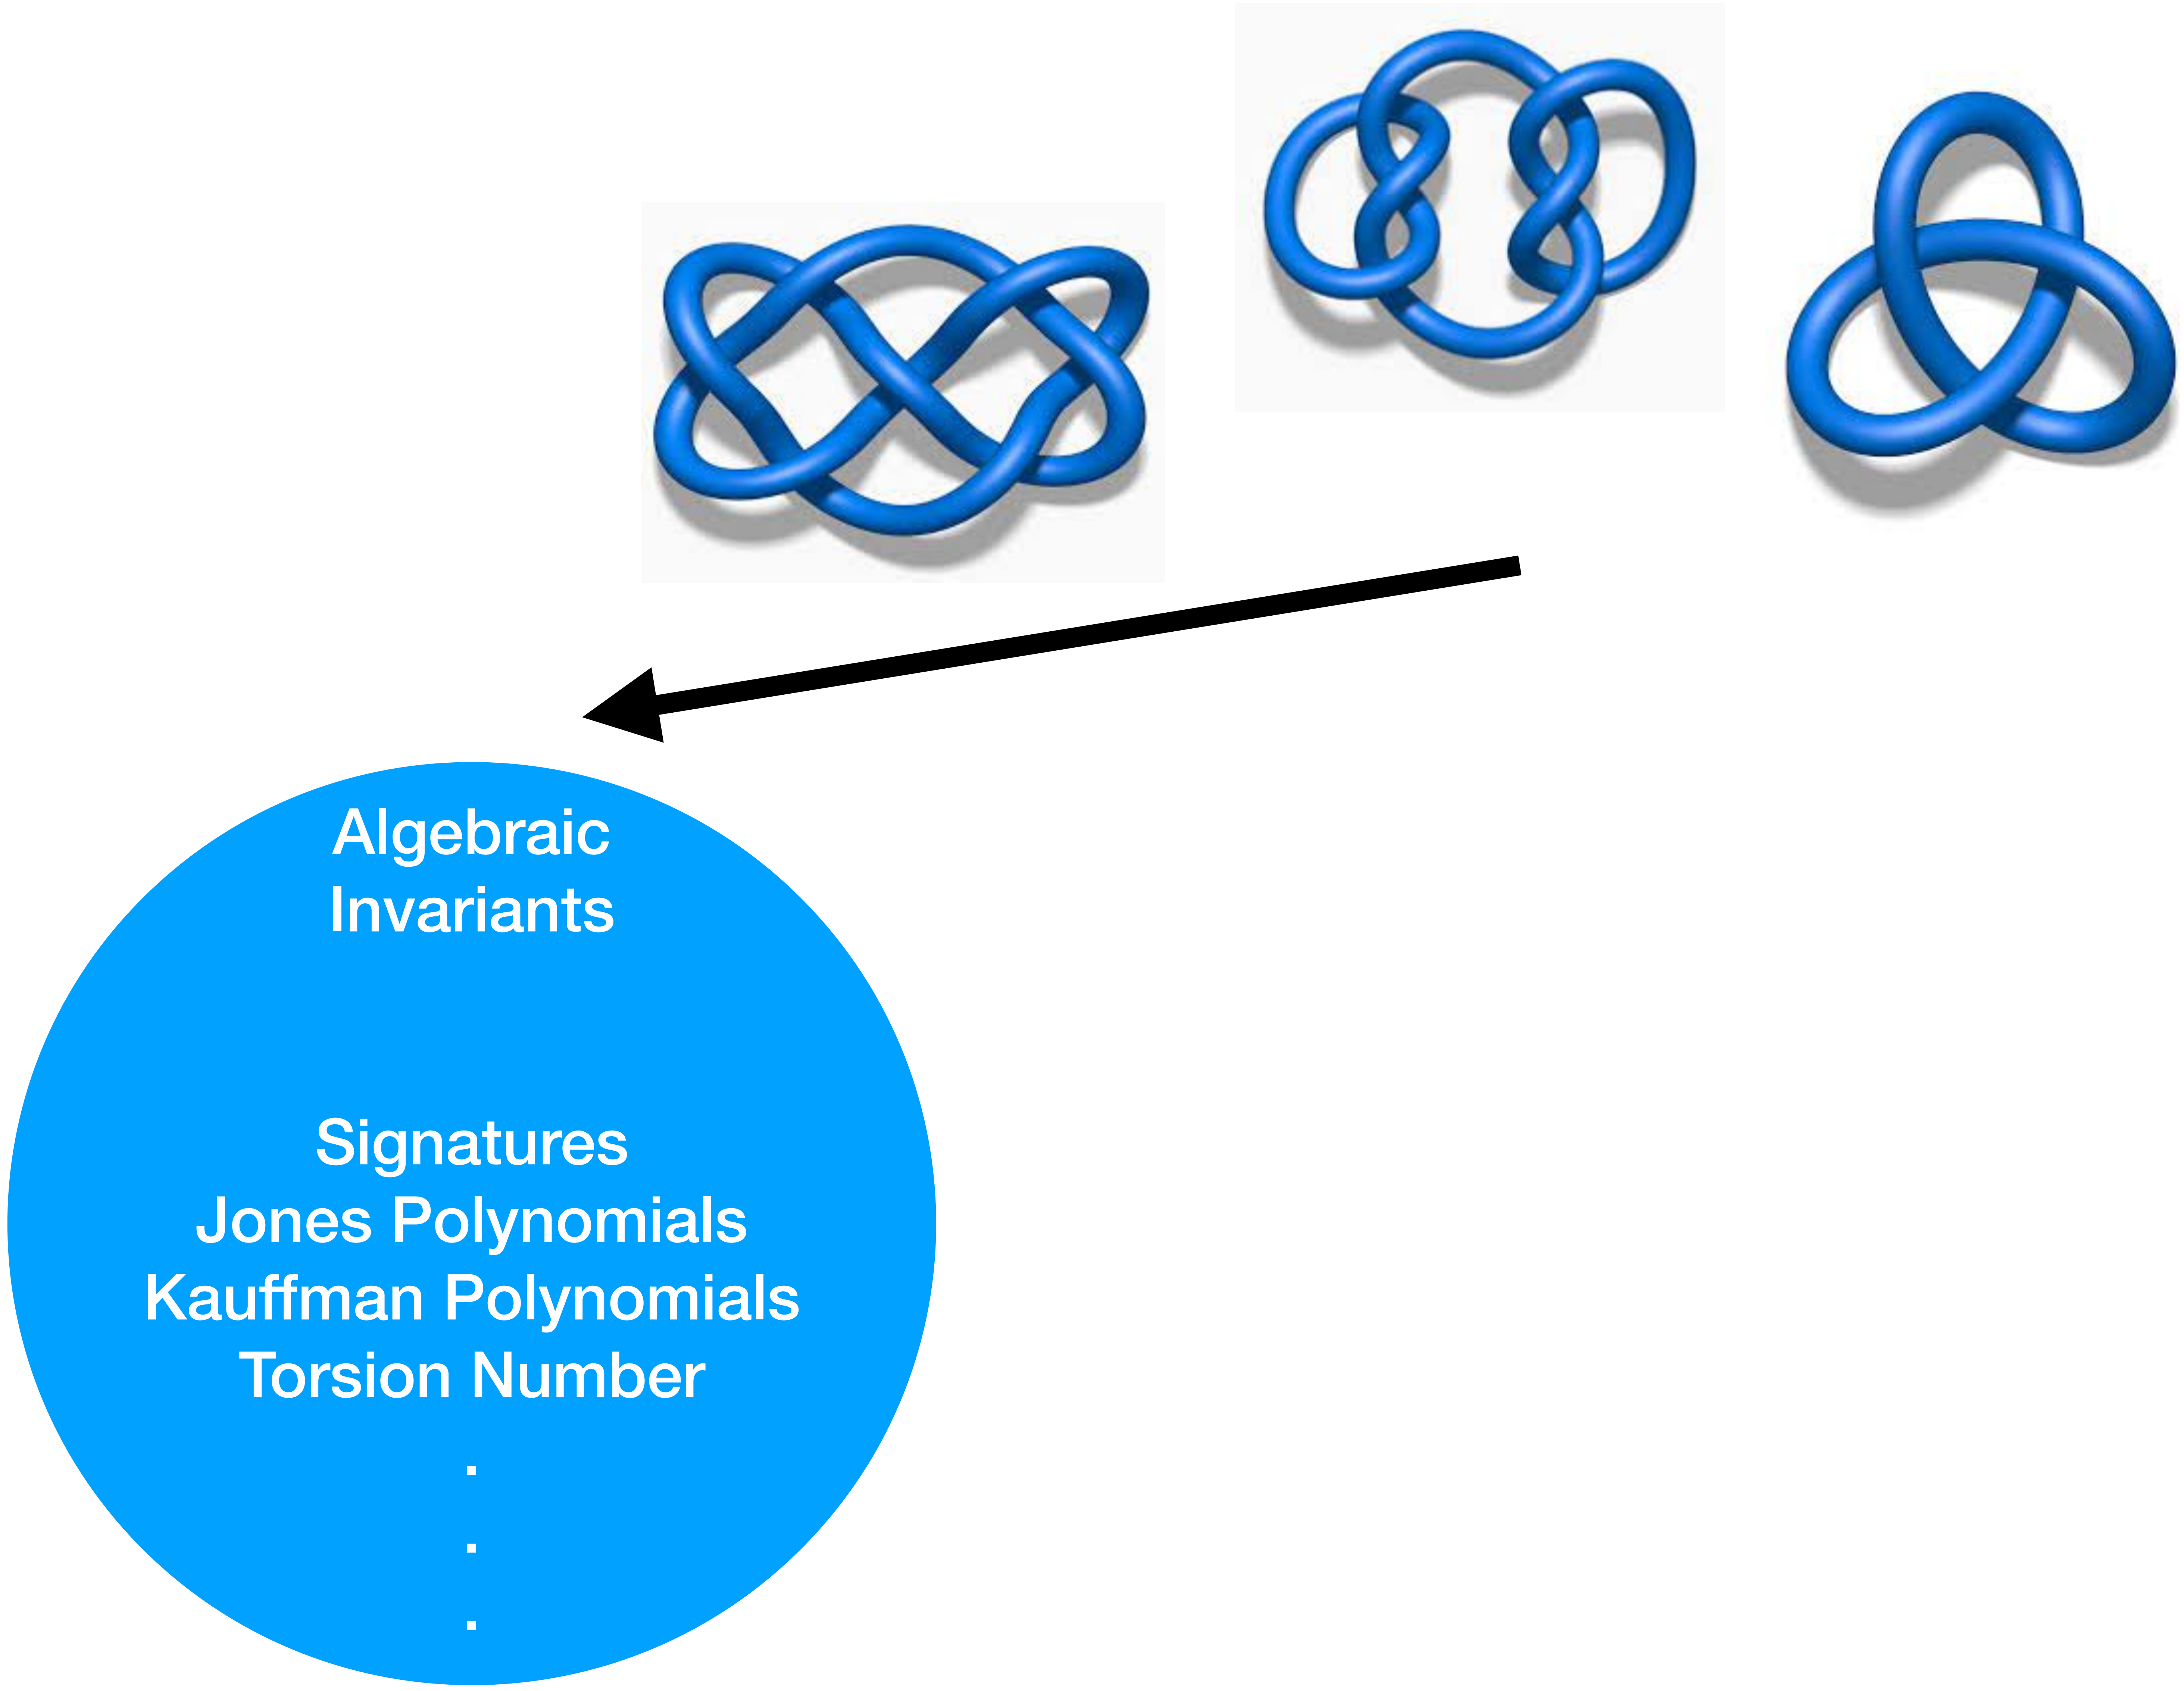
Algebraic
Invariants
Signatures
Jones Polynomials
Kauffman Polynomials
Torsion Number
.
.
.

Idea from [1]

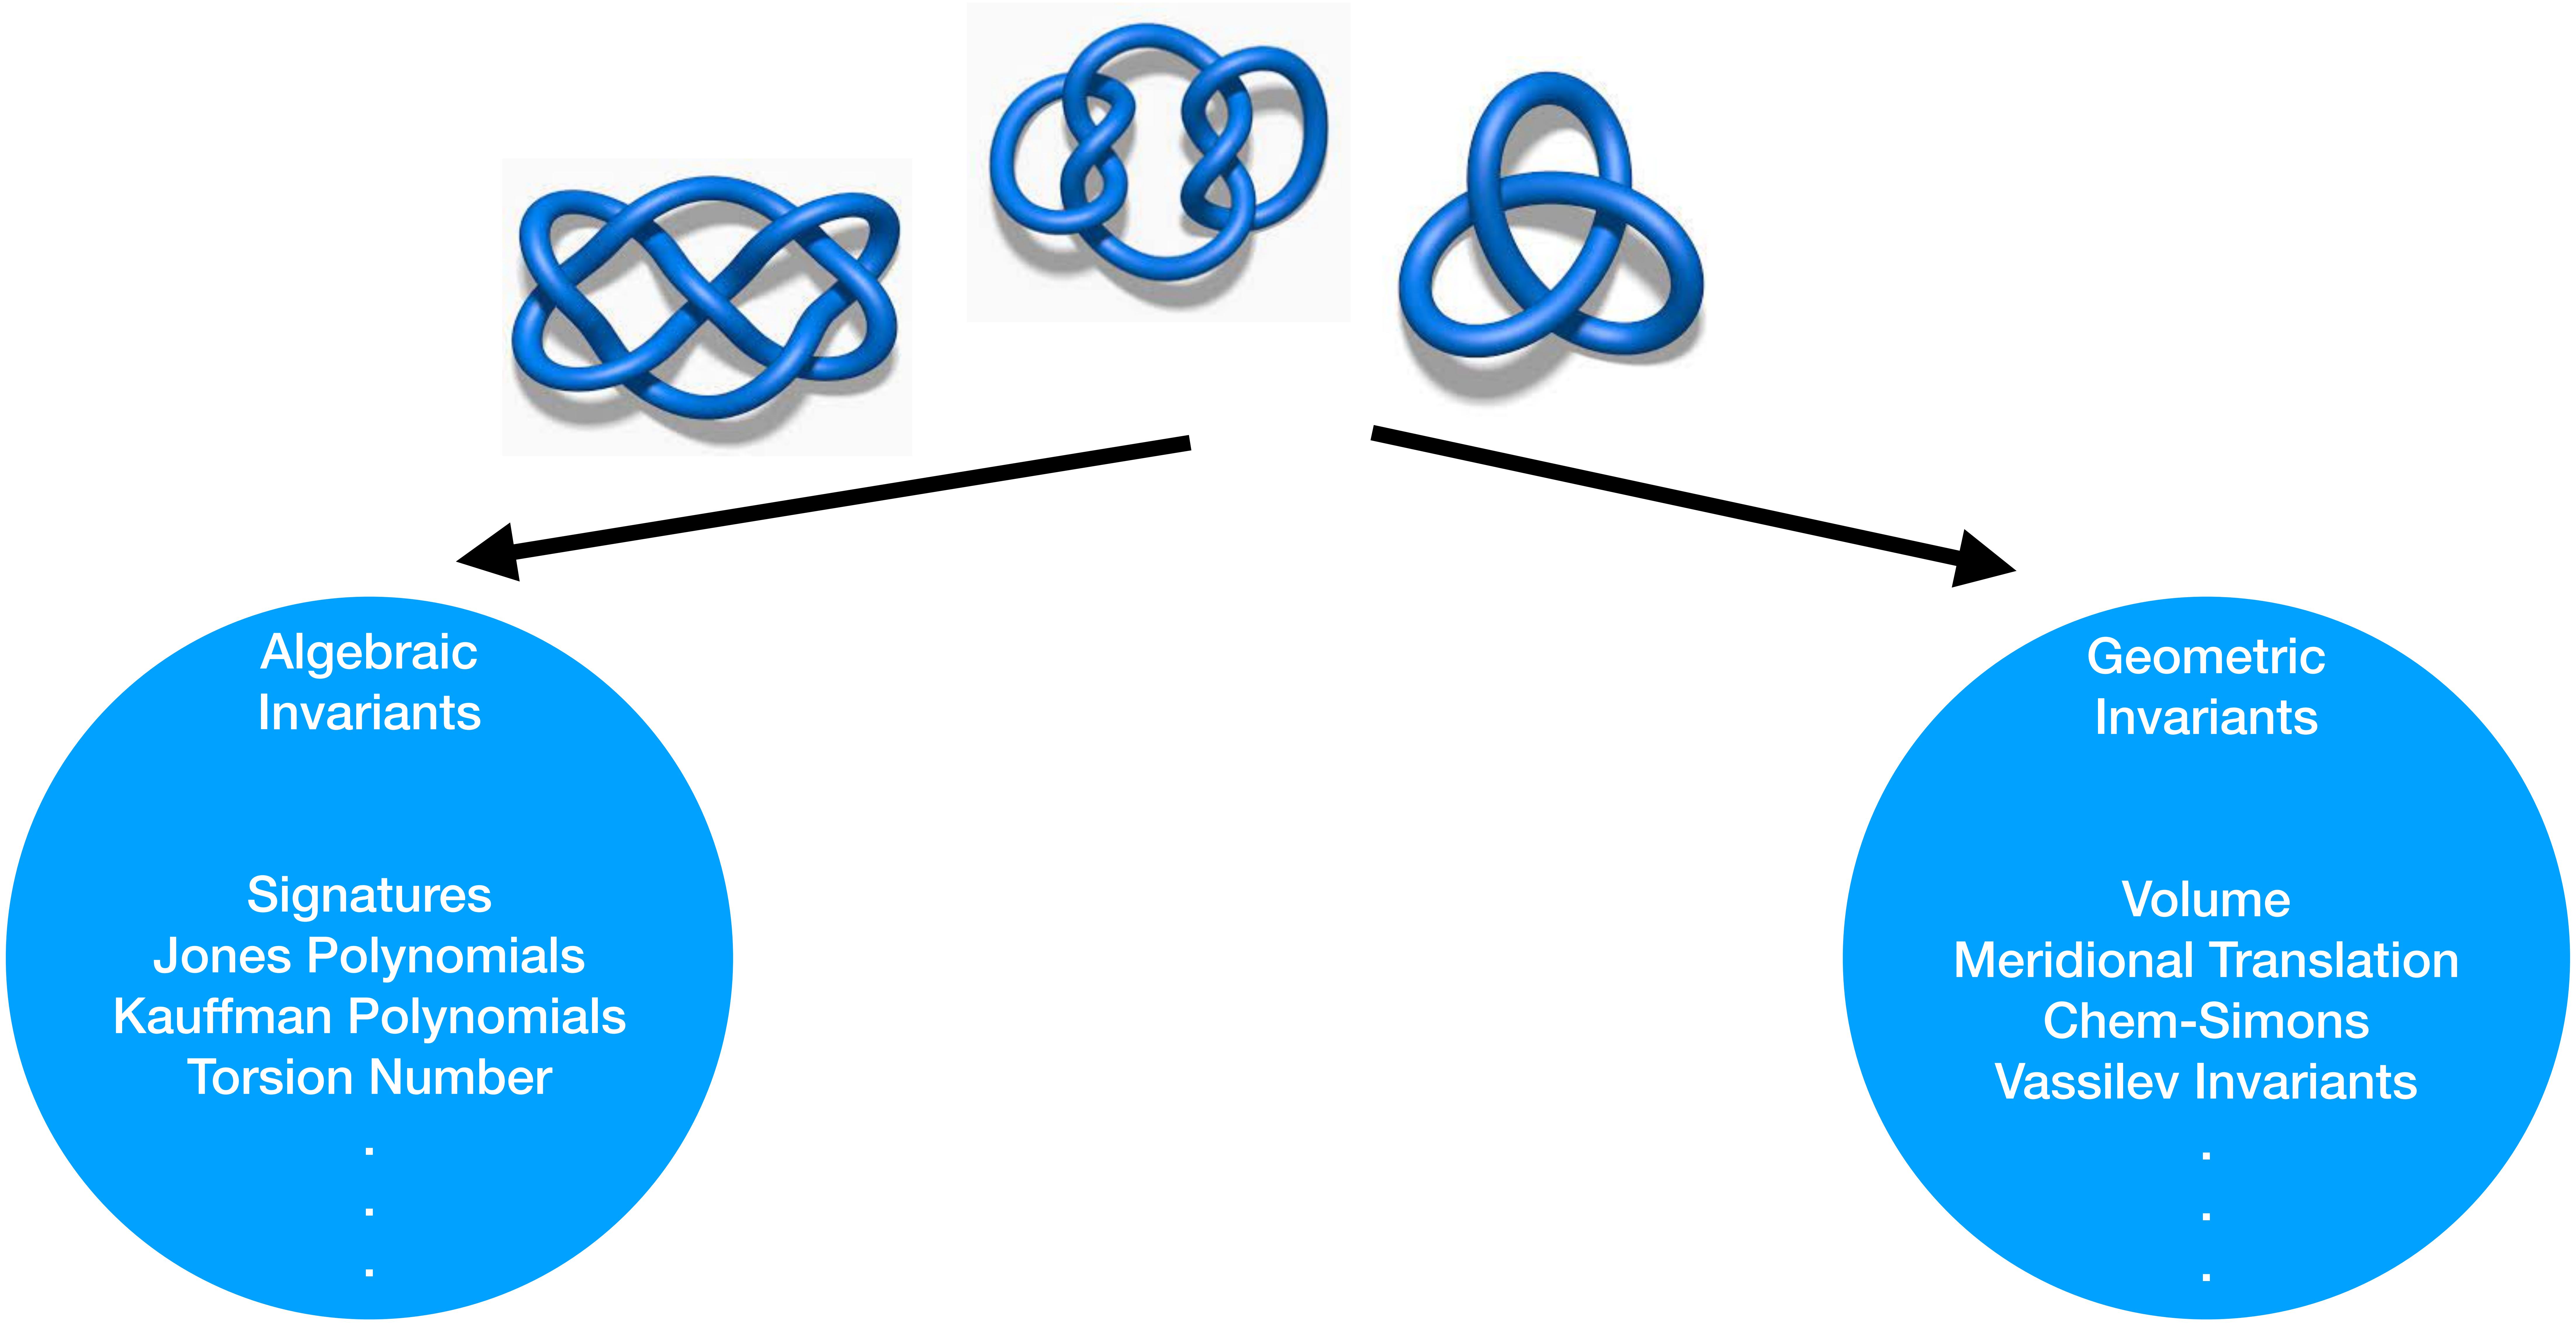
Algebraic
Invariants
Signatures
Jones Polynomials
Kauffman Polynomials
Torsion Number
.
.
.
Geometric
Invariants
Volume
Meridional Translation
Chem-Simons
Vassilev Invariants
.
.
.

Idea from [1]

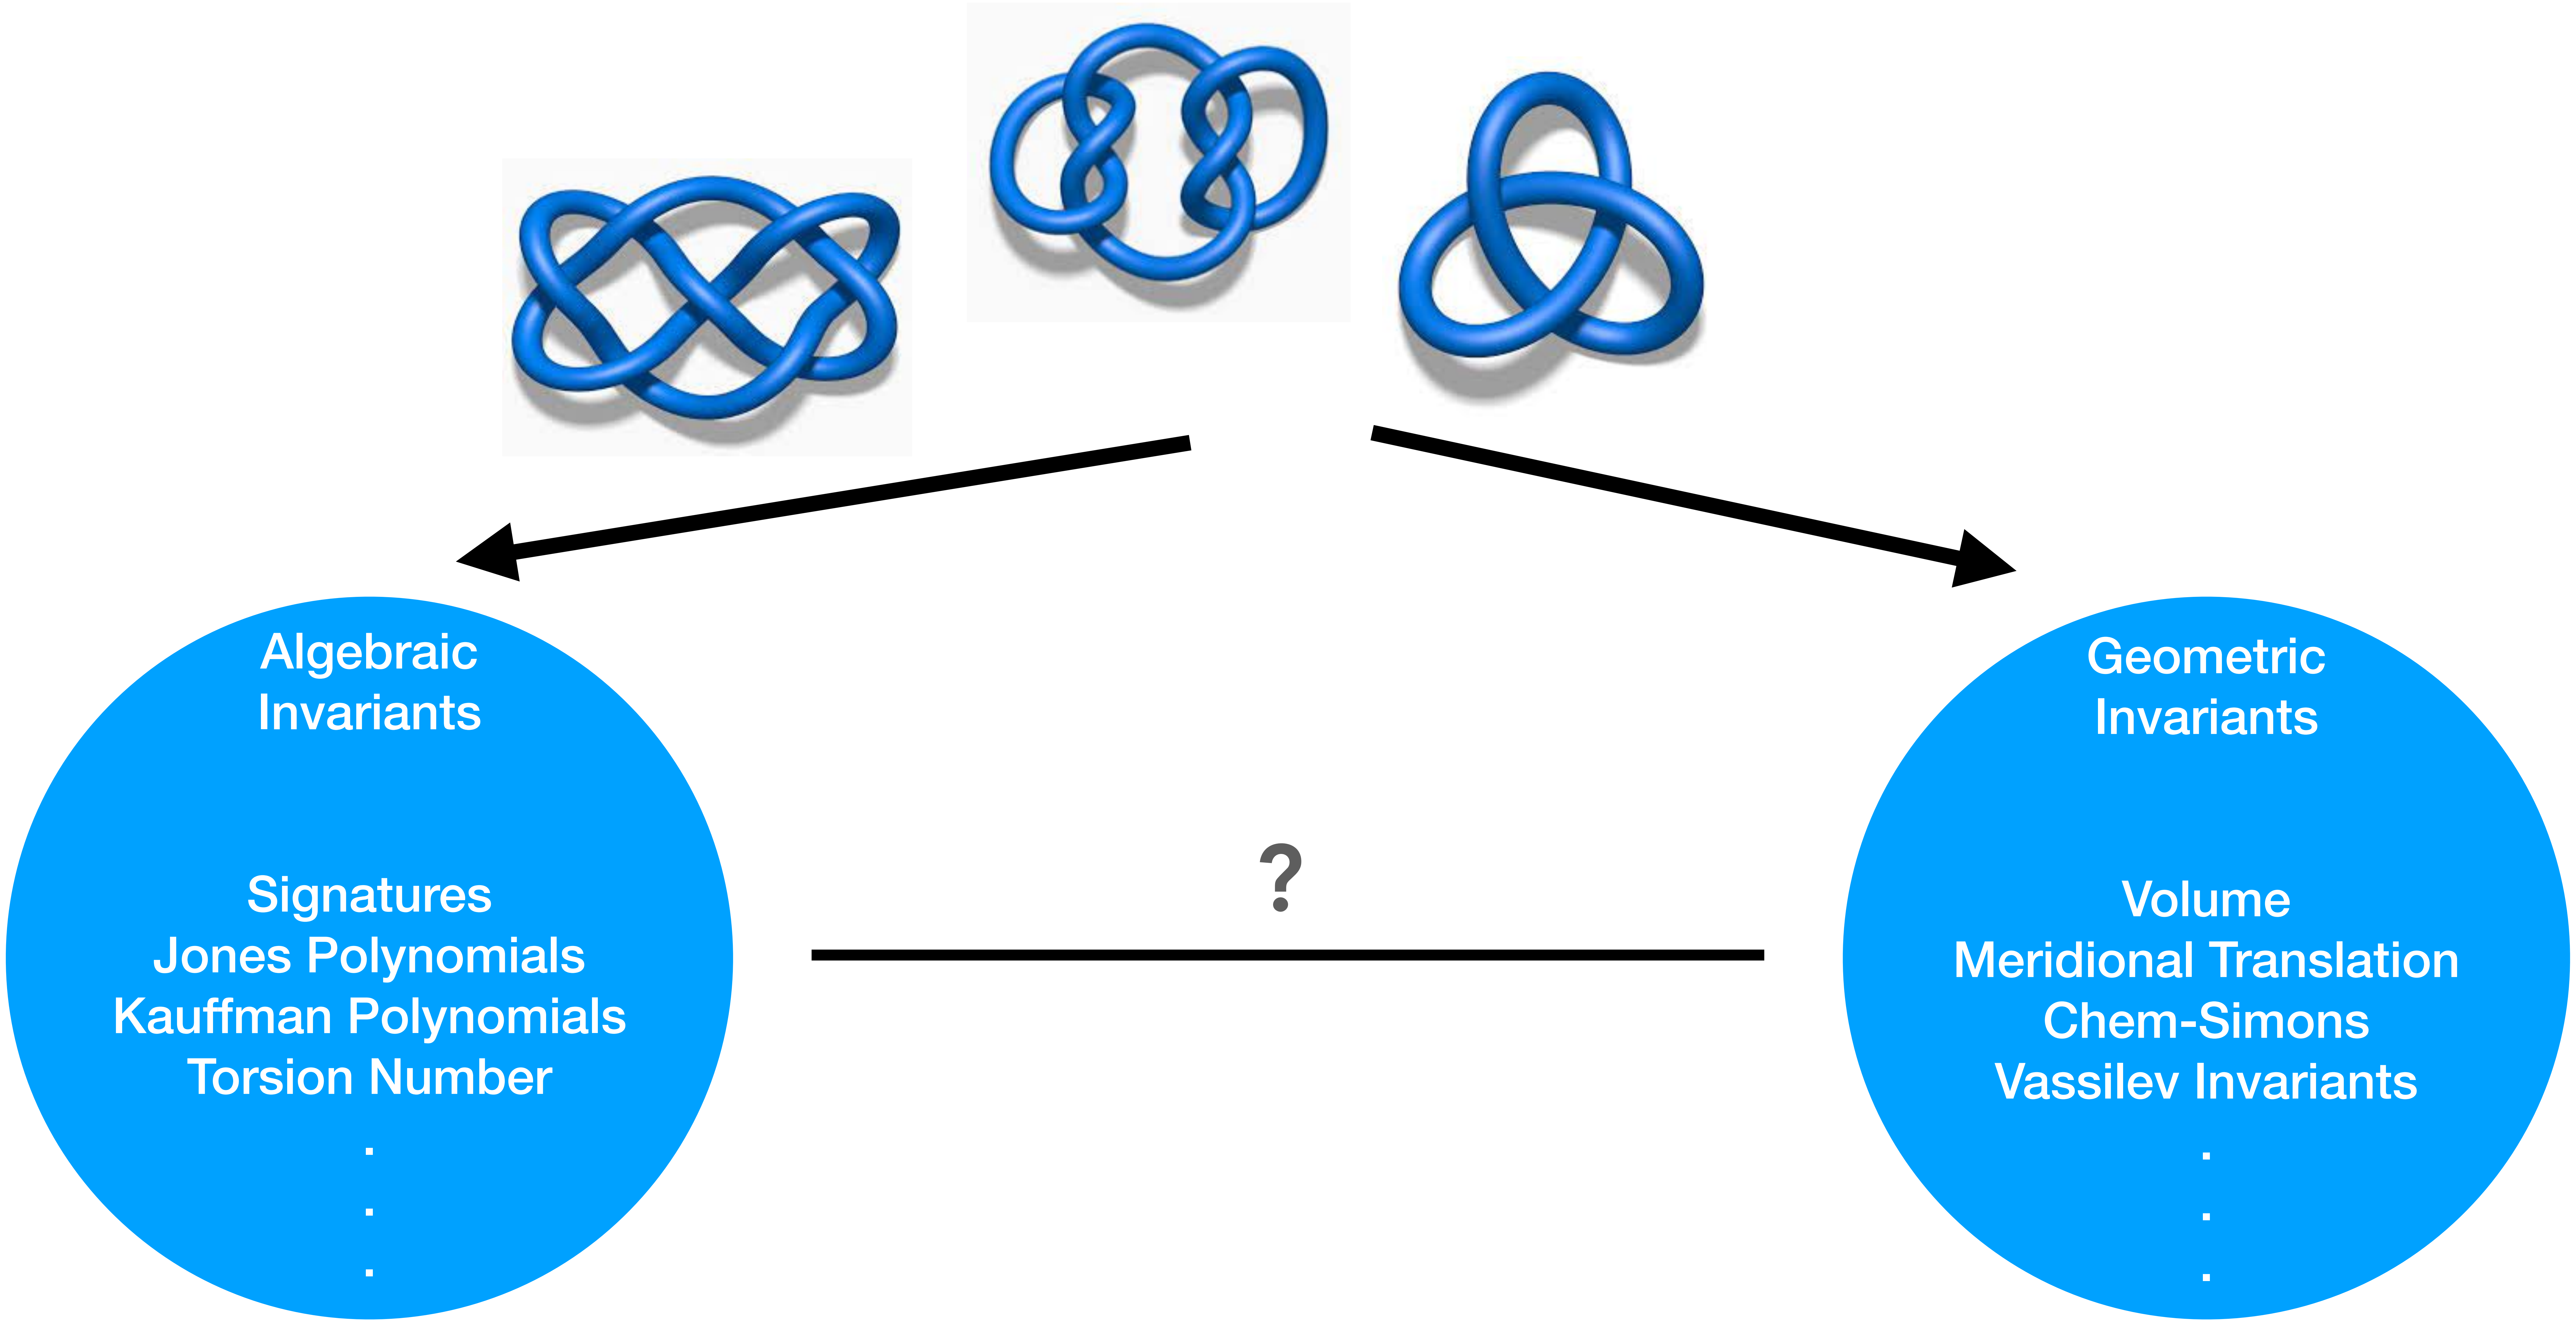

Algebraic Invariants
Signatures
Jones Polynomials
Kauffman Polynomials
Torsion Number
.
.
.
?
Geometric Invariants
Volume
Meridional Translation
Chem-Simons
Vassilev Invariants
.
.
.

Key Idea

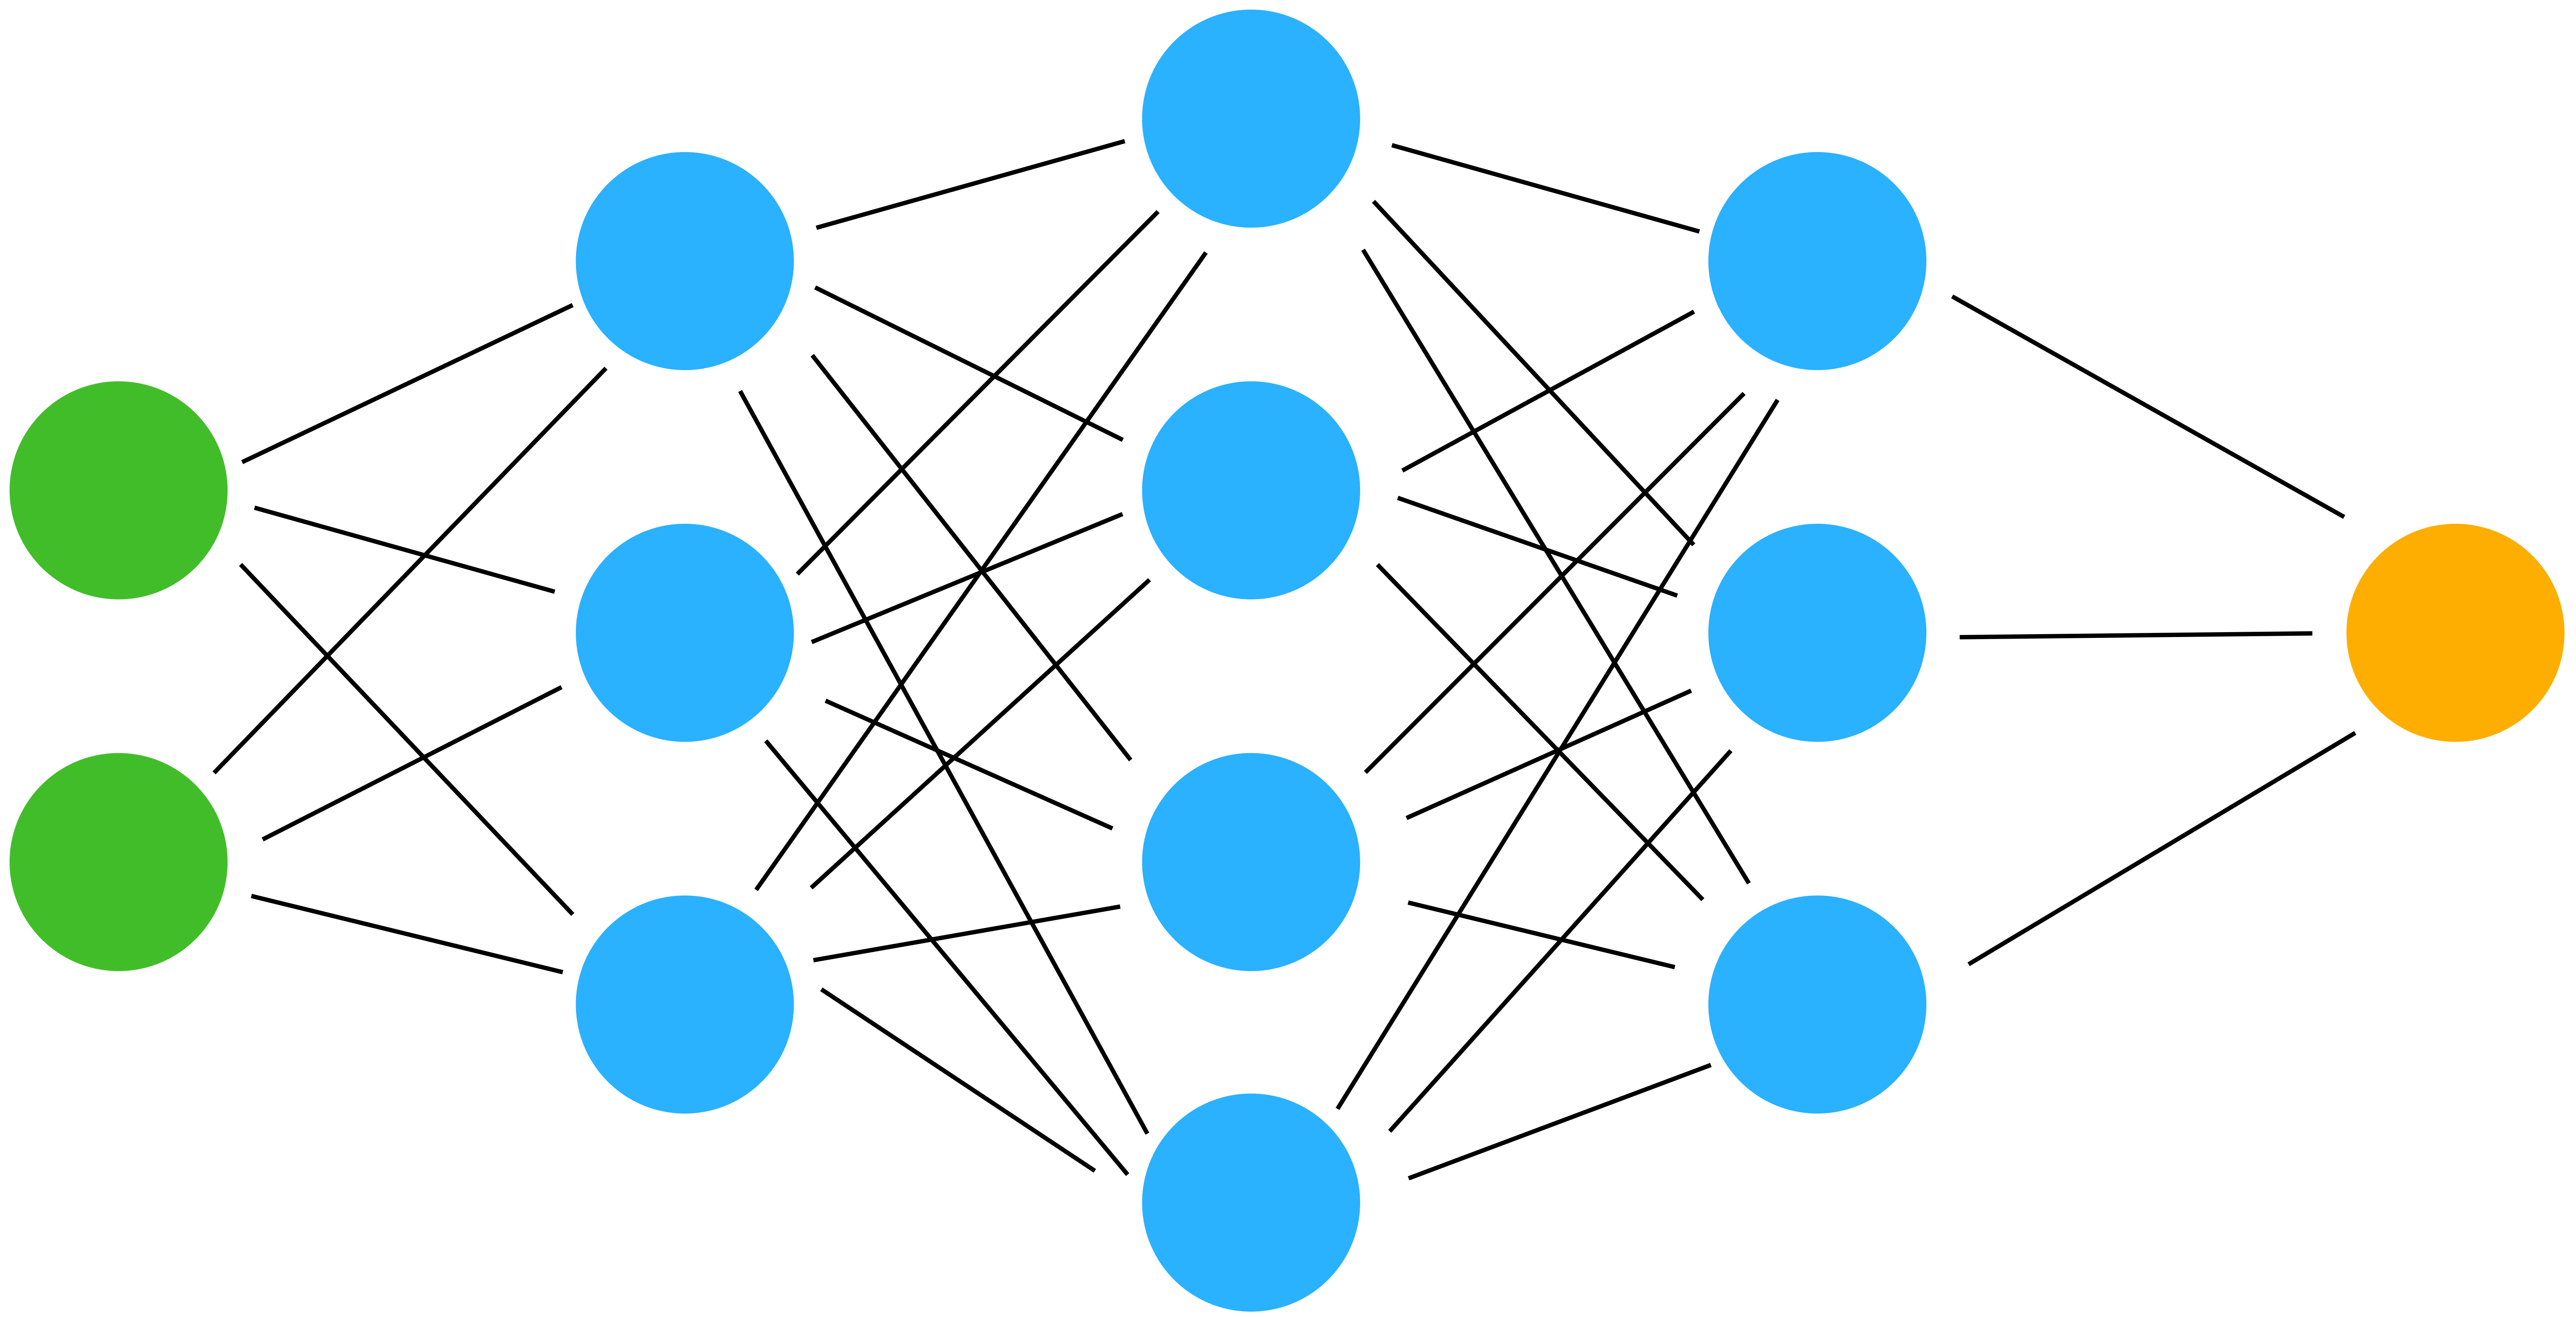

Key Idea

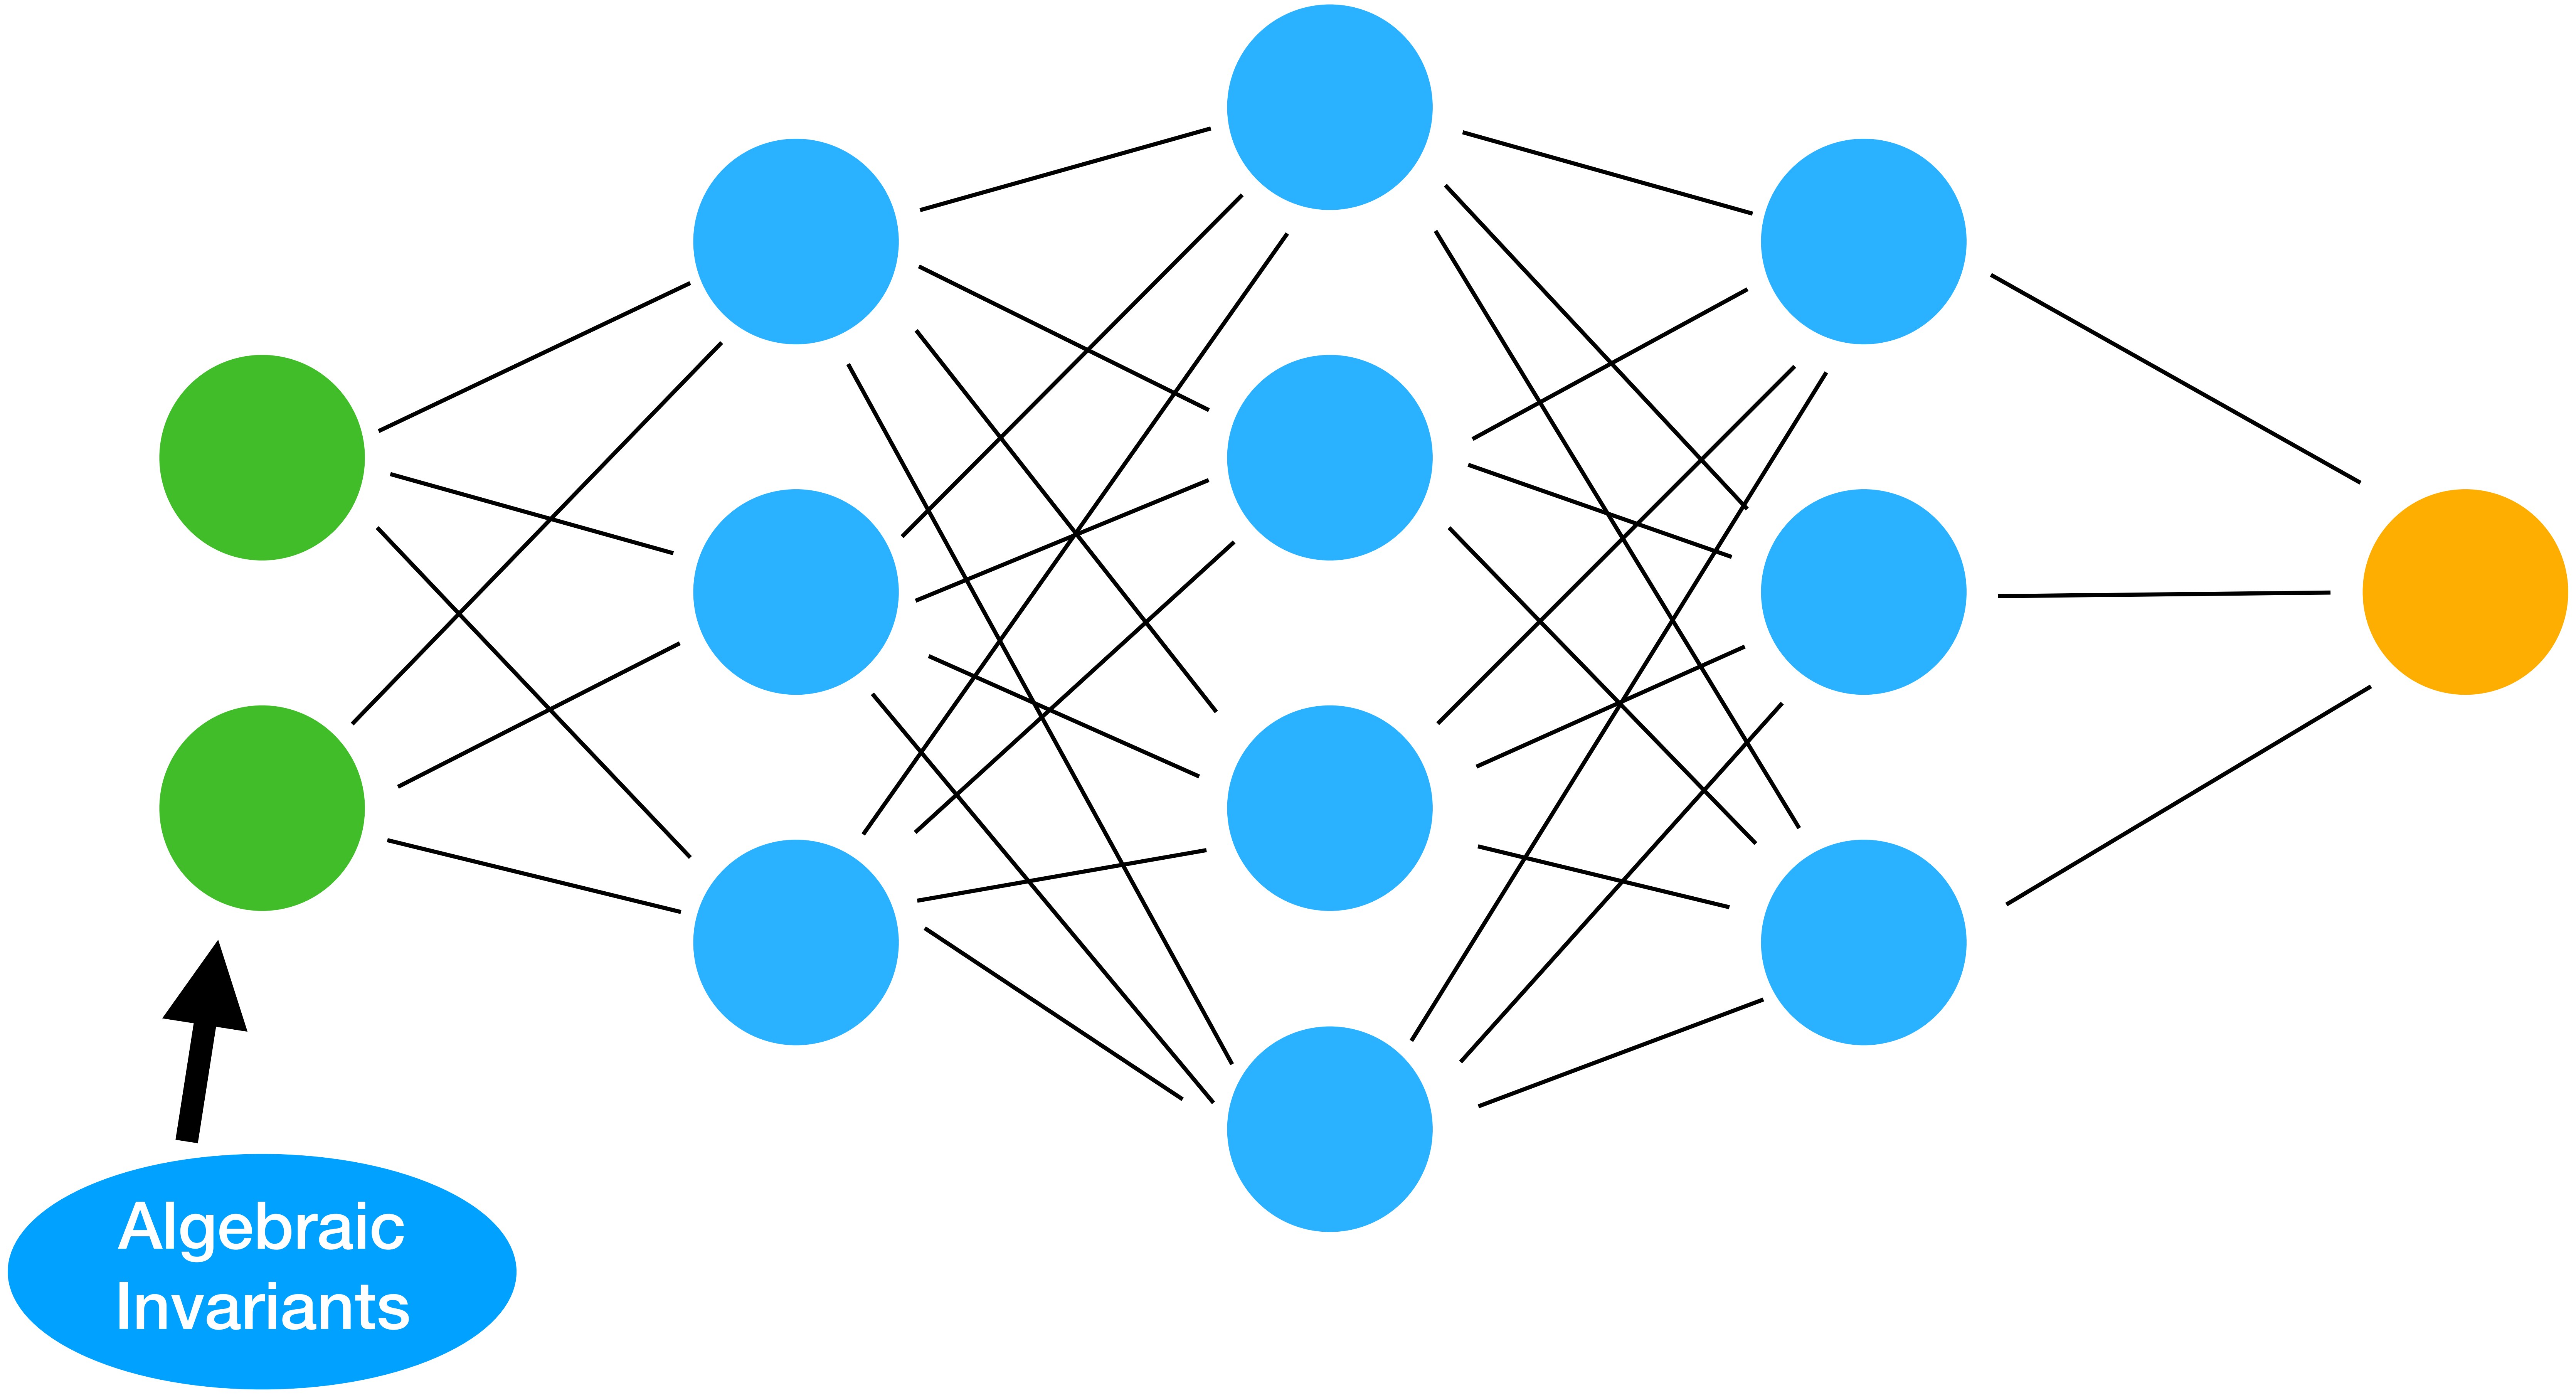

Key Idea

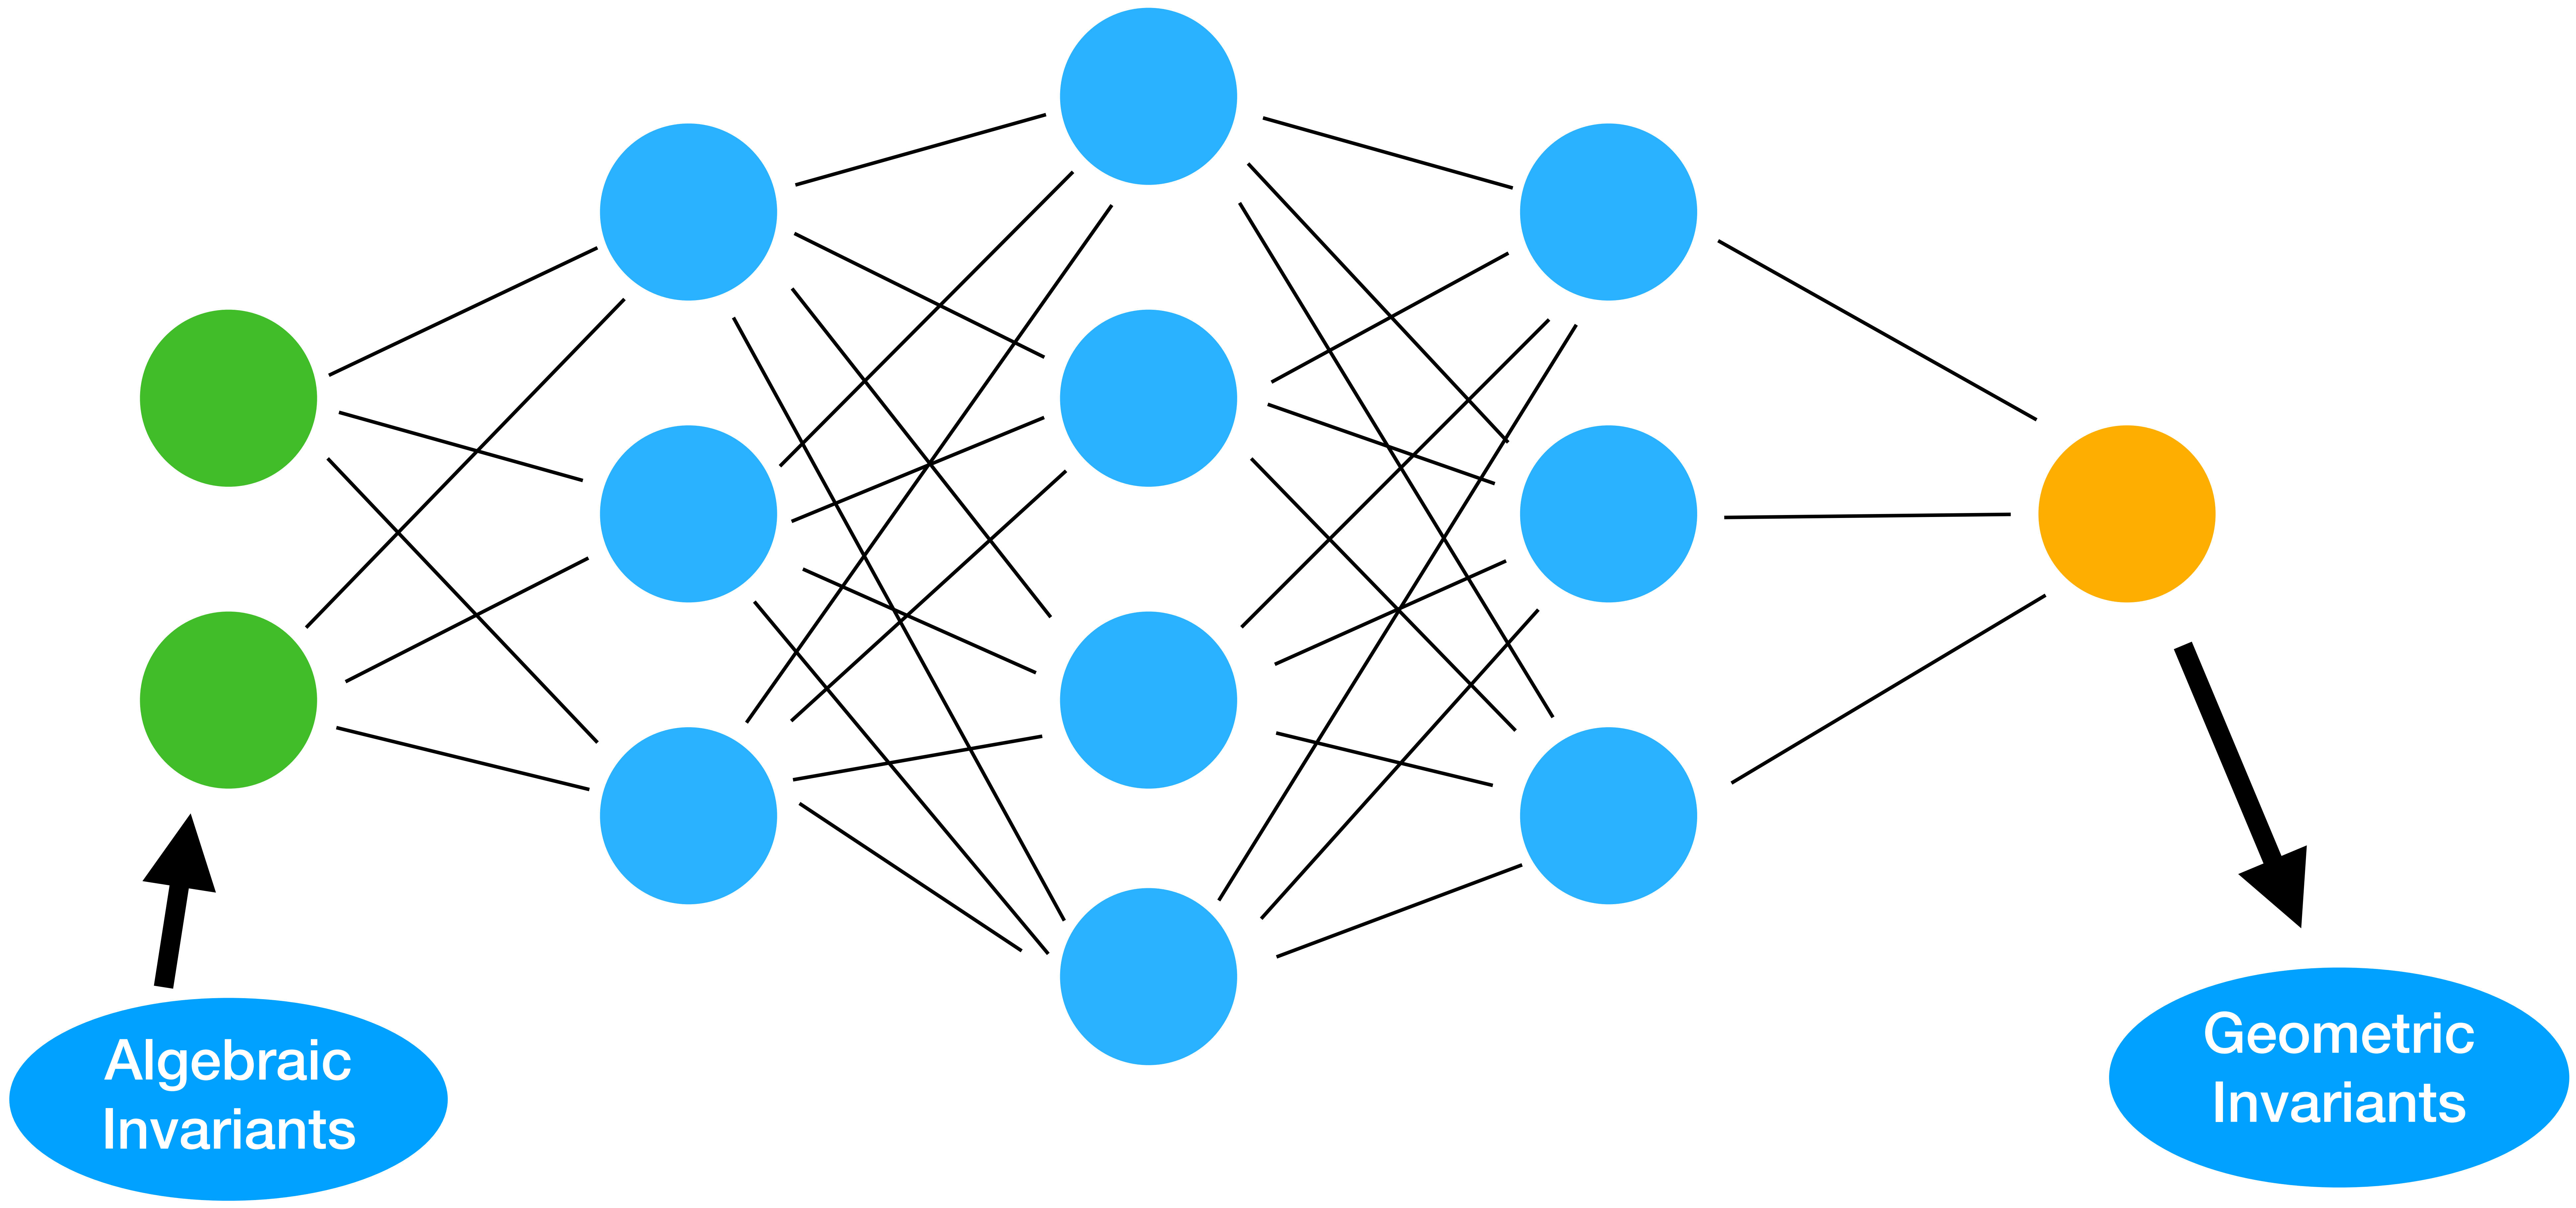

Key Idea

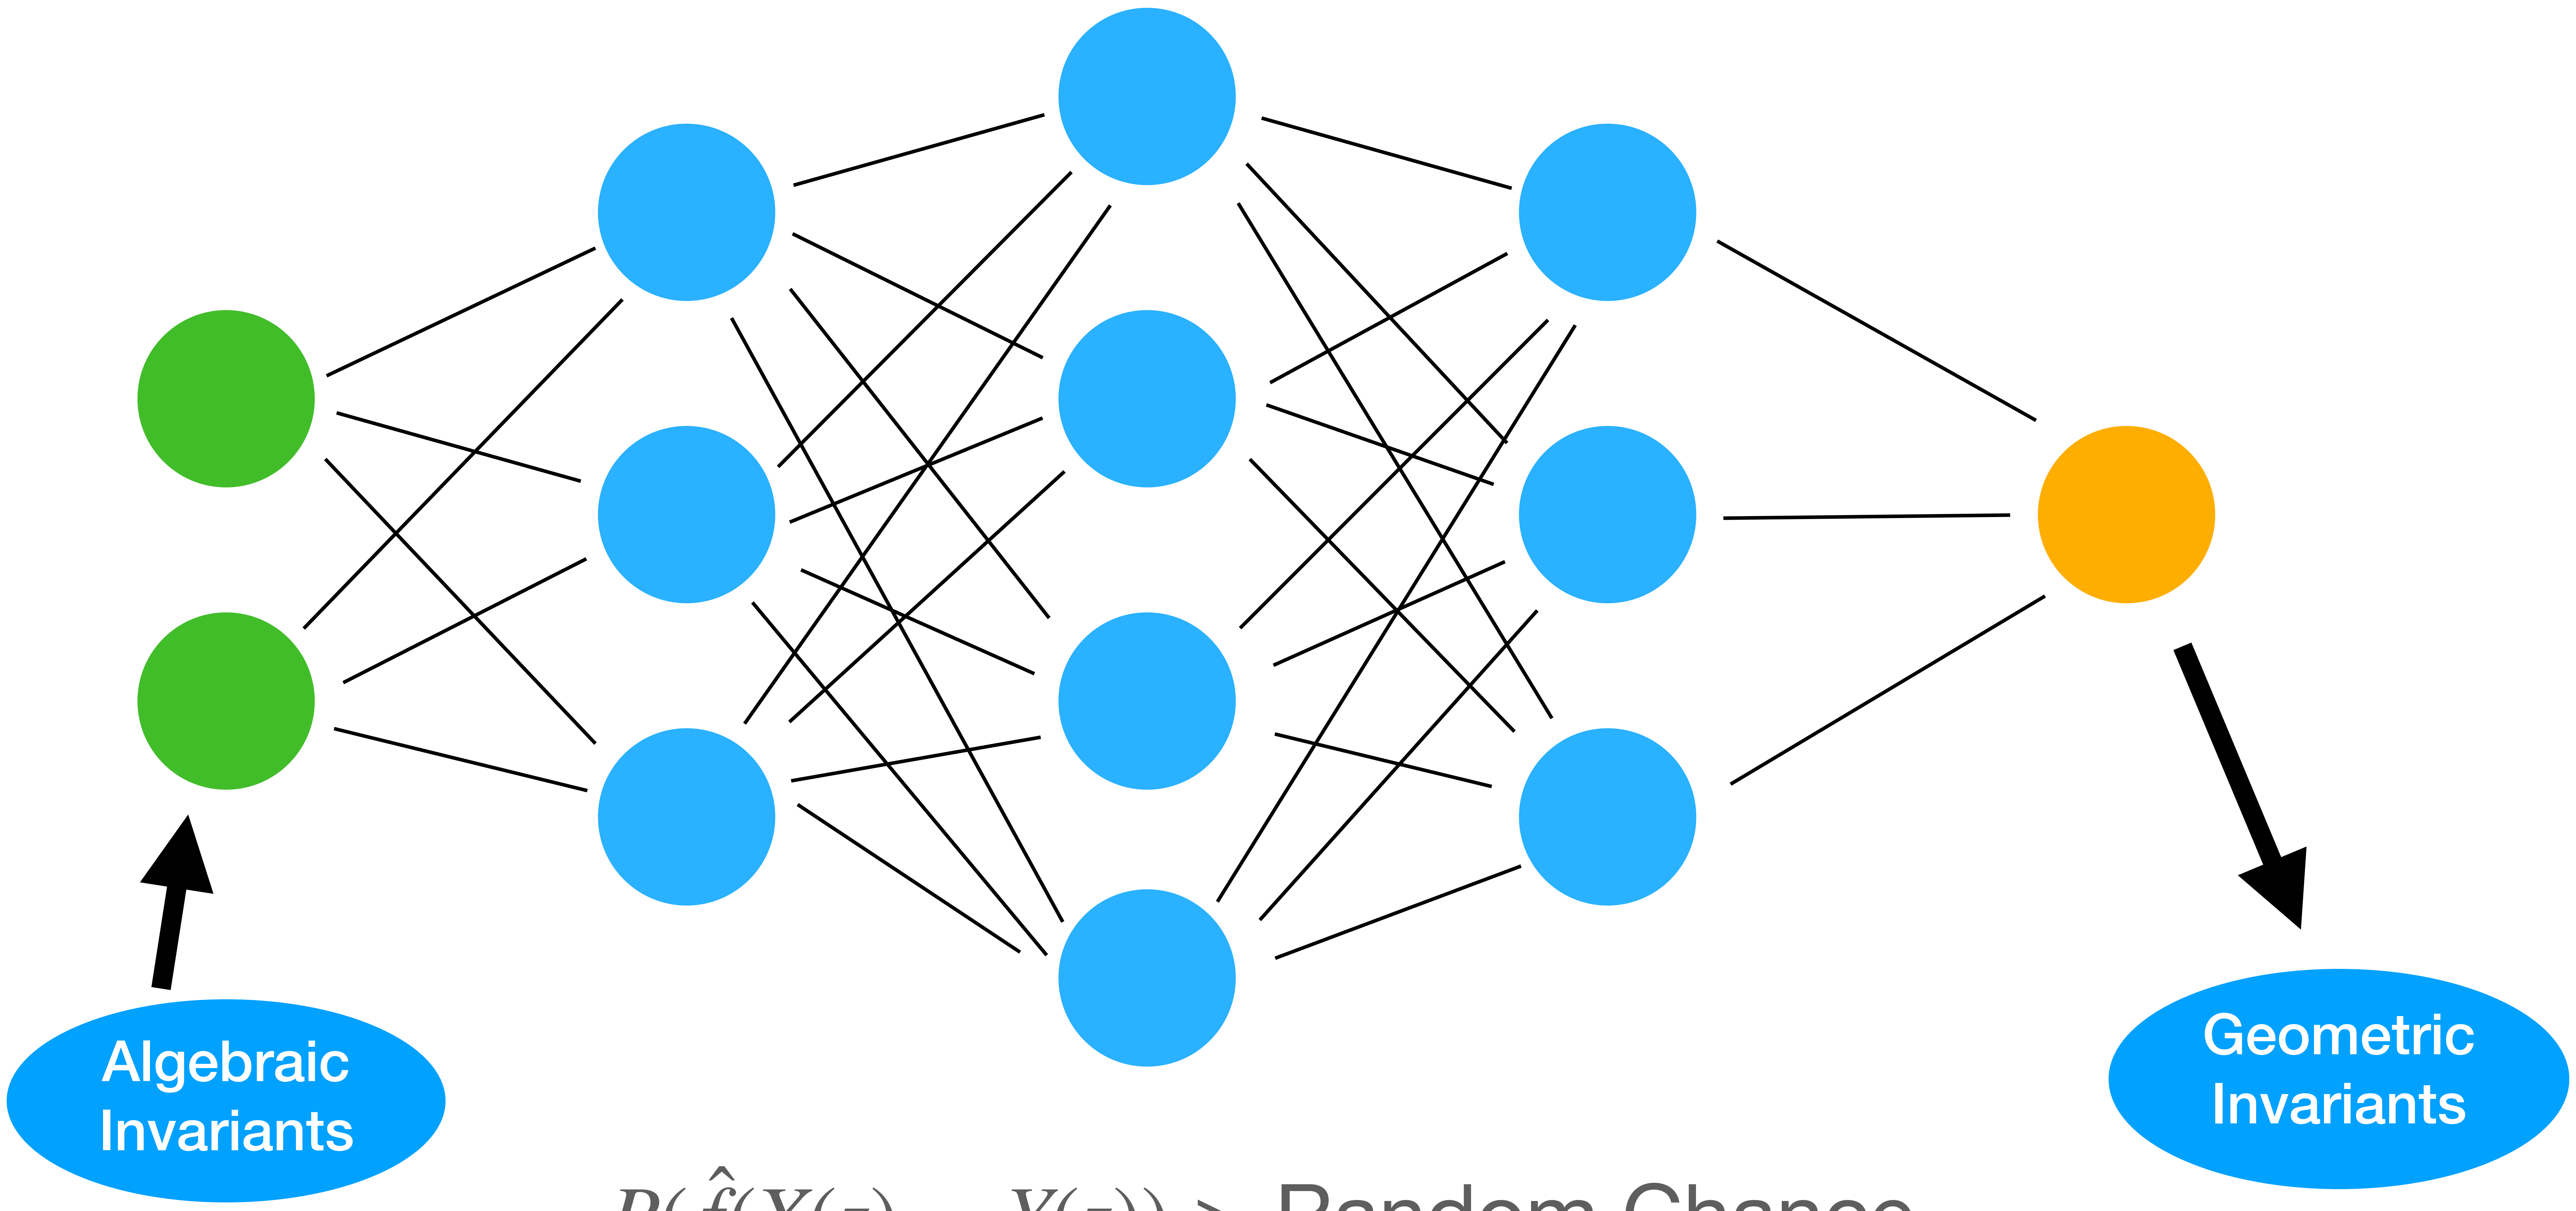

$P(\hat{f}(X(z) = Y(z)) >$ Random Chance

Idea from [1]

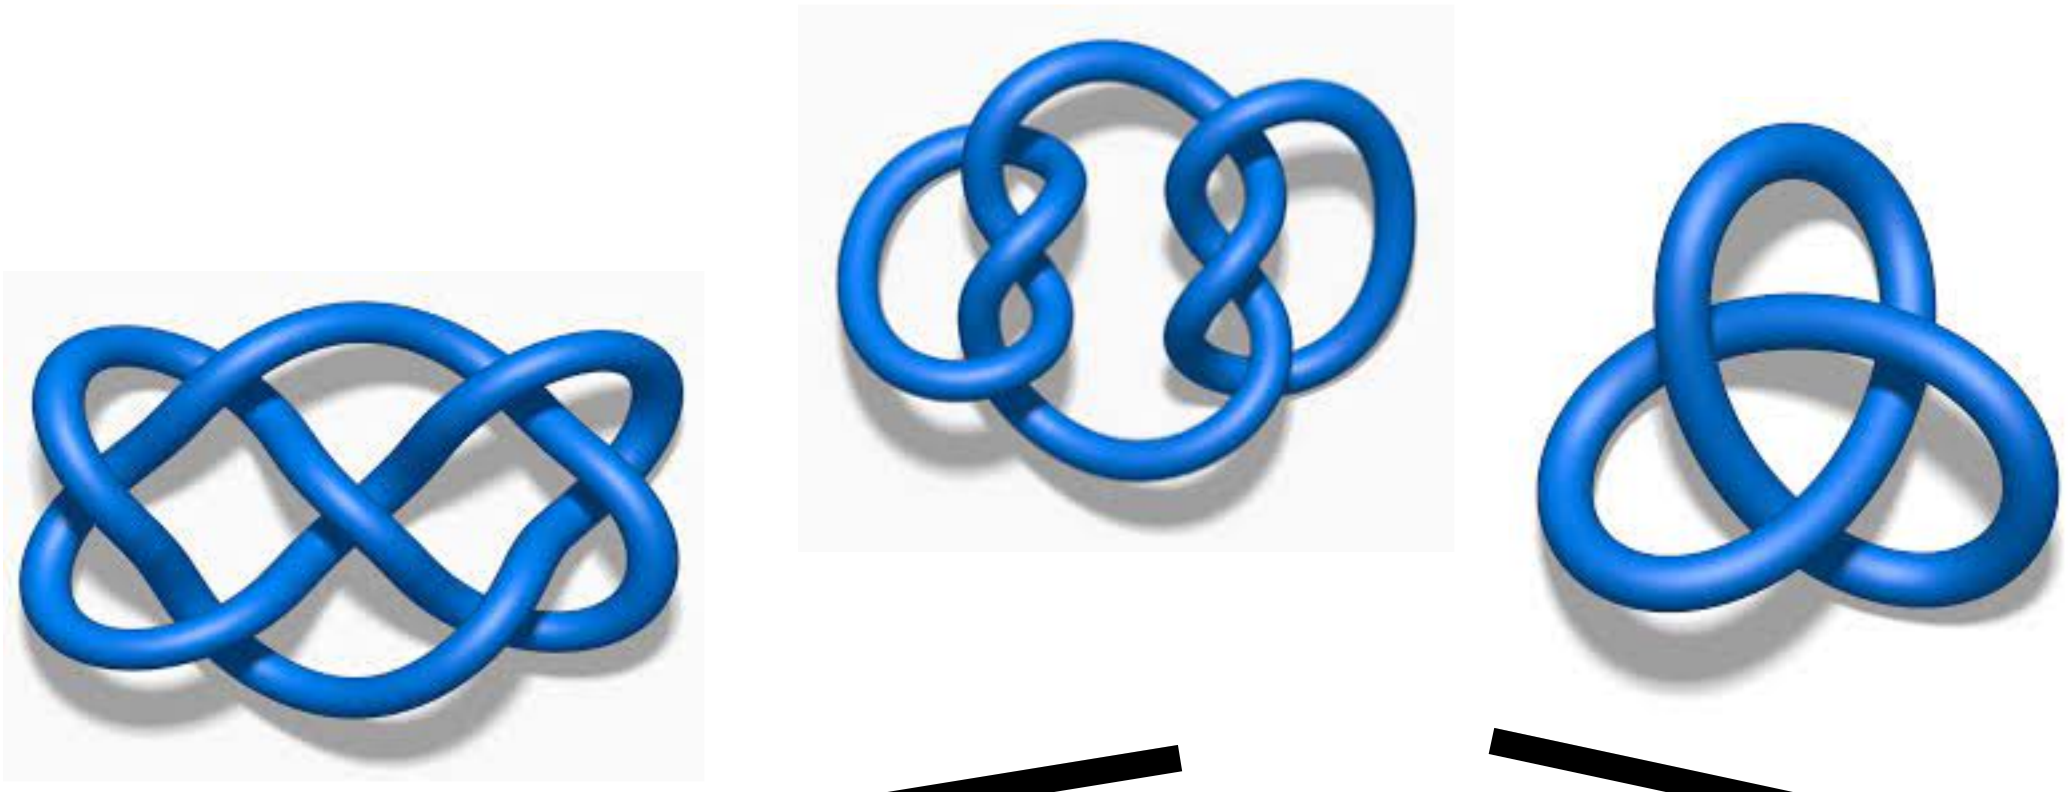

Algebraic Invariants

Signatures
Jones Polynomials
Kauffman Polynomials
Torsion Number
.
.
.

Geometric Invariants

Volume
Meridional Translation
Chem-Simons
Vassilev Invariants
.
.
.

Their view of the future

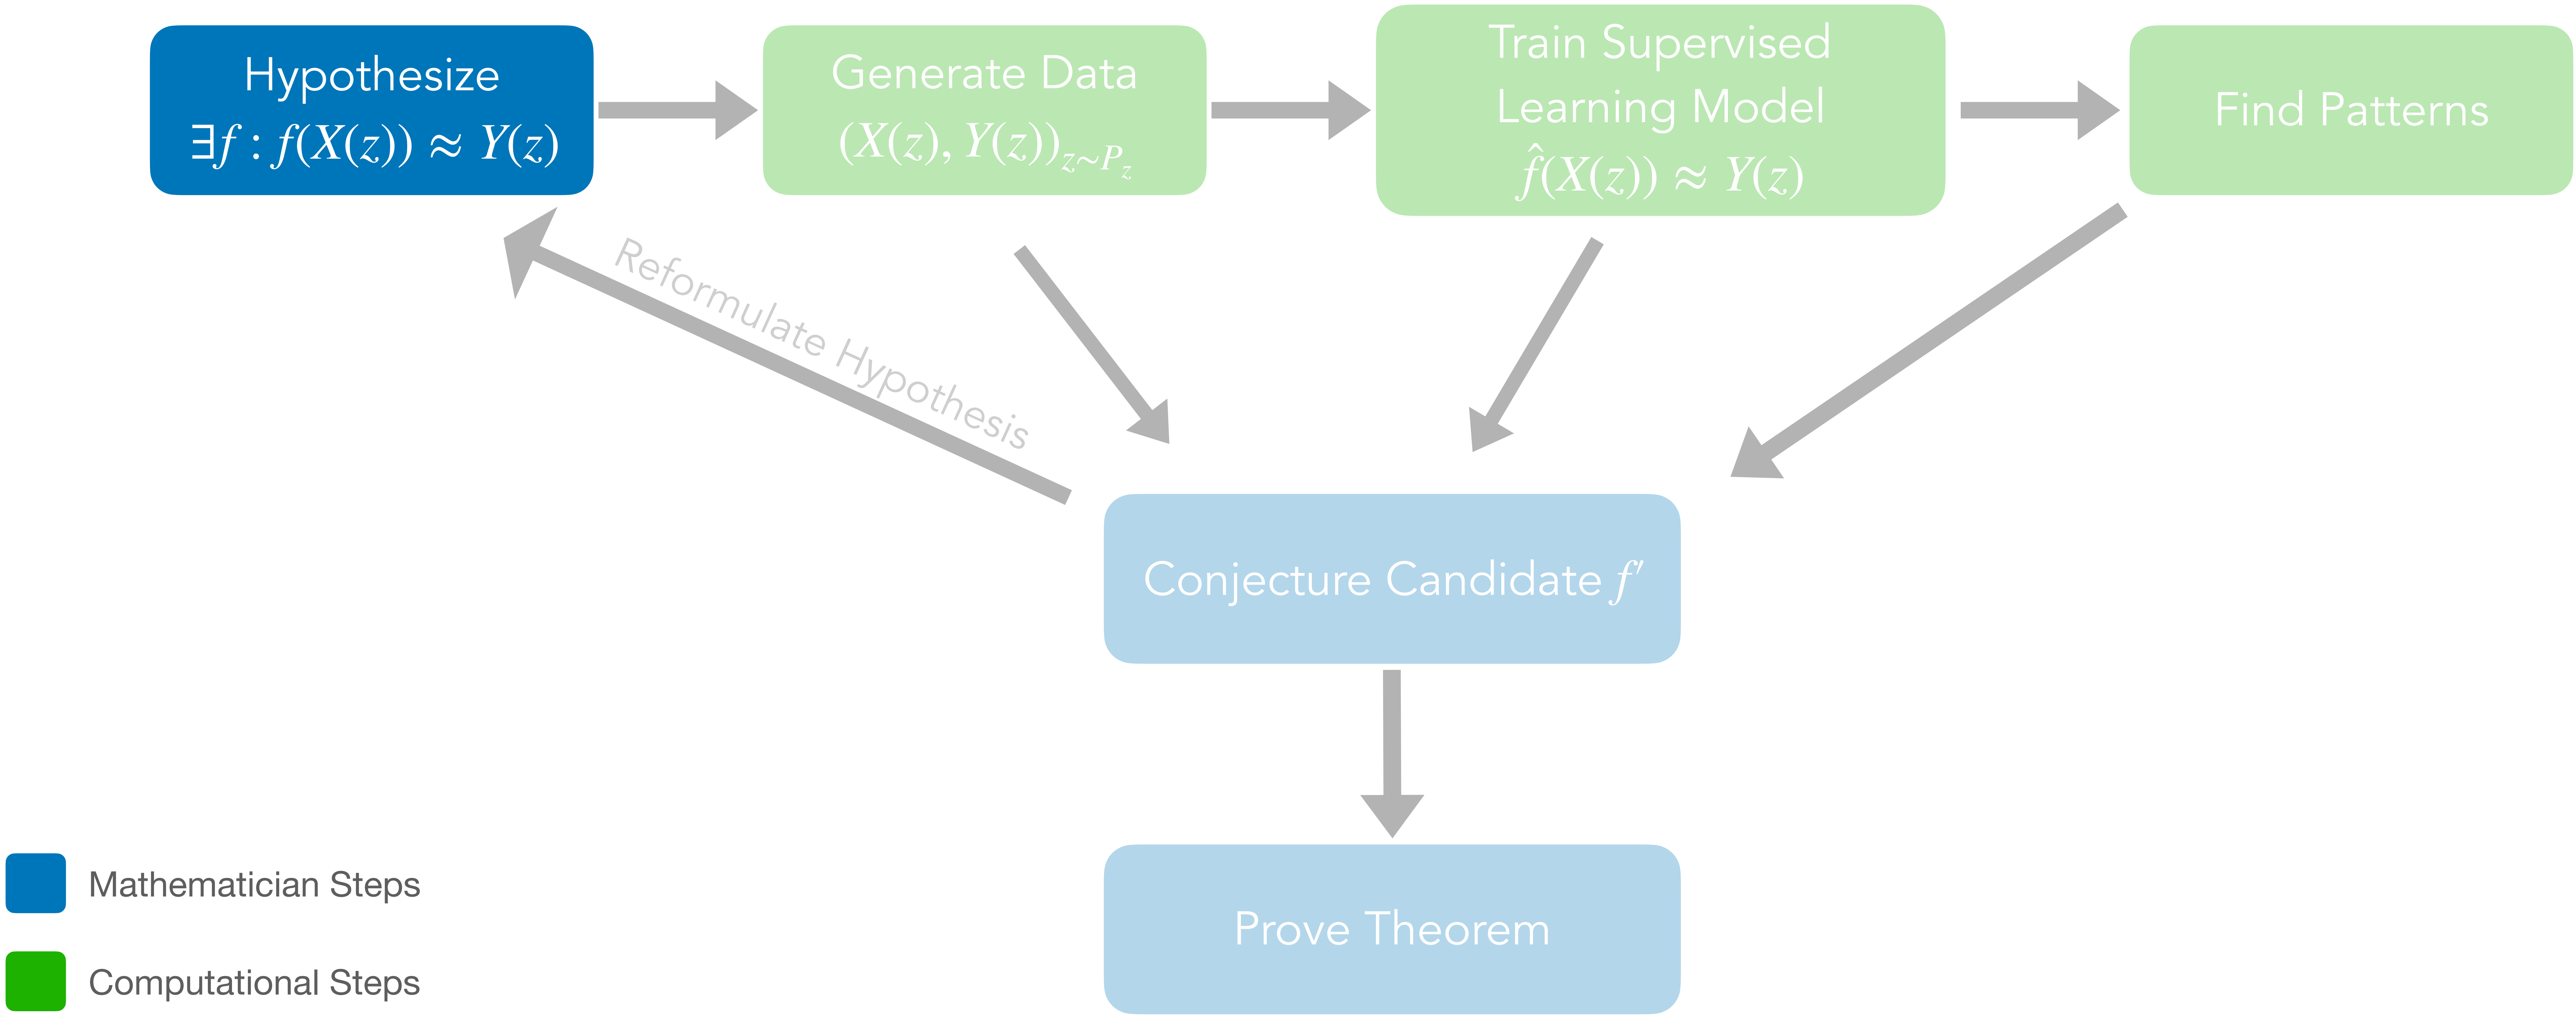

Their view of the future

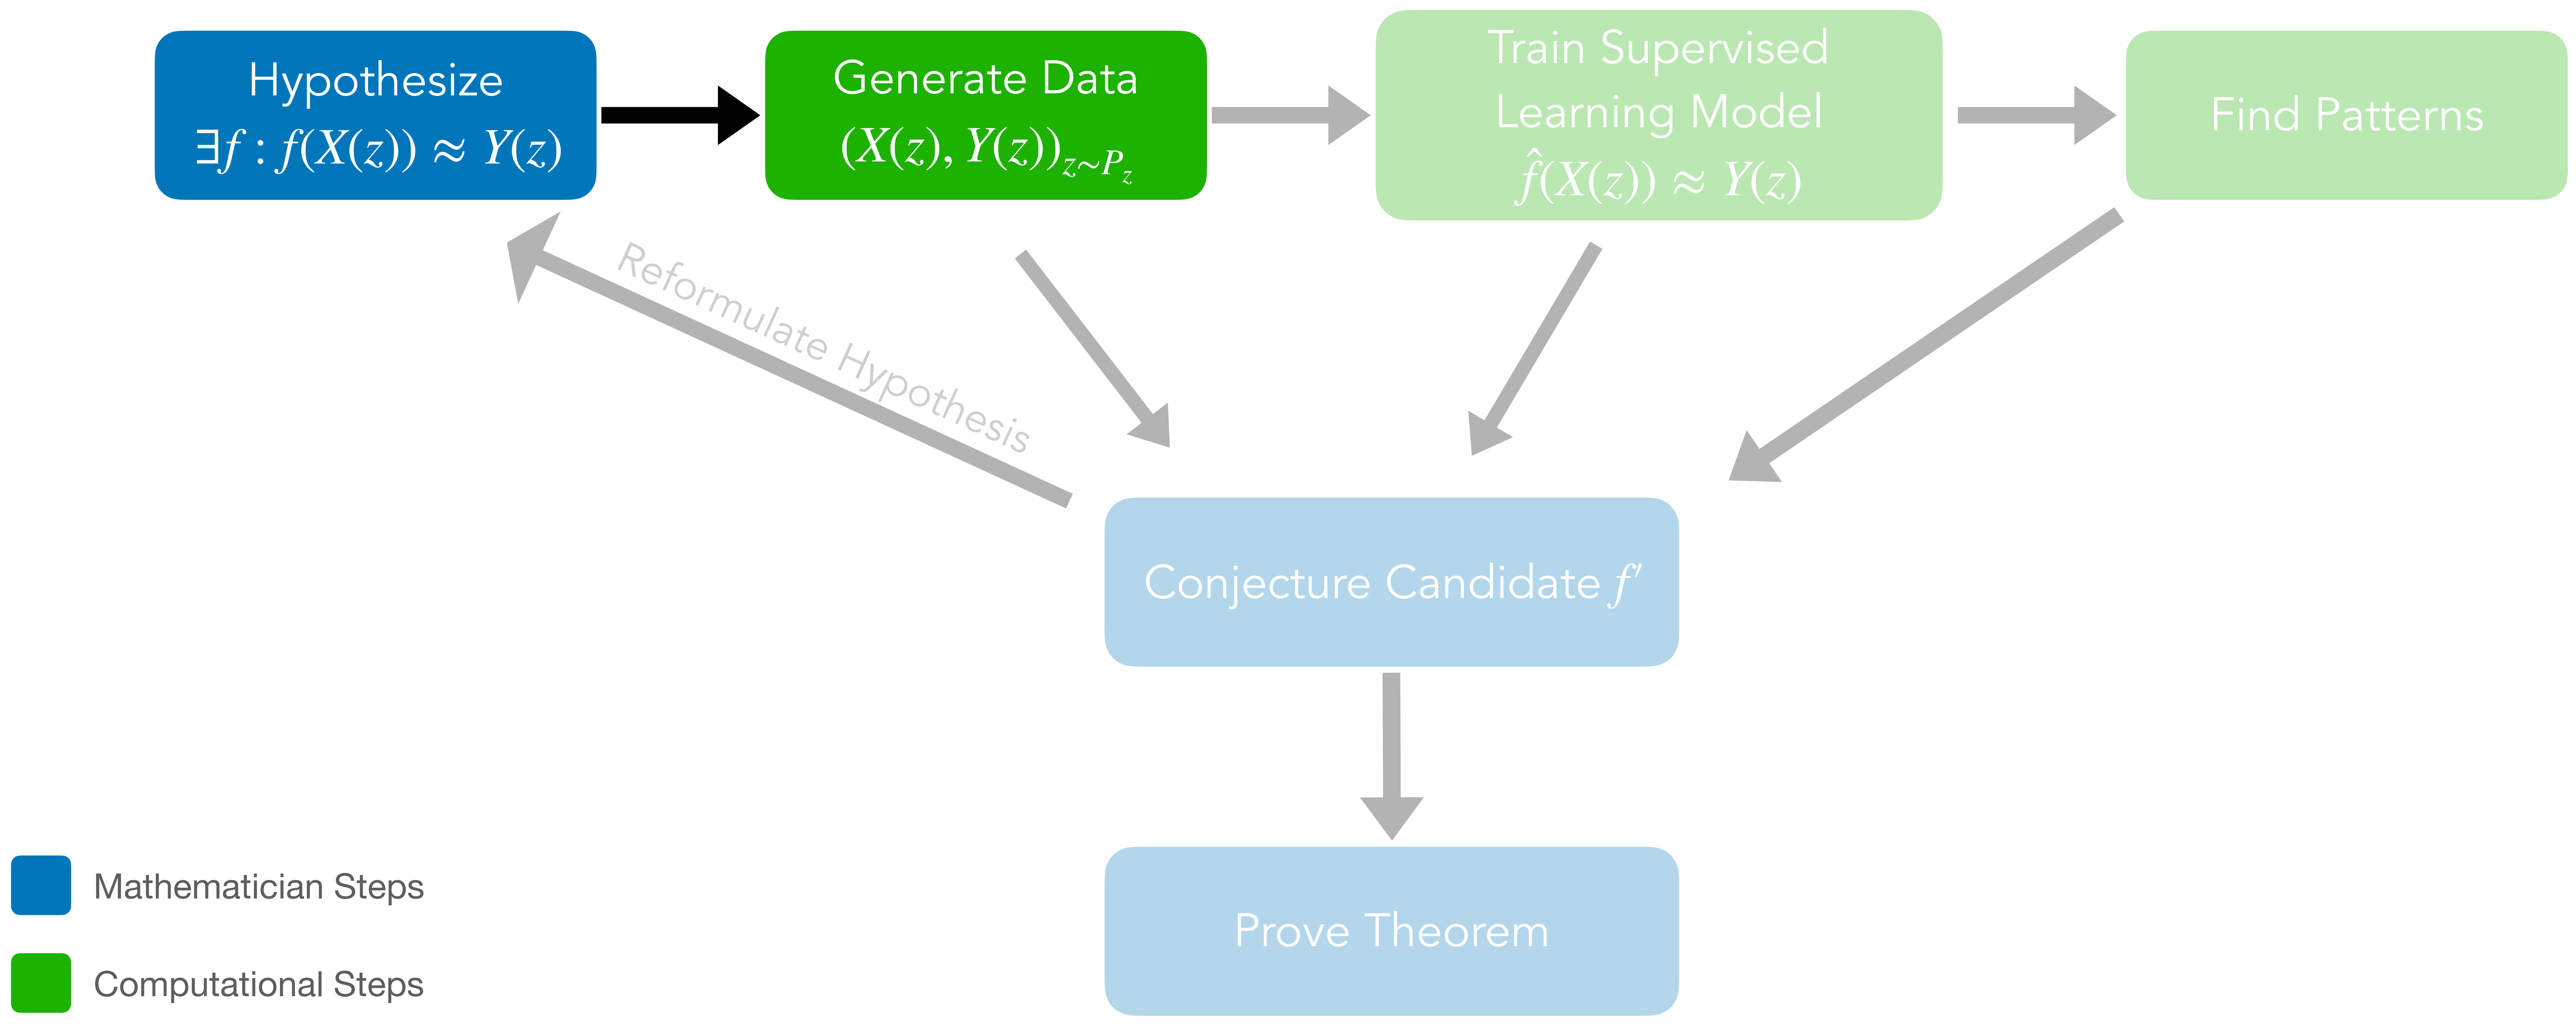

Their view of the future

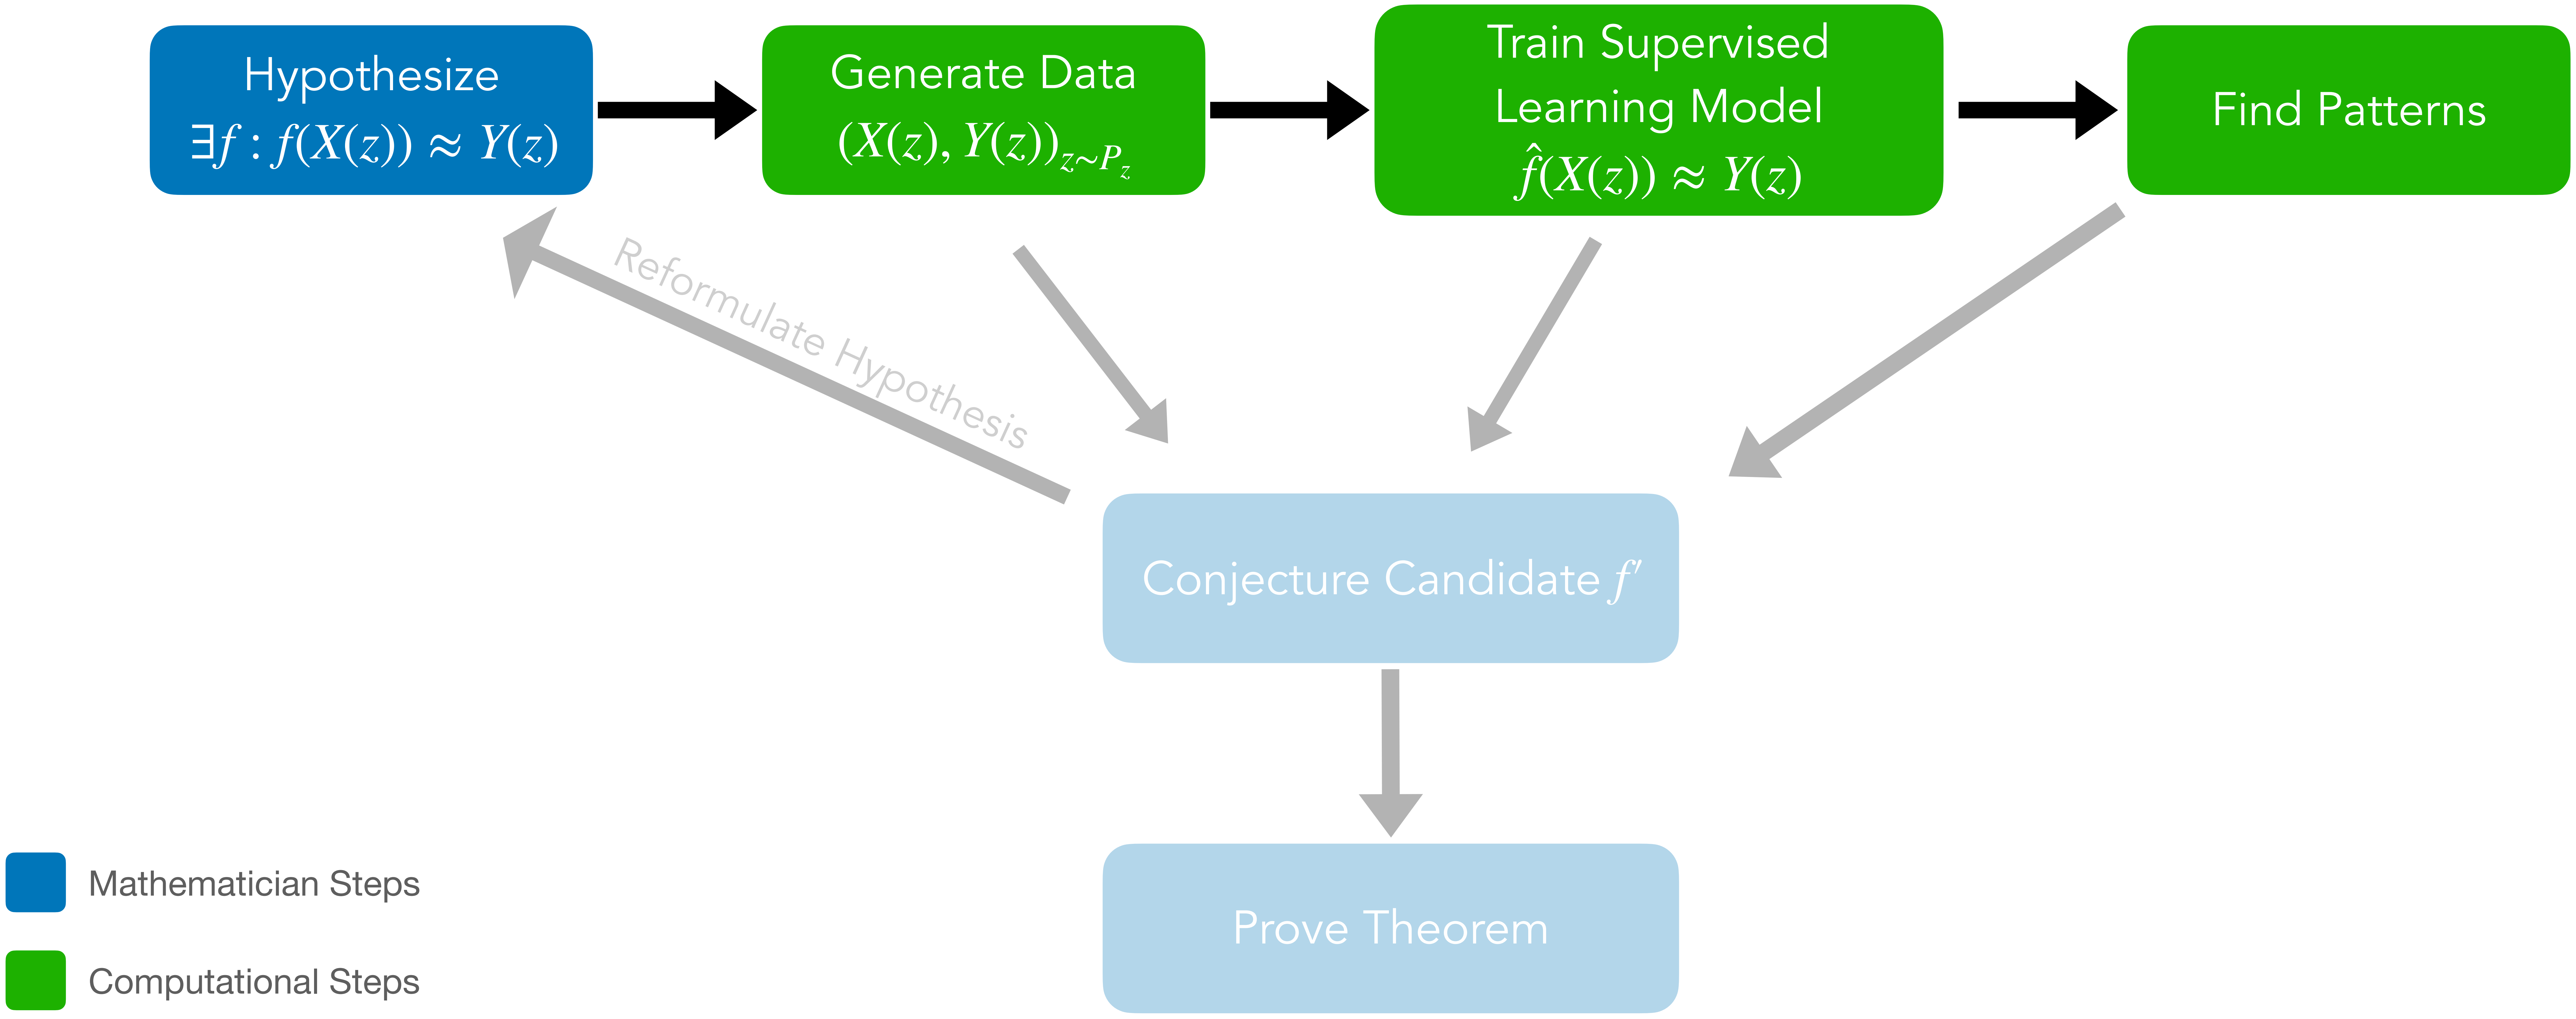

Their view of the future

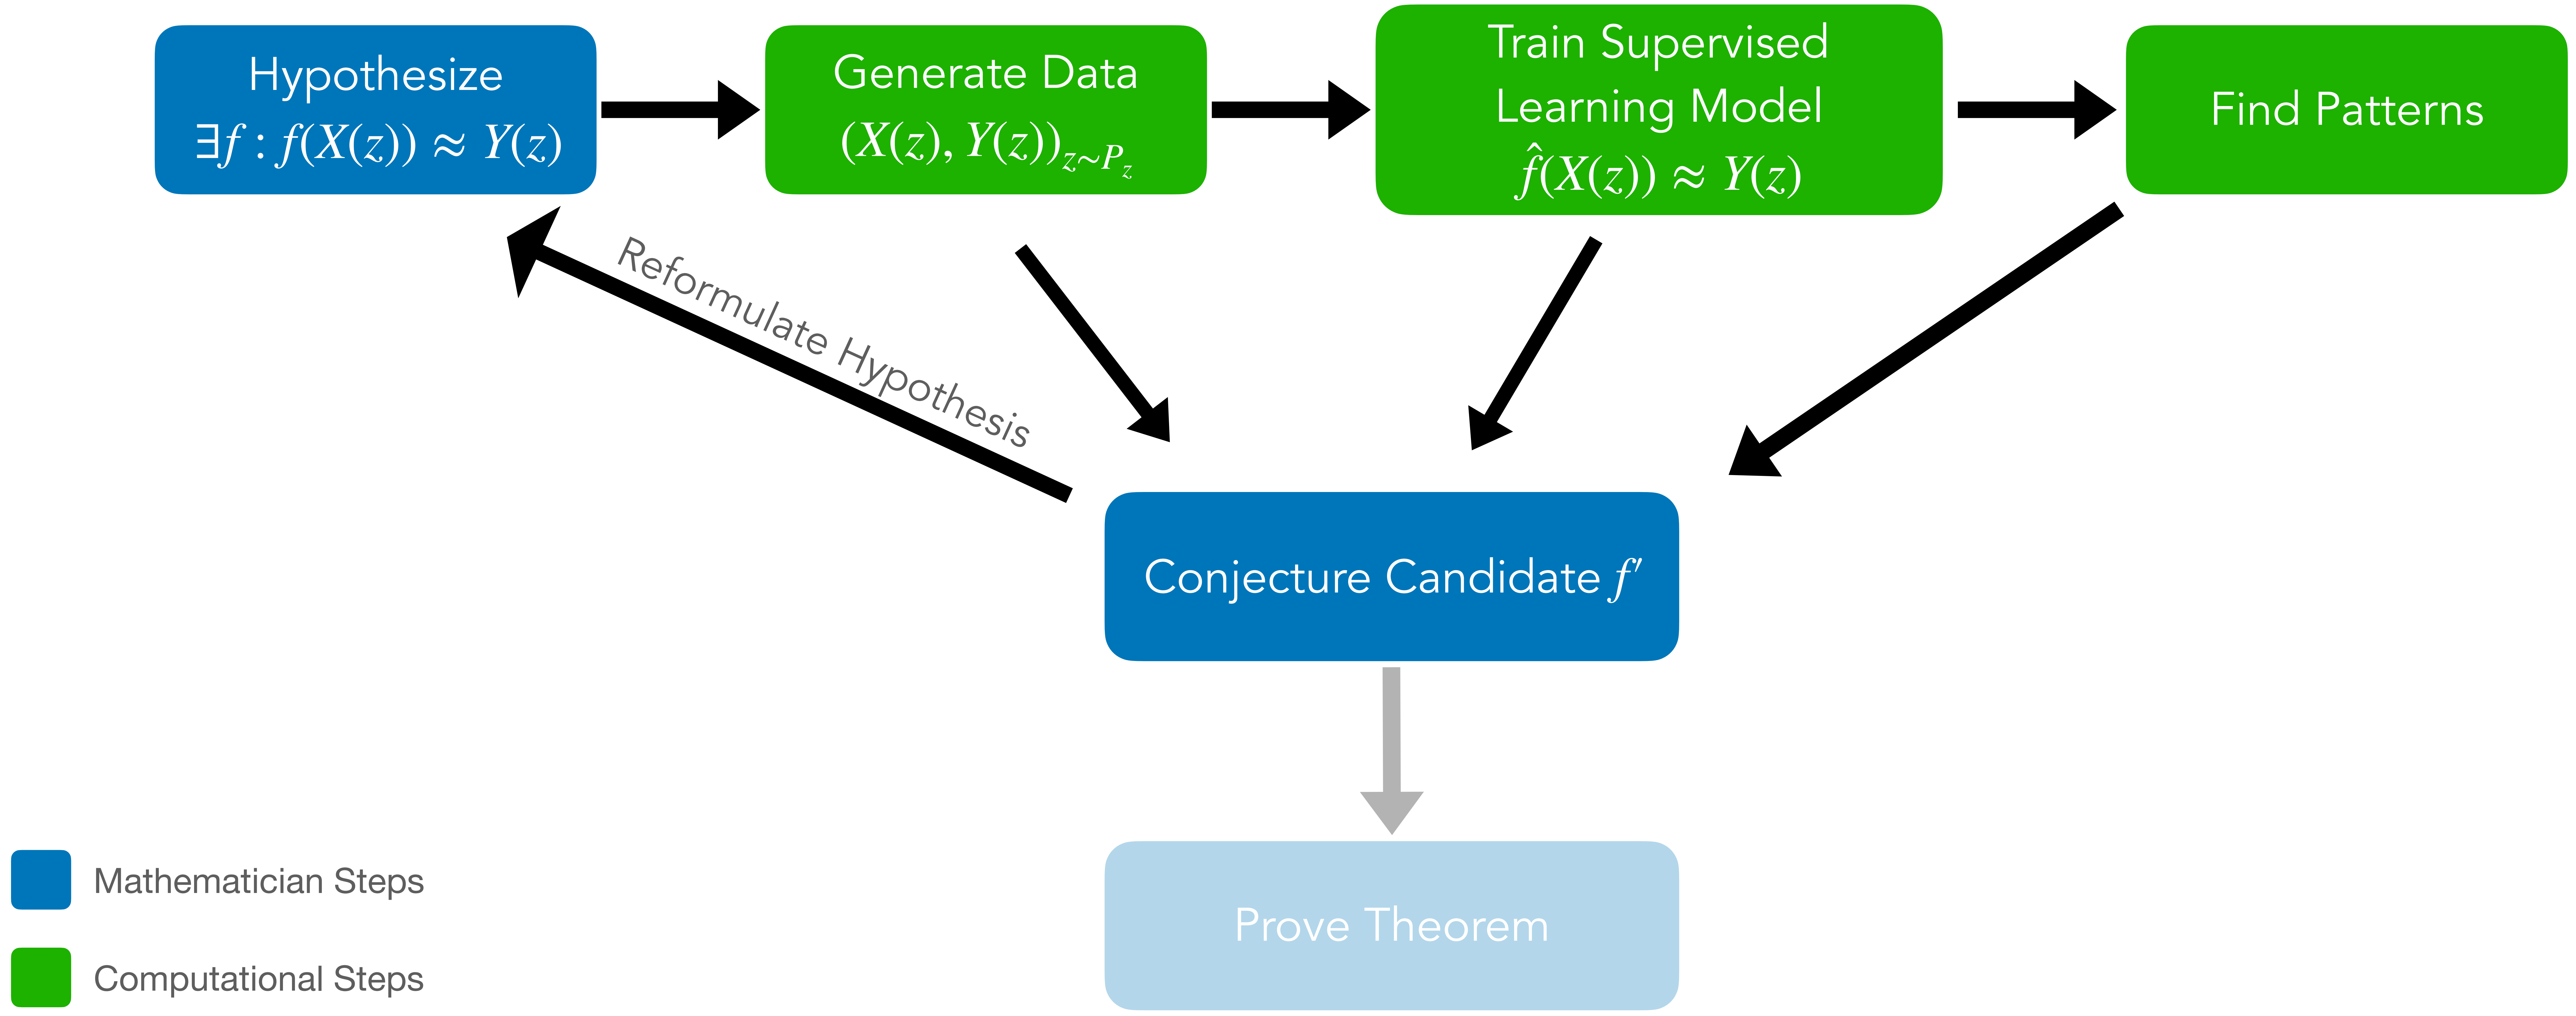

Their view of the future

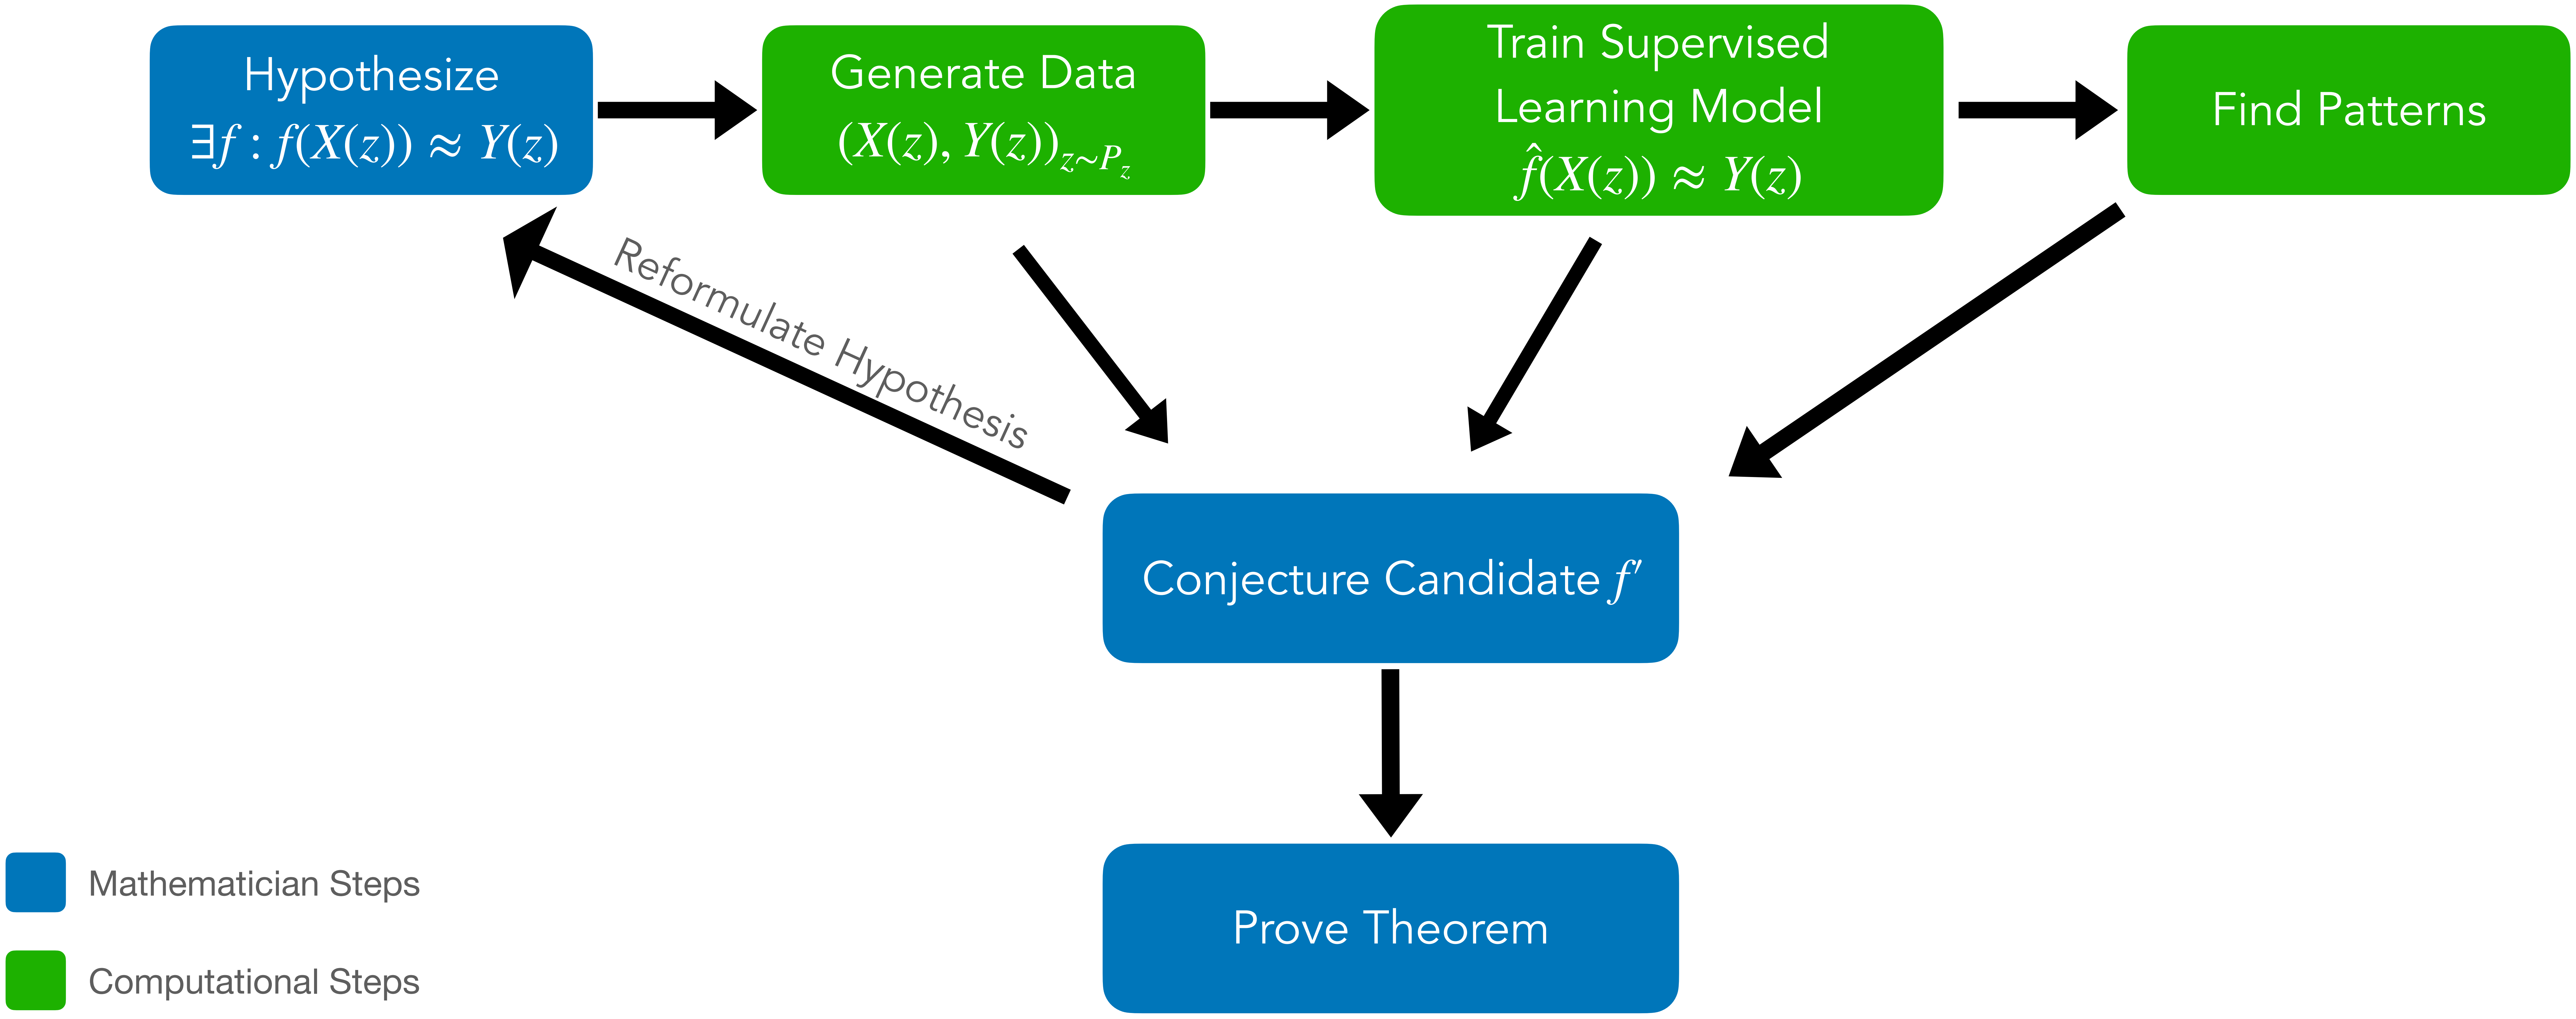

Their view of the future

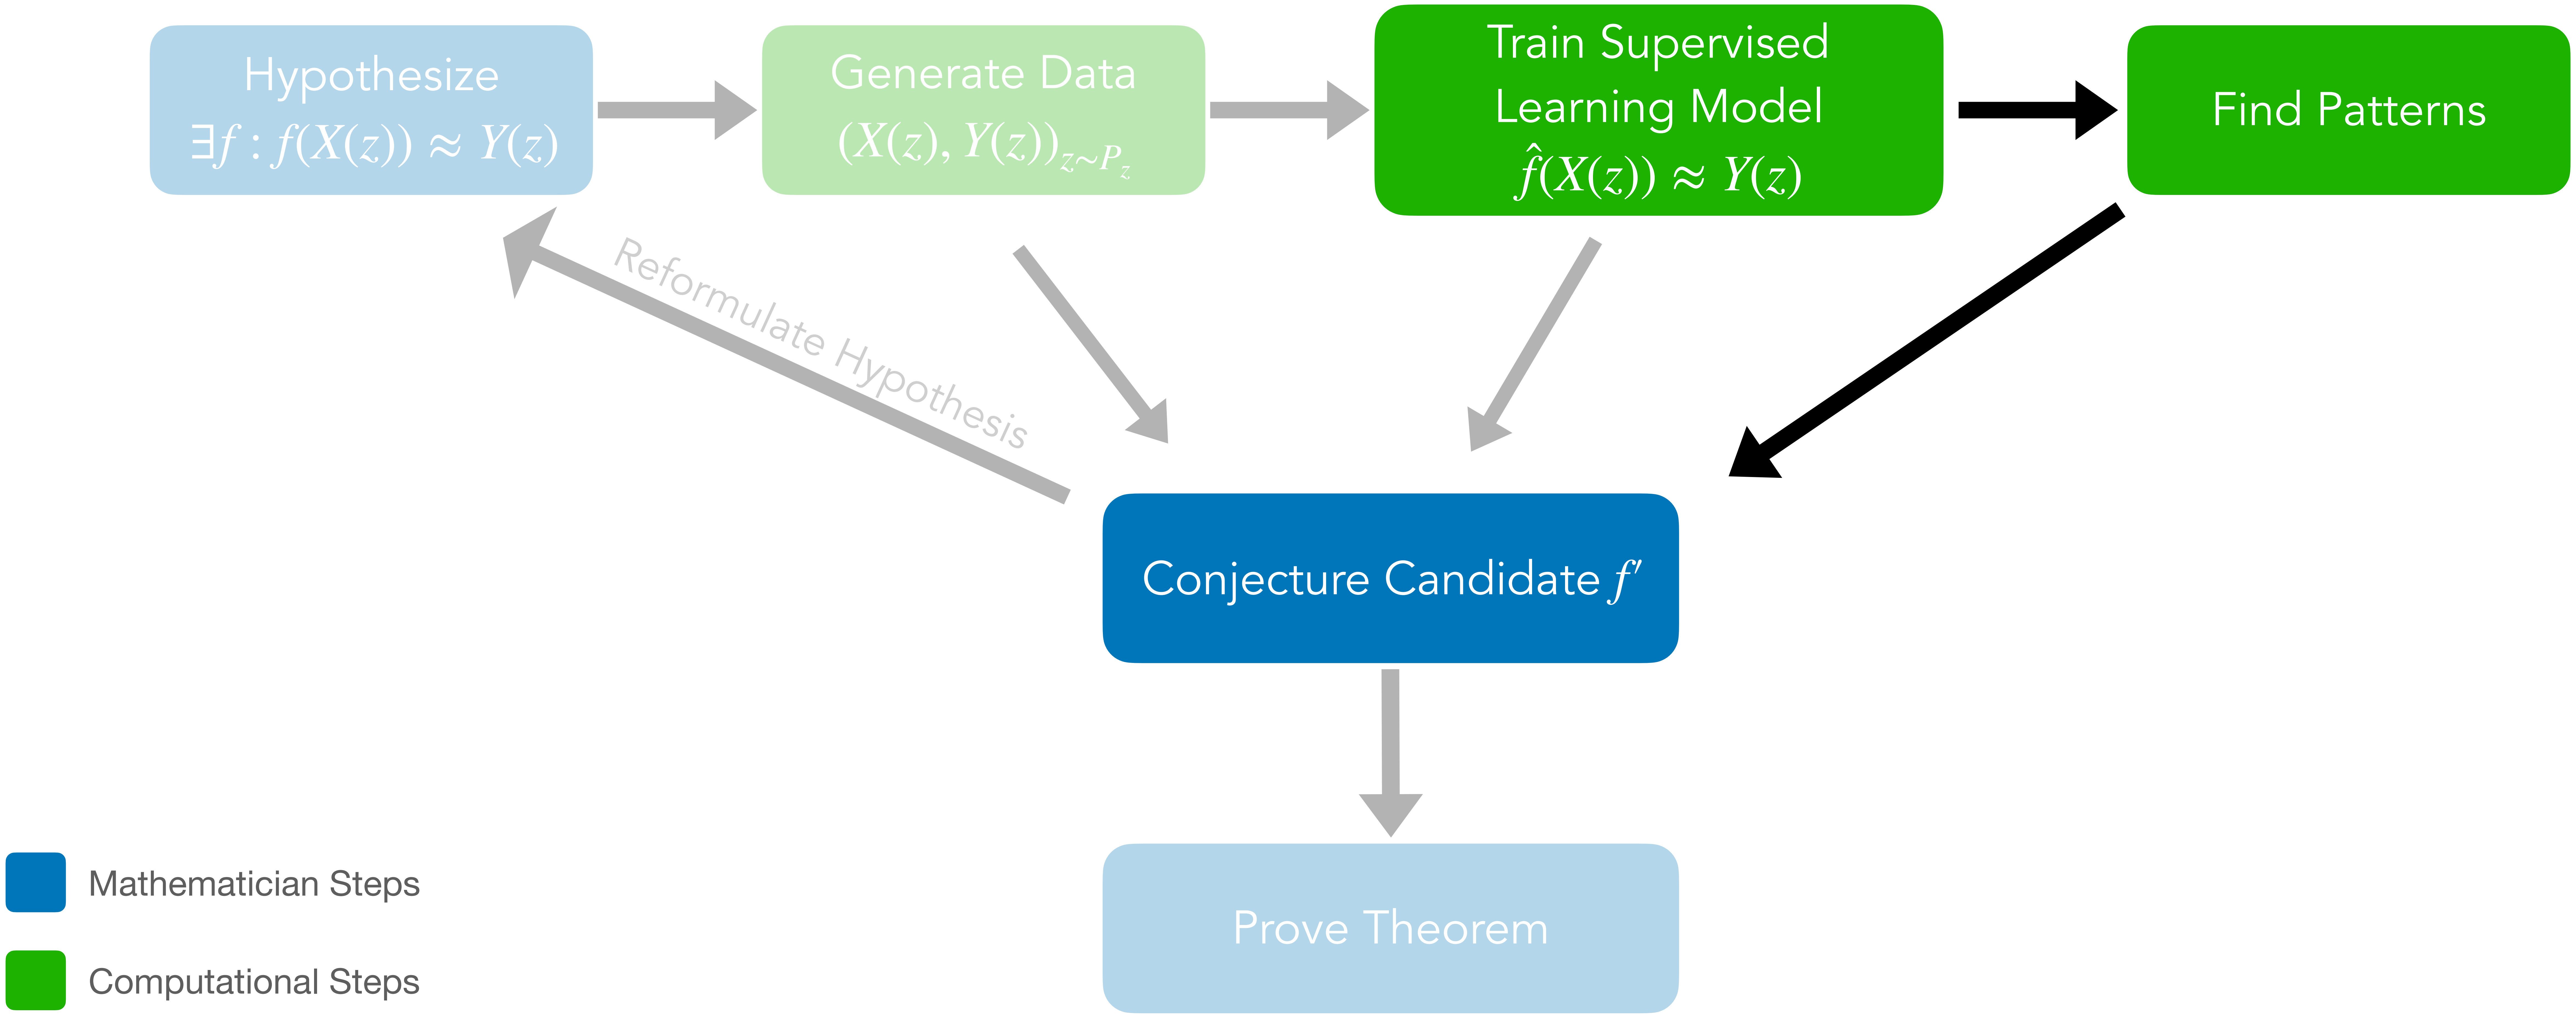

Their view of the future

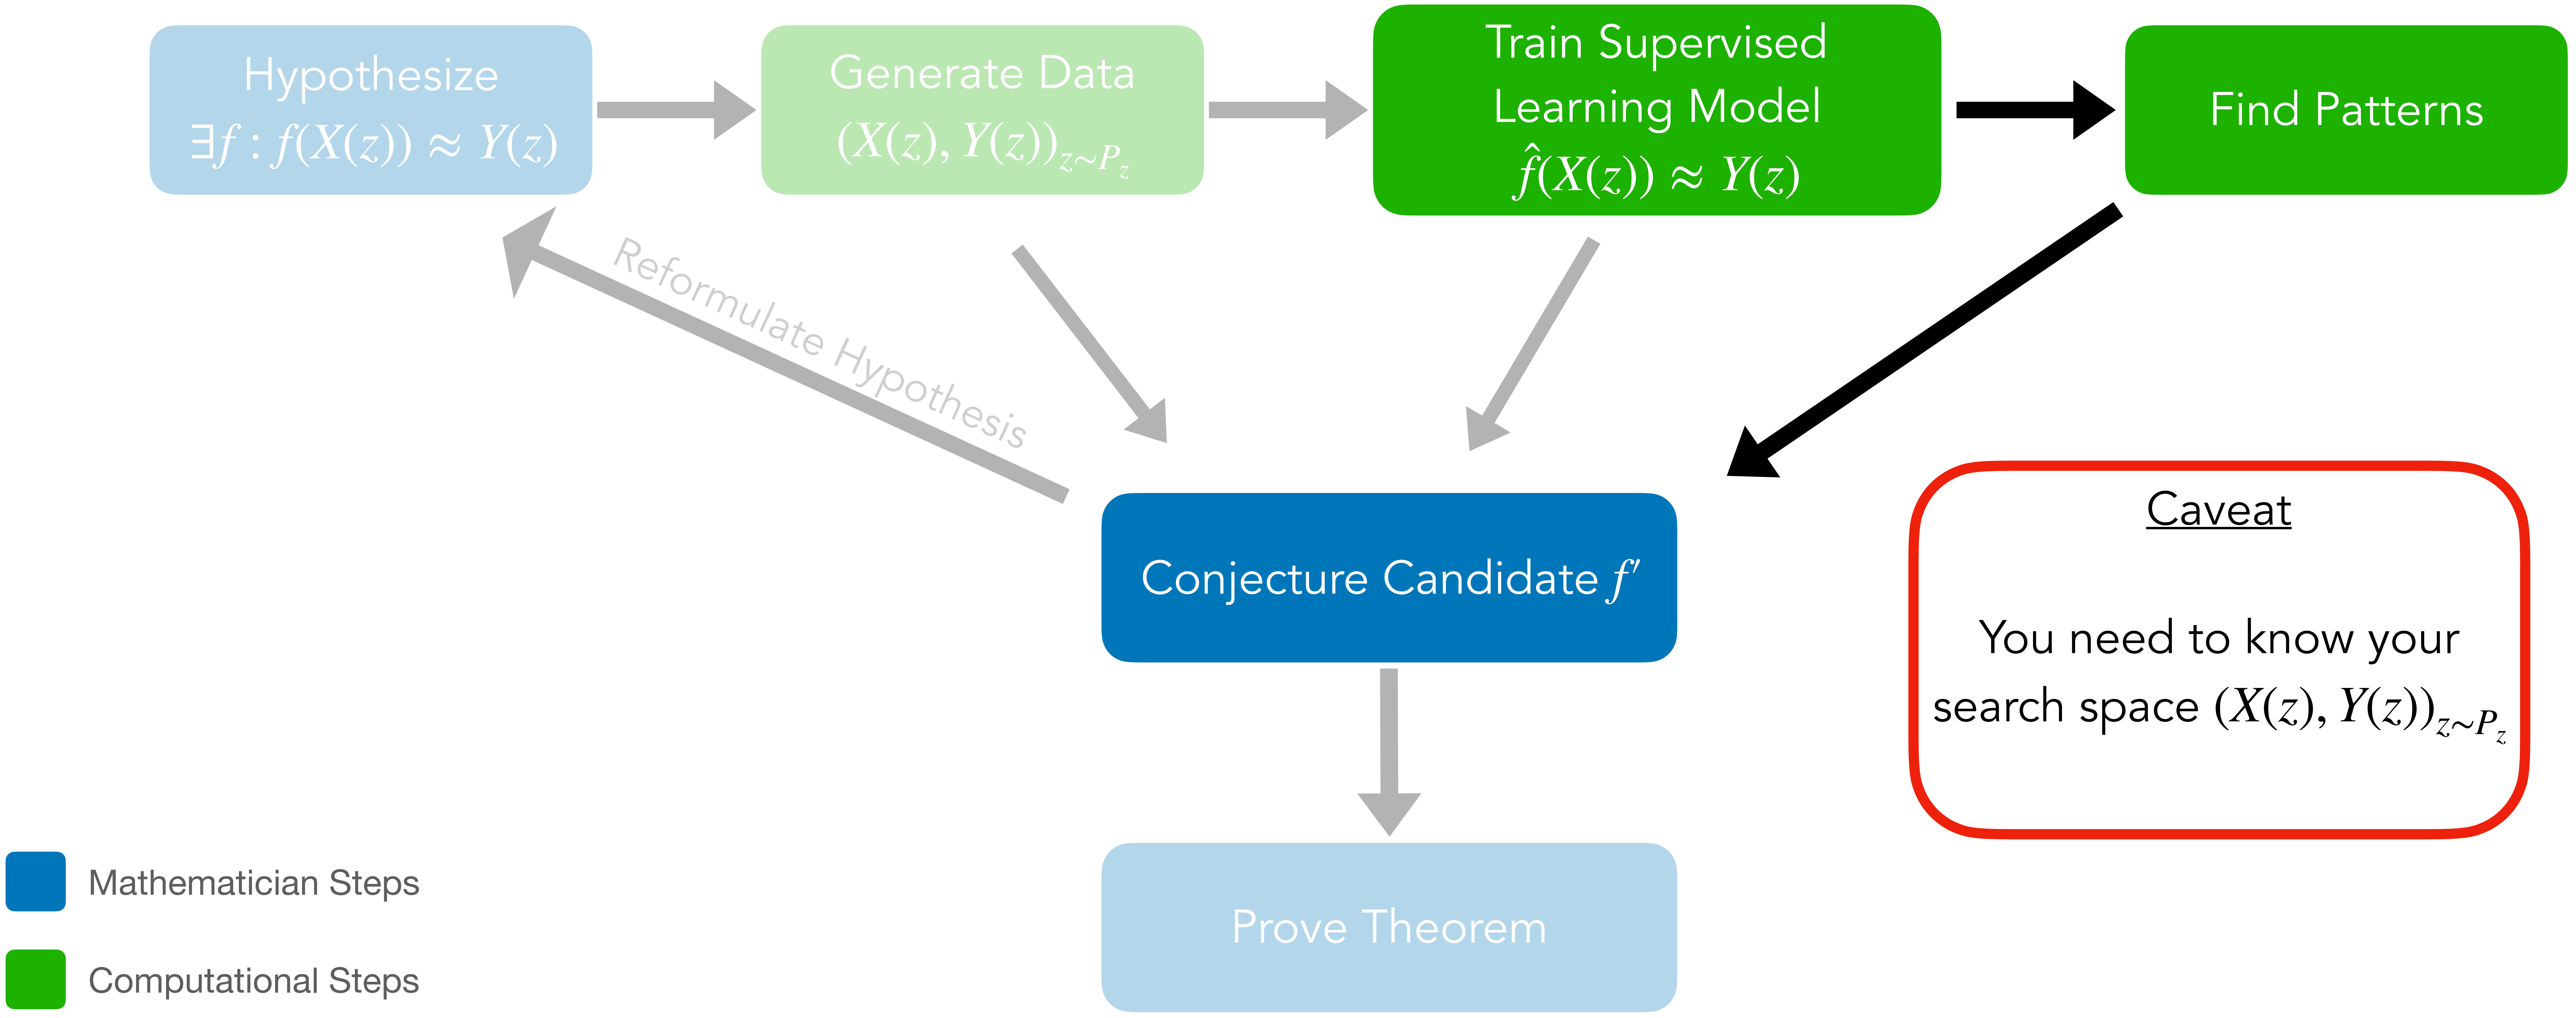

What I'm interested in

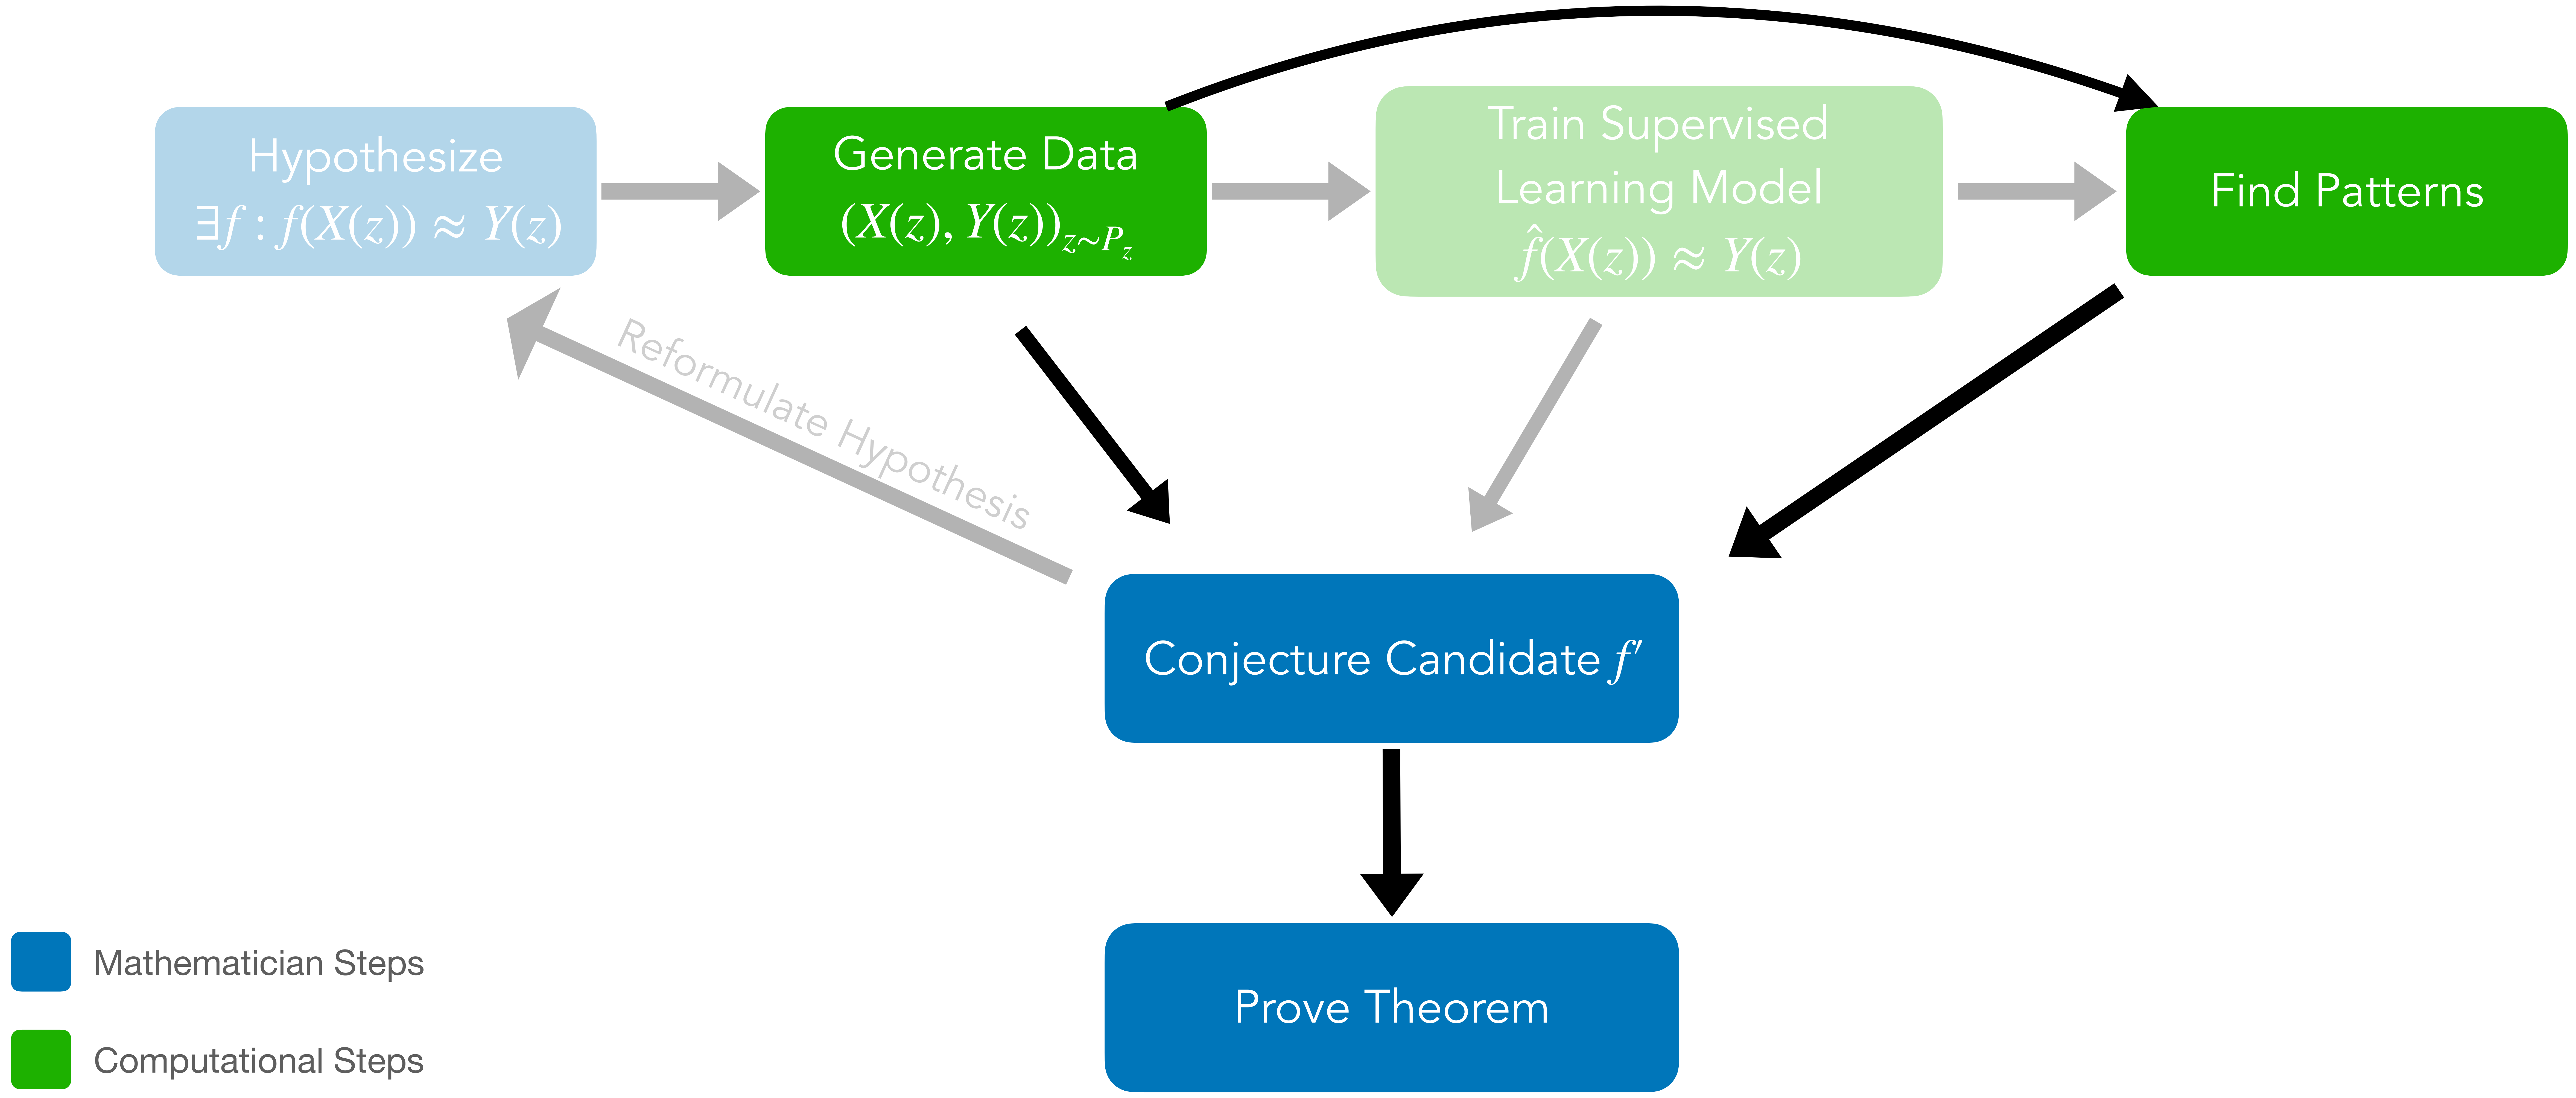

How can we use Reinforcement Learning or generative techniques?

Example Problems

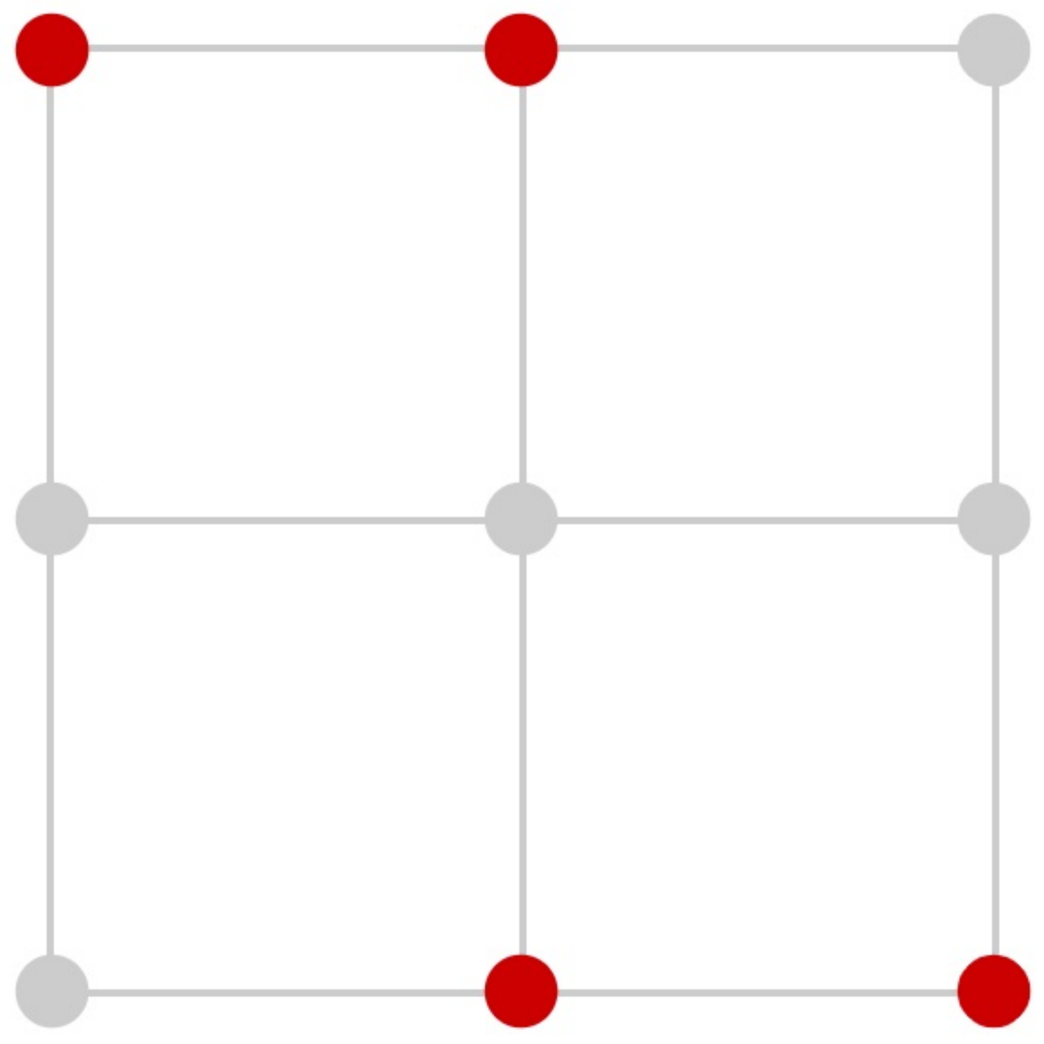

Combinatorics

Given an n x n finite integer lattice, what's the size of the largest subset such that no three points form an isosceles triangle?

Example Problems

Can we upper bound the number of points in the real plane
So that no empty convex-6-gons exist?

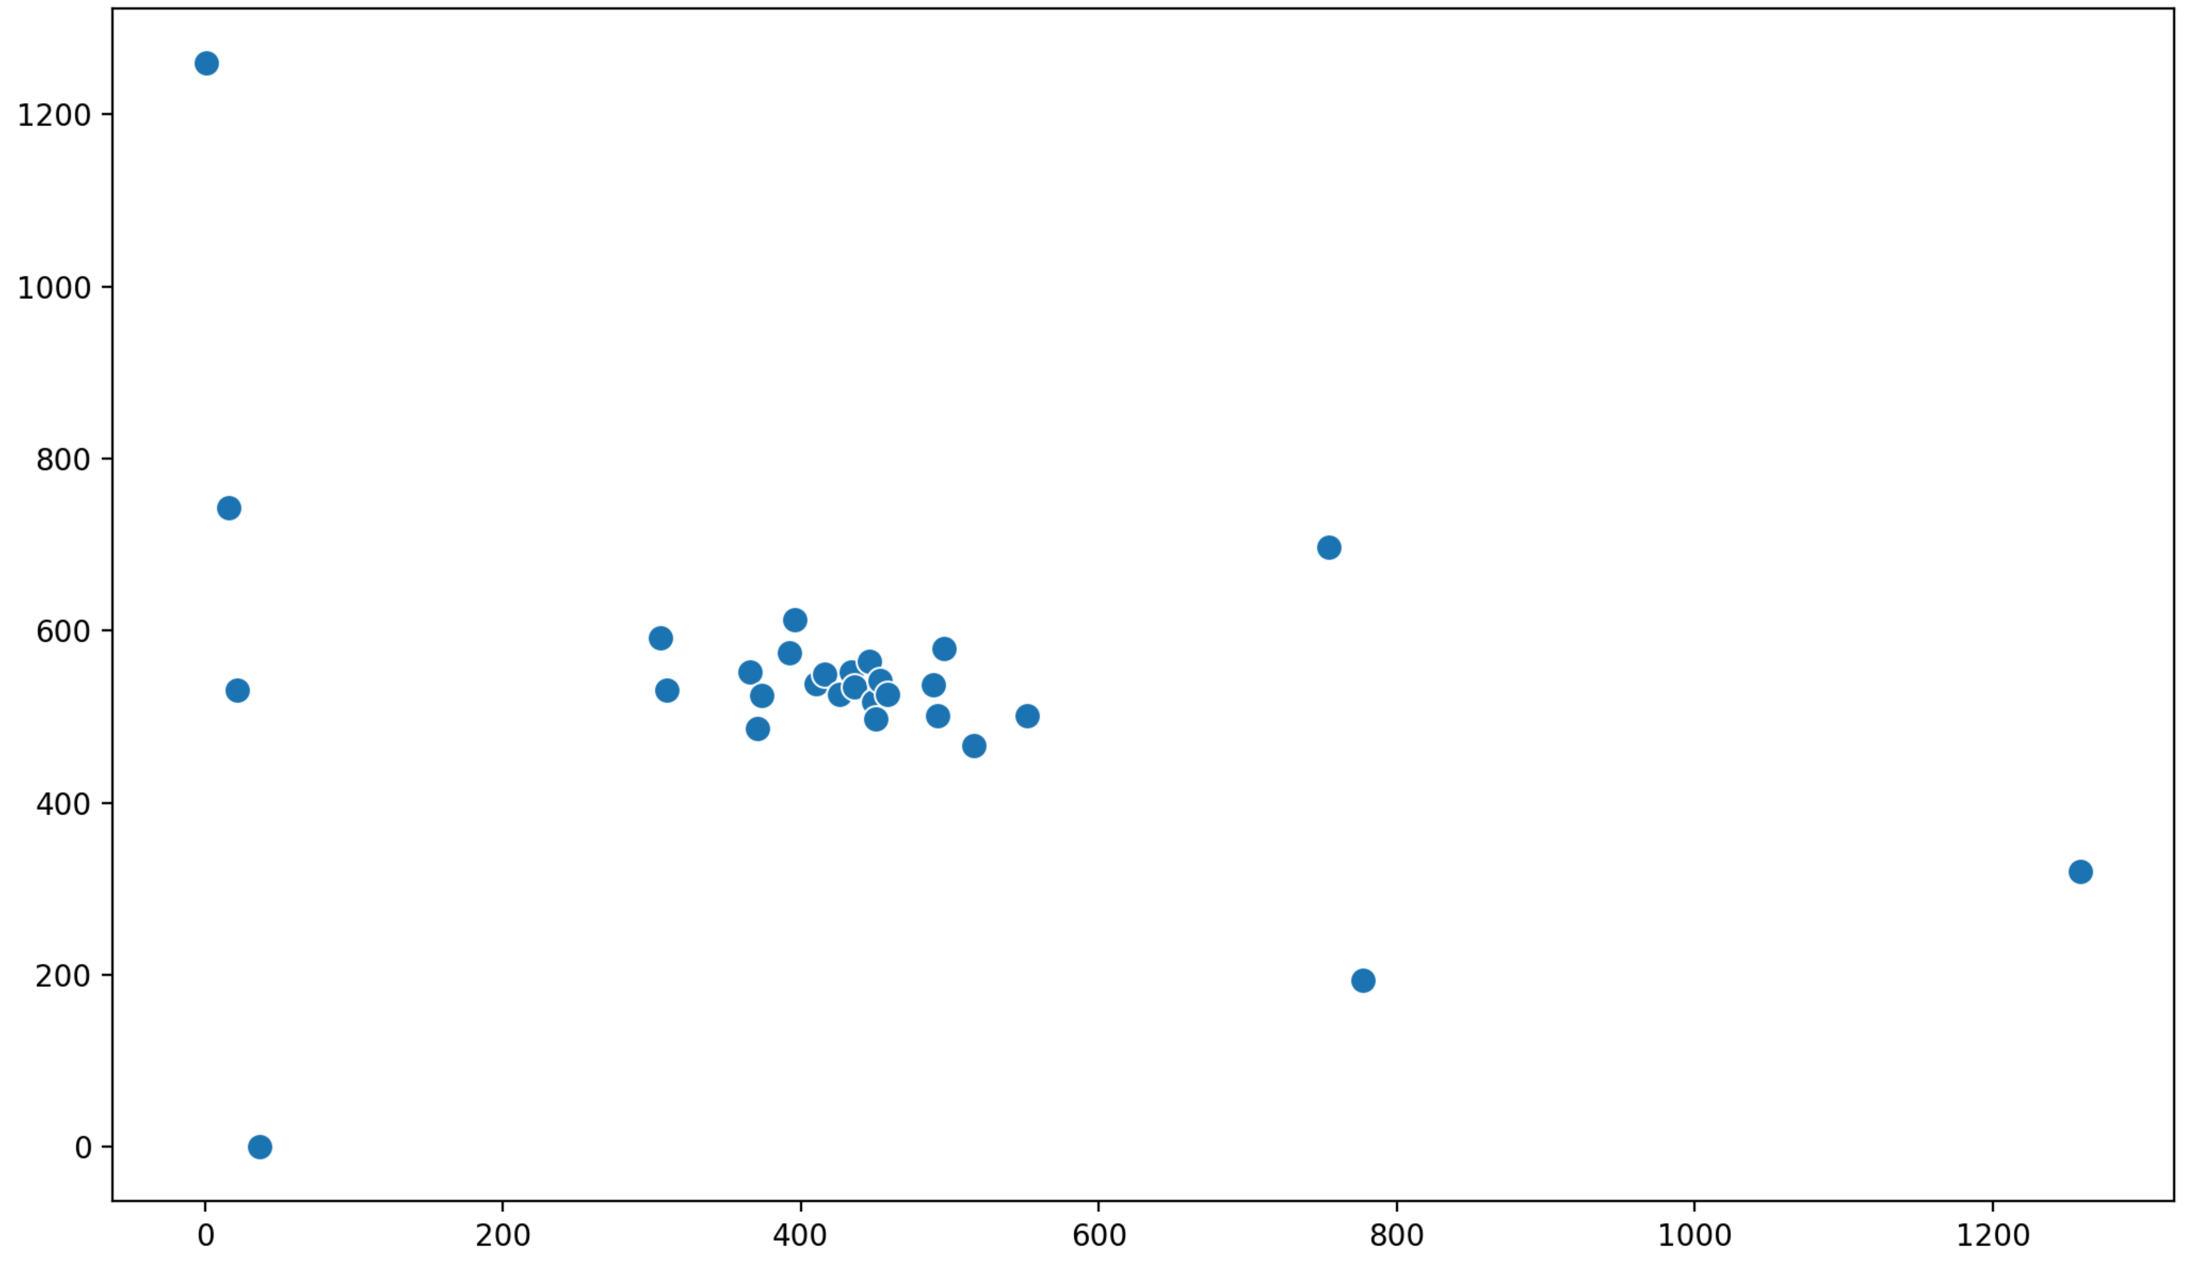

Convex Geometry

Example Problems

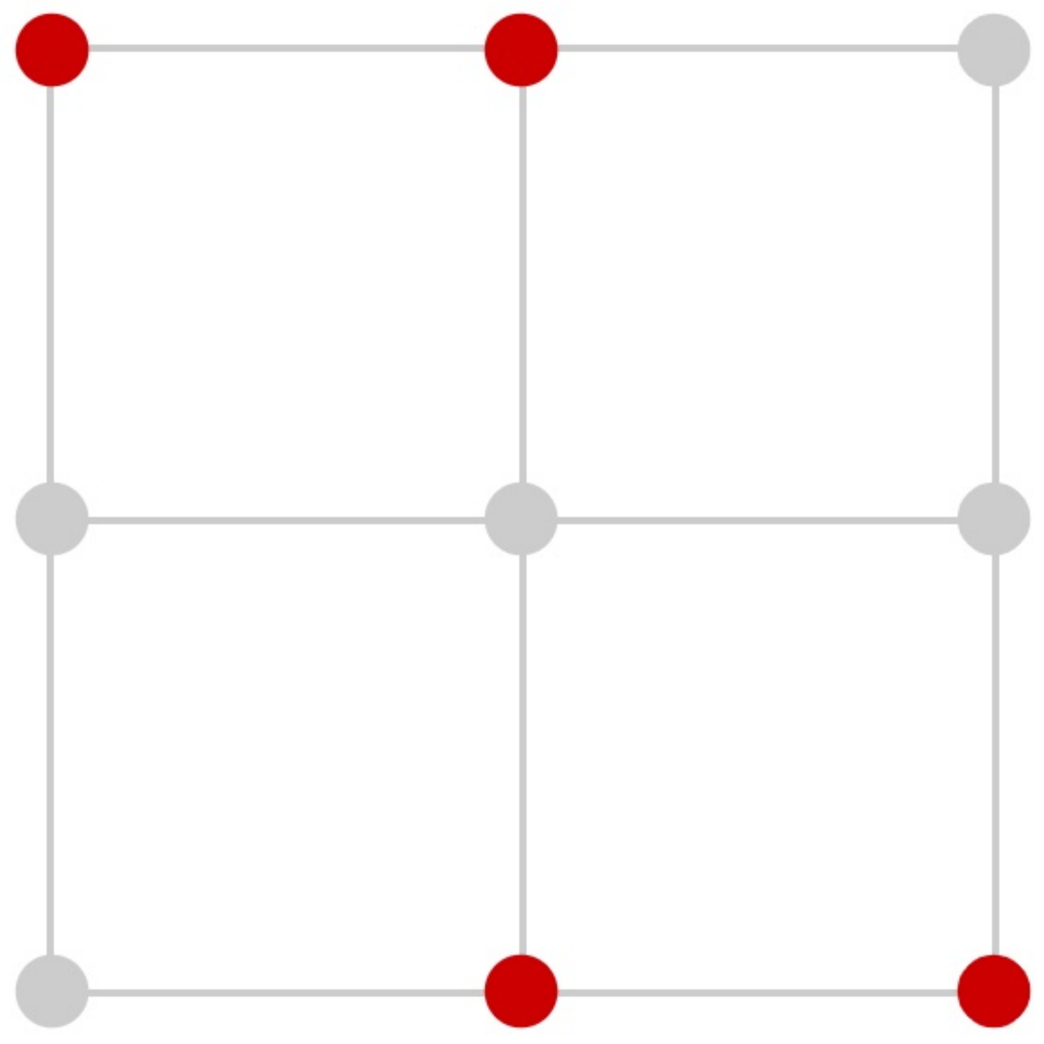

Combinatorics

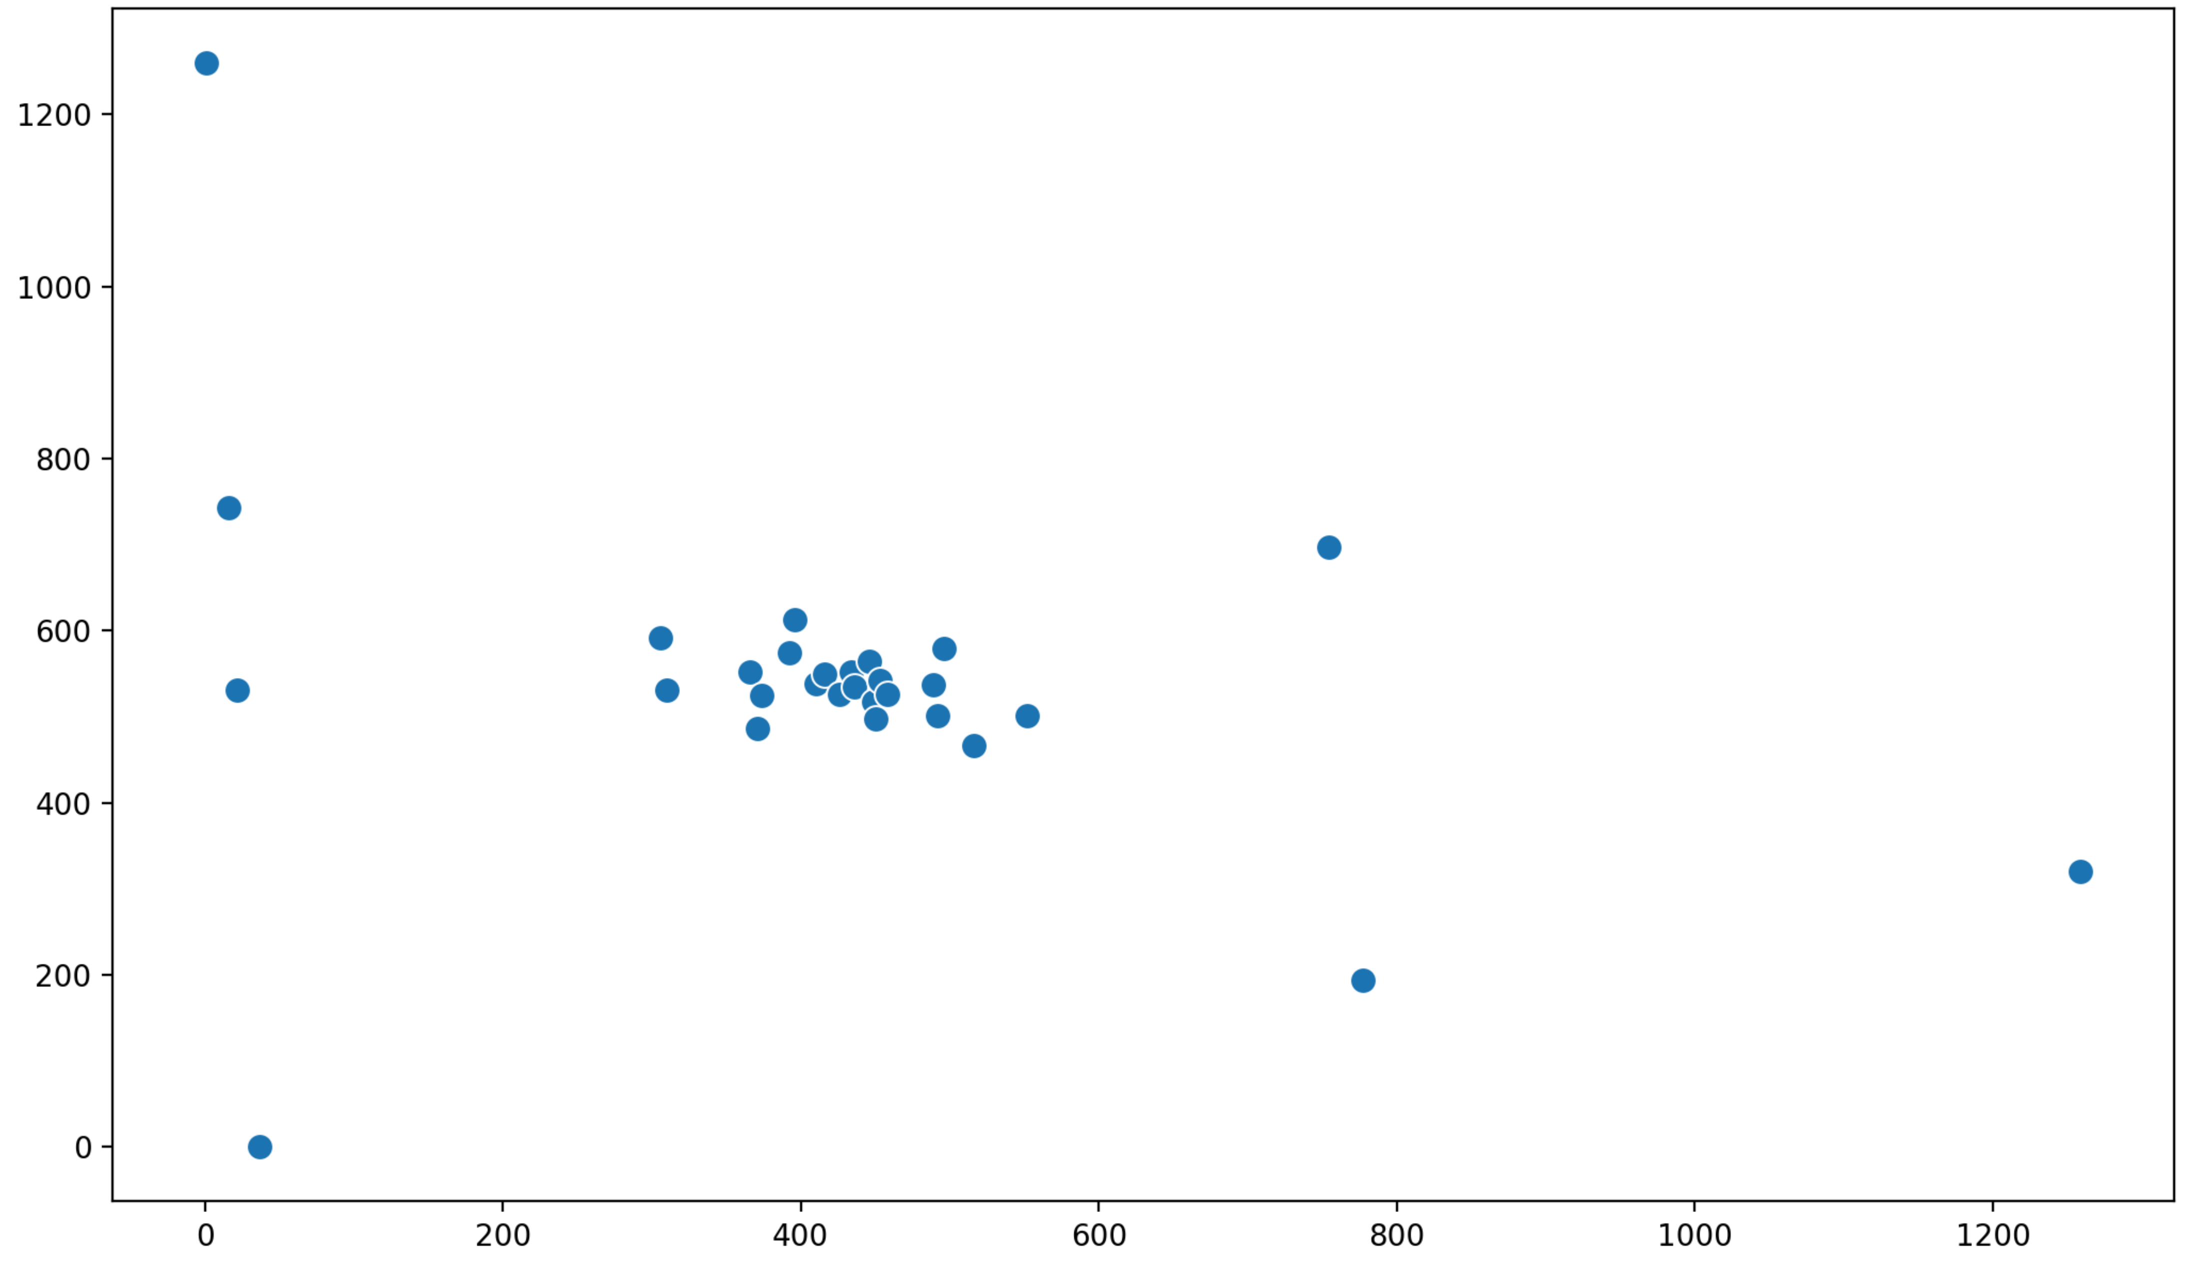

Convex Geometry

Aim

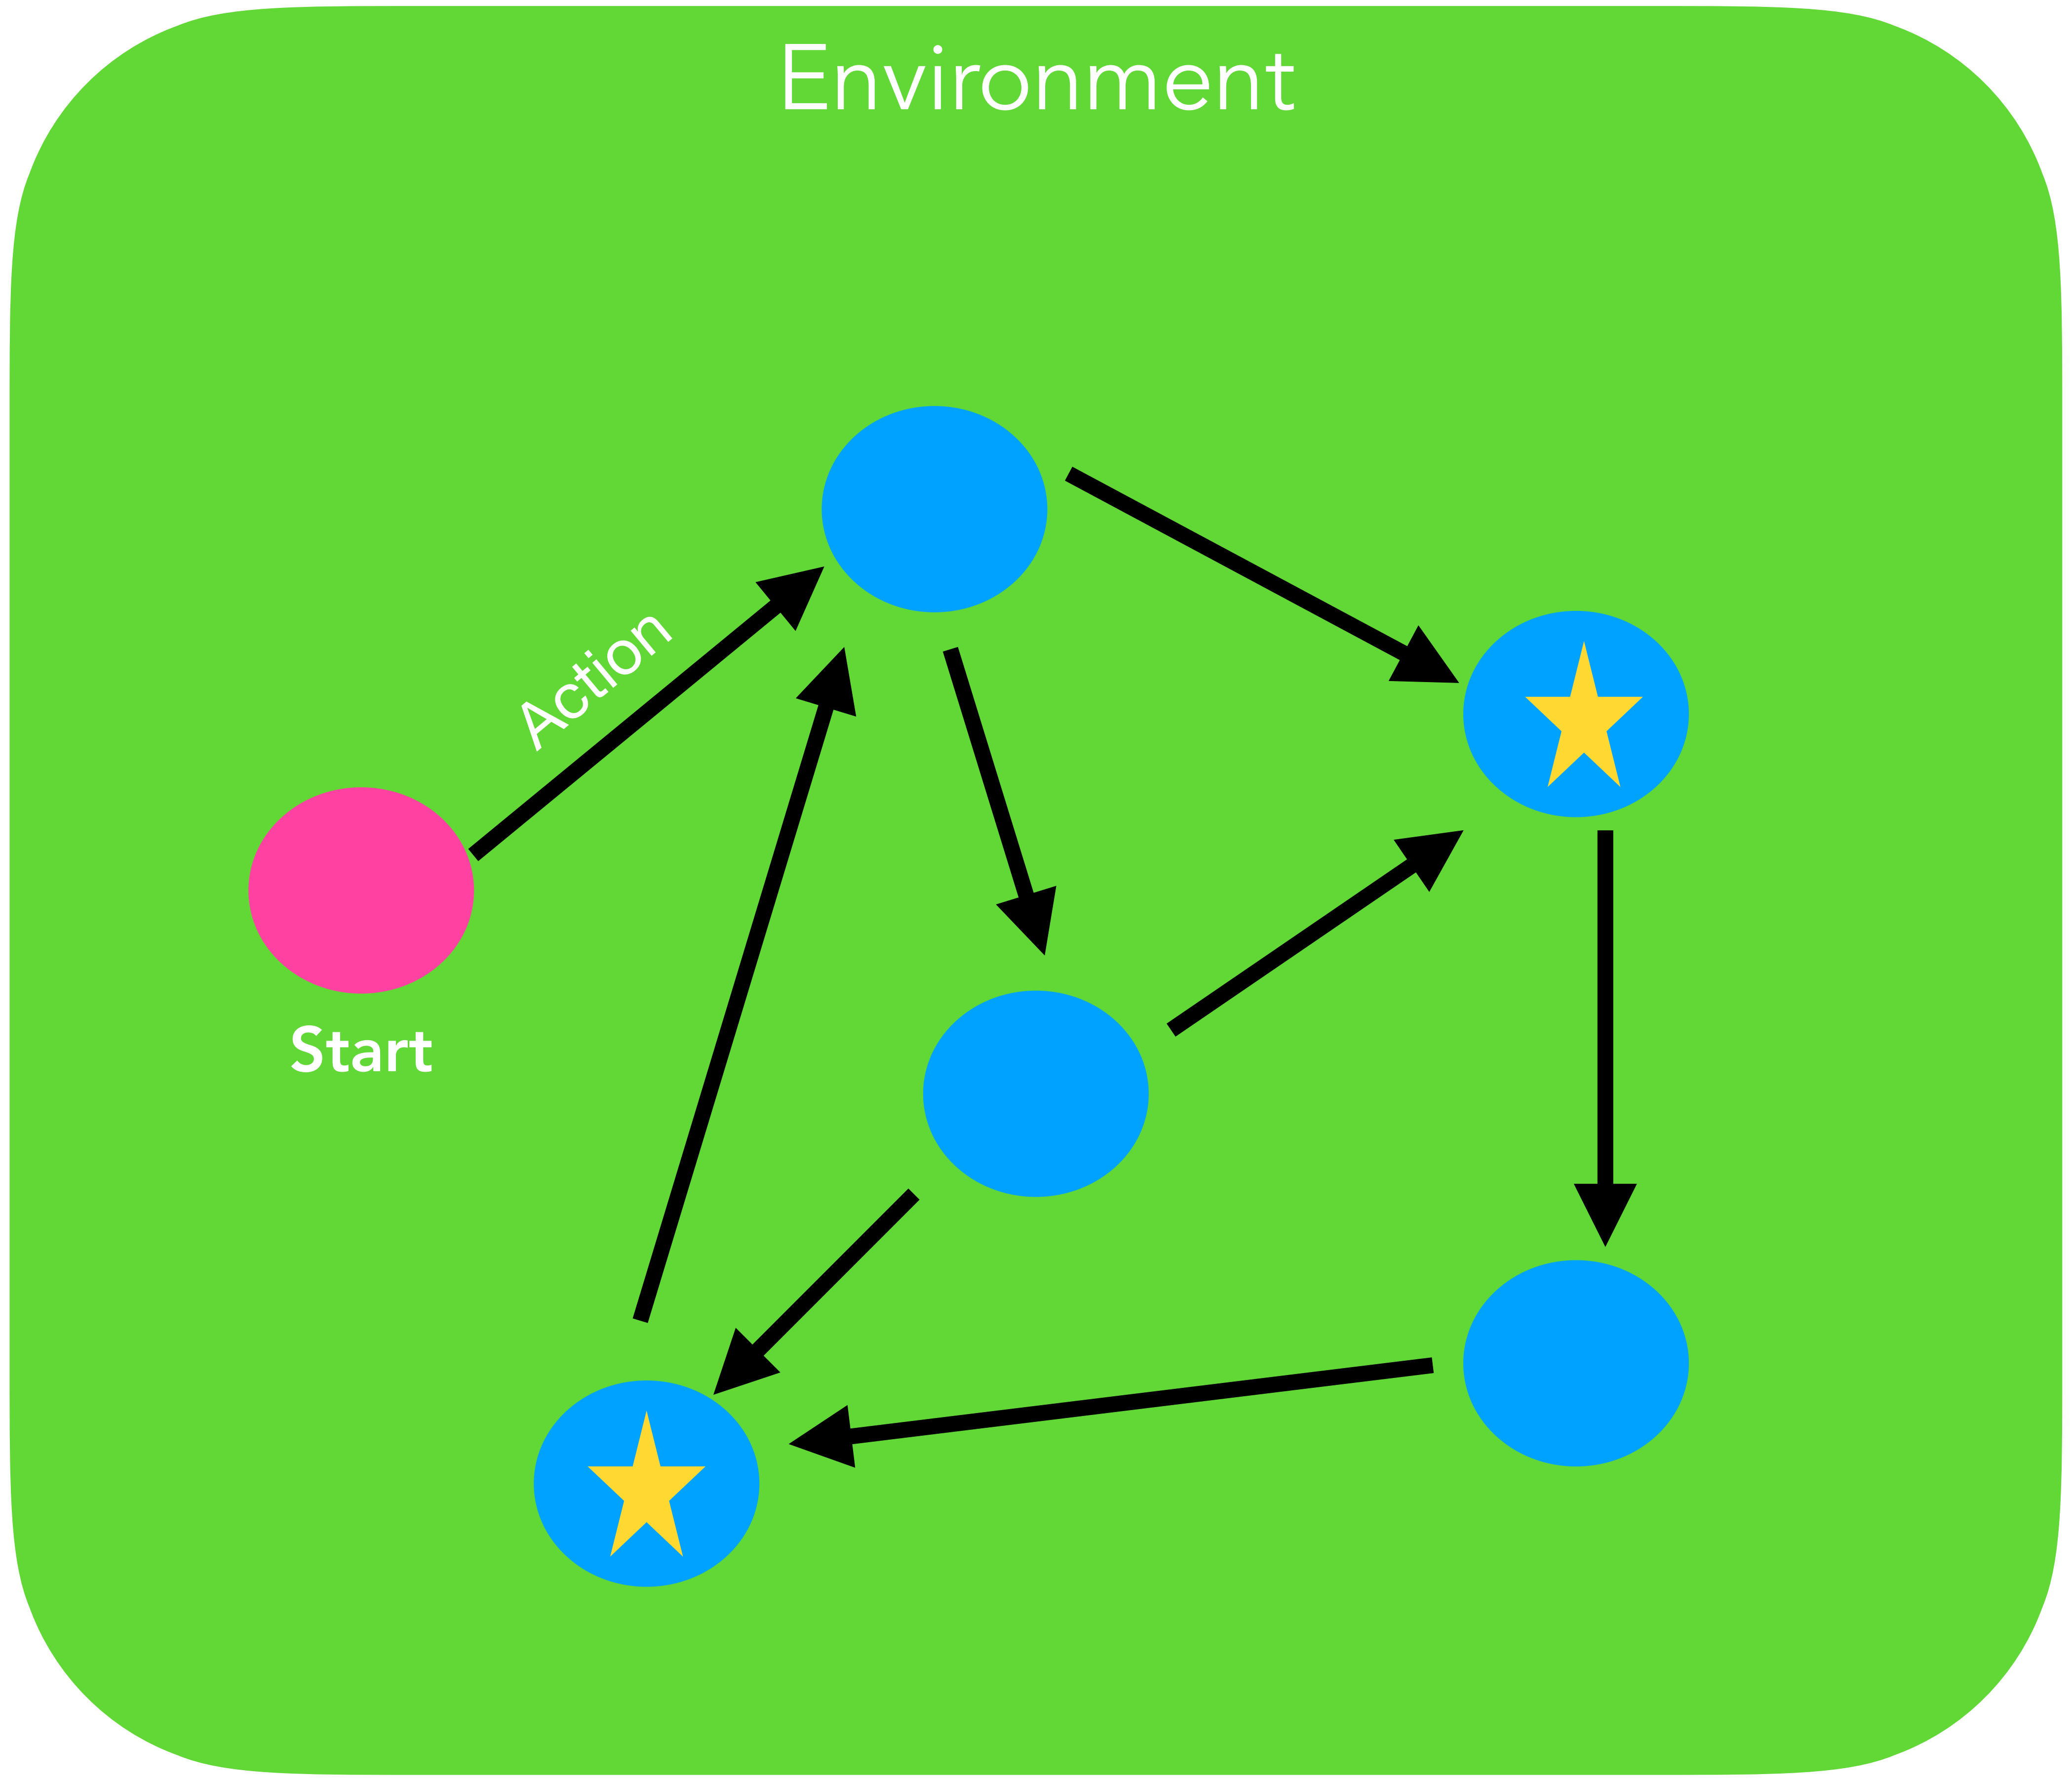

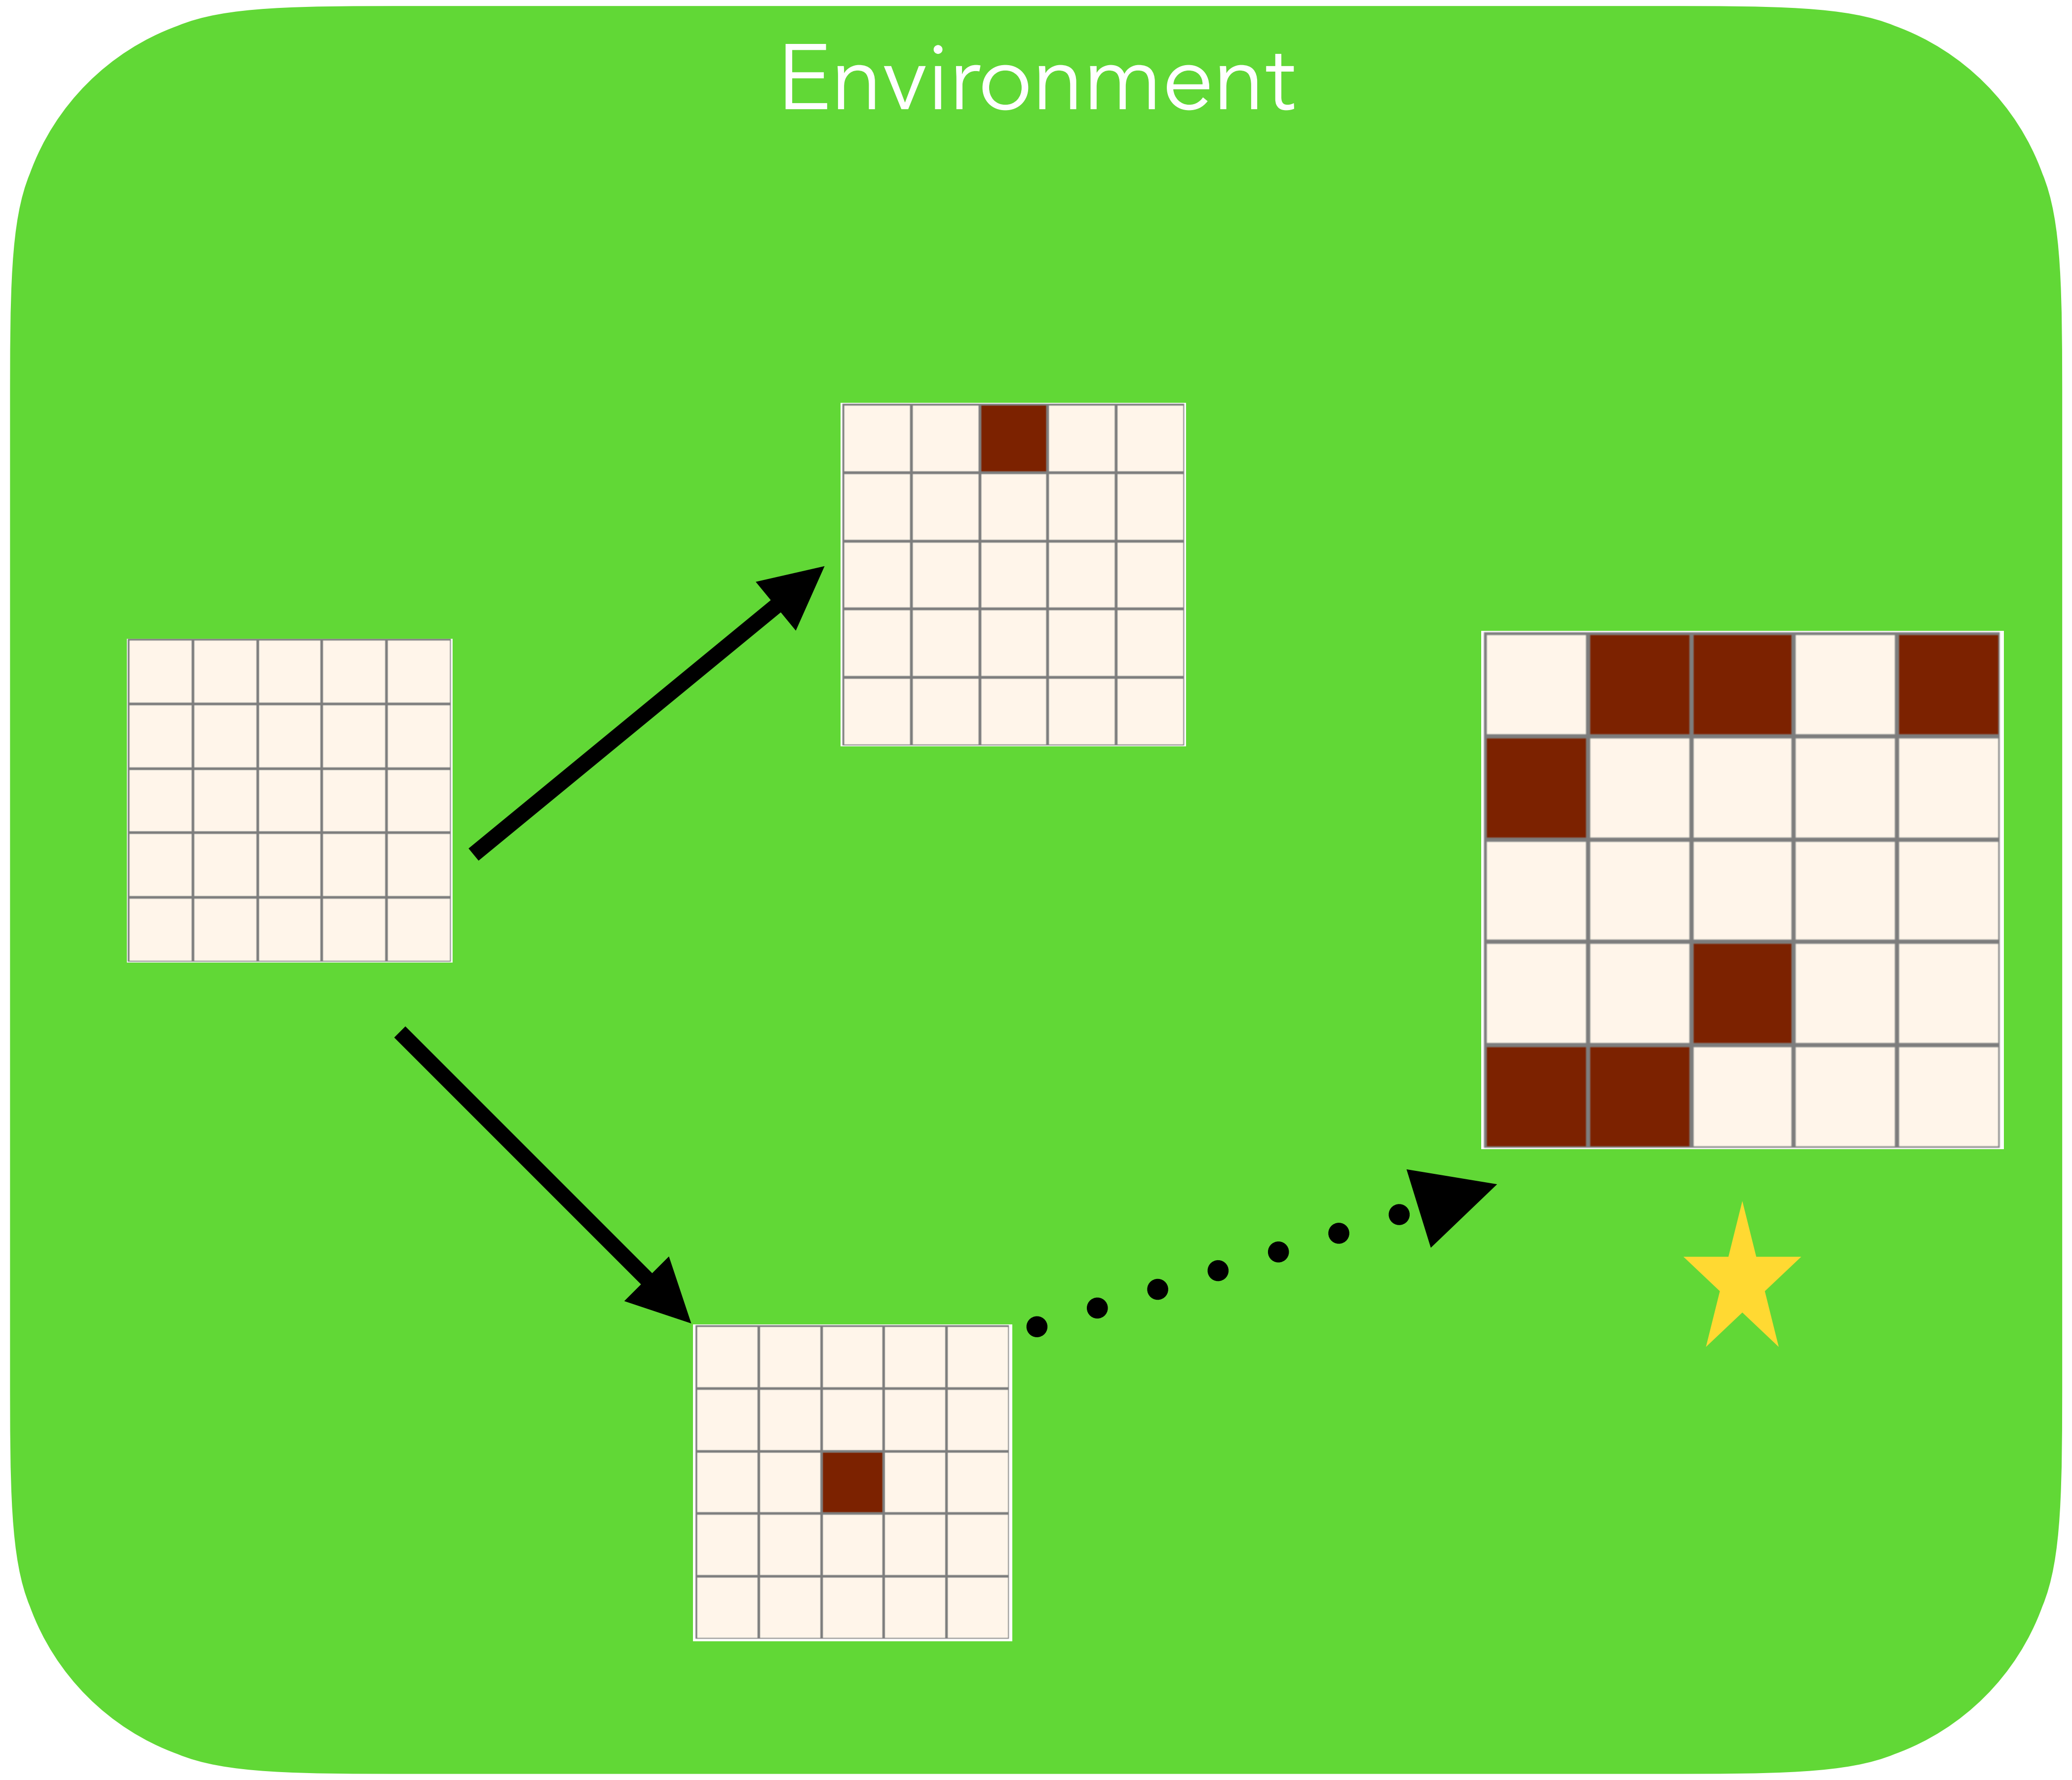

Aim

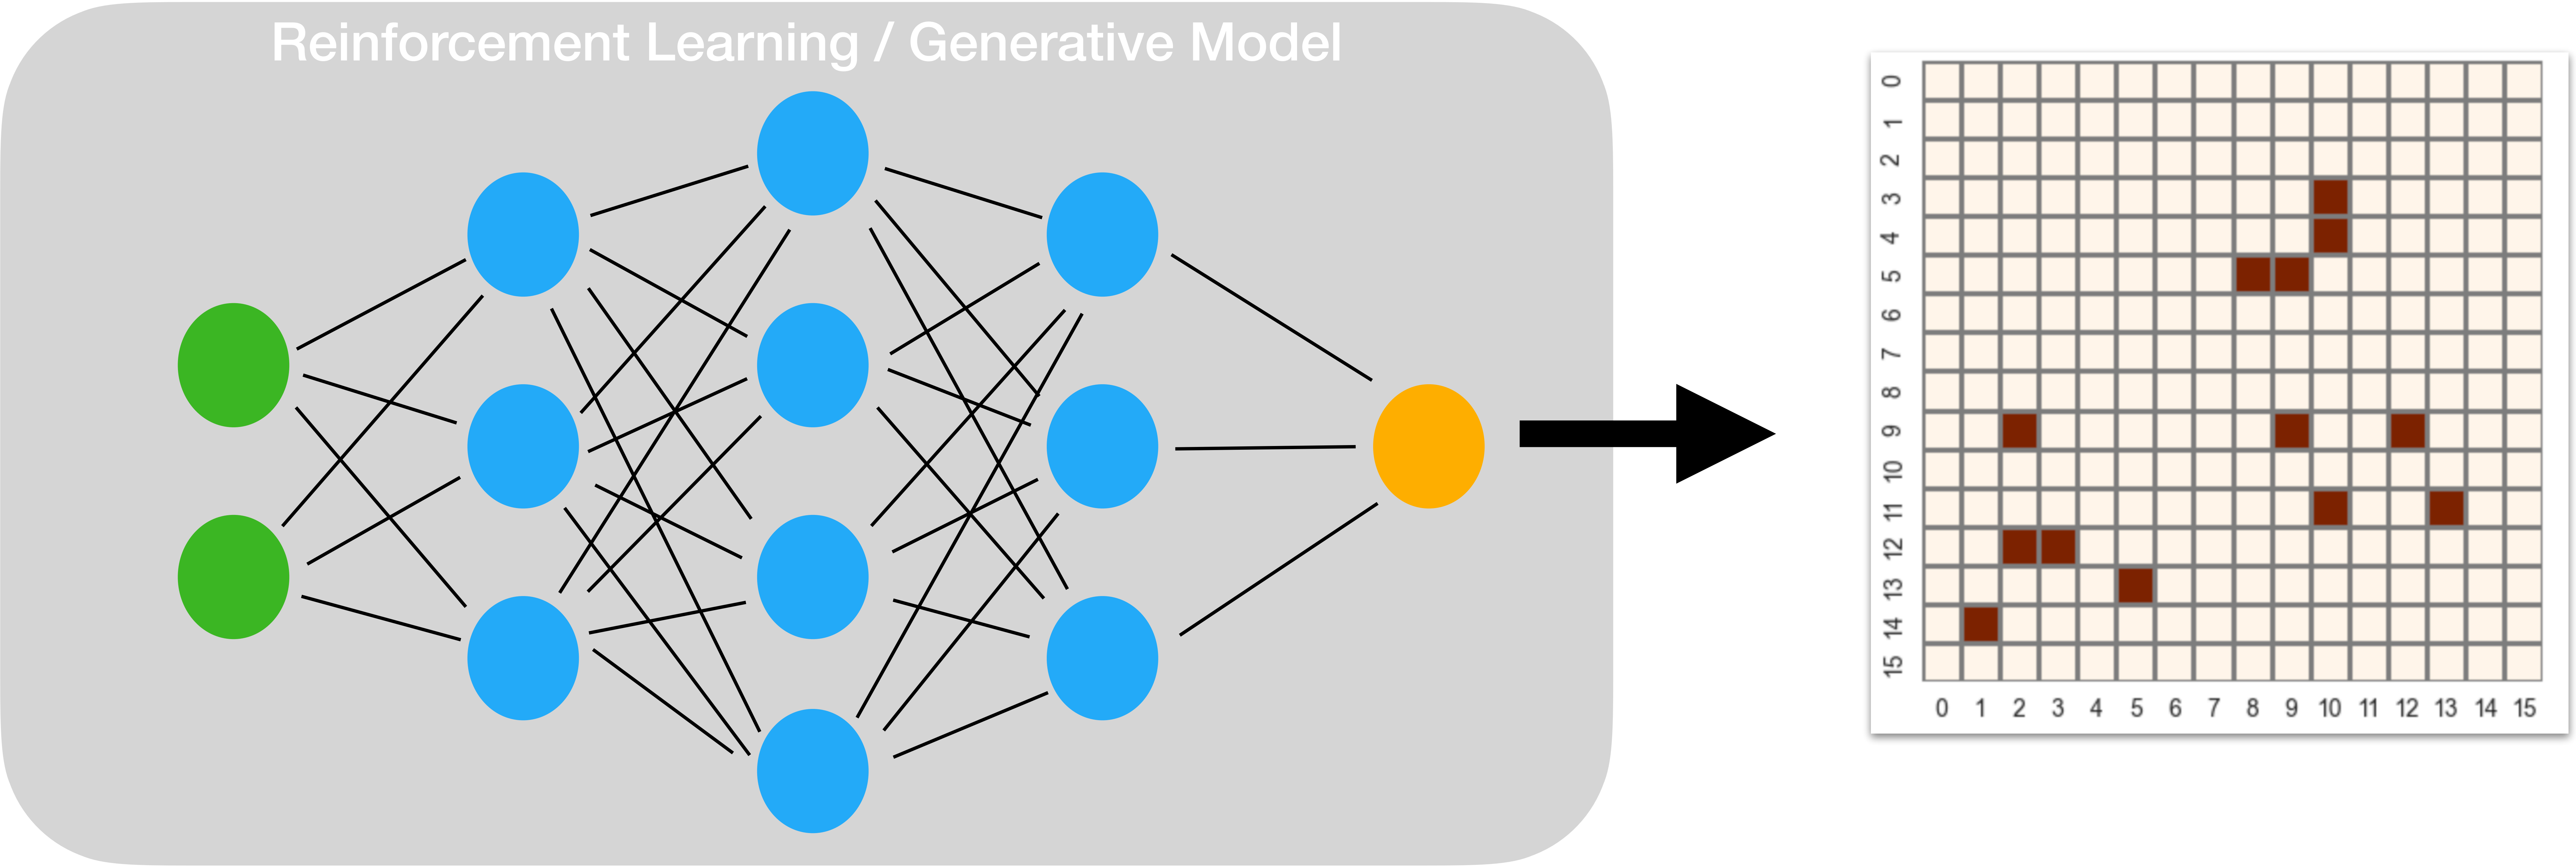
Reinforcement Learning / Generative Model
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15

Aim

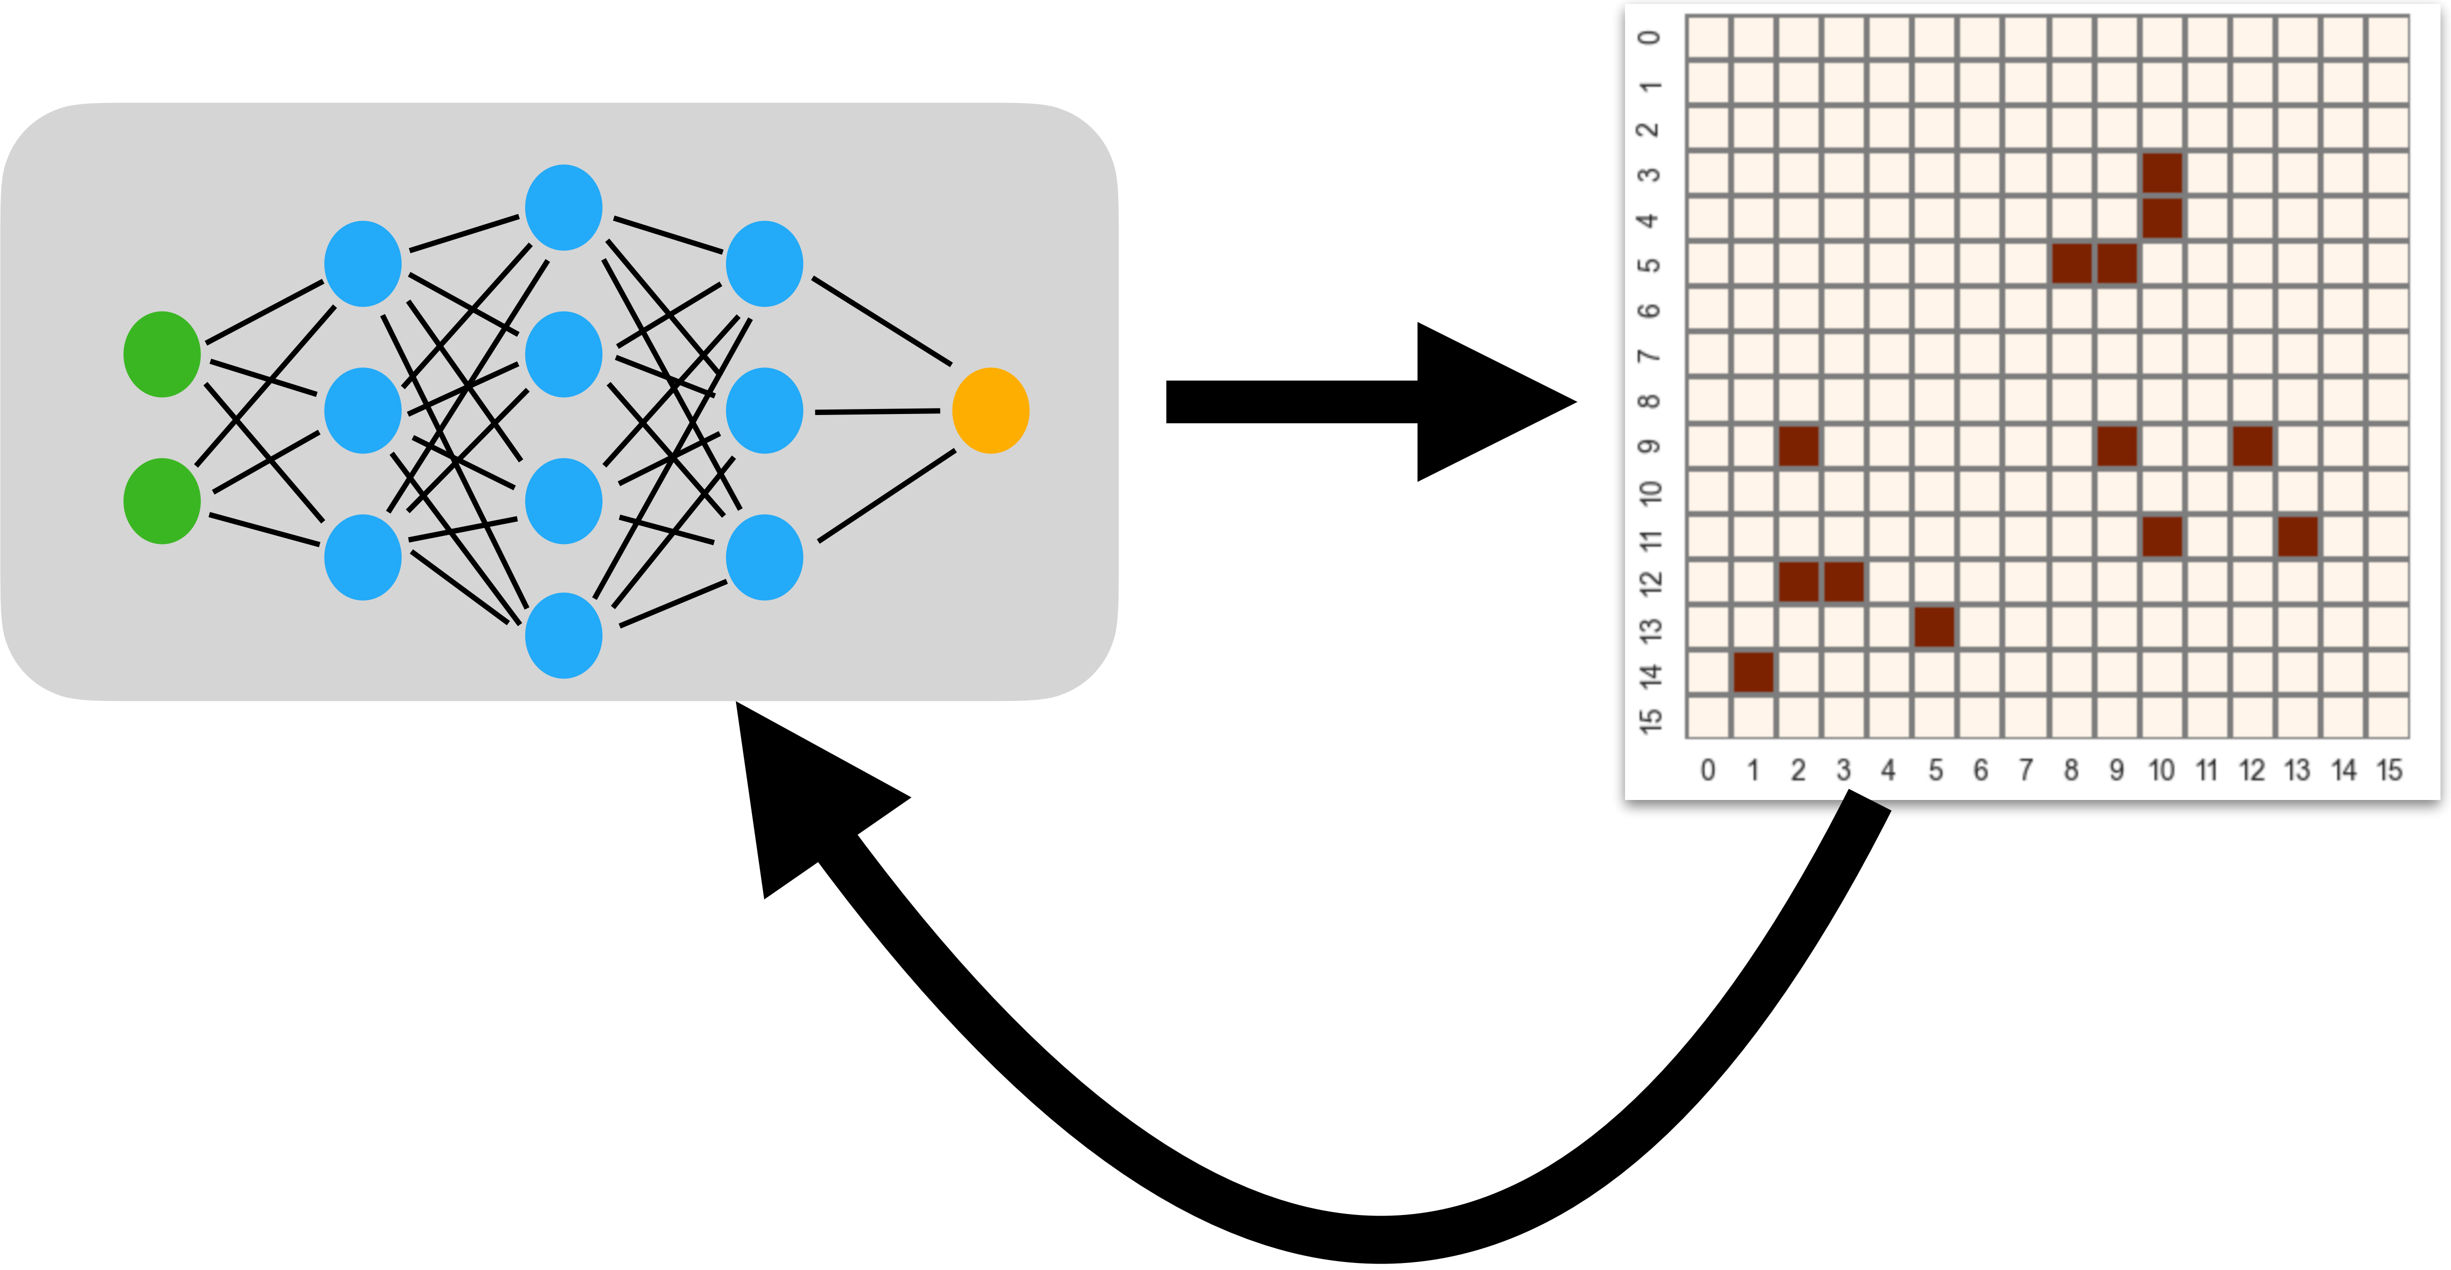

Aim

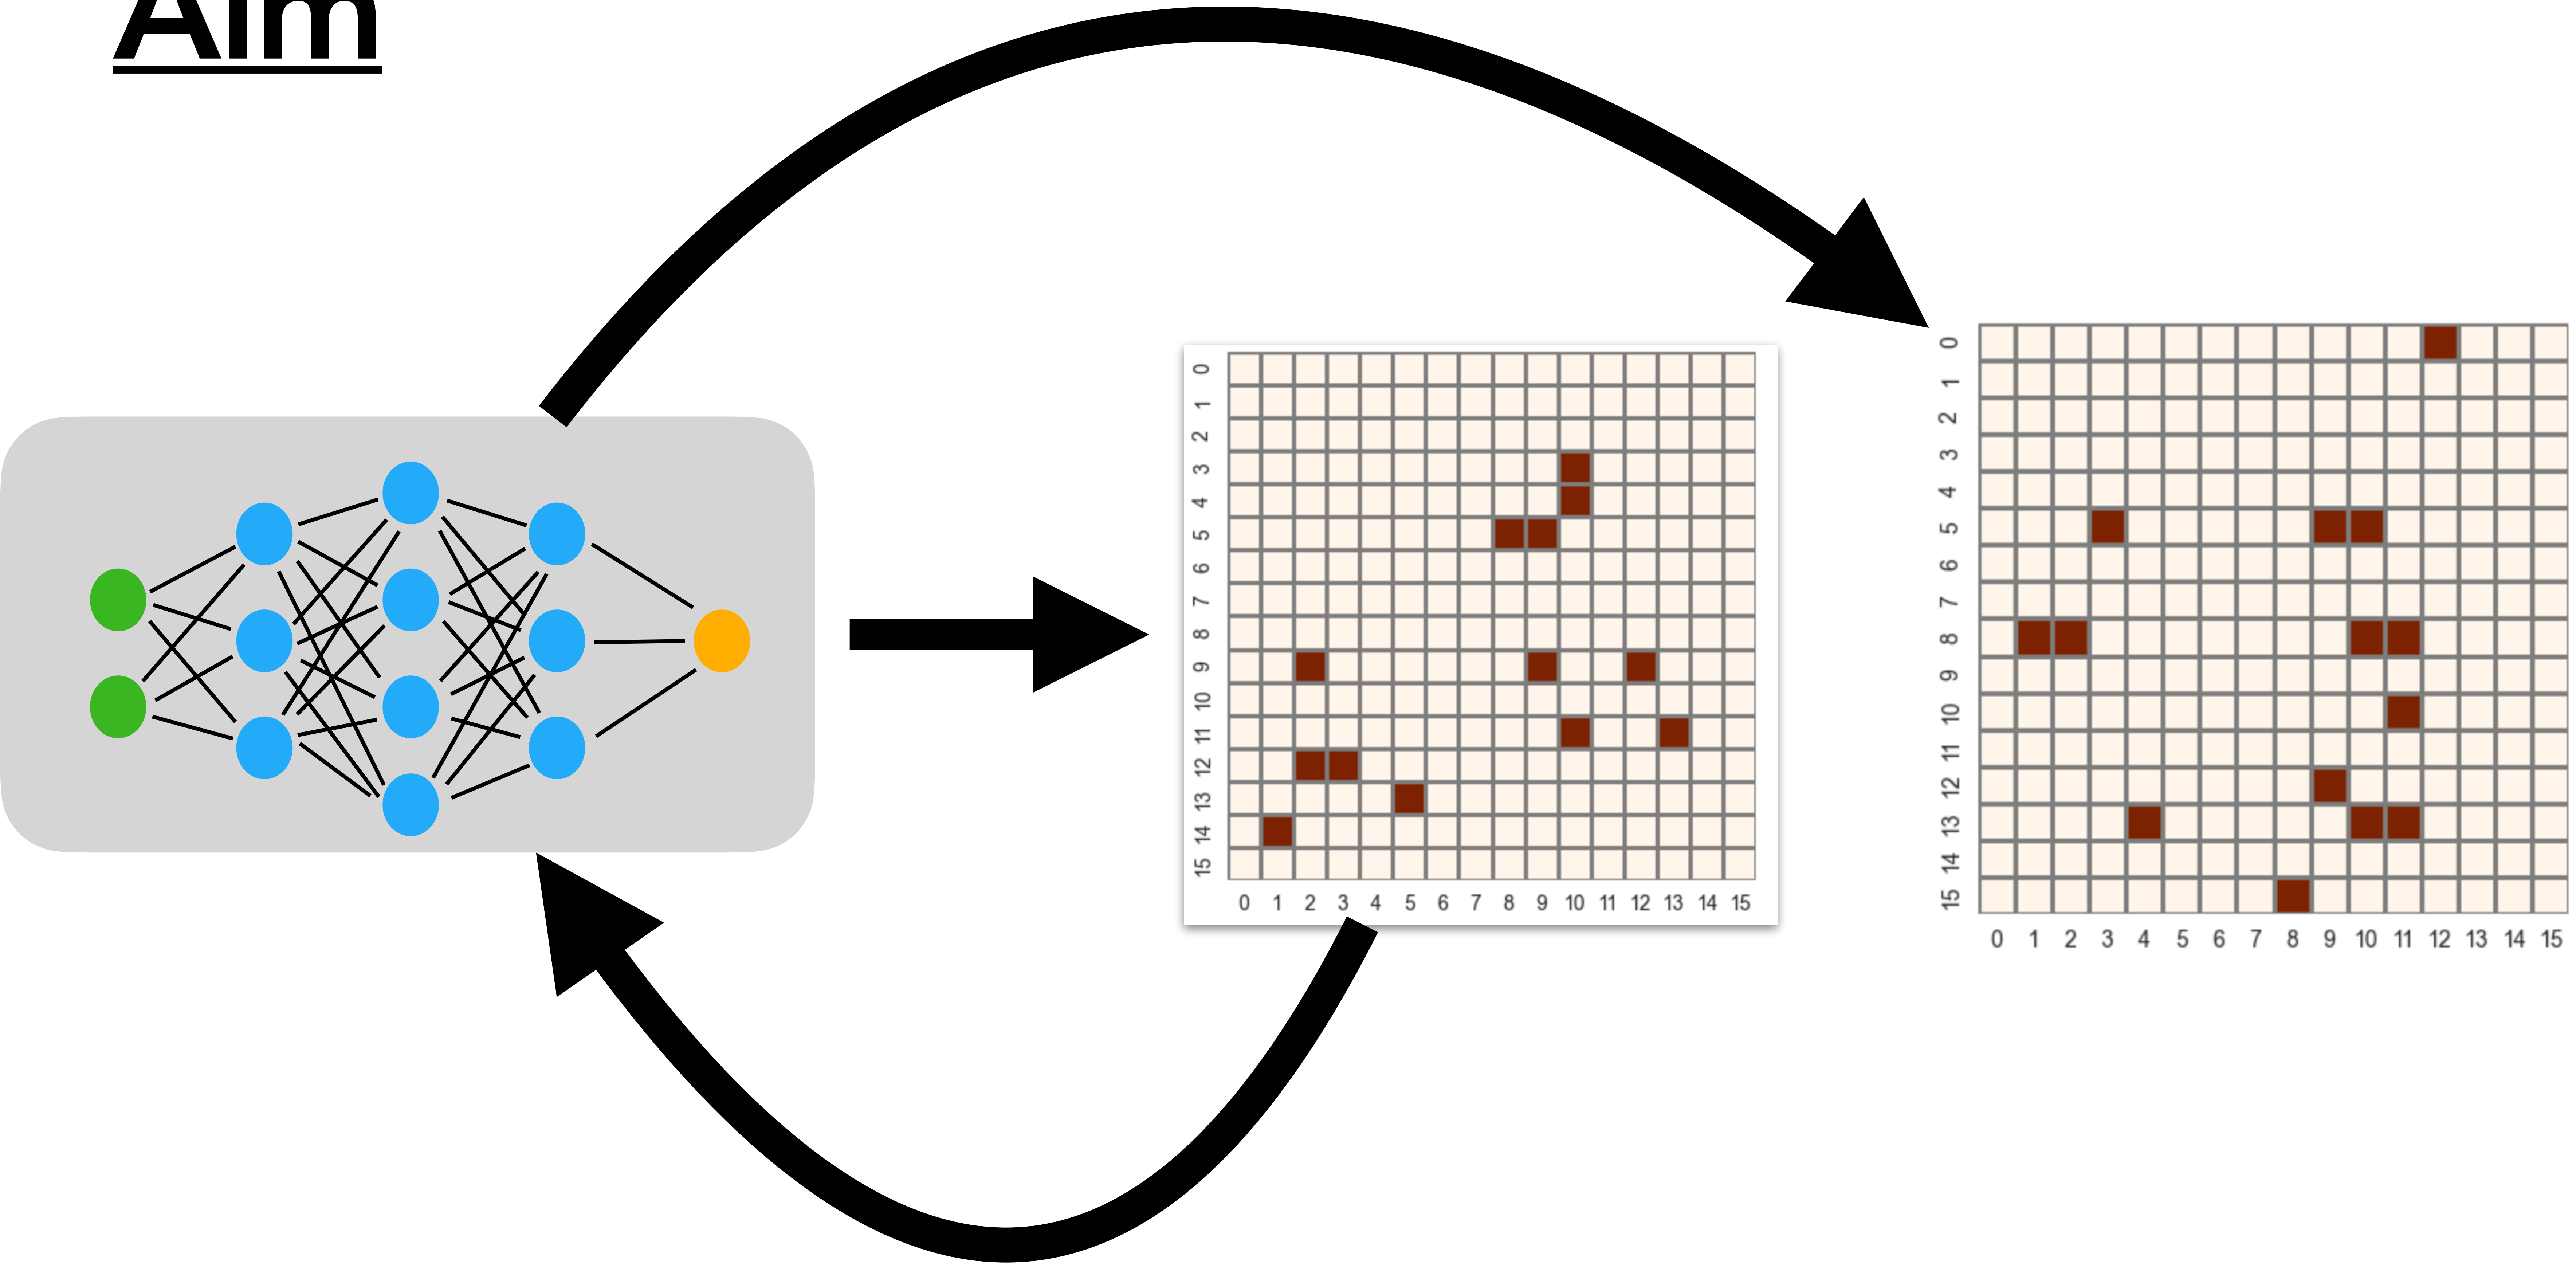
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15

Aim

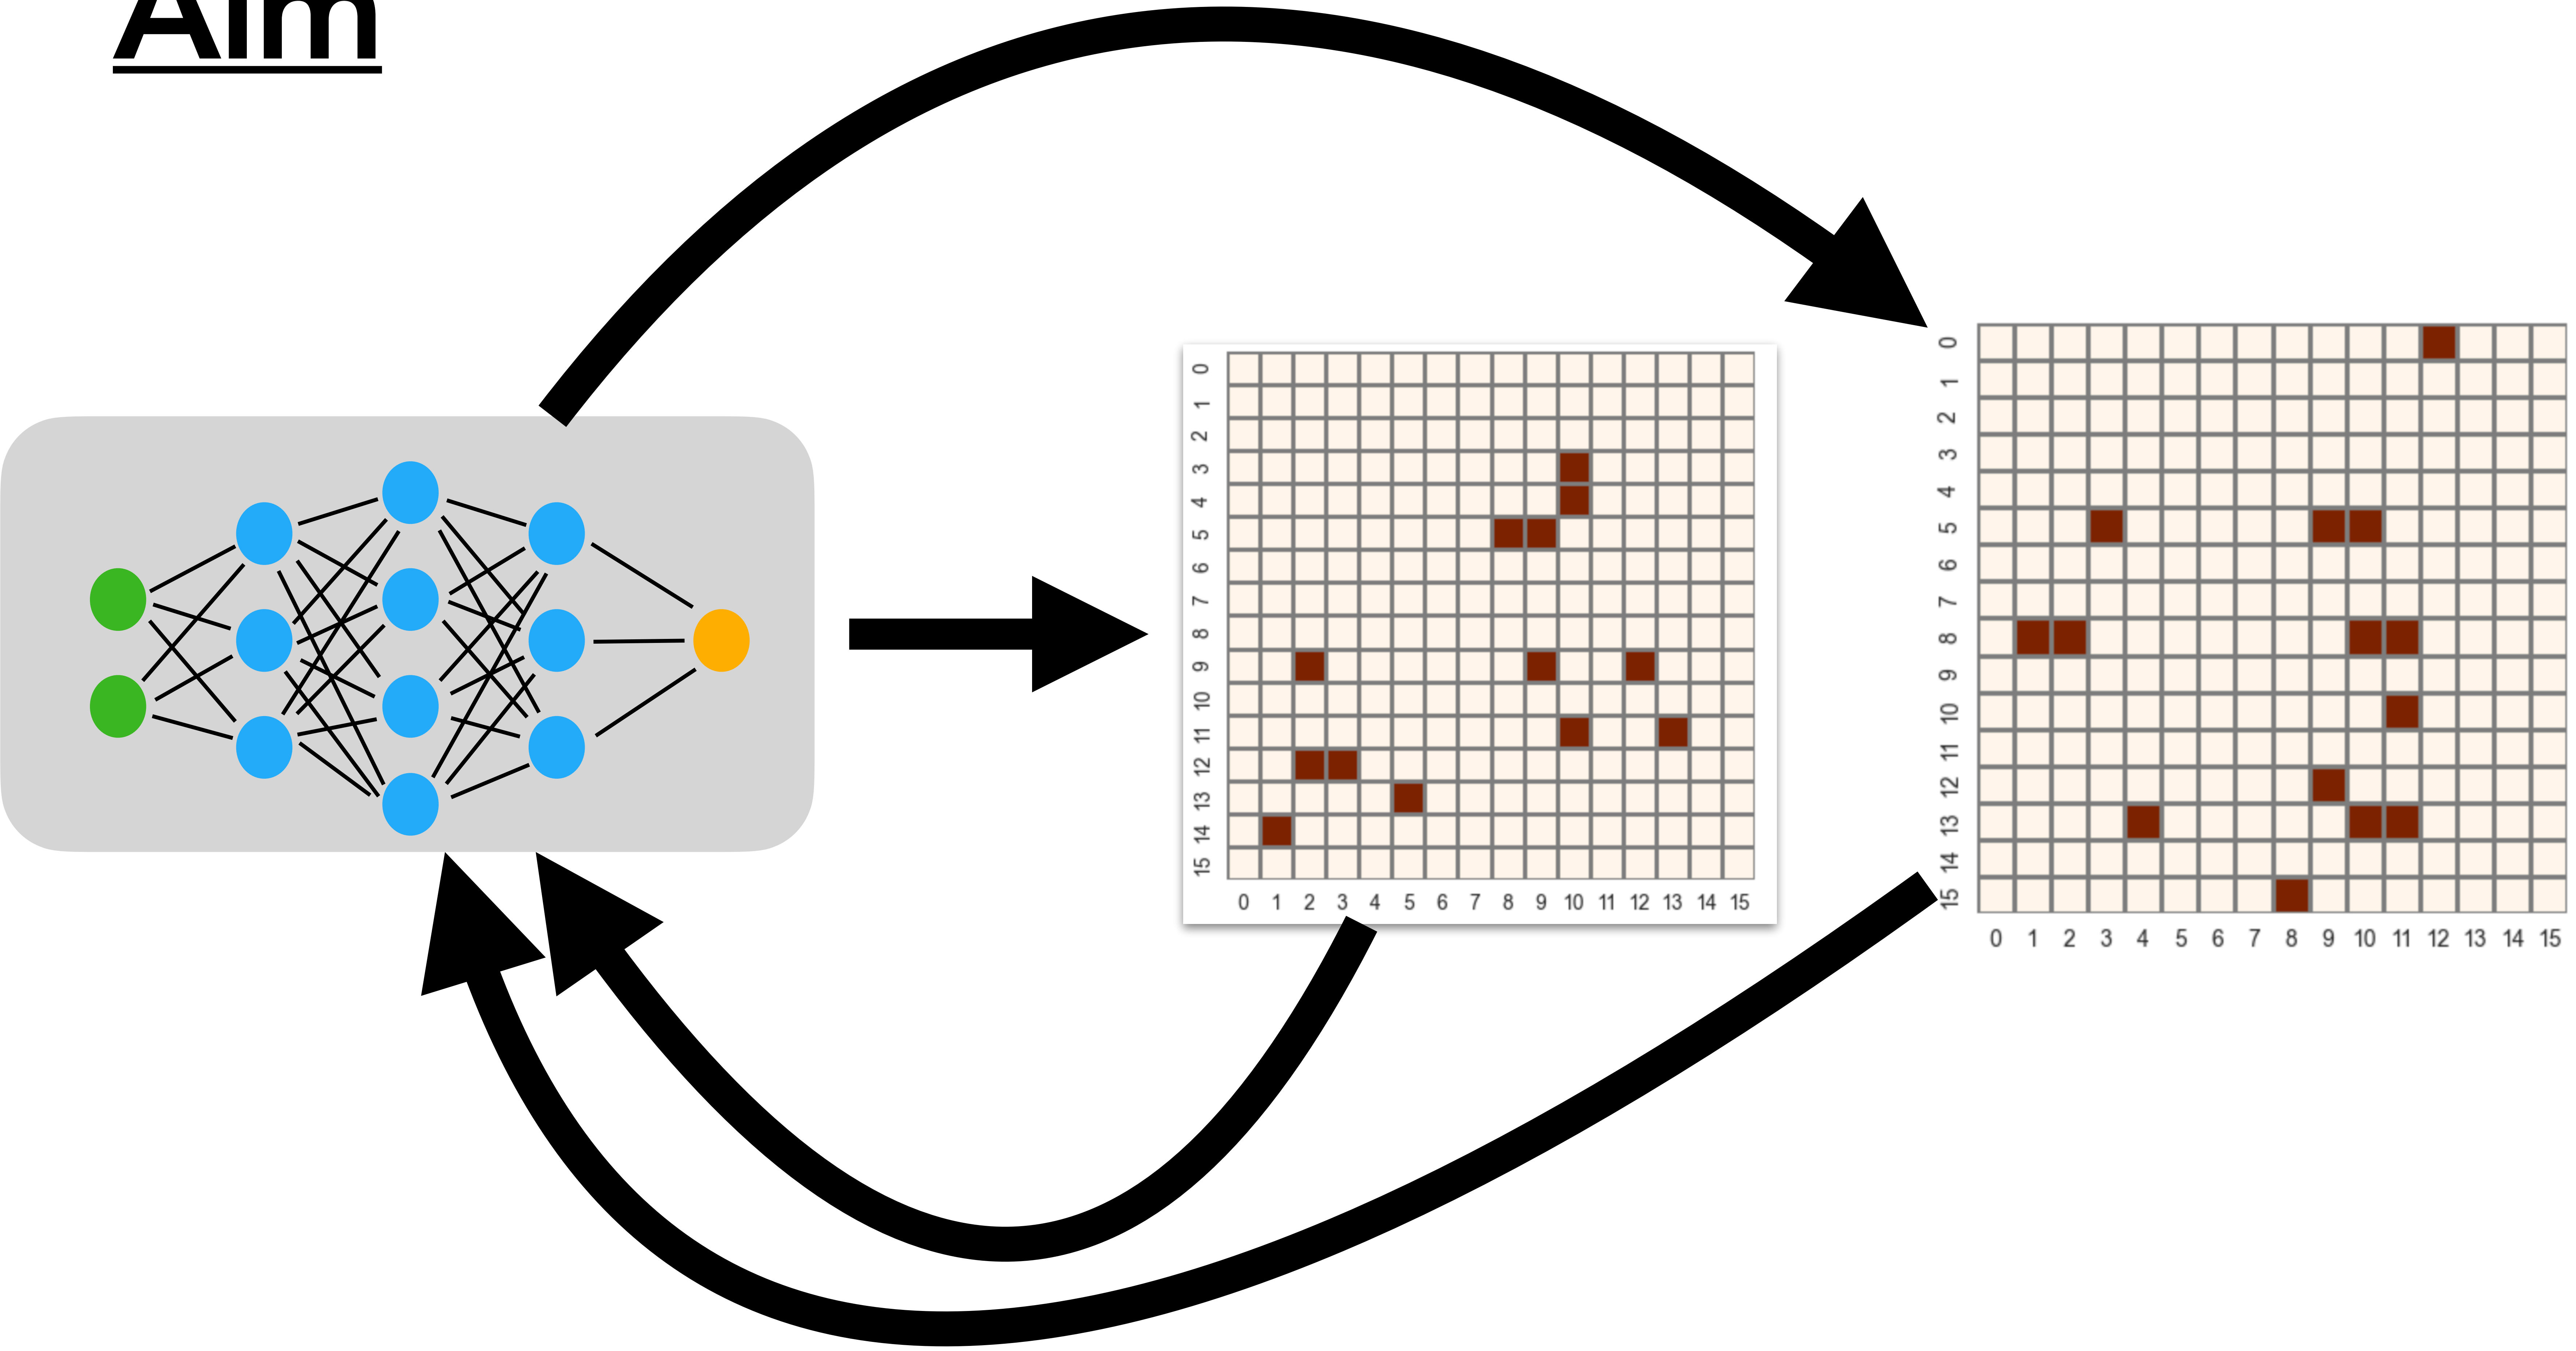
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15

Aim

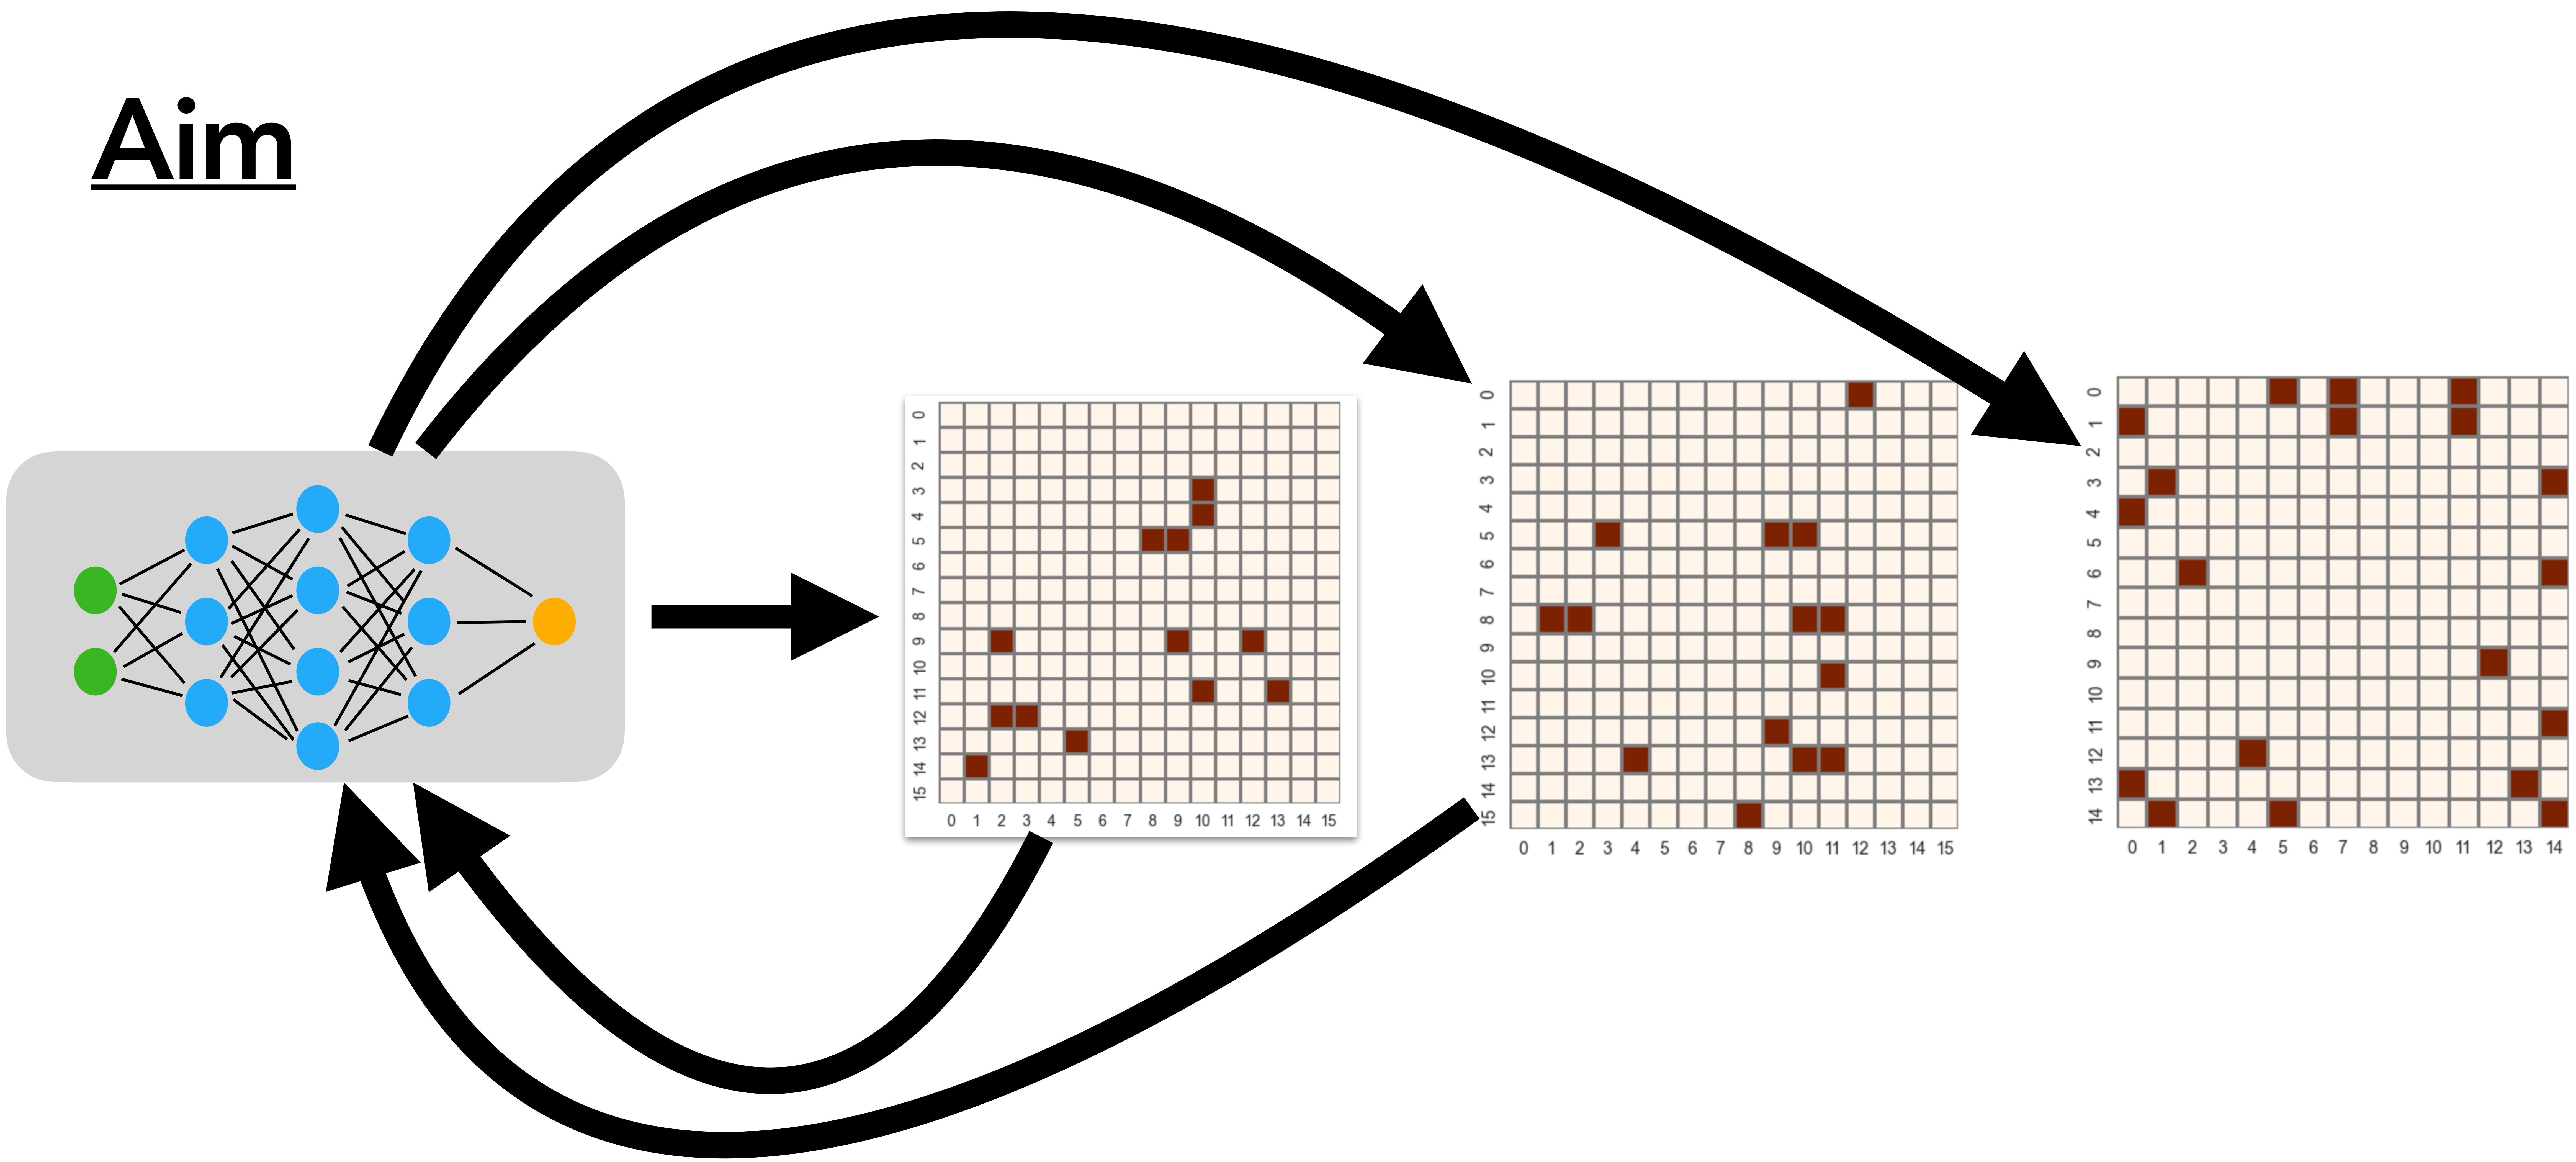
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14

My edits

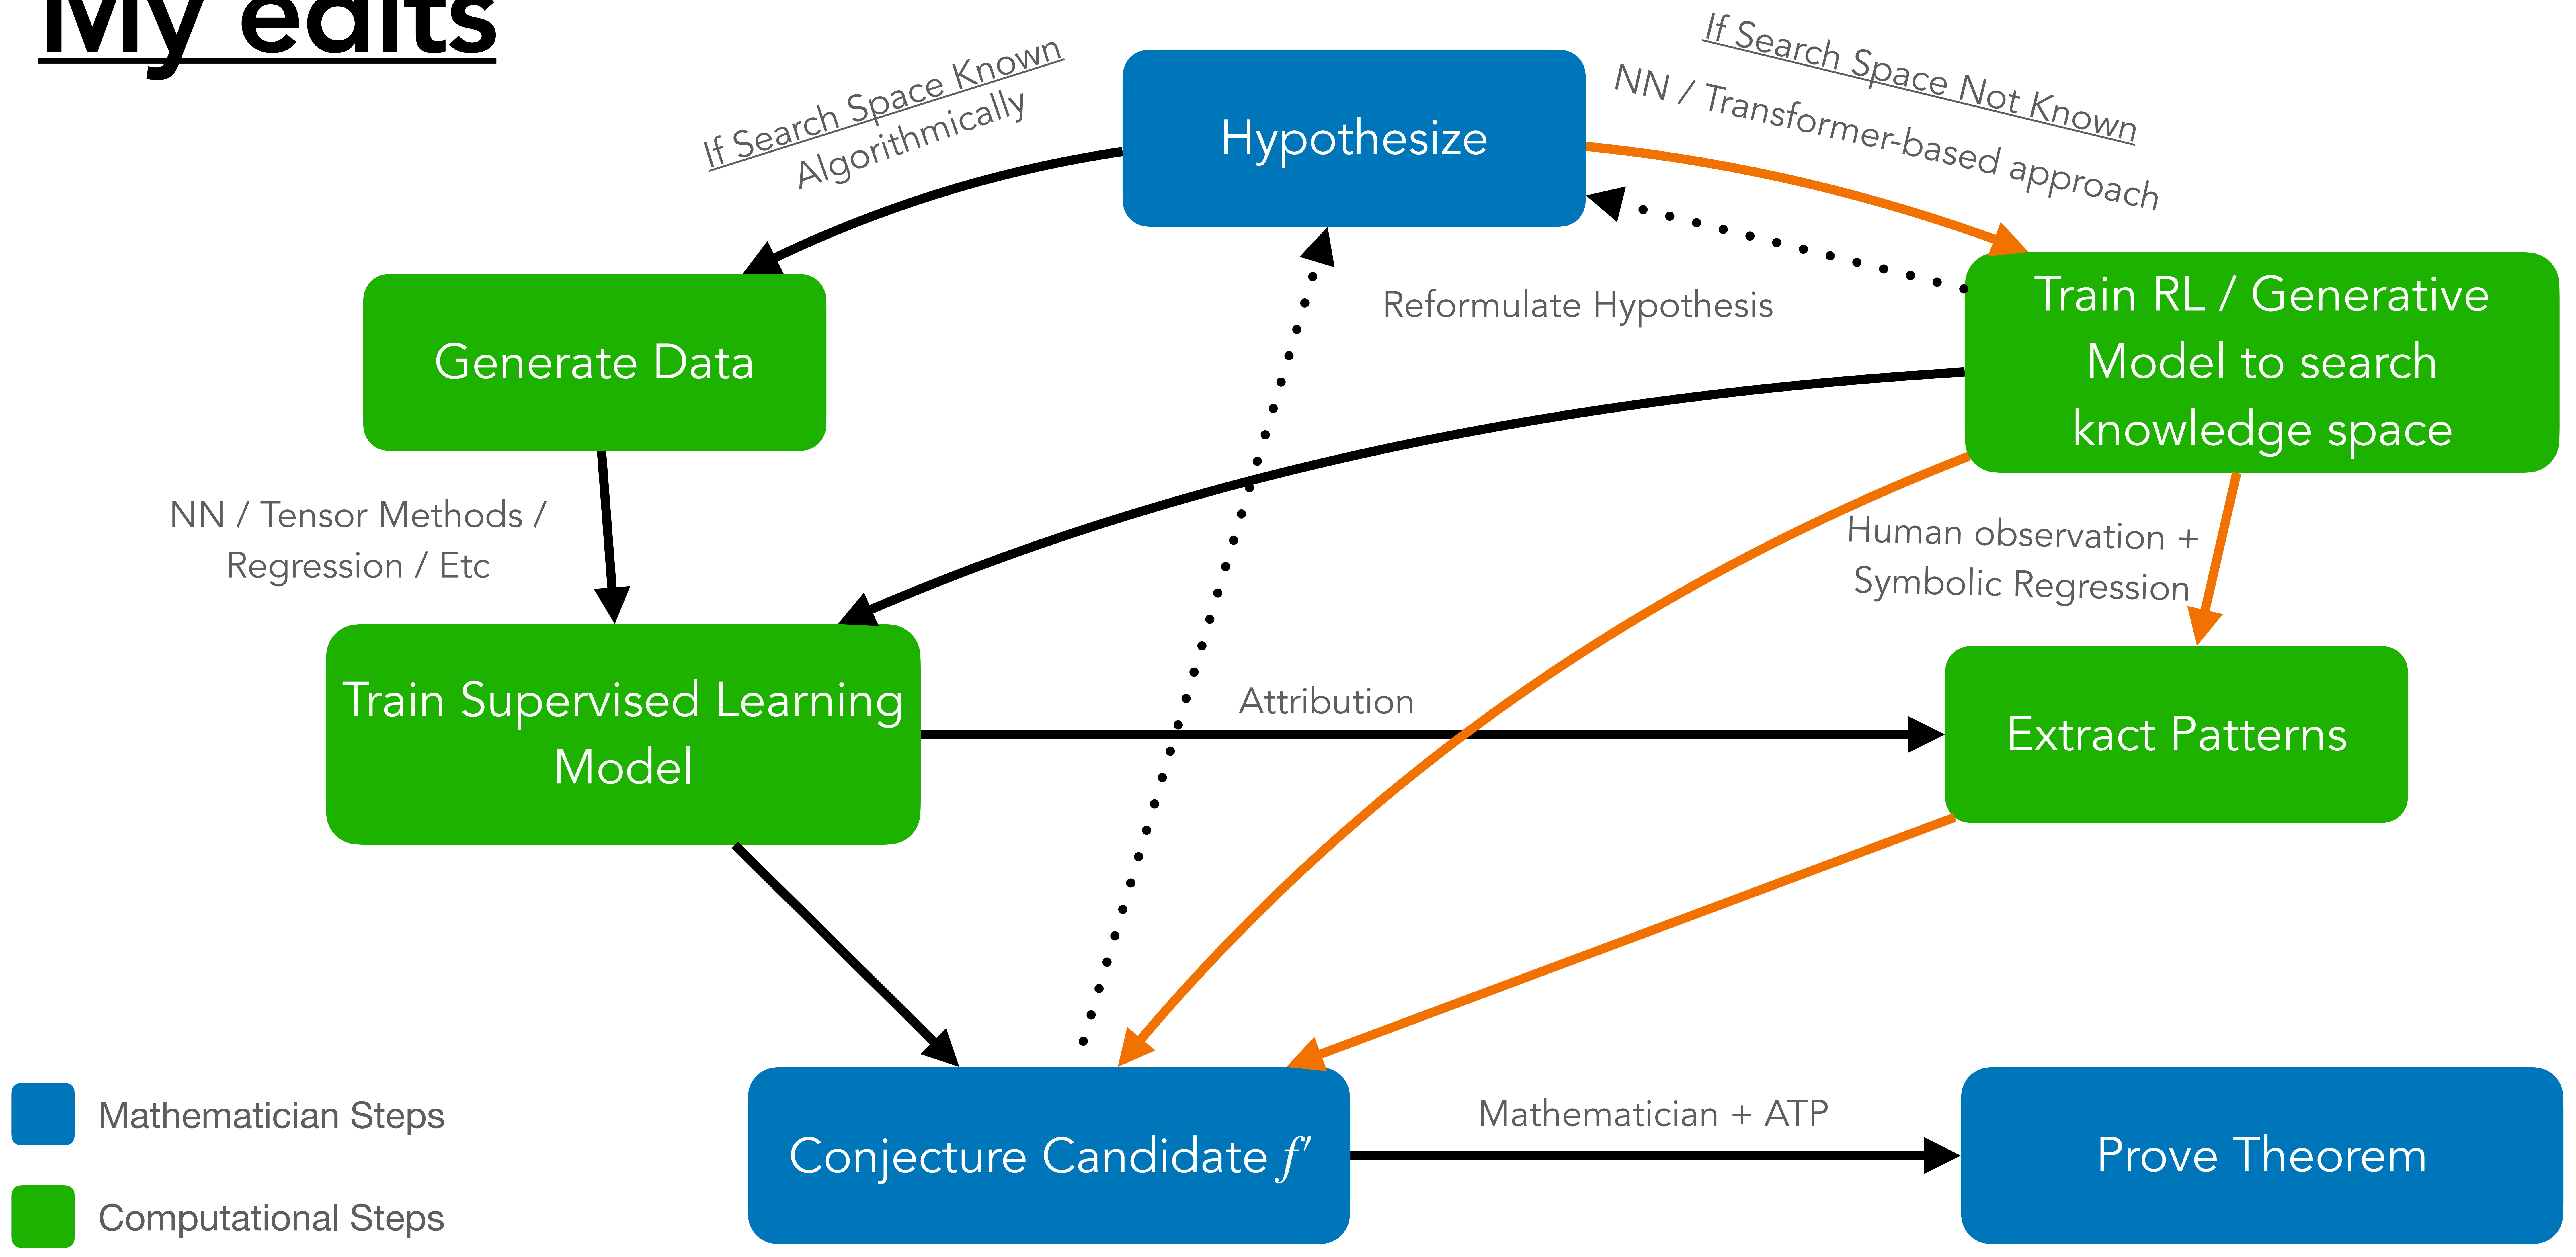

Hypothesize
If Search Space Known
Algorithmically
If Search Space Not Known
NN / Transformer-based approach
Generate Data
Reformulate Hypothesis
Train RL / Generative Model to search knowledge space
NN / Tensor Methods / Regression / Etc
Human observation + Symbolic Regression
Train Supervised Learning Model
Attribution
Extract Patterns
Conjecture Candidate f'
Mathematician + ATP
Prove Theorem
Mathematician Steps
Computational Steps

Things to keep in mind

Possible Problems

1. Reinforcement Learning requires very large amounts of data

Things to keep in mind

Possible Problems

1. Reinforcement Learning requires very large amounts of data
2. Reinforcement Learning is slow

Things to keep in mind

Possible Problems

1. Reinforcement Learning requires very large amounts of data
2. Reinforcement Learning is slow

Some perspectives on solutions

1. Math data is synthetic, so we can synthesise a lot!

Things to keep in mind

Possible Problems

1. Reinforcement Learning requires very large amounts of data
2. Reinforcement Learning is slow

Some perspectives on solutions

1. Math data is synthetic, so we can synthesise a lot!
2. My university has lots of Nvidia GPUs!

Things to keep in mind

Possible Problems

1. Reinforcement Learning requires very large amounts of data
2. Reinforcement Learning is slow

Some perspectives on solutions

1. Math data is synthetic, so we can synthesise a lot!
2. My university has lots of Nvidia GPUs! (But yes, we do need to engineer carefully).

Cool Papers

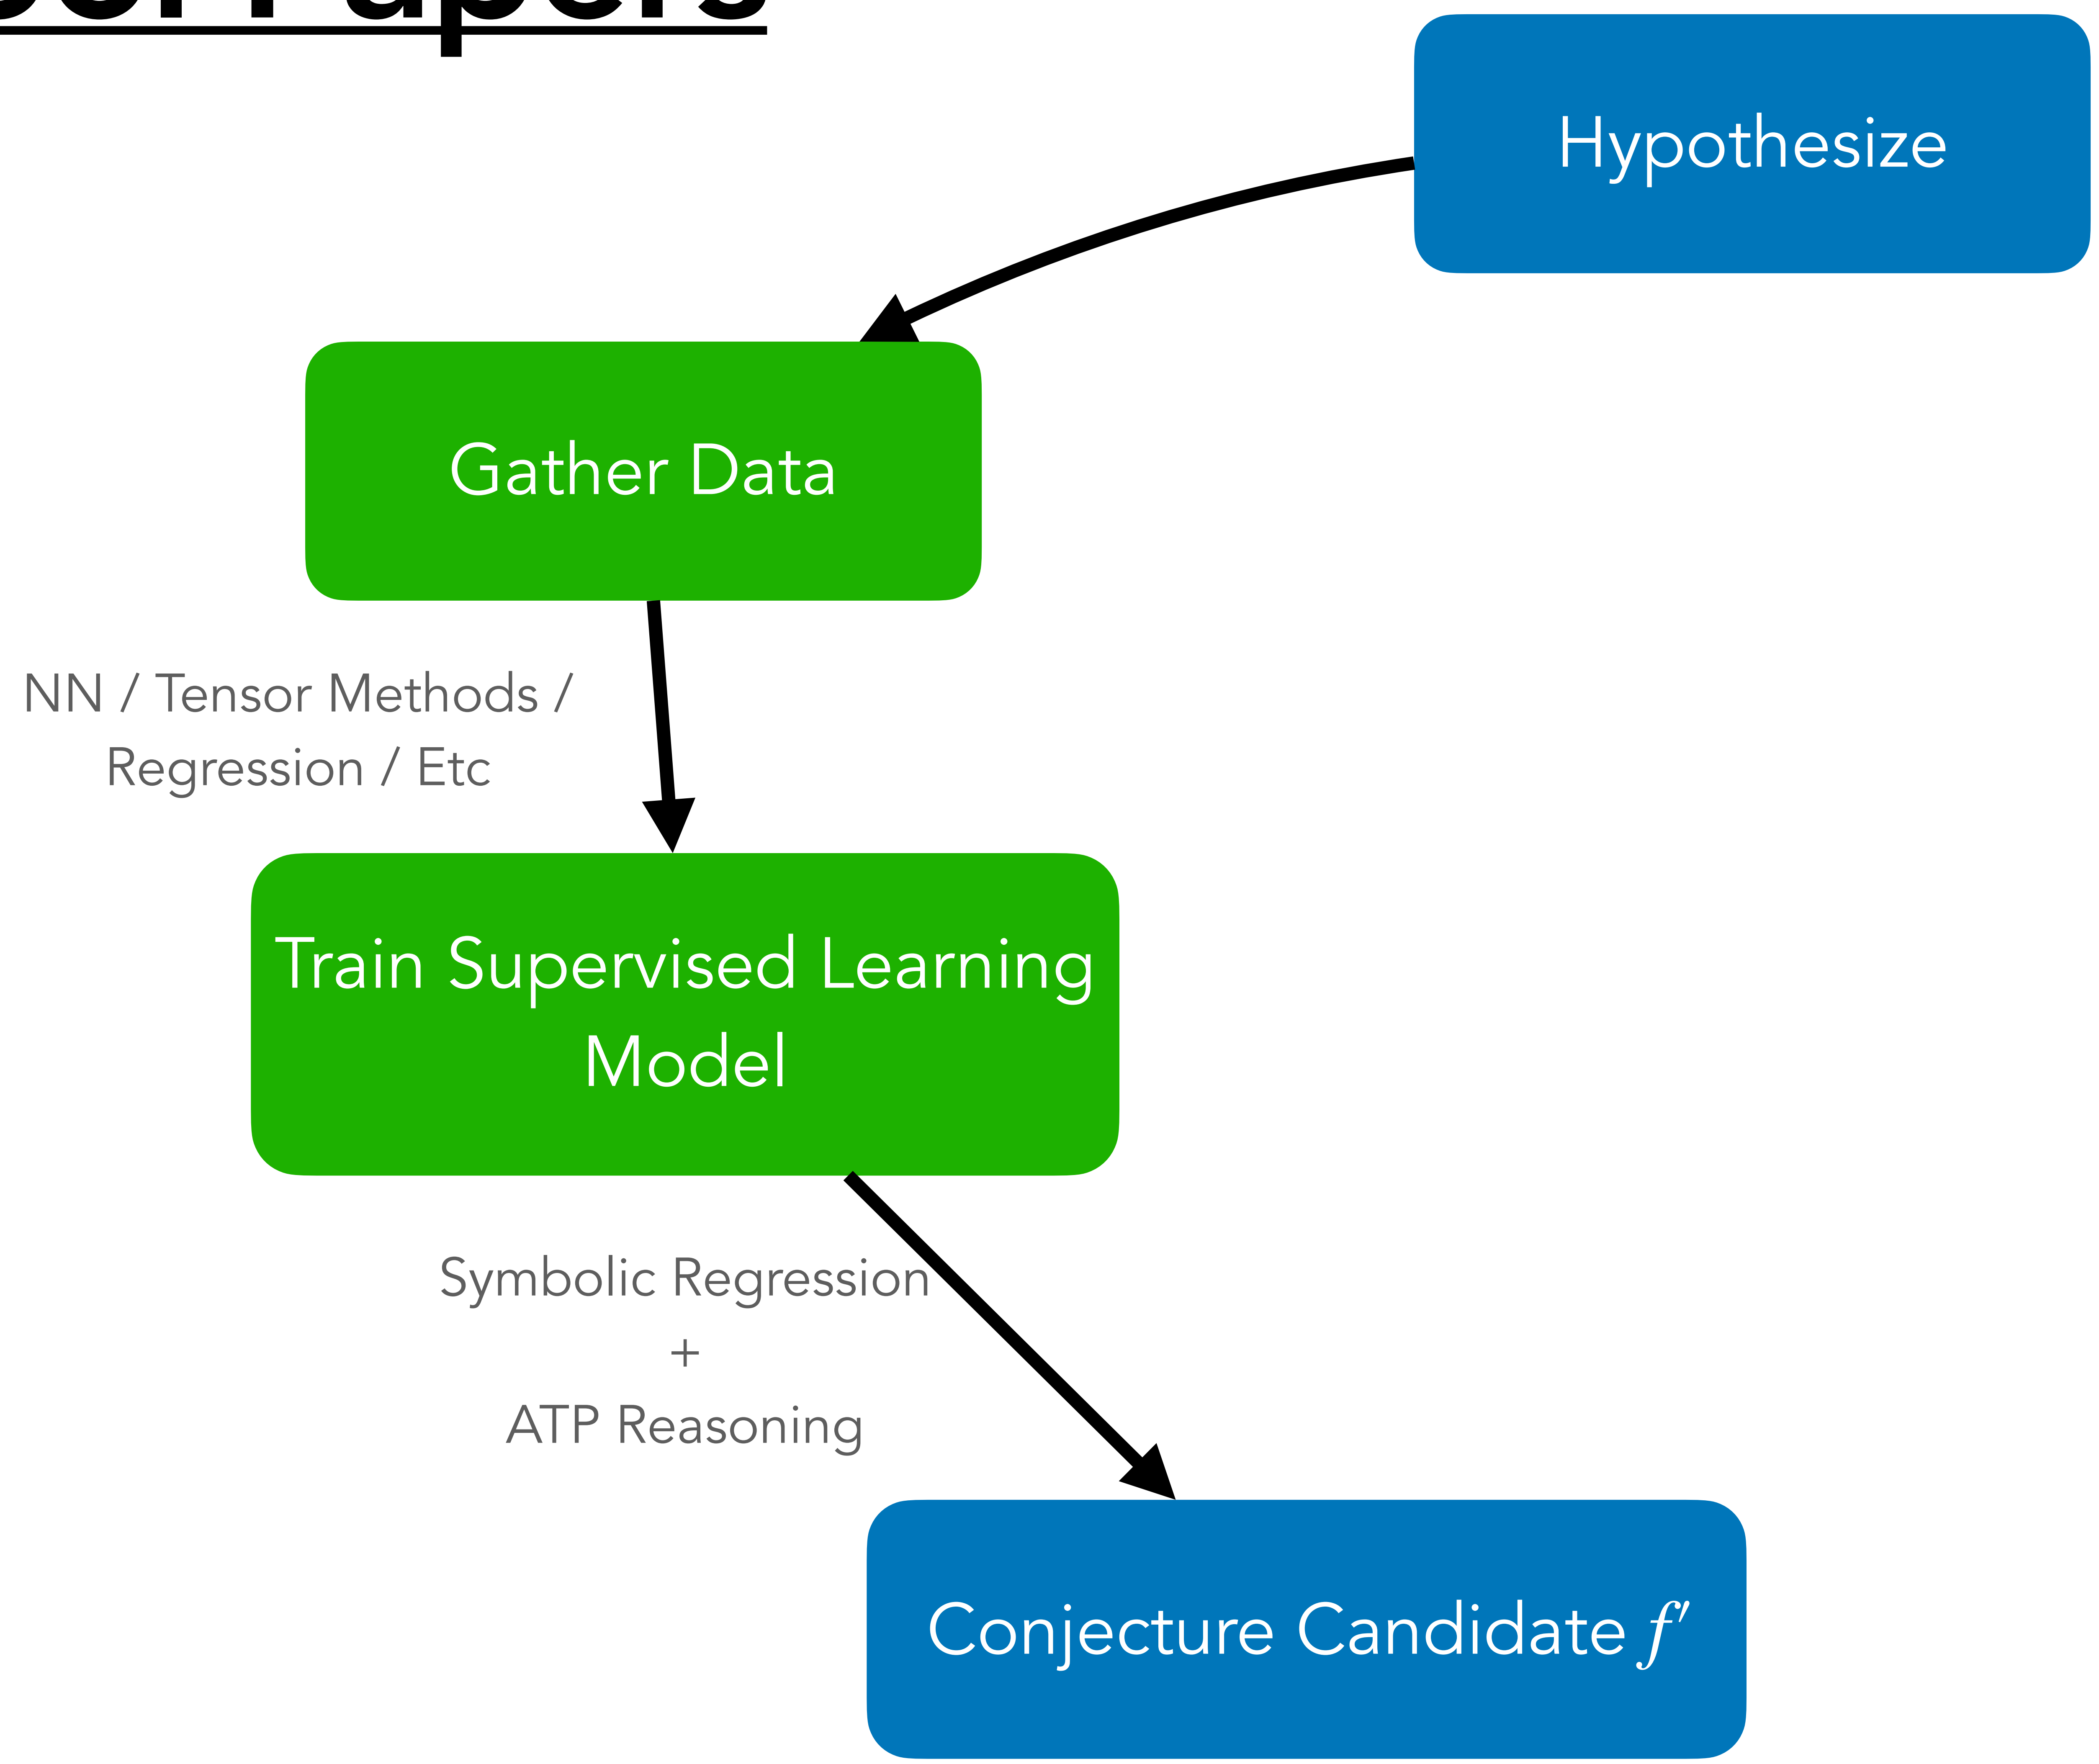

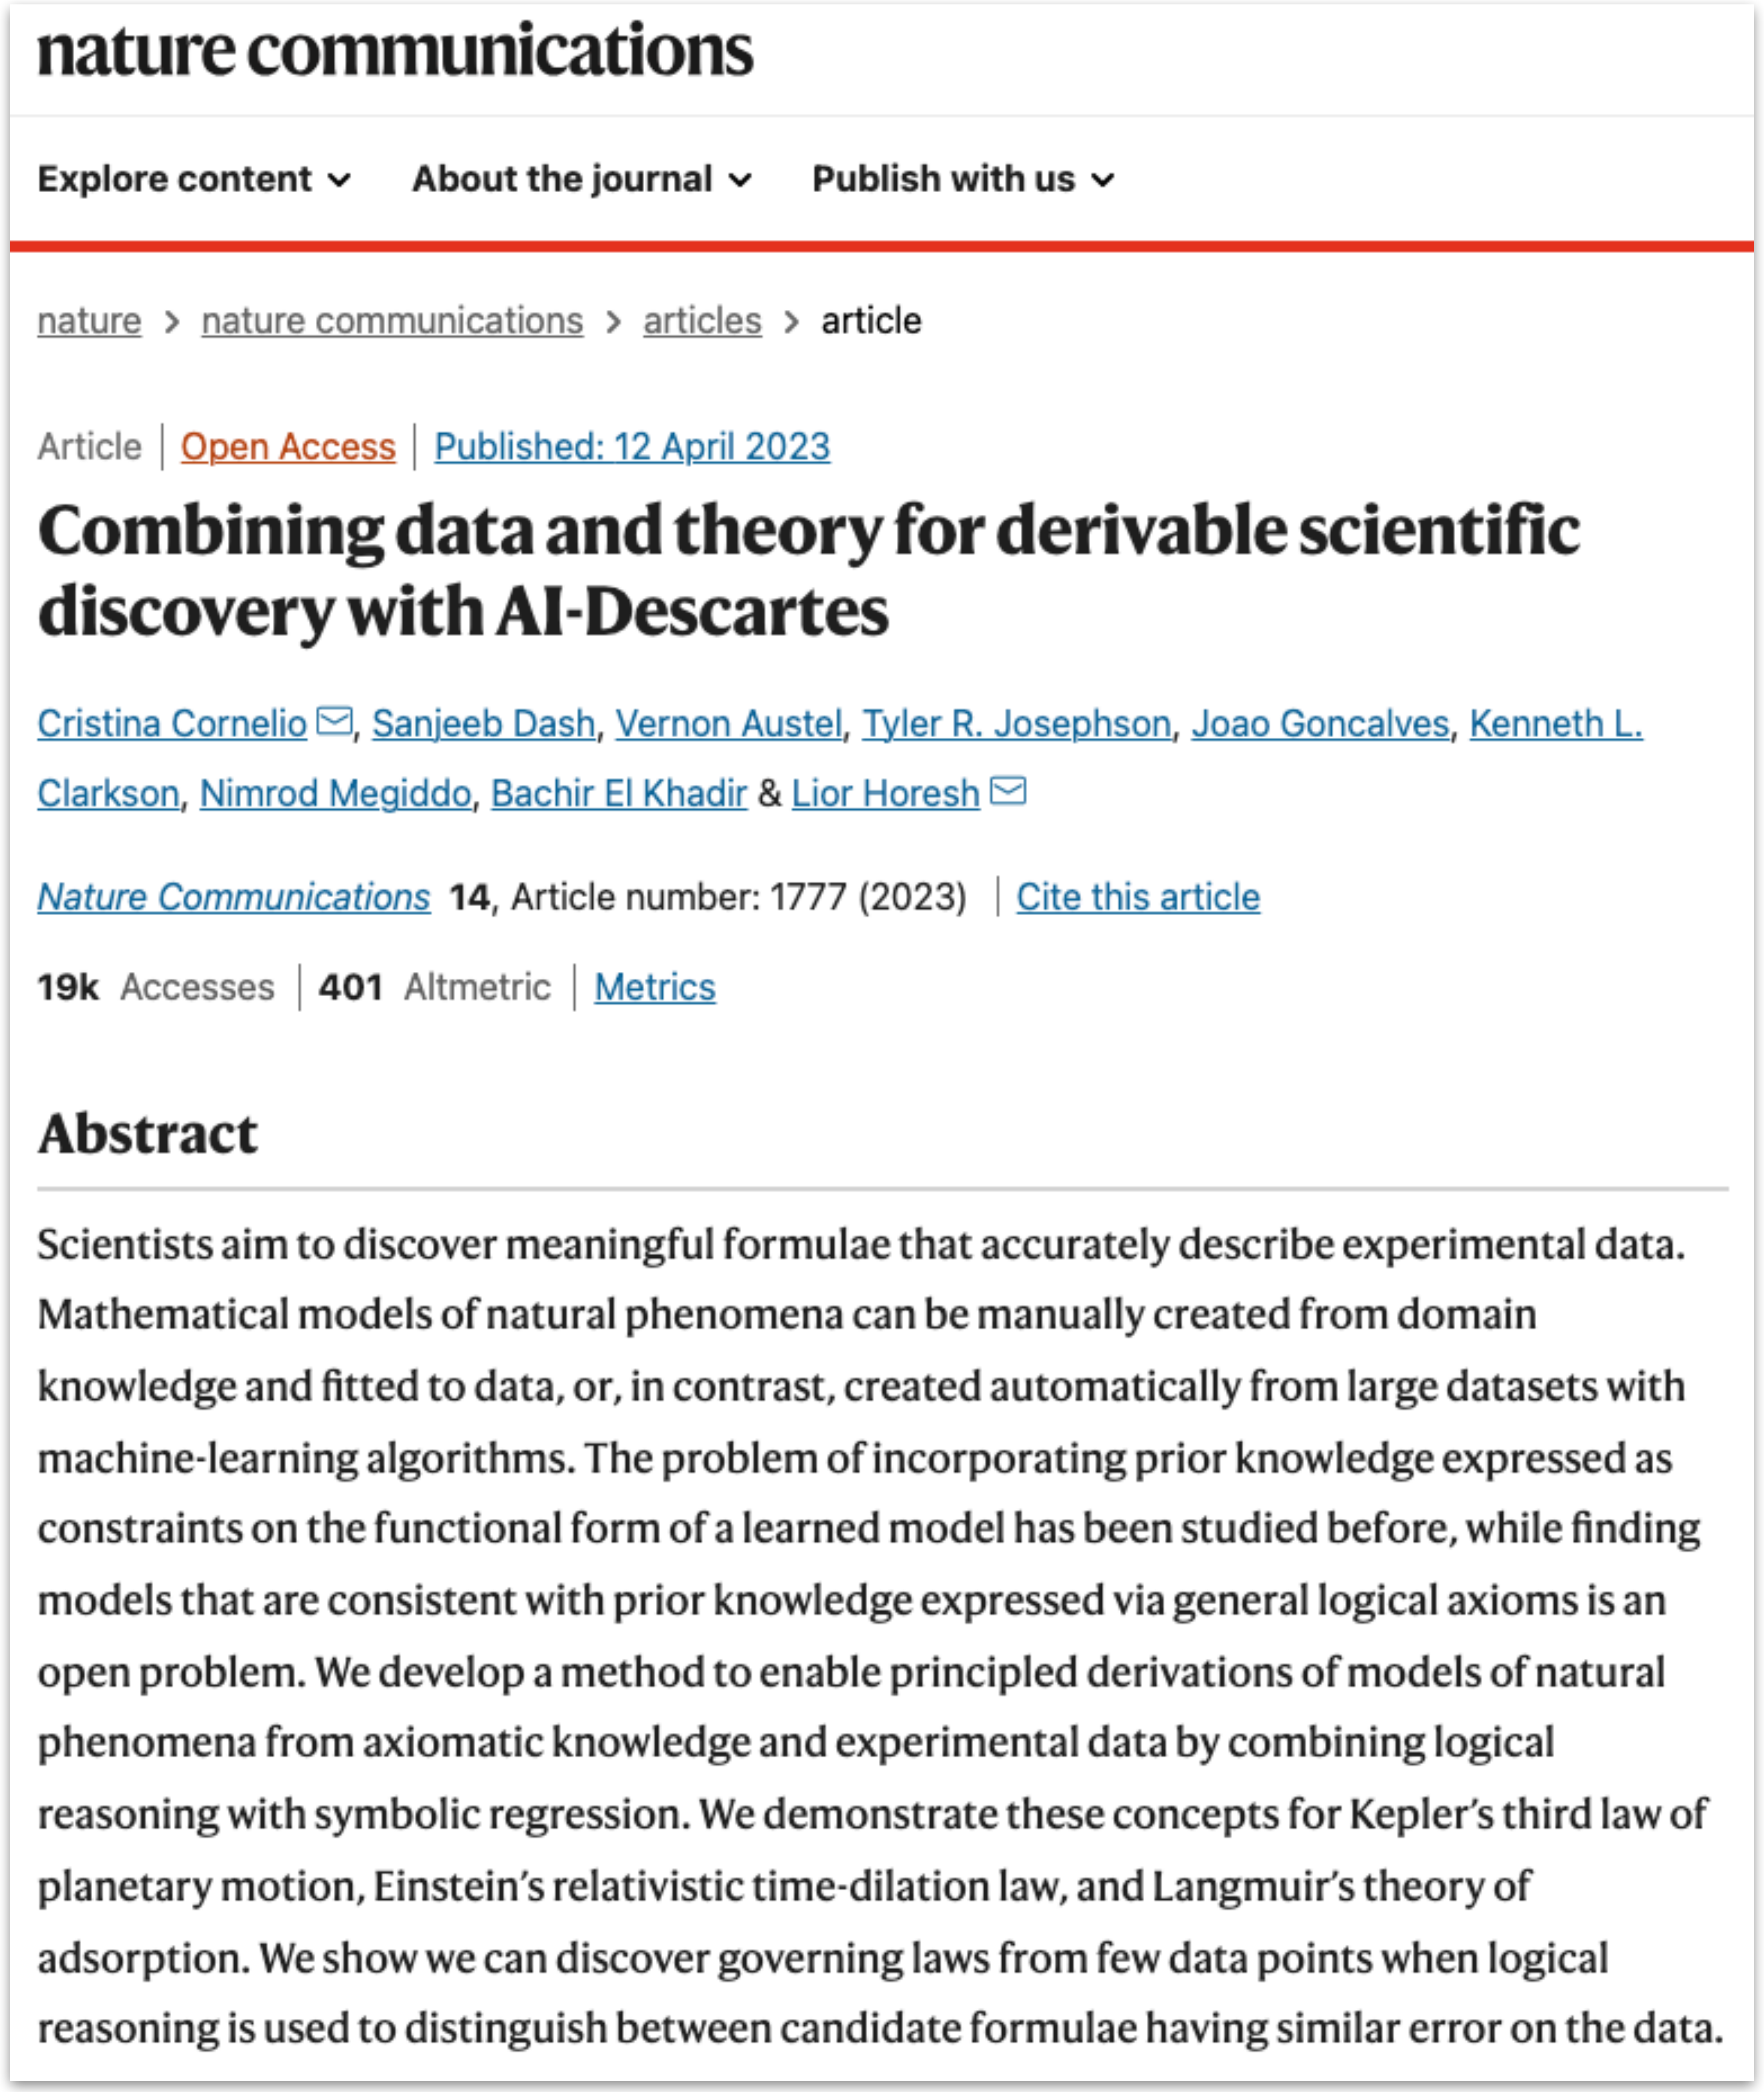
nature communications

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature > nature communications > articles > article

Article | Open Access | Published: 12 April 2023

Combining data and theory for derivable scientific discovery with AI-Descartes

Cristina Cornelio, Sanjeeb Dash, Vernon Austel, Tyler R. Josephson, Joao Goncalves, Kenneth L. Clarkson, Nimrod Megiddo, Bachir El Khadir & Lior Horesh

Nature Communications 14, Article number: 1777 (2023) | Cite this article

Abstract

Scientists aim to discover meaningful formulae that accurately describe experimental data. Mathematical models of natural phenomena can be manually created from domain knowledge and fitted to data, or, in contrast, created automatically from large datasets with machine-learning algorithms. The problem of incorporating prior knowledge expressed as constraints on the functional form of a learned model has been studied before, while finding models that are consistent with prior knowledge expressed via general logical axioms is an open problem. We develop a method to enable principled derivations of models of natural phenomena from axiomatic knowledge and experimental data by combining logical reasoning with symbolic regression. We demonstrate these concepts for Kepler's third law of planetary motion, Einstein's relativistic time-dilation law, and Langmuir's theory of adsorption. We show we can discover governing laws from few data points when logical reasoning is used to distinguish between candidate formulae having similar error on the data.

Cool Papers

Constructions in combinatorics via neural networks

Adam Zsolt Wagner*

Abstract

We demonstrate how by using a reinforcement learning algorithm, the deep cross-entropy method, one can find explicit constructions and counterexamples to several open conjectures in extremal combinatorics and graph theory. Amongst the conjectures we refute are a question of Brualdi and Cao about maximizing permanents of pattern avoiding matrices, and several problems related to the adjacency and distance eigenvalues of graphs.

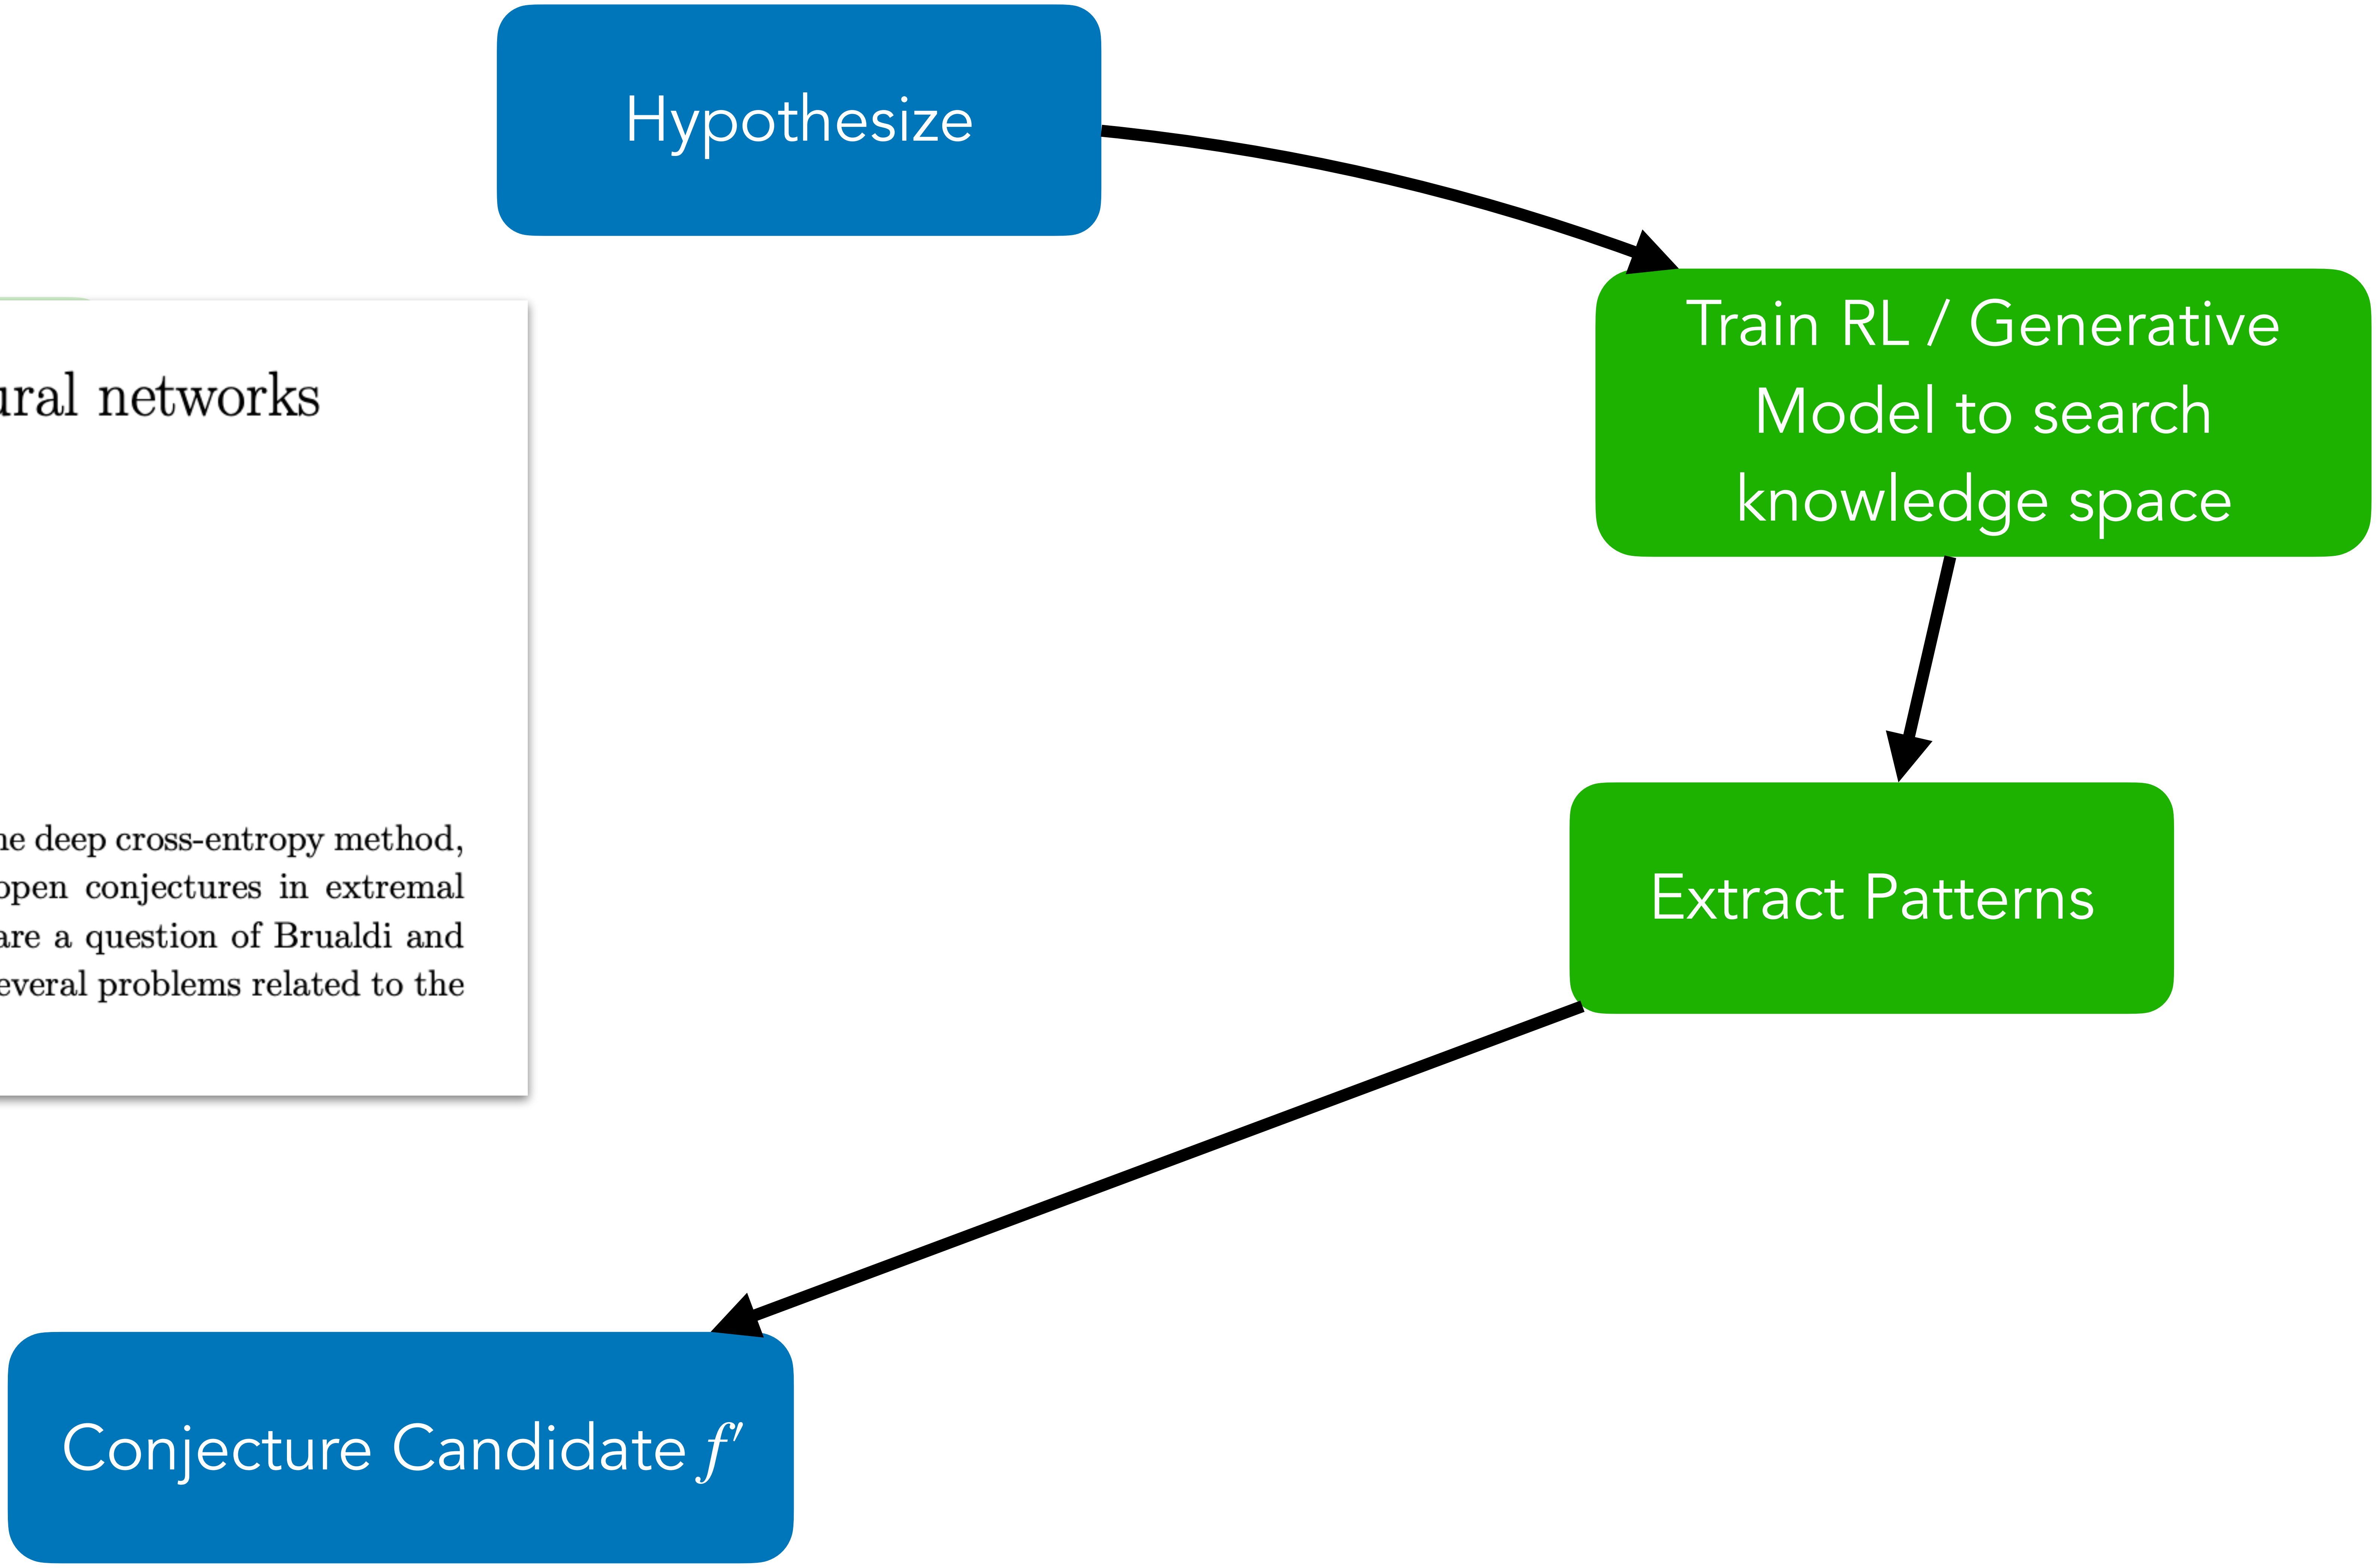

My edits

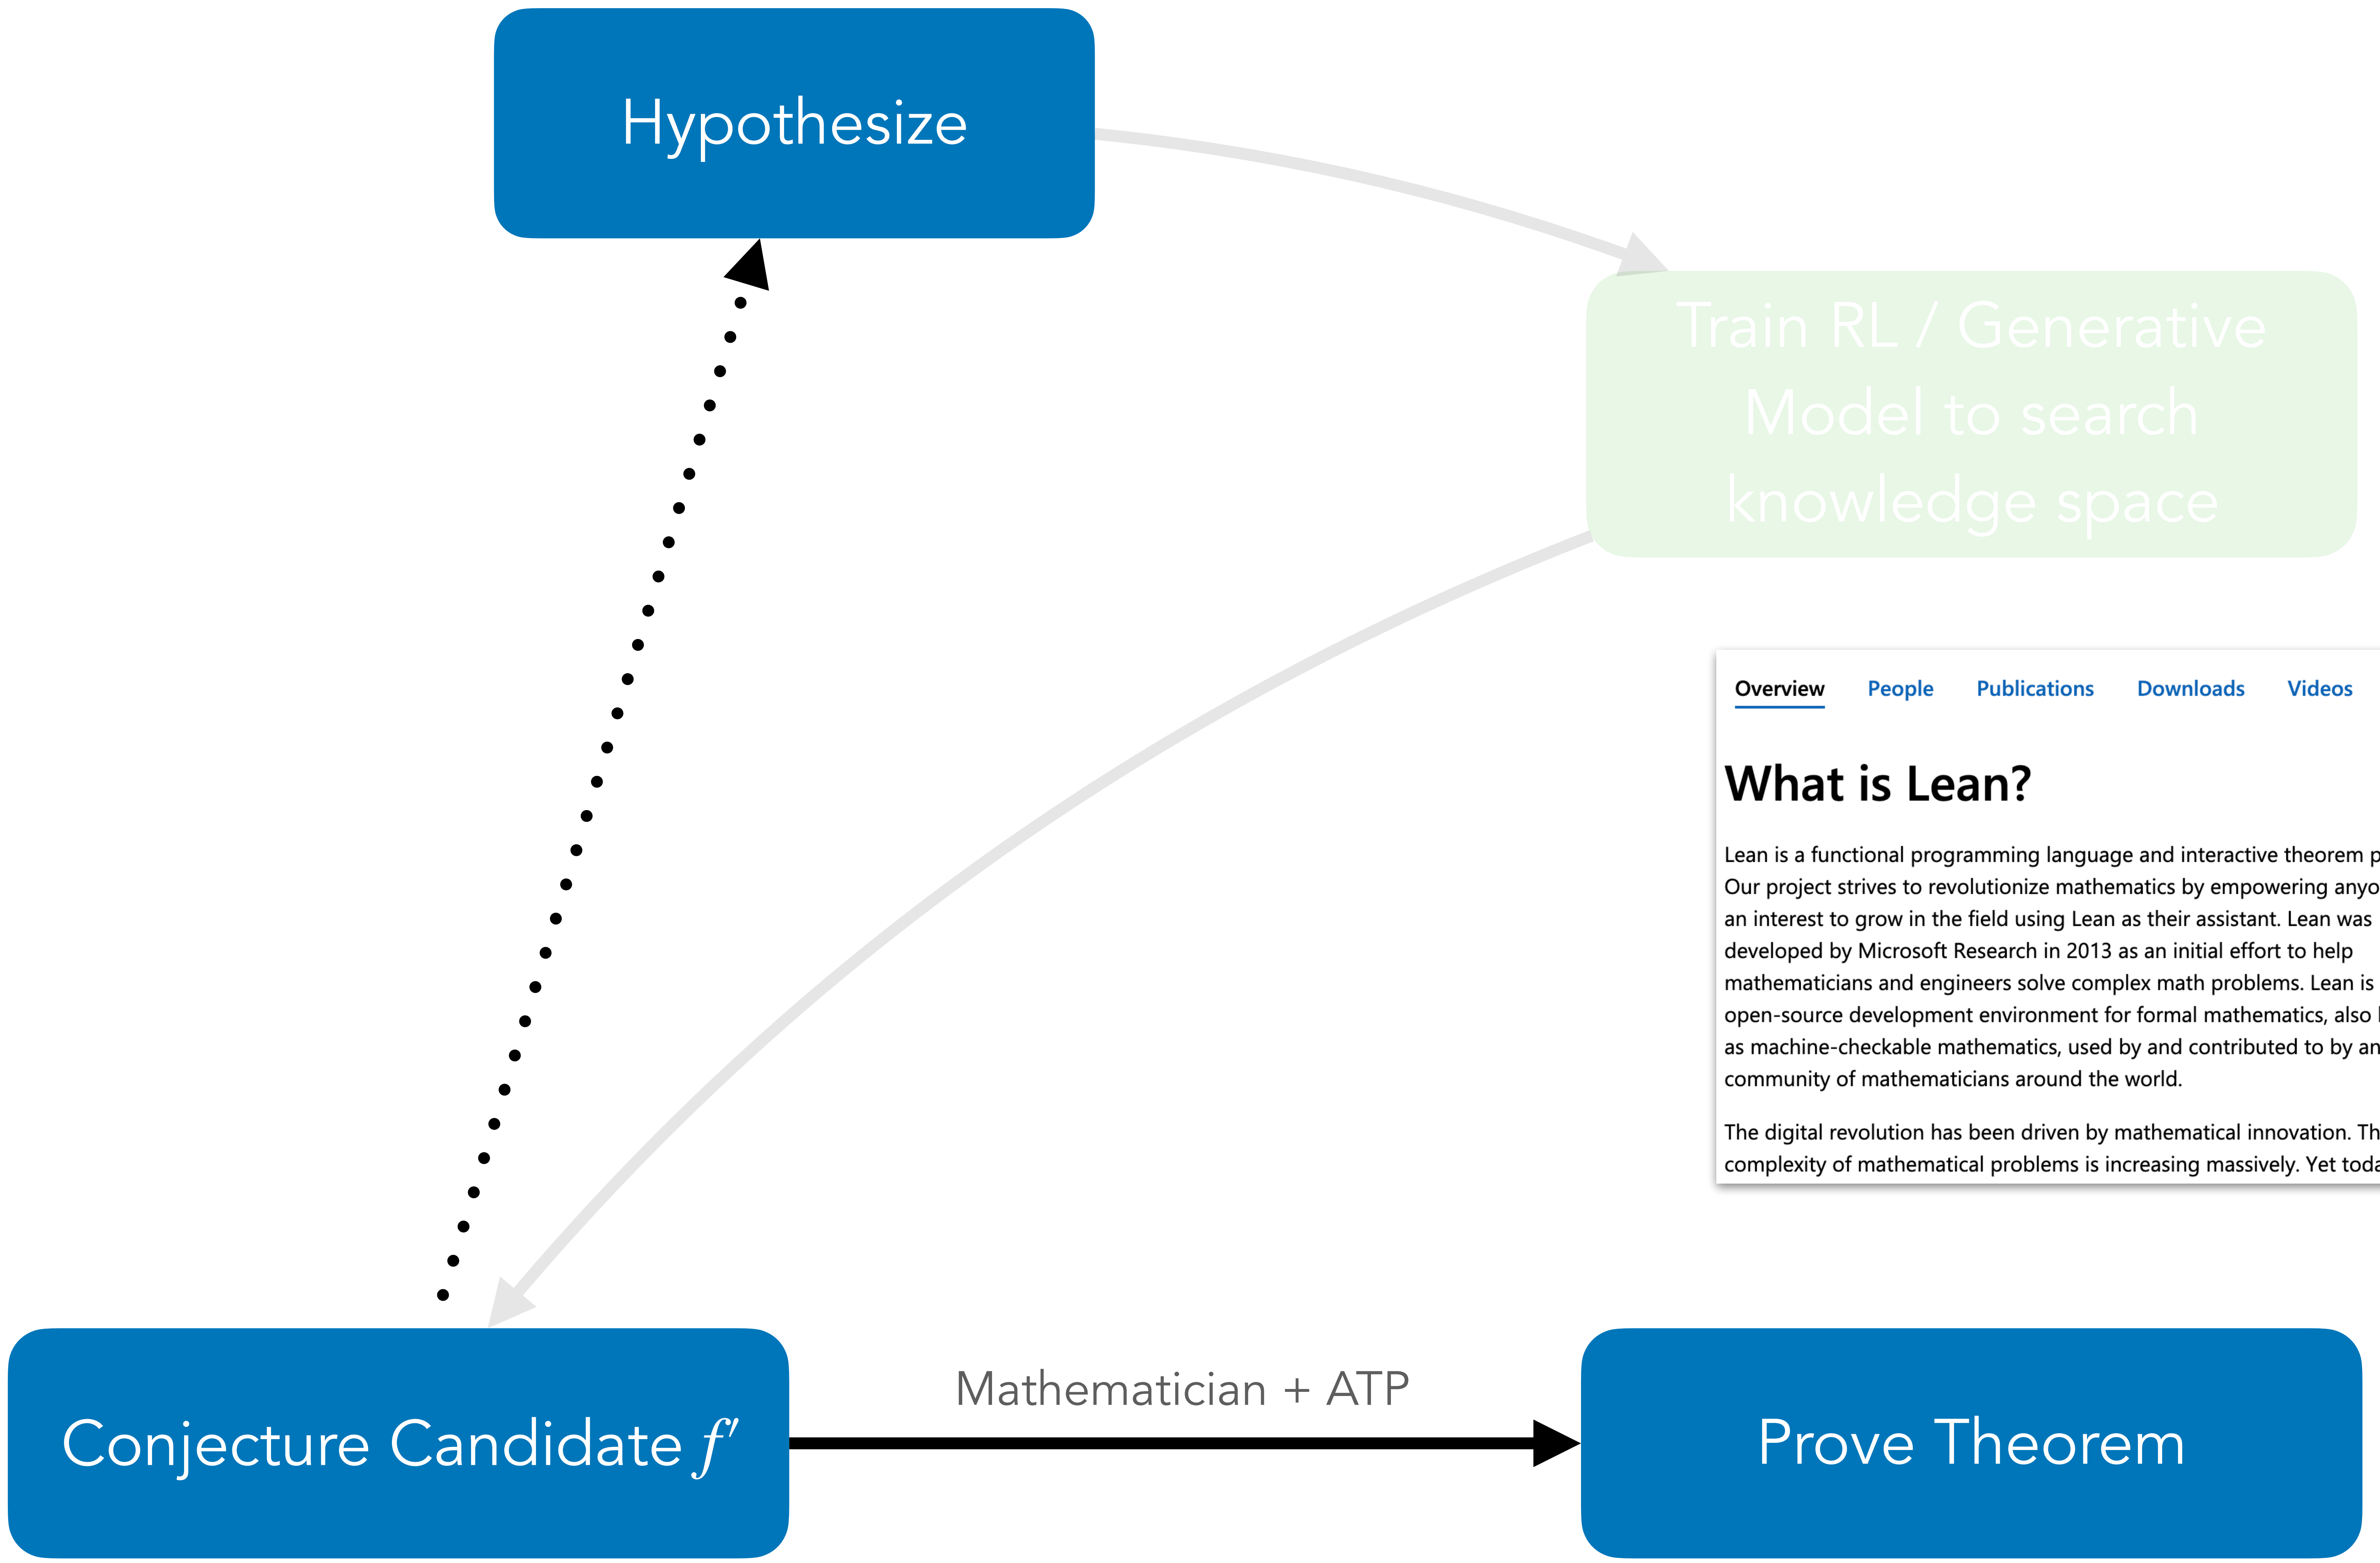

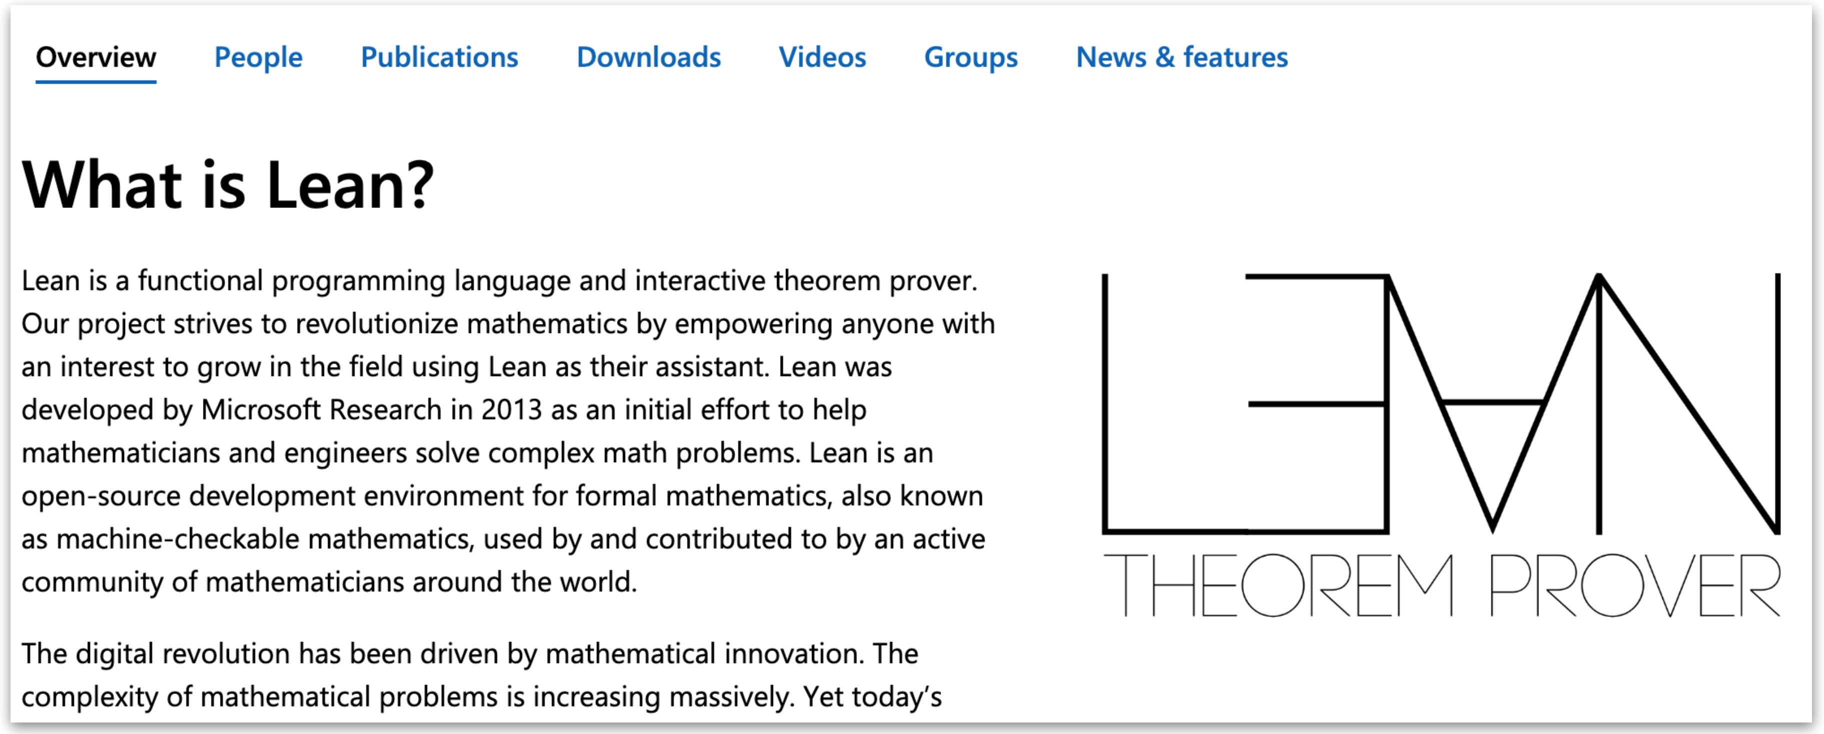
Overview People Publications Downloads Videos Groups News & features

What is Lean?

Lean is a functional programming language and interactive theorem prover. Our project strives to revolutionize mathematics by empowering anyone with an interest to grow in the field using Lean as their assistant. Lean was developed by Microsoft Research in 2013 as an initial effort to help mathematicians and engineers solve complex math problems. Lean is an open-source development environment for formal mathematics, also known as machine-checkable mathematics, used by and contributed to by an active community of mathematicians around the world.

The digital revolution has been driven by mathematical innovation. The complexity of mathematical problems is increasing massively. Yet today's

LEAN THEOREM PROVER

My edits

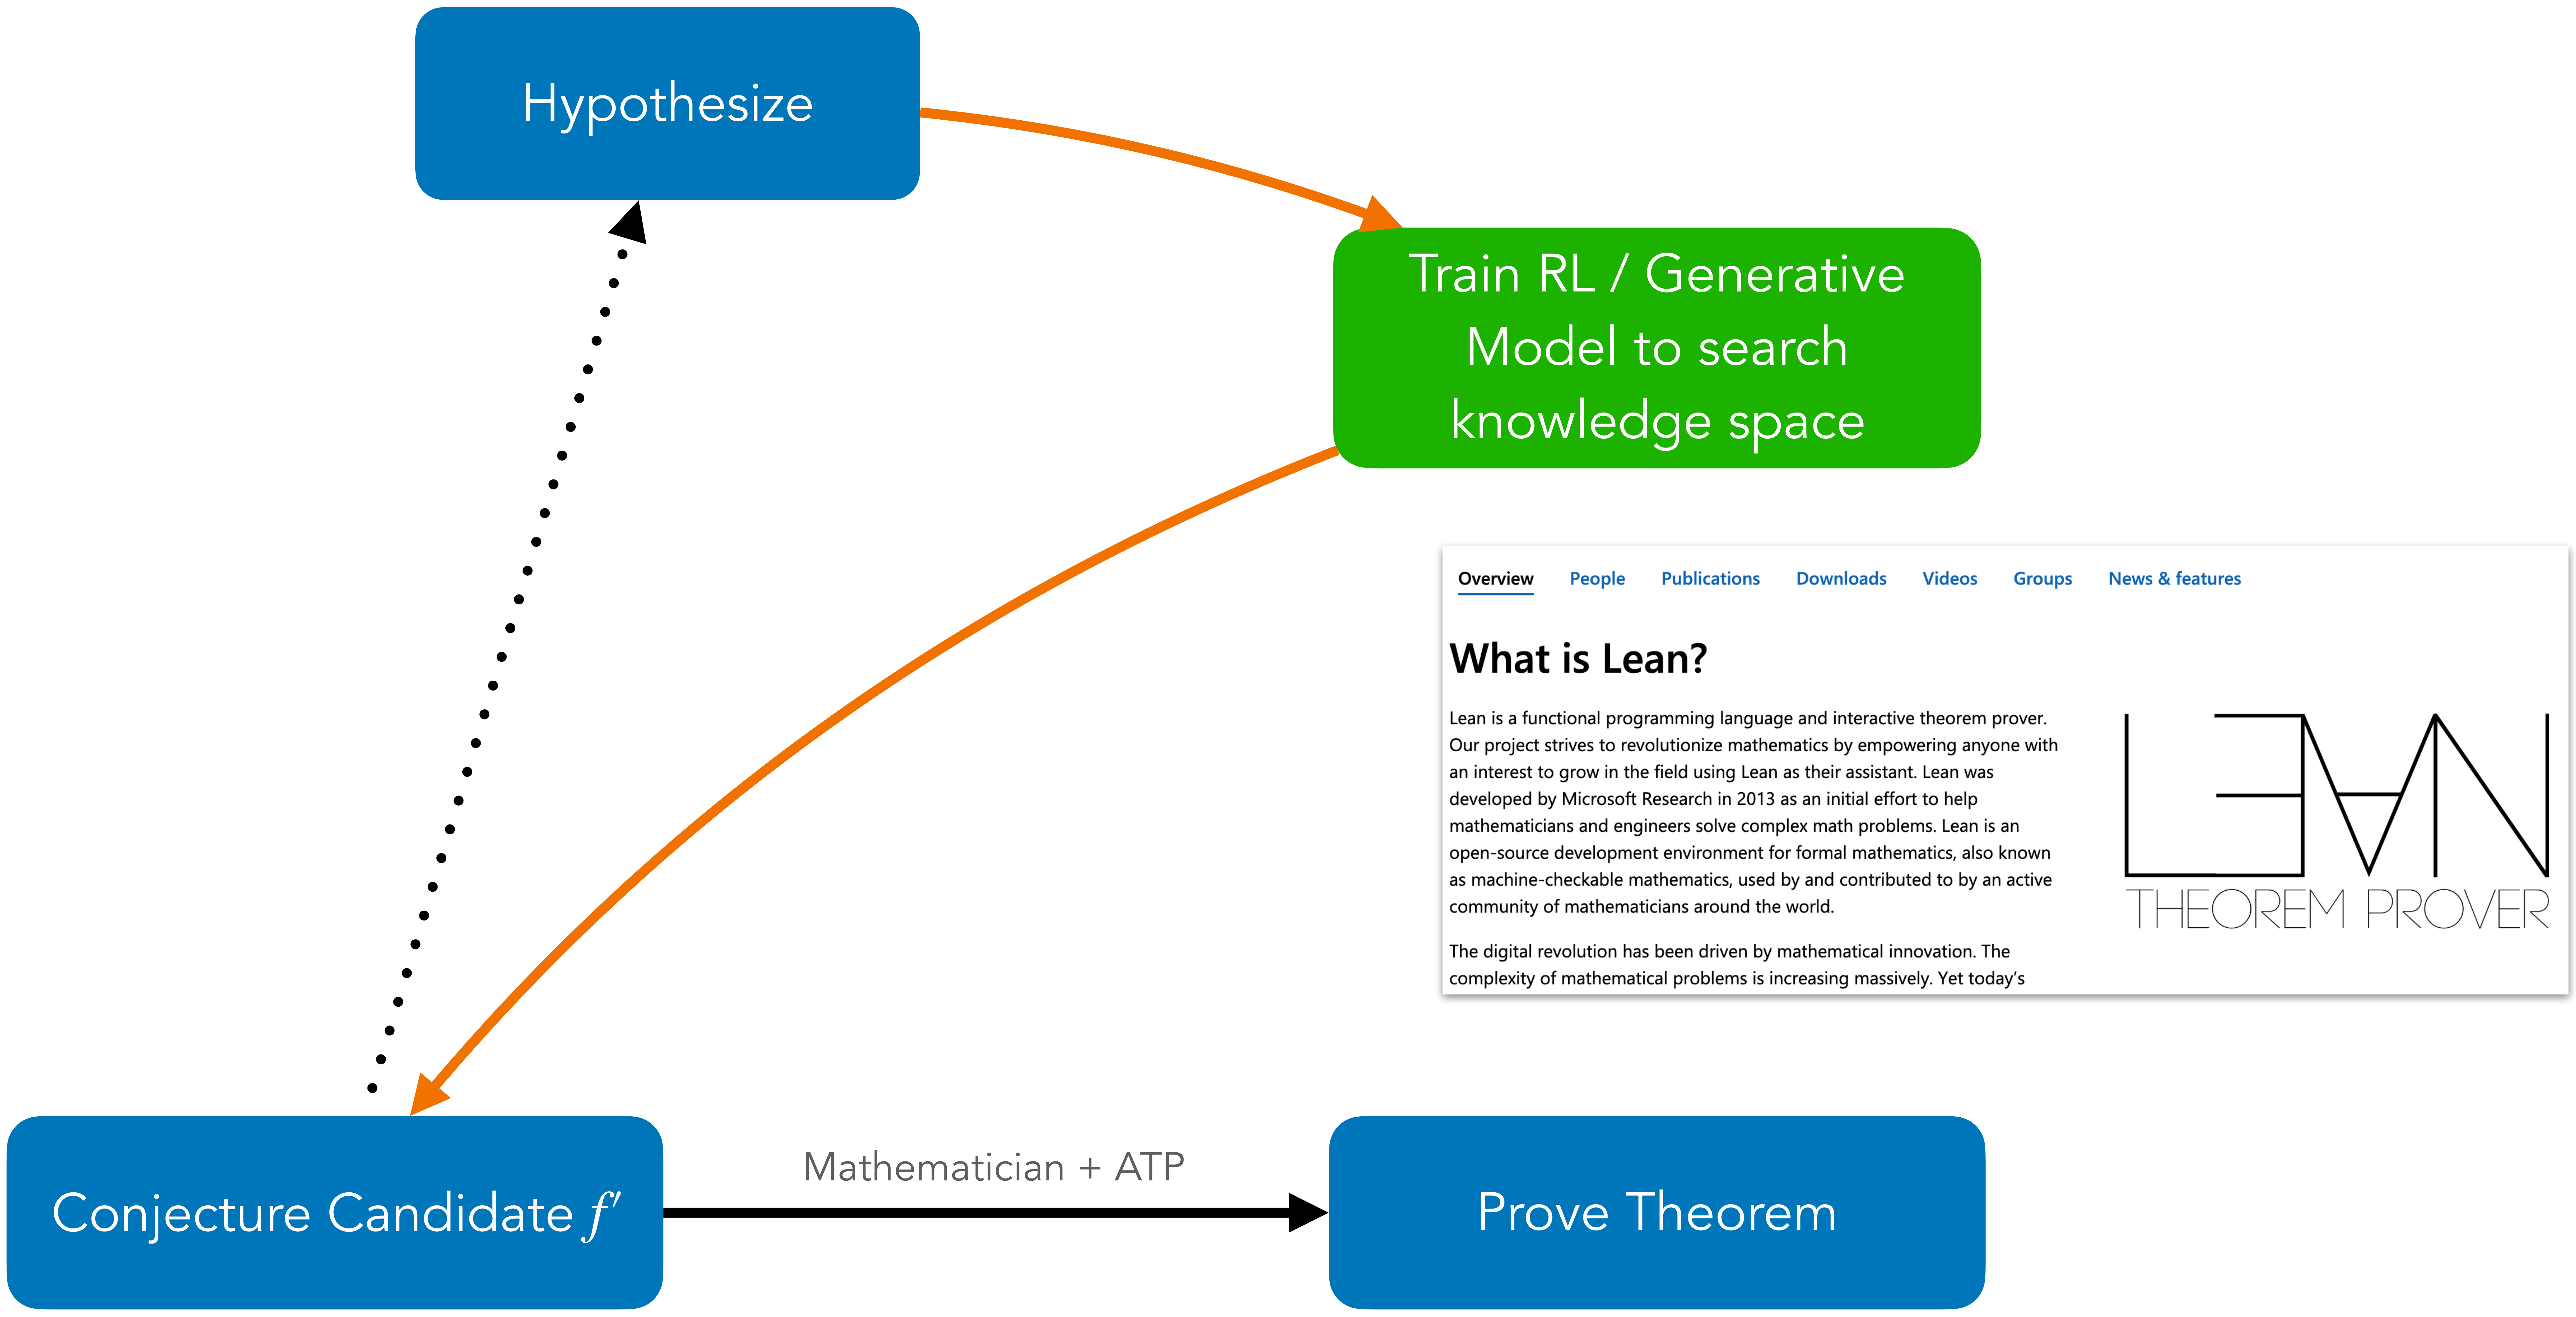
Hypothesize
Train RL / Generative Model to search knowledge space
Conjecture Candidate f'
Mathematician + ATP
Prove Theorem
Overview
People
Publications
Downloads
Videos
Groups
News & features
What is Lean?
Lean is a functional programming language and interactive theorem prover. Our project strives to revolutionize mathematics by empowering anyone with an interest to grow in the field using Lean as their assistant. Lean was developed by Microsoft Research in 2013 as an initial effort to help mathematicians and engineers solve complex math problems. Lean is an open-source development environment for formal mathematics, also known as machine-checkable mathematics, used by and contributed to by an active community of mathematicians around the world.
The digital revolution has been driven by mathematical innovation. The complexity of mathematical problems is increasing massively. Yet today's
LEAN
THEOREM PROVER

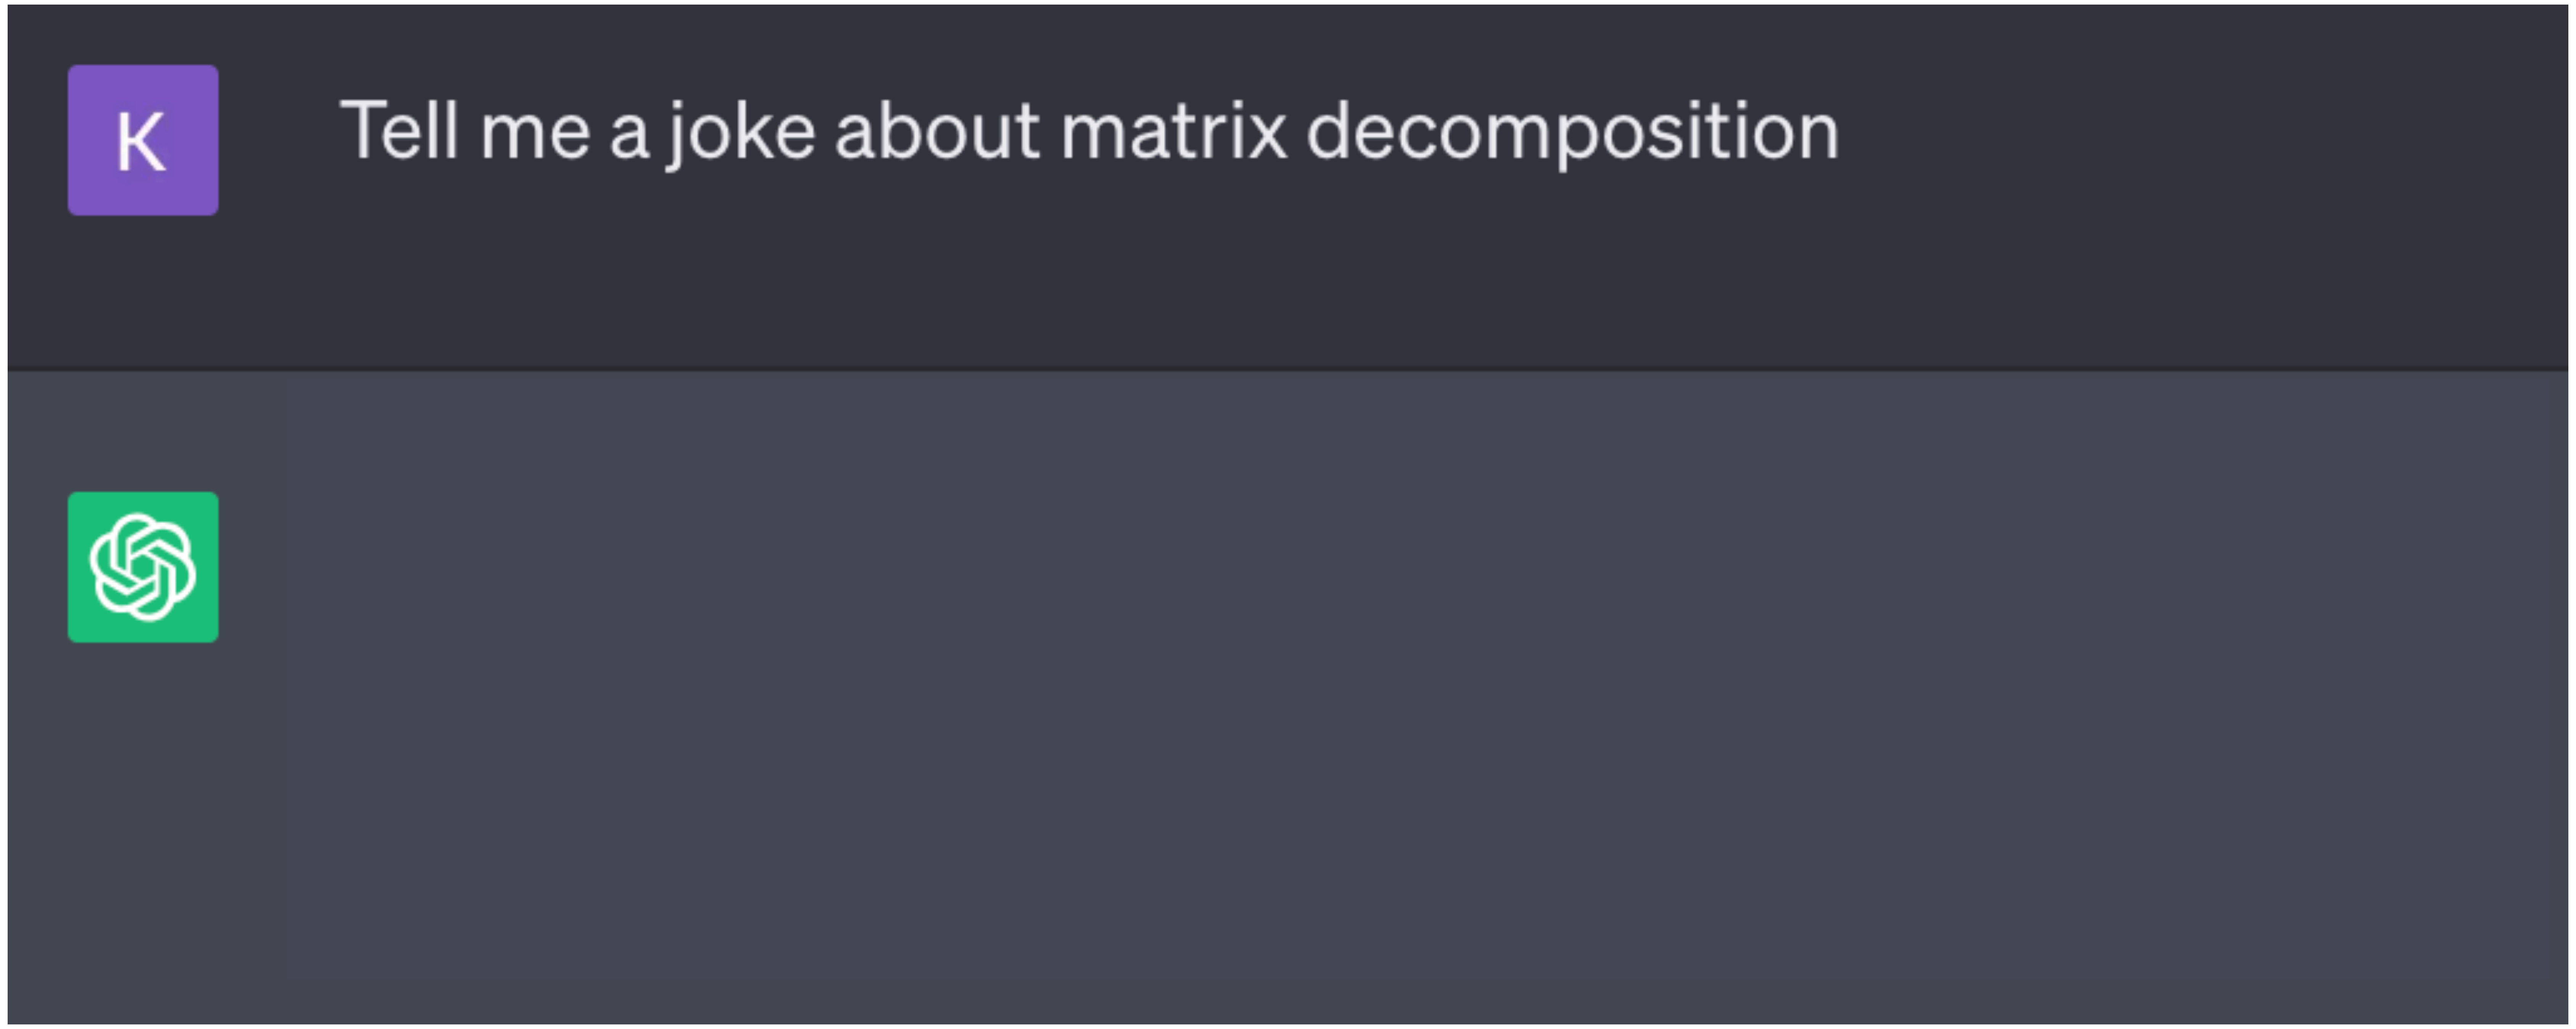
K
Tell me a joke about matrix decomposition

K

Tell me a joke about matrix decomposition

Why did the matrix refuse to go on vacation?

K

Tell me a joke about matrix decomposition

Why did the matrix refuse to go on vacation?

It didn't want to decompose under all that pressure! 😄🏝️

References

[1] Davies, A., Veličković, P., Buesing, L. *et al.* Advancing mathematics by guiding human intuition with AI. *Nature* **600**, 70–74 (2021).

[2] Cornelio, C., Dash, S., Austel, V. *et al.* Combining data and theory for derivable scientific discovery with AI-Descartes. *Nat Commun* **14**, 1777 (2023).

[3] Fawzi, A., Balog, M., Huang, A. *et al.* Discovering faster matrix multiplication algorithms with reinforcement learning. *Nature* **610**, 47–53 (2022).

Karan Srivastava
ksrivastava4@wisc.edu