General Principles of Human and Machine Learning

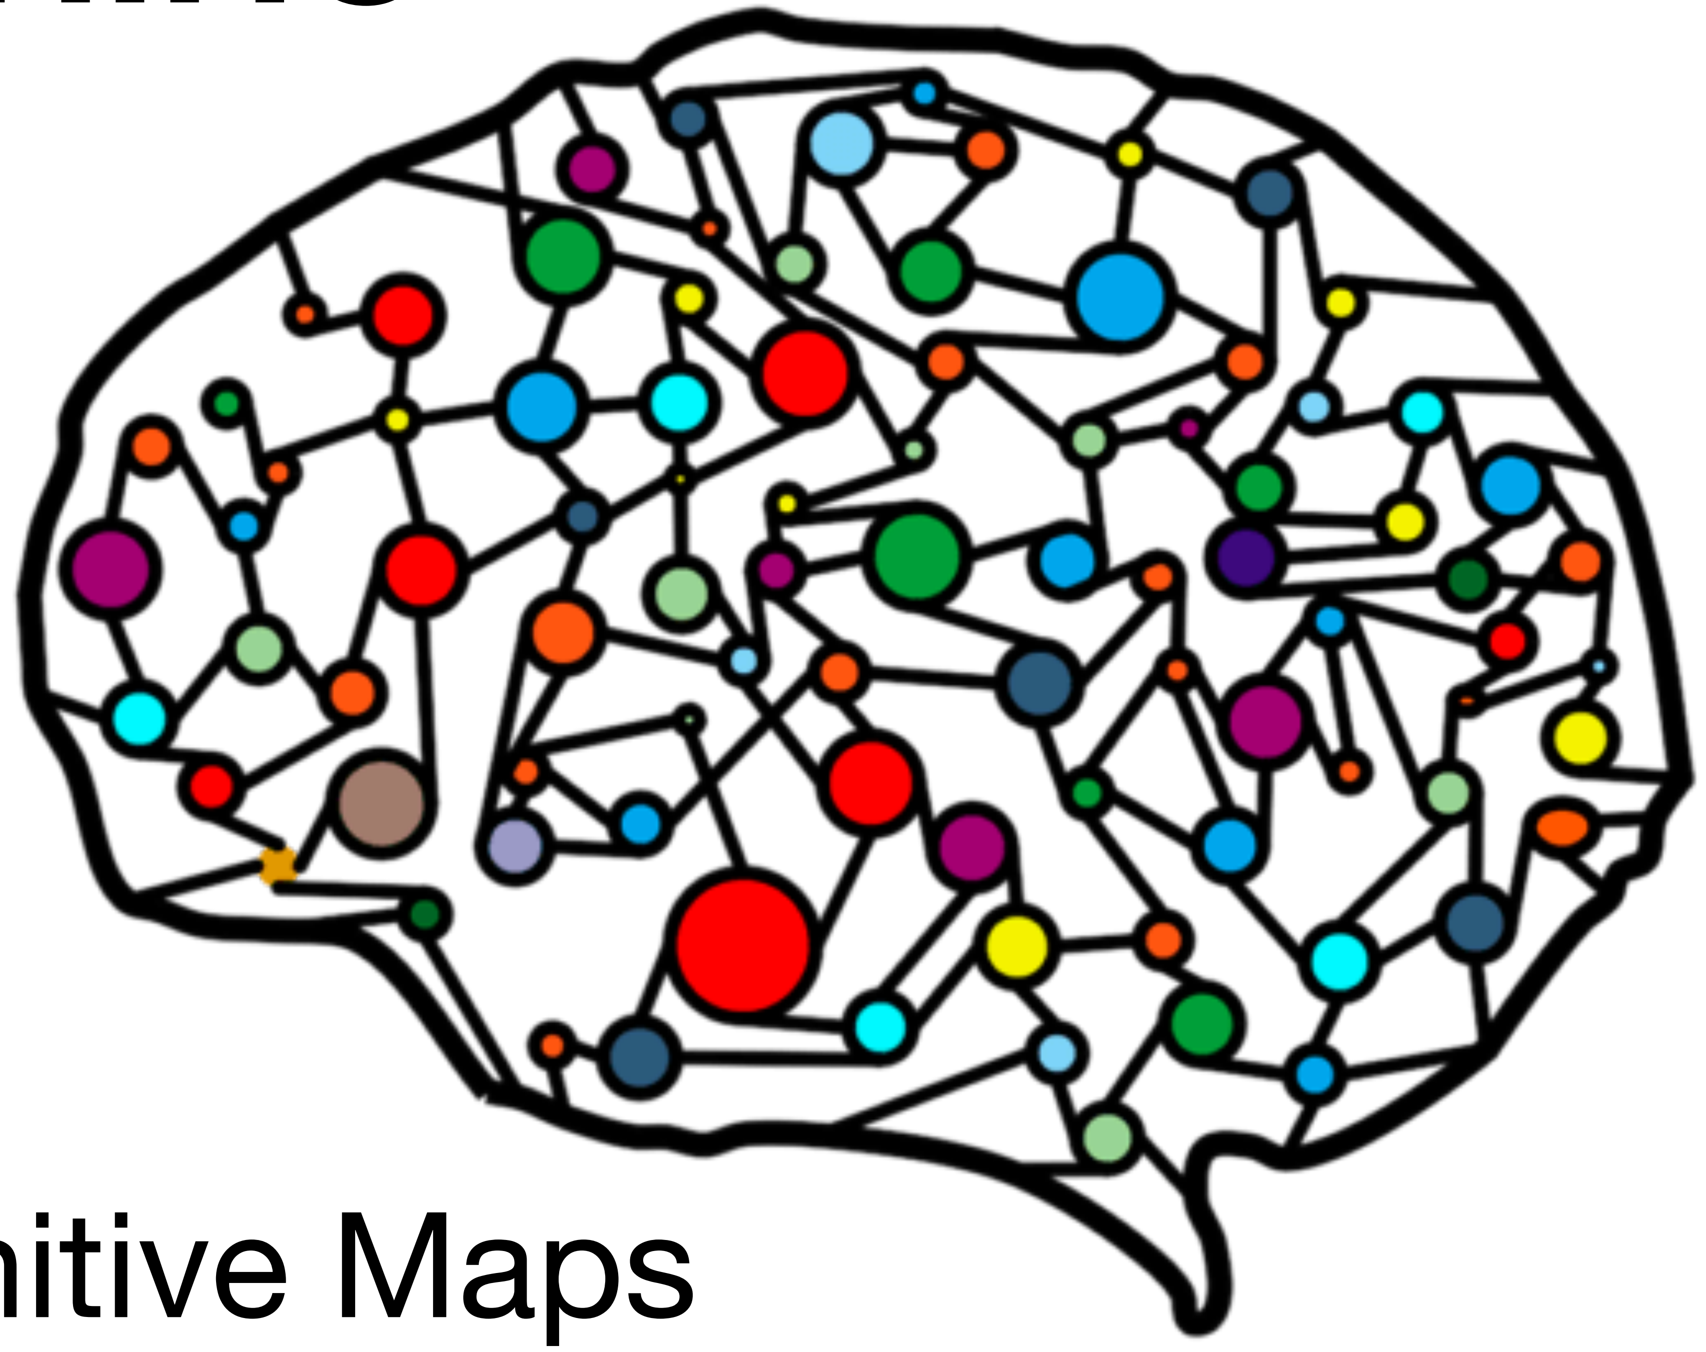

Lecture 3: Symbolic AI and Cognitive Maps

Dr. Charley Wu

https://hmc-lab.com/GPHML.html

Clarification from last week's tutorial

- Rescorla Wagner updates: Weights are only updated when the stimuli is present

For i where $CS_i = 1$:

$$w_i \leftarrow w_i + \eta(r_t - \hat{r}_t)$$

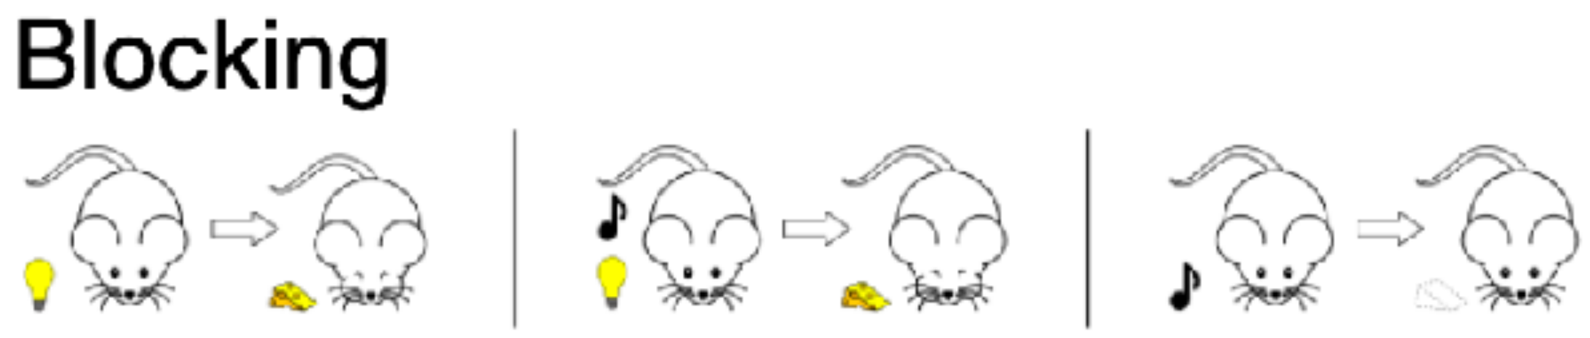

Clarification from last week's tutorial

- Rescorla Wagner updates: Weights are only updated when the stimuli is present

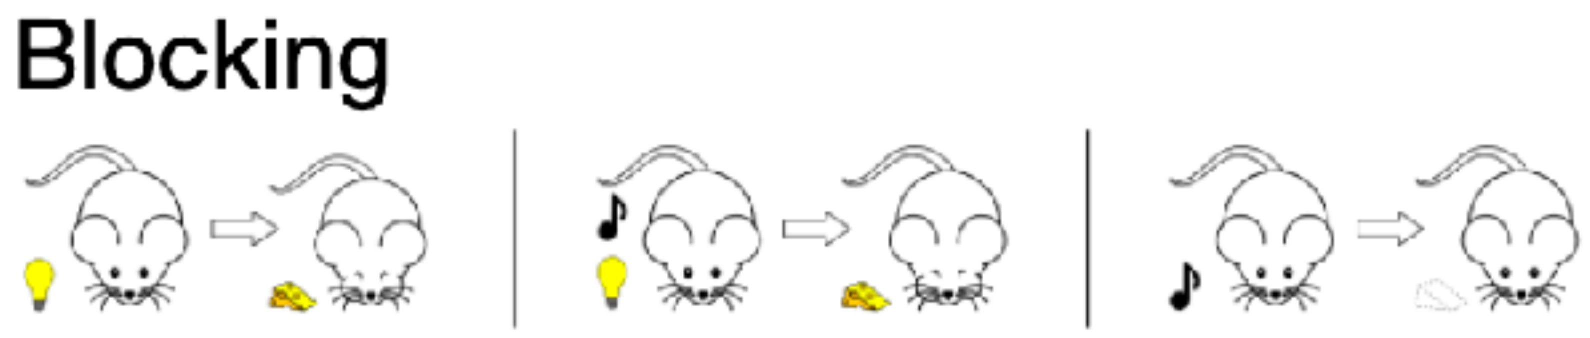

For i where $CS_i = 1$:

$$w_i \leftarrow w_i + \eta(r_t - \hat{r}_t)$$

Learning rate

Observed outcome

Predicted outcome

Lecture Plan

Symbolic AI

- What happened during the AI winter?
- Intelligence as manipulating symbols through rules and logical operations
- Learning as search

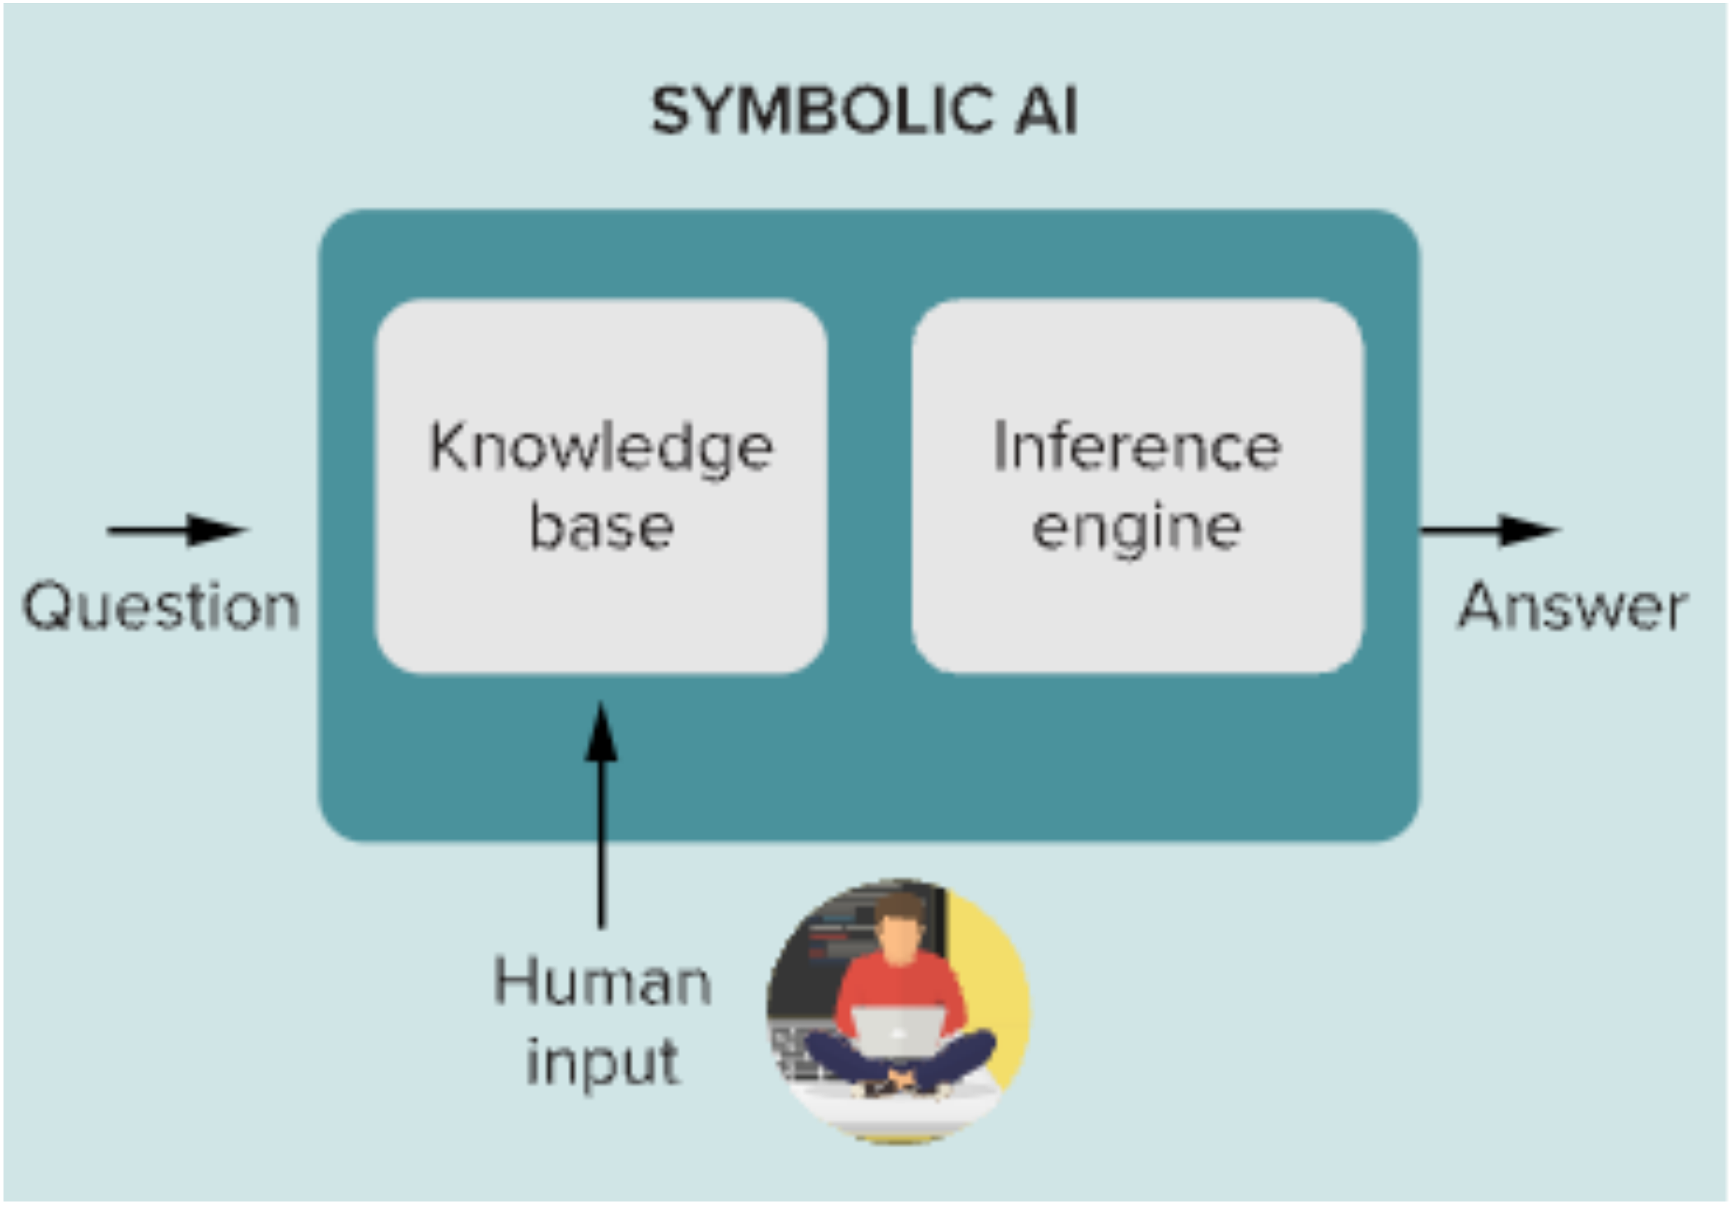

Cognitive Maps

- From Stimulus-Response learning to Stimulus-Stimulus learning
- Constructing a mental representation of the environment
- Neurological evidence for cognitive maps in the brain

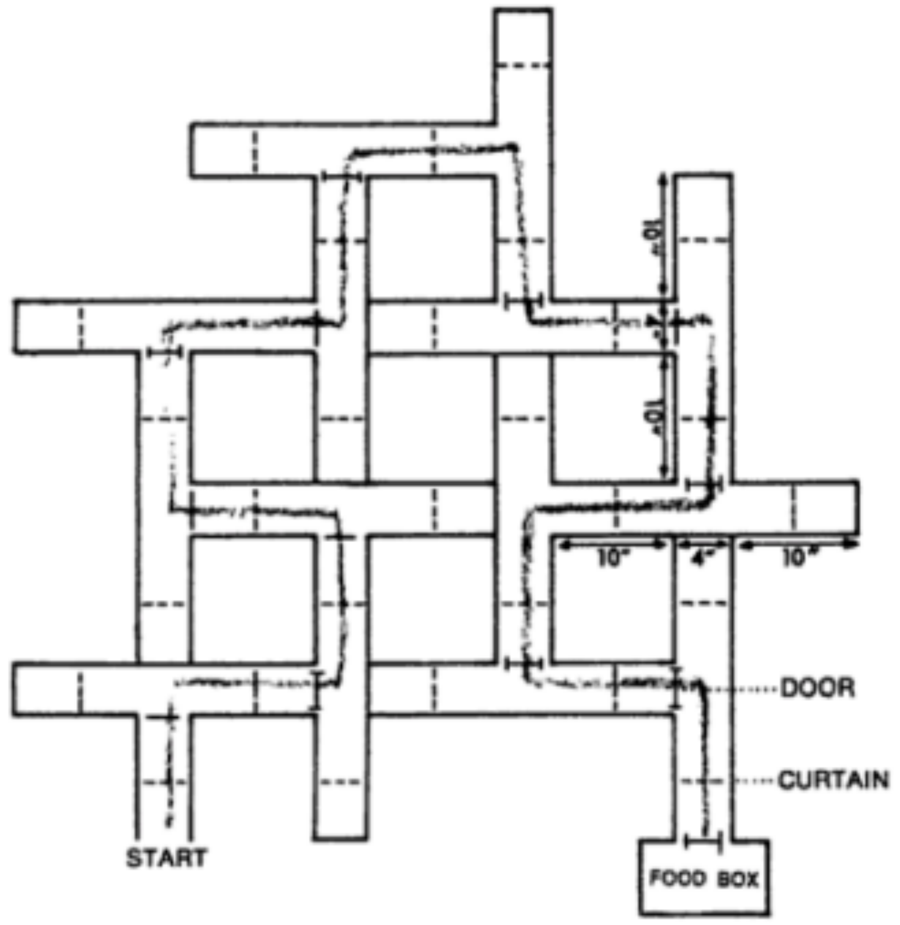

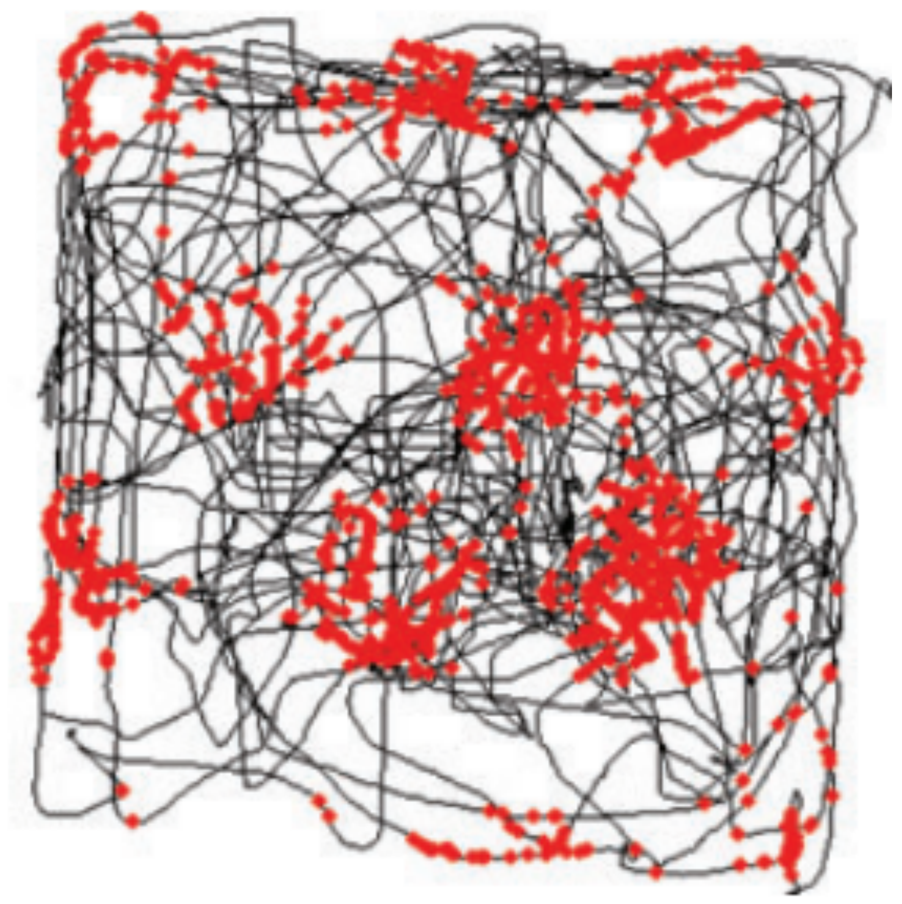

Timeline of AI

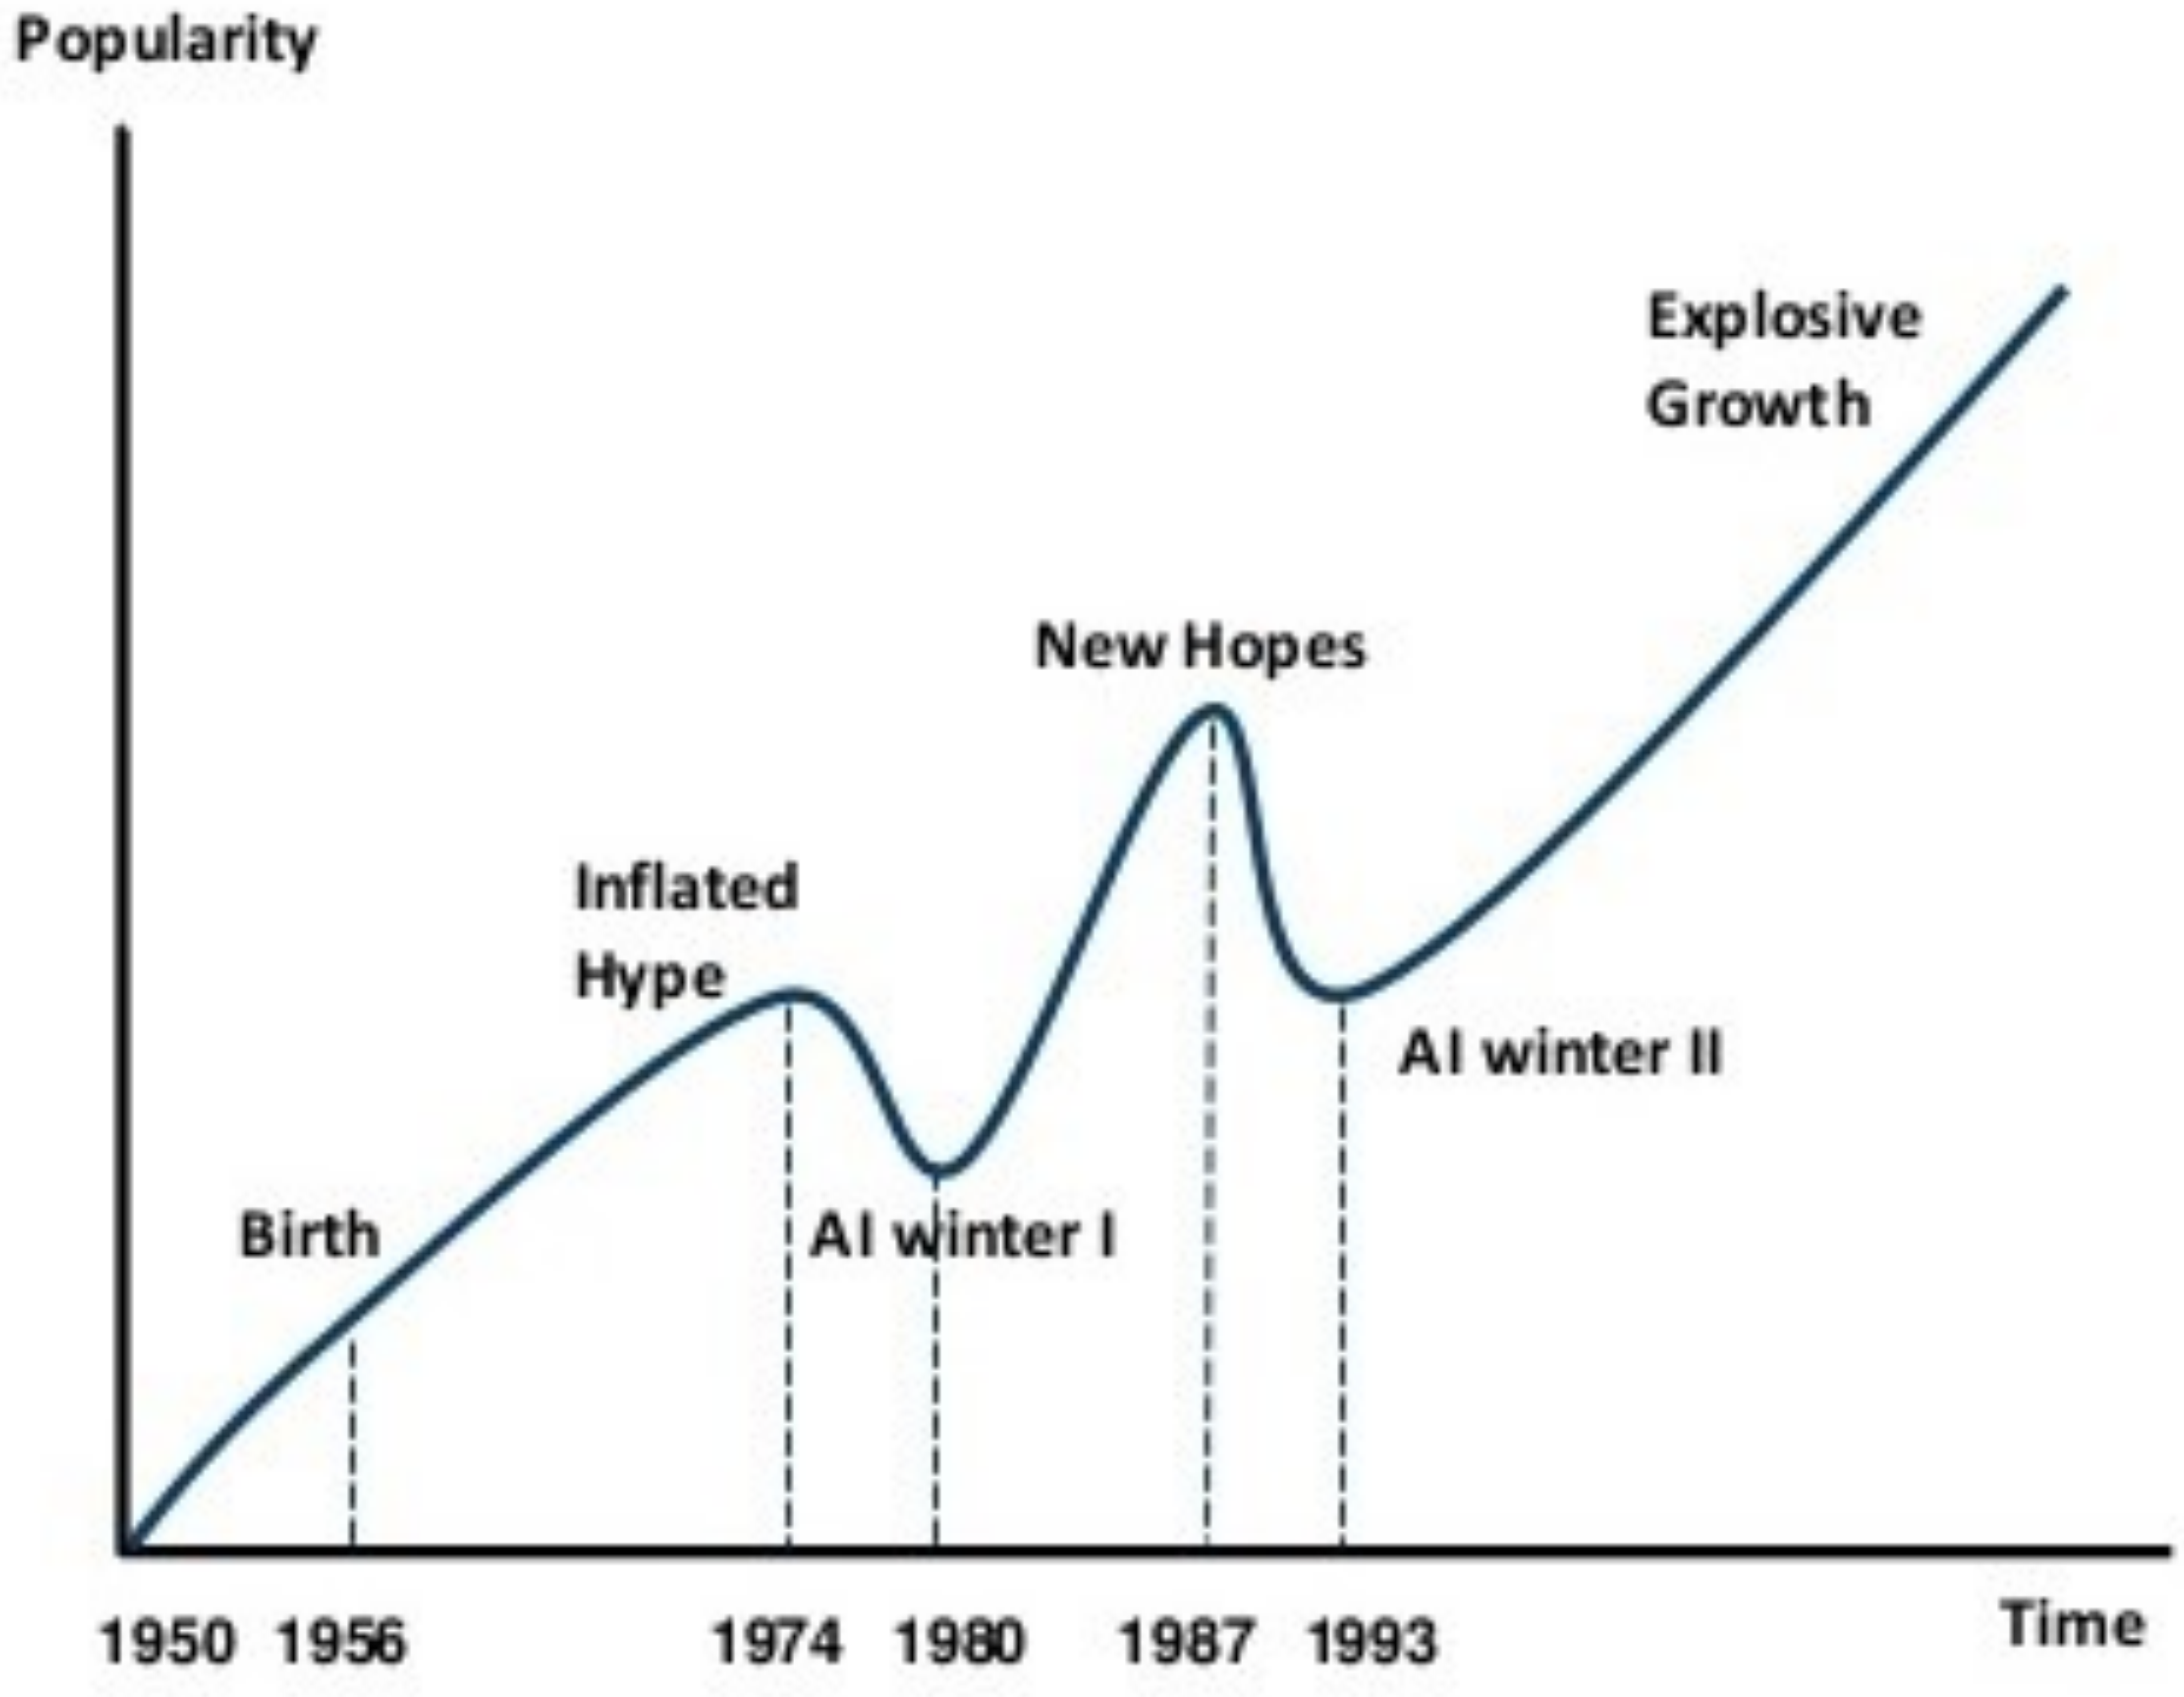

Timeline of AI

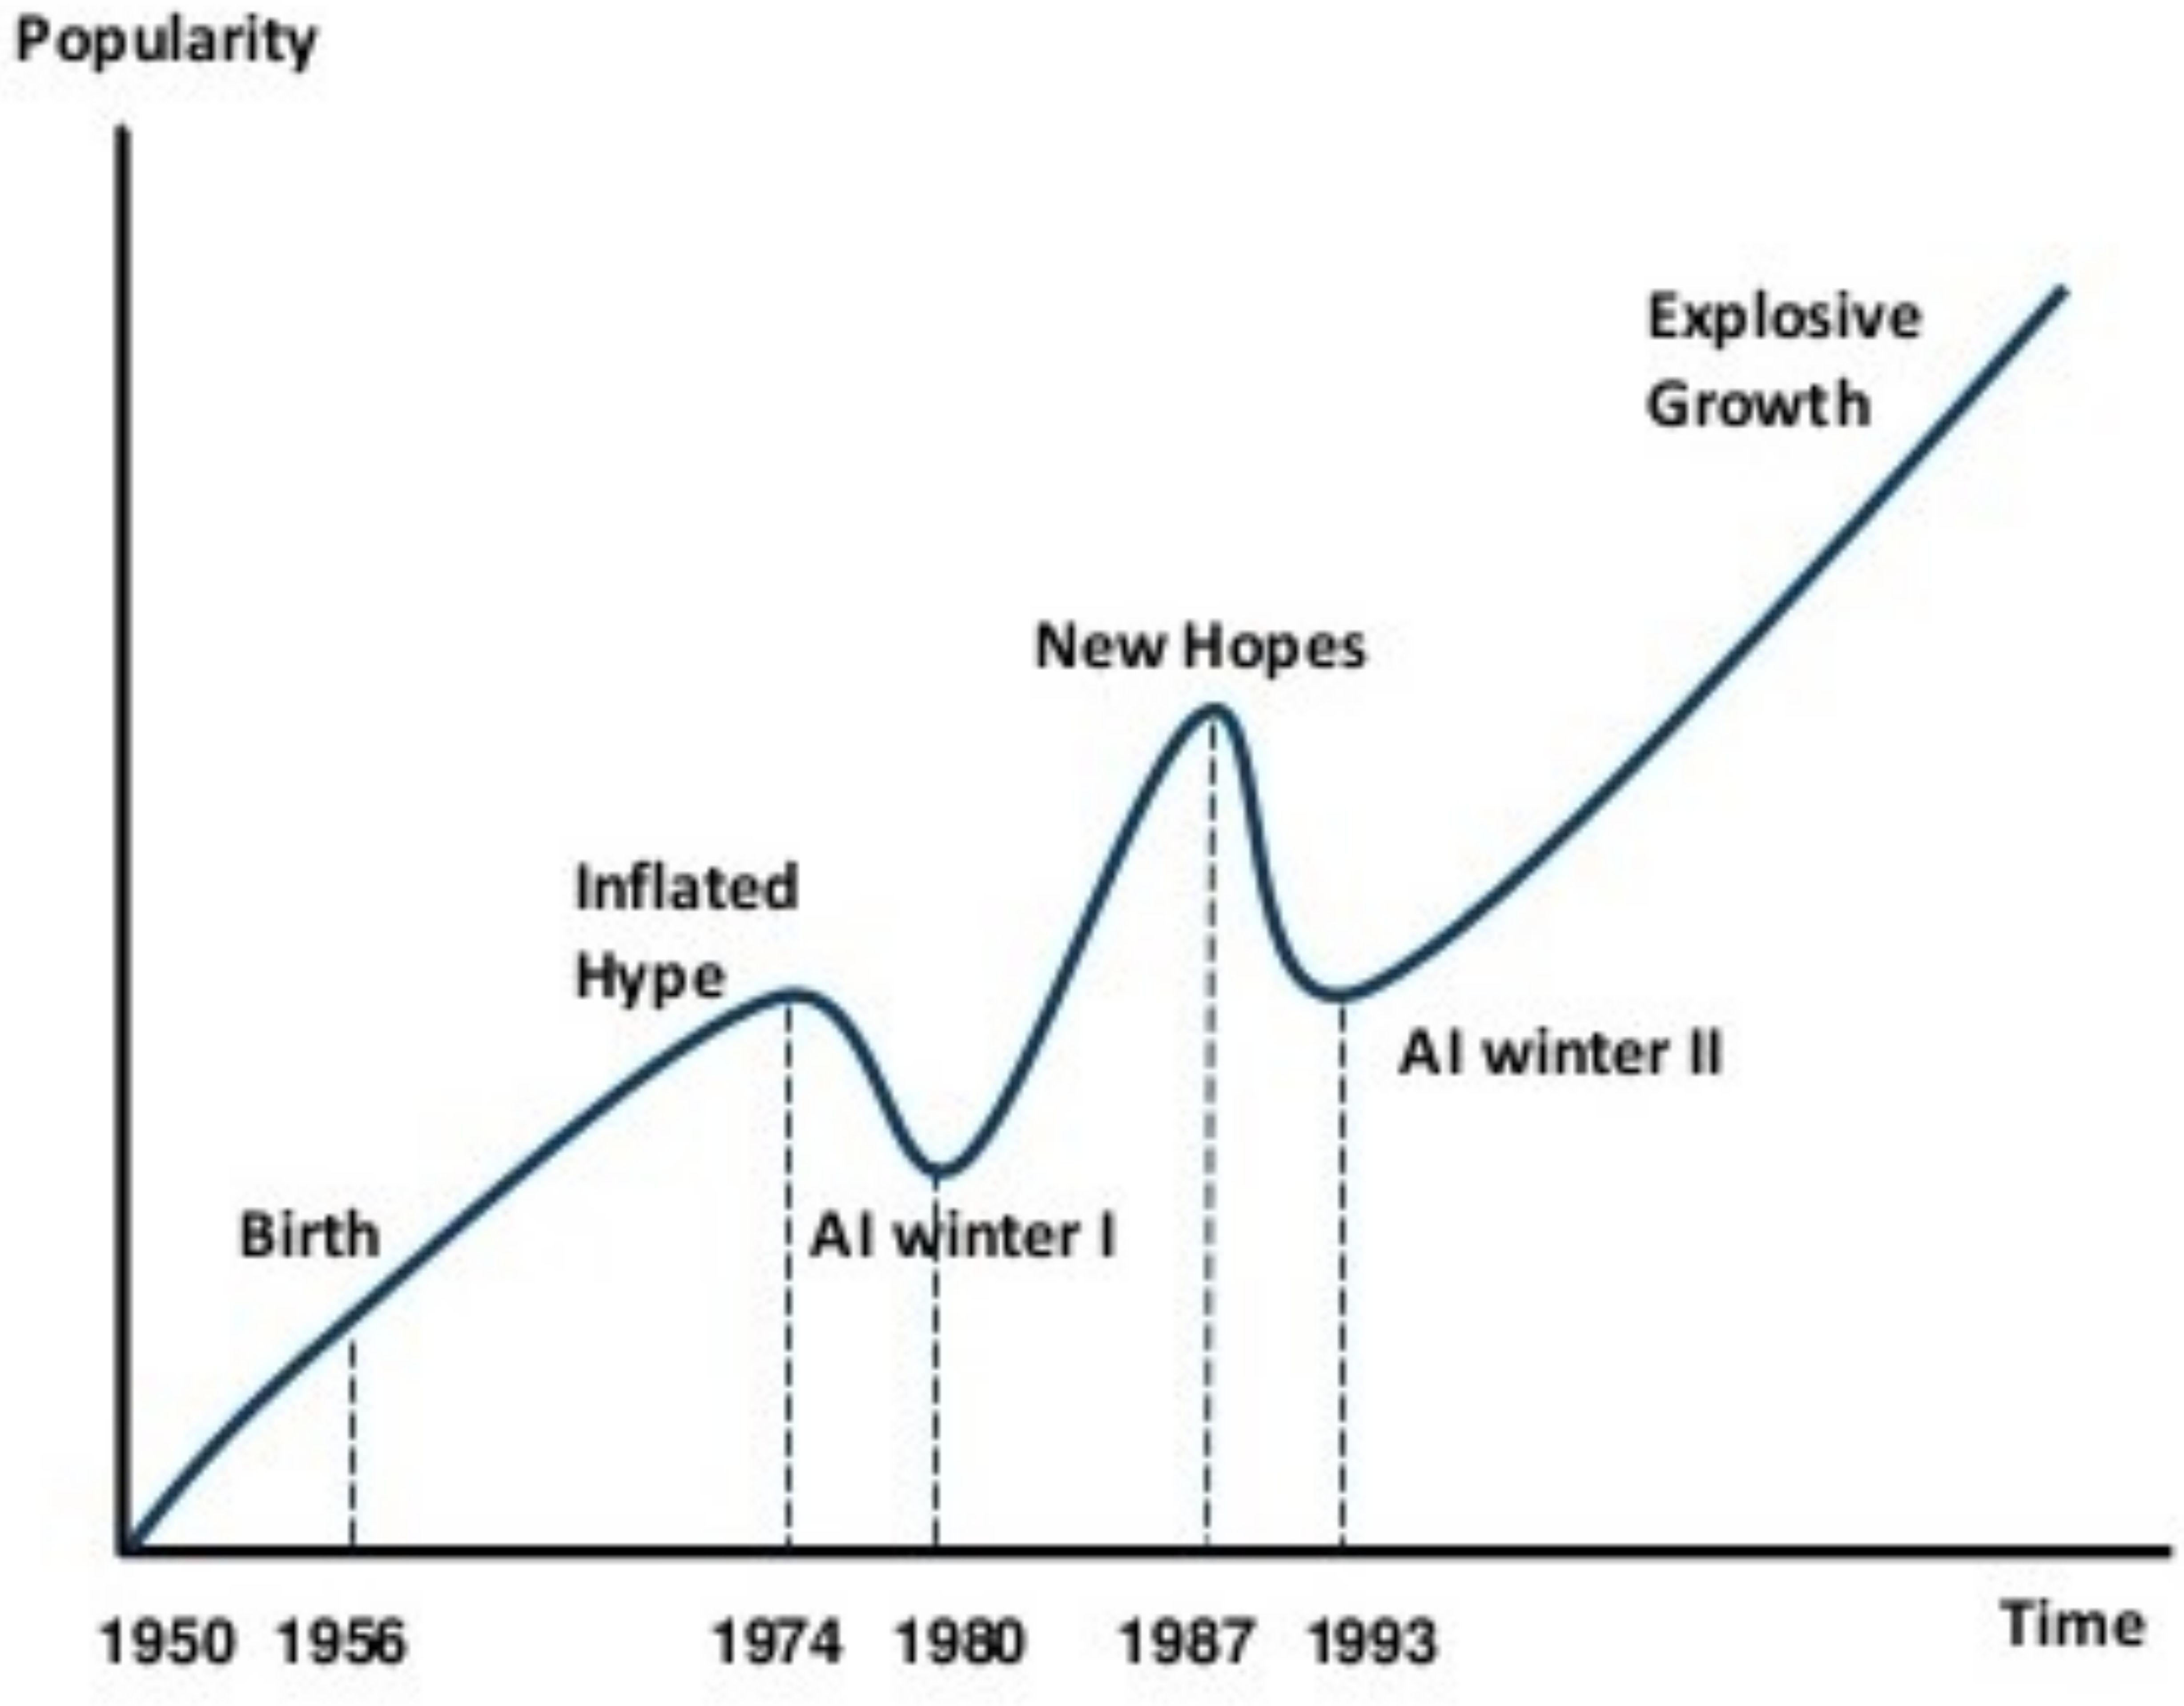

- AI has a long history of being “the next big thing”

Timeline of AI

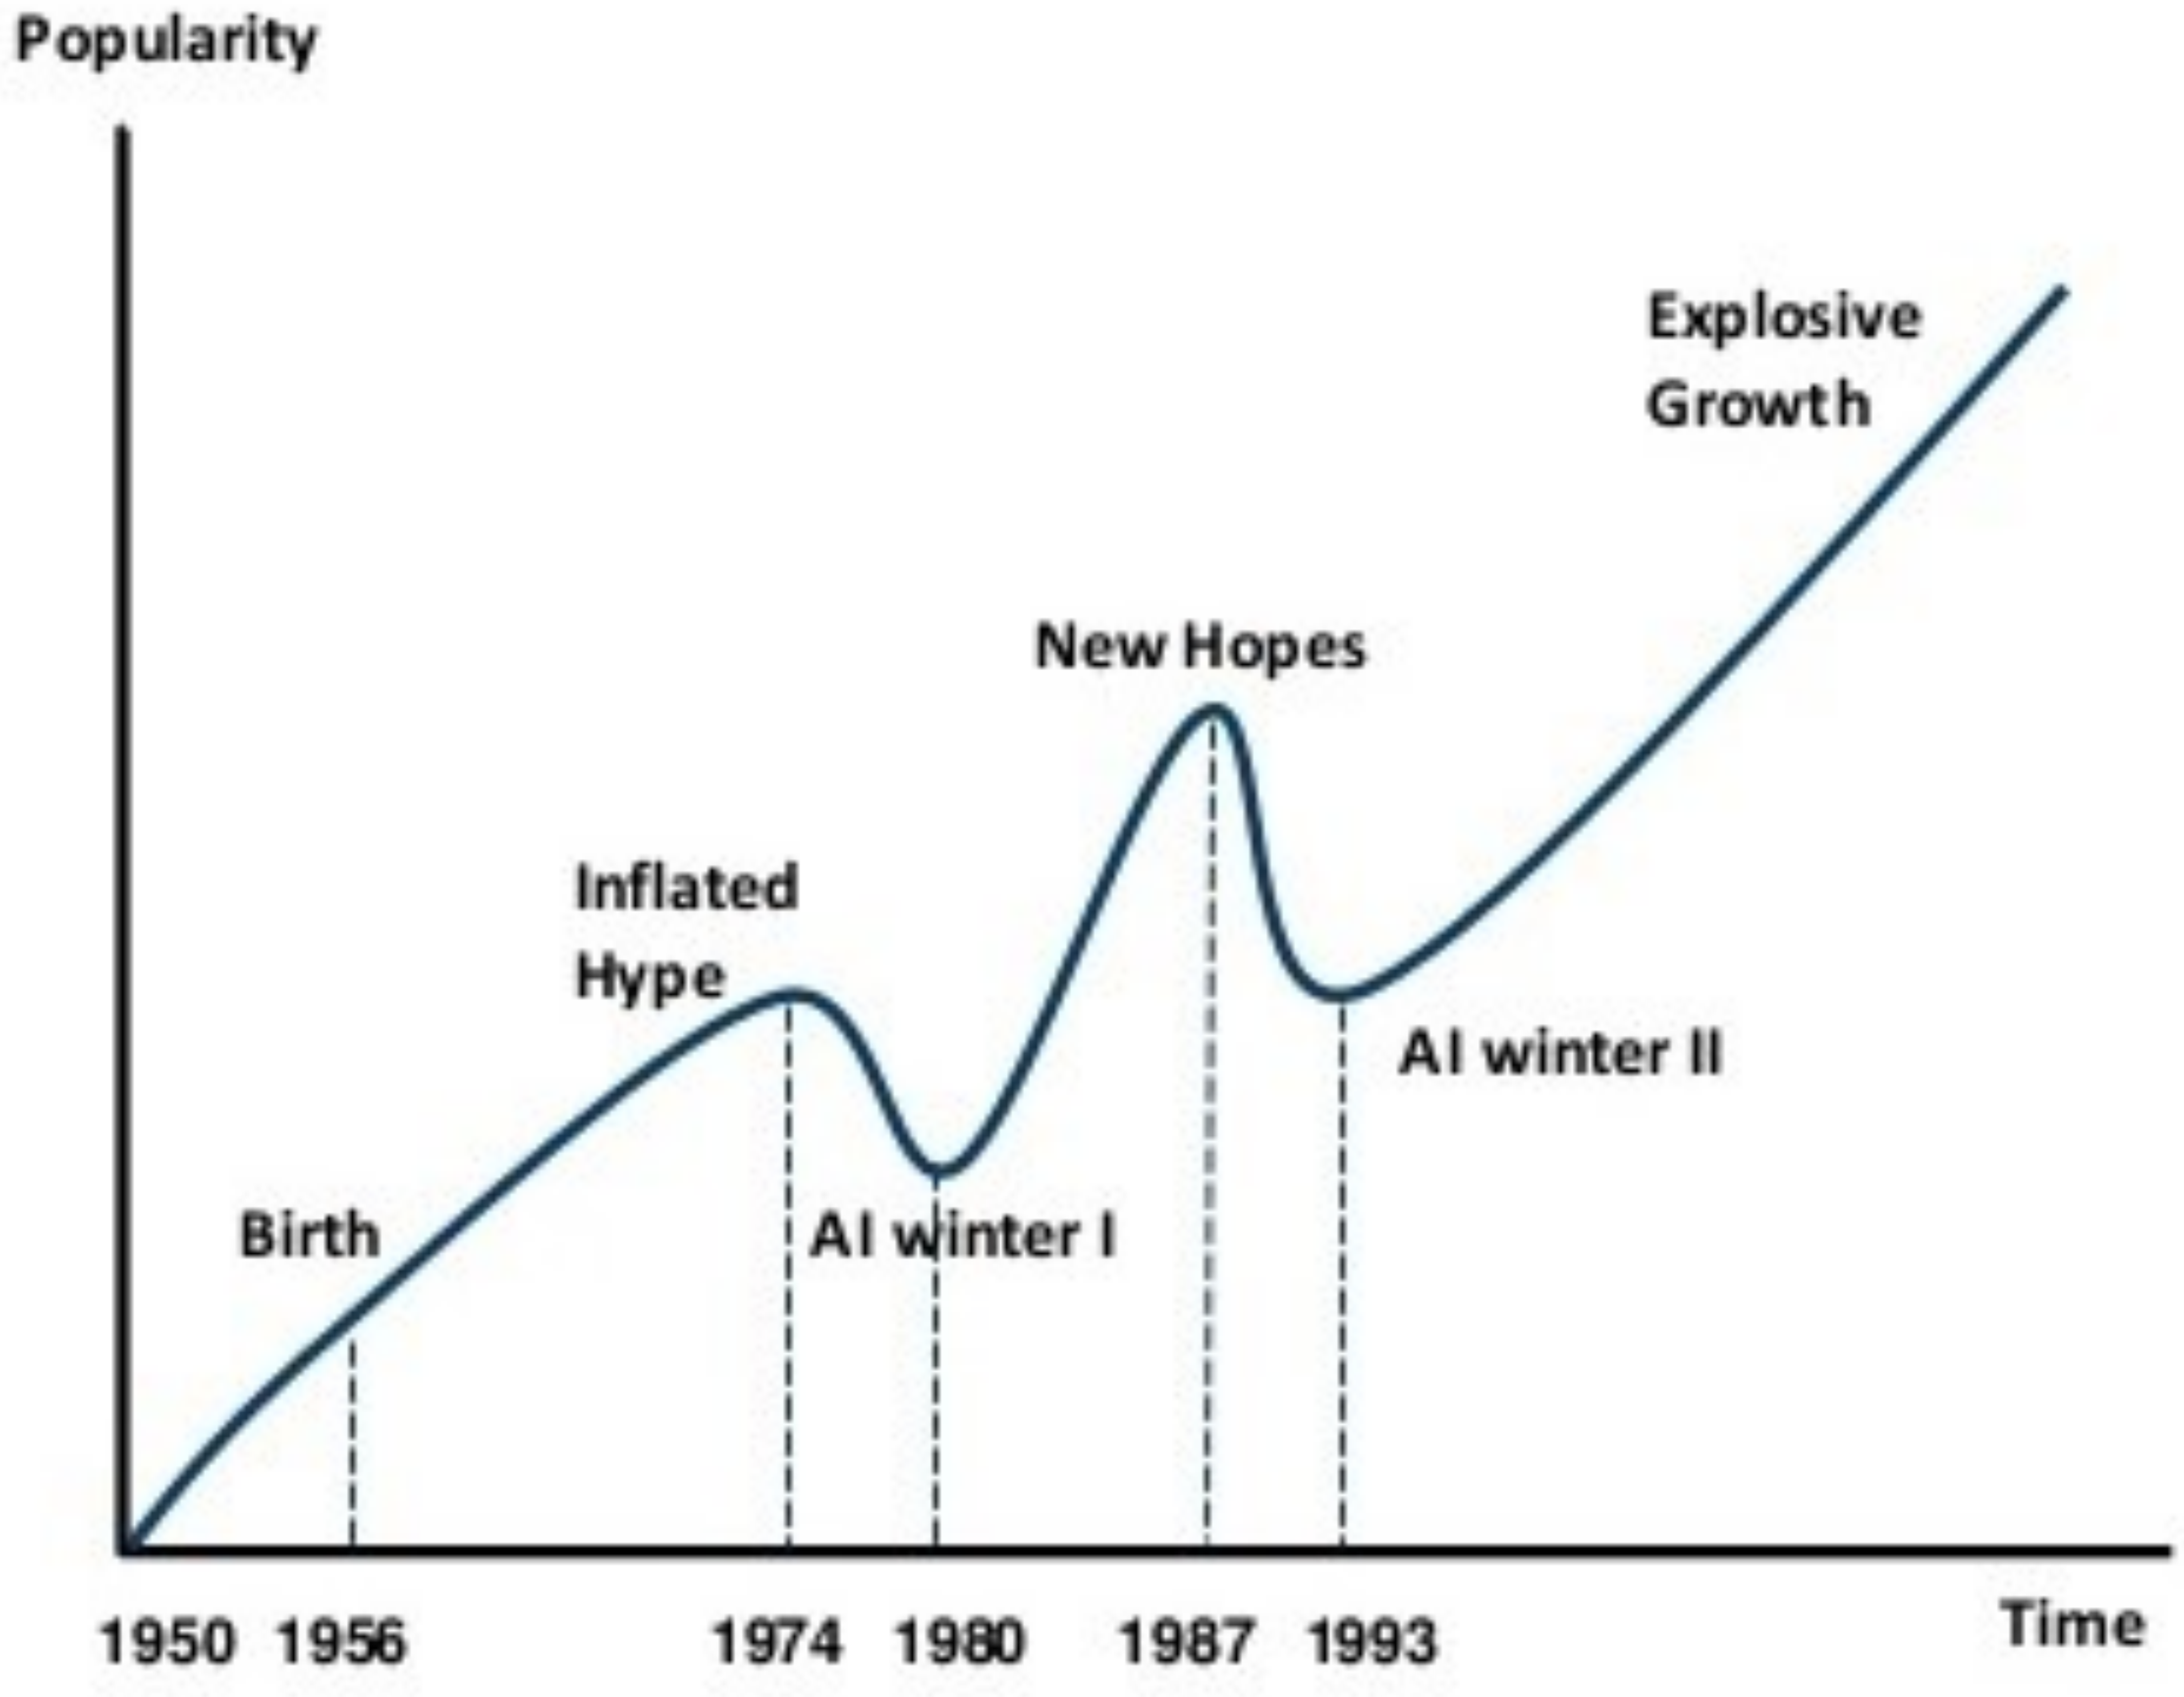

- AI has a long history of being "the next big thing"
- Multiple periods of "boom" and "bust"… including not one but two AI winters

Timeline of AI

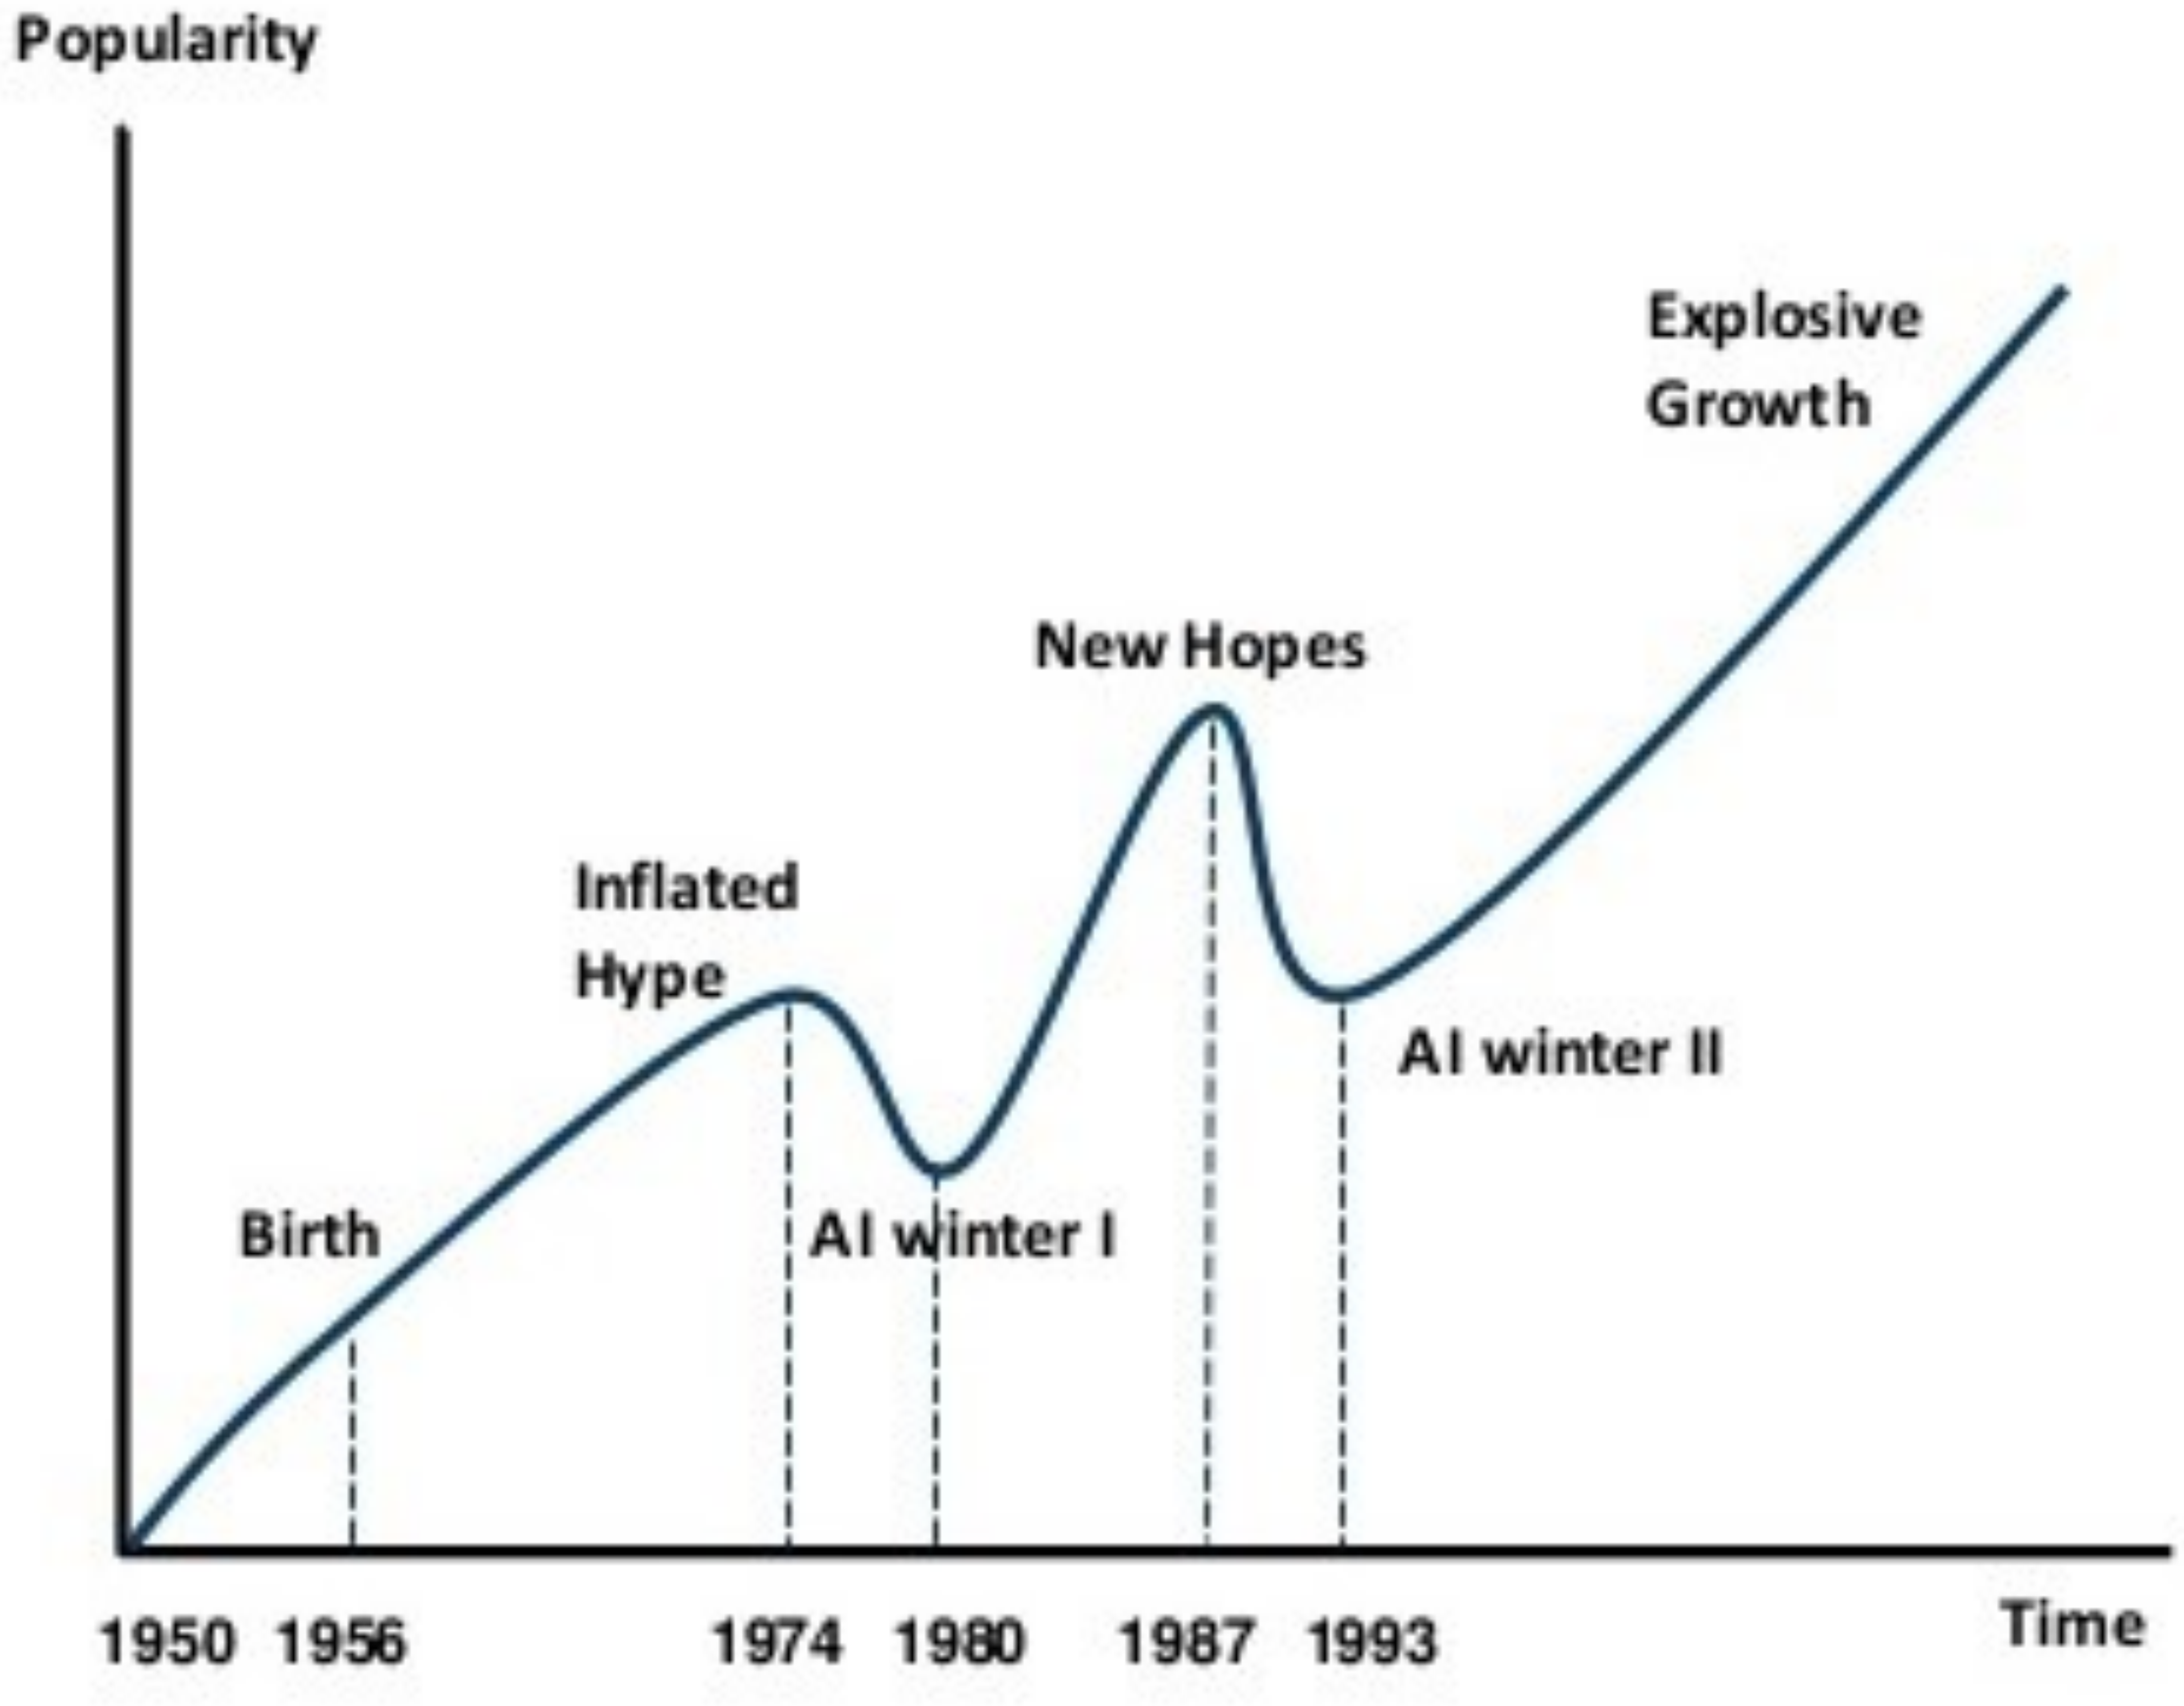

- AI has a long history of being “the next big thing”
- Multiple periods of “boom” and “bust”… including not one but two AI winters
- … so far

Timeline of AI

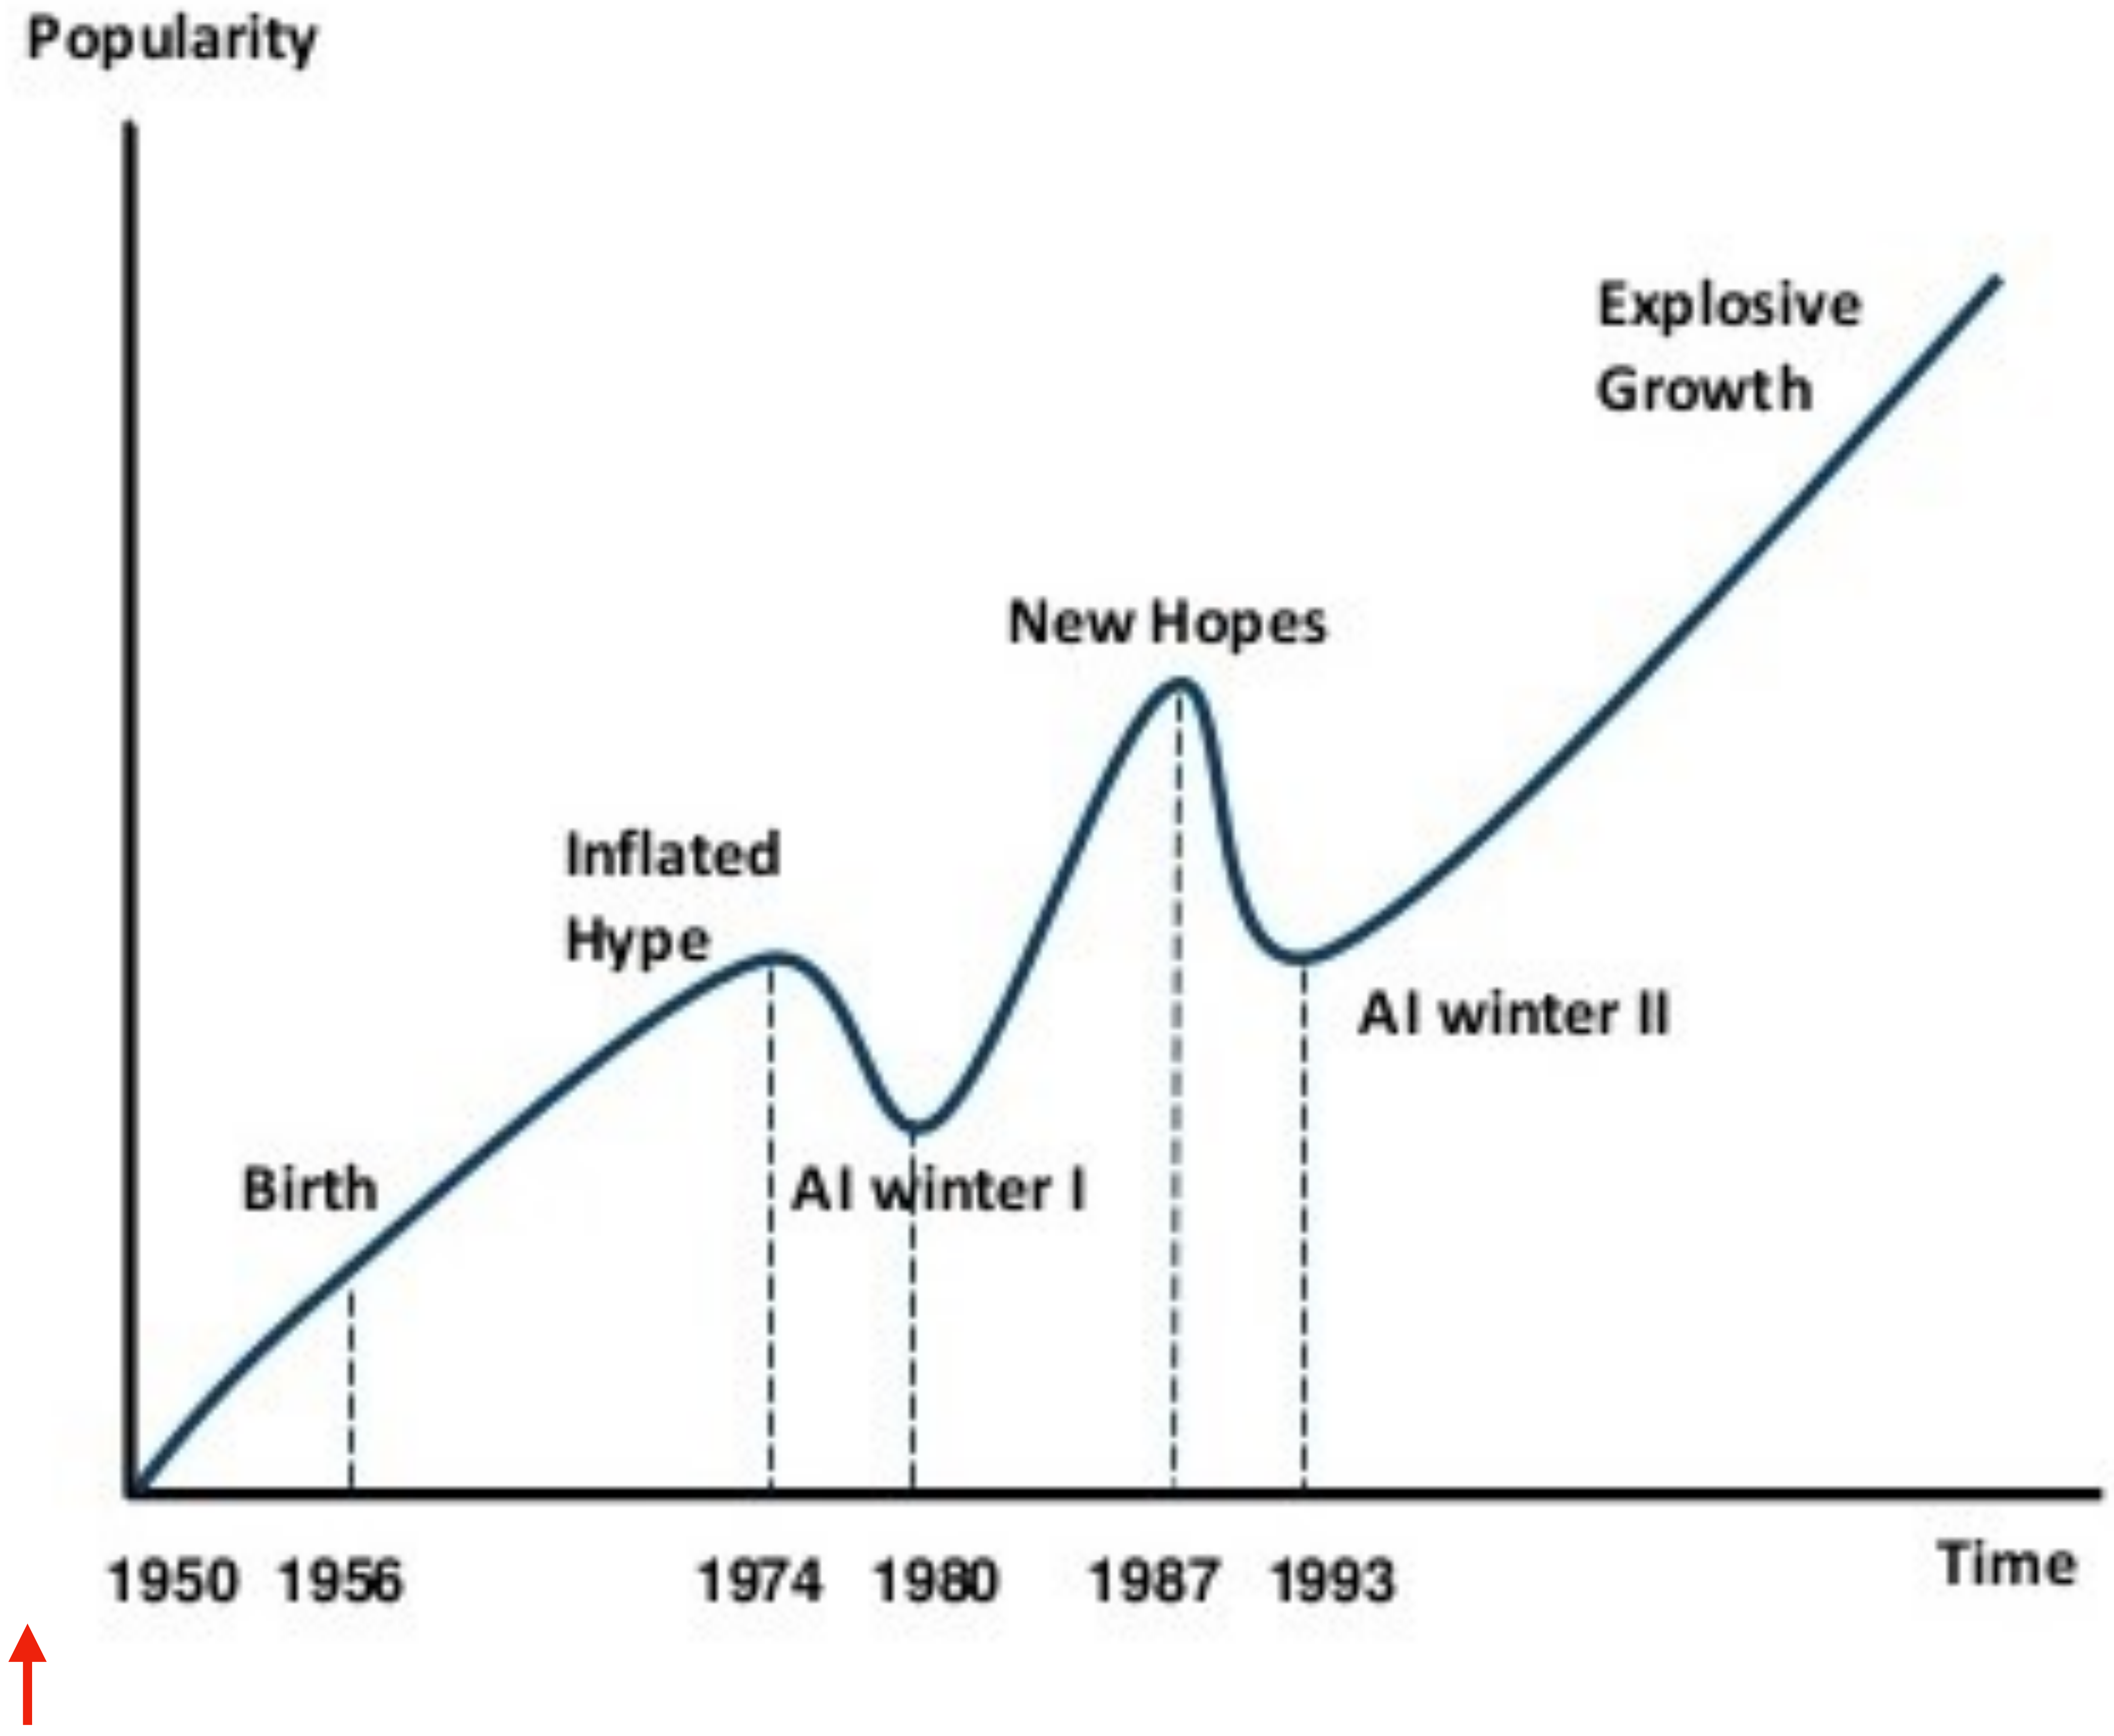

- AI has a long history of being “the next big thing”
- Multiple periods of “boom” and “bust”… including not one but two AI winters
- … so far

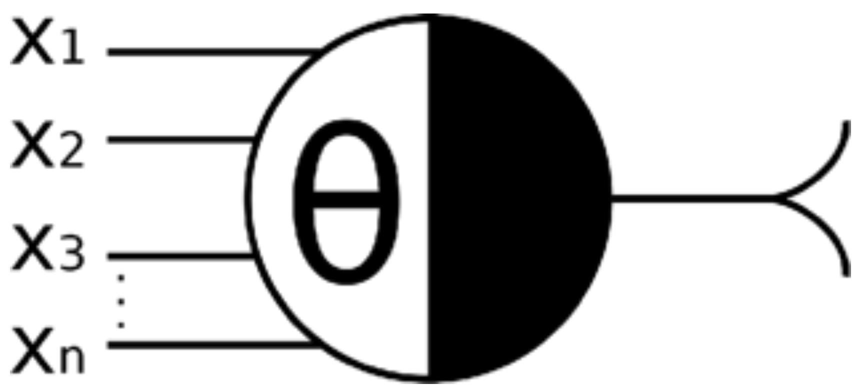

McCulloch & Pitts (1943)
Perceptron

Timeline of AI

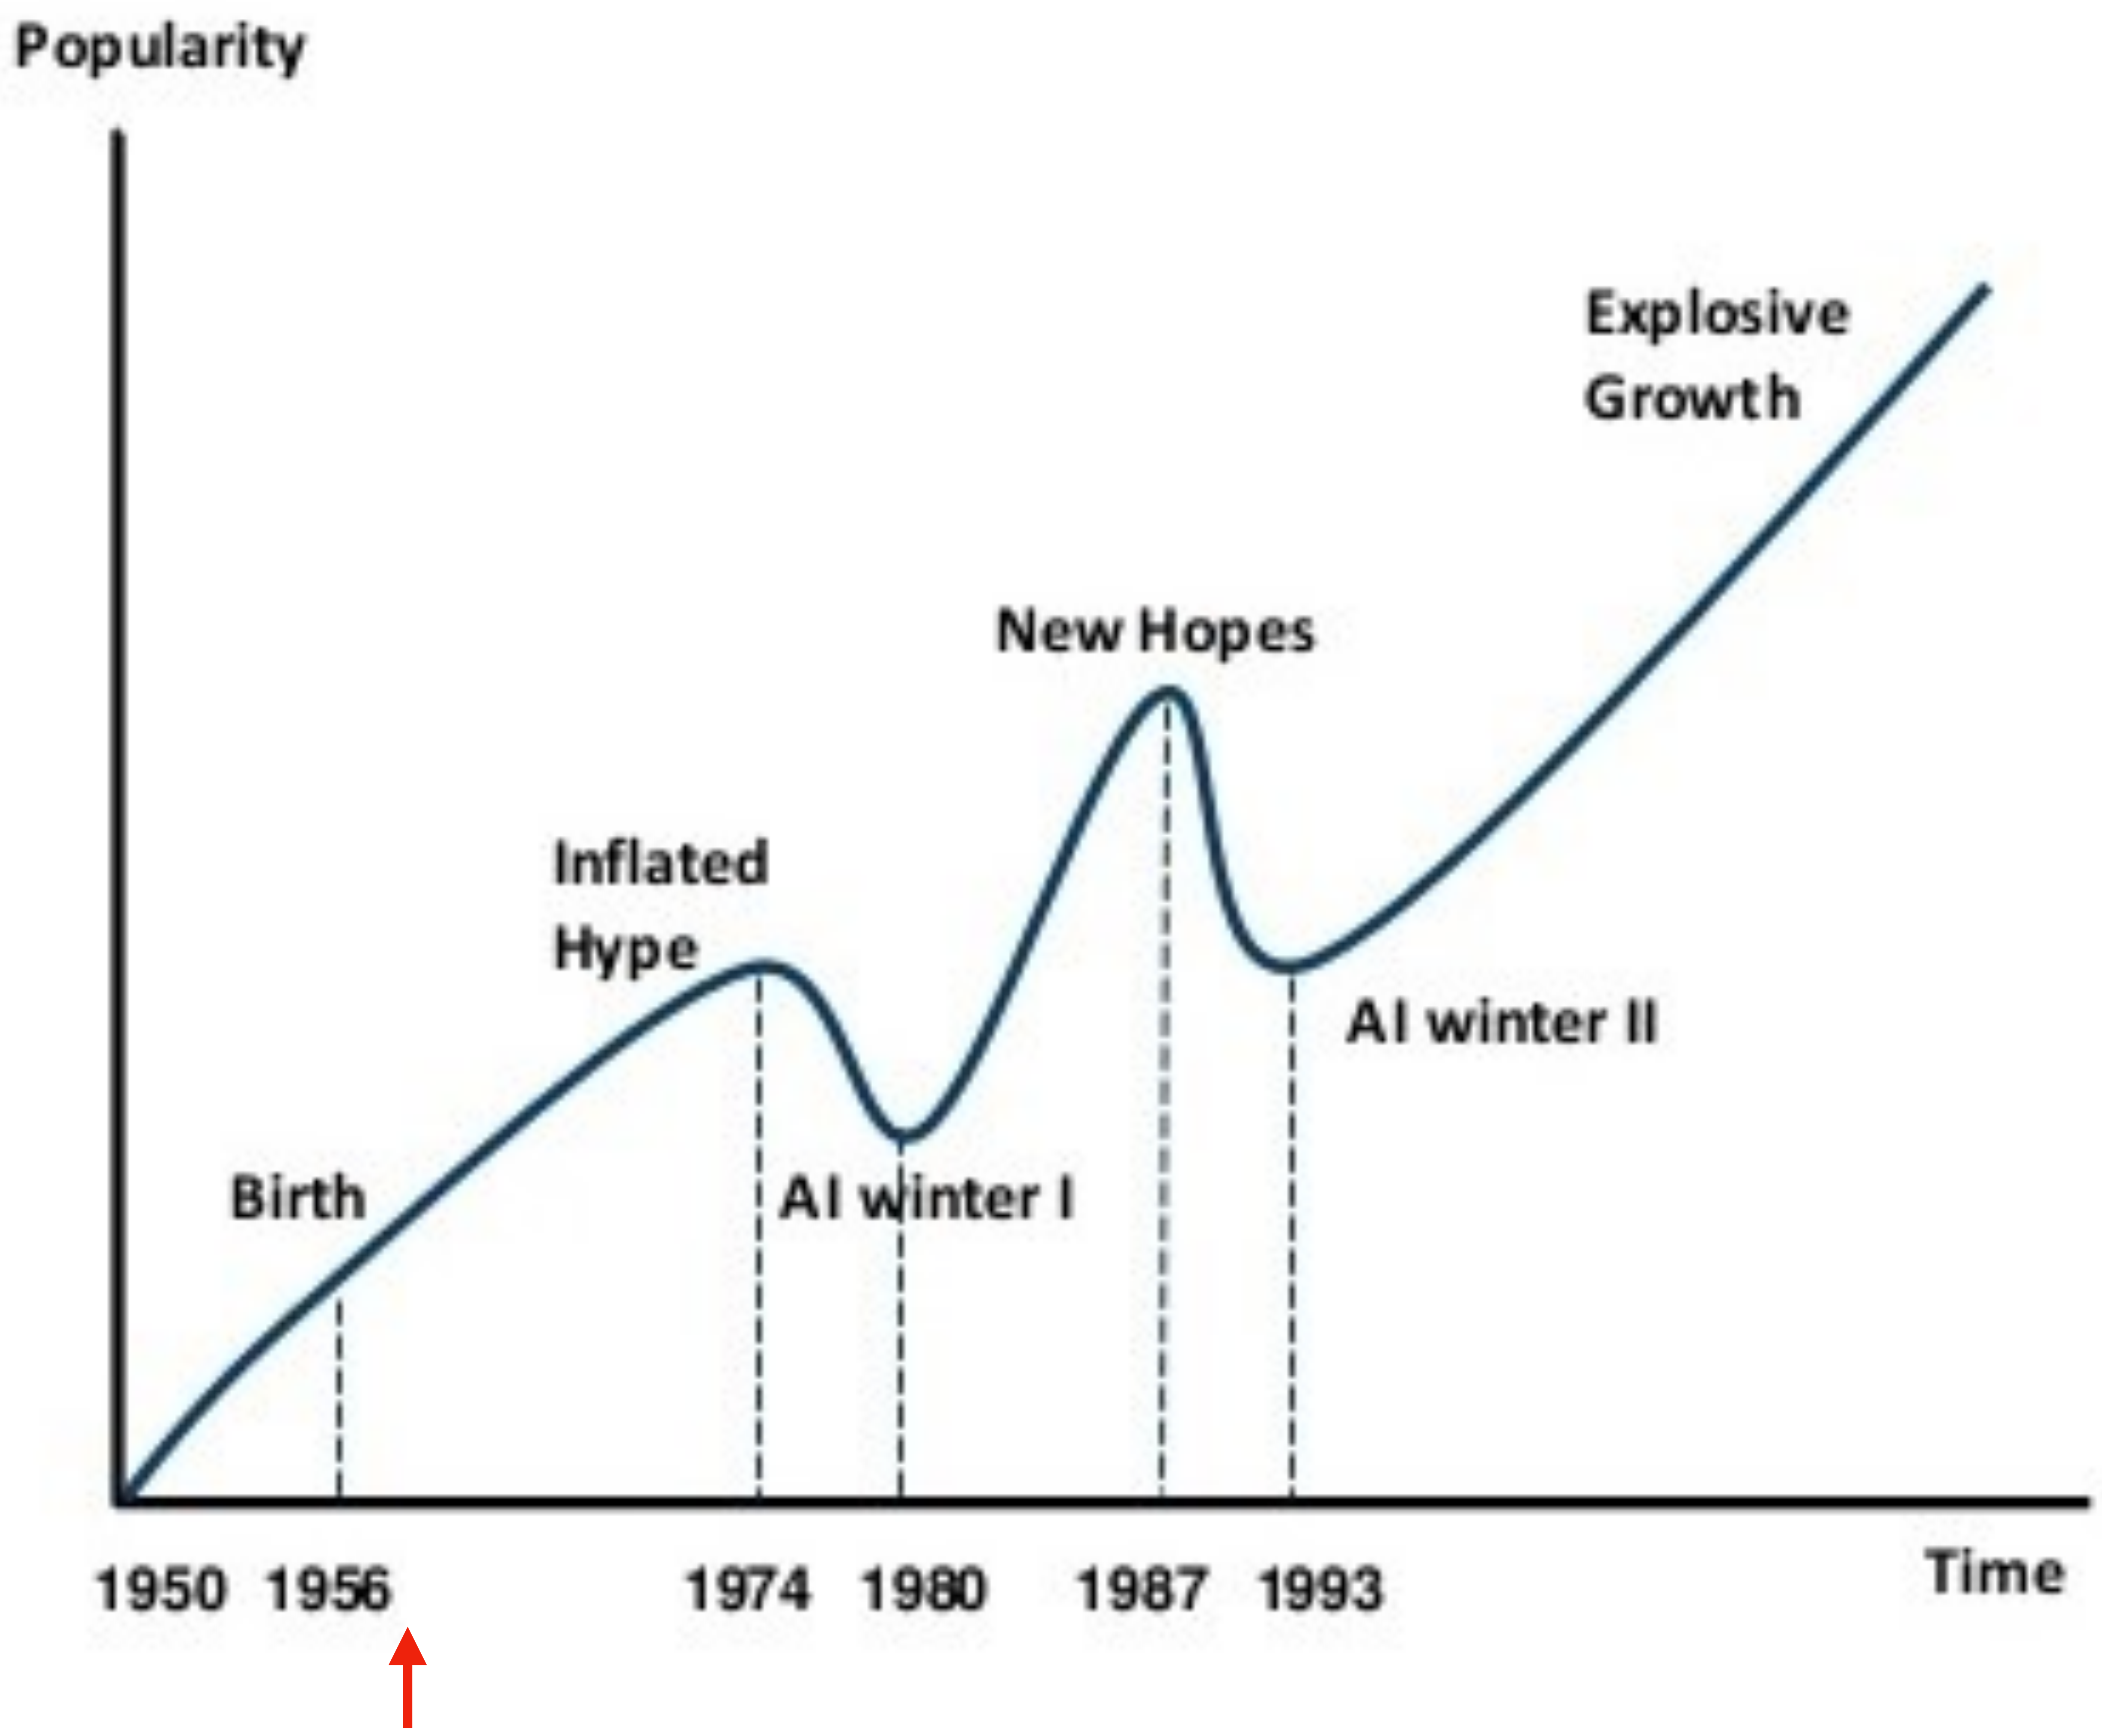

- AI has a long history of being “the next big thing”
- Multiple periods of “boom” and “bust”… including not one but two AI winters
- … so far

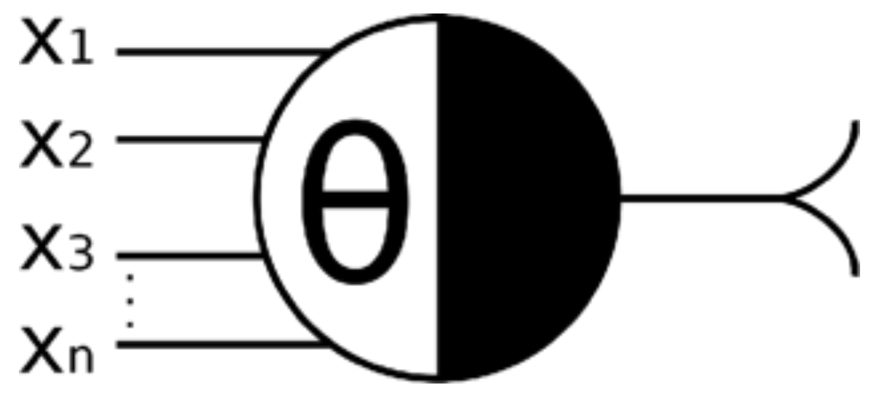

McCulloch & Pitts (1943) Perceptron

Rosenblatt (1958) Perceptron

Timeline of AI

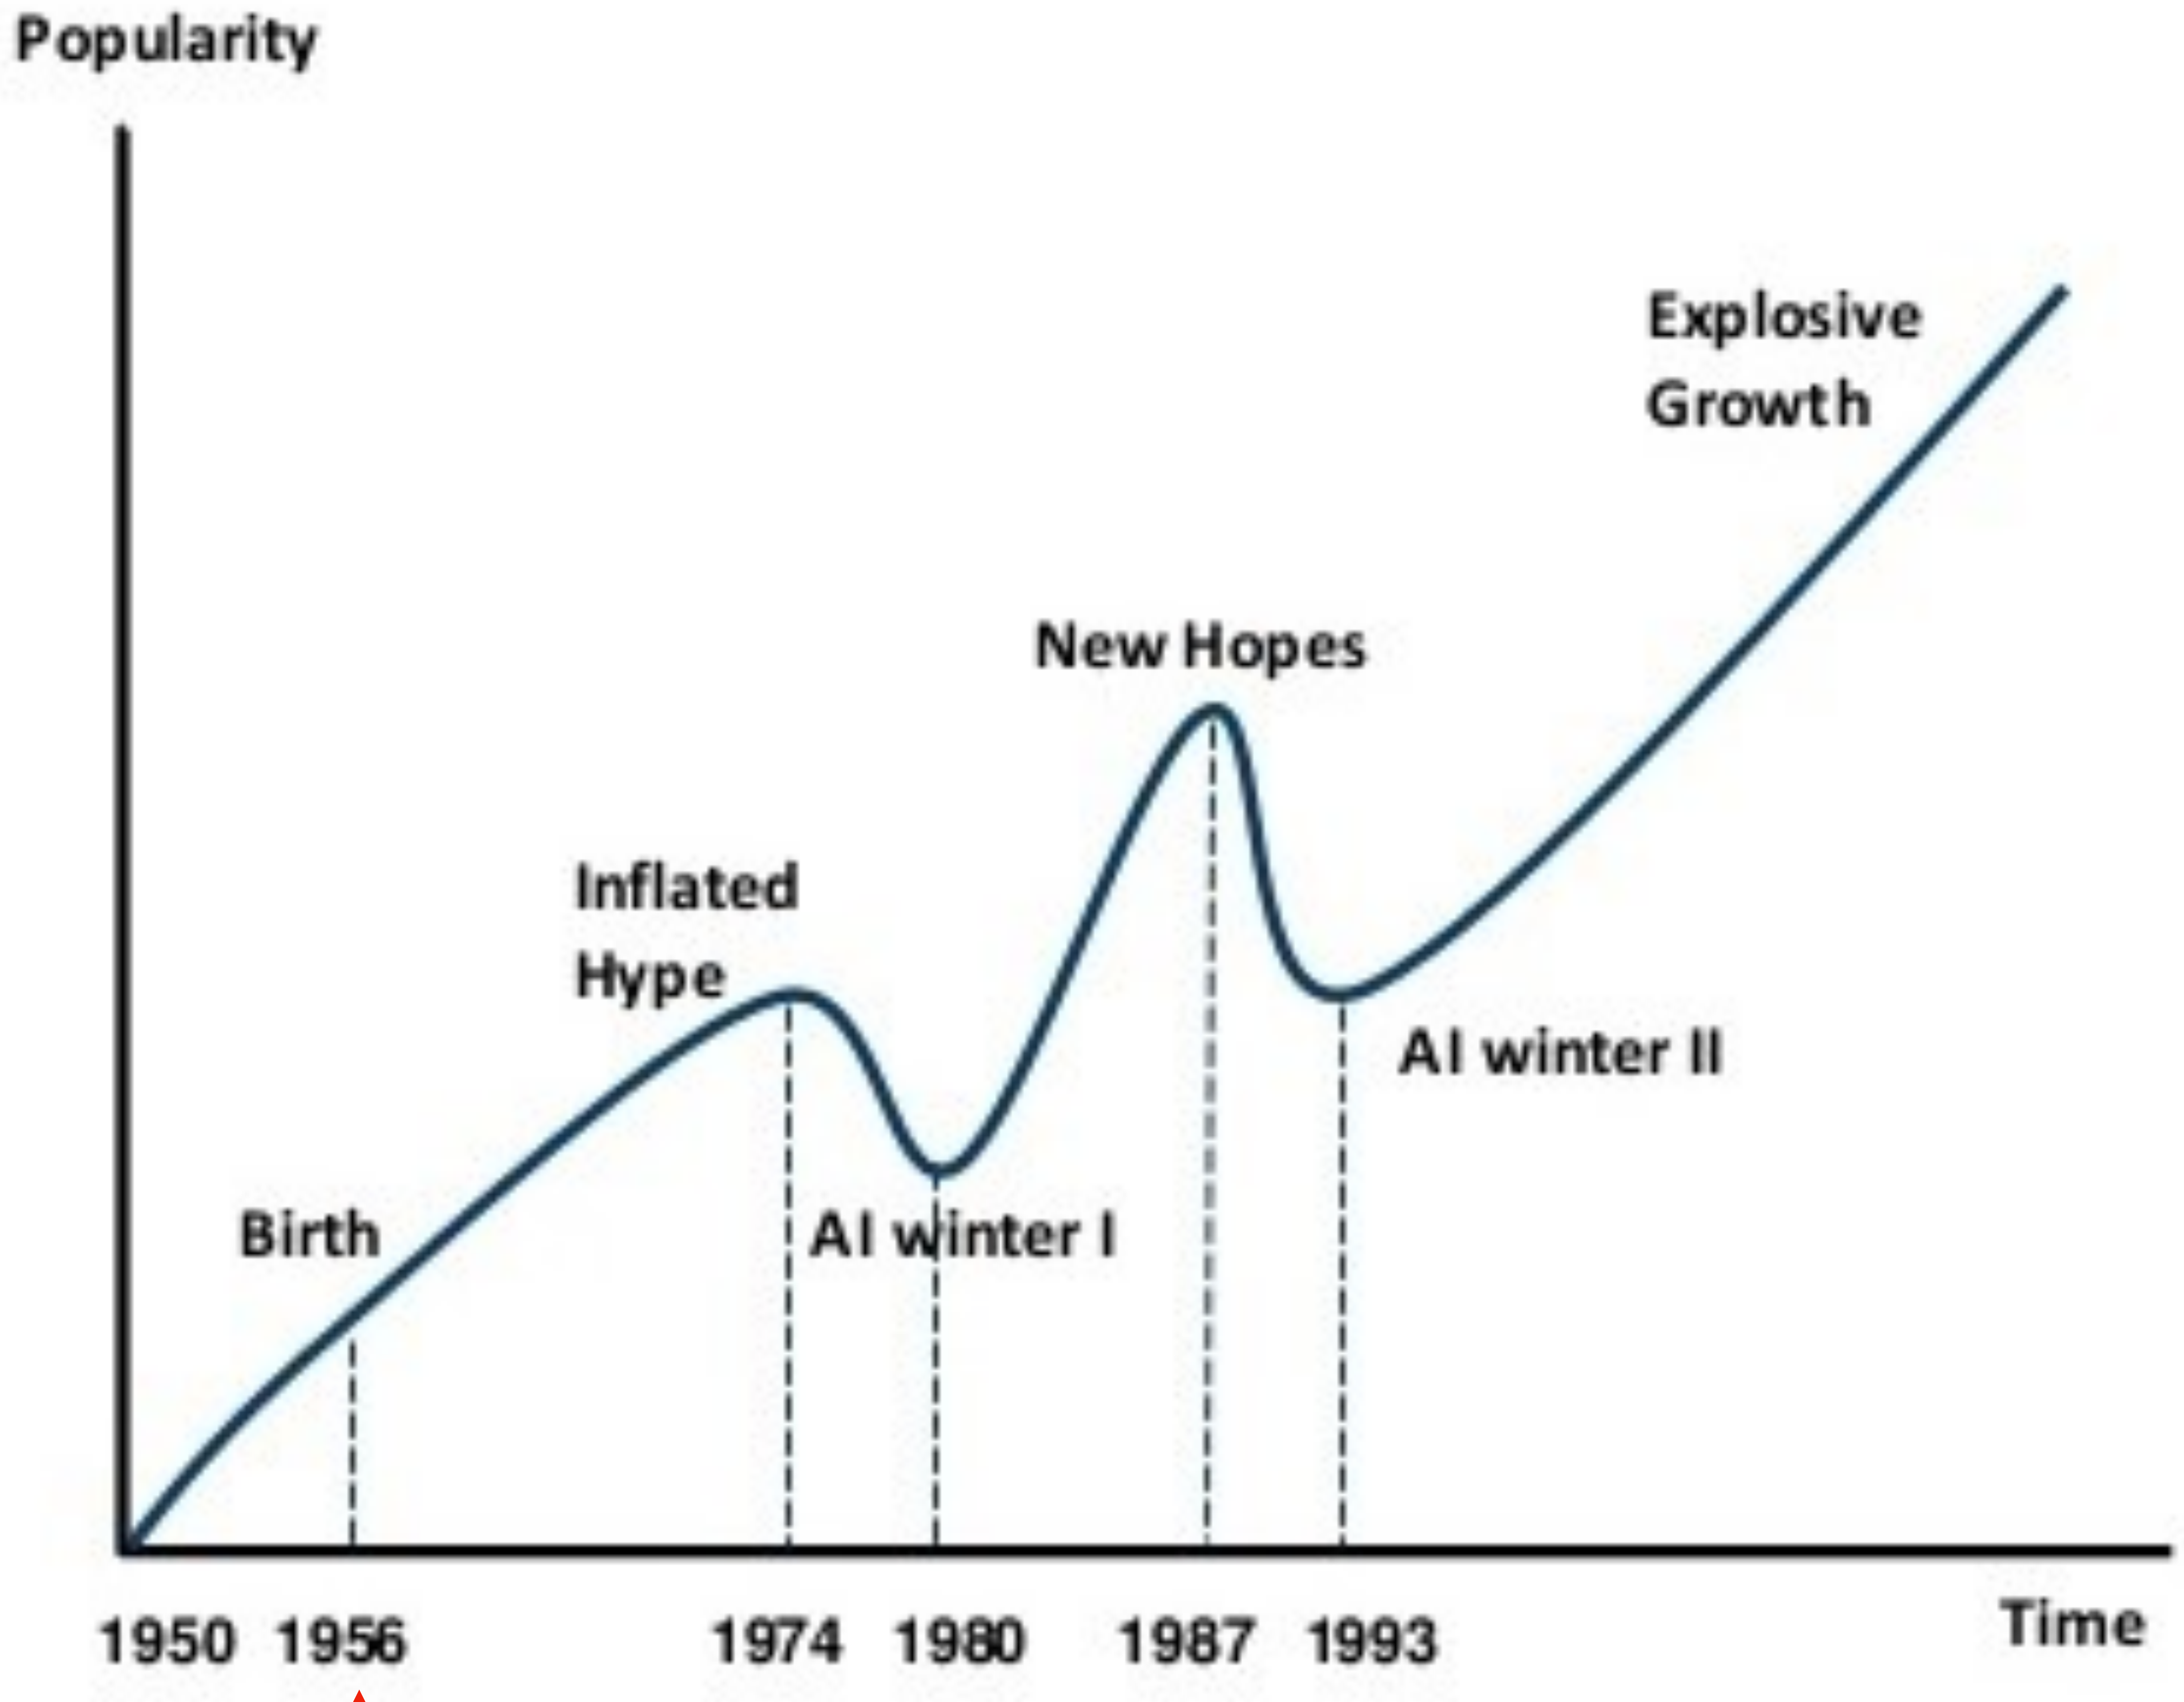

- **1956** Dartmouth workshop considered to be the founding event for AI as a field

Timeline of AI

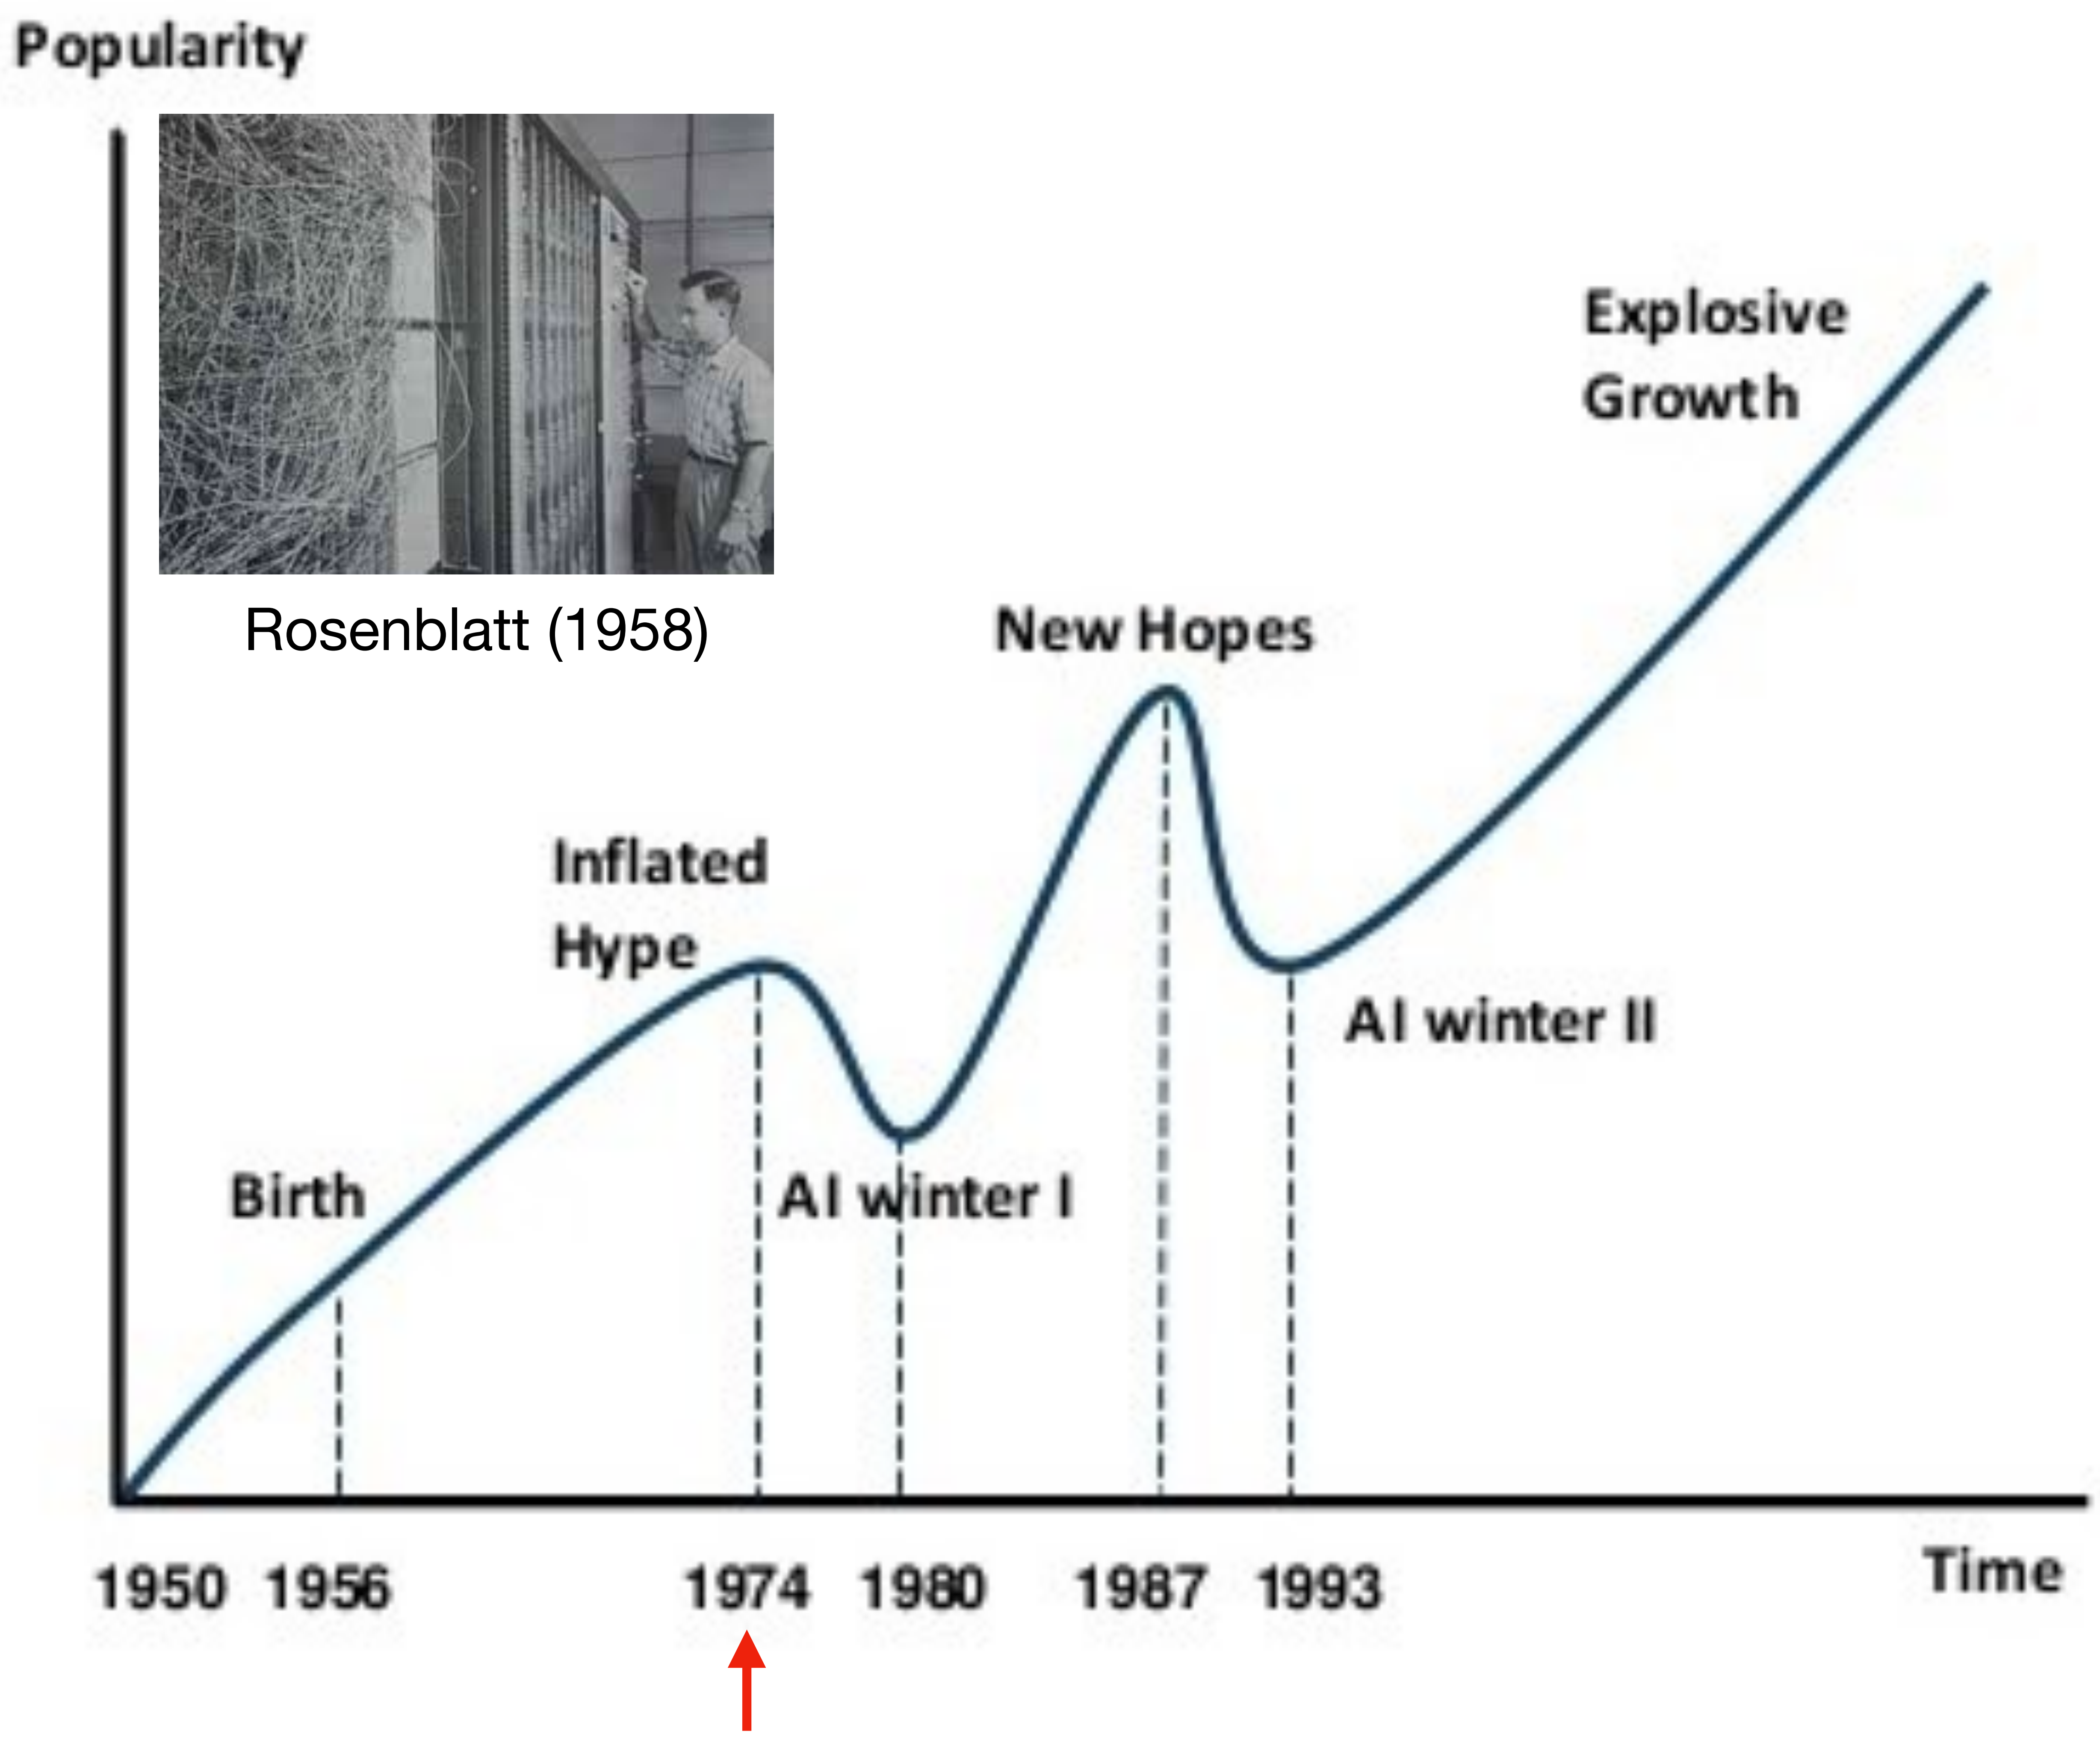

- **1956** Dartmouth workshop considered to be the founding event for AI as a field

Timeline of AI

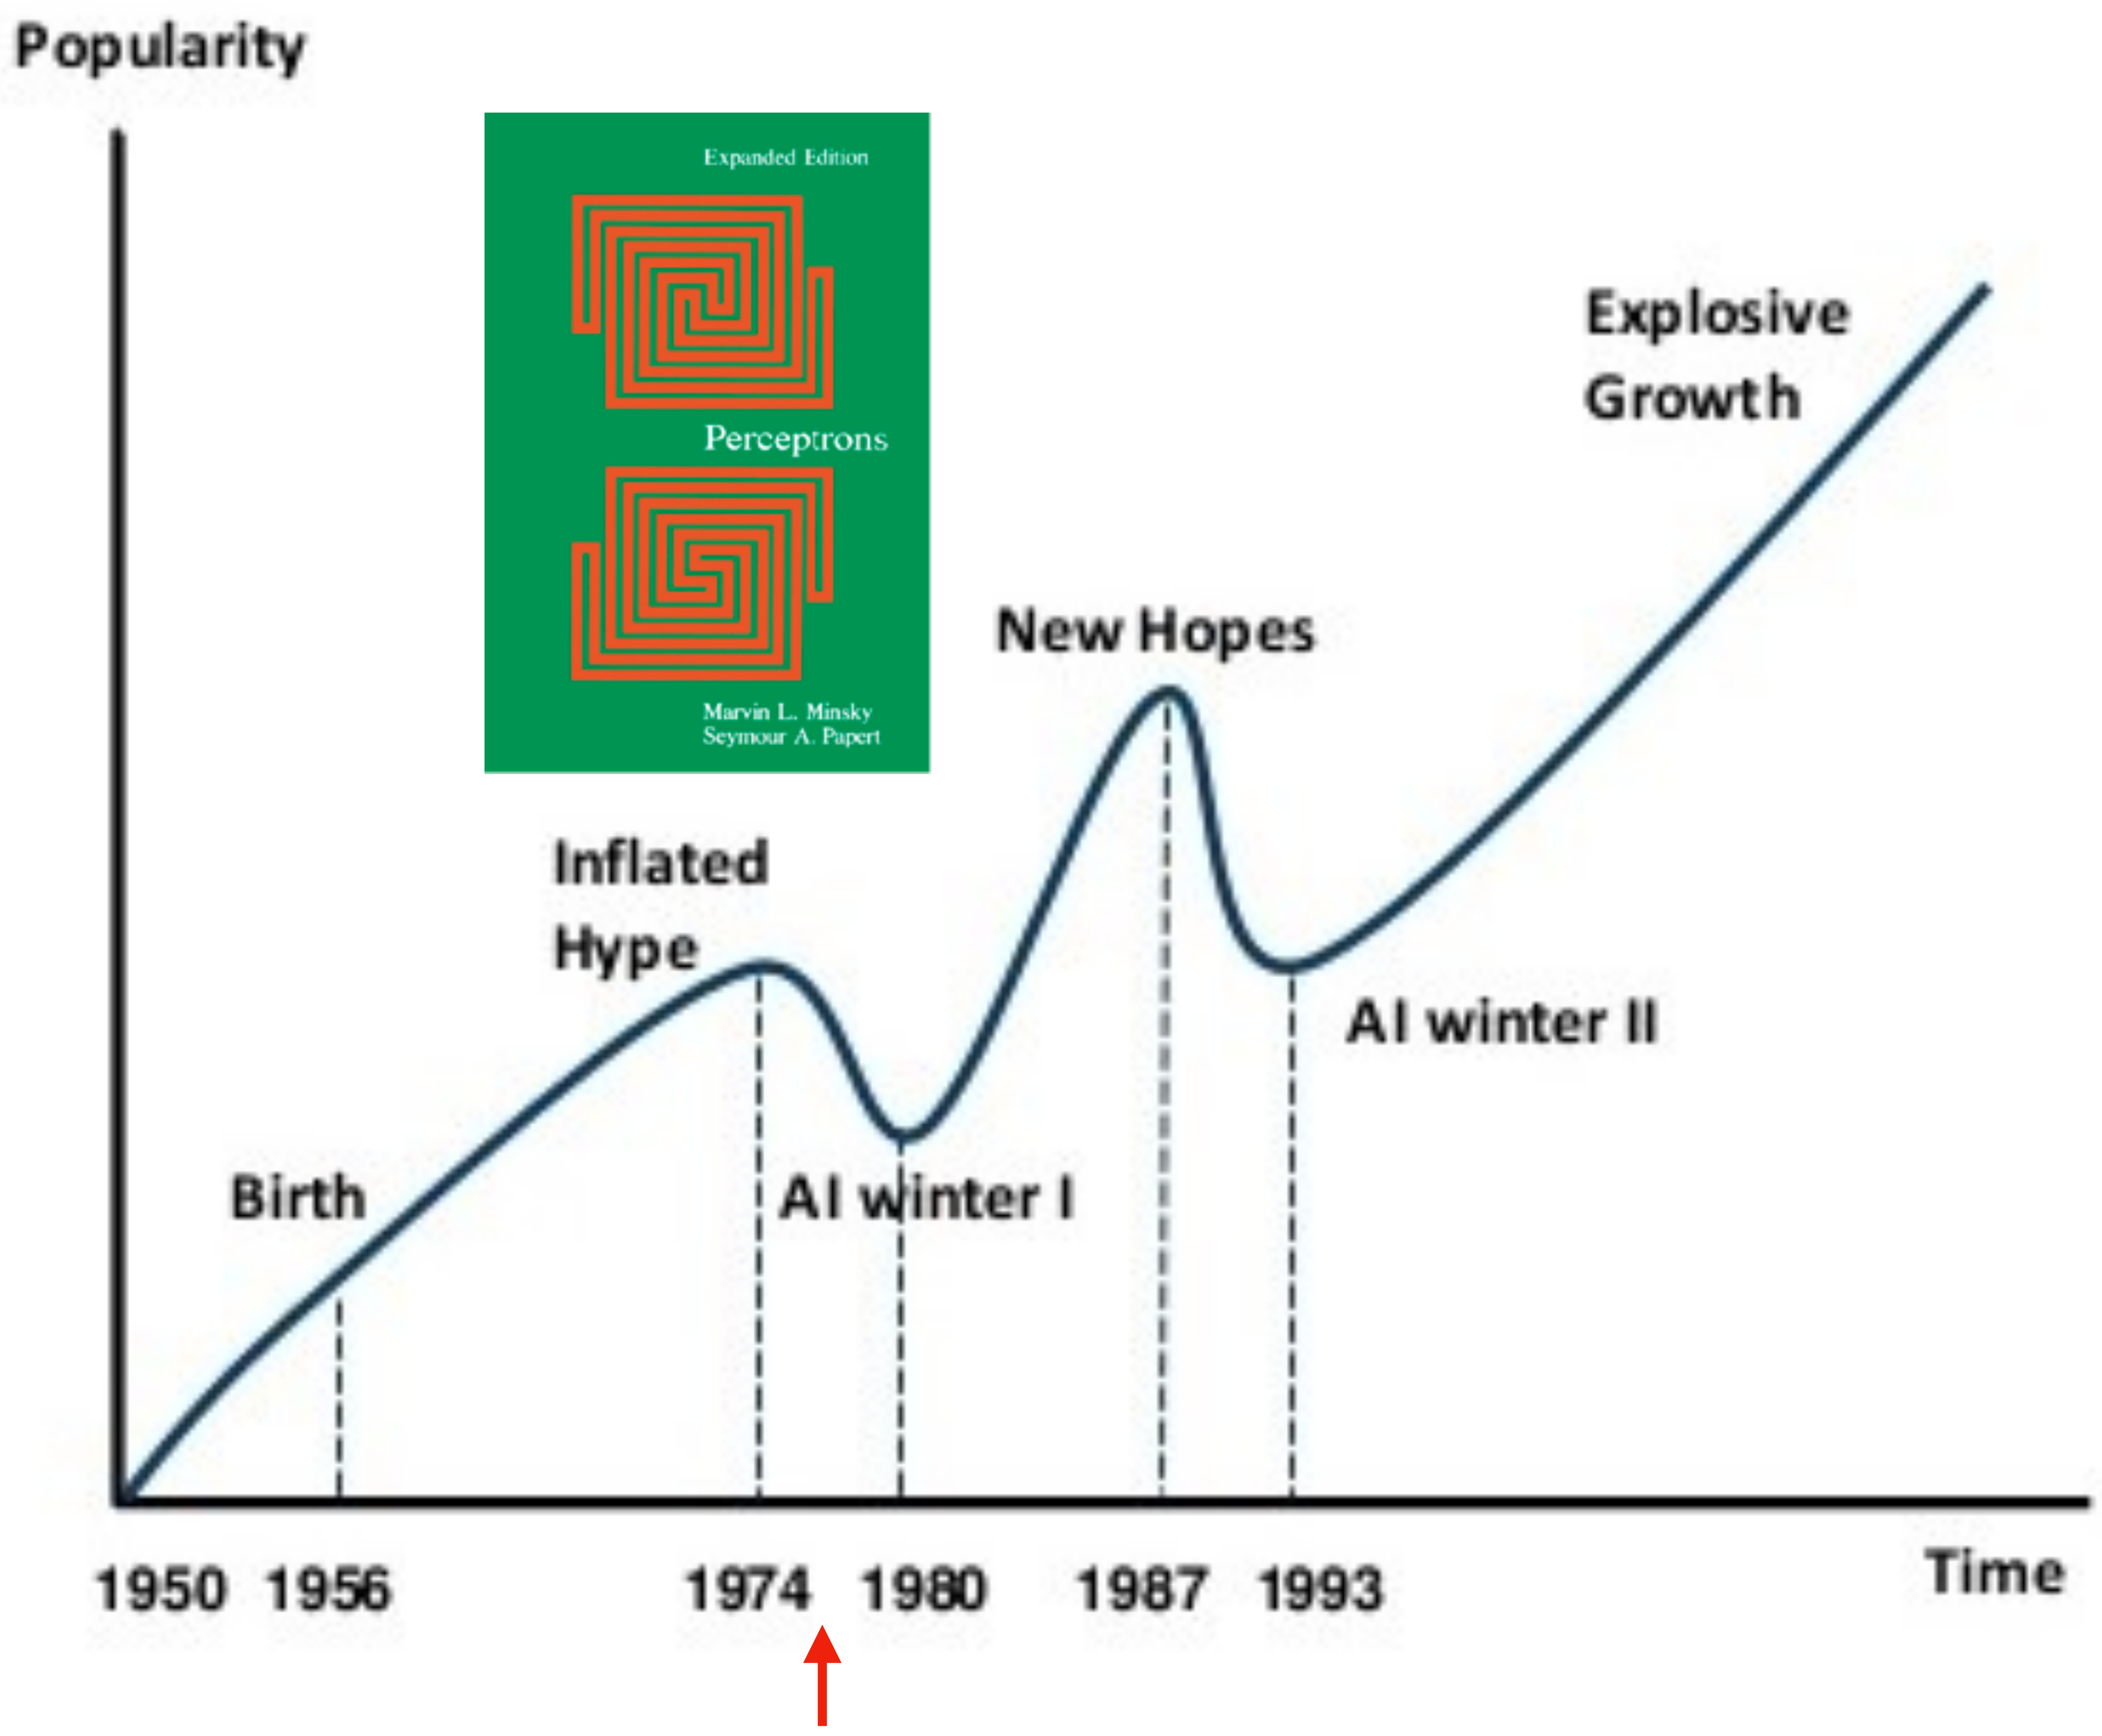

- Skepticism about Perceptrons not being able to solve XOR problems led to the first AI winter

Timeline of AI

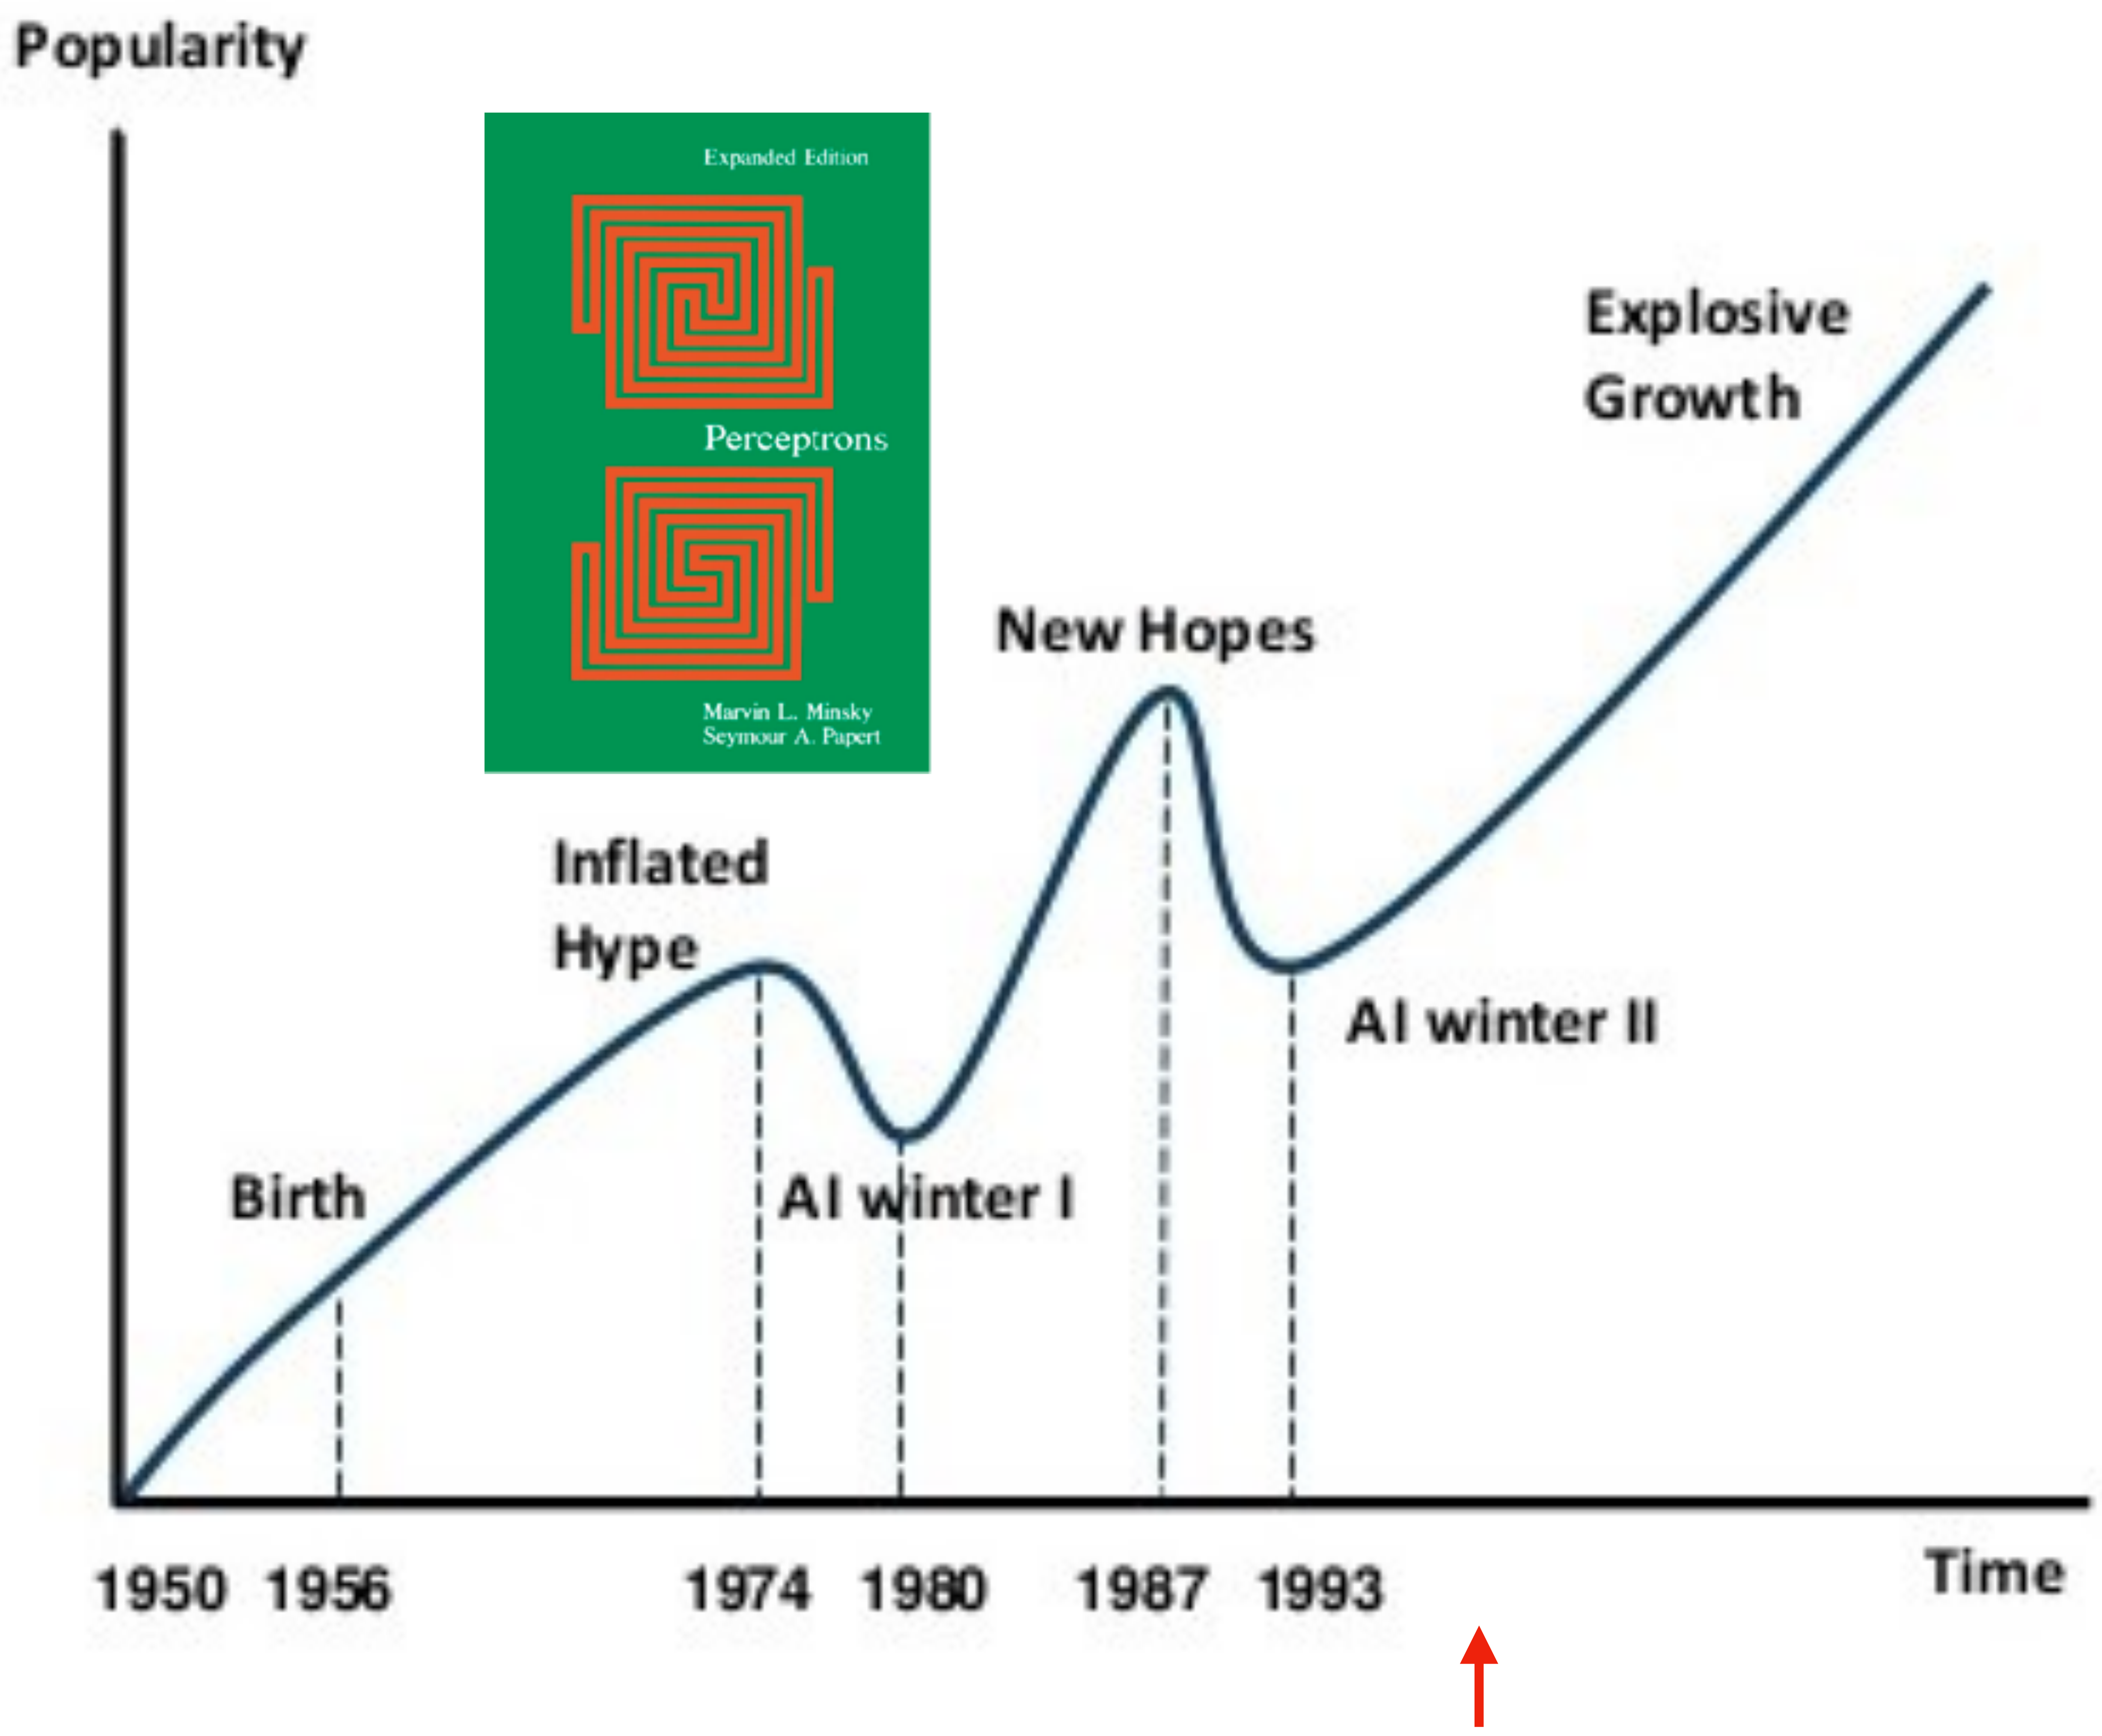

- Skepticism about Perceptrons not being able to solve XOR problems led to the first AI winter
- It wouldn't be until the deep learning revolution (~2006) that artificial neural networks would experience the same level of popularity

Timeline of AI

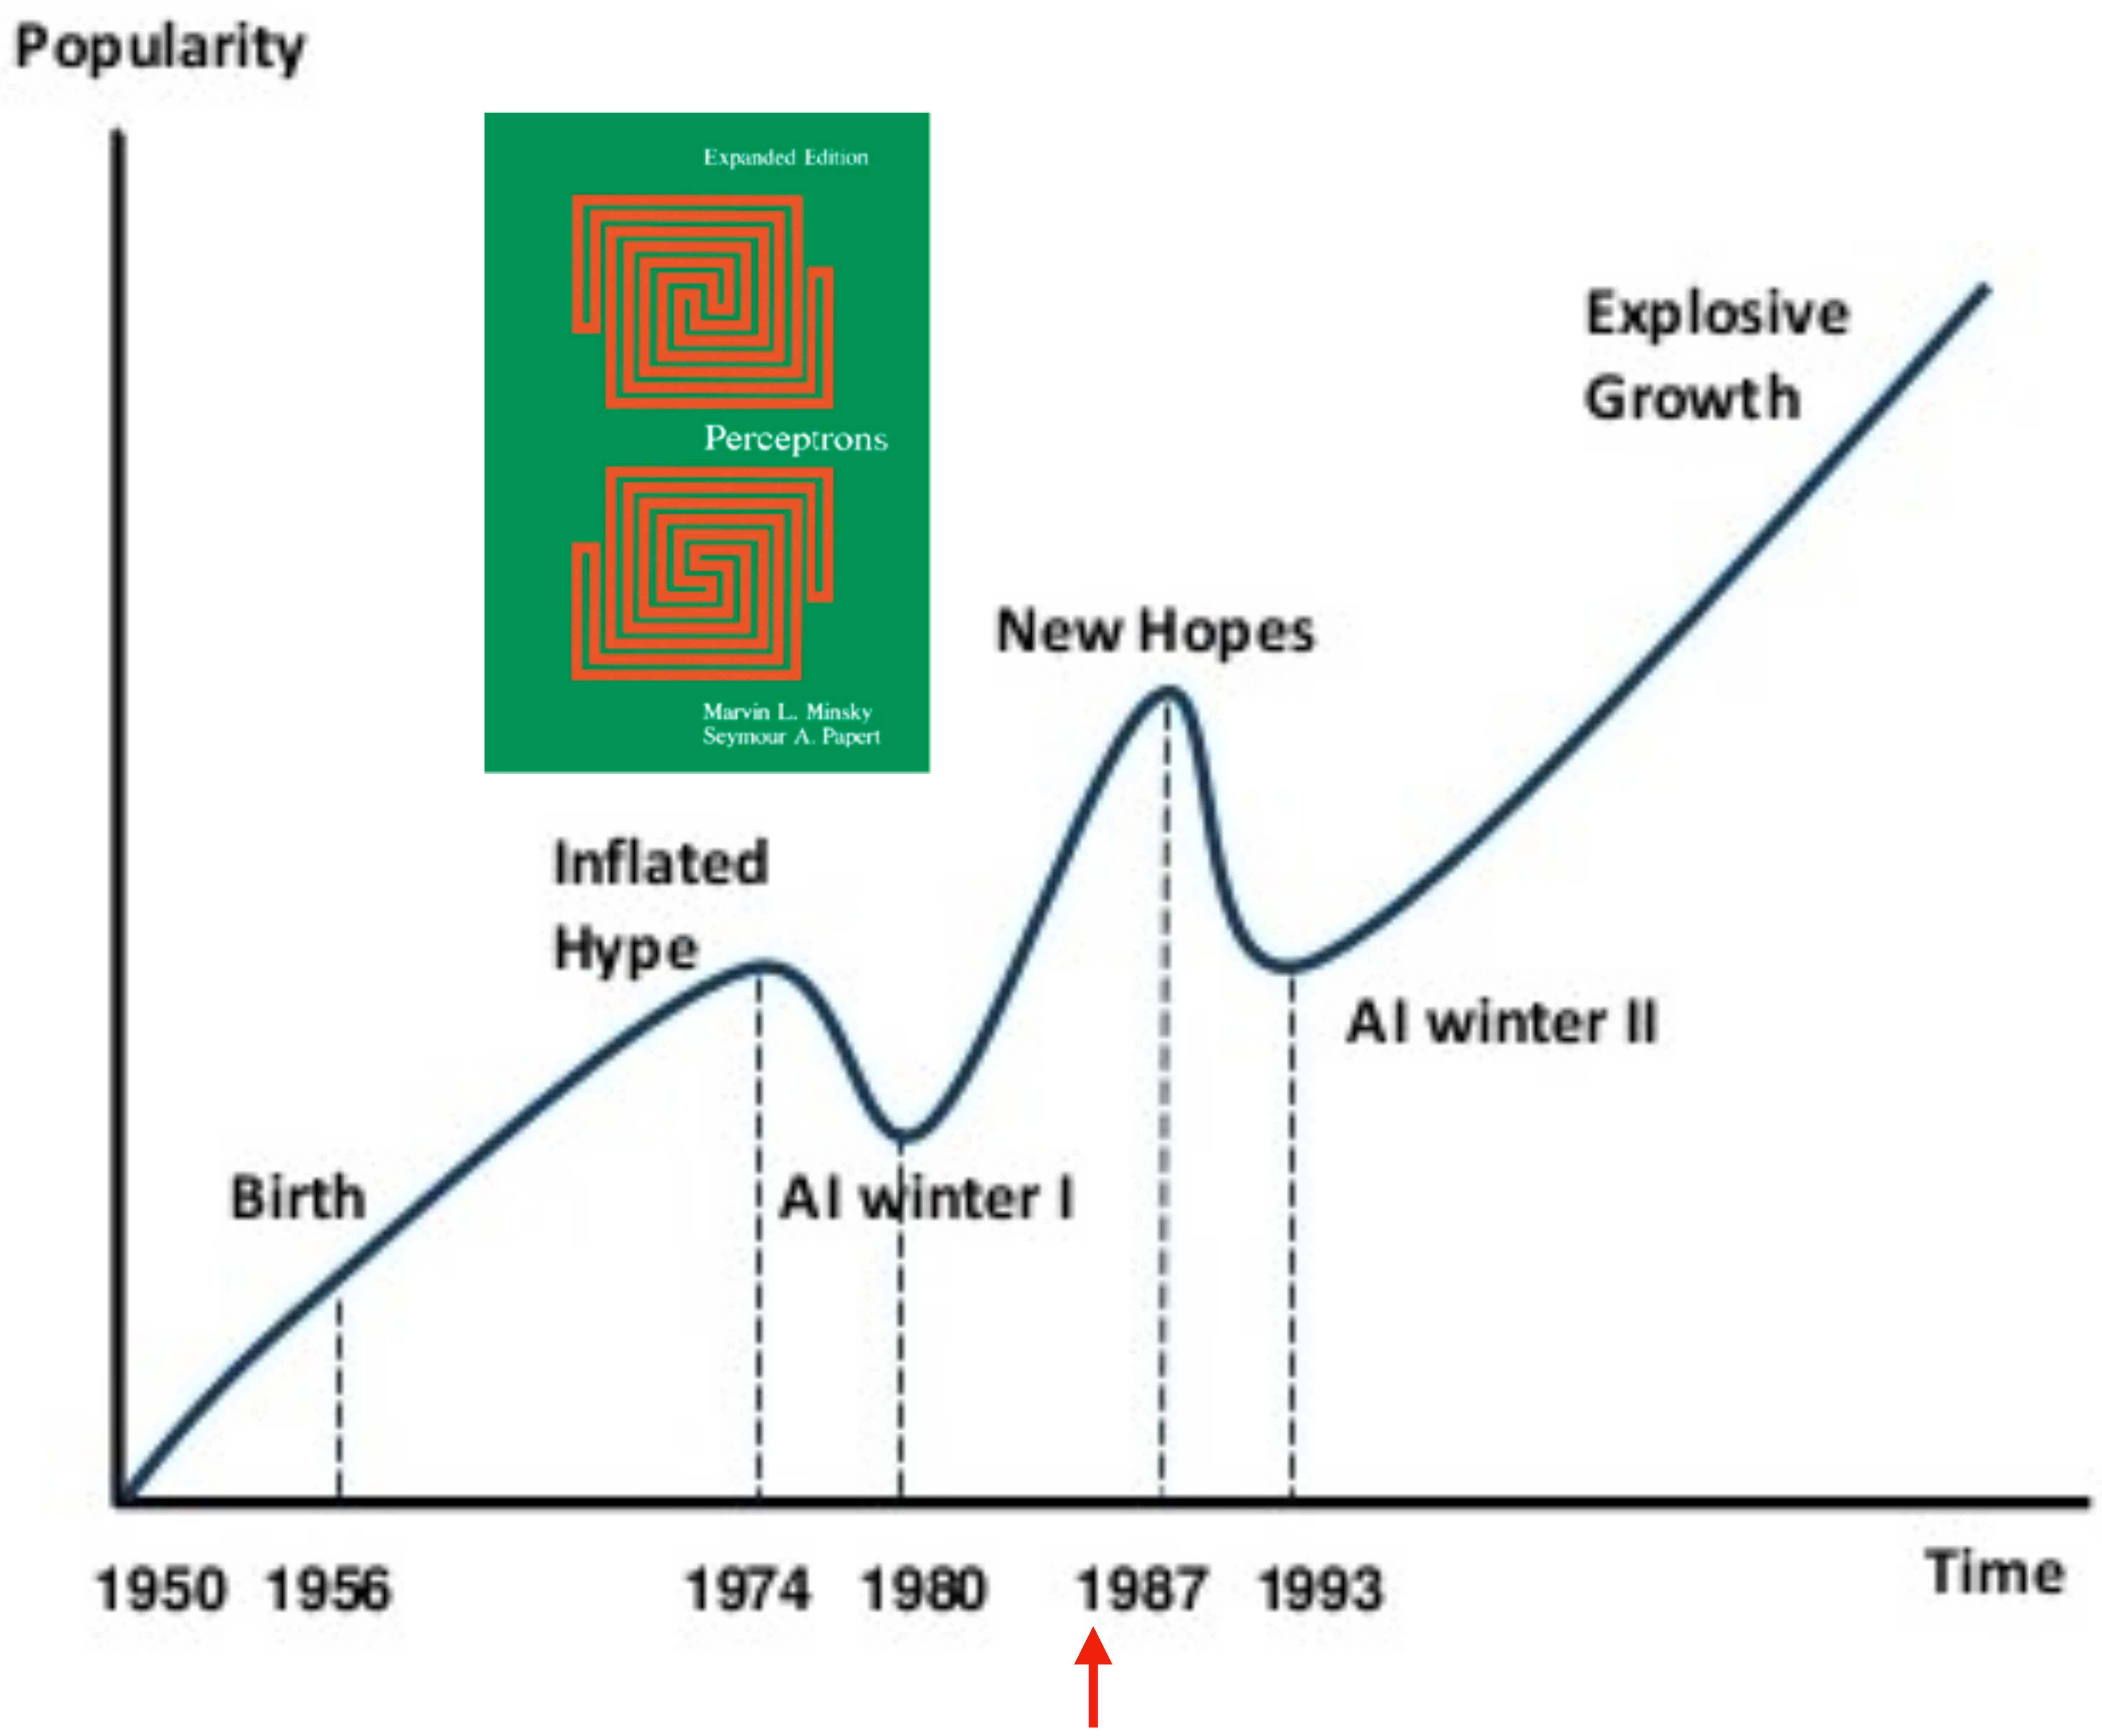

- Skepticism about Perceptrons not being able to solve XOR problems led to the first AI winter
- It wouldn't be until the deep learning revolution (~2006) that artificial neural networks would experience the same level of popularity
- But what happened in the 1980s when AI was more popular than ever? And why was there a 2nd AI winter?

Symbolic AI

- **Physical Symbol System hypothesis**:
  "A physical symbol system has the necessary and sufficient means for general intelligent action - Allen Newell and Herbert Simon (1976)"

Herbert Simon
& Allen Newell

Symbolic AI

- **Physical Symbol System hypothesis**:
  "A physical symbol system has the necessary and sufficient means for general intelligent action - Allen Newell and Herbert Simon (1976)"
  - **Symbols** can represent things in the world
    - e.g., (Apple), (ChatGPT), (Charley), etc…

Herbert Simon
& Allen Newell

Symbolic AI

- **Physical Symbol System hypothesis**:
  "A physical symbol system has the necessary and sufficient means for general intelligent action - Allen Newell and Herbert Simon (1976)"
  - **Symbols** can represent things in the world
    - e.g., (Apple), (ChatGPT), (Charley), etc…

Herbert Simon
& Allen Newell

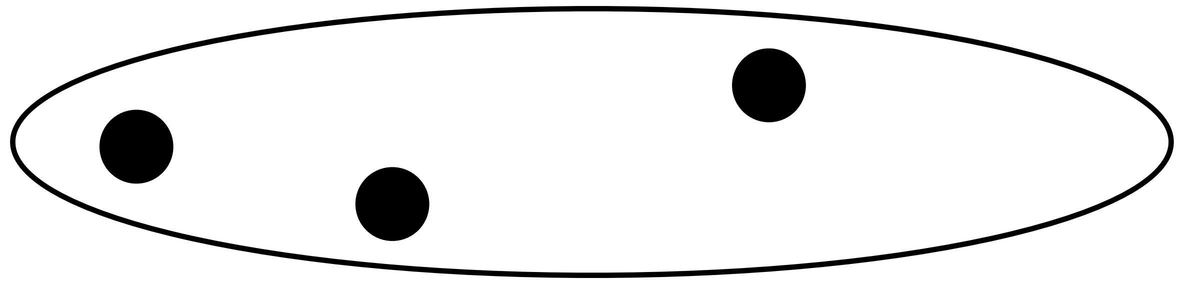

Symbolic AI

- **Physical Symbol System hypothesis**:
  "A physical symbol system has the necessary and sufficient means for general intelligent action - Allen Newell and Herbert Simon (1976)"
  - **Symbols** can represent things in the world
    - e.g., (Apple), (ChatGPT), (Charley), etc…

Herbert Simon
& Allen Newell

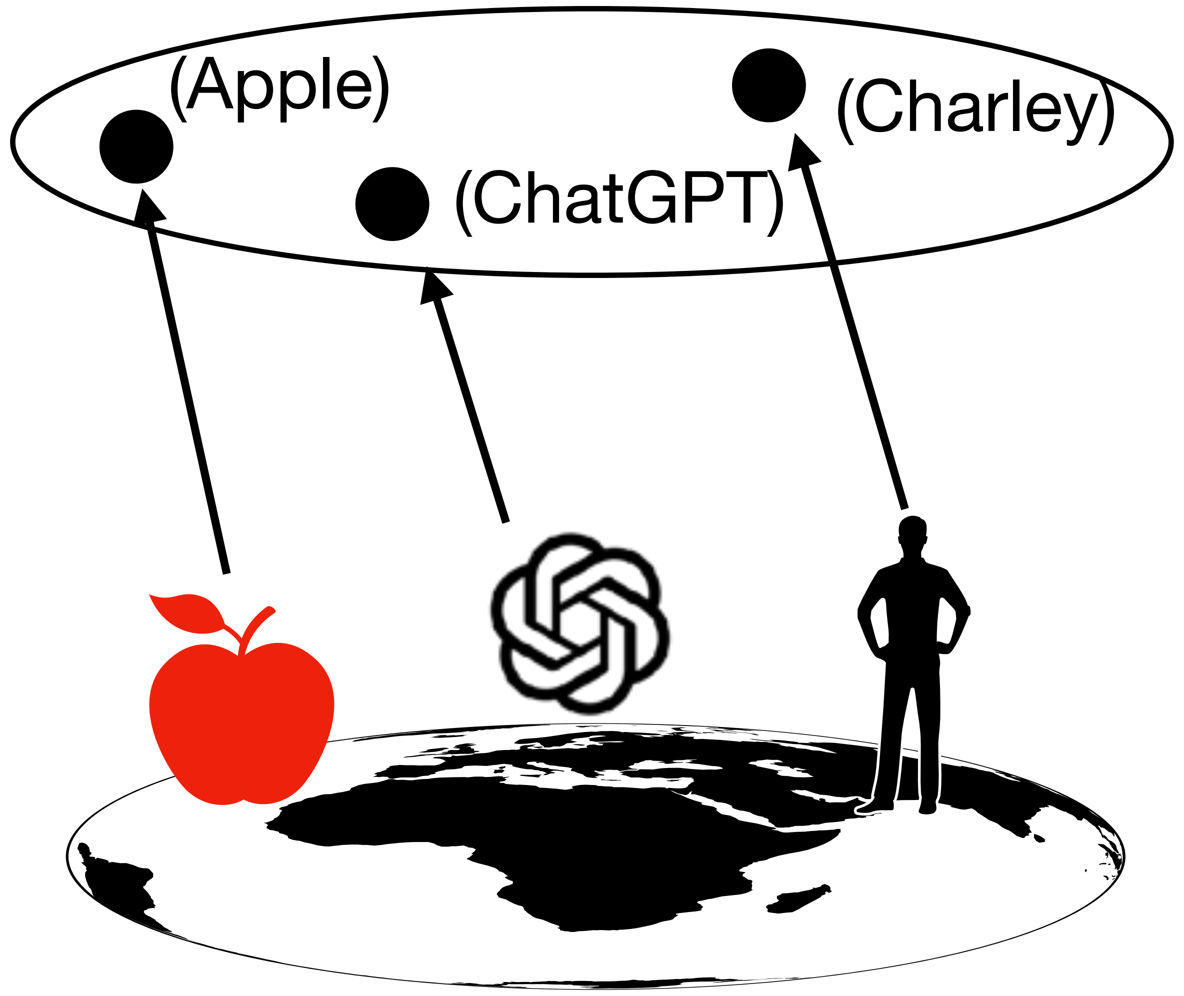

Symbolic AI

- **Physical Symbol System hypothesis**:
  "A physical symbol system has the necessary and sufficient means for general intelligent action - Allen Newell and Herbert Simon (1976)"
  - **Symbols** can represent things in the world
    - e.g., (Apple), (ChatGPT), (Charley), etc…
  - **Relations** can be i) predicates that describes a symbol or ii) verbs describing how symbols interact with other symbols
    - i) red(Apple), unreliable(ChatGPT), instructor(Charley)
    - ii) eat(Charley, Apple), generatePicture(ChatGPT, Apple)

Herbert Simon
& Allen Newell

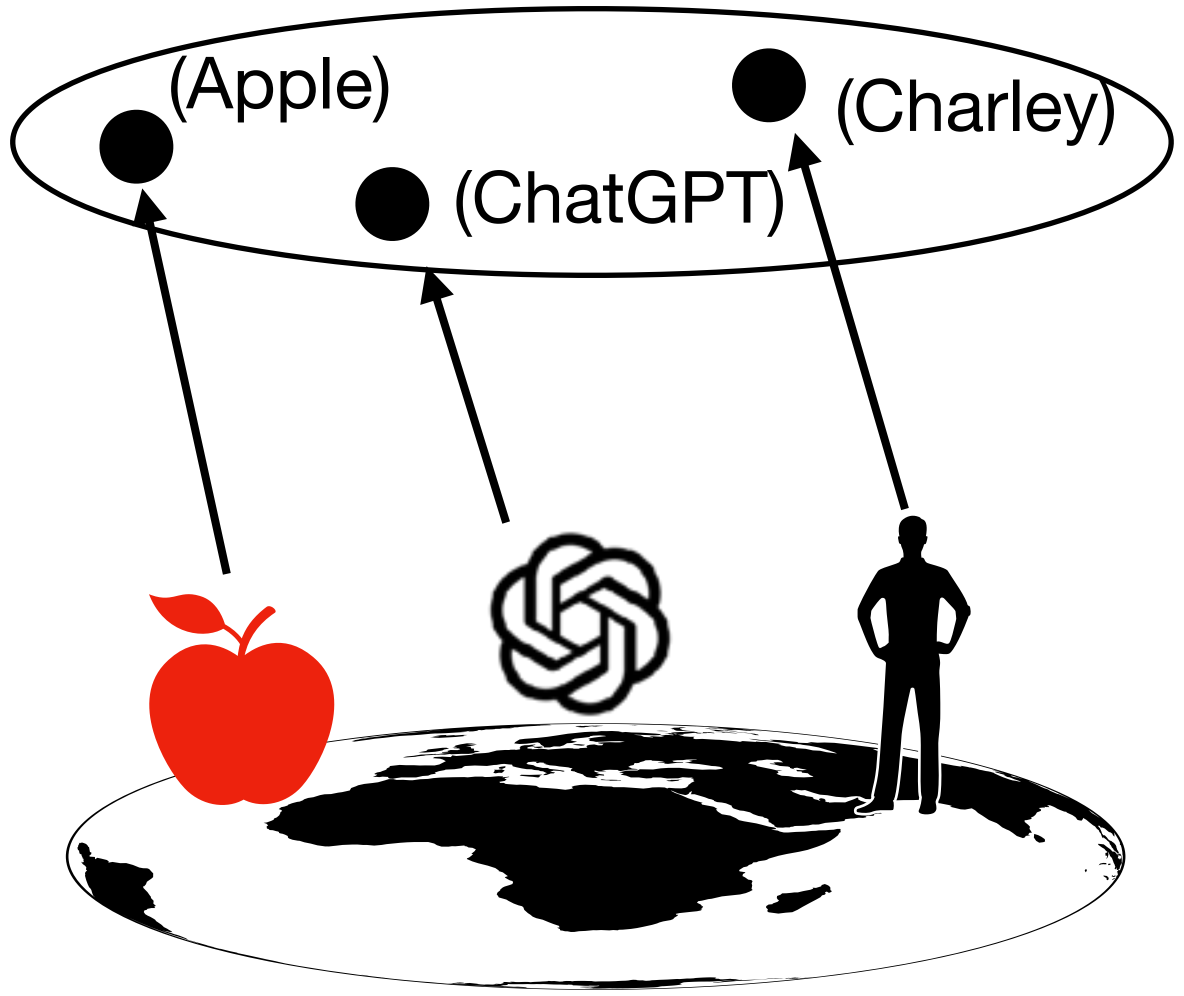

Symbolic AI

- **Physical Symbol System hypothesis**:
  “A physical symbol system has the necessary and sufficient means for general intelligent action - Allen Newell and Herbert Simon (1976)”
  - **Symbols** can represent things in the world
    - e.g., (Apple), (ChatGPT), (Charley), etc…
  - **Relations** can be i) predicates that describes a symbol or ii) verbs describing how symbols interact with other symbols
    - i) red(Apple), unreliable(ChatGPT), instructor(Charley)
    - ii) eat(Charley, Apple), generatePicture(ChatGPT, Apple)

Herbert Simon
& Allen Newell

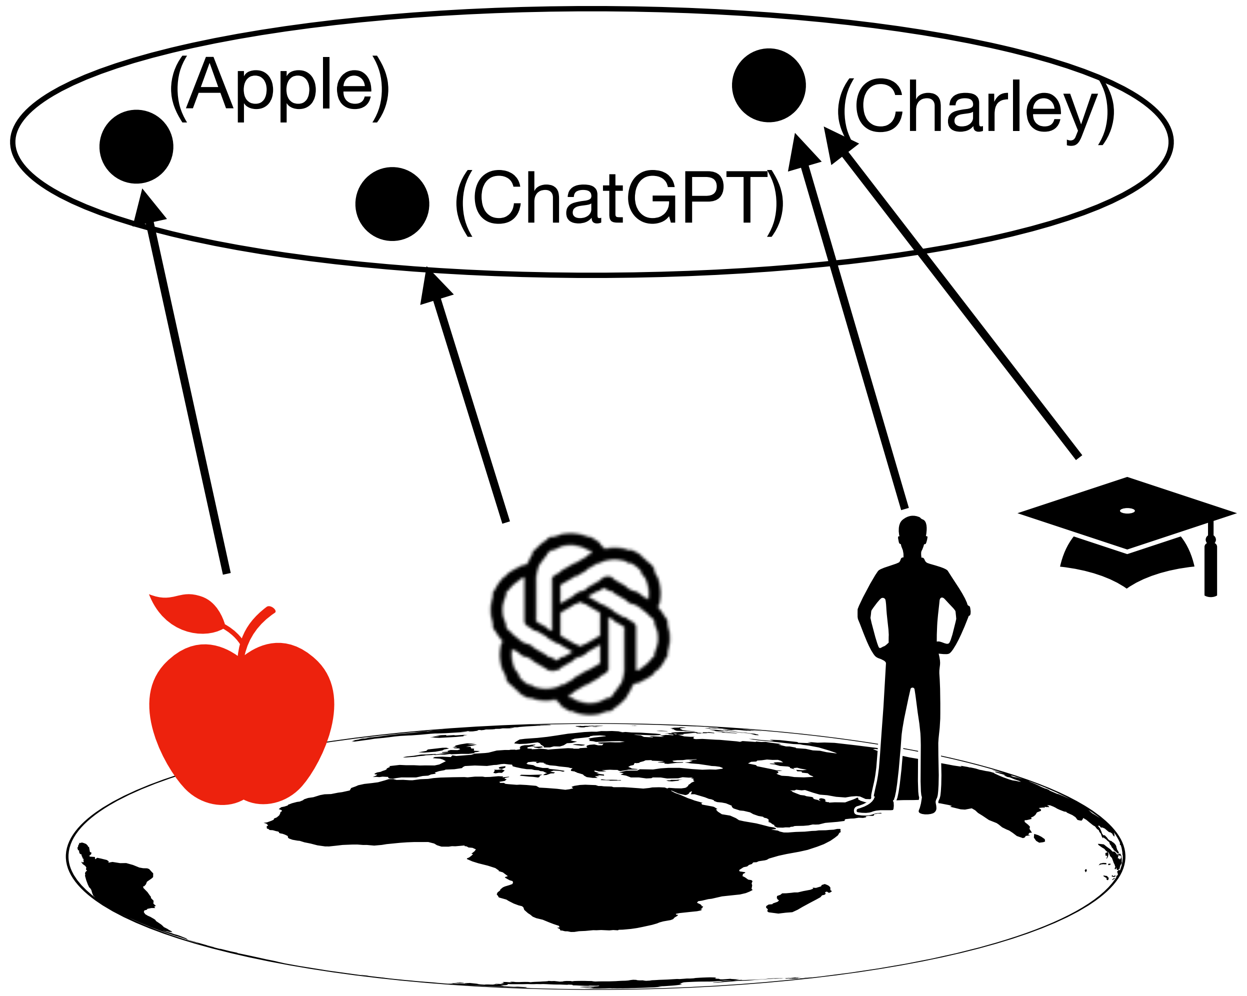

Symbolic AI

- **Physical Symbol System hypothesis**:
  "A physical symbol system has the necessary and sufficient means for general intelligent action - Allen Newell and Herbert Simon (1976)"
  - **Symbols** can represent things in the world
    - e.g., (Apple), (ChatGPT), (Charley), etc…
  - **Relations** can be i) predicates that describes a symbol or ii) verbs describing how symbols interact with other symbols
    - i) red(Apple), unreliable(ChatGPT), instructor(Charley)
    - ii) eat(Charley, Apple), generatePicture(ChatGPT, Apple)

Herbert Simon
& Allen Newell

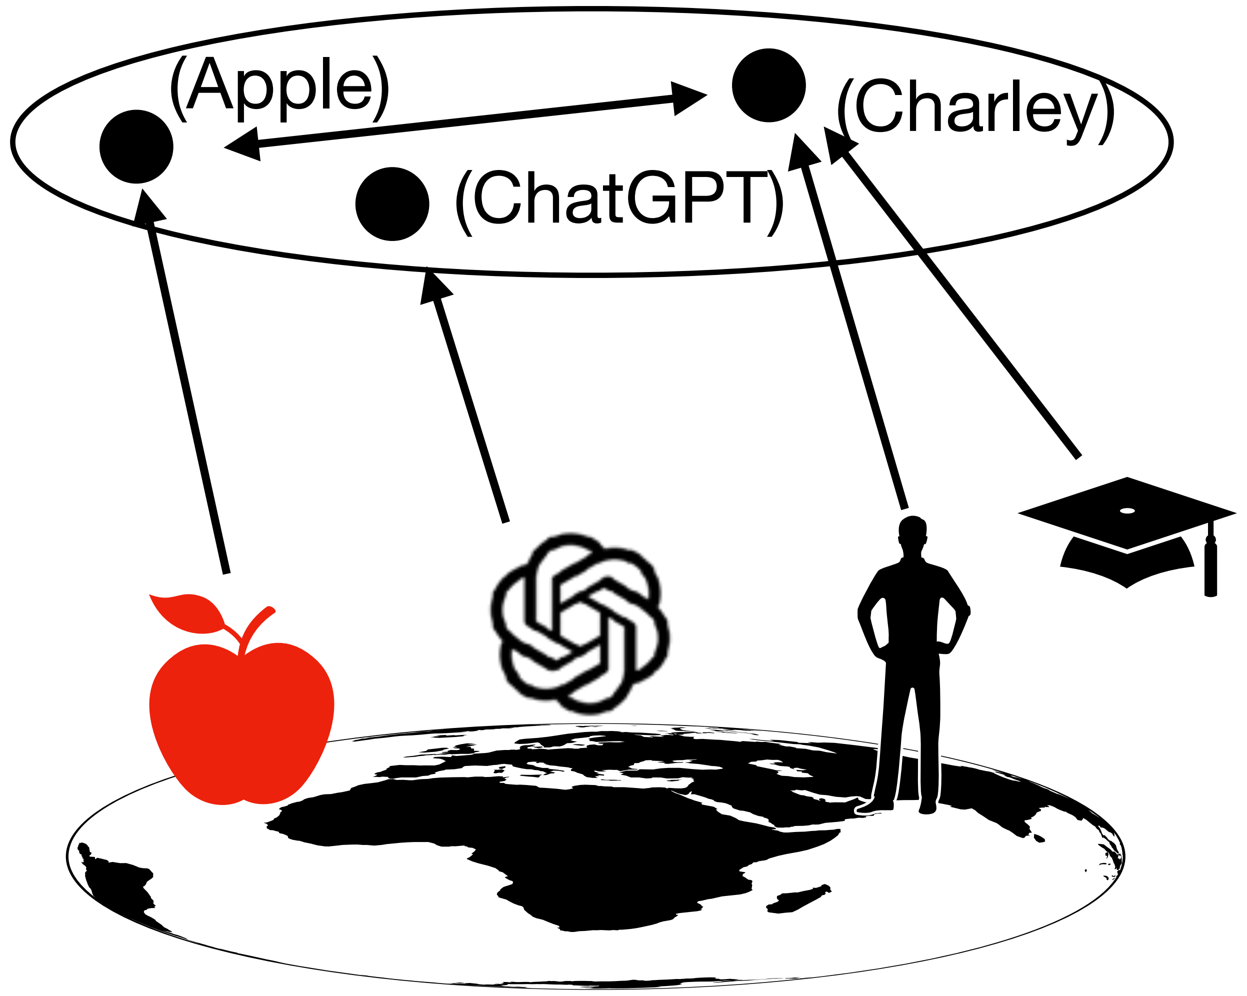

Symbolic AI

- **Physical Symbol System hypothesis**:
  "A physical symbol system has the necessary and sufficient means for general intelligent action - Allen Newell and Herbert Simon (1976)"
  - **Symbols** can represent things in the world
    - e.g., (Apple), (ChatGPT), (Charley), etc…
  - **Relations** can be i) predicates that describes a symbol or ii) verbs describing how symbols interact with other symbols
    - i) red(Apple), unreliable(ChatGPT), instructor(Charley)
    - ii) eat(Charley, Apple), generatePicture(ChatGPT, Apple)
- By populating a **knowledge base** with symbols and relations, we can use a program to find new propositions (*inference*)
  - General Problem Solver (Simon, Shaw, & Newell, 1957)
  - Expert systems: popularized in the 1980s as the future of AI

Herbert Simon
& Allen Newell

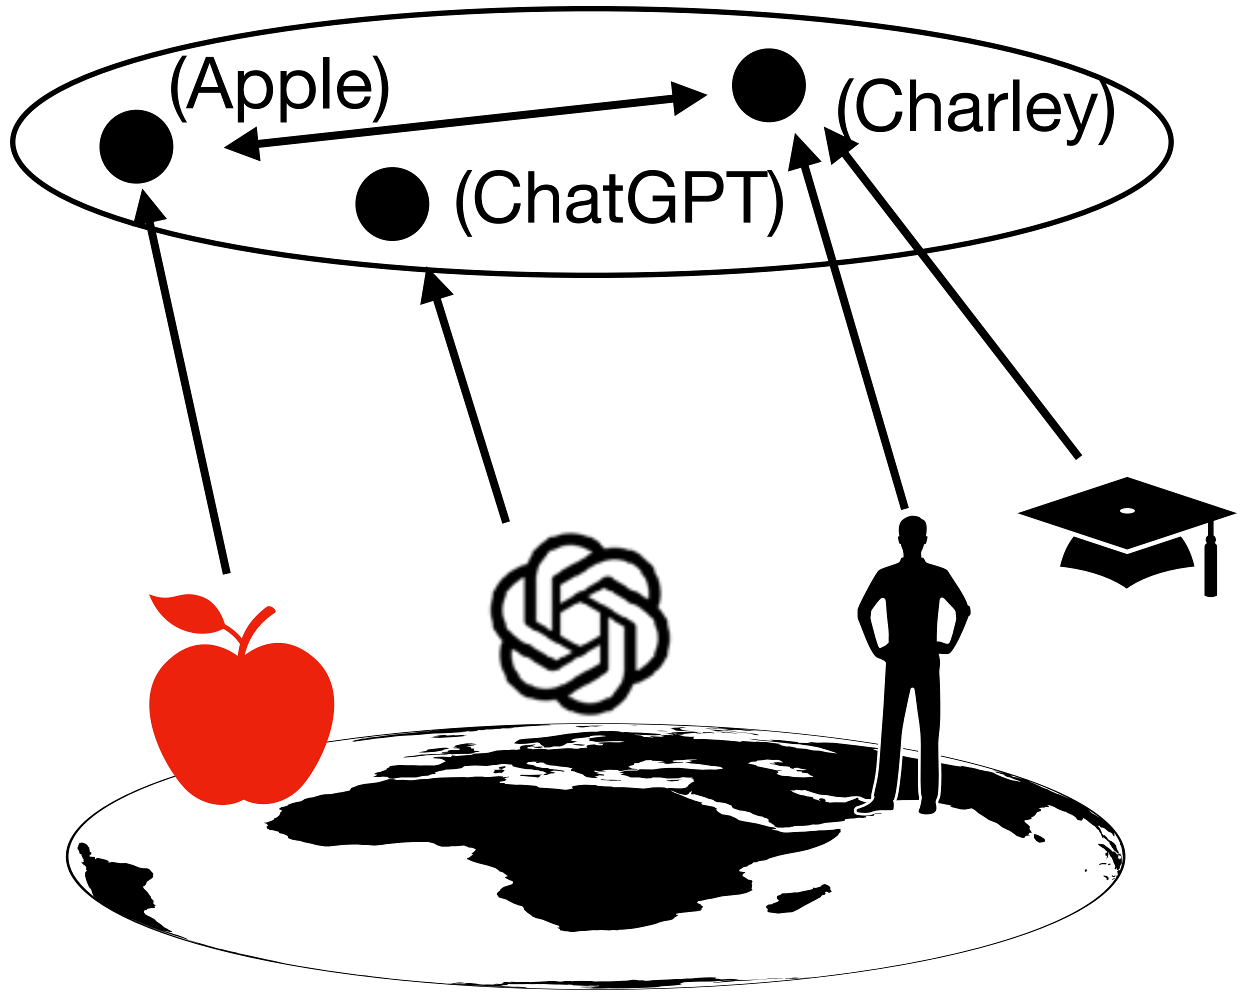

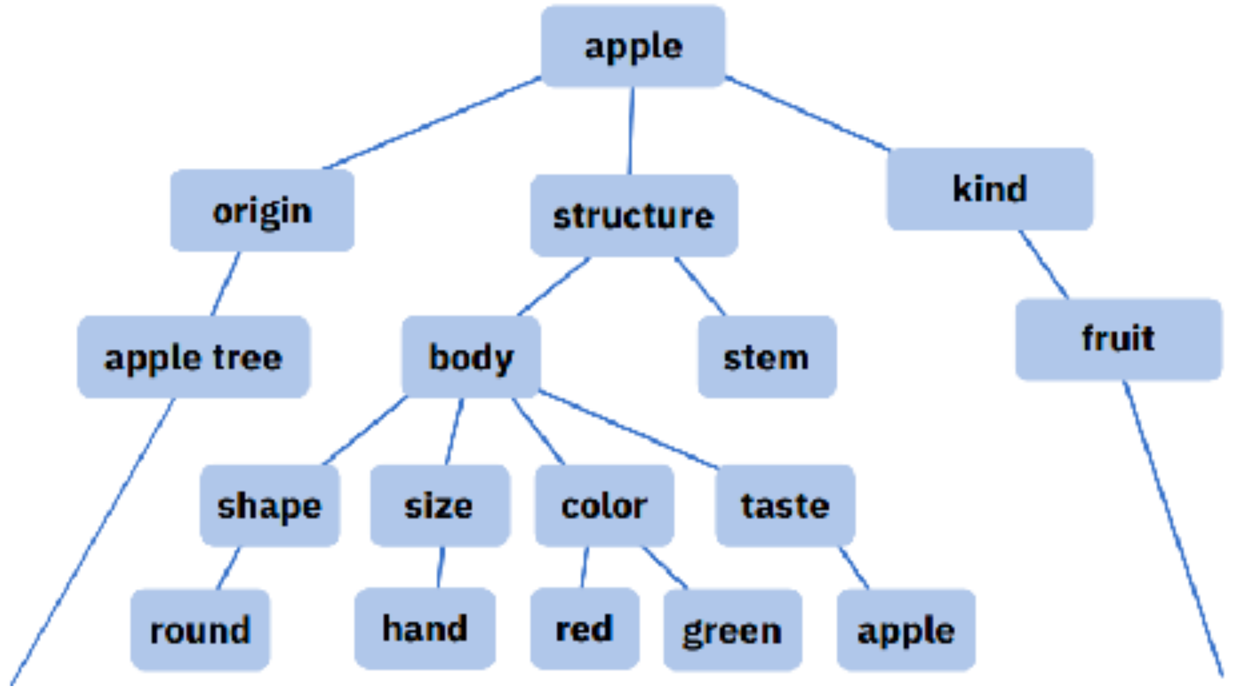

Expert Systems

- The first truly successful forms of AI, widely applied in medicine, finance, and education
- Expert knowledge is codified in the form of facts and logical rules by a *knowledge engineer*
  - If X then Y
  - If Socrates is a man, then Socrates is mortal
- This forms the basis of an *inference engine,* which can apply known rules/facts to generates new facts (adding to the knowledge base) and resolve rule conflicts
- Two modes for solving problems

not on the exam

  - **Forward chaining**: What happens next?
    - Apply rules and facts to arrive at logical conclusions about outcomes
  - **Backwards chaining**: Why did it happen?
    - Starting from a desired outcome, figure out the set of antecedents that can aid in arriving at that outcome

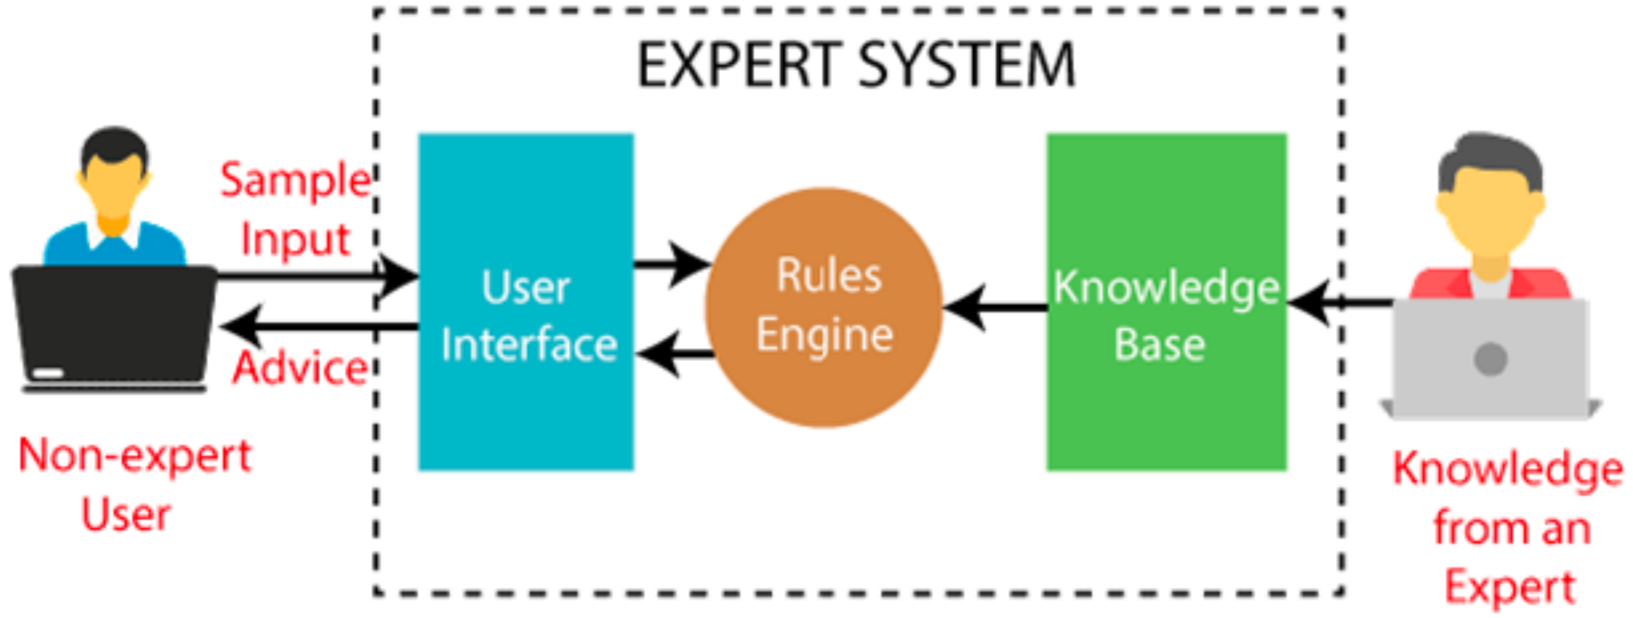

Forward chaining

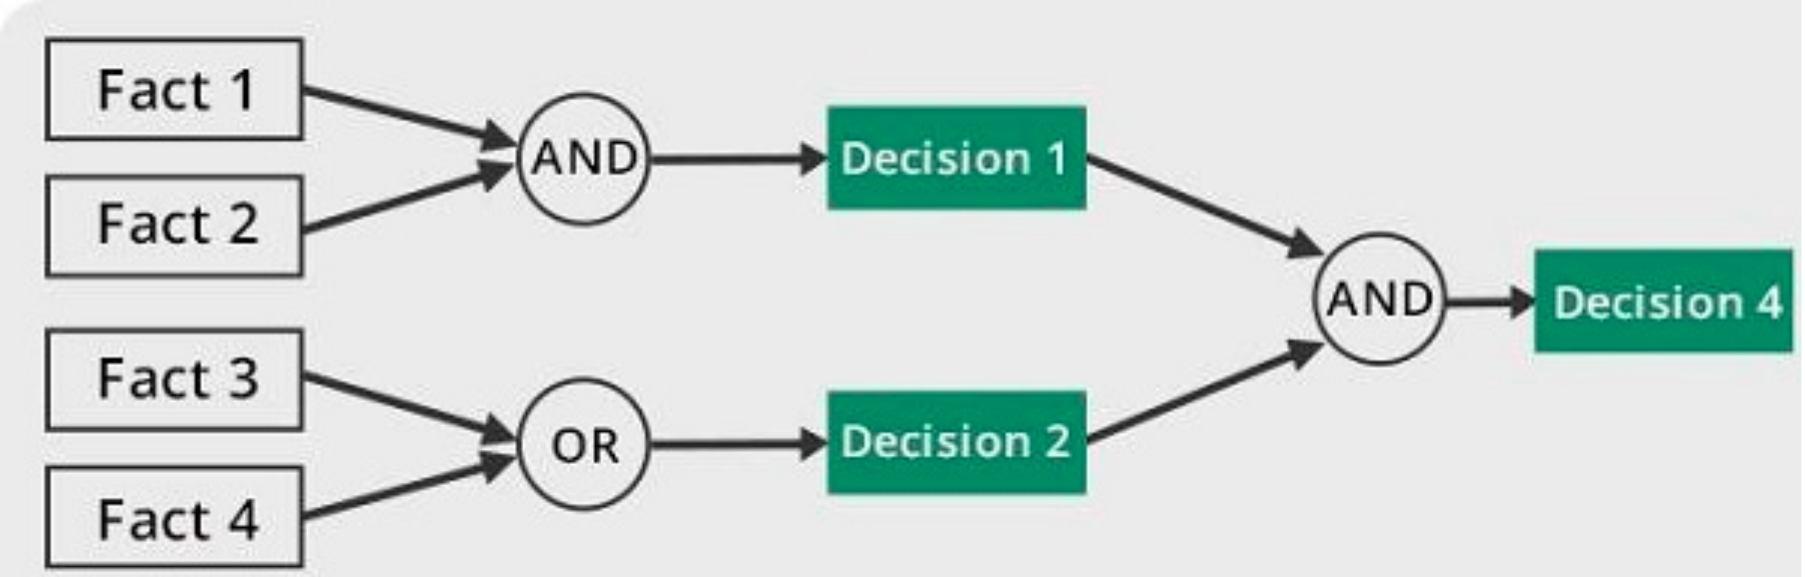

Backward chaining

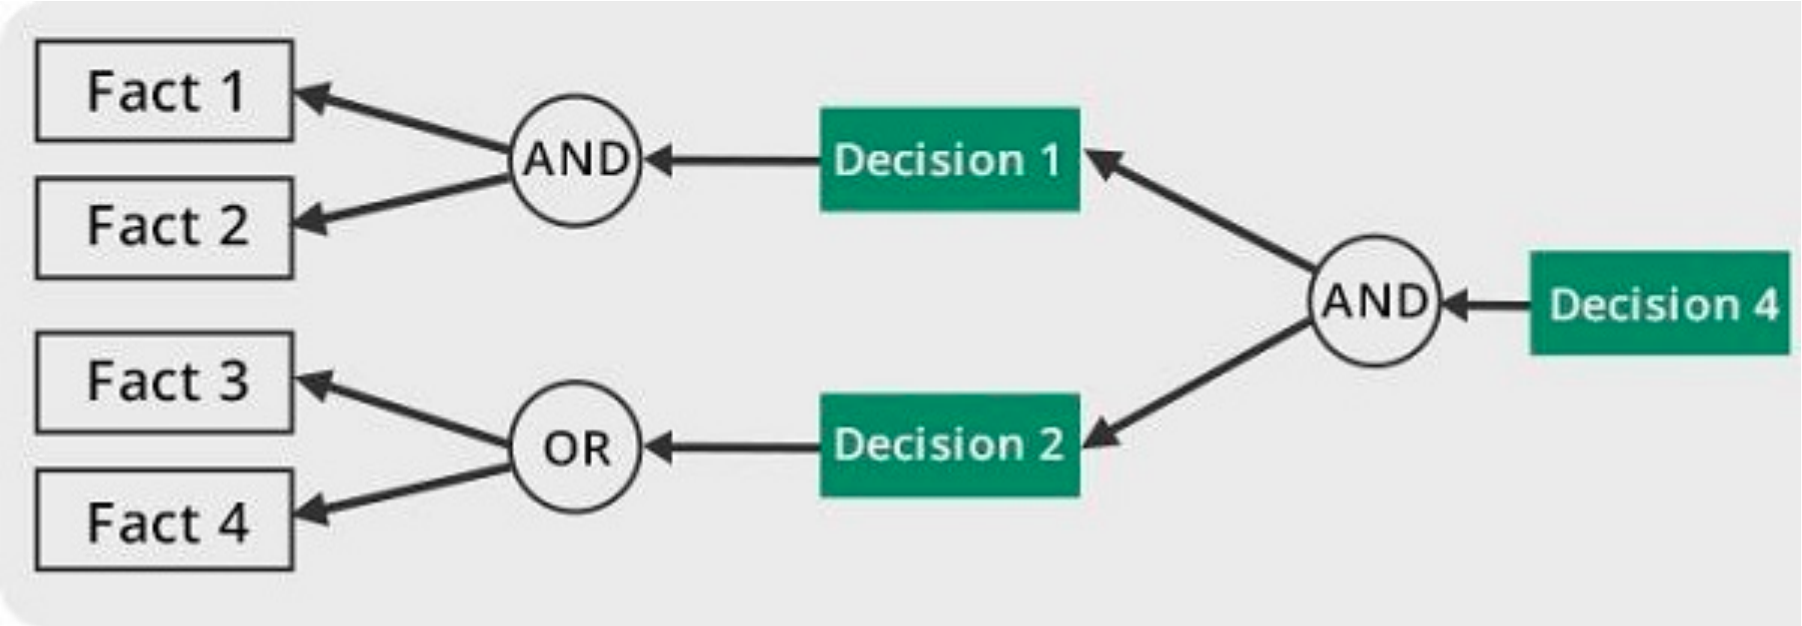

Strengths and Limitations

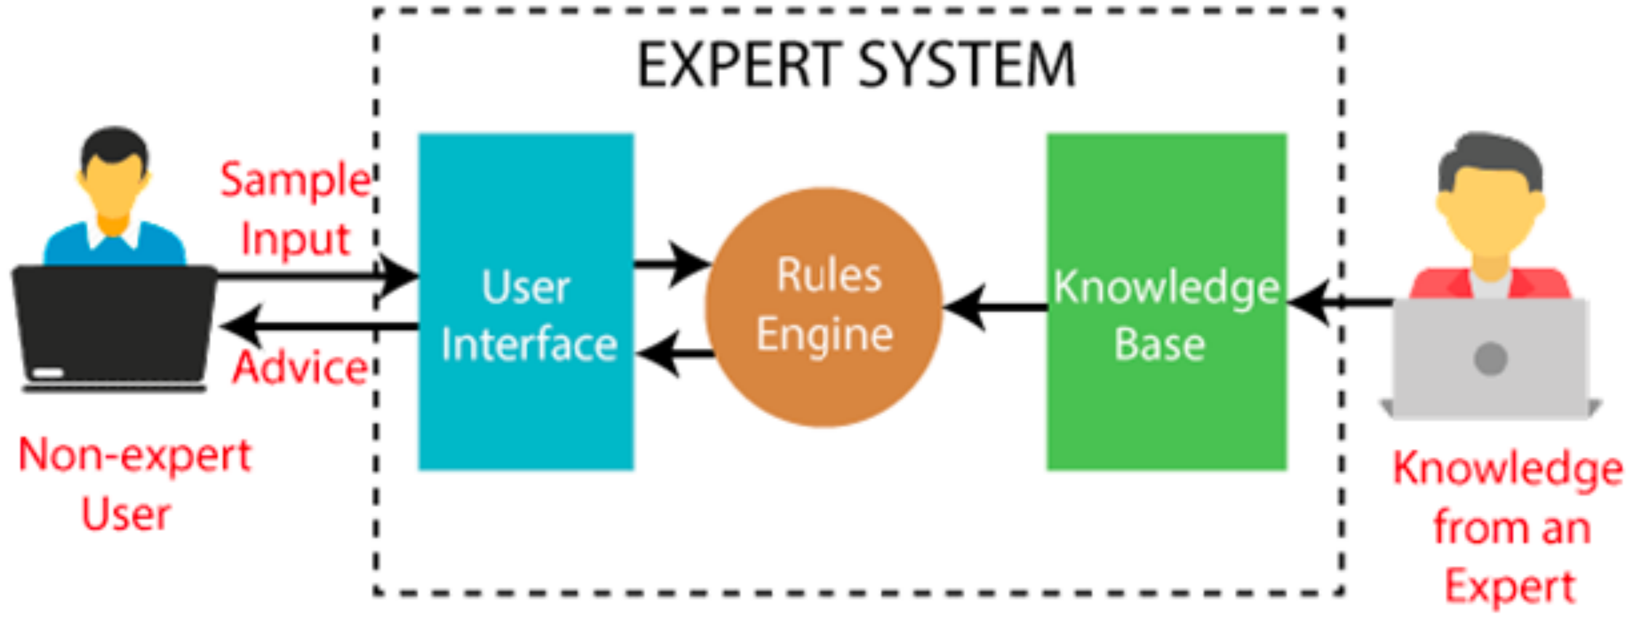

Strengths

- Knowledge is explicit rather than implicit (e.g., neural networks), allowing for interpretability
- Applying rules can be very fast and solutions were generated in real-time
- Rules offer rapid generalization, with a single instance
- Decisions are interpretable by following logic
- No hallucinations!

Limitations

- Cannot learn by itself!
- Require knowledge engineers to codify rules, with high maintenance and development costs
- Limited generalization to new situations, where existing rules don't apply exactly
- If-Then statements cannot capture all relationships without massive scaling problems

Symbolic vs. sub-symbolic AI

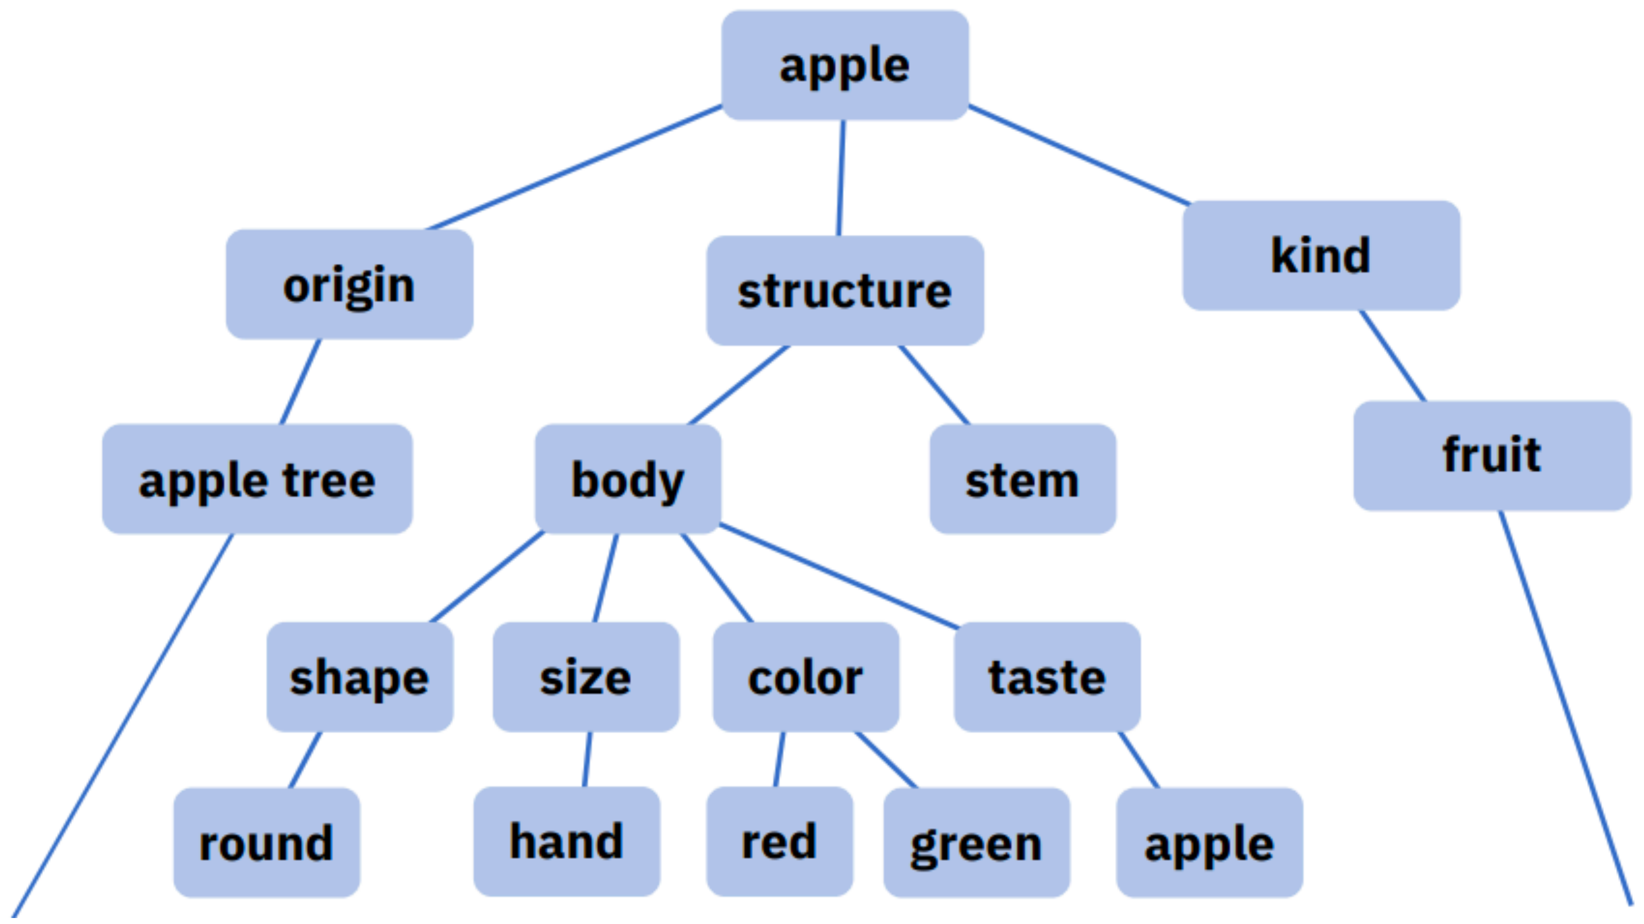

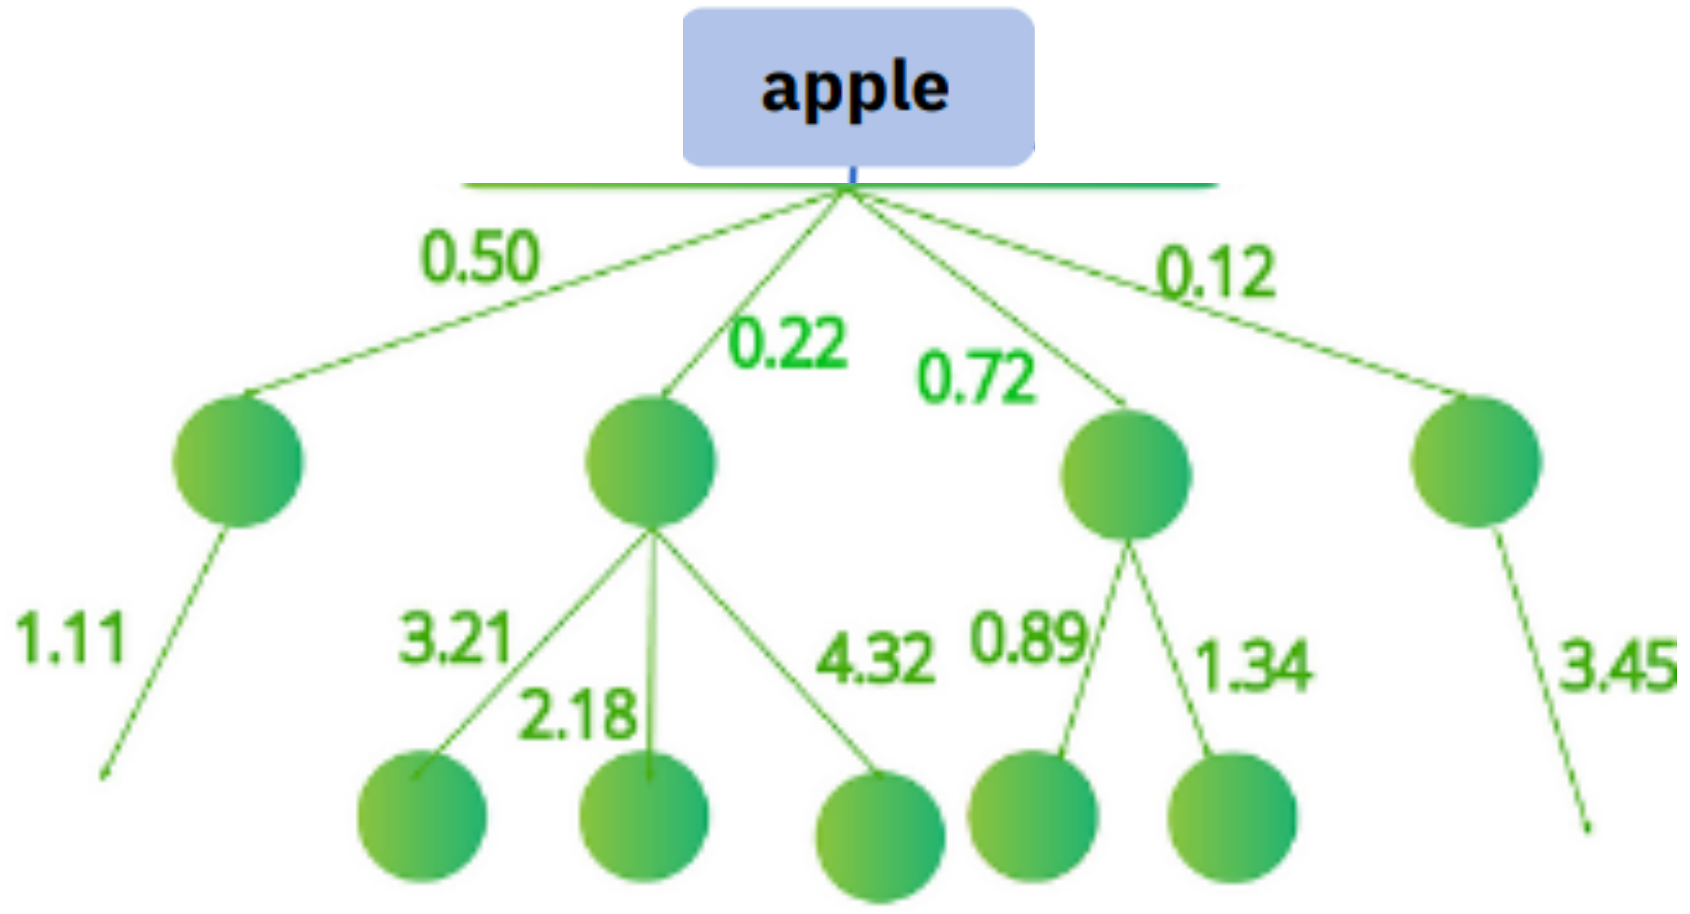

Symbolic AI

- Symbols, rules, and structured representations
- "**Language of thought**" (LoT) hypothesis (Fodor, 1975): concepts/knowledge represented by a language-like system
- Compositionality: symbols and rules can be combined to produce new representations
- Extracting symbolic representations and search over compositional hypothesis spaces is difficult

Sub-symbolic AI

- Representations encoded through connection weights
- No explicit representation of concepts or knowledge, but distributed throughout the network
- Efficiency: knowledge can be implicitly learned by capturing statistical patterns
- Interpretation of representations and behavior is difficult

Symbolic vs. sub-symbolic AI

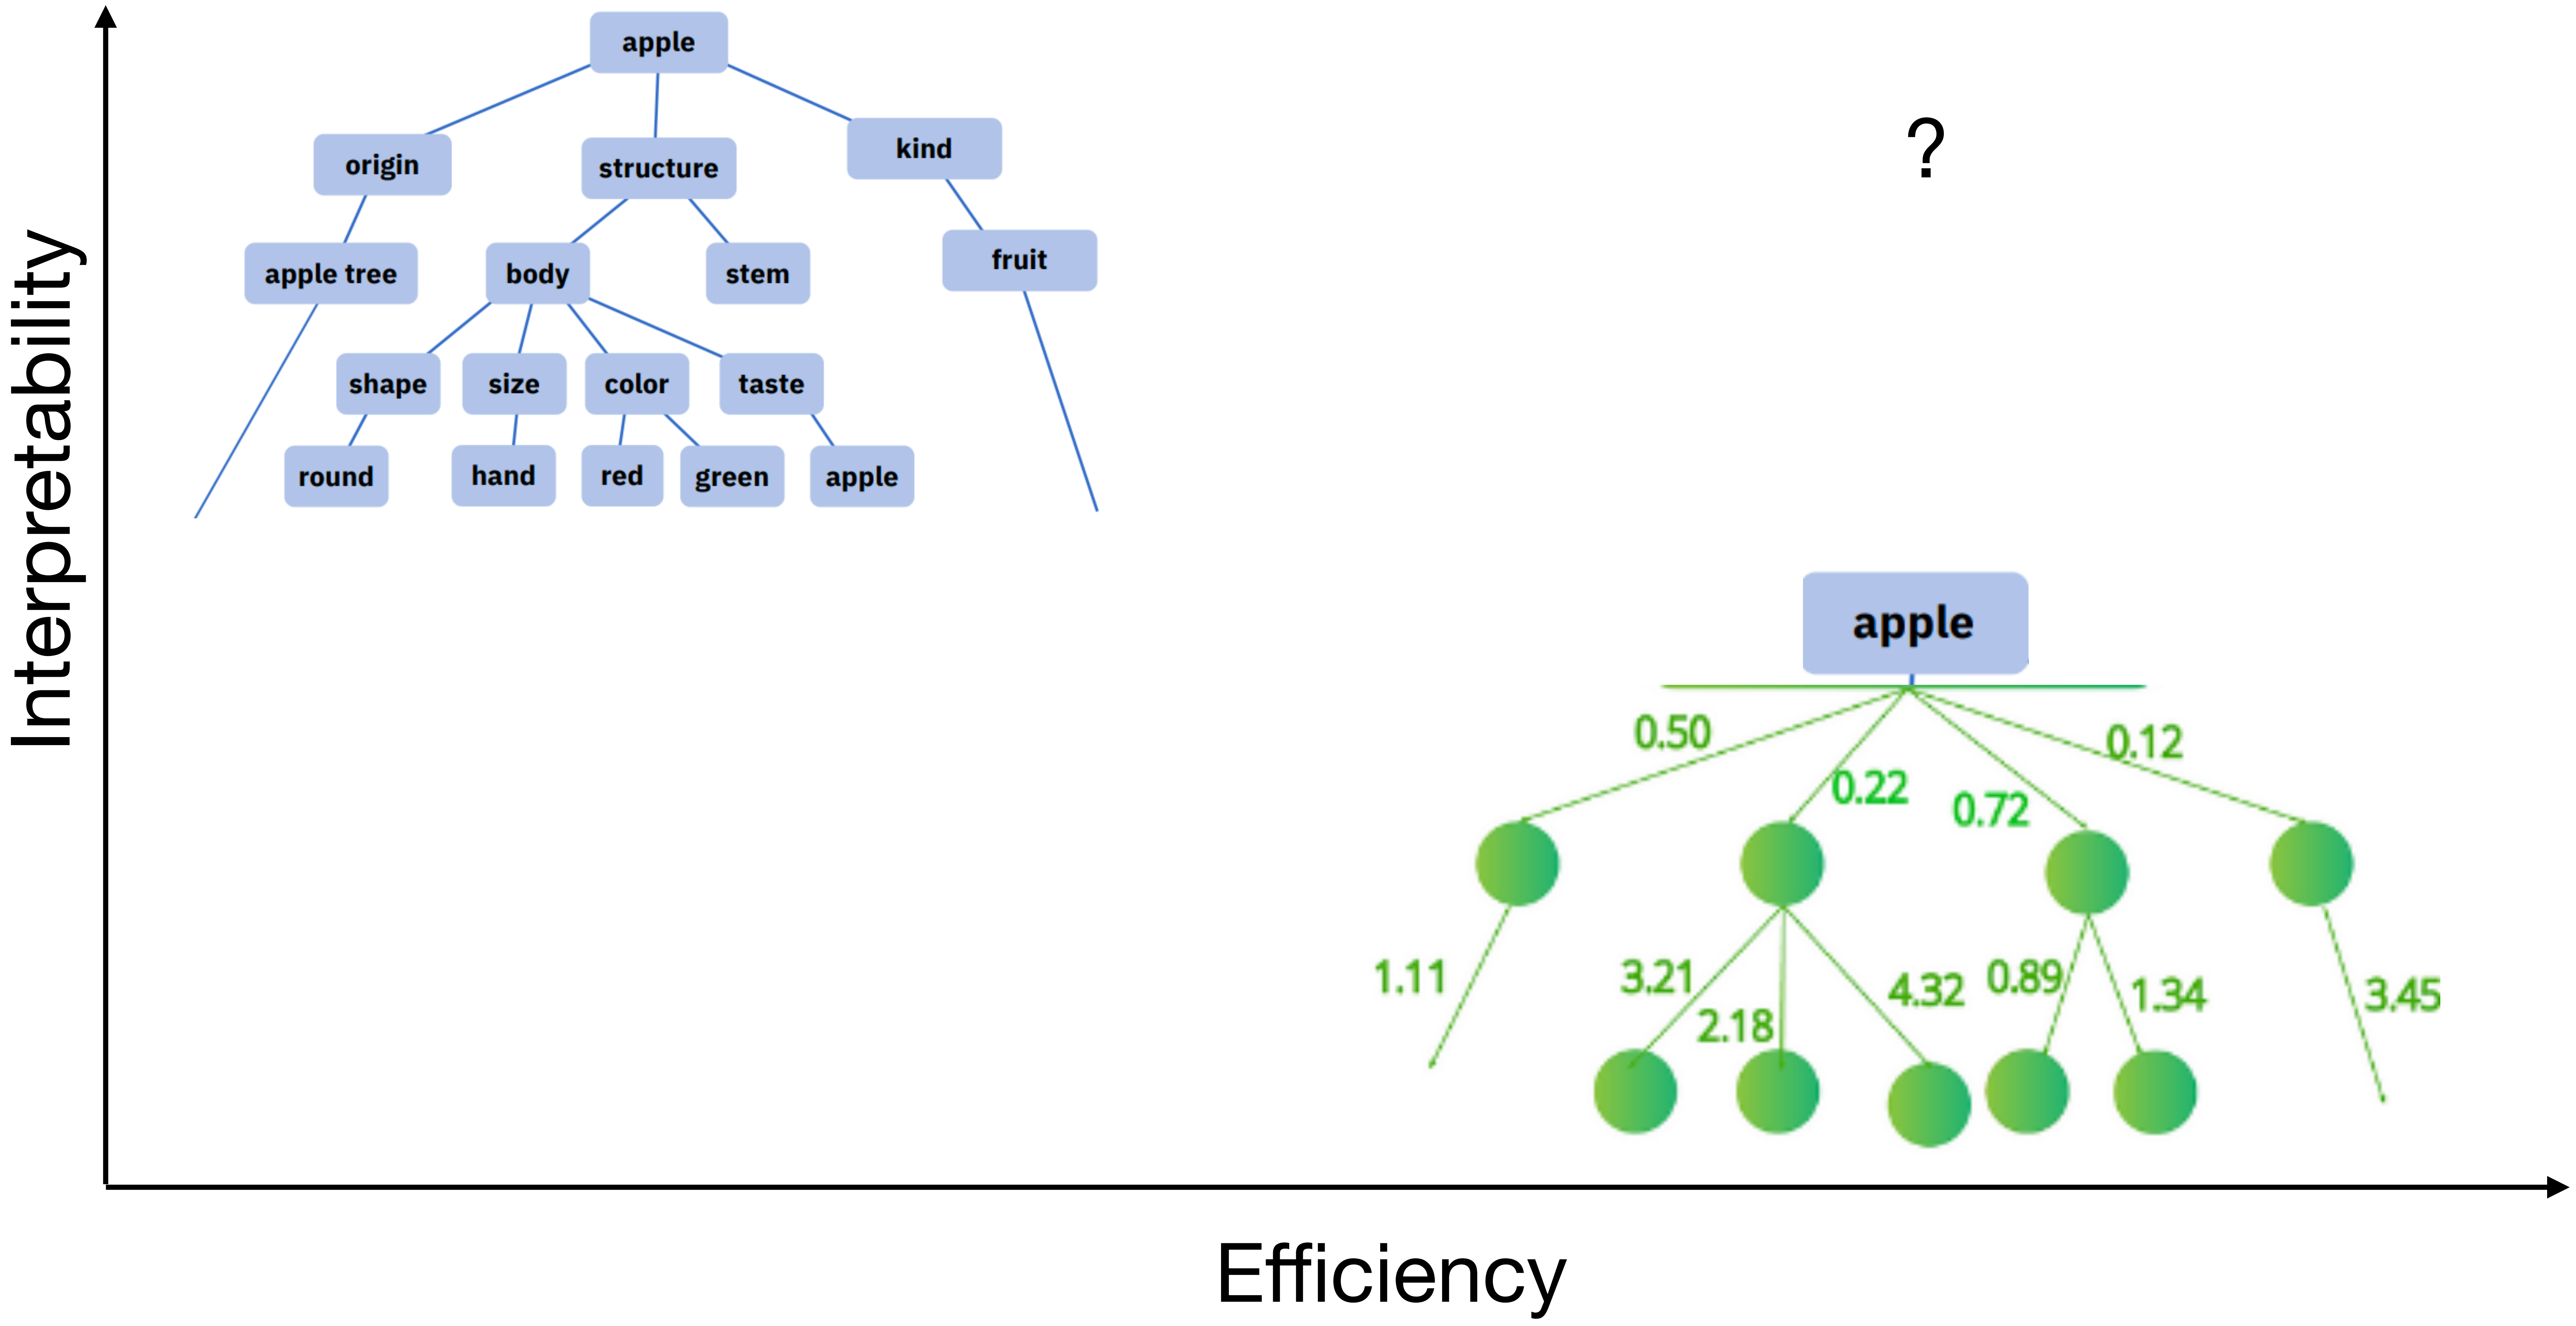

Neurosymbolic AI

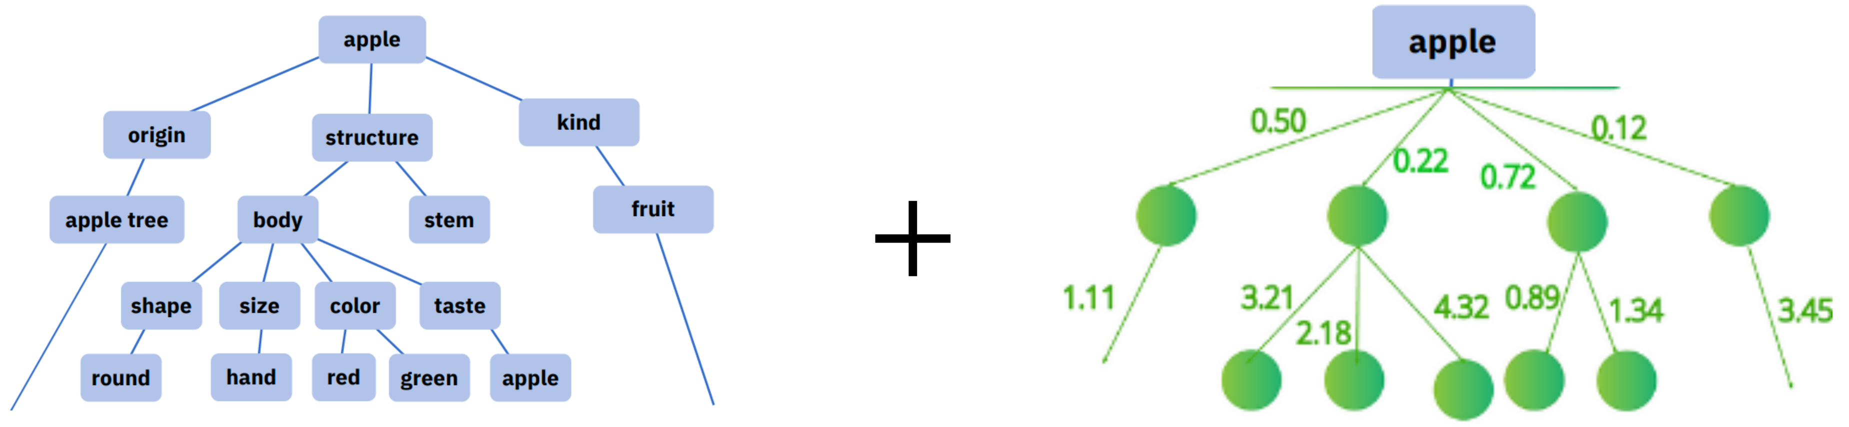

- Neurosymbolic AI aims to combine symbolic and subsymbolic approaches to get the best of both worlds
- Modern AI assistants (e.g., Siri, Google, Alexa) are essentially expert systems with ANN voice recognition and text-to-speech

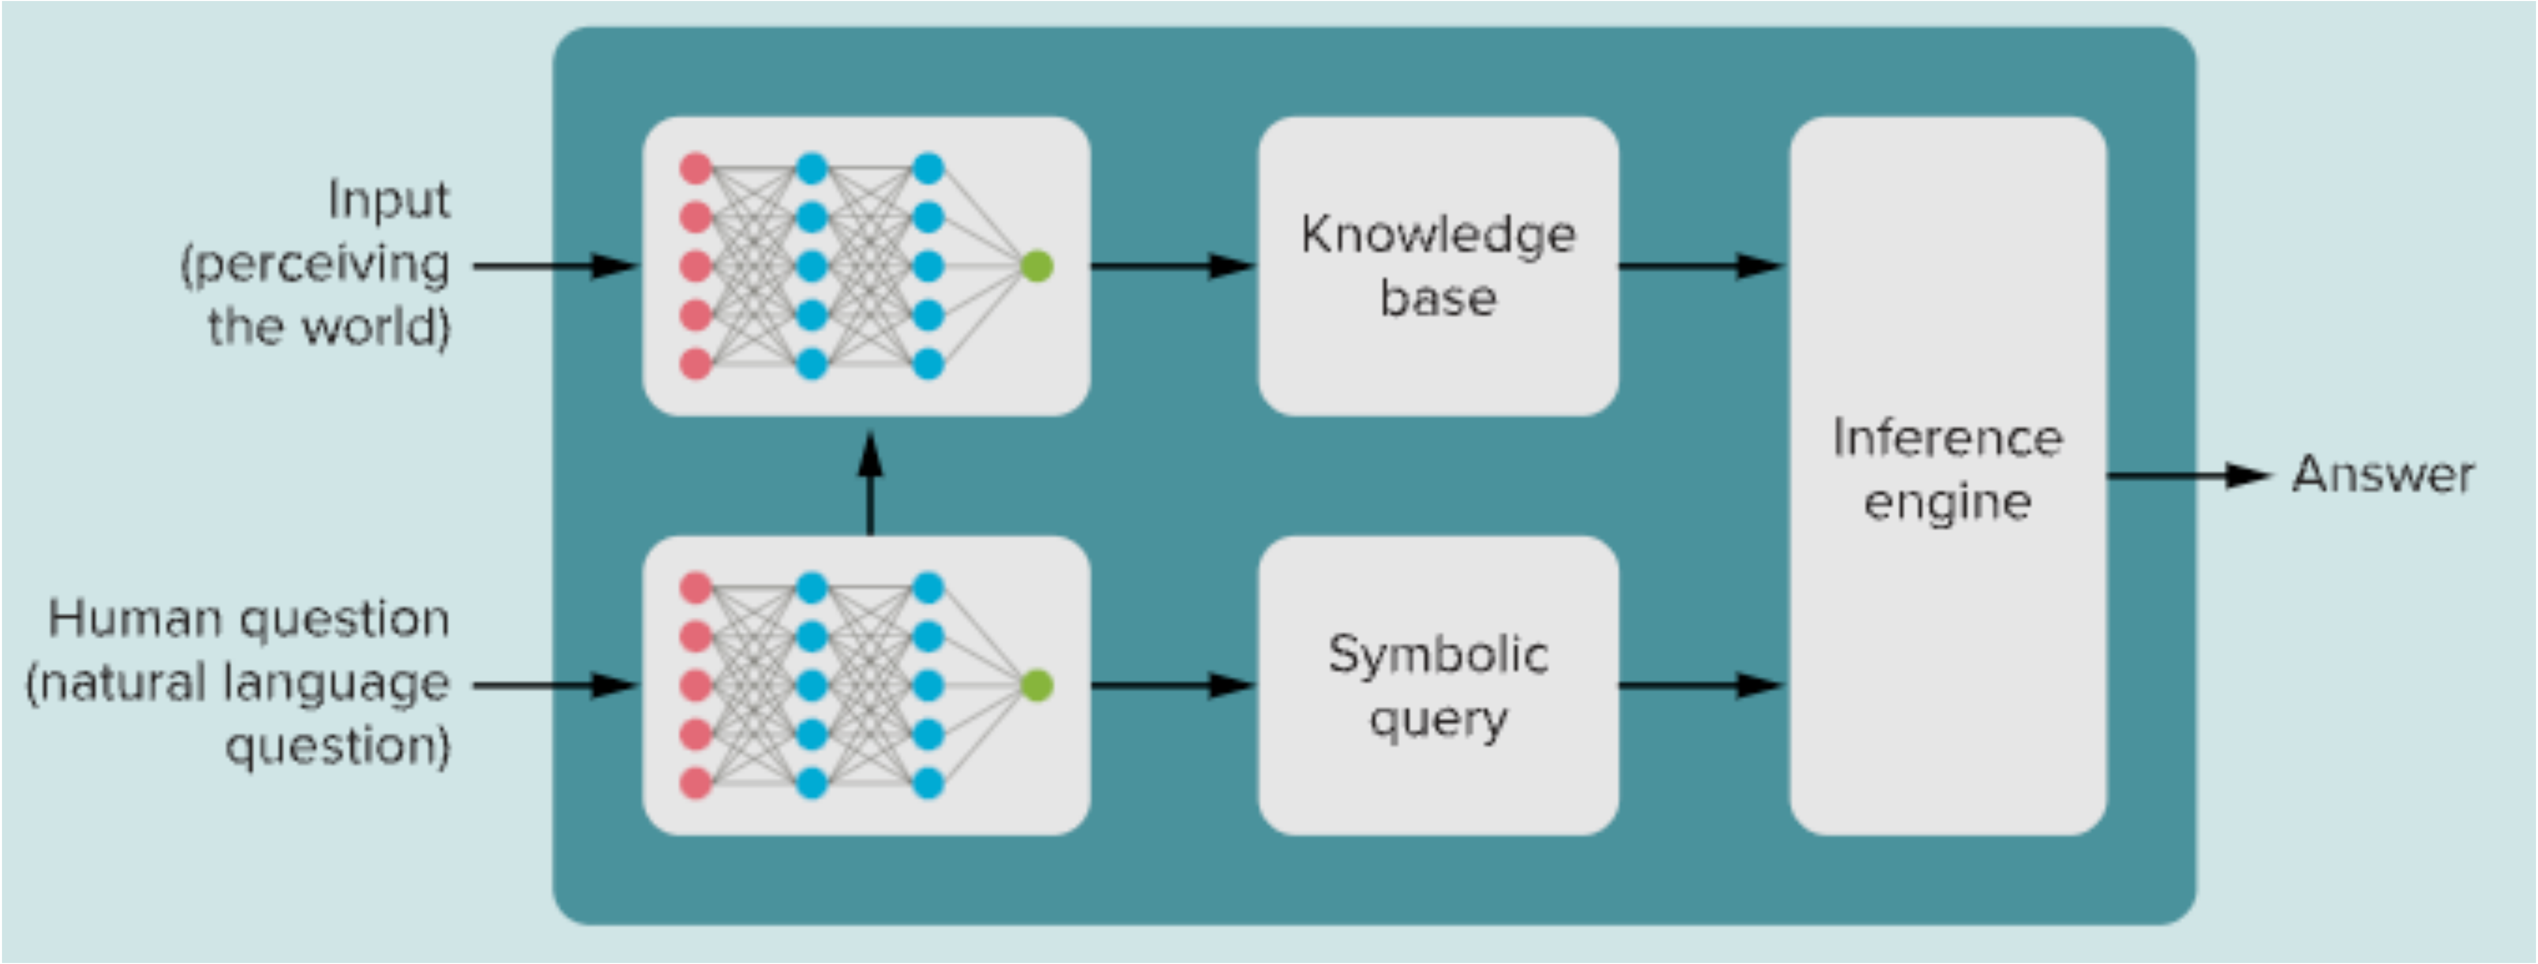

Neurosymbolic AI

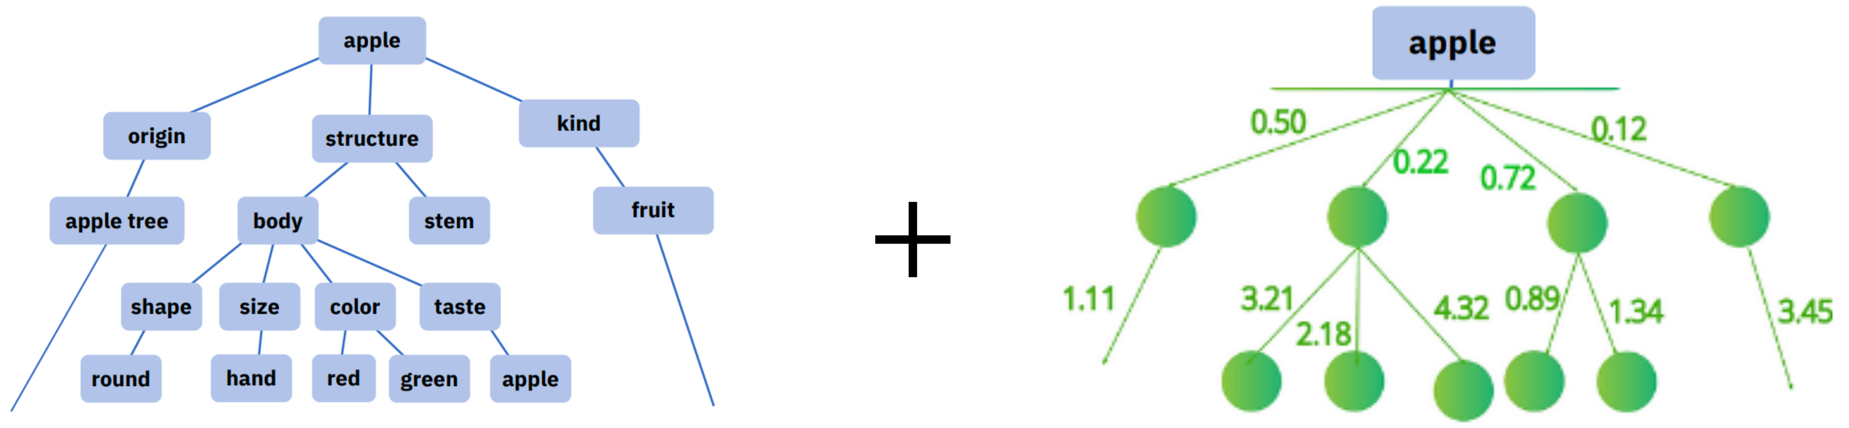

- Neurosymbolic AI aims to combine symbolic and subsymbolic approaches to get the best of both worlds
- Modern AI assistants (e.g., Siri, Google, Alexa) are essentially expert systems with ANN voice recognition and text-to-speech

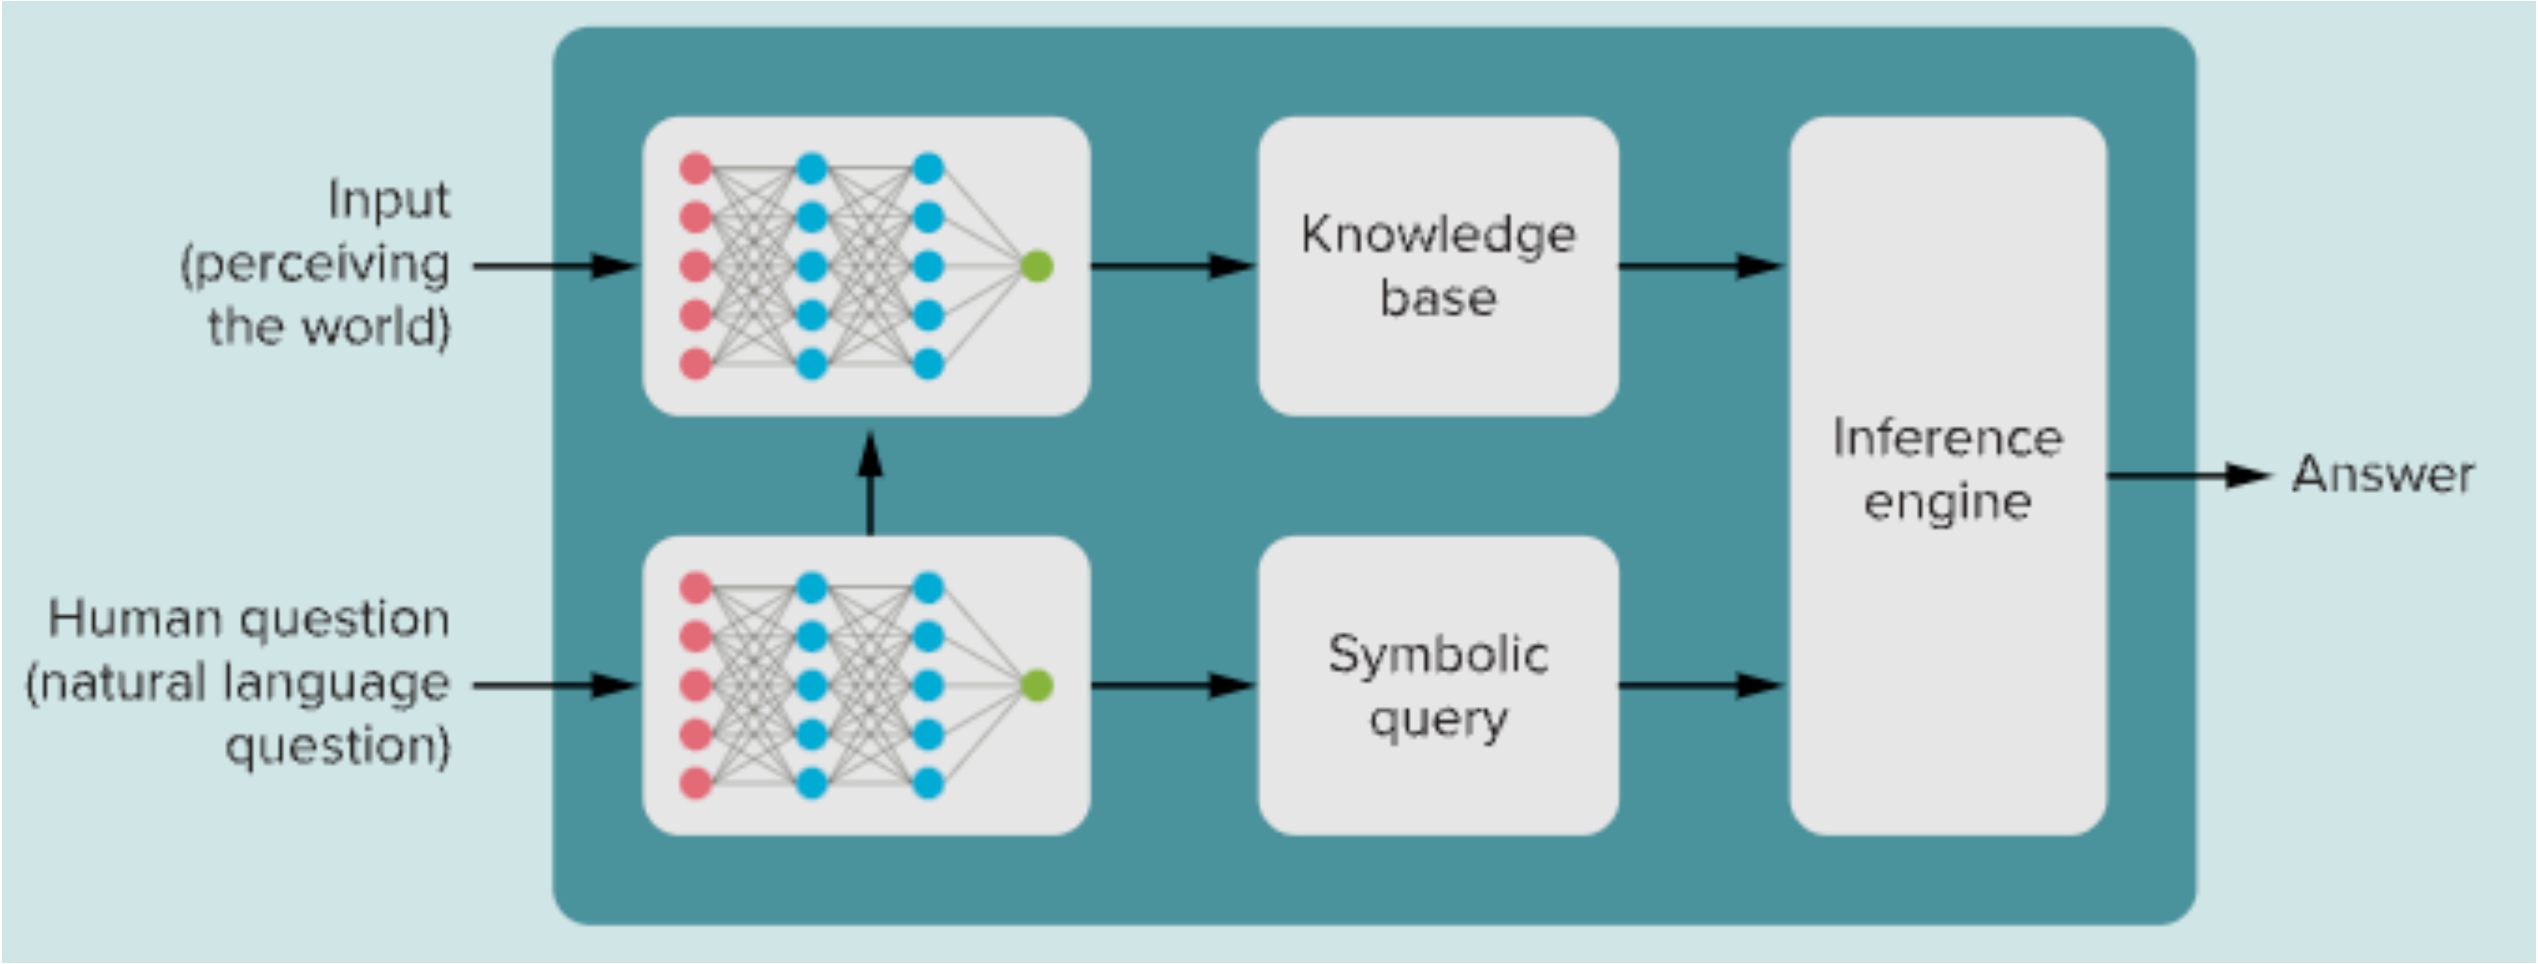

- Current challenges:
  - Learning the knowledge base through data
  - Relating messy real-world data to neat (and limited) symbols/relations in a knowledge base

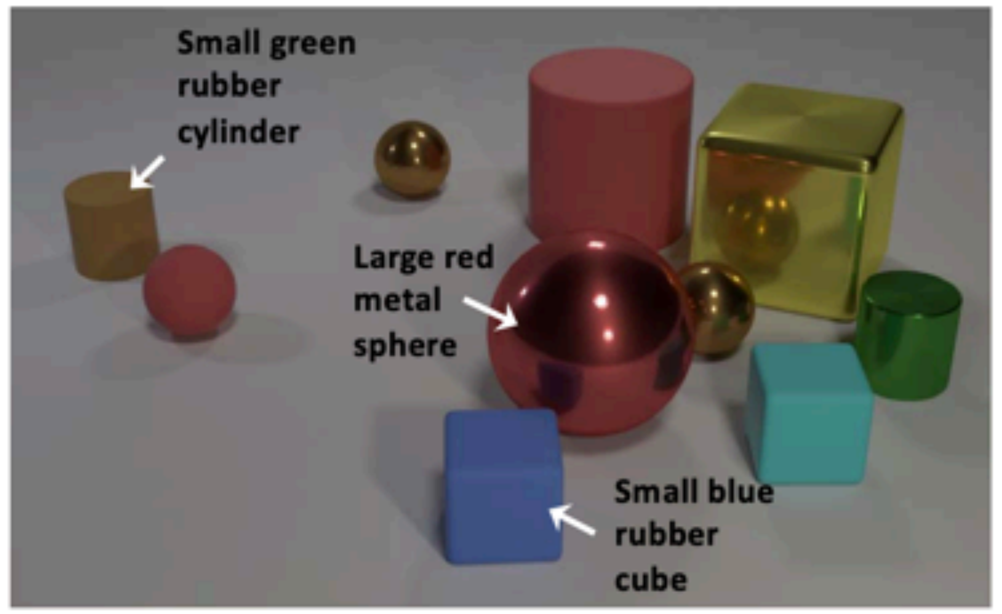

Yi et al., (2018)

Question: *Are there an equal number of large things and metal spheres?*

Program: equal_number(count(filter_size(Scene, Large)), count(filter_material(filter_shape(Scene, Sphere), Metal)))

Answer: *Yes*

How can symbolic knowledge be learned from data?

How can symbolic knowledge be learned from data?

One-shot generalization

Lake et al., (*Science* 2015)

How can symbolic knowledge be learned from data?

One-shot generalization

Lake et al., (*Science* 2015)

Parsing into parts and relations

How can symbolic knowledge be learned from data?

One-shot generalization

Lake et al., (*Science* 2015)

Parsing into parts and relations

Generalization from related concepts

How can symbolic knowledge be learned from data?

One-shot generalization

Lake et al., (*Science* 2015)

Parsing into parts and relations

Generalization from related concepts

Program Induction

the process of inferring ***rules or instructions*** *that generate an observed pattern of data*

How can symbolic knowledge be learned from data?

One-shot generalization

Lake et al., (*Science* 2015)

Parsing into parts and relations

Generalization from related concepts

Program Induction

the process of inferring ***rules or instructions*** *that generate an observed pattern of data*

Data $\mathscr{D}$

How can symbolic knowledge be learned from data?

One-shot generalization

Lake et al., (*Science* 2015)

Parsing into parts and relations

Generalization from related concepts

Program Induction

the process of inferring ***rules or instructions*** *that generate an observed pattern of data*

Data $\mathscr{D}$

Program π

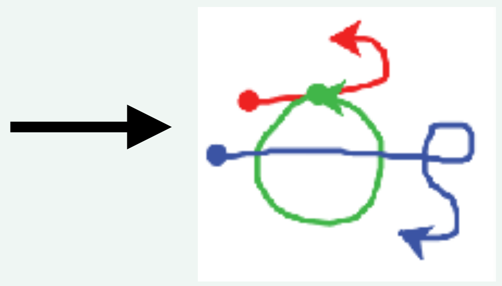

How can symbolic knowledge be learned from data?

One-shot generalization

Lake et al., (*Science* 2015)

Parsing into parts and relations

Generalization from related concepts

Program Induction

the process of inferring ***rules or instructions*** *that generate an observed pattern of data*

Data $\mathscr{D}$ Program π

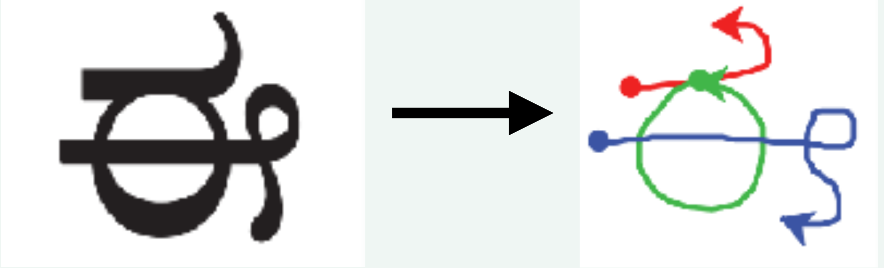

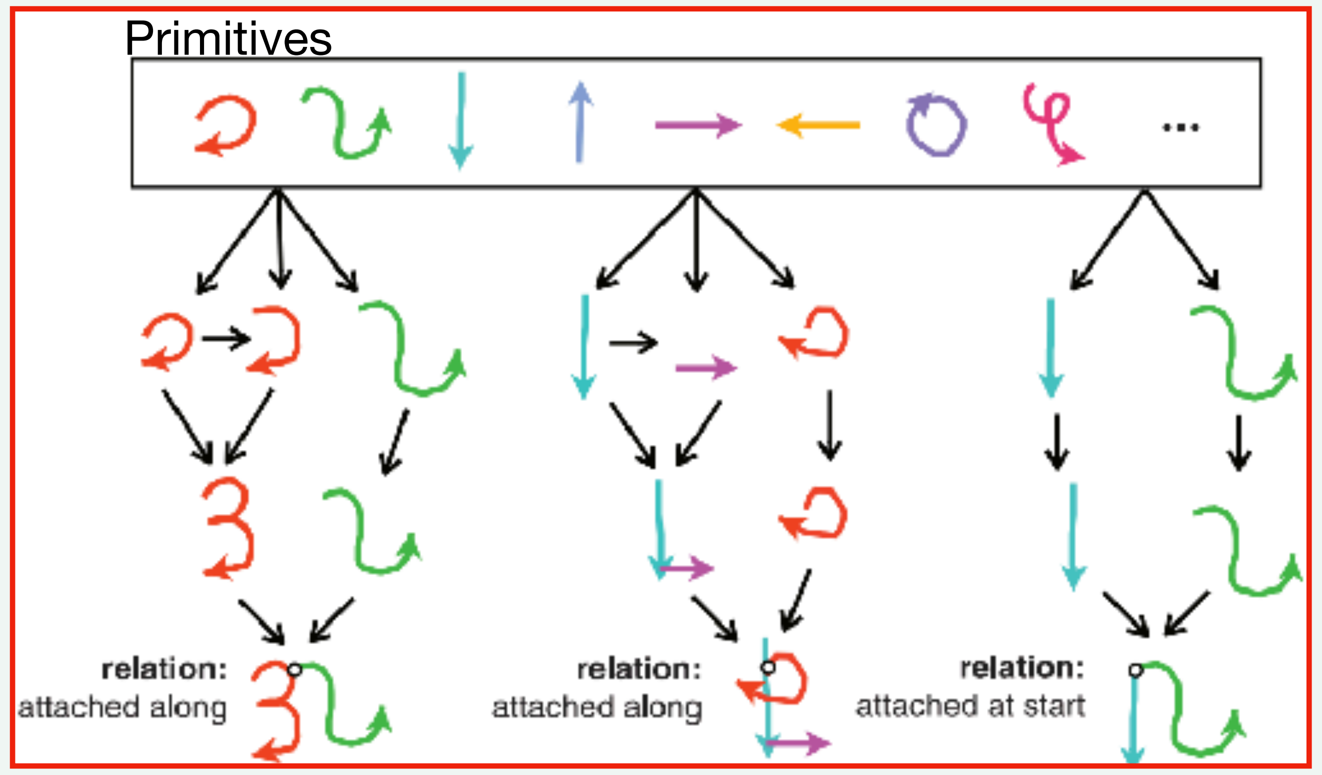

How can symbolic knowledge be learned from data?

One-shot generalization

Lake et al., (*Science* 2015)

Parsing into parts and relations

Generalization from related concepts

Program Induction

$$P(\pi \,|\, \mathscr{D}) \propto P(\mathscr{D} \,|\, \pi)P(\pi)$$

*the process of inferring **rules or instructions** that generate an observed pattern of data*

Data $\mathscr{D}$ Program π

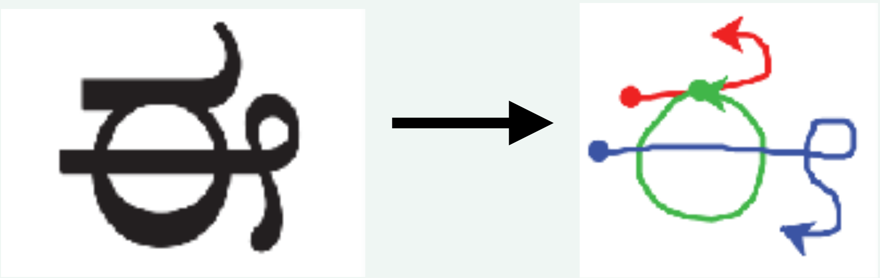

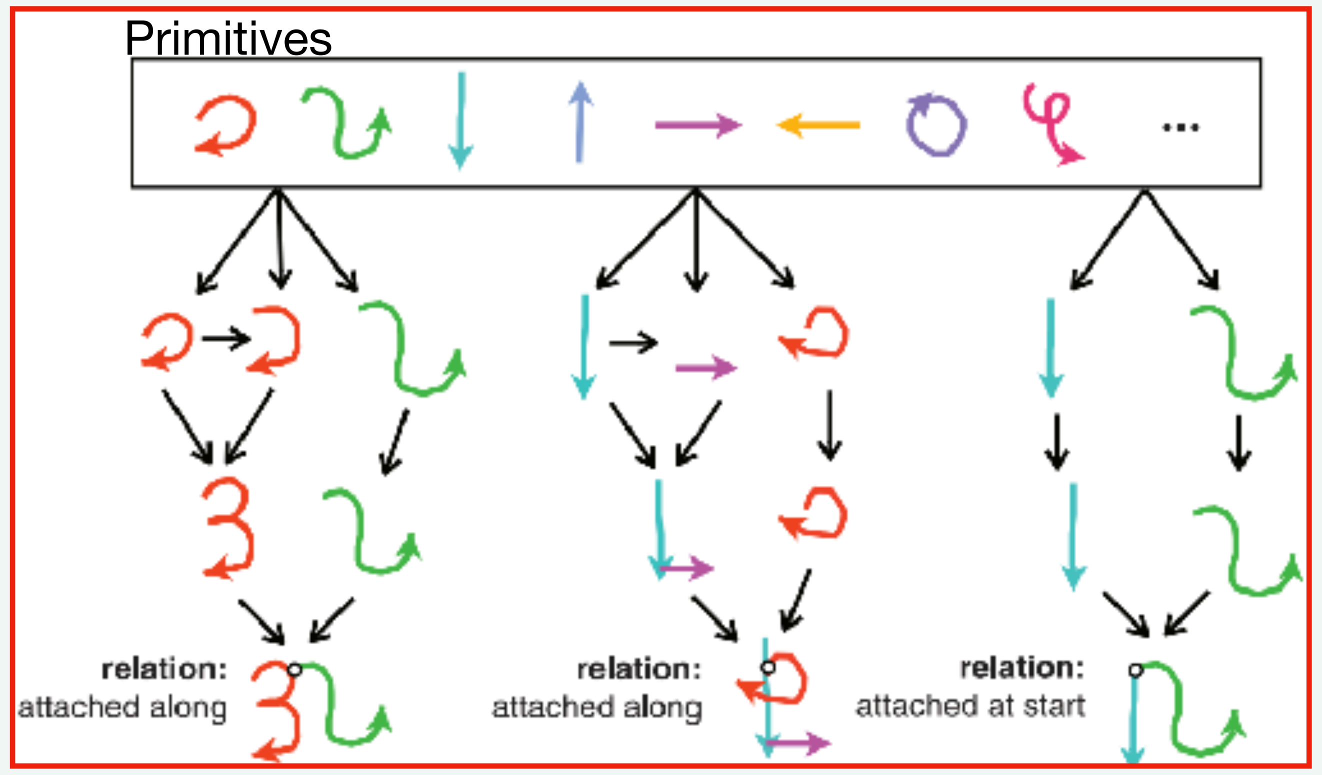

How can symbolic knowledge be learned from data?

One-shot generalization

Lake et al., (*Science* 2015)

Parsing into parts and relations

Generalization from related concepts

Program Induction

$$P(\pi \,|\, \mathscr{D}) \propto P(\mathscr{D} \,|\, \pi)P(\pi)$$

the process of inferring ***rules or instructions*** *that generate an observed pattern of data*

Data $\mathscr{D}$

Program π

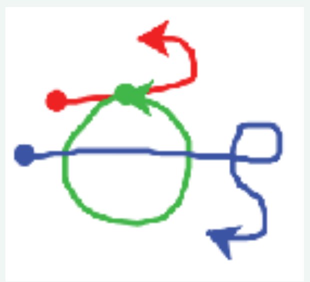

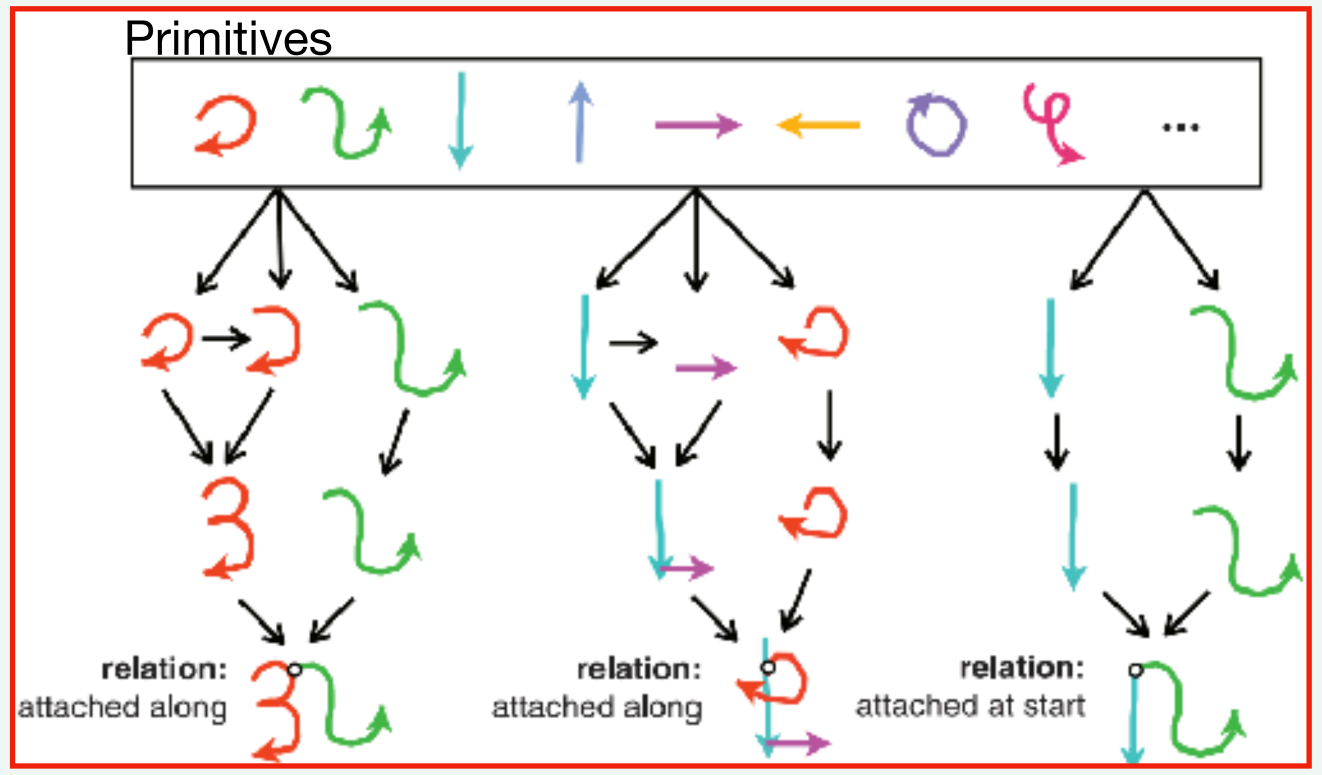


```
procedure GENERATETYPE
  κ ← P(κ)                          ▷ Sample number of parts
  for i = 1 ... κ do
    n_i ← P(n_i|κ)                  ▷ Sample number of sub-parts
    for j = 1 ... n_i do
      s_ij ← P(s_ij|s_i(j−1))       ▷ Sample sub-part sequence
    end for
    R_i ← P(R_i|S_1, ..., S_i−1)    ▷ Sample relation
  end for
  ψ ← {κ, R, S}
  return @GENERATETOKEN(ψ)          ▷ Return program
```

How can symbolic knowledge be learned from data?

One-shot generalization

Parsing into parts and relations

Generalization from related concepts

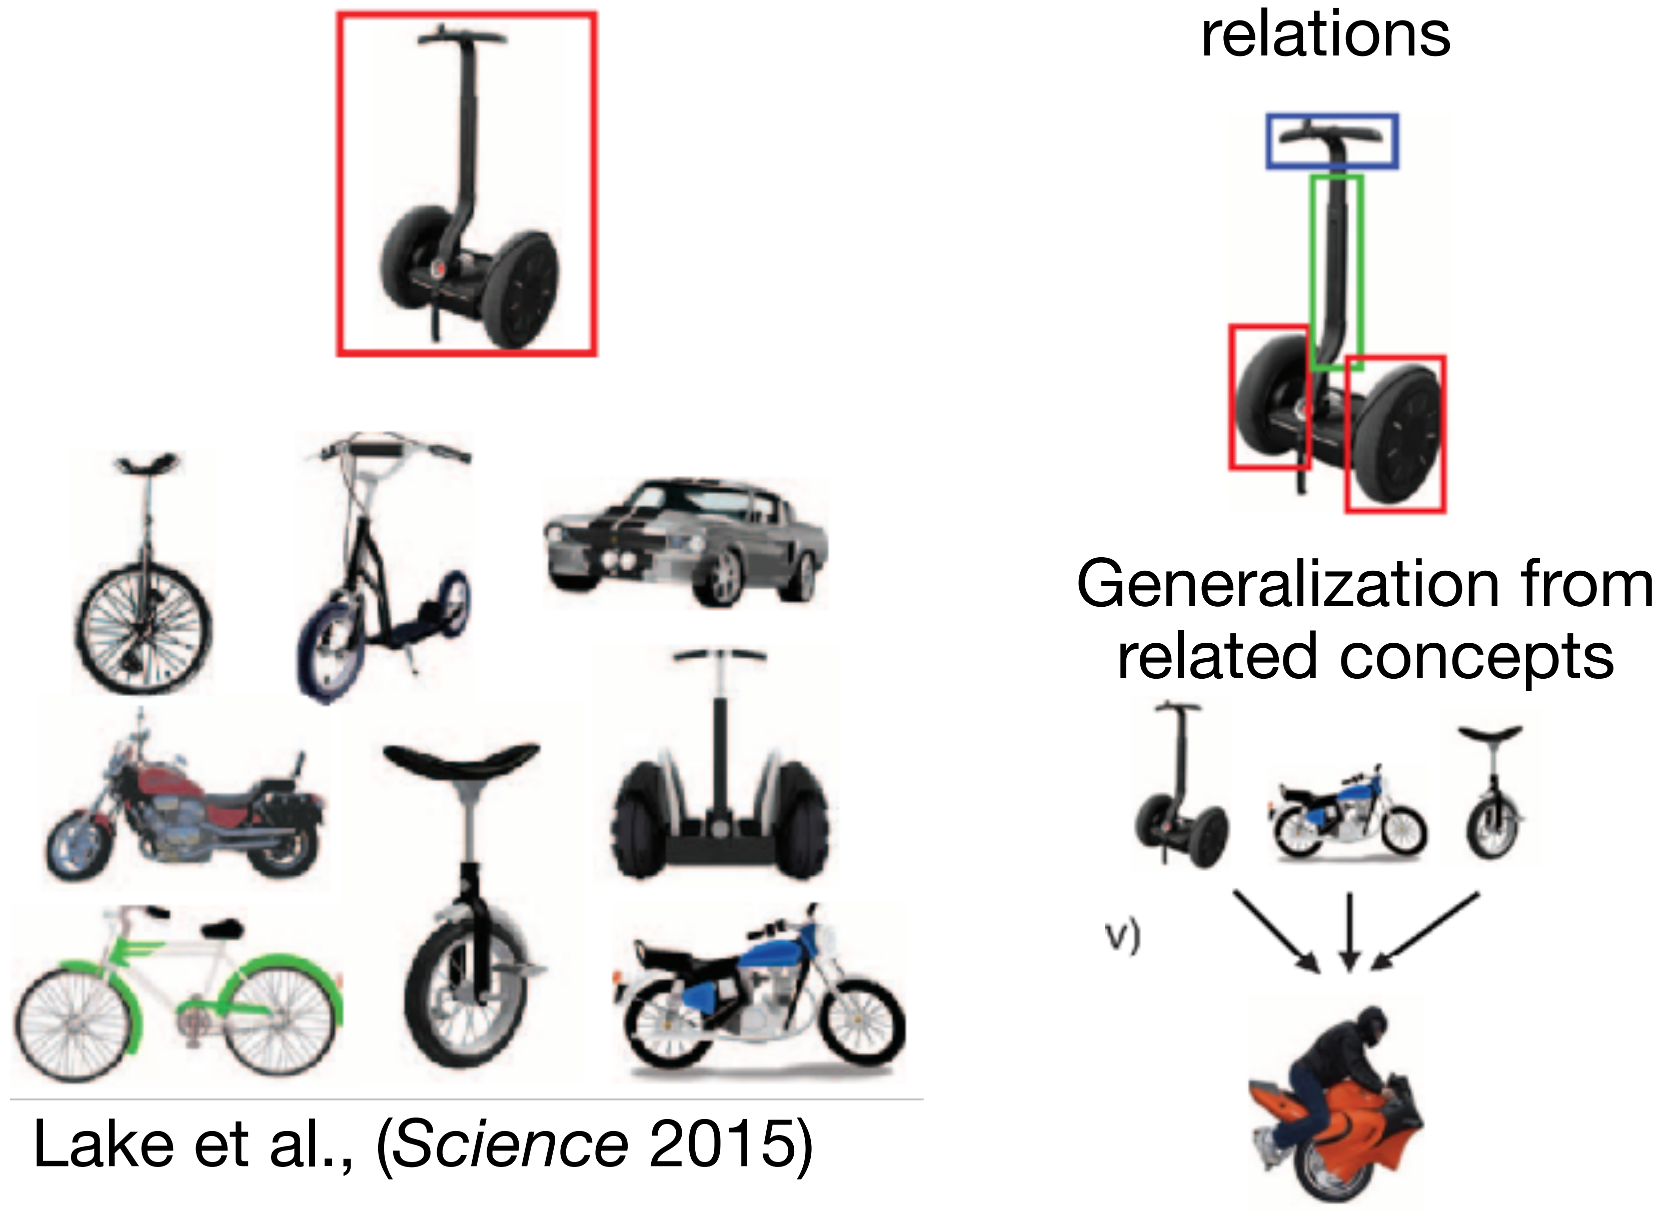

Lake et al., (*Science* 2015)

Program Induction

$$P(\pi \mid \mathscr{D}) \propto P(\mathscr{D} \mid \pi) P(\pi)$$

*the process of inferring **rules or instructions** that generate an observed pattern of data*

Data $\mathscr{D}$ Program π

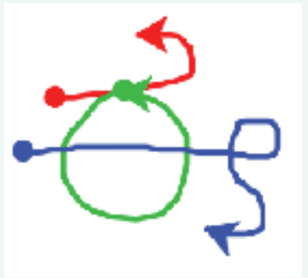

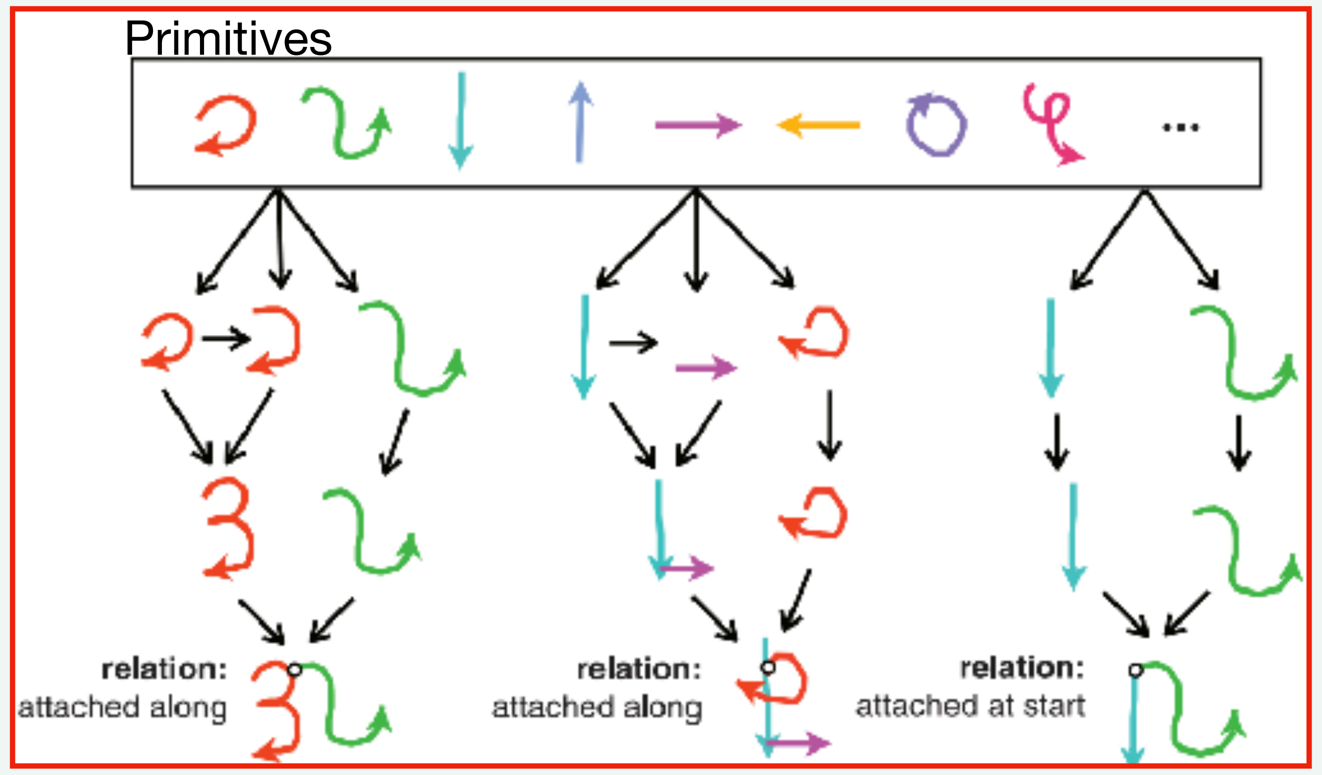


```
procedure GENERATETYPE
  κ ← P(κ)                              ▷ Sample number of parts
  for i = 1 ... κ do
    n_i ← P(n_i | κ)                    ▷ Sample number of sub-parts
    for j = 1 ... n_i do
      s_ij ← P(s_ij | s_i(j−1))          ▷ Sample sub-part sequence
    end for
    R_i ← P(R_i | S_1, ..., S_{i−1})     ▷ Sample relation
  end for
  ψ ← {κ, R, S}
  return @GENERATETOKEN(ψ)              ▷ Return program
```

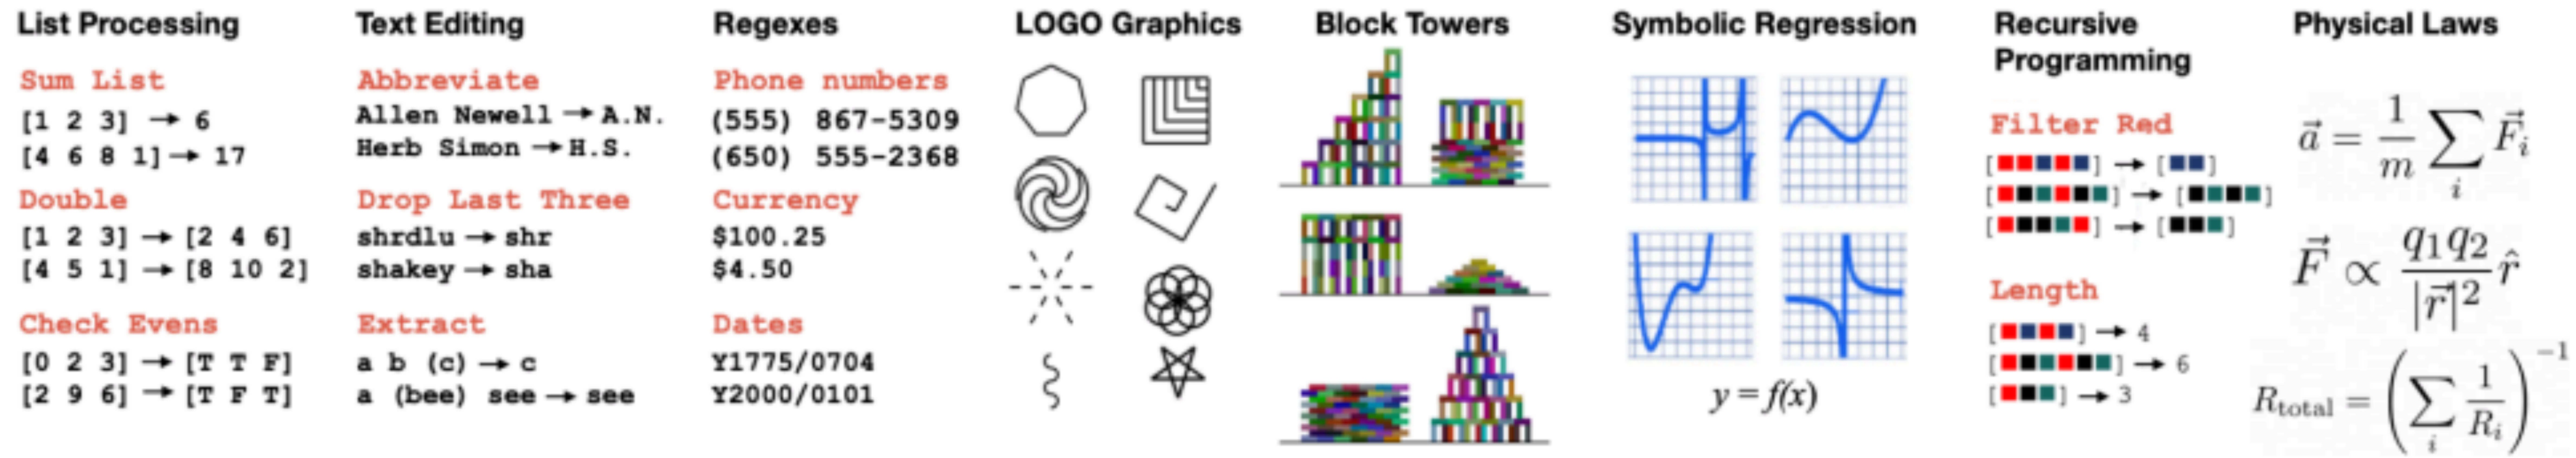

Dreamcoder (Ellis et al., 2020)

How can symbolic knowledge be learned from data?

One-shot generalization

Lake et al., (*Science* 2015)

Parsing into parts and relations

Generalization from related concepts

Program Induction

$$P(\pi \mid \mathscr{D}) \propto P(\mathscr{D} \mid \pi)P(\pi)$$

the process of inferring ***rules or instructions*** *that generate an observed pattern of data*

Data $\mathscr{D}$

Program π

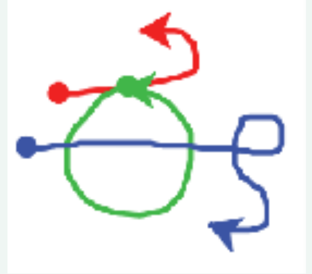

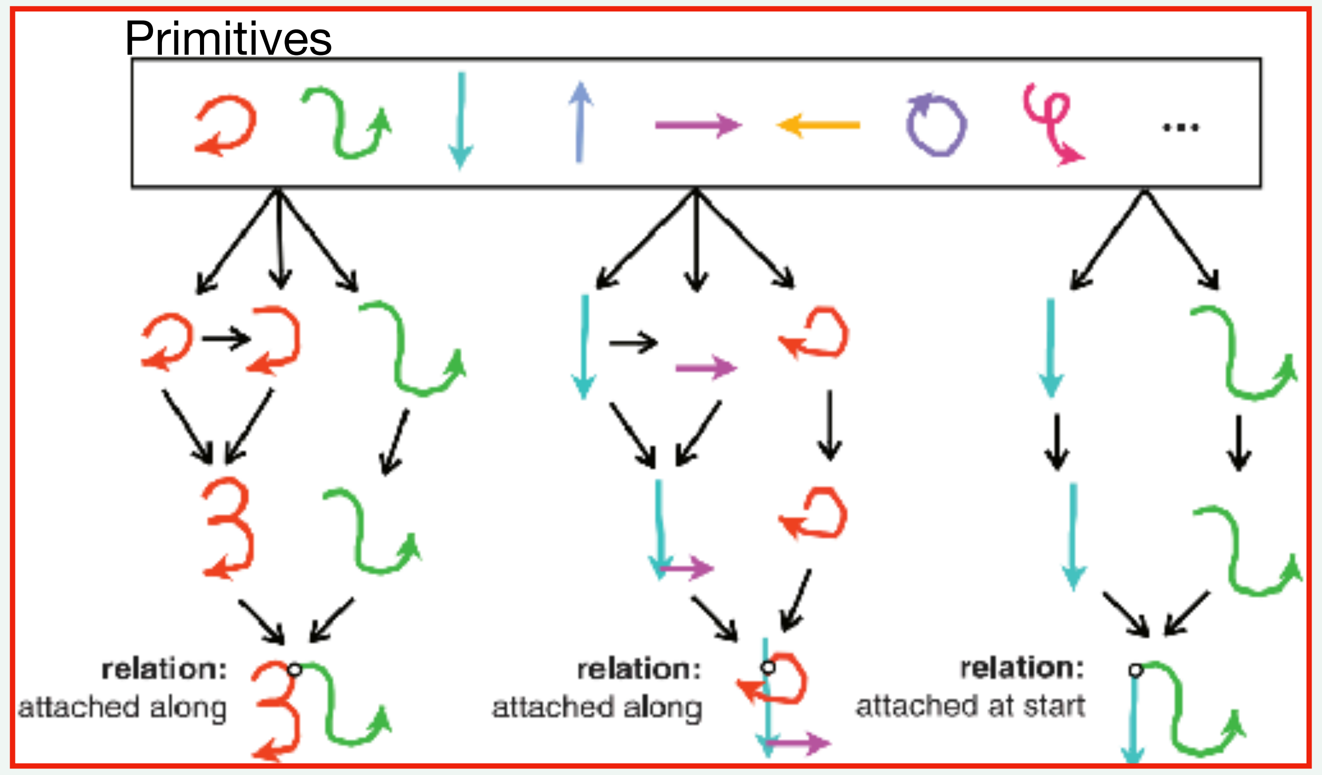

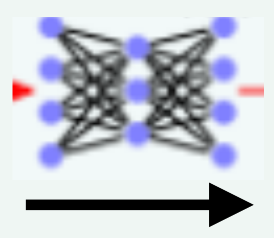

```
procedure GENERATETYPE
  κ ← P(κ)                                  ▷ Sample number of parts
  for i = 1 ... κ do
    n_i ← P(n_i|κ)                          ▷ Sample number of sub-parts
    for j = 1 ... n_i do
      s_ij ← P(s_ij|s_i(j−1))               ▷ Sample sub-part sequence
    end for
    R_i ← P(R_i|S_1, ..., S_i−1)            ▷ Sample relation
  end for
  ψ ← {κ, R, S}
  return @GENERATETOKEN(ψ)                  ▷ Return program
```

List Processing

Sum List
[1 2 3] → 6
[4 6 8 1] → 17

Double
[1 2 3] → [2 4 6]
[4 5 1] → [8 10 2]

Check Evens
[0 2 3] → [T T F]
[2 9 6] → [T F T]

Text Editing

Abbreviate
Allen Newell → A.N.
Herb Simon → H.S.

Drop Last Three
shrdlu → shr
shakey → sha

Extract
a b (c) → c
a (bee) see → see

Regexes

Phone numbers
(555) 867-5309
(650) 555-2368

Currency
$100.25
$4.50

Dates
Y1775/0704
Y2000/0101

LOGO Graphics

Block Towers

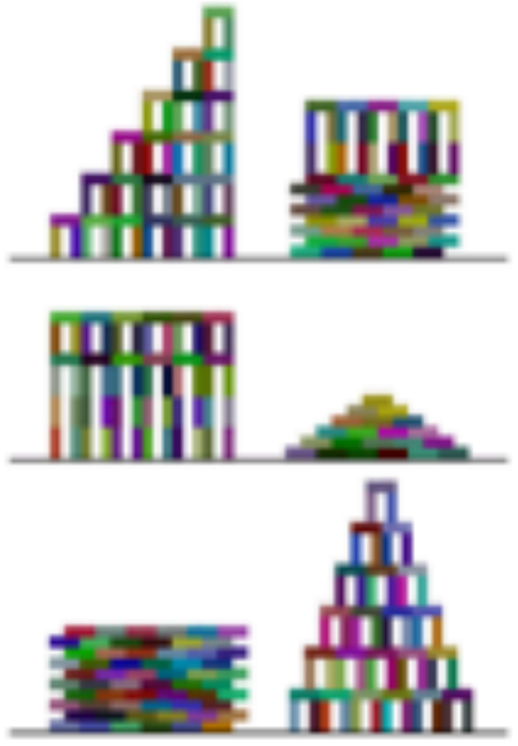

Symbolic Regression

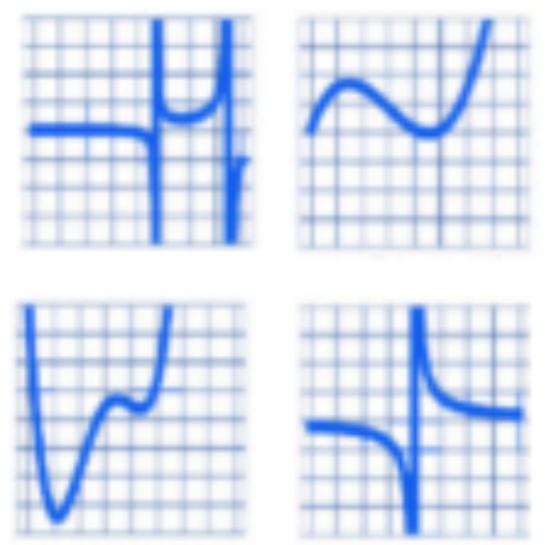

$y = f(x)$

Recursive Programming

Filter Red

Length

Physical Laws

$$\vec{a} = \frac{1}{m}\sum_i \vec{F}_i$$

$$\vec{F} \propto \frac{q_1 q_2}{|\vec{r}|^2}\hat{r}$$

$$R_{\text{total}} = \left(\sum_i \frac{1}{R_i}\right)^{-1}$$

Dreamcoder (Ellis et al., 2020)

not on the exam

Wake-Sleep Algorithm

- Inspired by Hinton et al., (1995)
- **Wake**: find the best program to solve the current task using a recognition model (neural network)

$$\arg\max_{\pi} P(\mathscr{D} \mid \pi)$$

- **Sleep:** Update $P(\pi)$
  - ***Abstraction***: Grow library to find more compressible programs
  - ***Dreaming***: Train recognition model by sampling programs that solved previous experienced tasks (*replays*) and by sampling tasks that can be solved by programs in the current library (*fantasies*)

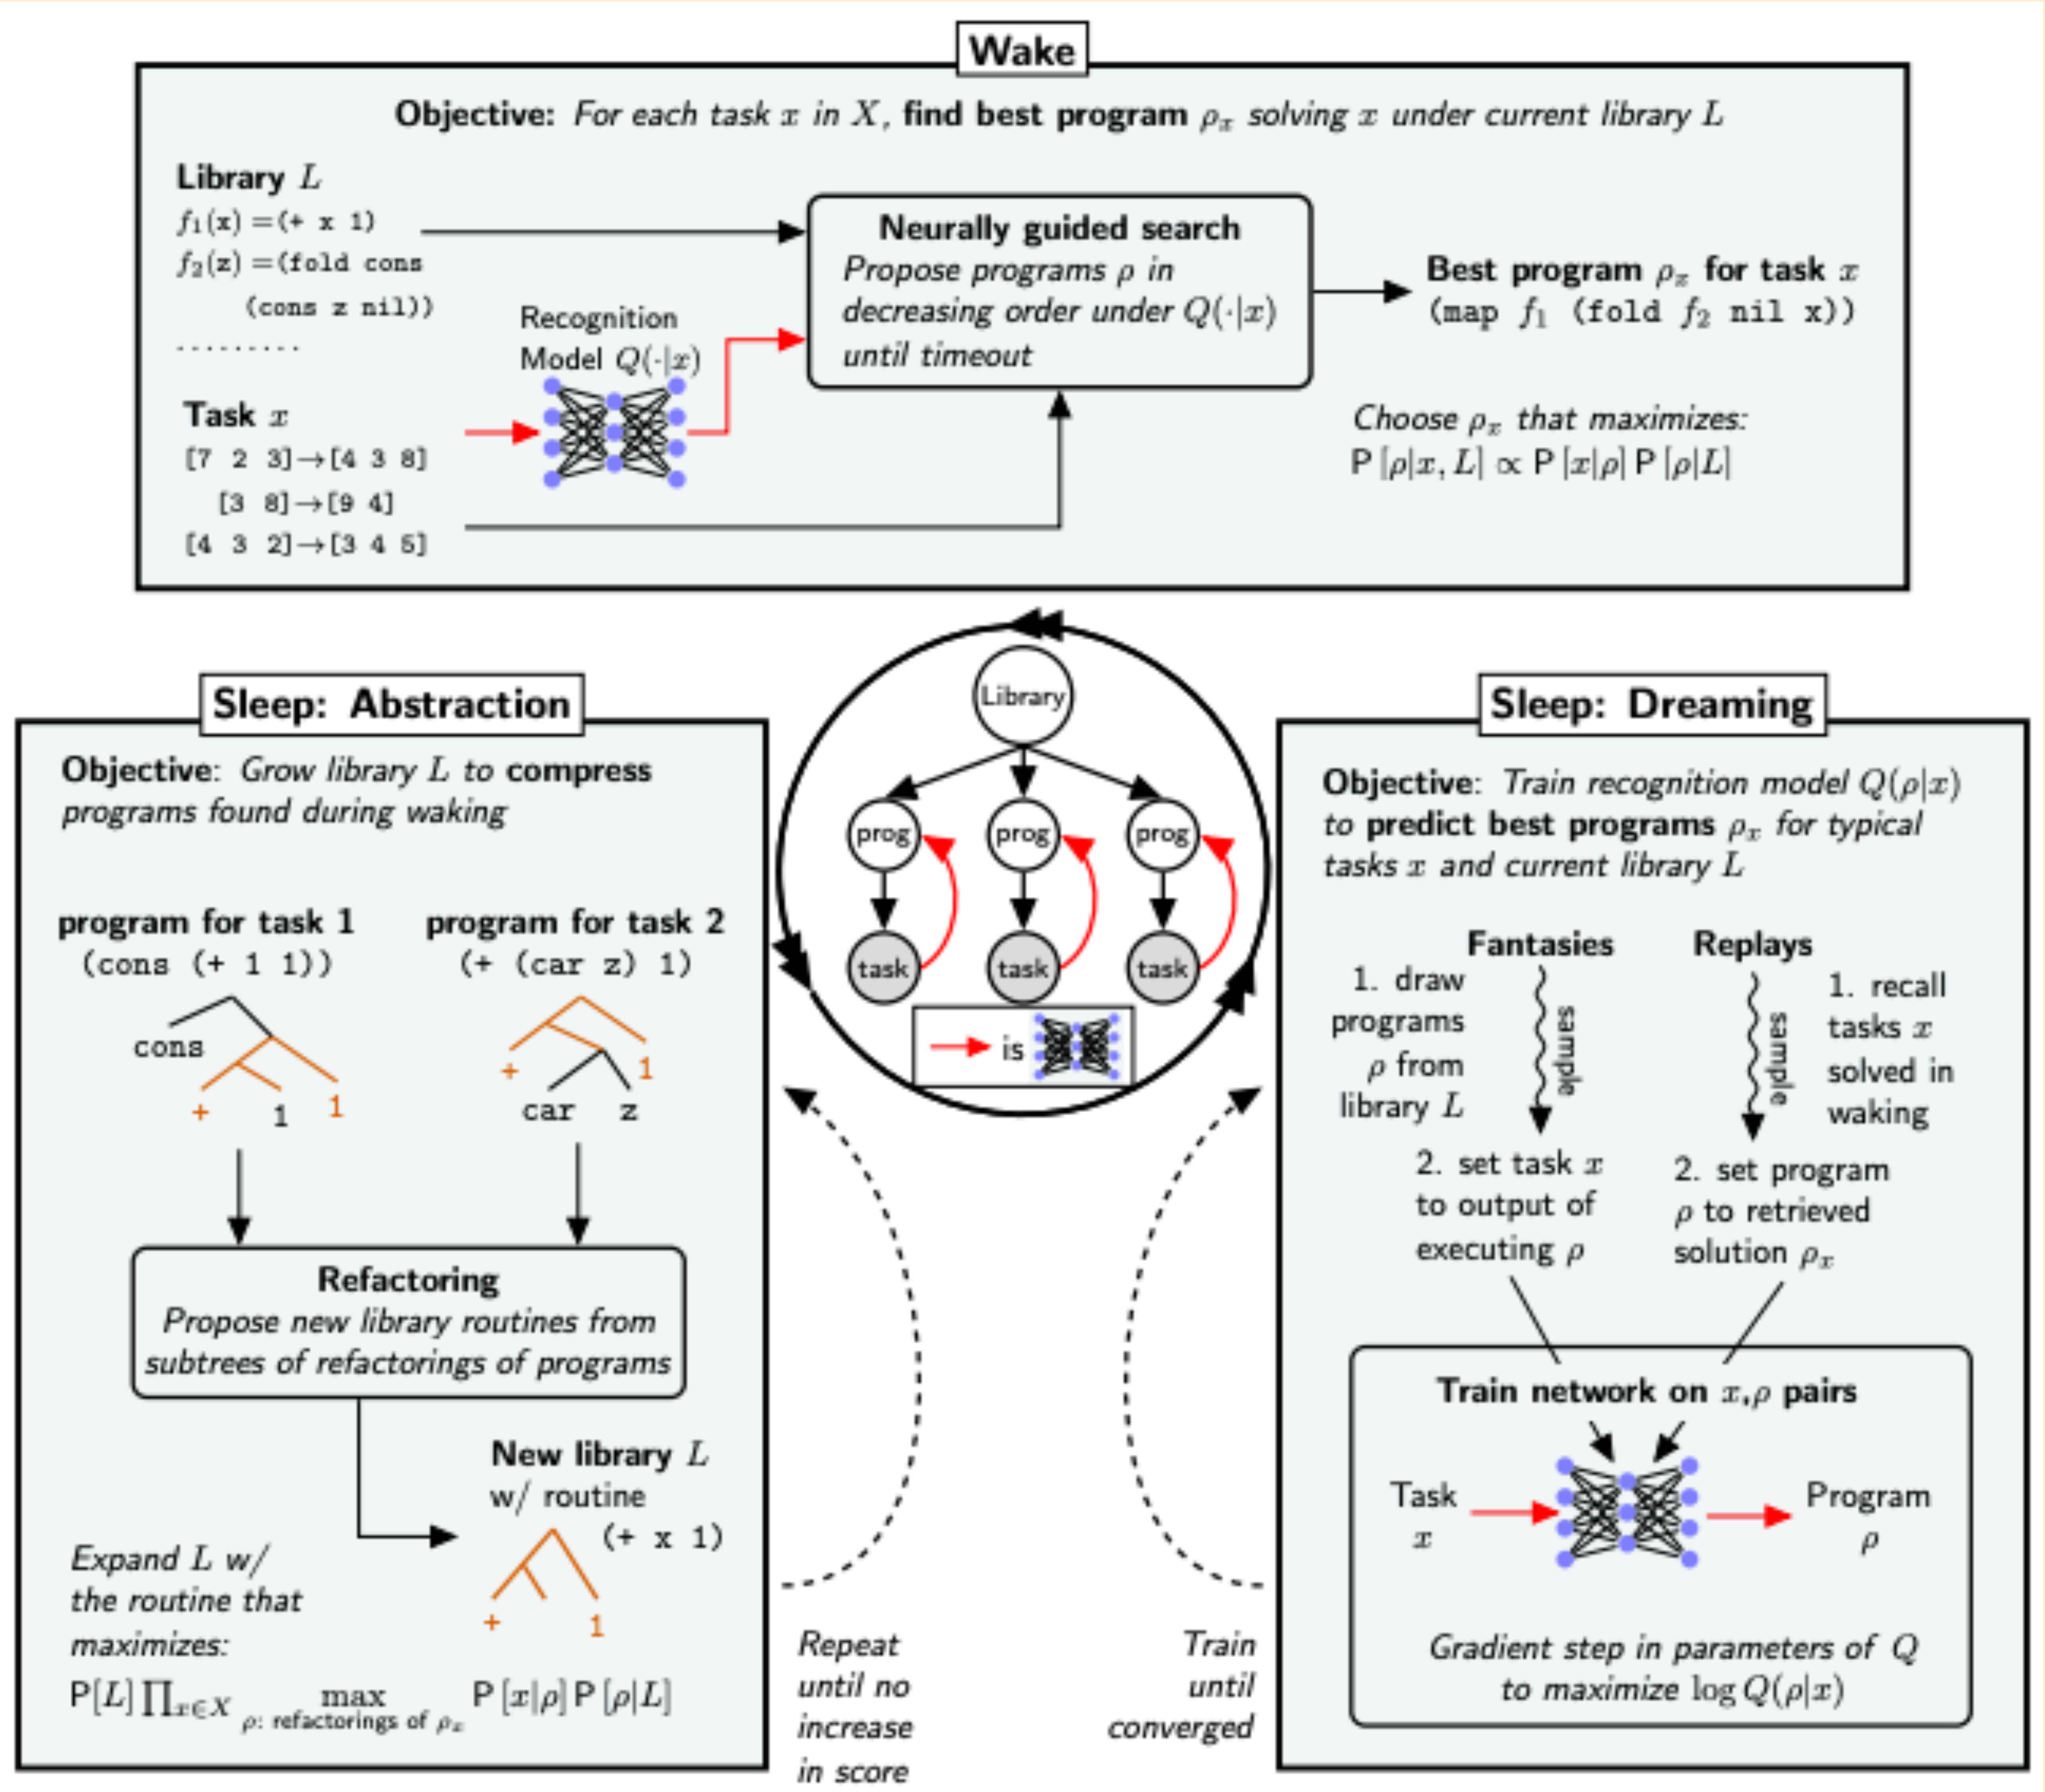

Dreamcoder (Ellis et al., 2020)

not on the exam

$$P(\pi \mid \mathscr{D}) \propto P(\mathscr{D} \mid \pi) P(\pi)$$

Wake-Sleep Algorithm

- Inspired by Hinton et al., (1995)
- **Wake**: find the best program to solve the current task using a recognition model (neural network)

  $$\arg\max_{\pi} P(\mathscr{D} \mid \pi)$$

- **Sleep:** Update $P(\pi)$
  - ***Abstraction***: Grow library to find more compressible programs
  - ***Dreaming***: Train recognition model by sampling programs that solved previous experienced tasks (*replays*) and by sampling tasks that can be solved by programs in the current library (*fantasies*)

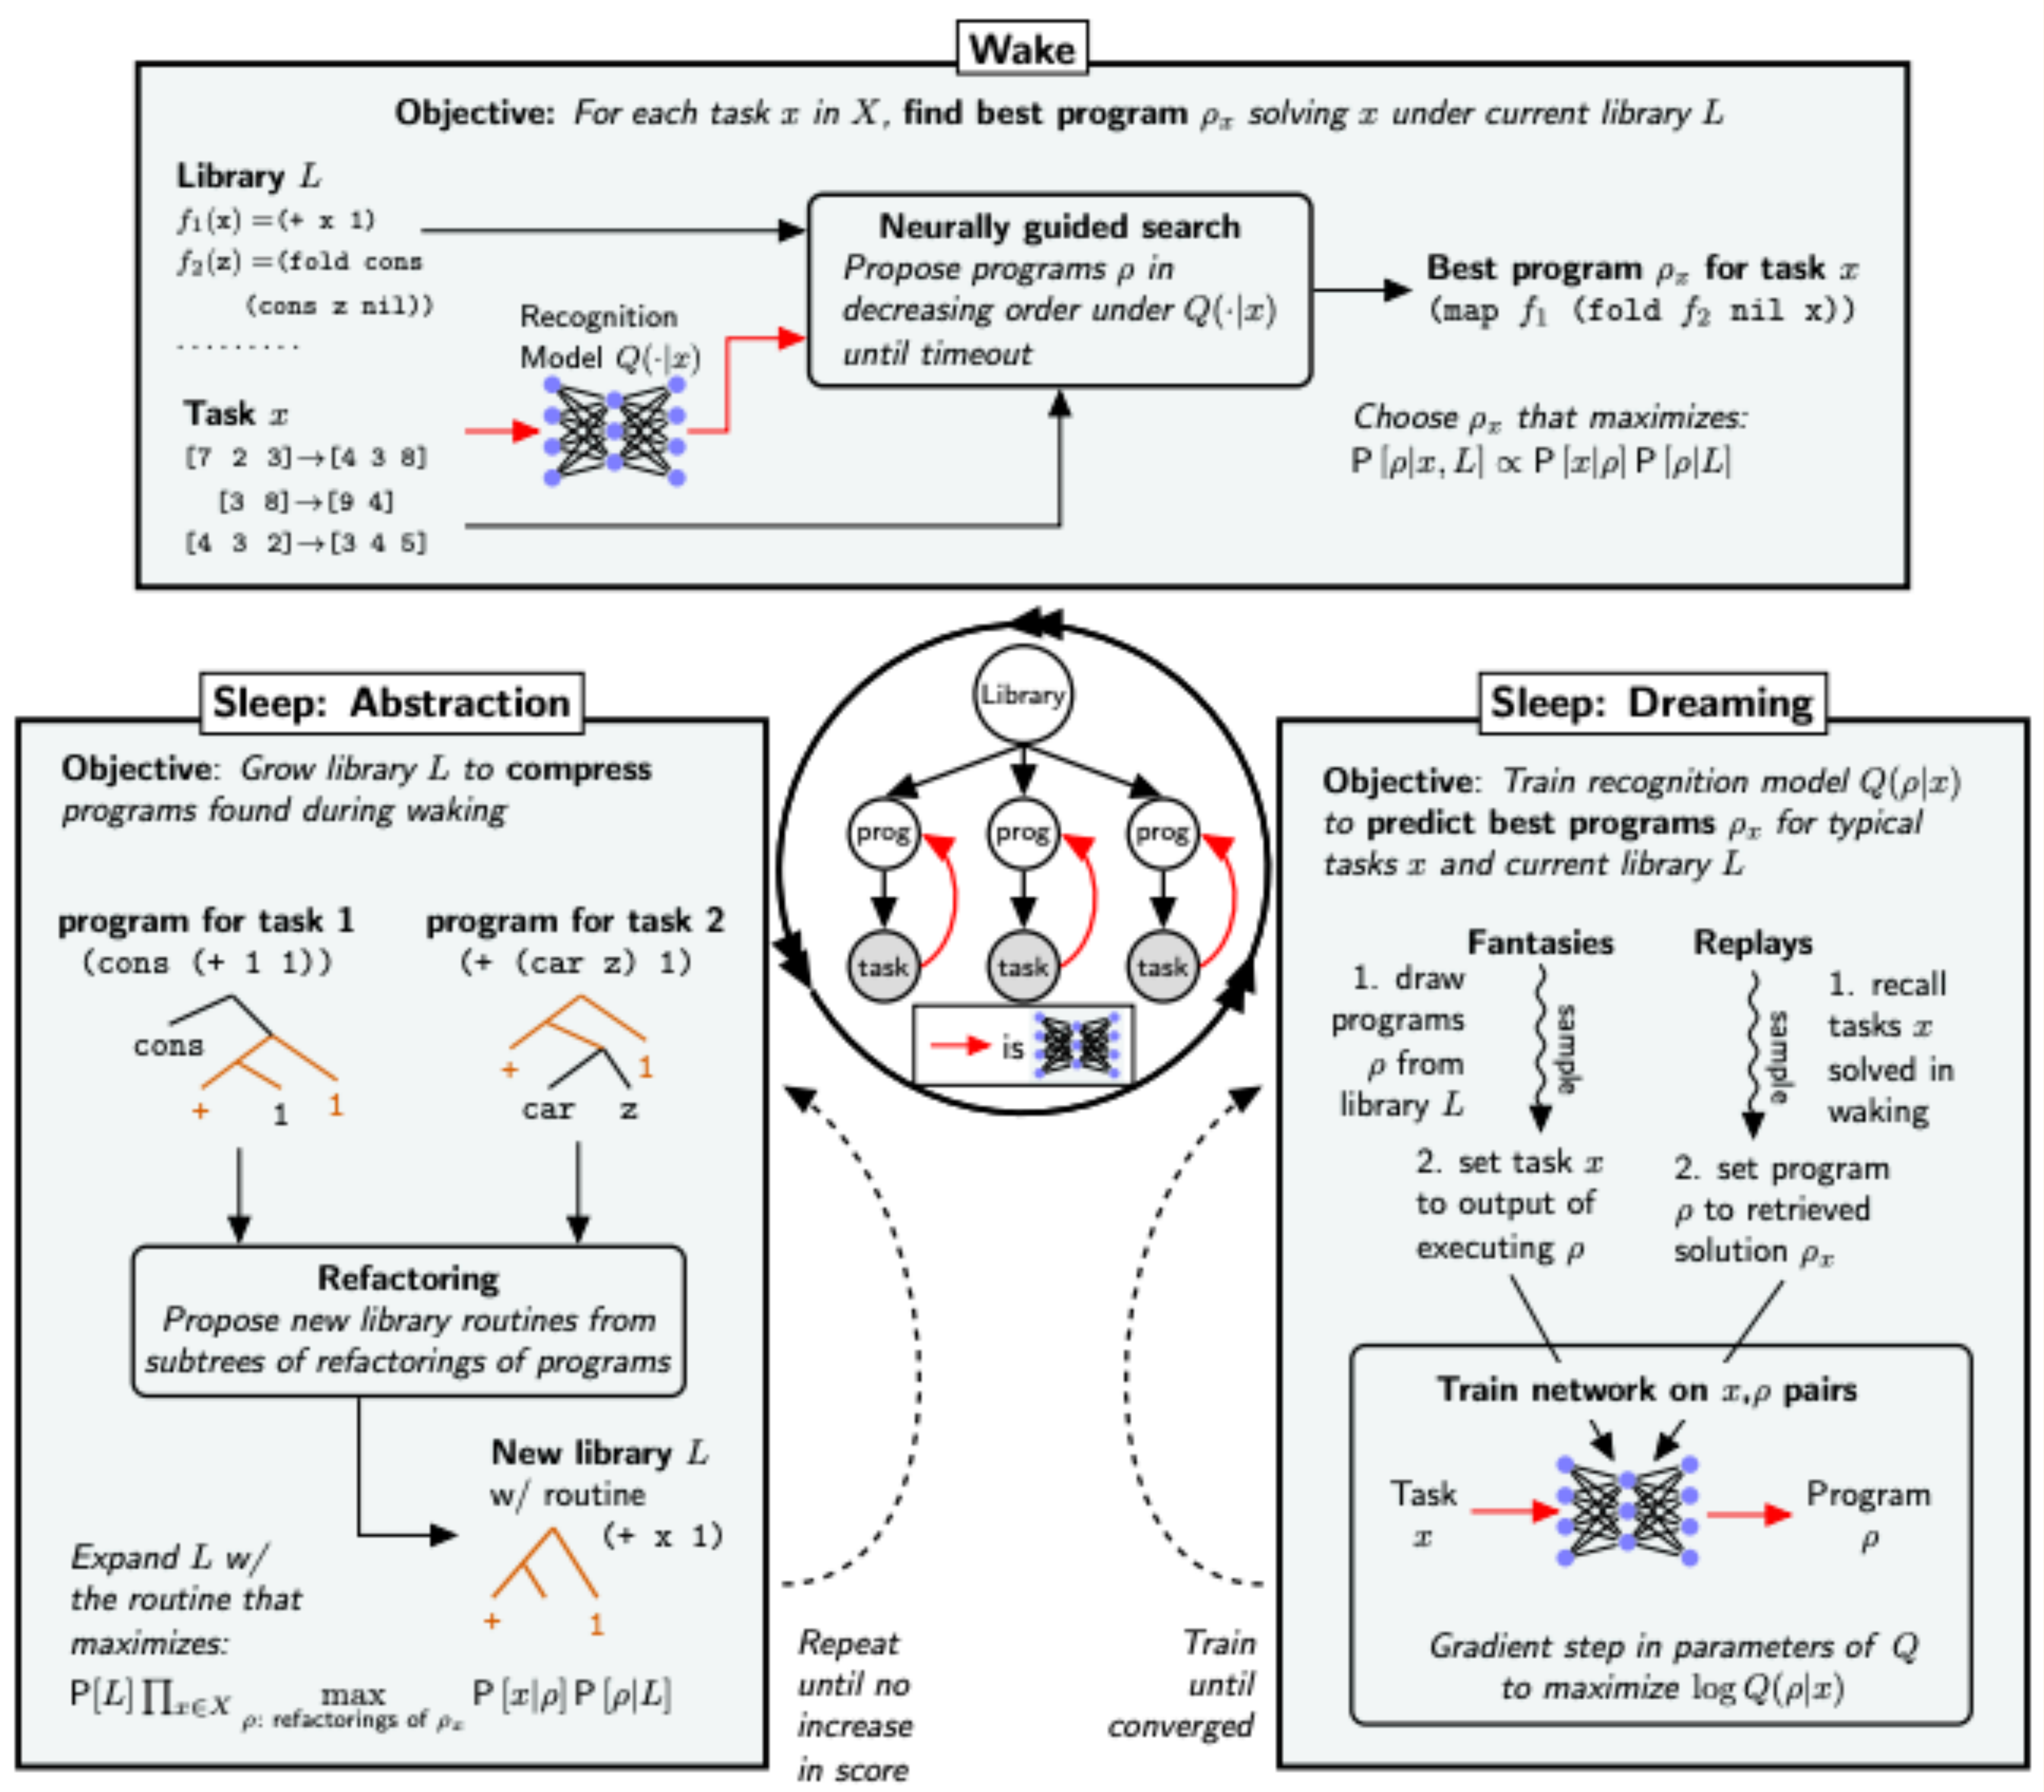

Dreamcoder (Ellis et al., 2020)

Current work in my lab

Current work in my lab

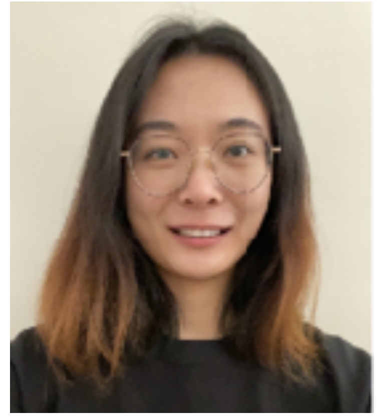
Hanqi Zhou


David Nagy

- Can program induction inform us about how people represent the world? (Fodor, 1975; Piantadosi et al., 2016; Dehaene et al., 2022)

Current work in my lab

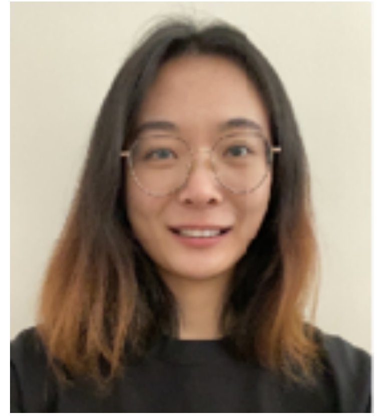
Hanqi Zhou


David Nagy

- Can program induction inform us about how people represent the world? (Fodor, 1975; Piantadosi et al., 2016; Dehaene et al., 2022)

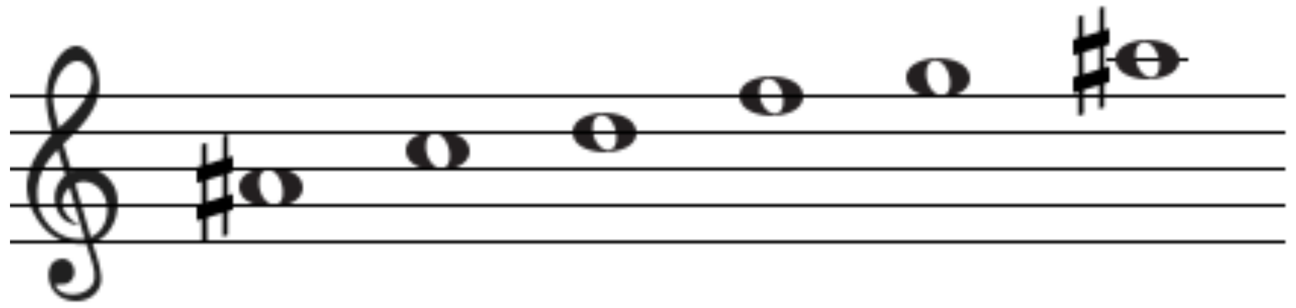

Current work in my lab

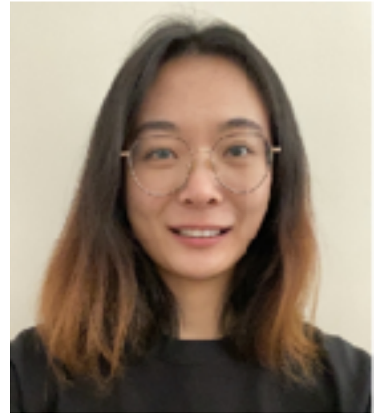

Hanqi Zhou

David Nagy

- Can program induction inform us about how people represent the world? (Fodor, 1975; Piantadosi et al., 2016; Dehaene et al., 2022)

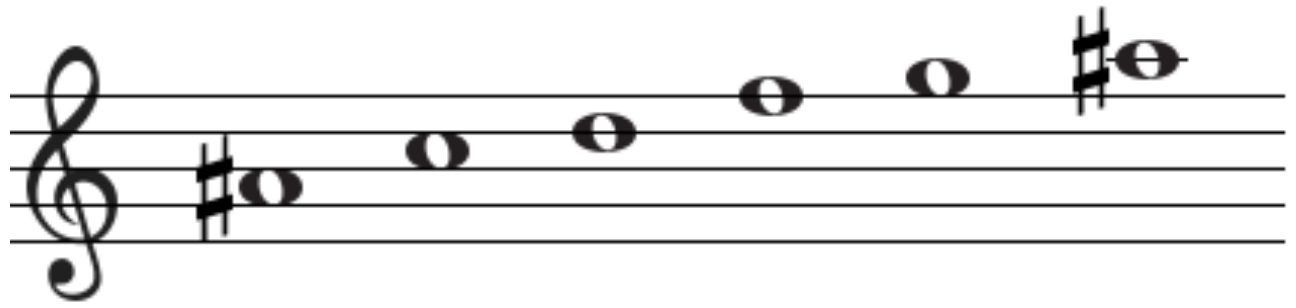

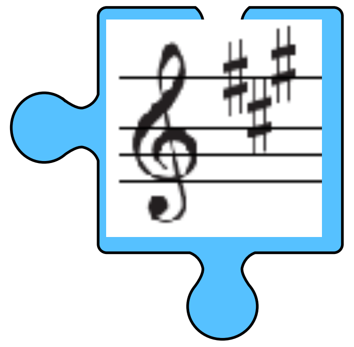

Current work in my lab

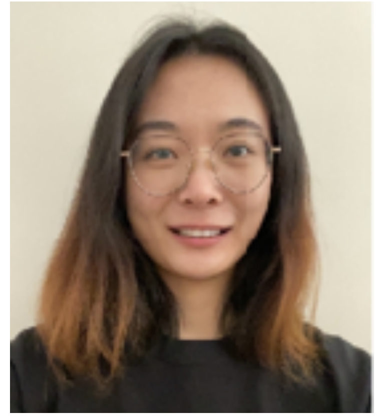
Hanqi Zhou


David Nagy

- Can program induction inform us about how people represent the world? (Fodor, 1975; Piantadosi et al., 2016; Dehaene et al., 2022)

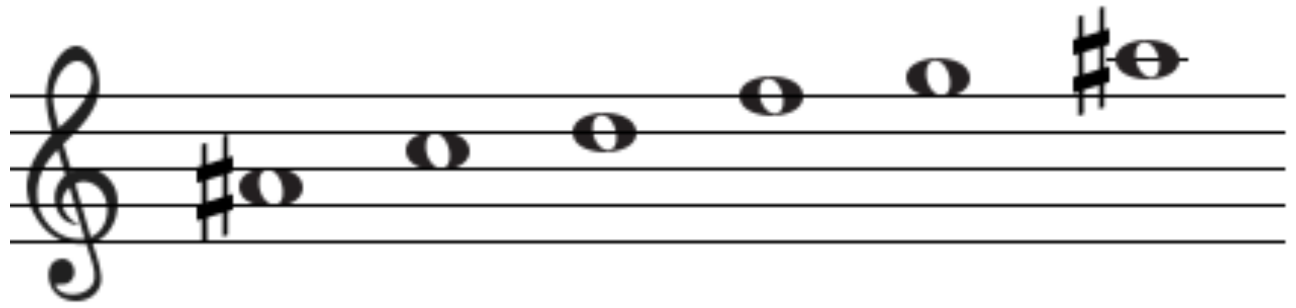

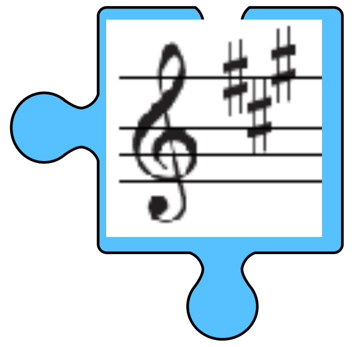

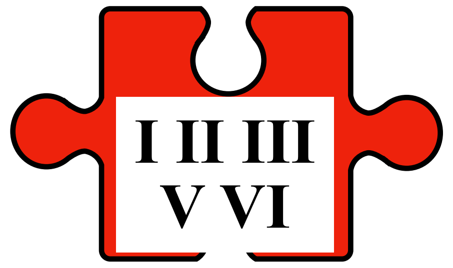

Current work in my lab

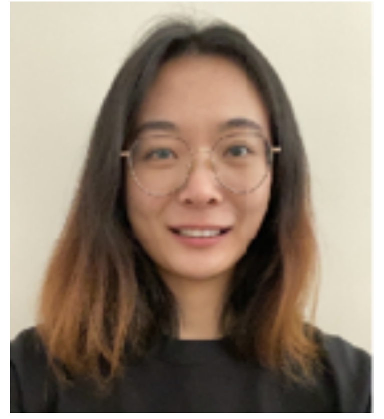

Hanqi Zhou

David Nagy

- Can program induction inform us about how people represent the world? (Fodor, 1975; Piantadosi et al., 2016; Dehaene et al., 2022)

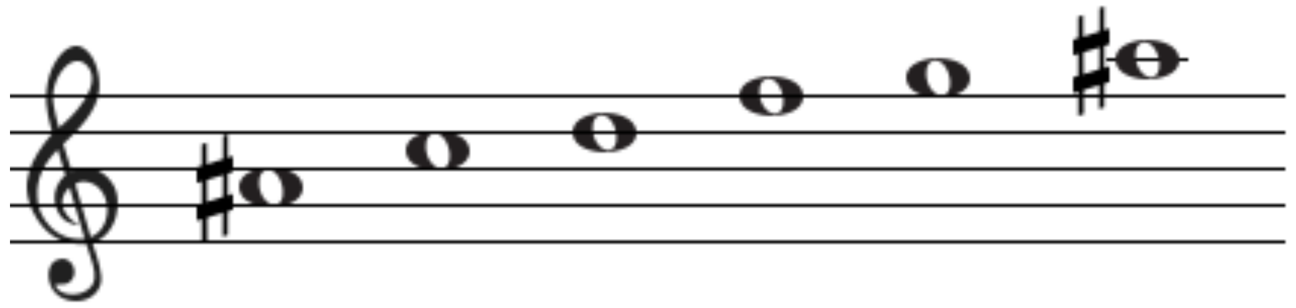

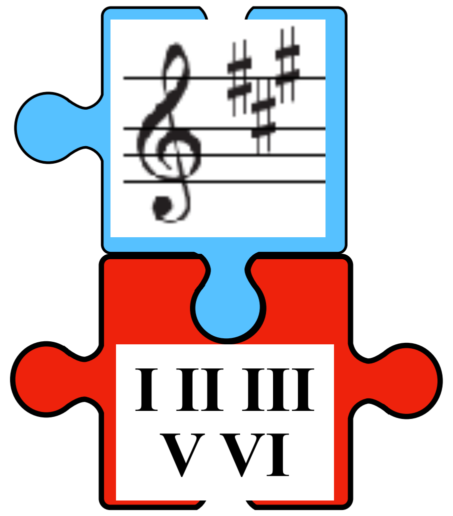

Current work in my lab

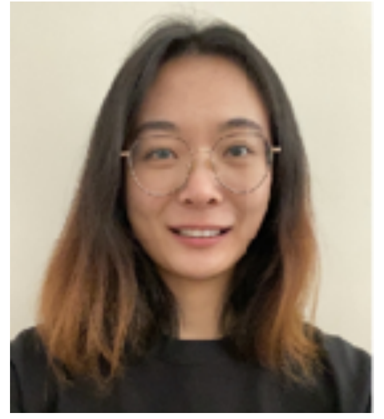

- Can program induction inform us about how people represent the world? (Fodor, 1975; Piantadosi et al., 2016; Dehaene et al., 2022)

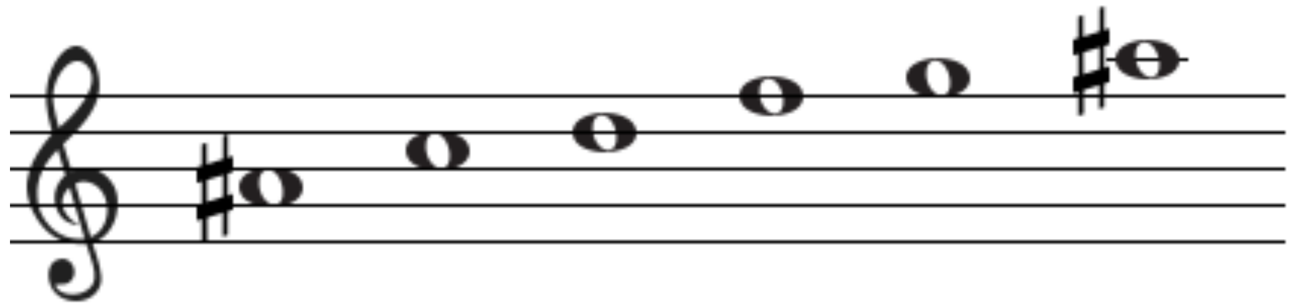

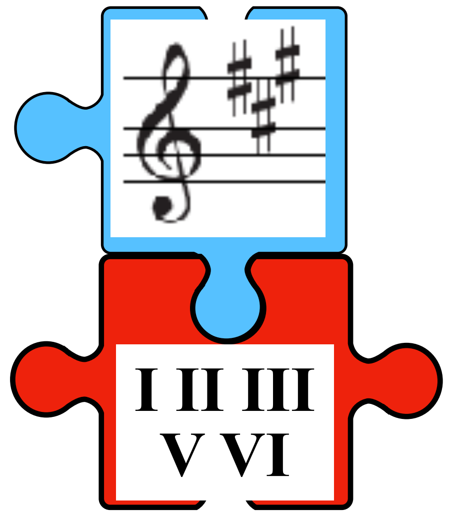

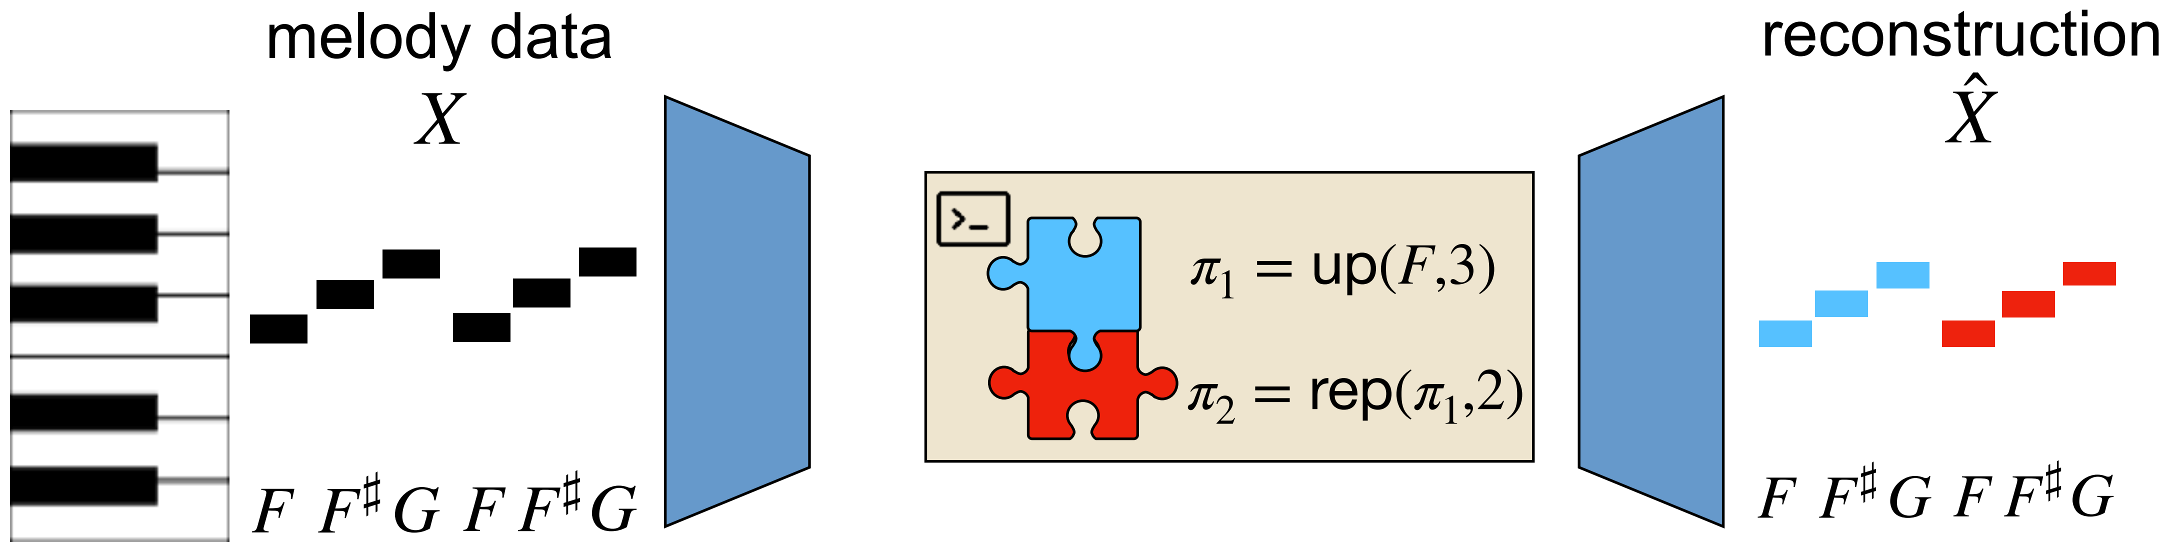

- **Program induction under constraints**
  (Zhou, Nagy & Wu, 2024)

Current work in my lab

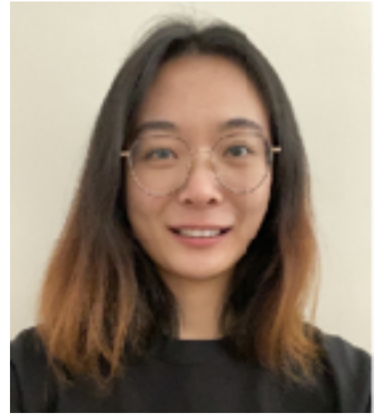

- Can program induction inform us about how people represent the world? (Fodor, 1975; Piantadosi et al., 2016; Dehaene et al., 2022)

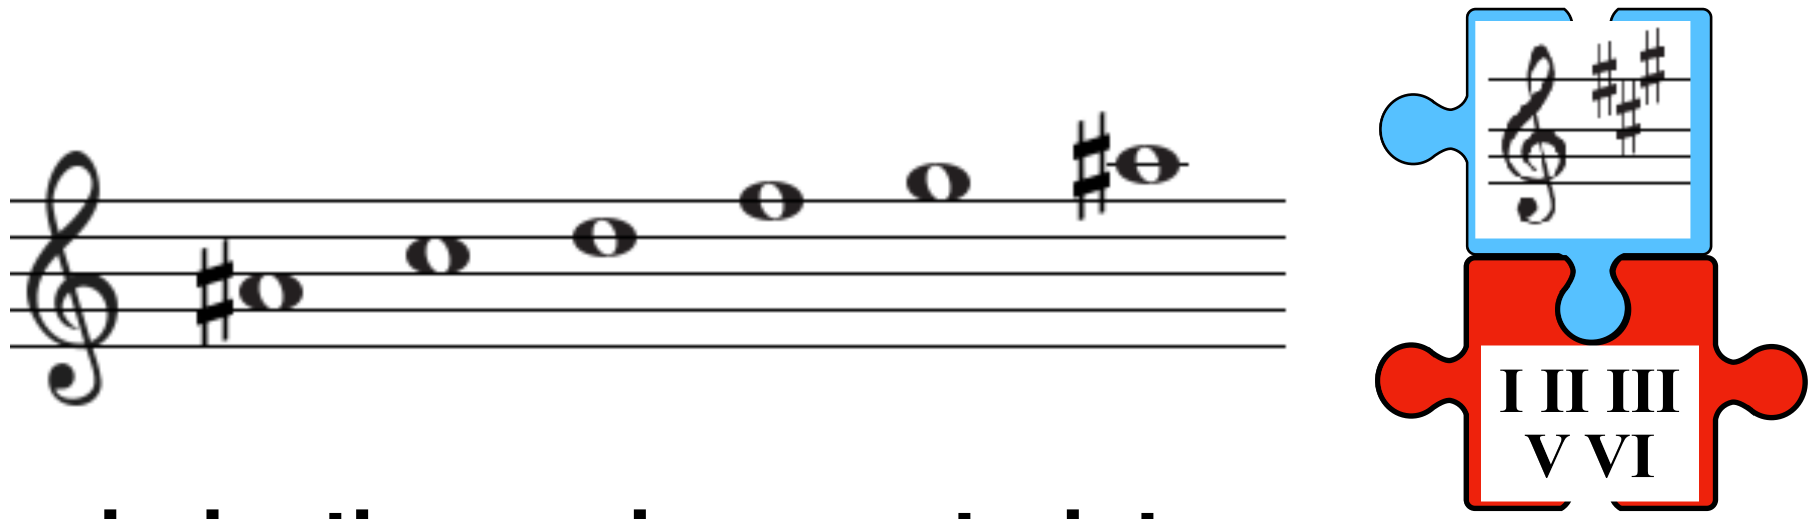

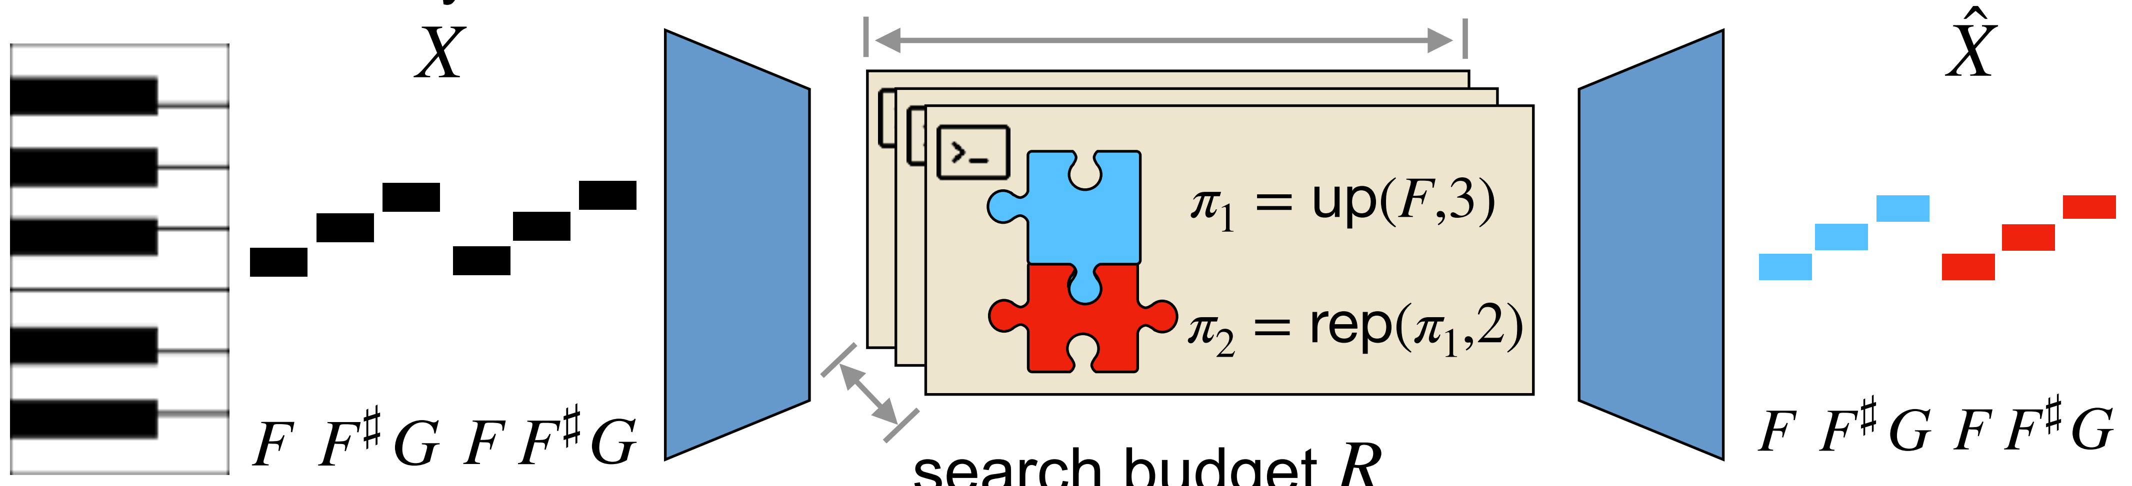

- **Program induction under constraints** (Zhou, Nagy & Wu, 2024)
    - Encoding melodies as programs under limited resources: description length R_l and search budget R_s

Current work in my lab

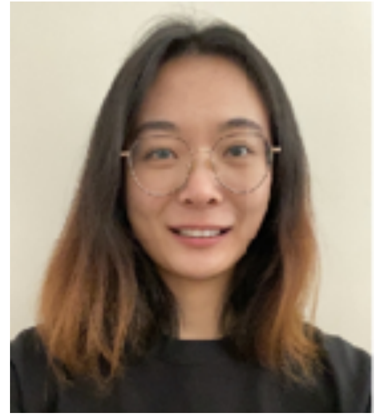

- Can program induction inform us about how people represent the world? (Fodor, 1975; Piantadosi et al., 2016; Dehaene et al., 2022)

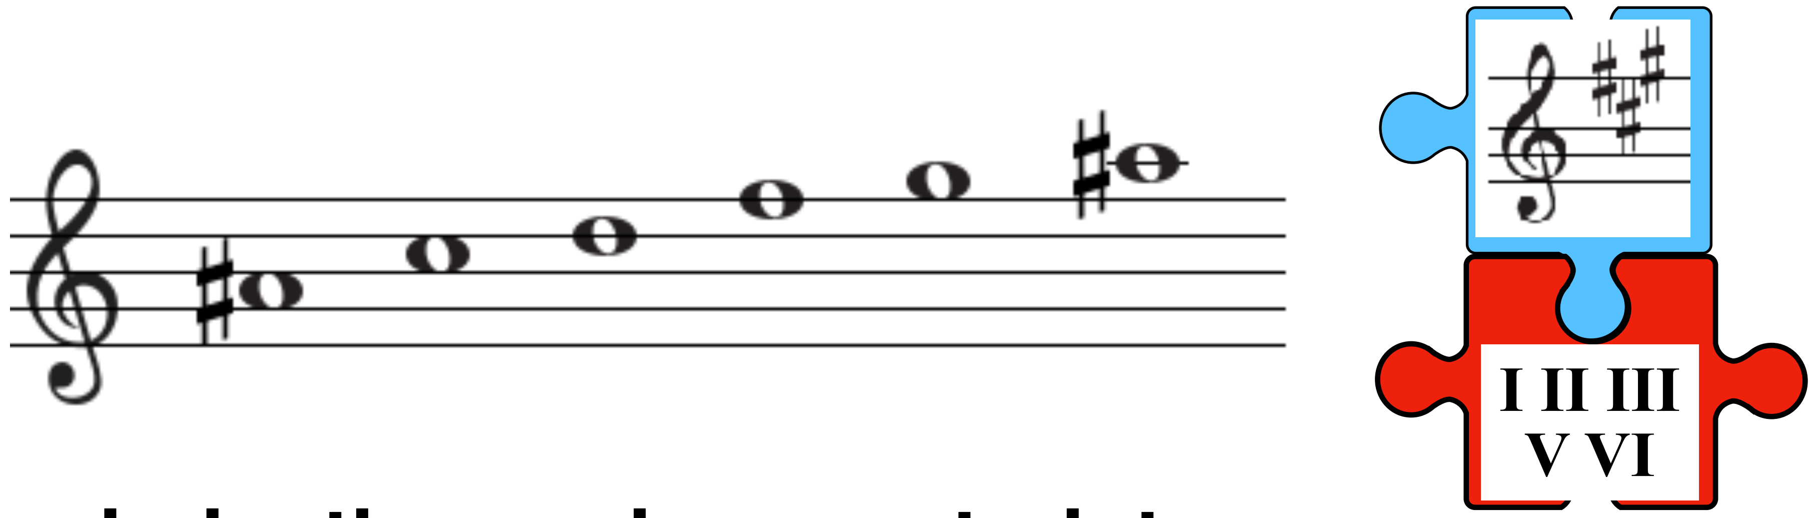

- **Program induction under constraints** (Zhou, Nagy & Wu, 2024)
  - Encoding melodies as programs under limited resources: description length R_l and search budget R_s
  - We show cognitive constraints drive curriculum effects …due to storing/reusing programs in a concept library

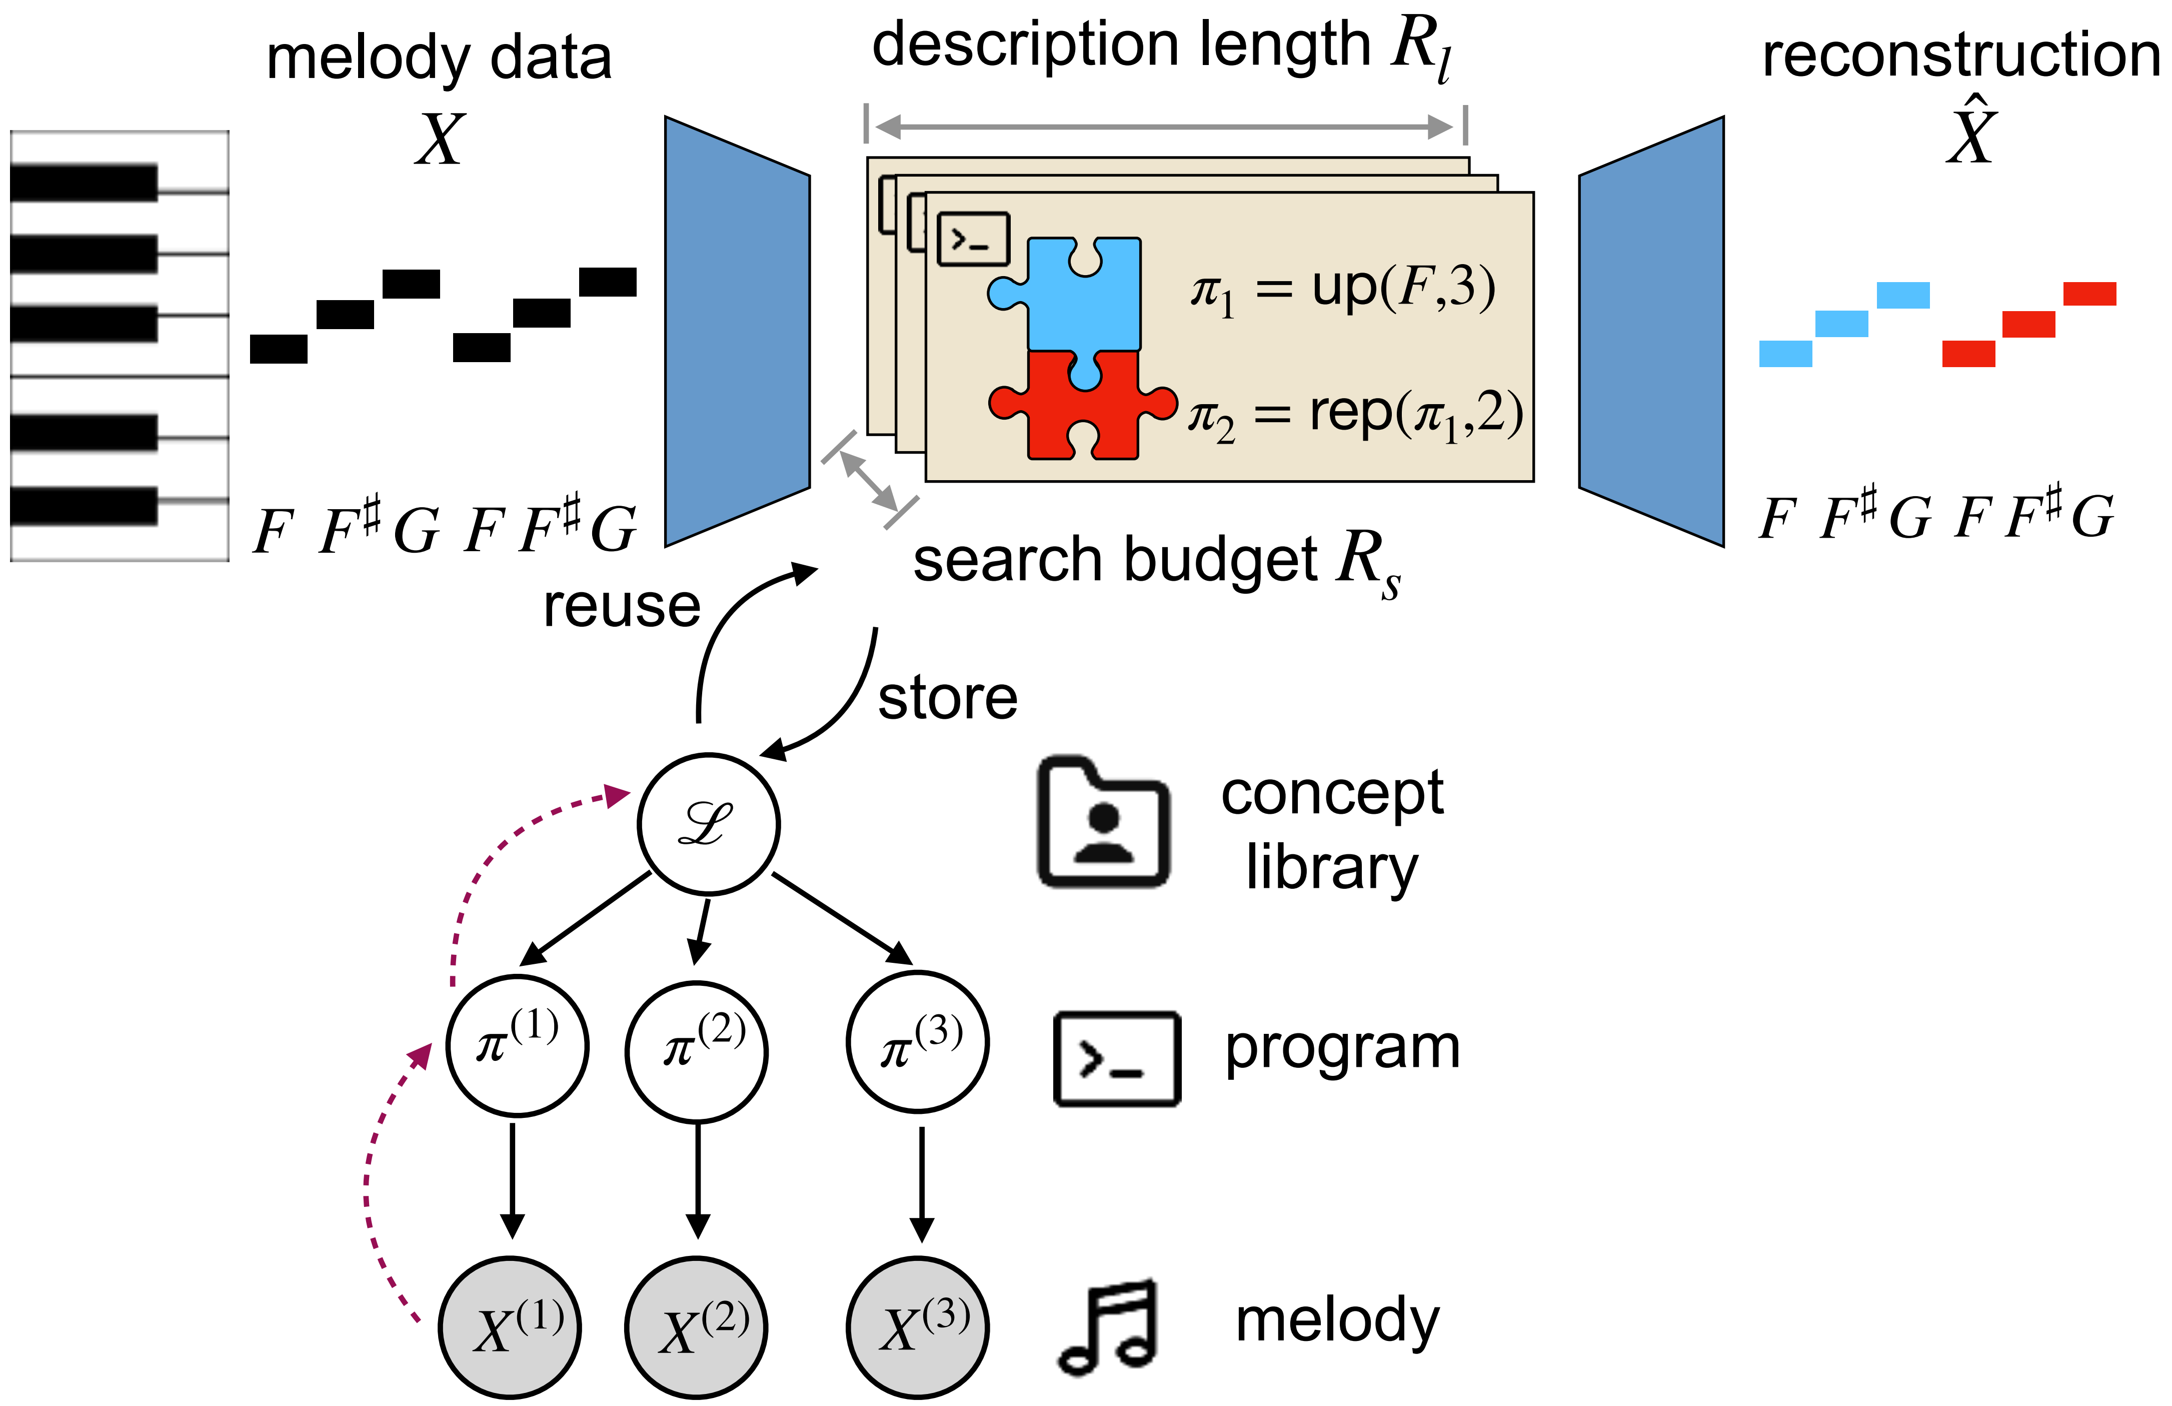

Current work in my lab

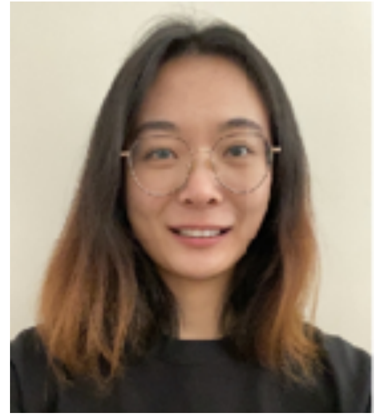
Hanqi Zhou


David Nagy

- Can program induction inform us about how people represent the world? (Fodor, 1975; Piantadosi et al., 2016; Dehaene et al., 2022)

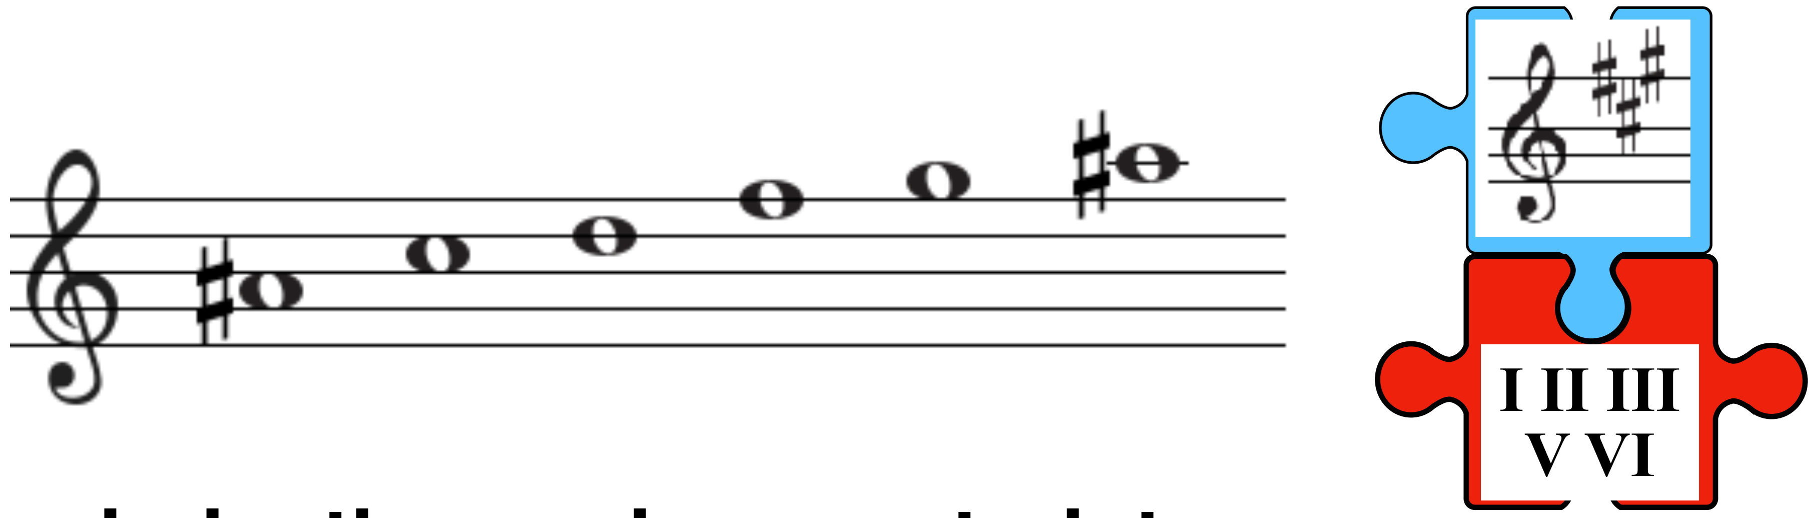

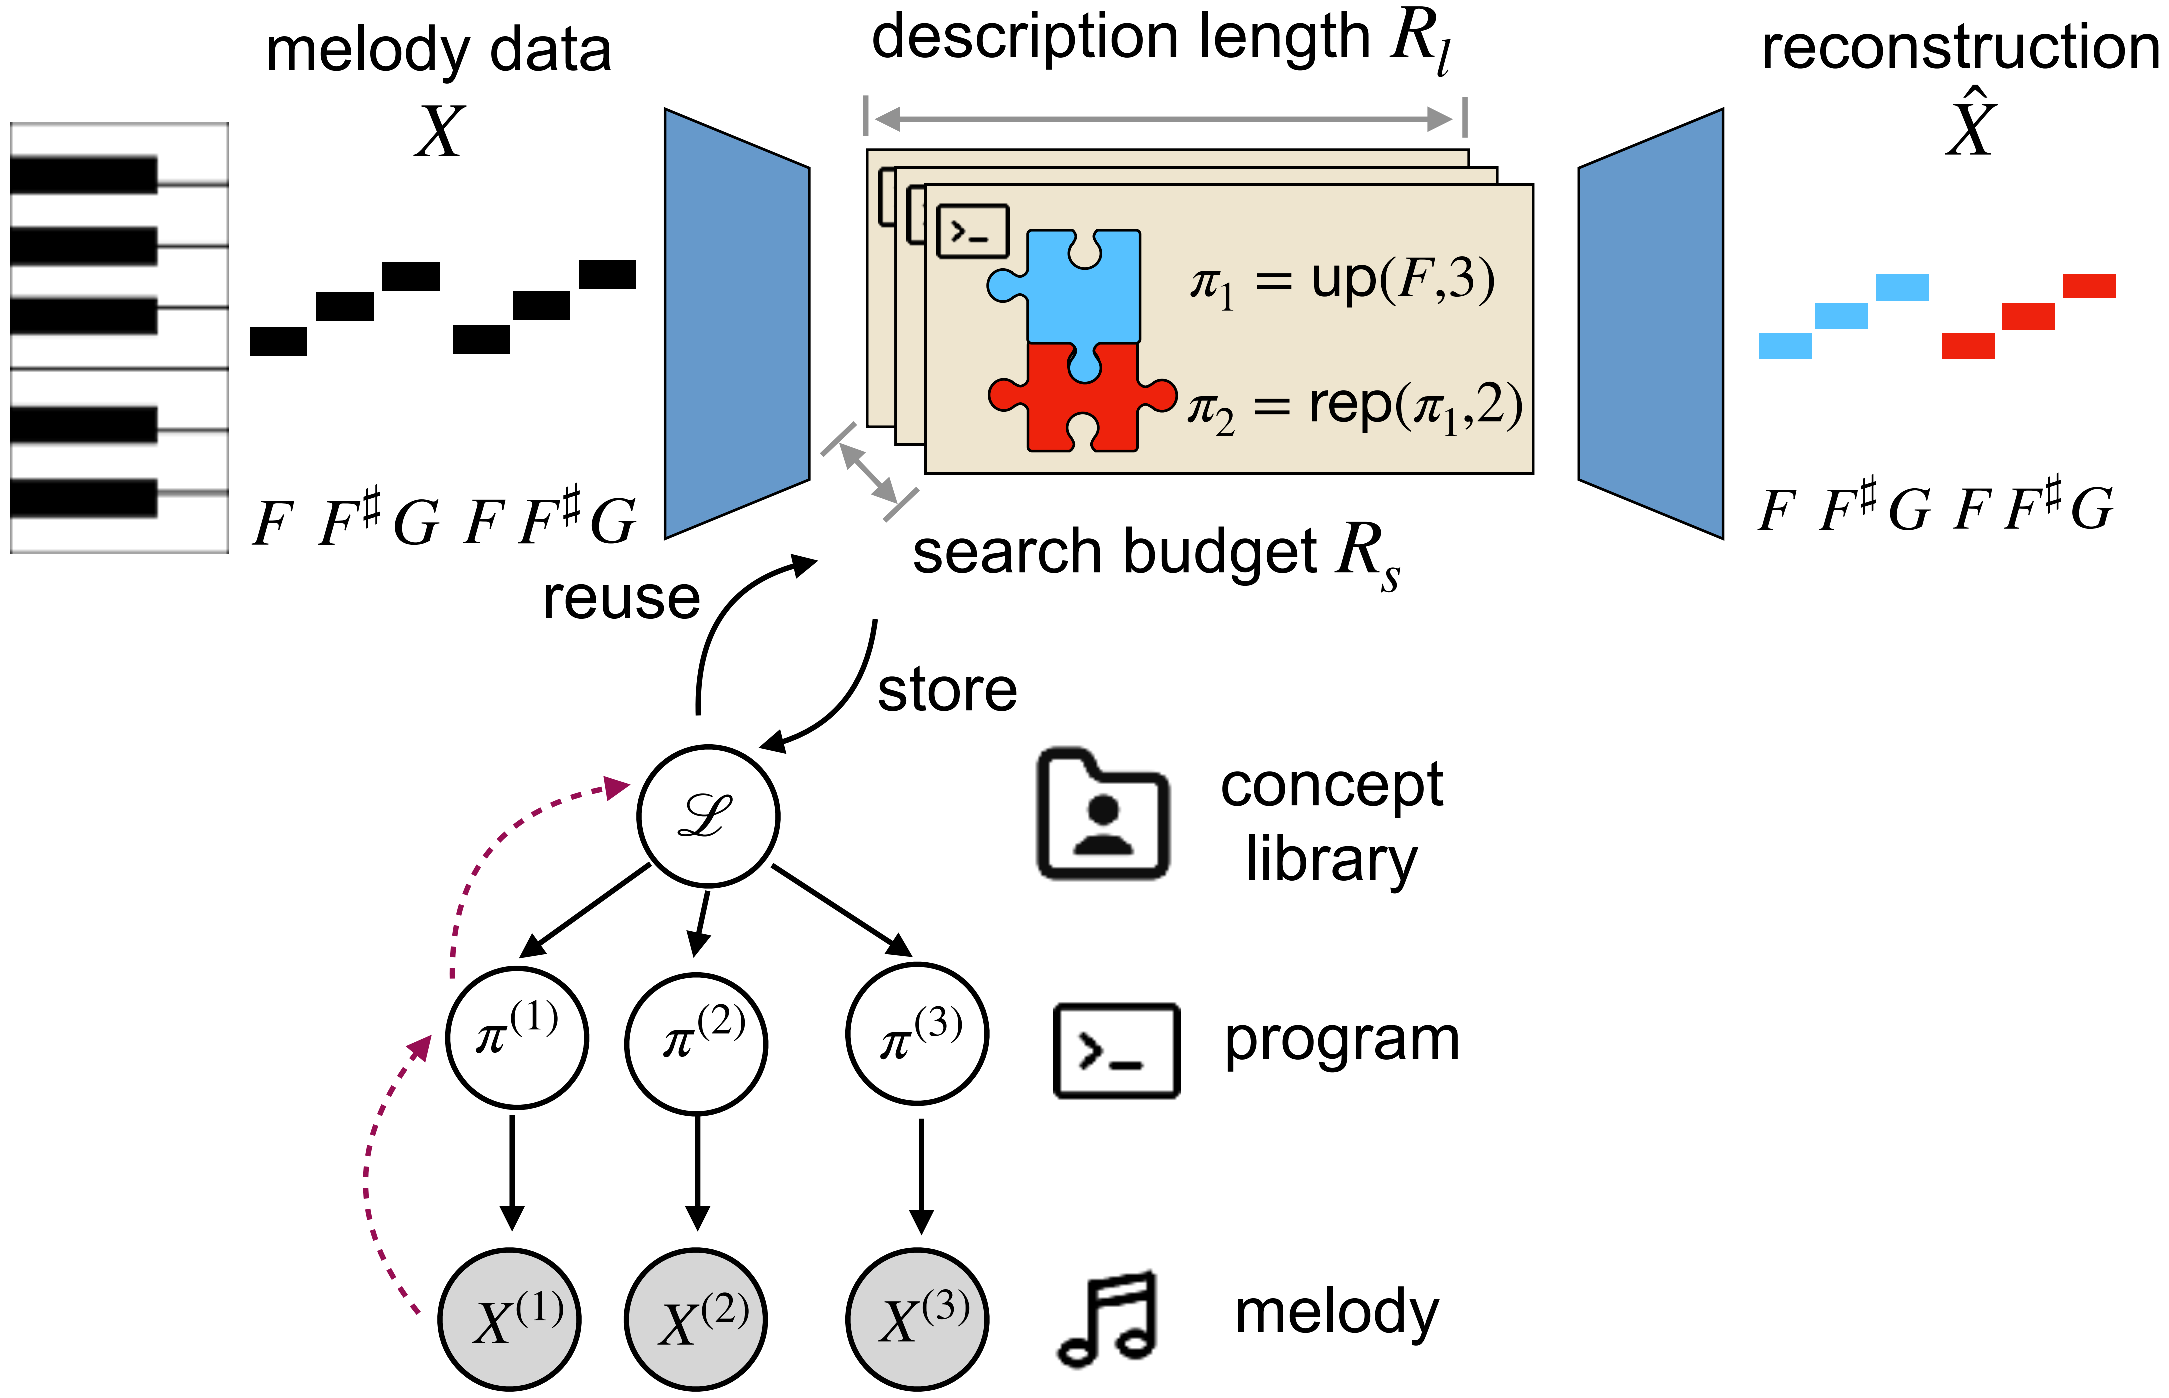

- **Program induction under constraints** (Zhou, Nagy & Wu, 2024)
  - Encoding melodies as programs under limited resources: description length R_l and search budget R_s
  - We show cognitive constraints drive curriculum effects …due to storing/reusing programs in a concept library
  - We can *a priori* generate **synergistic curricula** that result in better generalization (for the model)

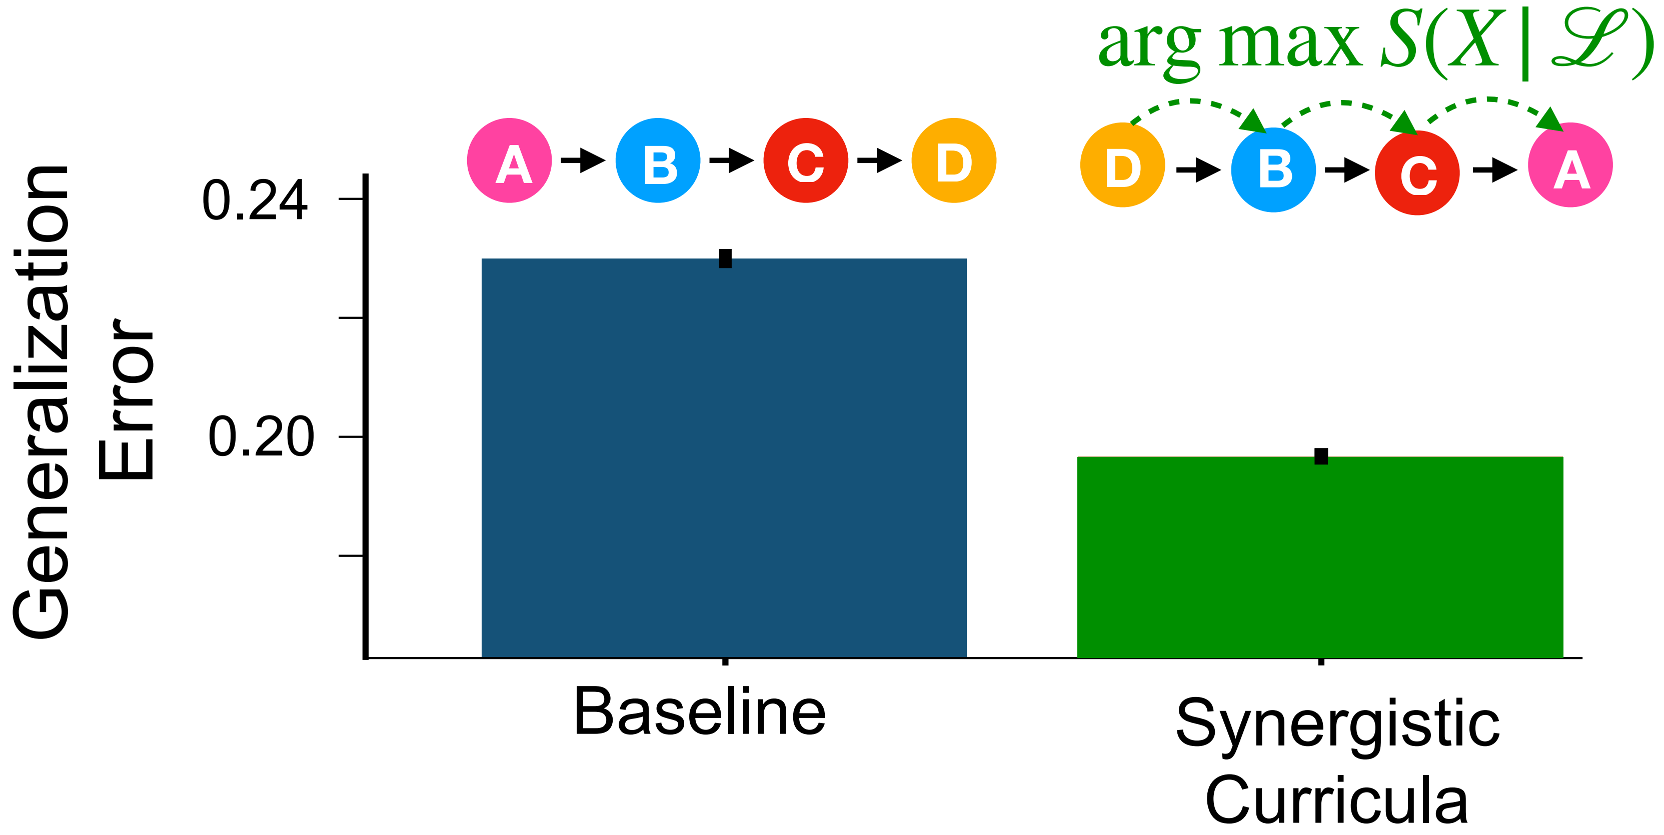

Current work in my lab

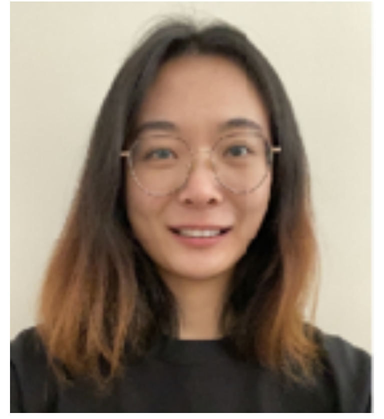
Hanqi Zhou


David Nagy

- Can program induction inform us about how people represent the world? (Fodor, 1975; Piantadosi et al., 2016; Dehaene et al., 2022)

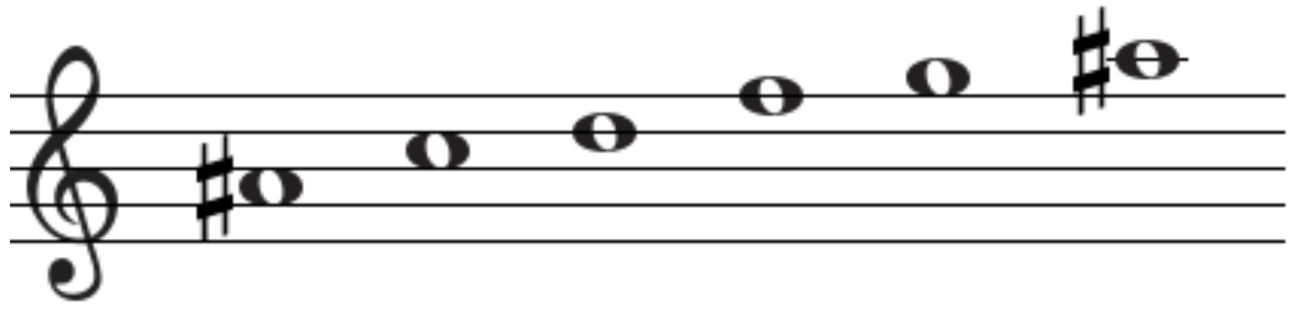

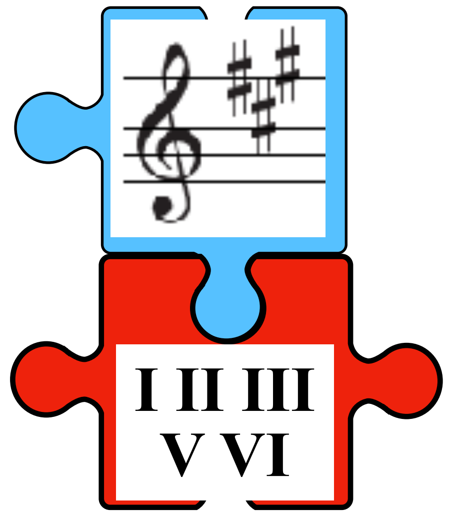

- **Program induction under constraints**
  (Zhou, Nagy & Wu, 2024)
  - Encoding melodies as programs under limited resources: description length R_l and search budget R_s
  - We show cognitive constraints drive curriculum effects …due to storing/reusing programs in a concept library
  - We can *a priori* generate **synergistic curricula** that result in better generalization (for the model)

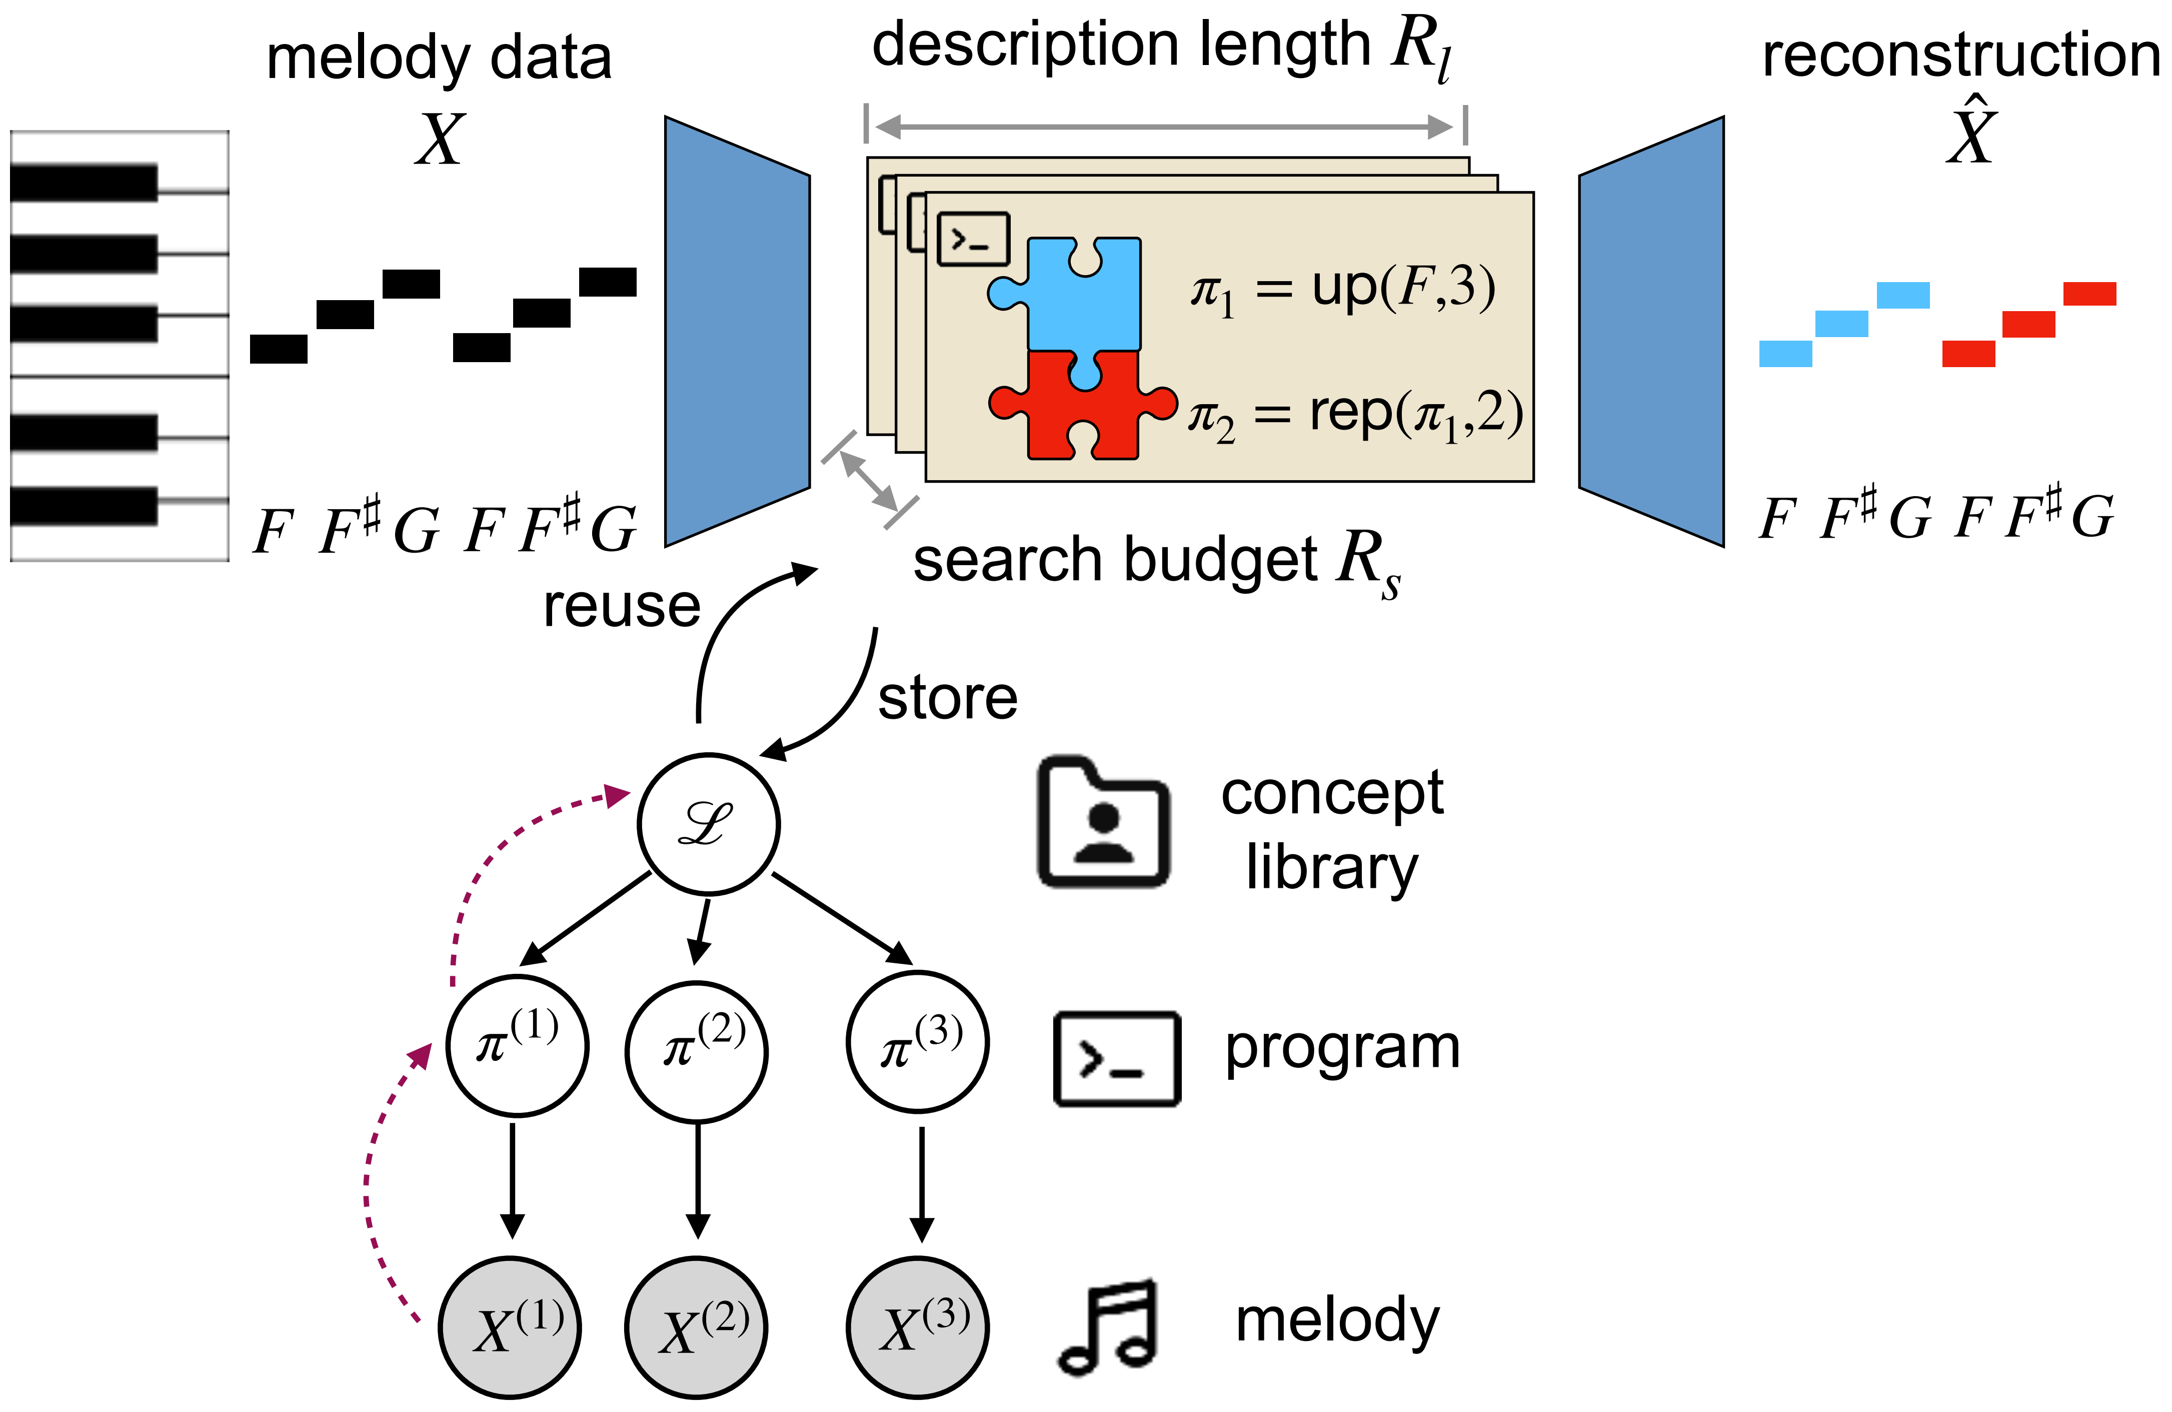

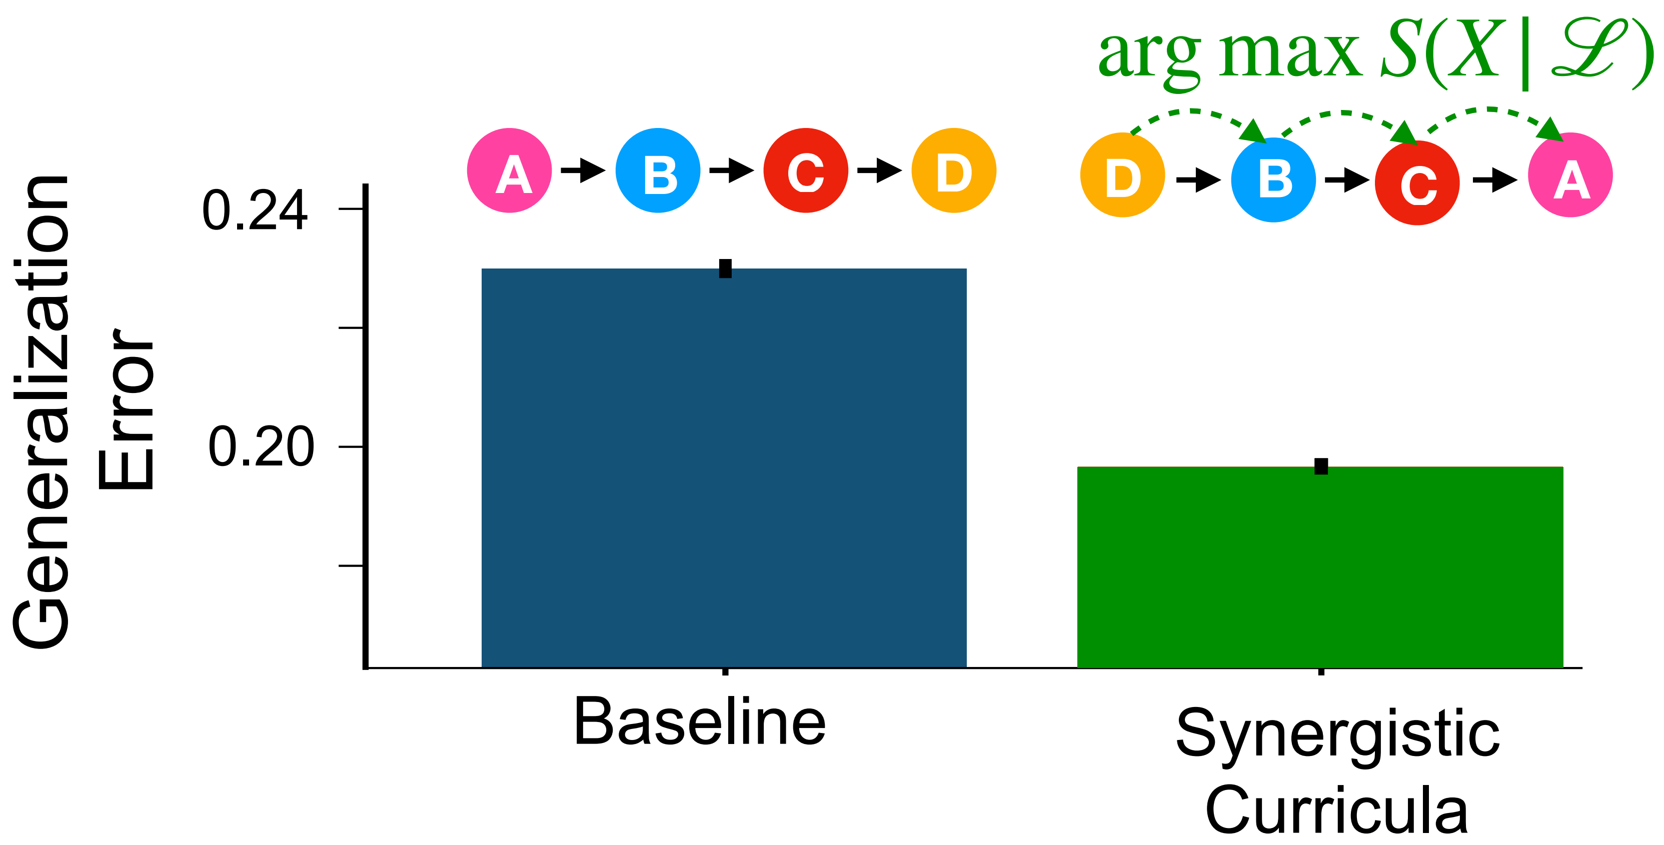

- **Next steps**
  - Test our predictions on human learners

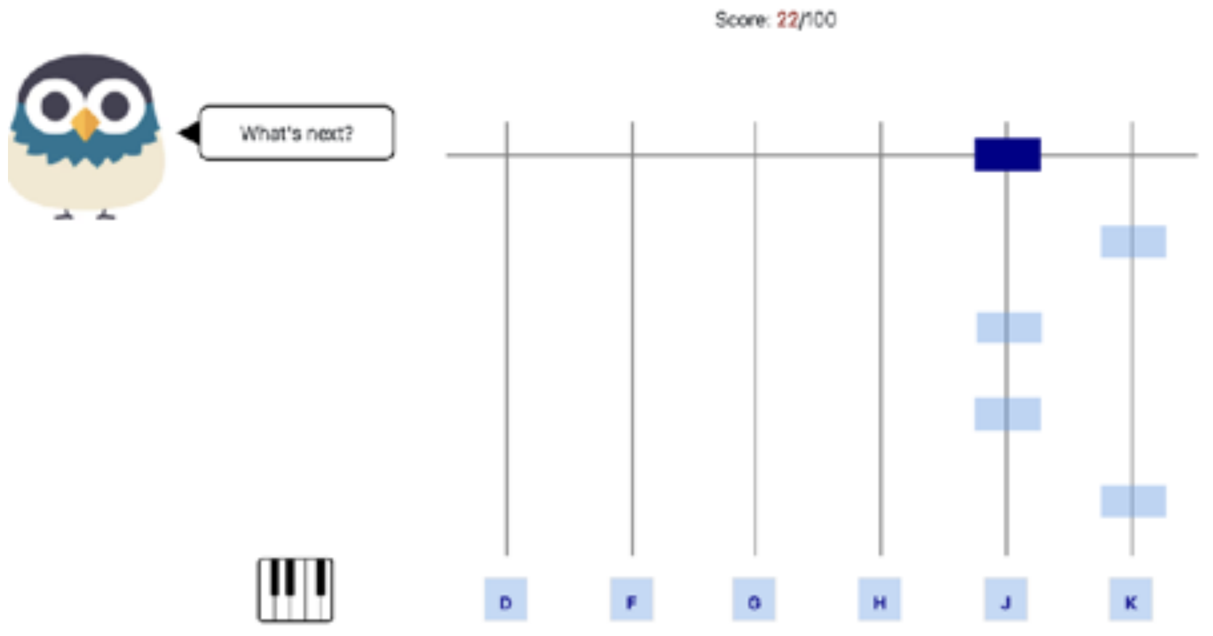

Learning as Search

- A big part of what makes symbolic AI difficult is ***search***
  - Representing relations between all possible symbols creates a combinatorial explosion
  - There are (typically) no gradients for symbolic representations
- Learning can thus be understood as a search problem
  - Finding which rules/programs capture data
  - Finding which hypotheses to test
- One of the major contributions of symbolic AI research was developing search algorithms
  - A*
  - Montecarlo Tree Search

A* Heuristic Search

- One of the most popular methods for path-finding and search over graphs (Hart et al., 1968)
- Expand the path by choosing candidate node n that minimizes cost function $f(n) = g(n) + h(n)$
  - *Keep the current path short*: $g(n)$ is the *cost of the path so far* from the start to n
    - Costs can also represent complexity (i.e., the number of symbolic operations)
  - *Move towards the goal*: $h(n)$ is a **heuristic** that estimates the cost of the cheapest remaining path from n to the goal (often Euclidean distance)
    - The heuristic avoids calculating the actual remaining cost to the goal, which is very costly
- More efficient than **backwards induction**, but intractable for any interesting program induction problems

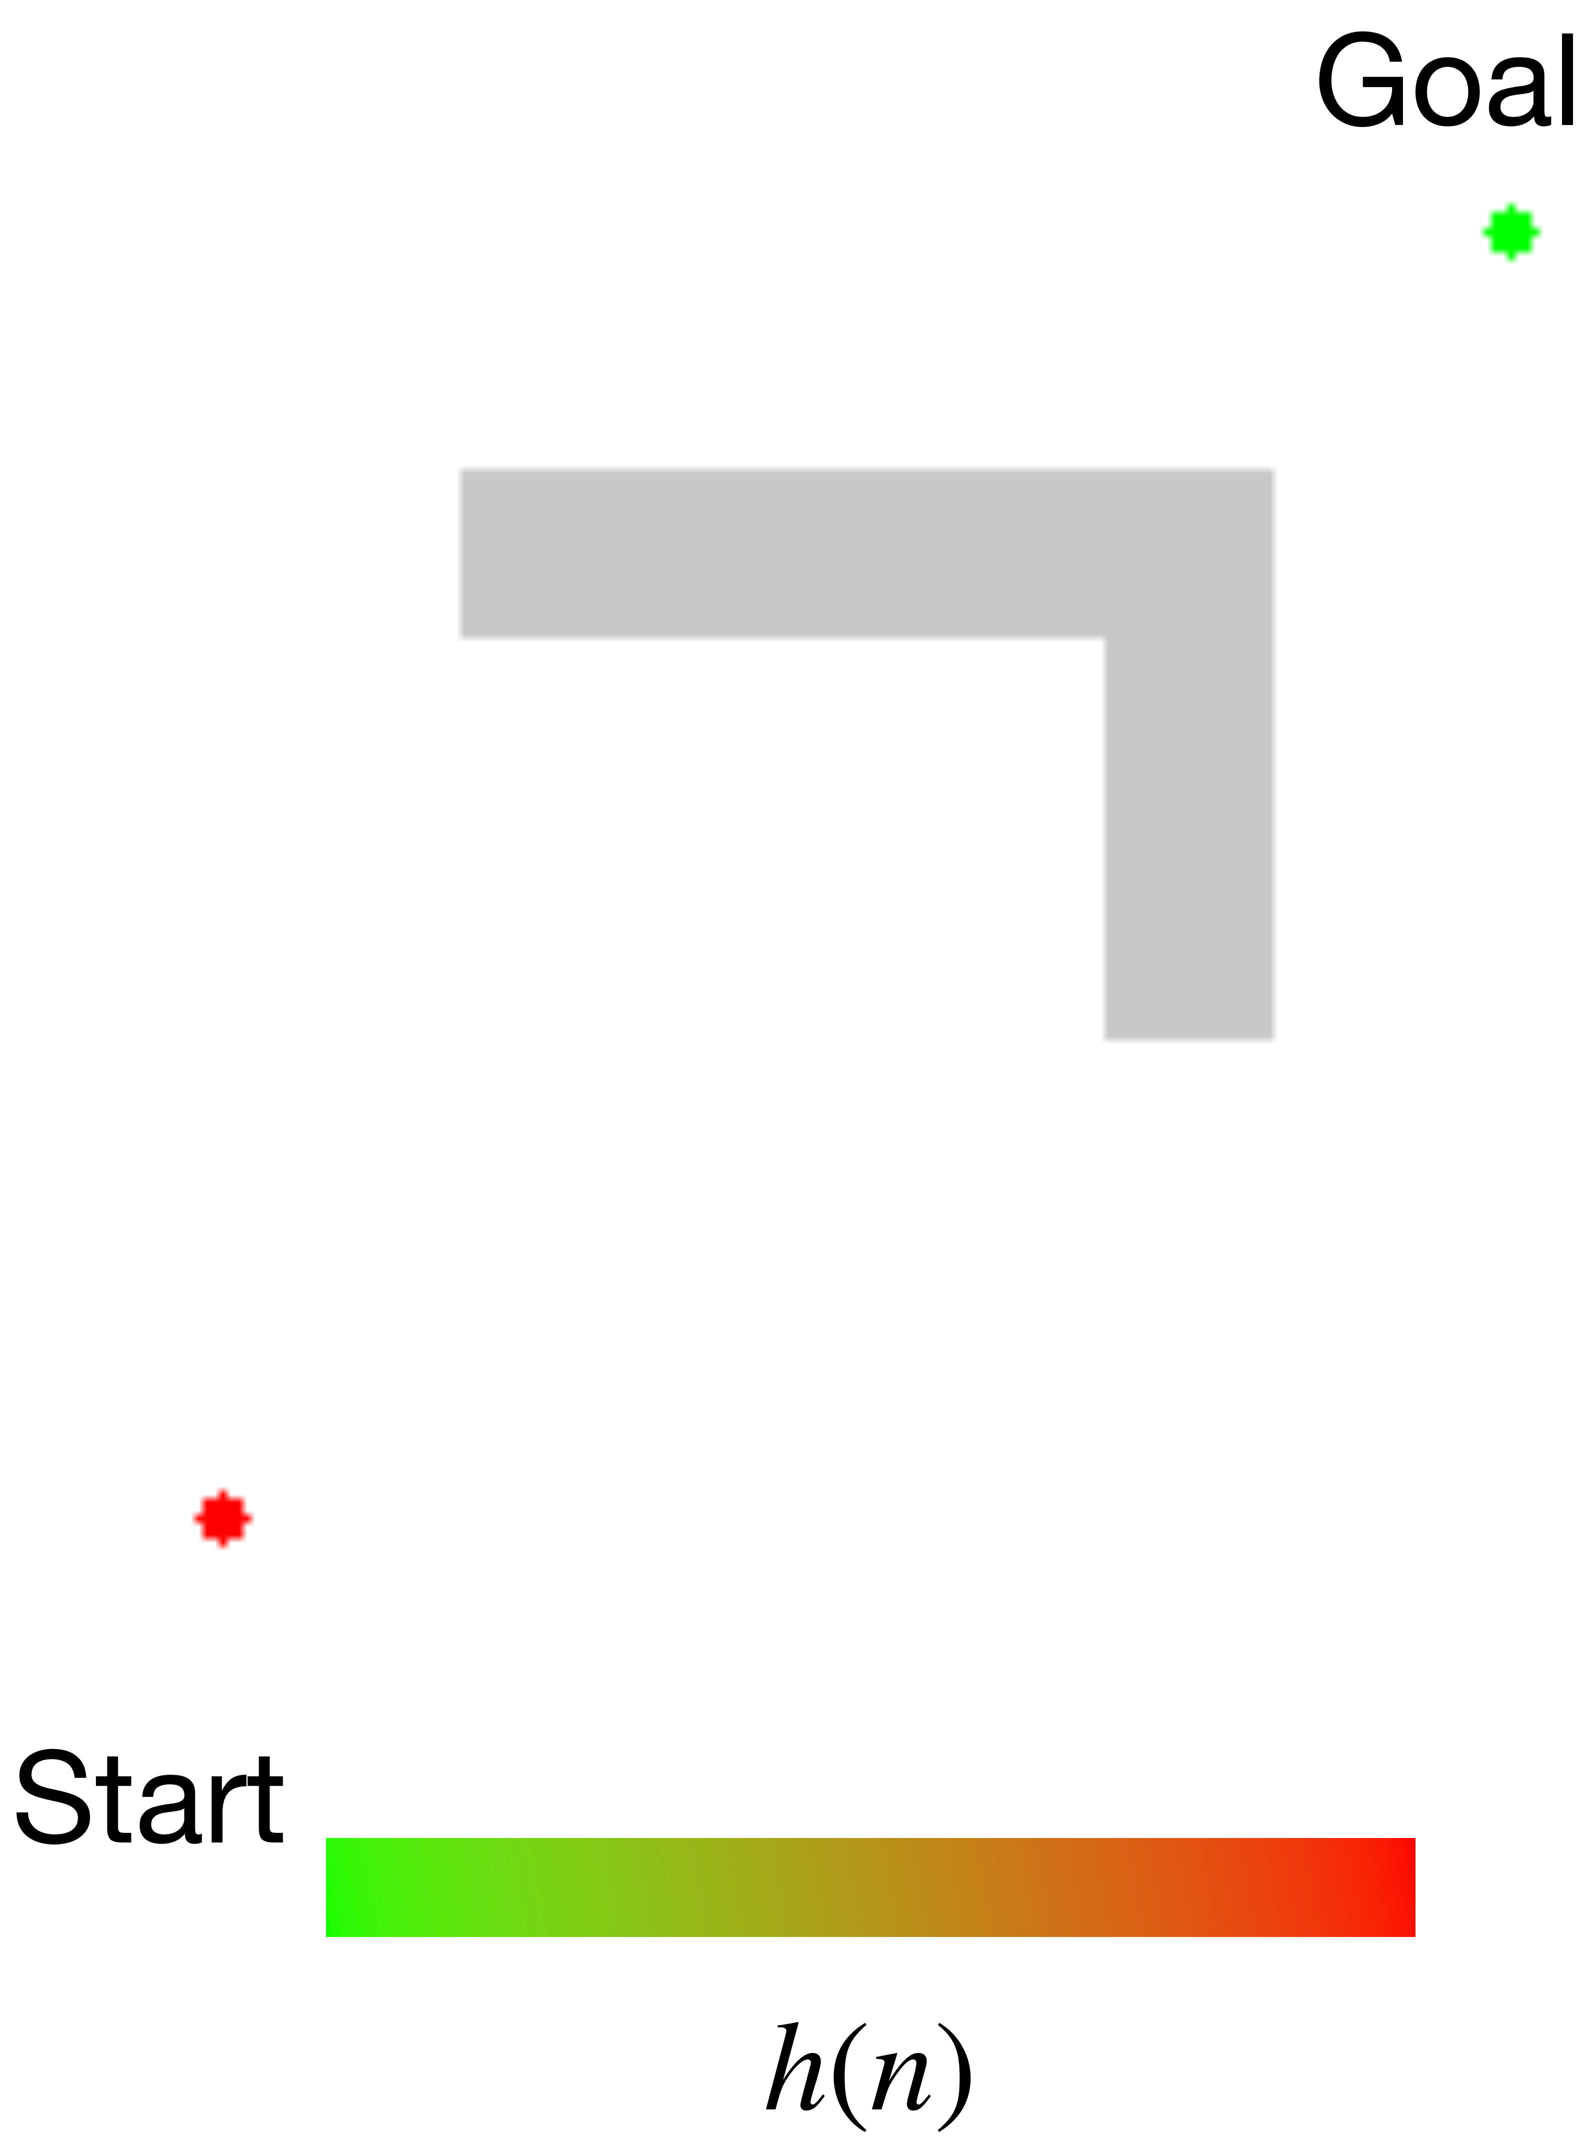

A* Heuristic Search

- One of the most popular methods for path-finding and search over graphs (Hart et al., 1968)
- Expand the path by choosing candidate node n that minimizes cost function $f(n) = g(n) + h(n)$
  - *Keep the current path short*: $g(n)$ is the *cost of the path so far* from the start to n
    - Costs can also represent complexity (i.e., the number of symbolic operations)
  - *Move towards the goal*: $h(n)$ is a **heuristic** that estimates the cost of the cheapest remaining path from n to the goal (often Euclidean distance)
    - The heuristic avoids calculating the actual remaining cost to the goal, which is very costly
- More efficient than **backwards induction**, but intractable for any interesting program induction problems

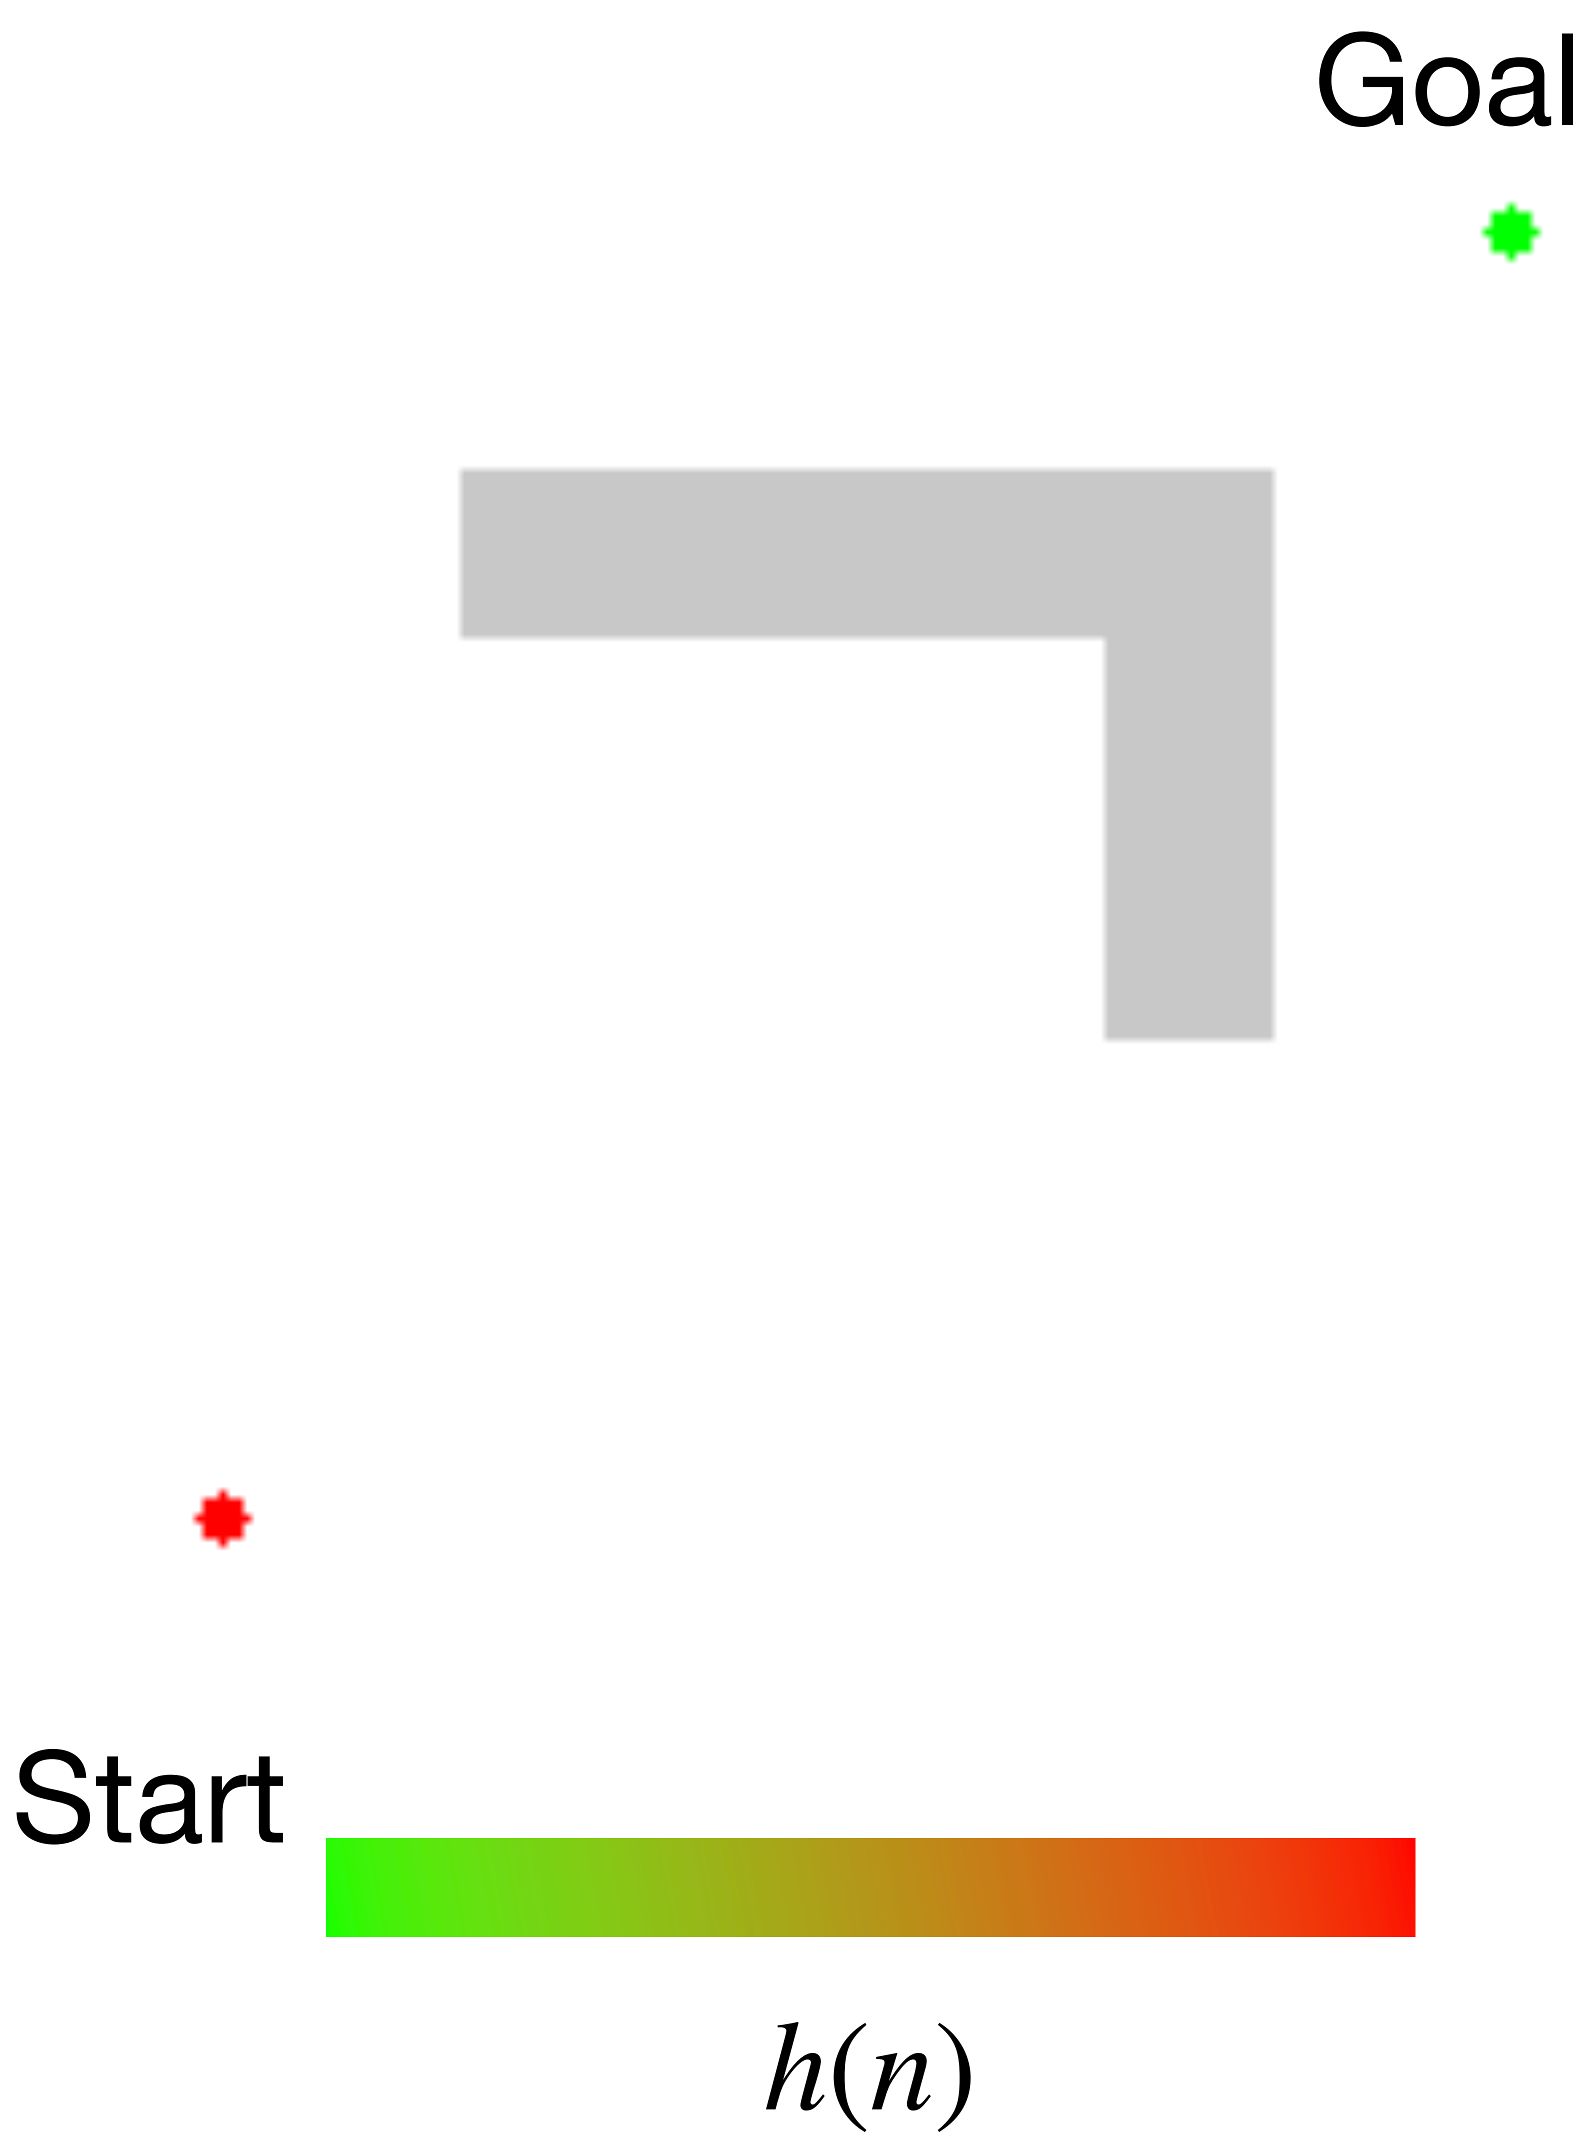

A* Heuristic Search

- One of the most popular methods for path-finding and search over graphs (Hart et al., 1968)
- Expand the path by choosing candidate node n that minimizes cost function $f(n) = g(n) + h(n)$
  - *Keep the current path short*: $g(n)$ is the *cost of the path so far* from the start to n
    - Costs can also represent complexity (i.e., the number of symbolic operations)
  - *Move towards the goal*: $h(n)$ is a **heuristic** that estimates the cost of the cheapest remaining path from n to the goal (often Euclidean distance)
    - The heuristic avoids calculating the actual remaining cost to the goal, which is very costly
- More efficient than **backwards induction**, but intractable for any interesting program induction problems

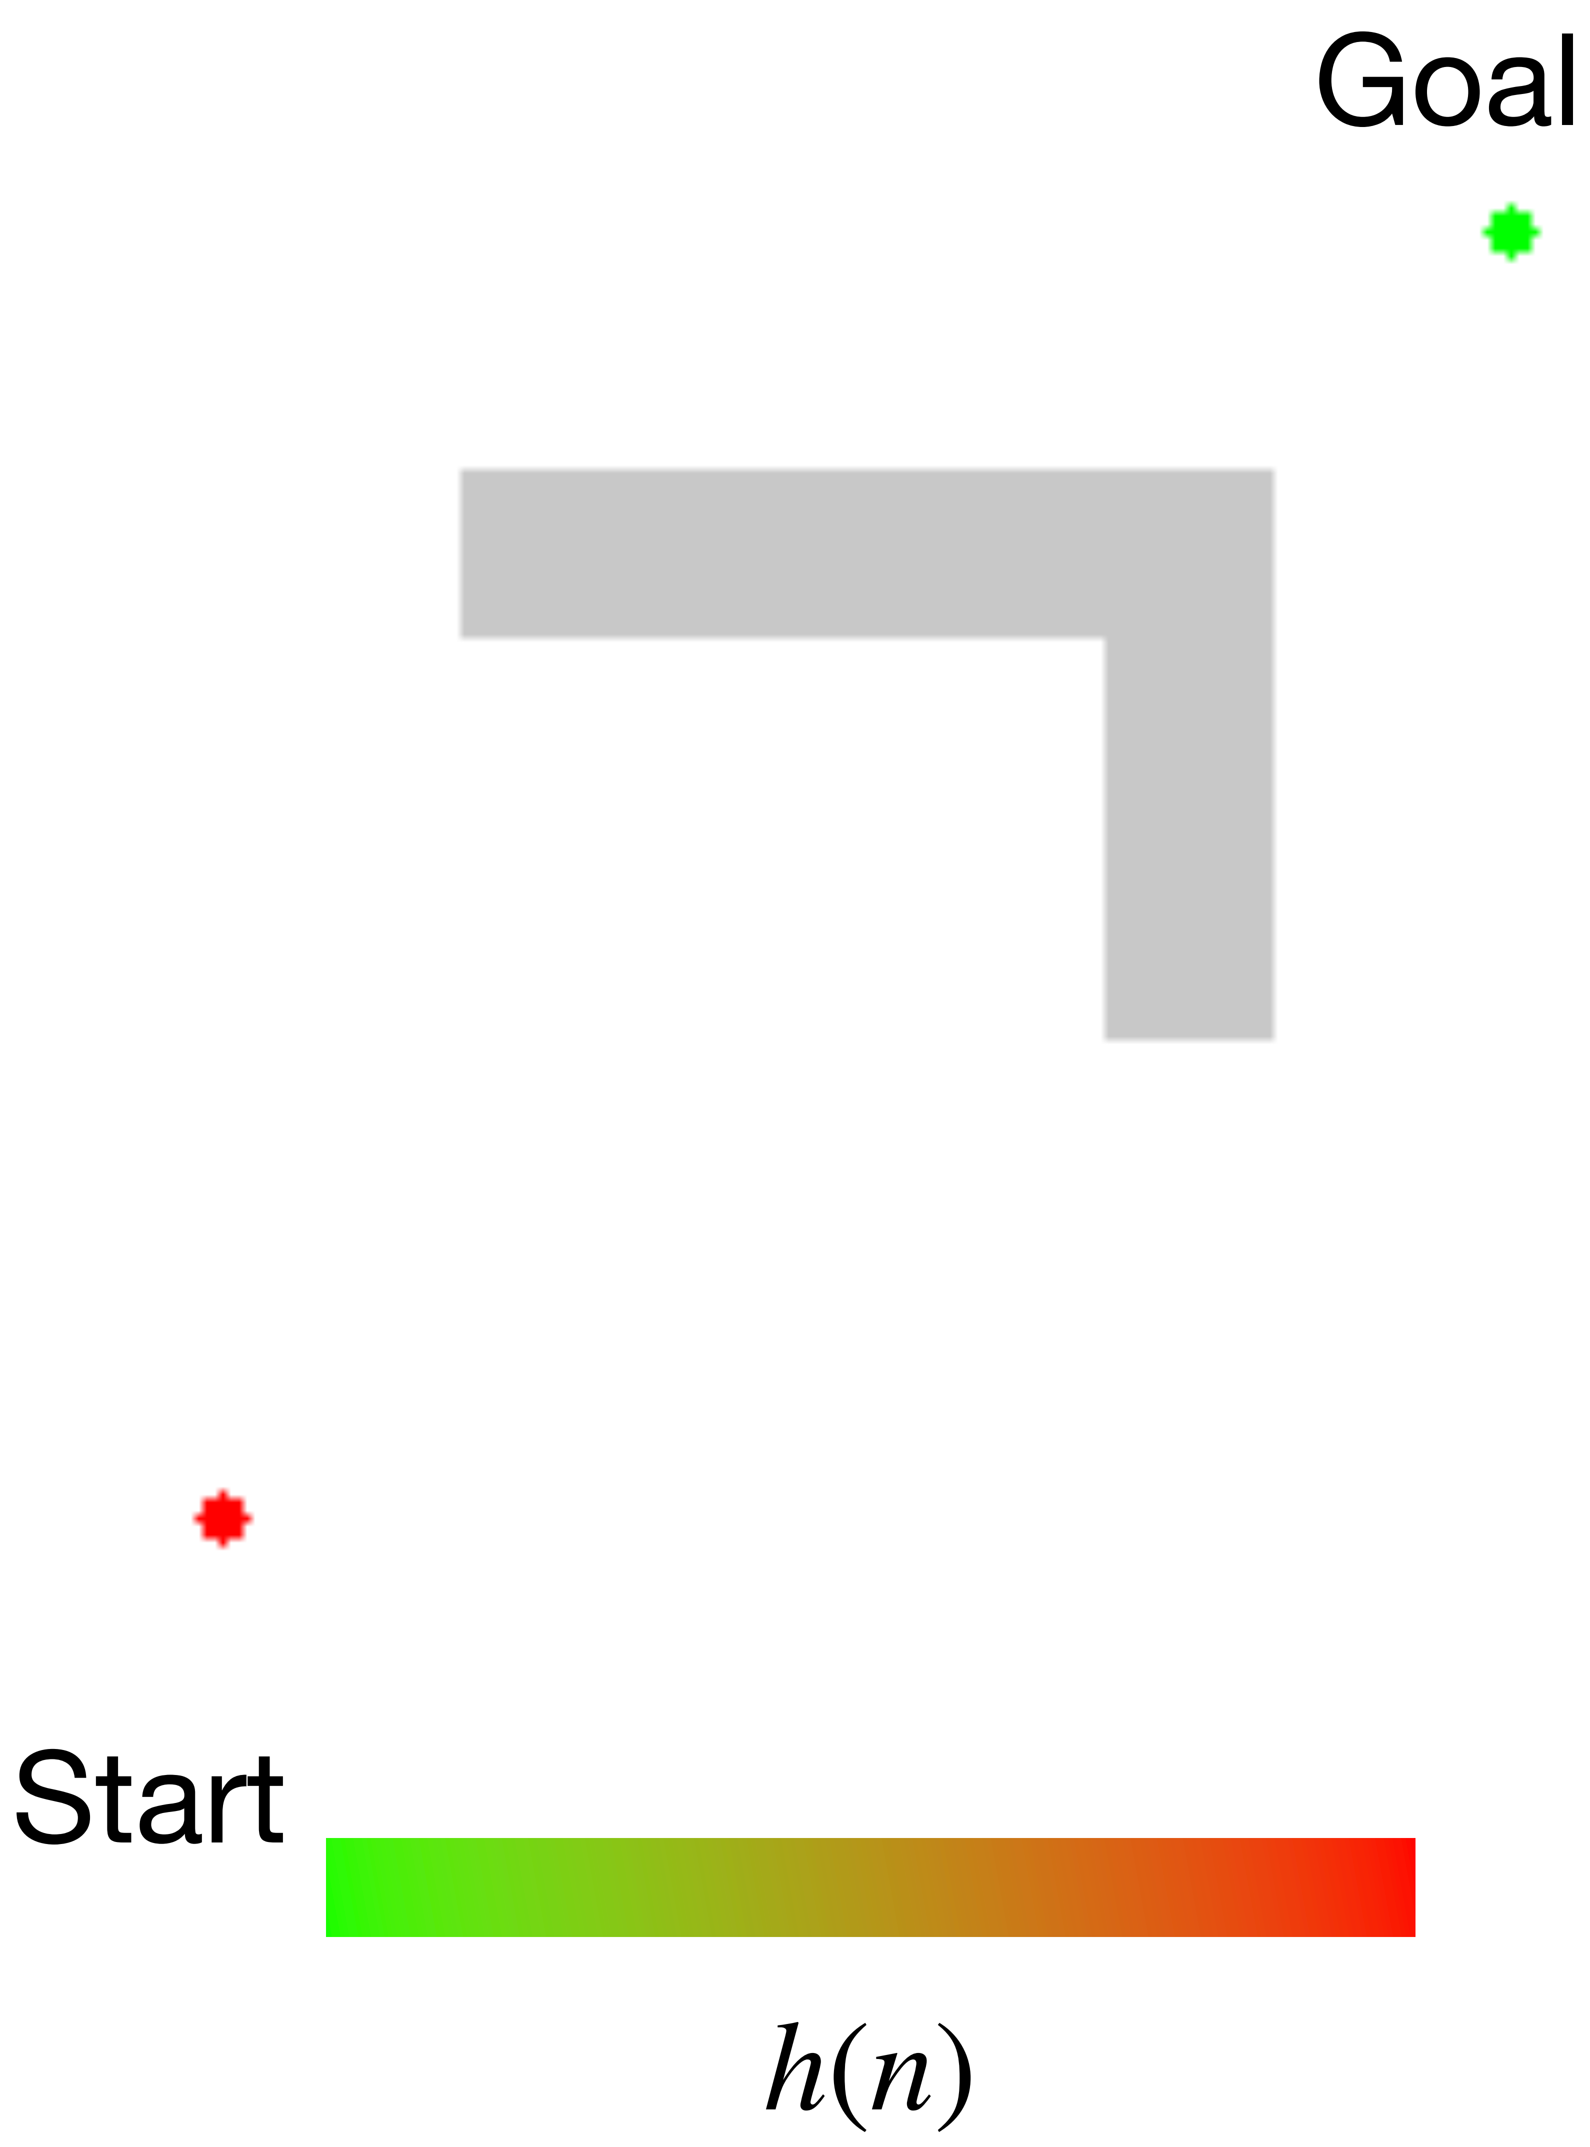

Heuristic: a problem-solving strategy or method that is not guaranteed to find the optimal solution, but is designed to find a satisfactory solution in a reasonable amount of time

A* Heuristic Search

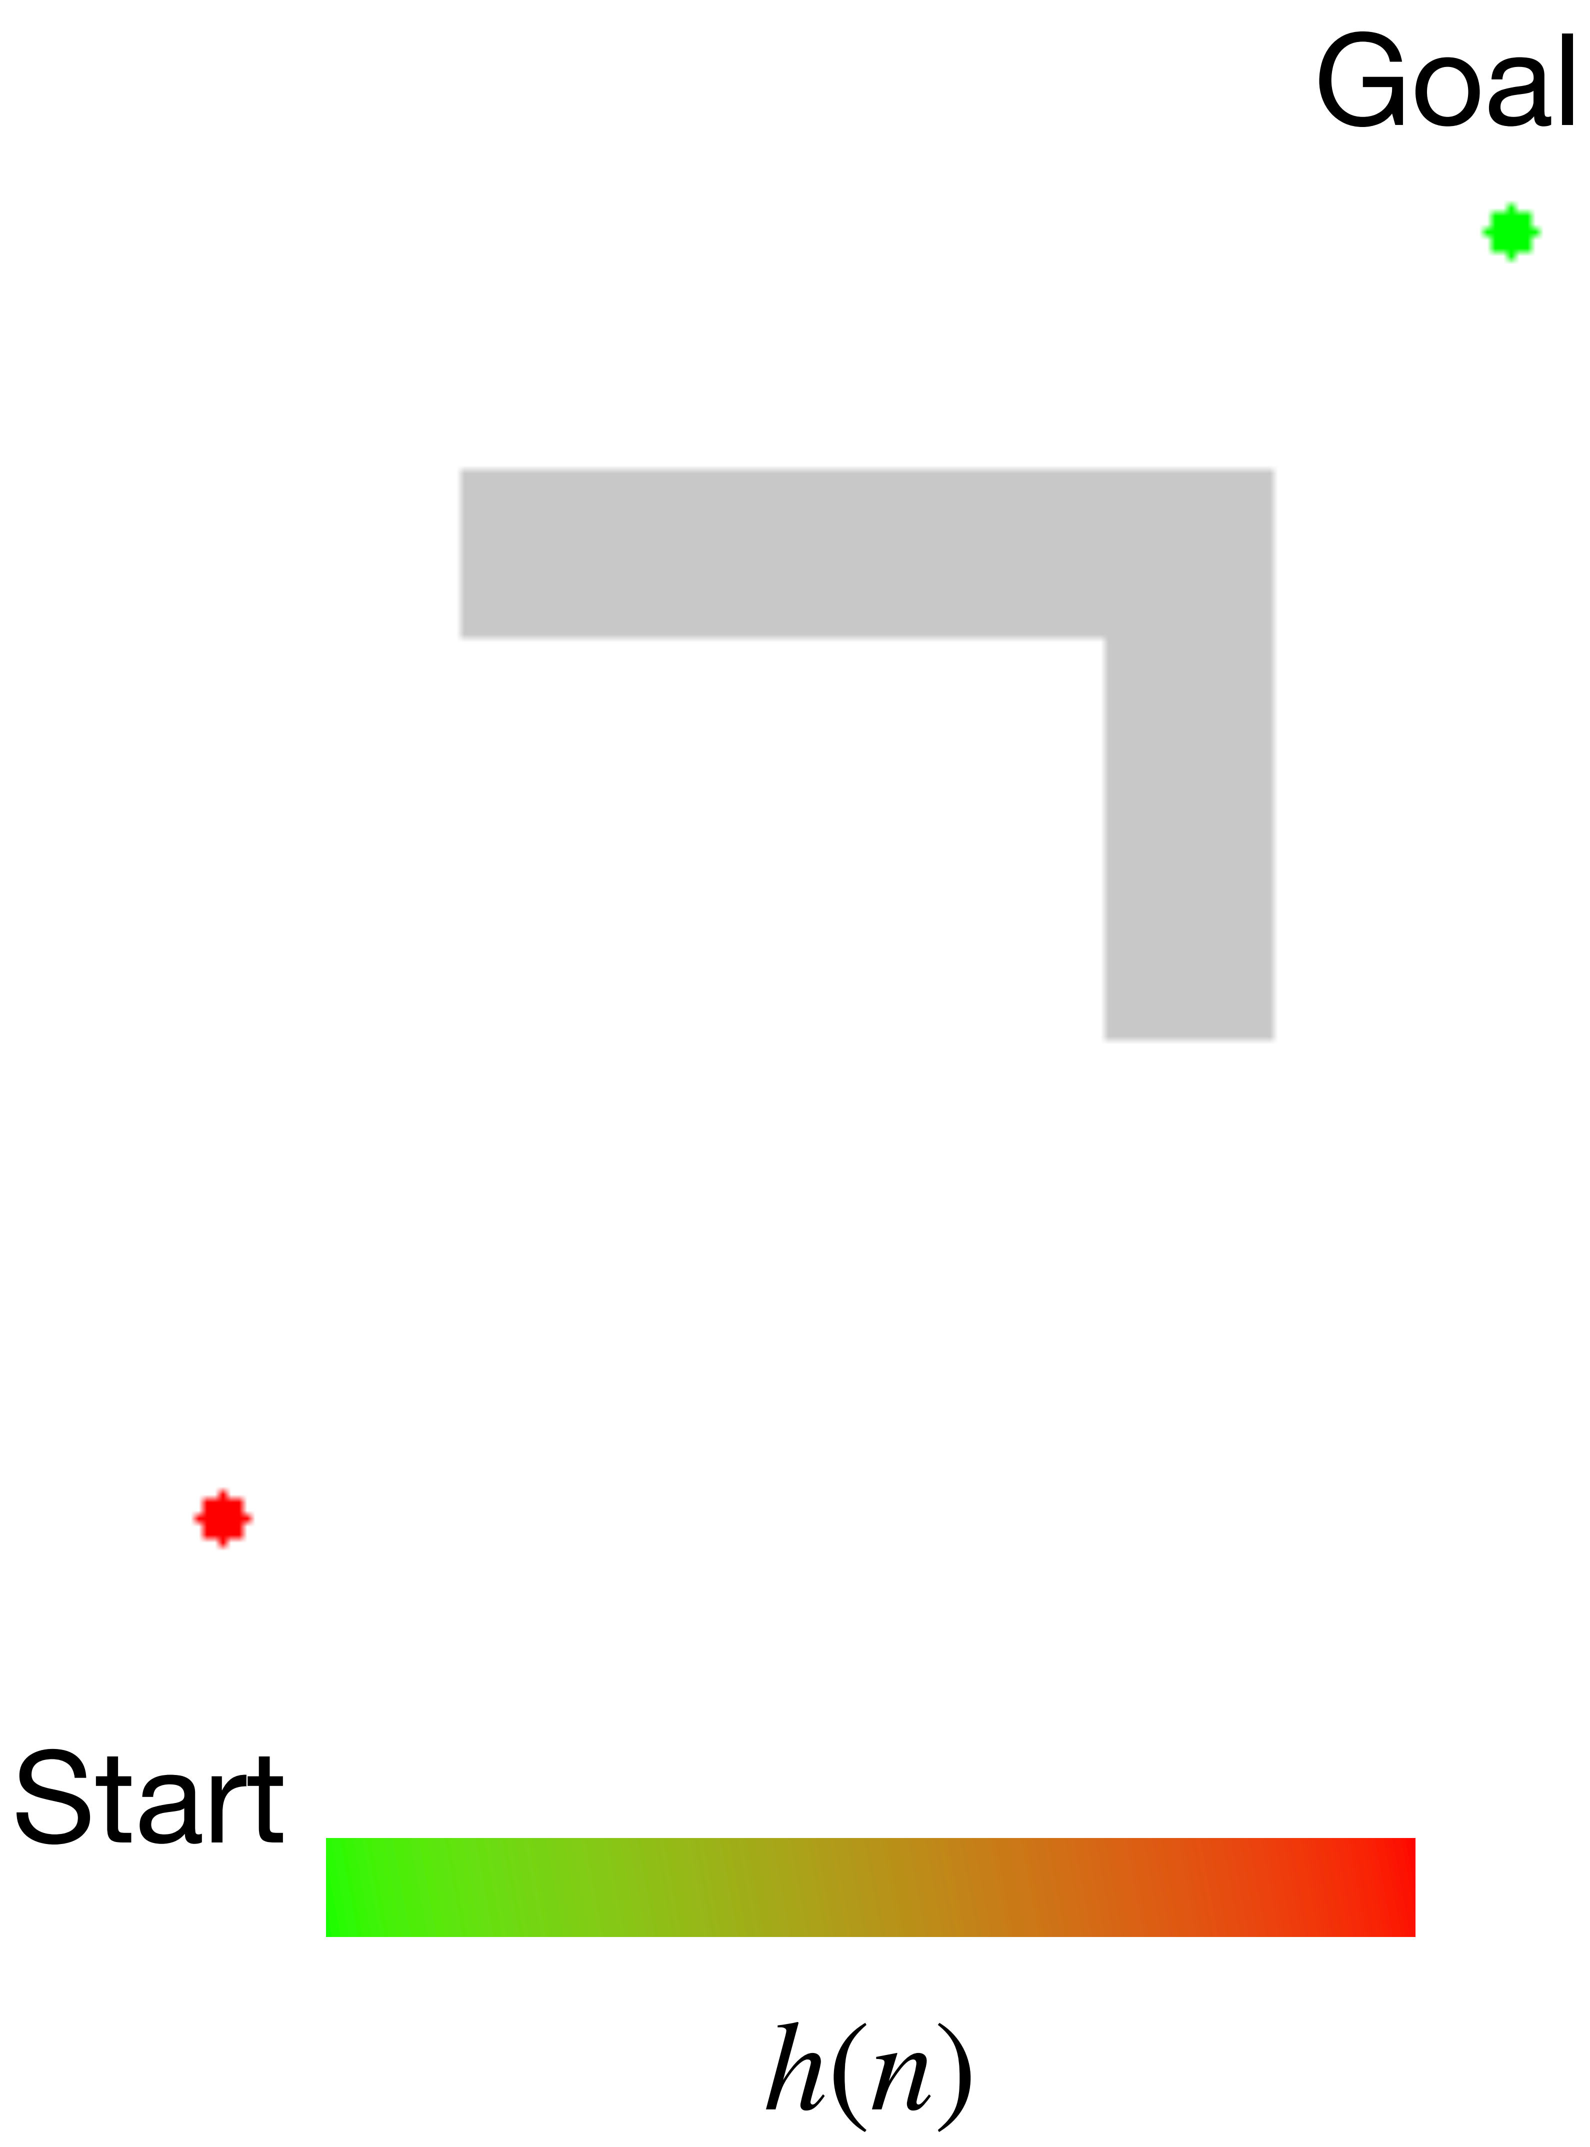

- One of the most popular methods for path-finding and search over graphs (Hart et al., 1968)
- Expand the path by choosing candidate node n that minimizes cost function $f(n) = g(n) + h(n)$
  - *Keep the current path short*: $g(n)$ is the *cost of the path so far* from the start to n
    - Costs can also represent complexity (i.e., the number of symbolic operations)
  - *Move towards the goal*: $h(n)$ is a **heuristic** that estimates the cost of the cheapest remaining path from n to the goal (often Euclidean distance)
    - The heuristic avoids calculating the actual remaining cost to the goal, which is very costly
- More efficient than **backwards induction**, but intractable for any interesting program induction problems

Heuristic: a problem-solving strategy or method that is not guaranteed to find the optimal solution, but is designed to find a satisfactory solution in a reasonable amount of time

Backwards induction: determining a sequence of optimal choices by reasoning from the endpoint of a problem back to the beginning

Monte Carlo Tree Search

- A key mechanism in AlphaGo (Silver et al., 2016) and other modern RL algorithms
- **Select** nodes for expansion (often using a heuristic based on reward + *information gain*)
- **Expand** node and perform **simulations**
- **Backpropogate** the value of the child to the parent node
    - This allows us to save a heuristic value for the parent node based on previous simulations over the children

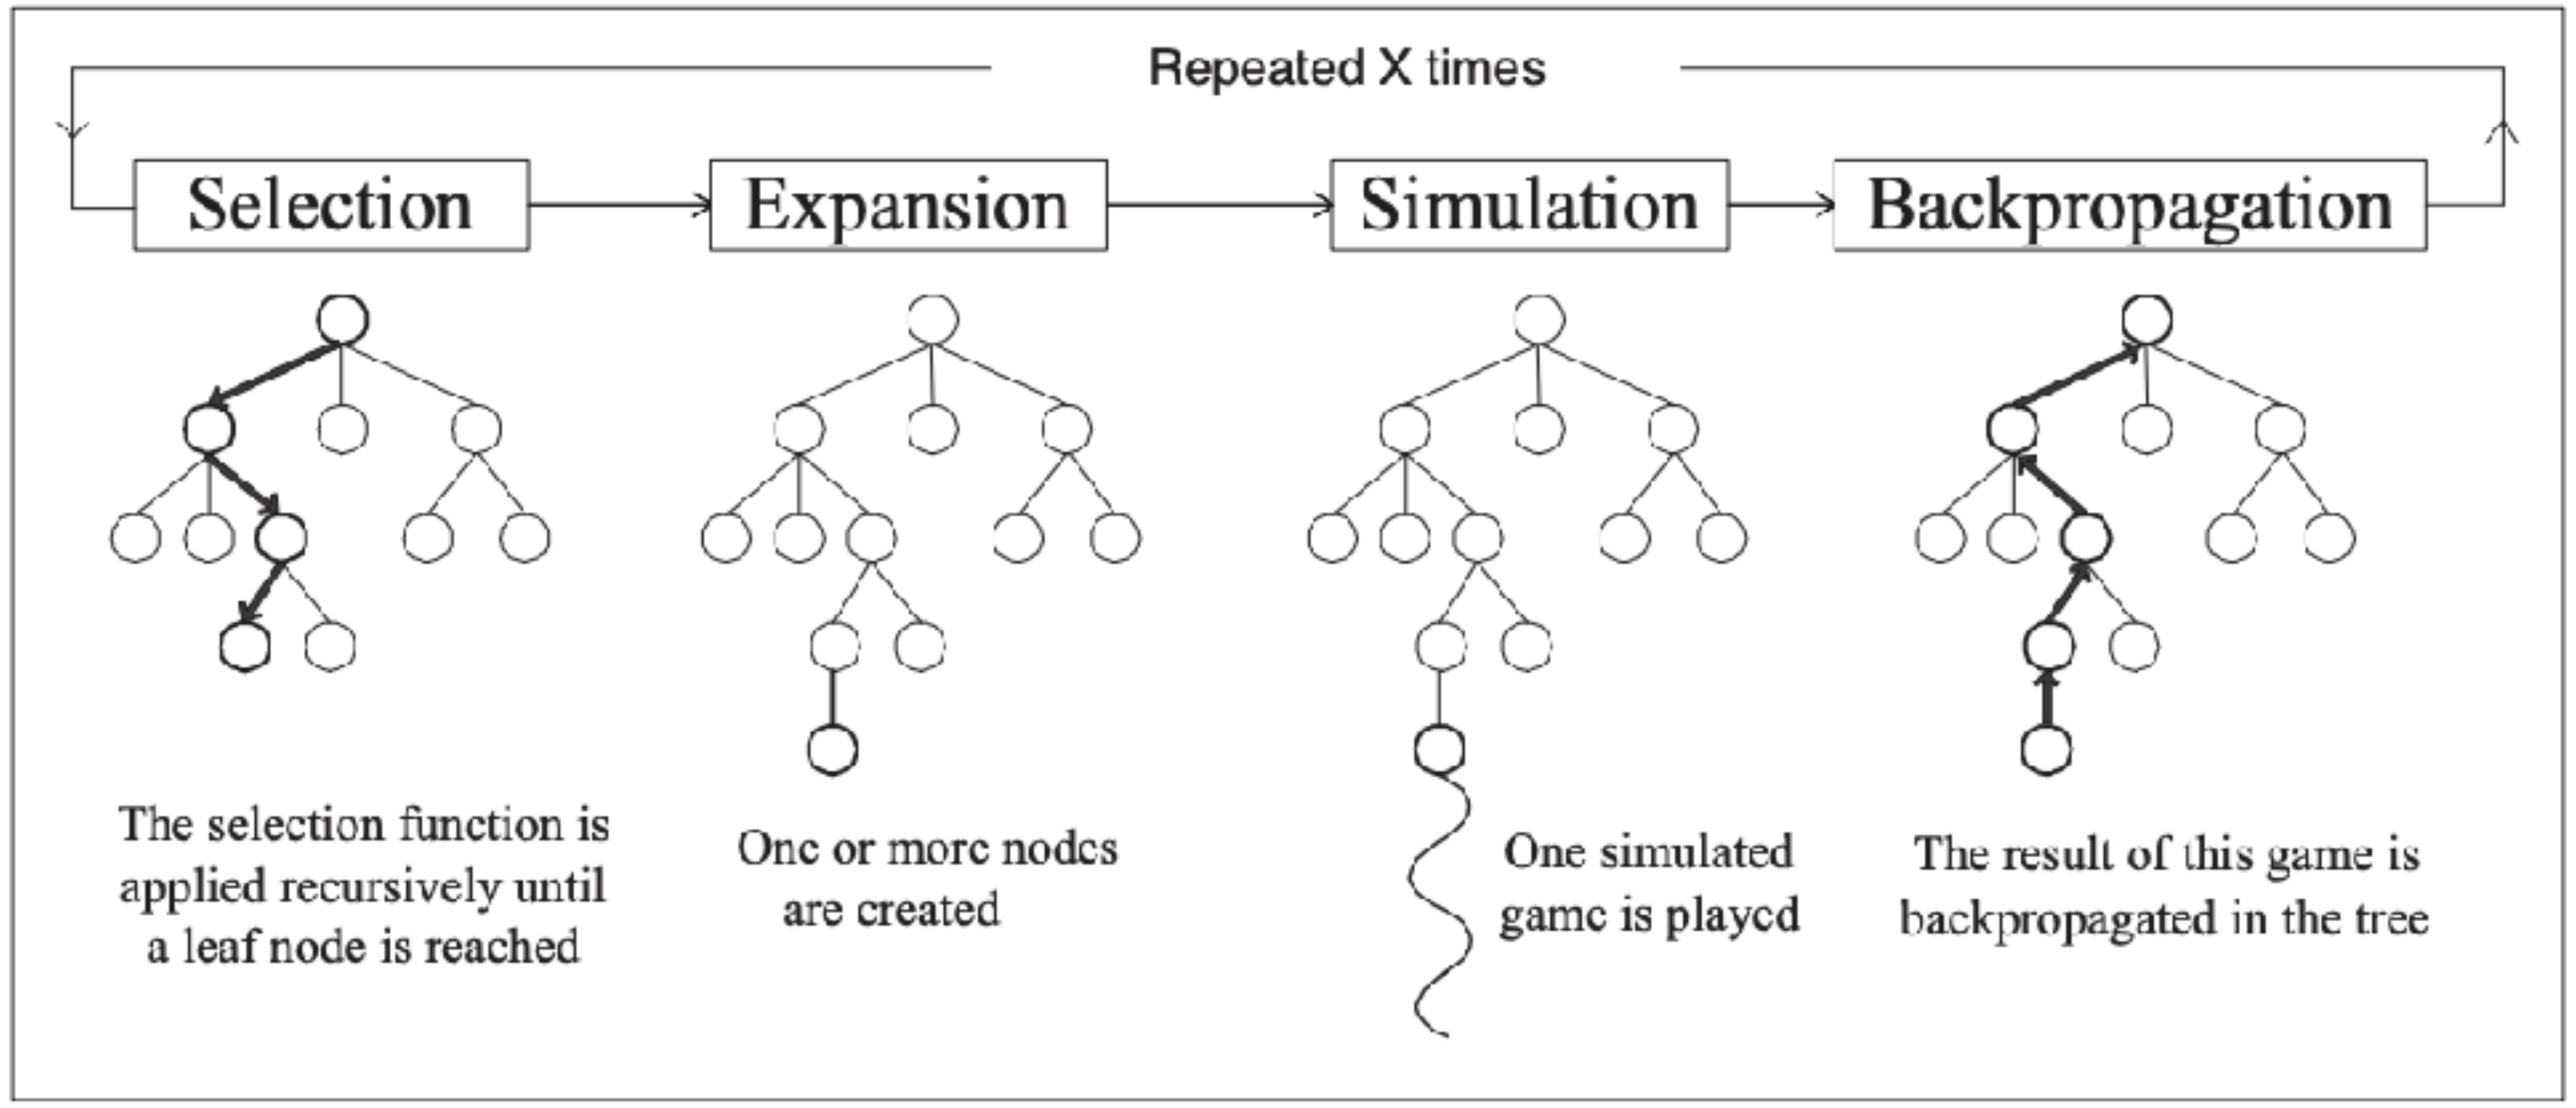

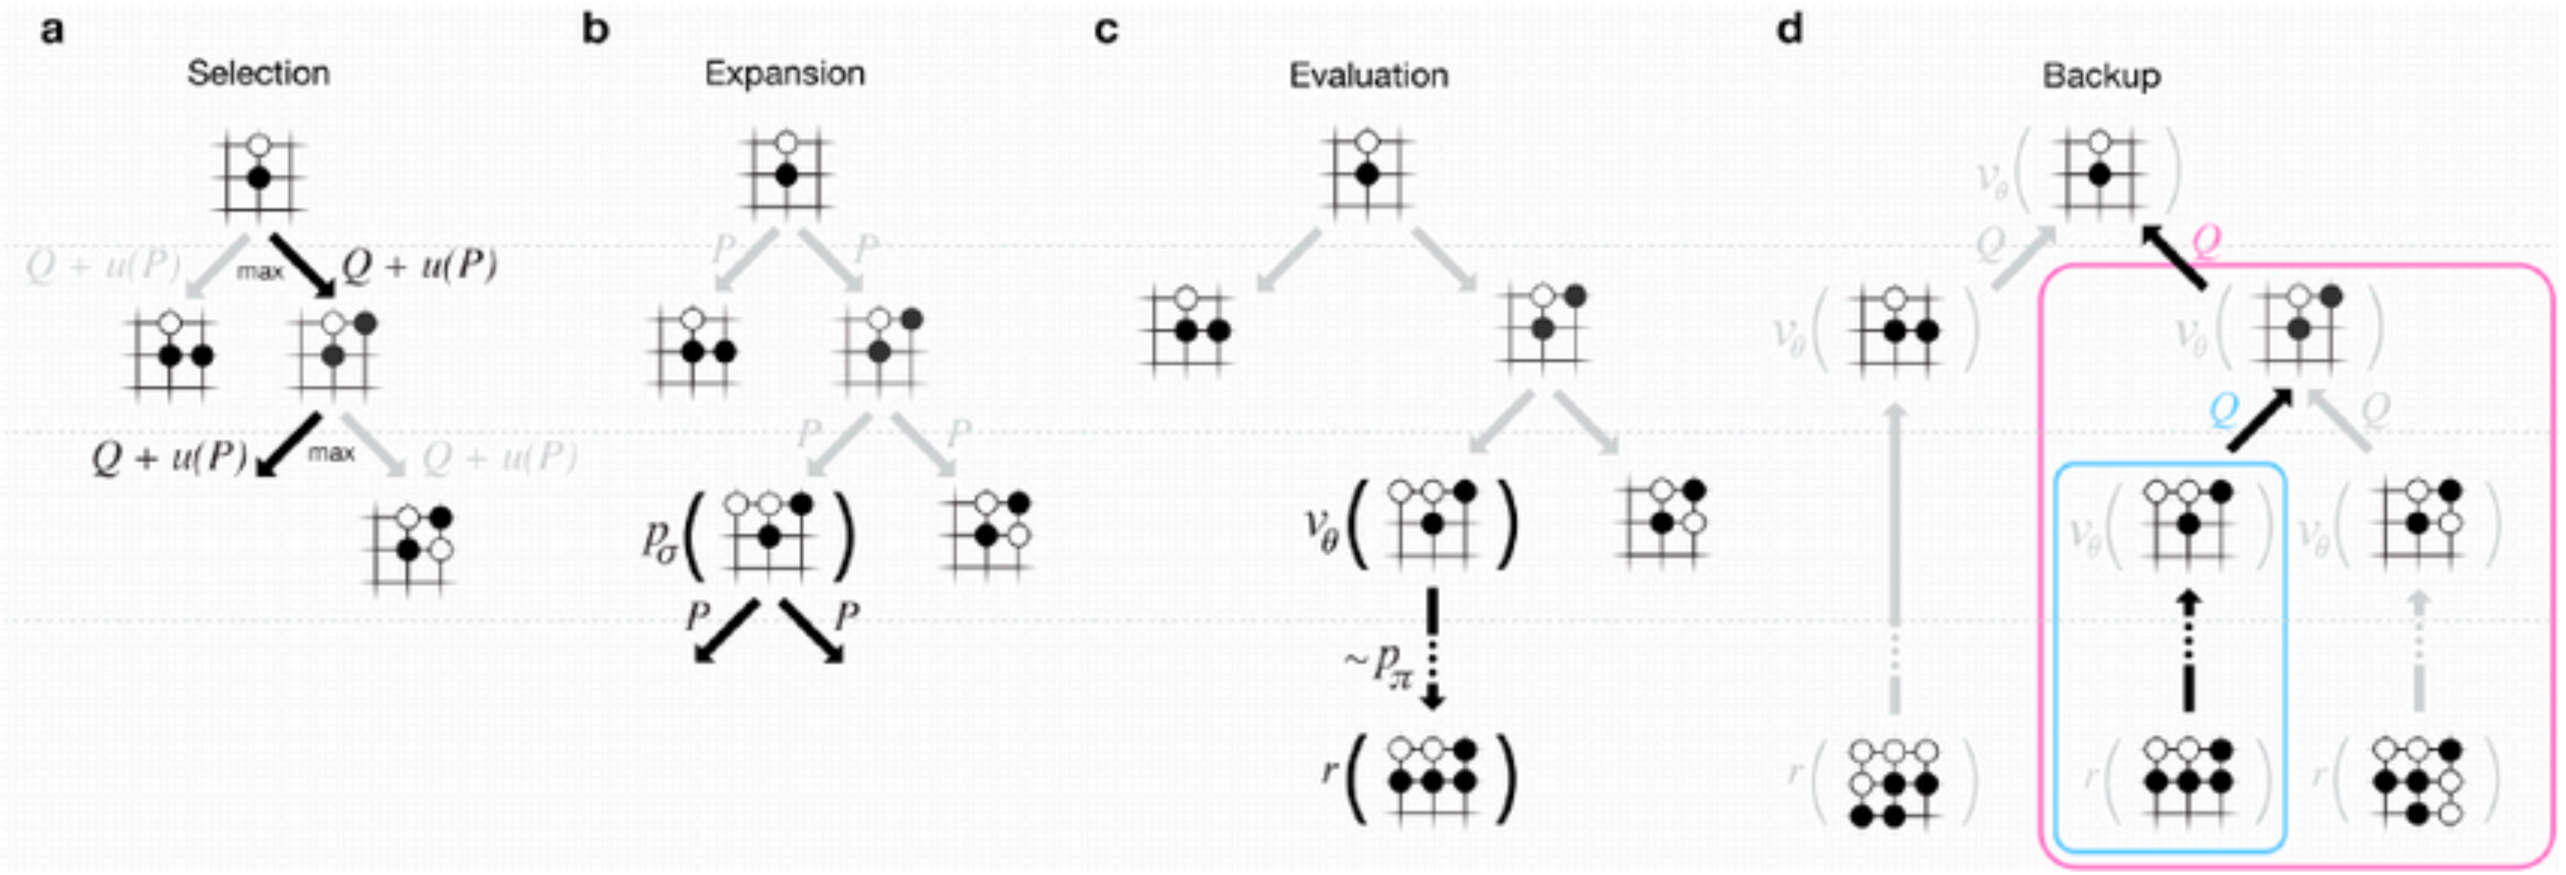

Monte Carlo Tree Search

- A key mechanism in AlphaGo (Silver et al., 2016) and other modern RL algorithms
- **Select** nodes for expansion (often using a heuristic based on reward + *information gain*)
- **Expand** node and perform **simulations**
- **Backpropogate** the value of the child to the parent node
  - This allows us to save a heuristic value for the parent node based on previous simulations over the children

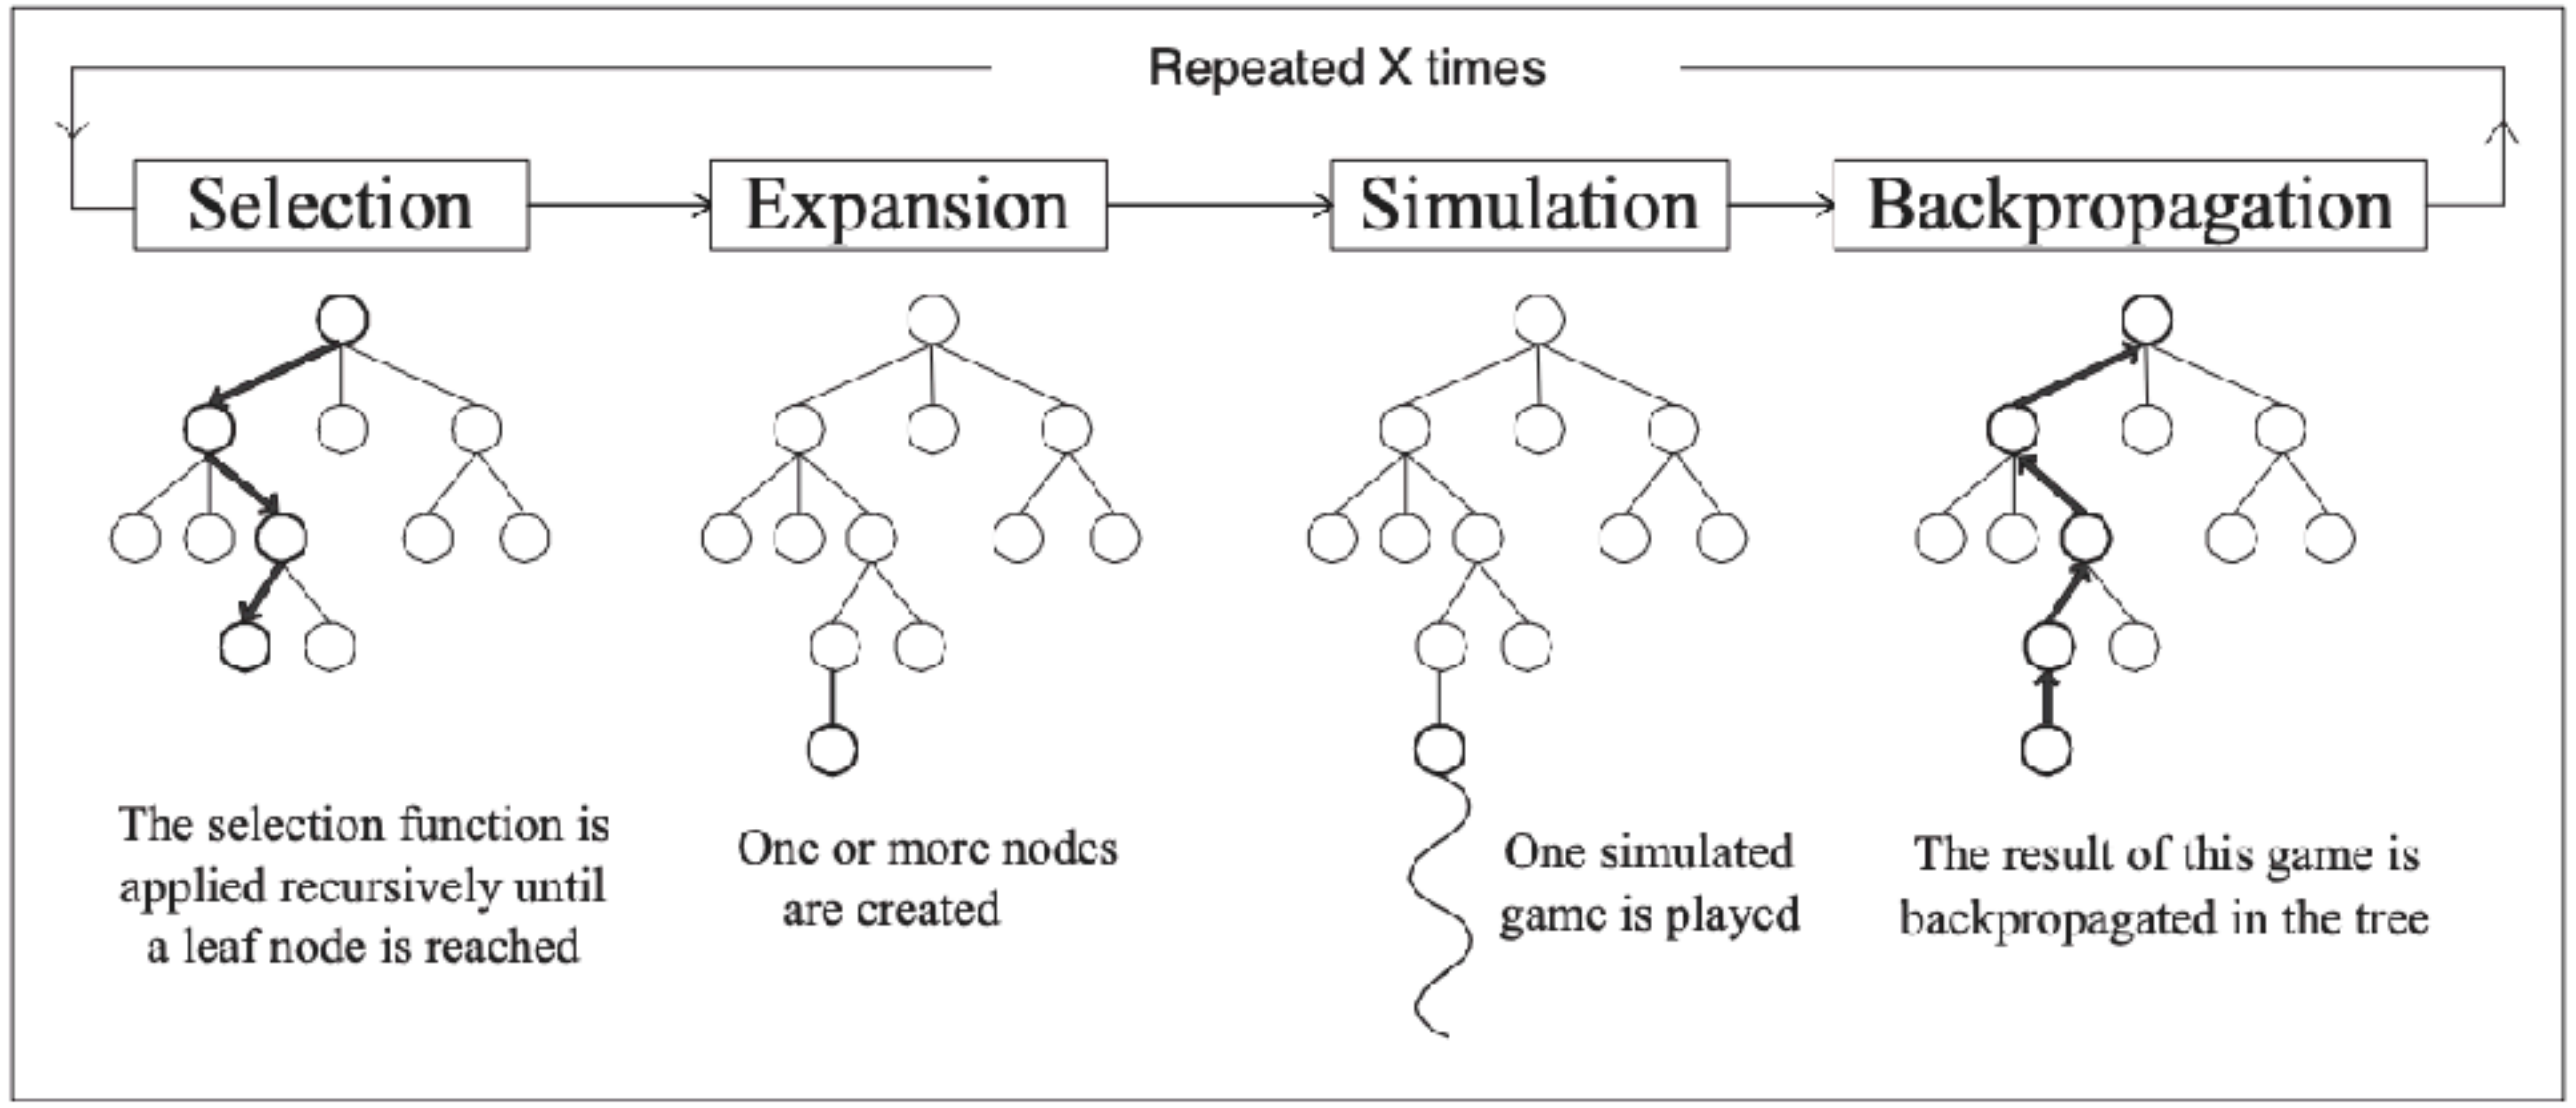

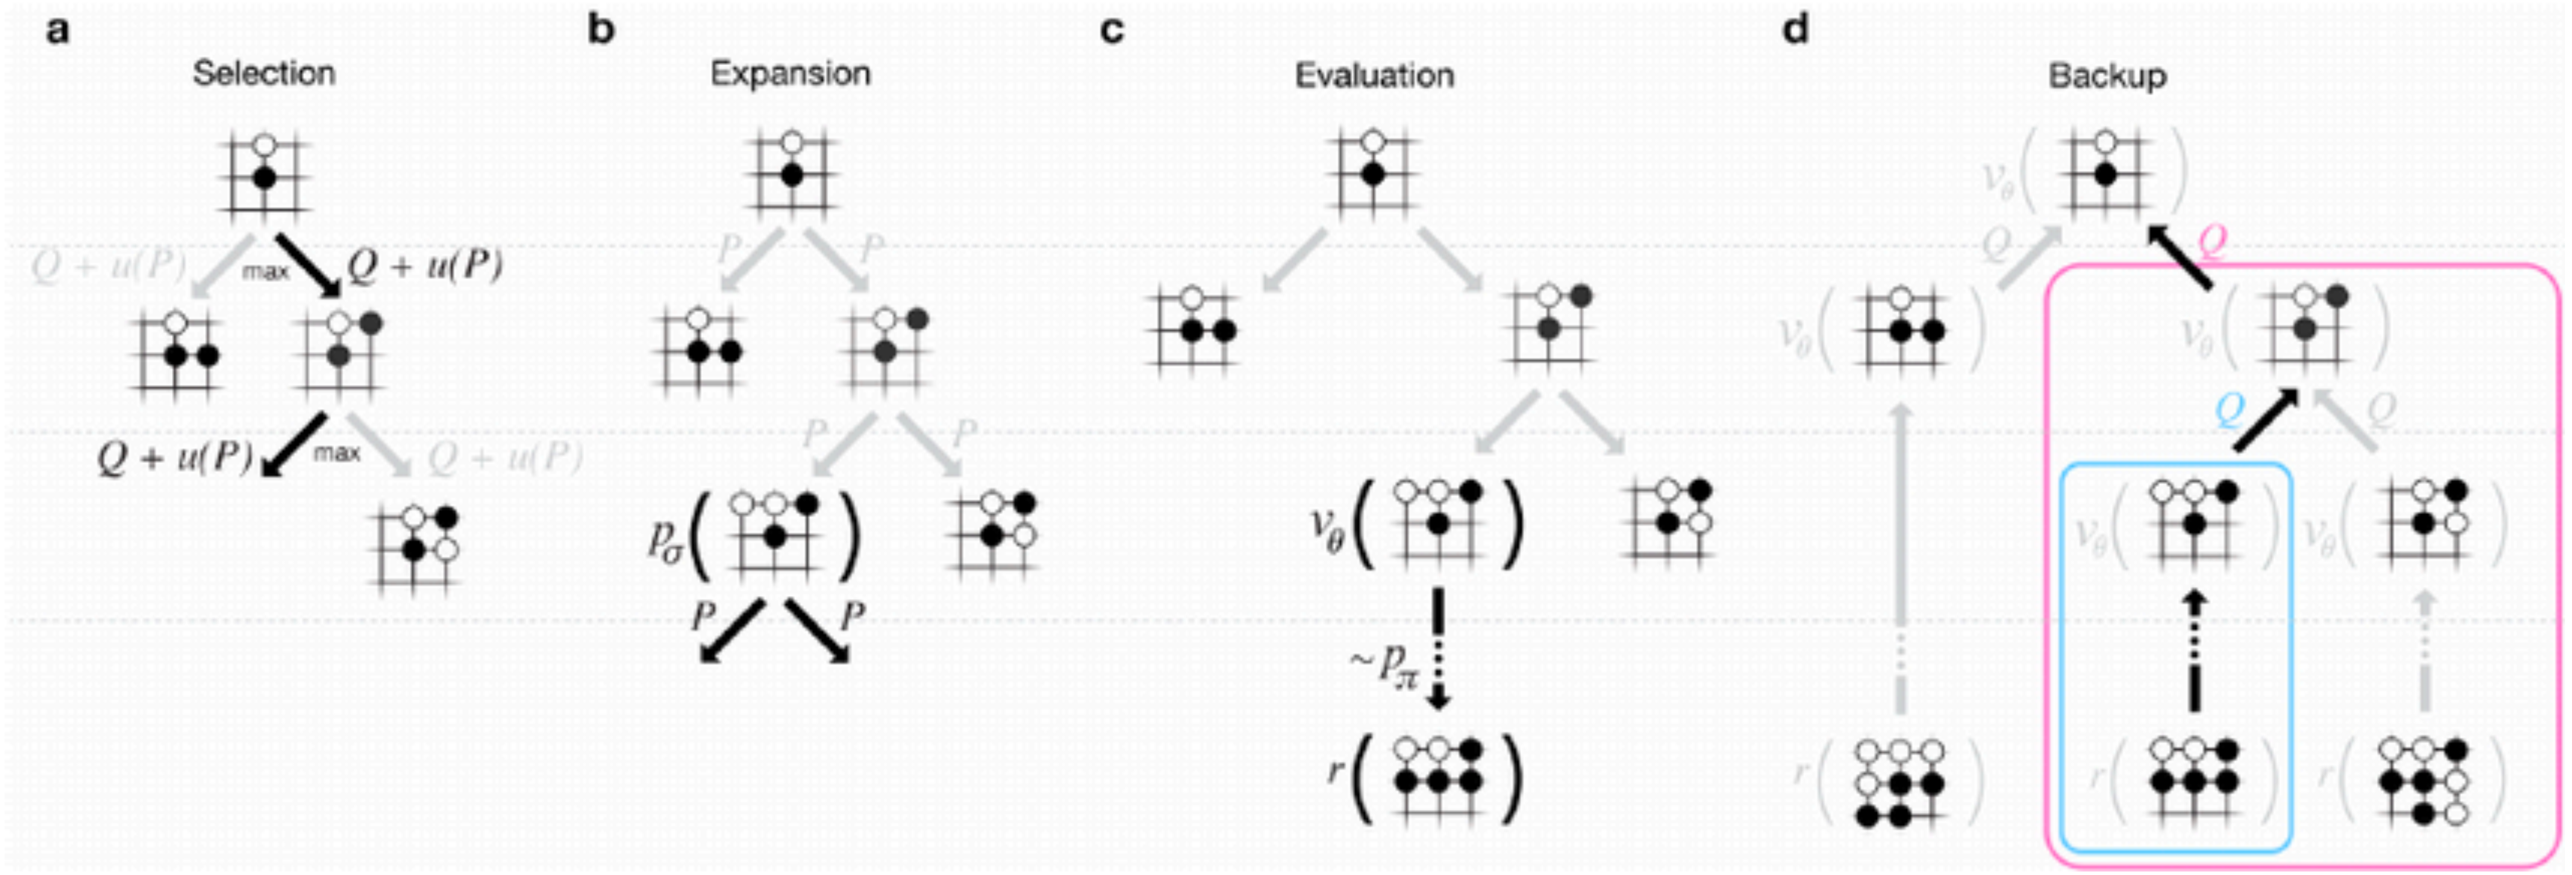

Information gain: The amount of information gained by an observation (i.e., expanding a node). Often approximated using count-based methods:
↑ info gain ∝ ↓fewer visits

Symbolic AI: Summary

- **Symbols** and relational **rules** are a powerful tool for describing the world
  - Capture rapid generalization and allow for compositional construction of new representations
  - Explicit formulation of relationships in the world that mirror our own Language of Thought and provides interpretable predictions
- **Learning is difficult** and rules can sometimes be too rigid
  - Compositional hypothesis space leads to a combinatorial explosion of possible symbolic representations, where search can be very costly
  - Learning is often framed as a search problem, where heuristic solutions provide a valuable aid
- **Neurosymbolic AI** might offer the best of both worlds by combining the fast learning of subsymbolic AI (i.e., neural networks) with the powerful abstractions of symbolic AI

5 minute break

Cognitive Maps

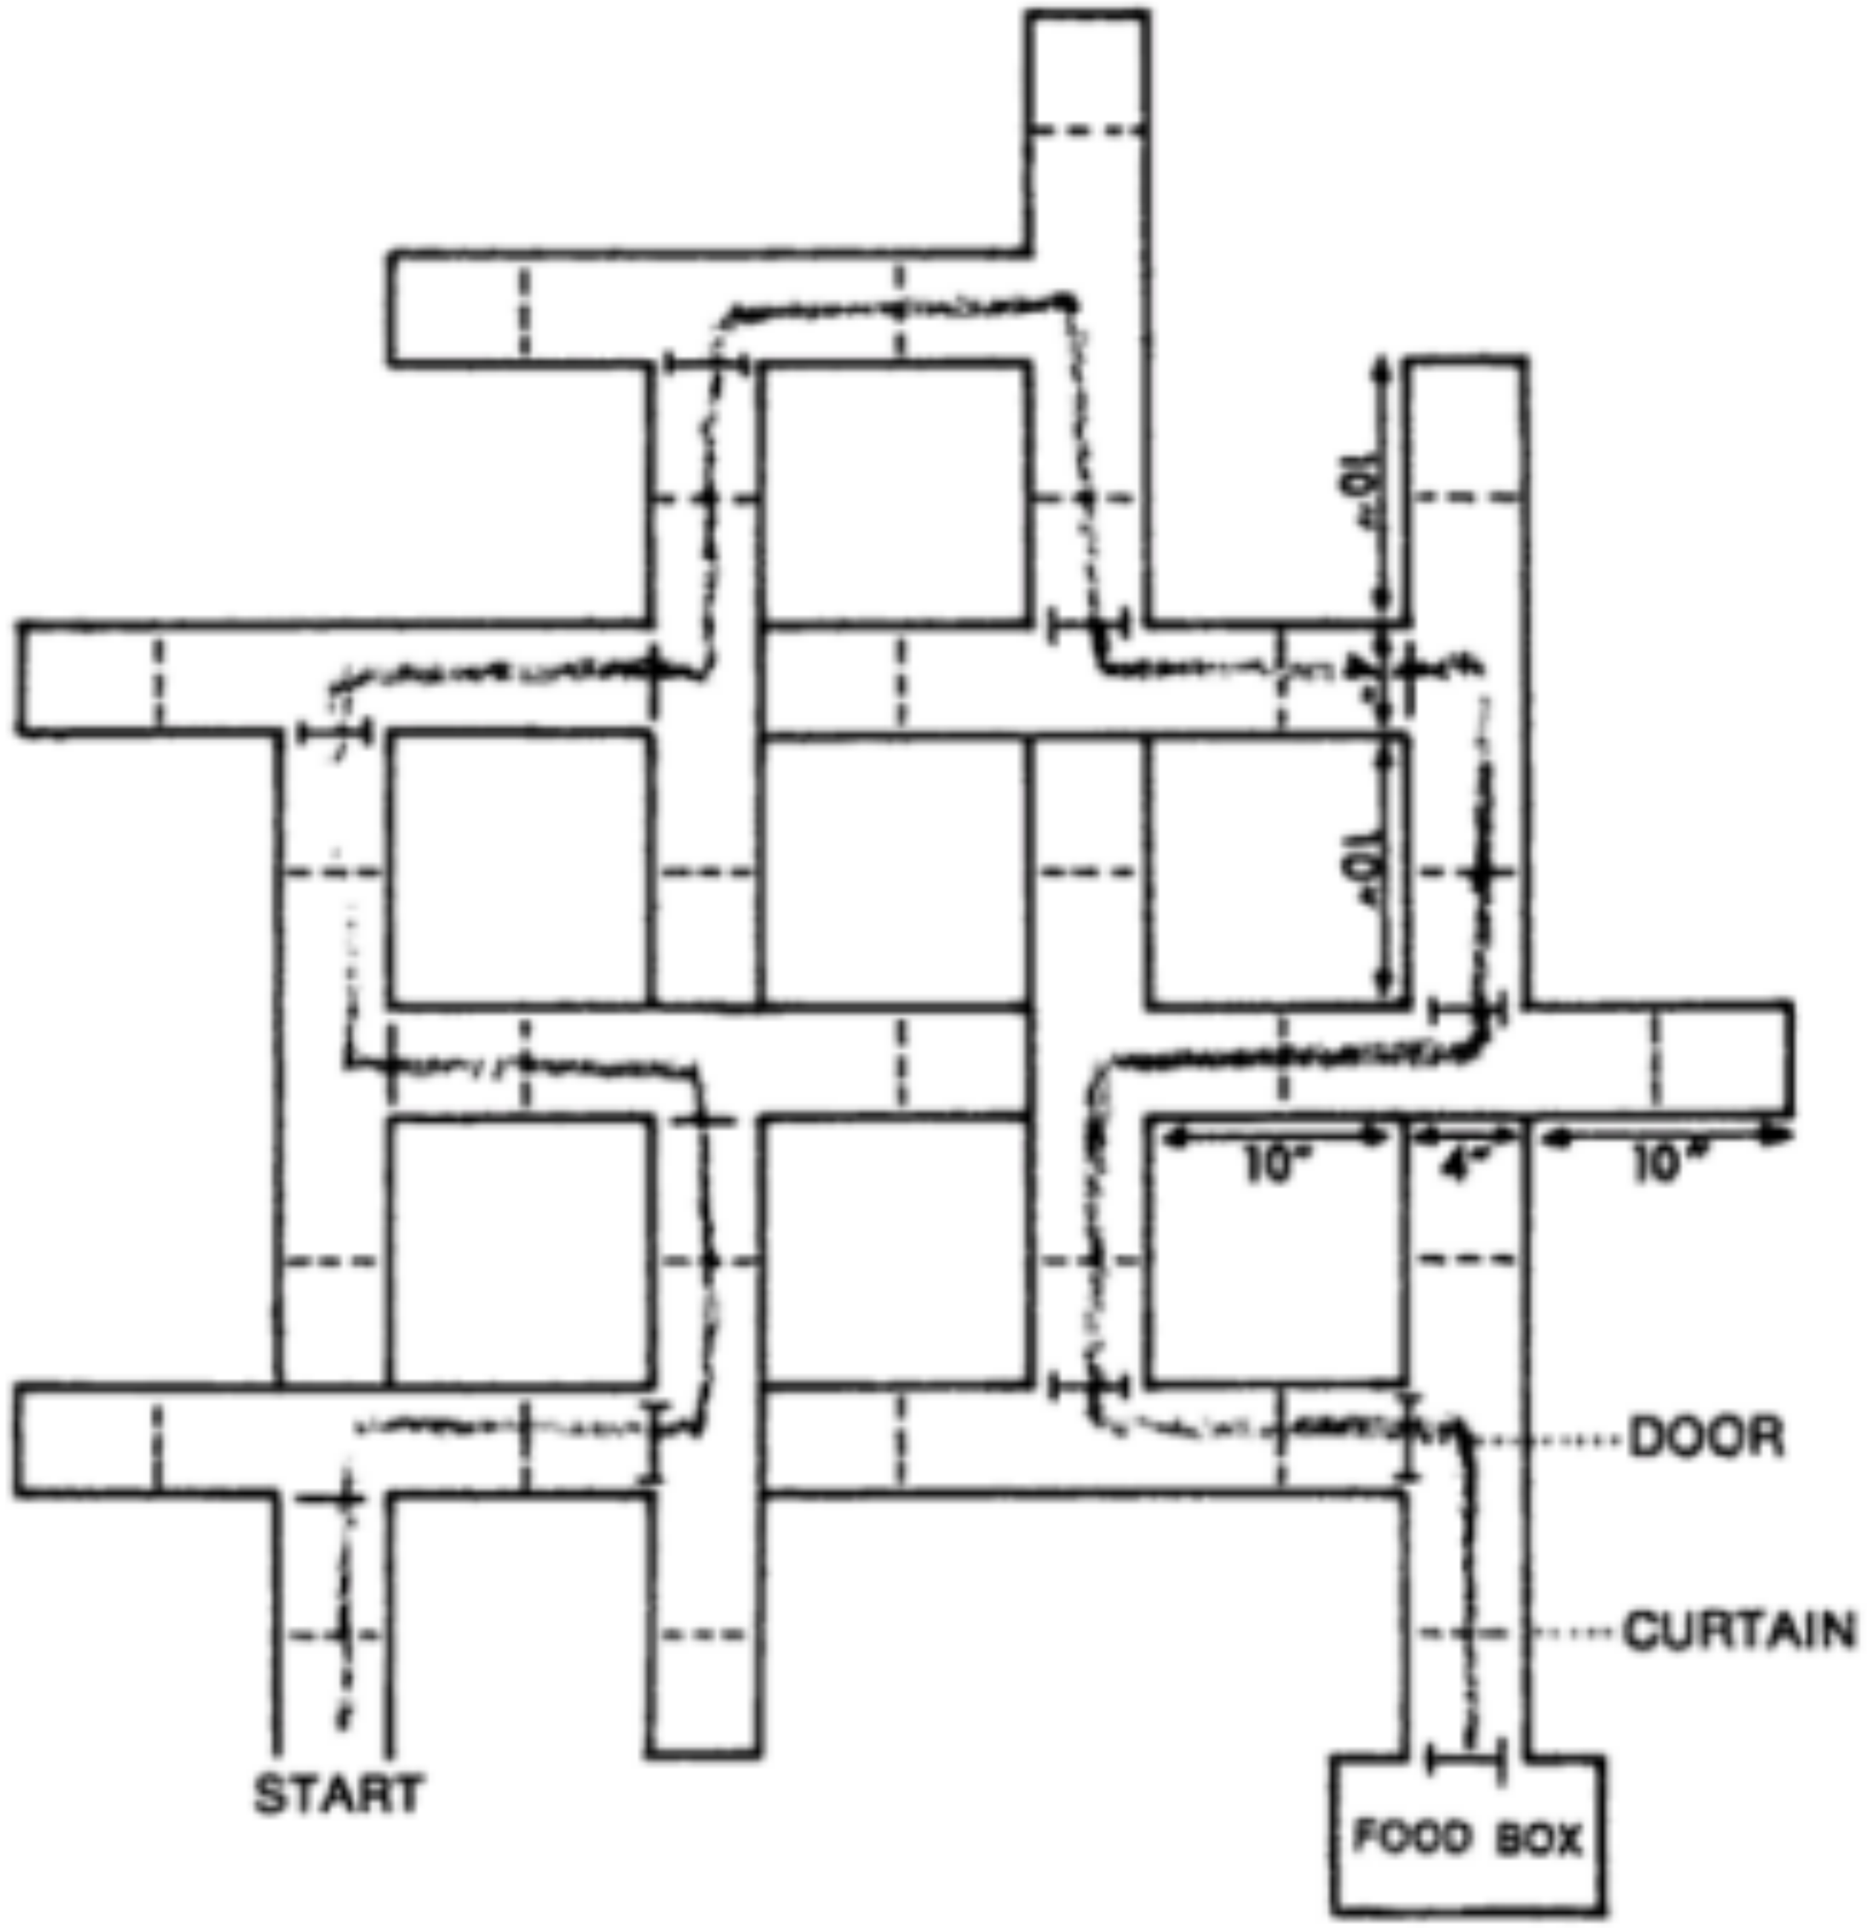

Tolman (1948)

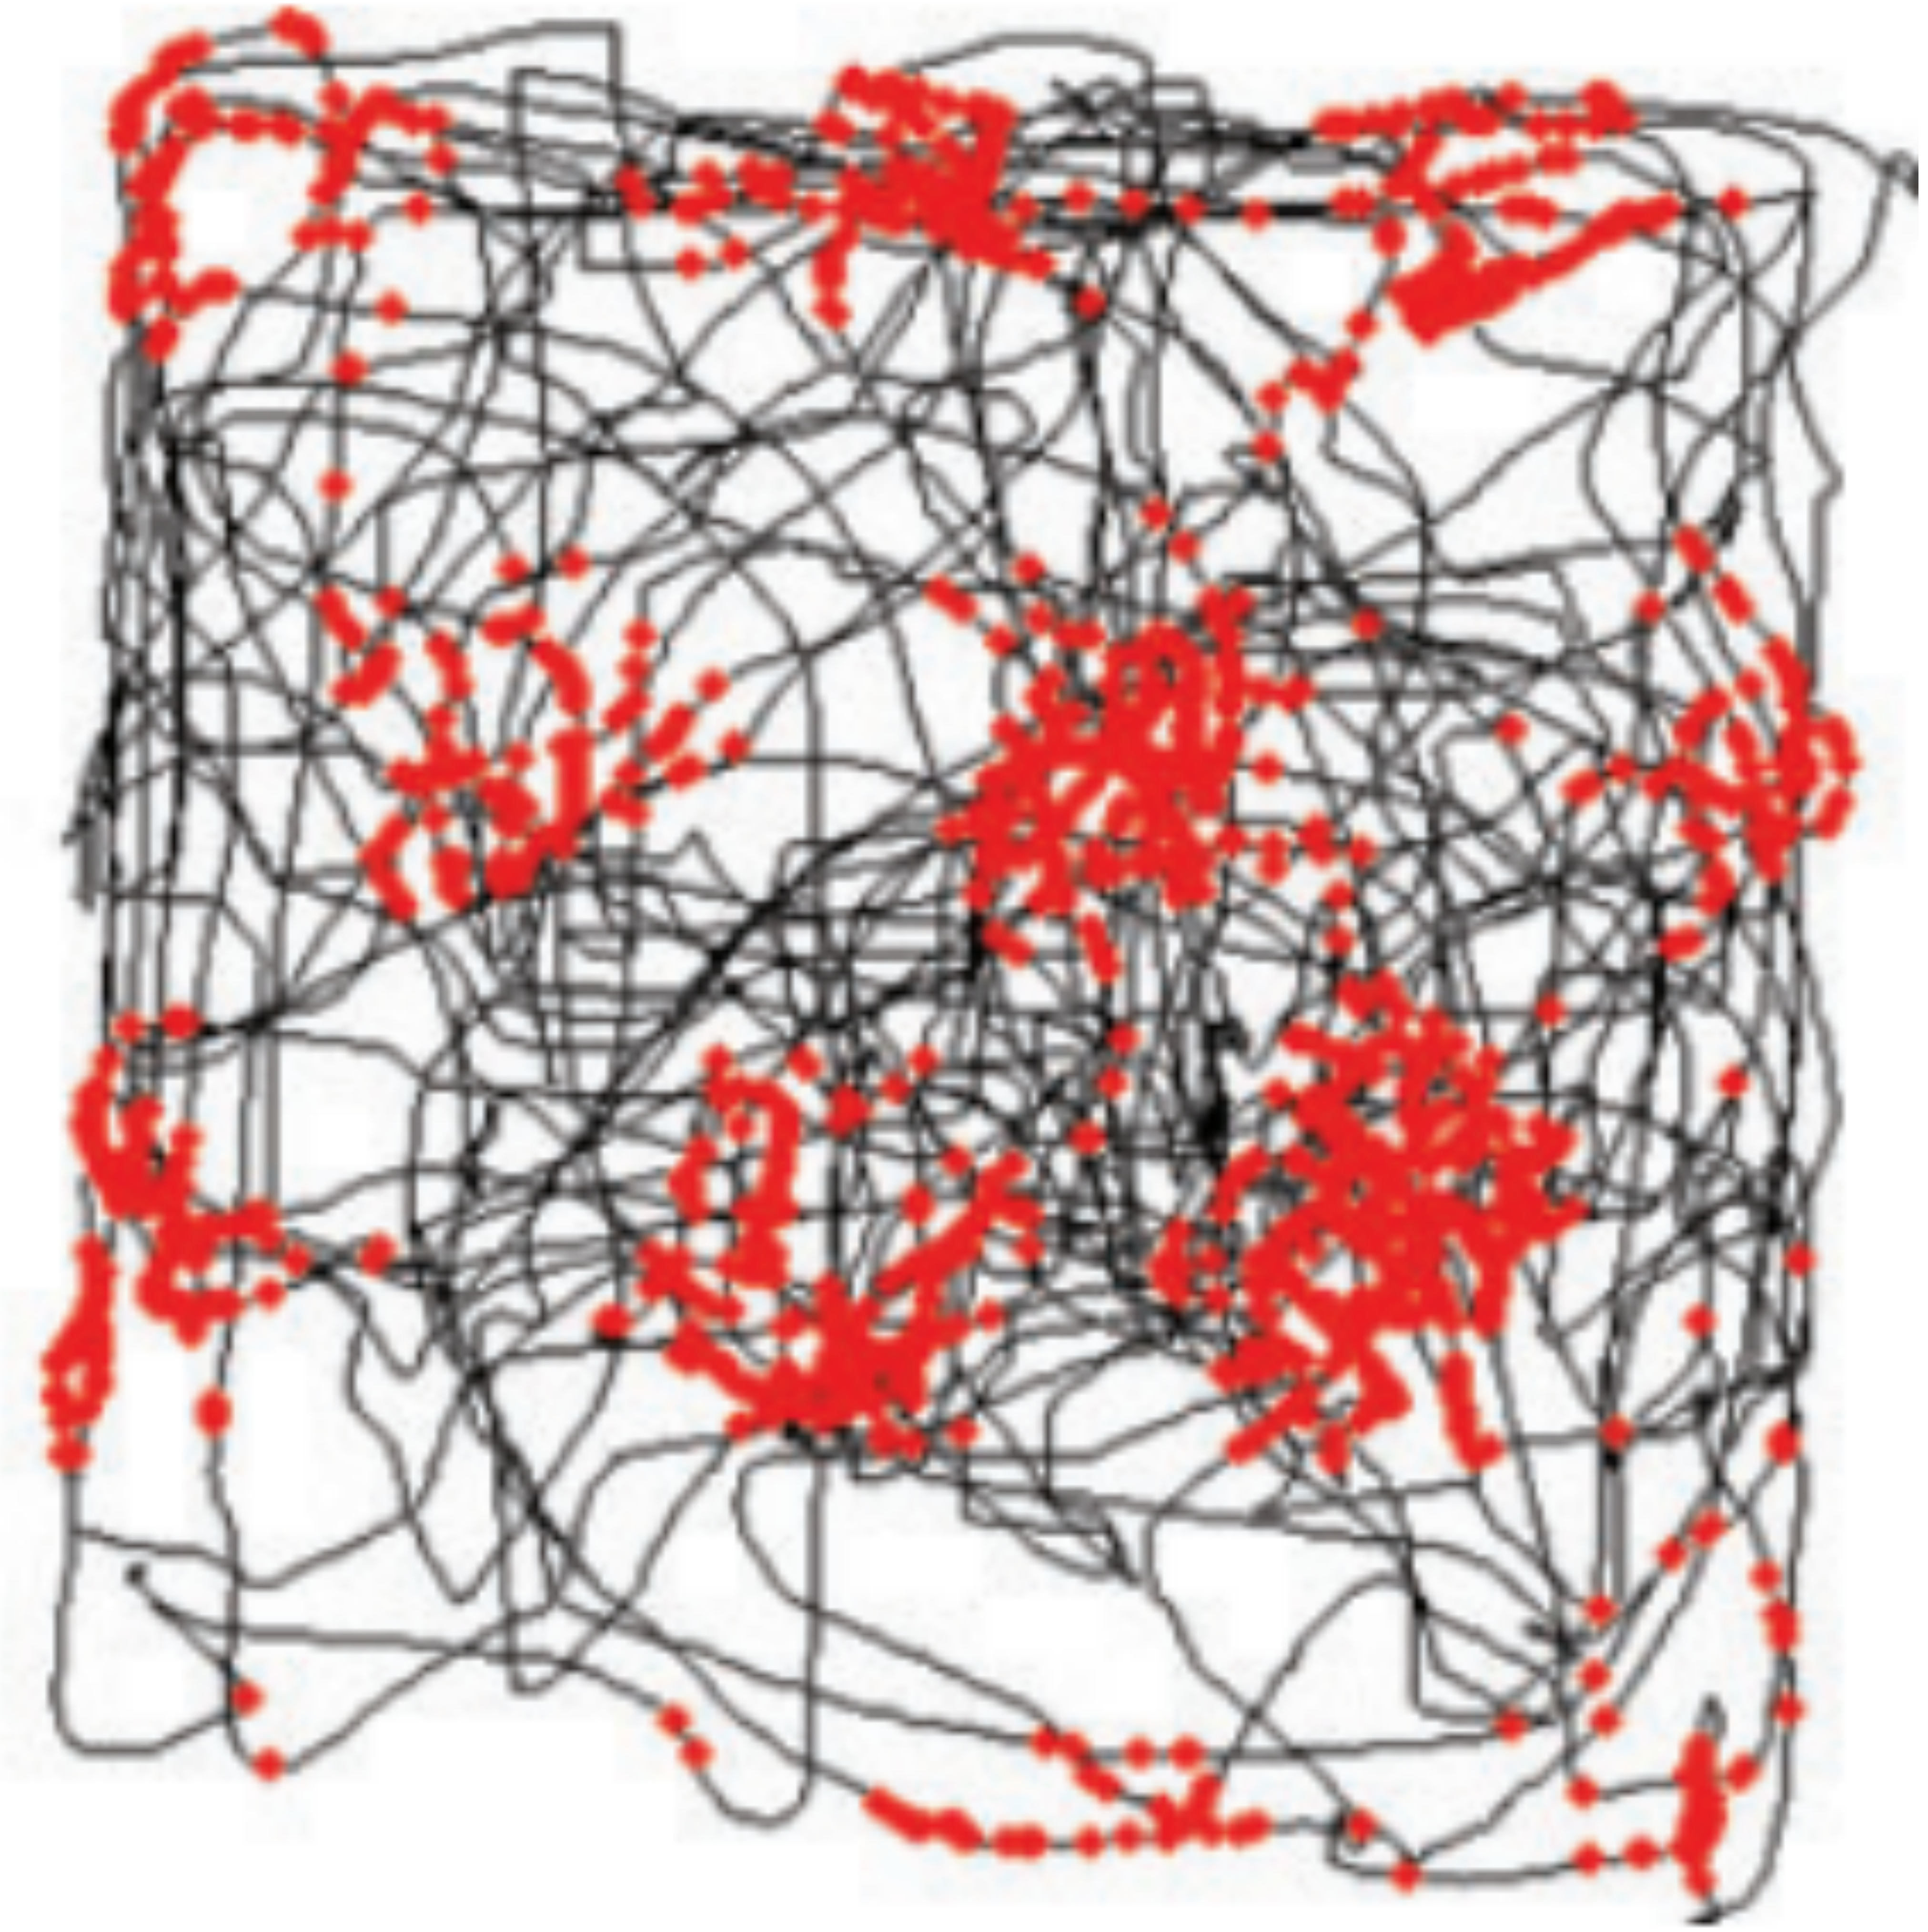

Moser et al., (2008)

The story so far …

Thorndike’s (1911) Law of Effect

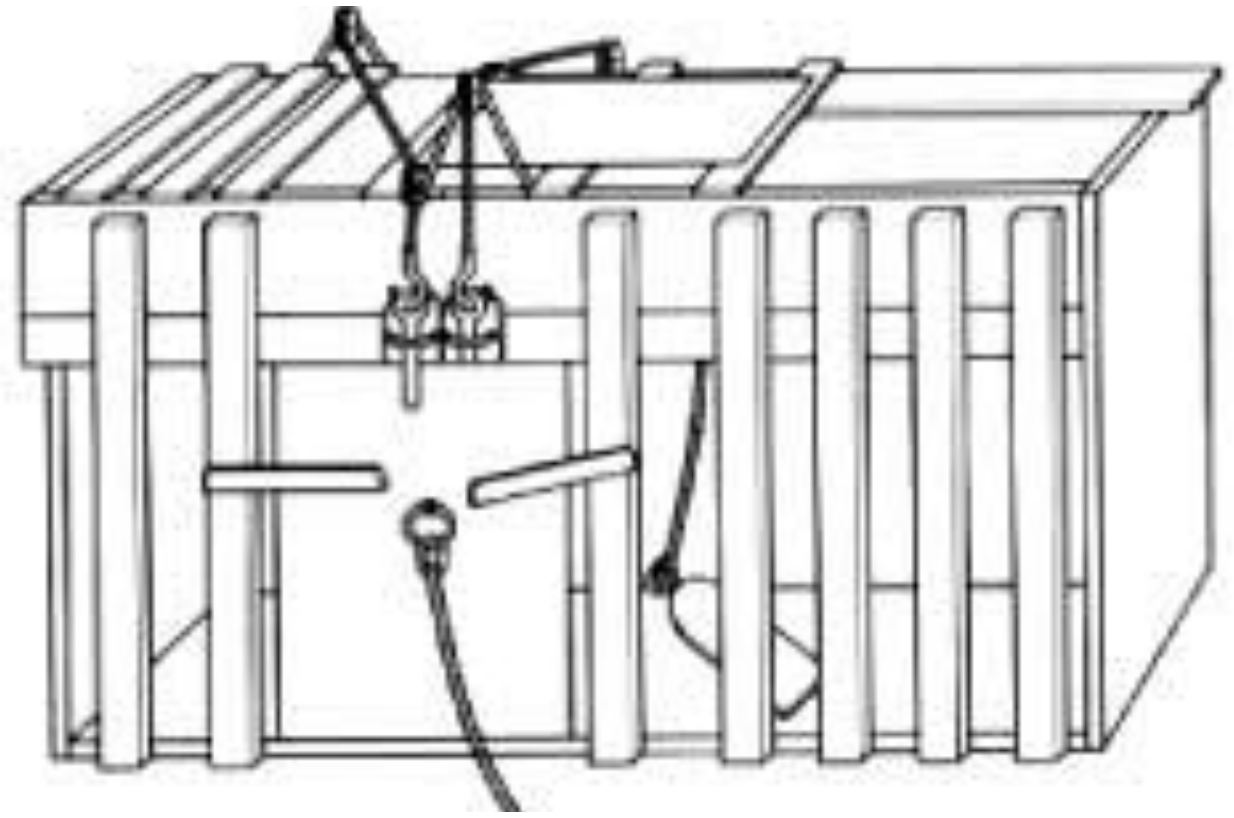

Puzzle Box

Thorndike’s (1911) Law of Effect

Cat

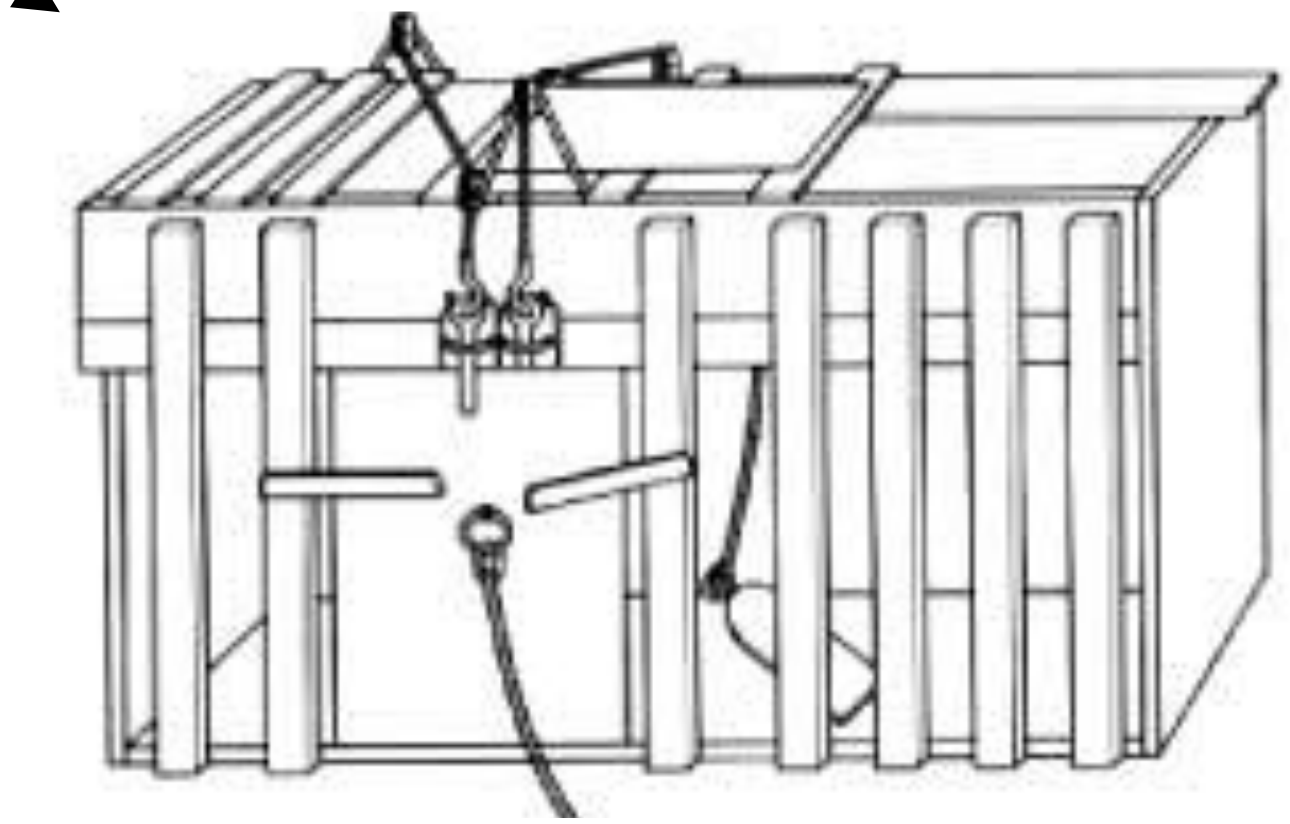

Puzzle Box

Thorndike's (1911) Law of Effect

Cat

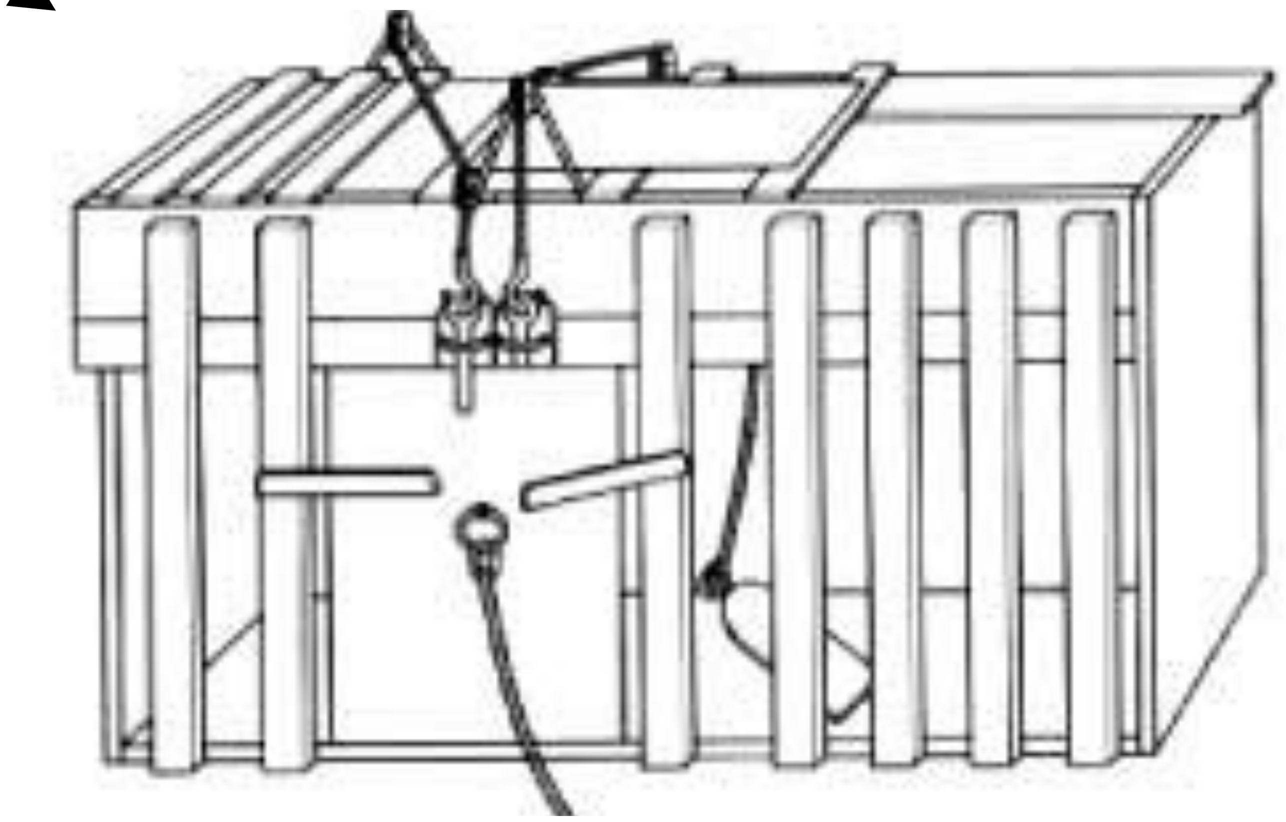

Puzzle Box

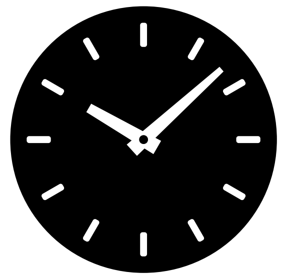

Time to escape

Thorndike's (1911) Law of Effect

Cat

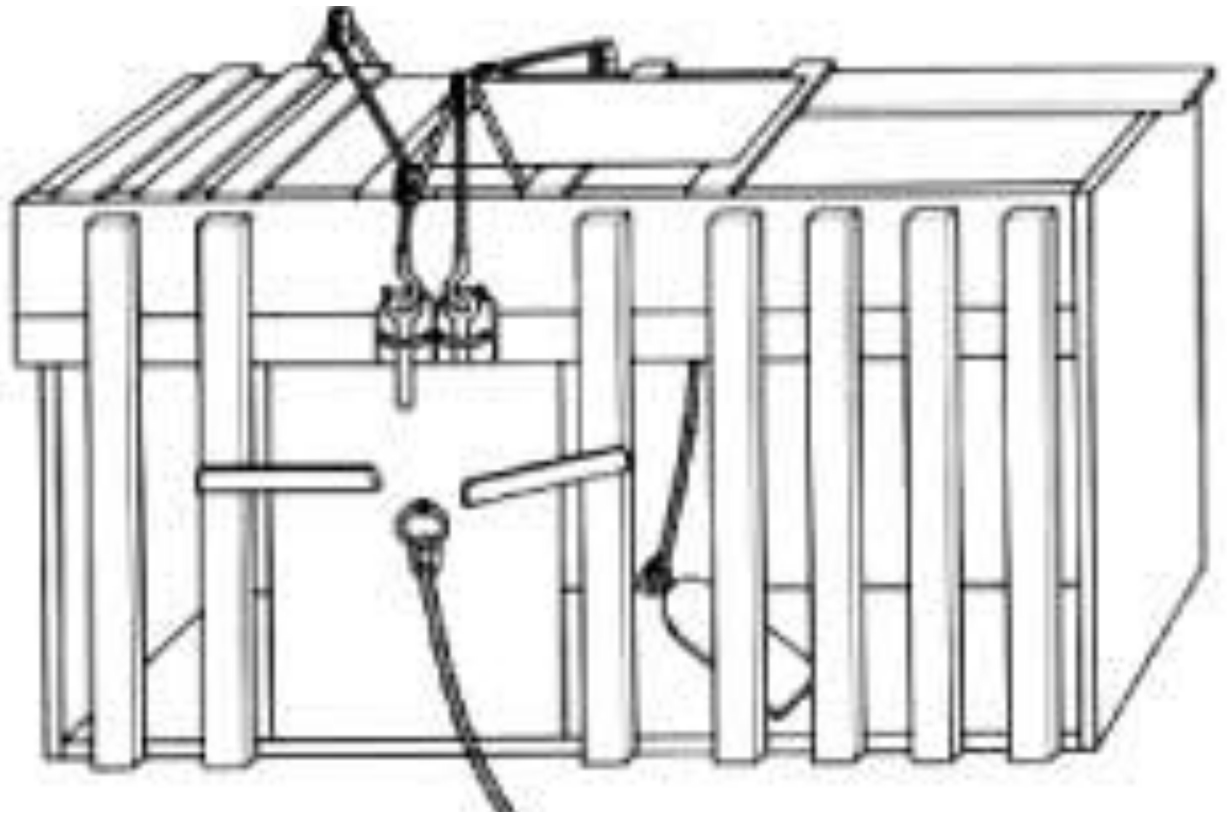

Puzzle Box

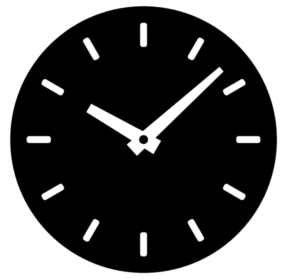

Time to escape

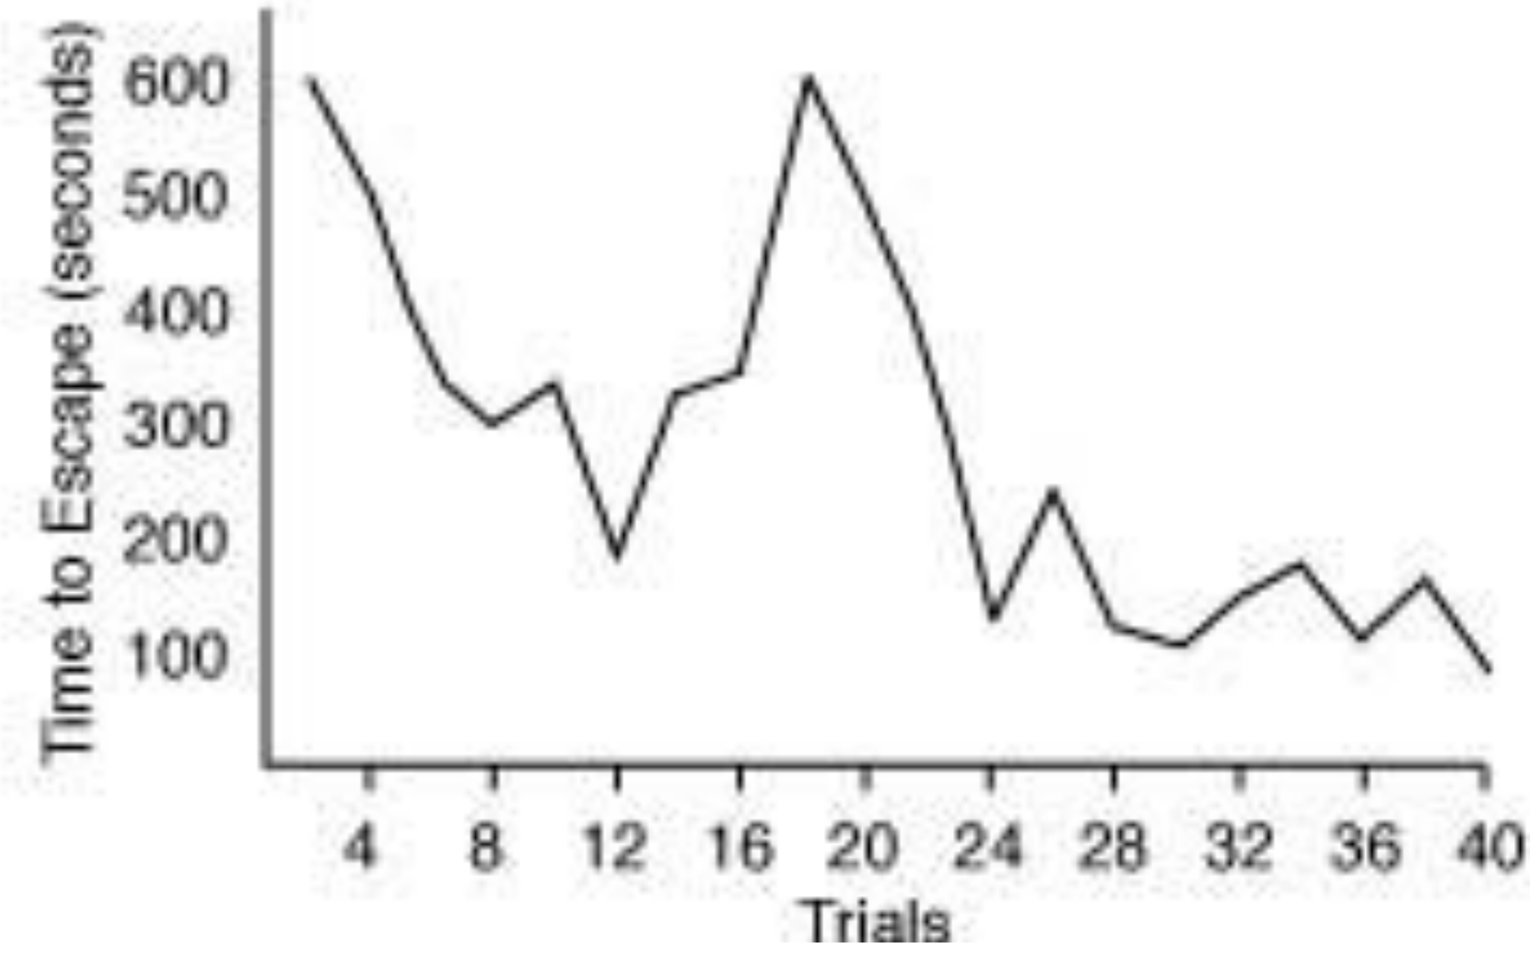

Actions associated with satisfaction are strengthened, while those associated with discomfort become weakened.

Classical and Operant Conditioning

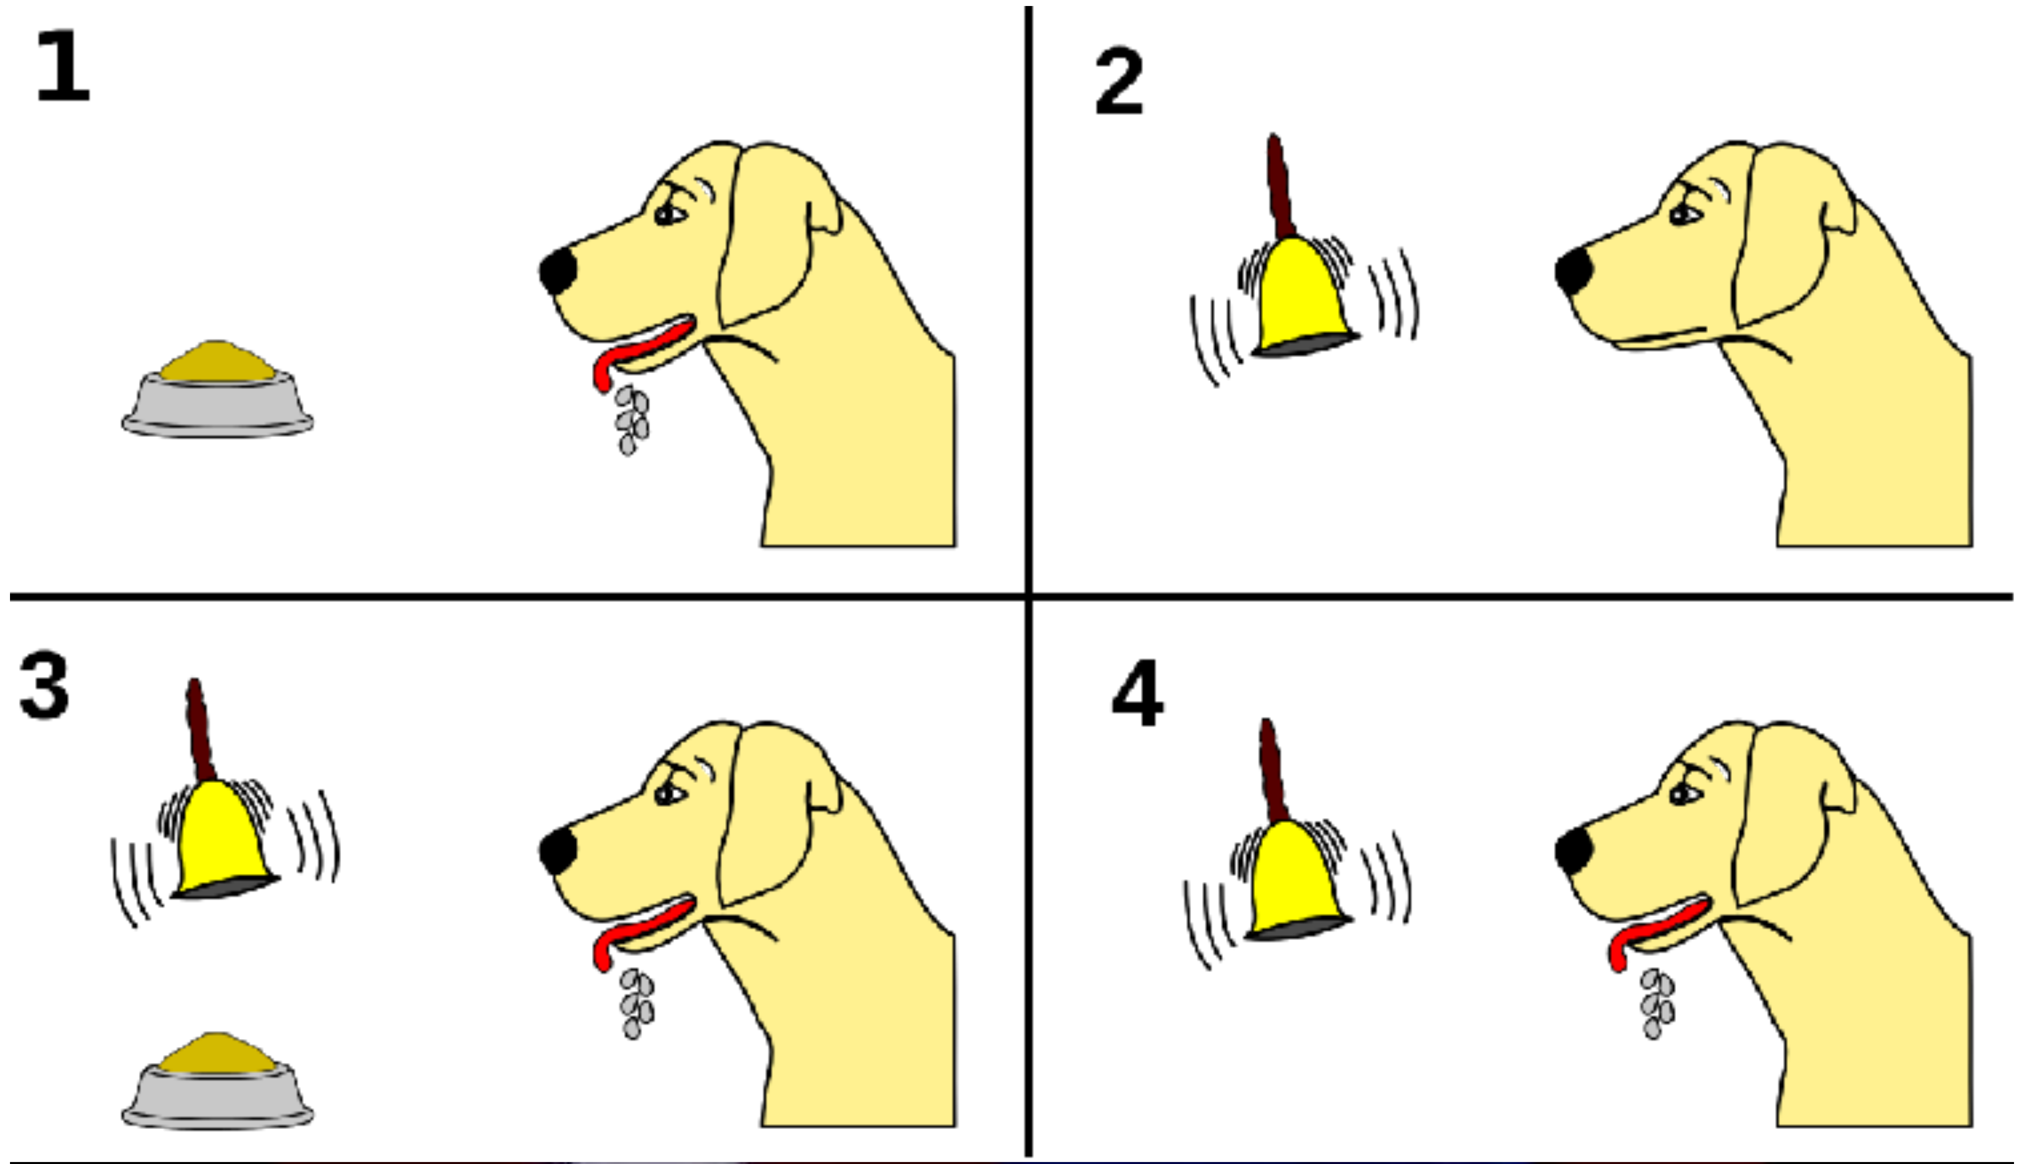

Classical Condition (Pavlov, 1927)

Learning as the *passive* coupling of stimulus (bell ringing) and response (salivation), anticipating future rewards

Operant Condition (Skinner, 1938)

Skinner (1938): Learning as the *active* shaping of behavior in response to rewards or punishments

Edward Tolman (1886 - 1959)

- Raised by an adament Quaker mother
- Studied at MIT, Harvard, and Giessen
- Inspired by Gestalt psychologists like Kurt Koffka and Kurt Lewin
- Coined "***Purposive Behaviorism***"
    - Behavior needs to be studied in the context of the purpose or goals of behavior
- In contrast to other **behaviorists** at the time, Tolman believed in latent learning and the need to talk about hidden mental states in how we make decisions

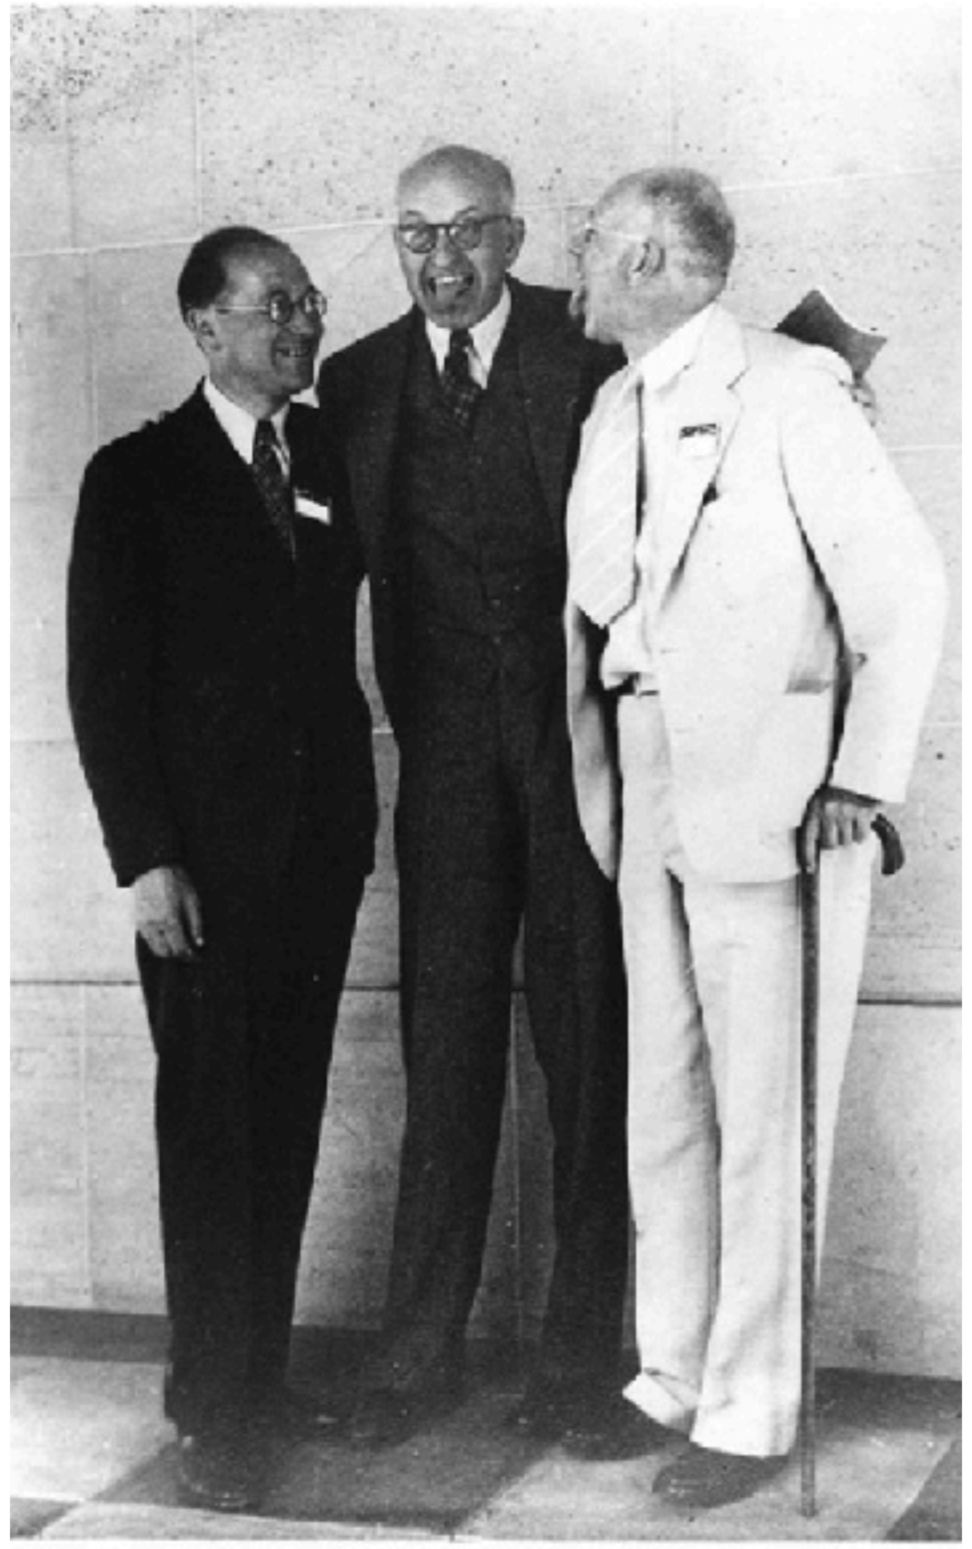

Lewin, Tolman, & Hull

Tolman and Cognitive maps

- Learning is not just a telephone switchboard connecting incoming sensory signals to outgoing responses (S-R Learning)
- Rather, “latent learning” establishes something like a “field map of the environment” gets etablished (S-S learning)

Stimulus-Response (S-R) Learning

Stimulus-Stimulus (S-S) Learning

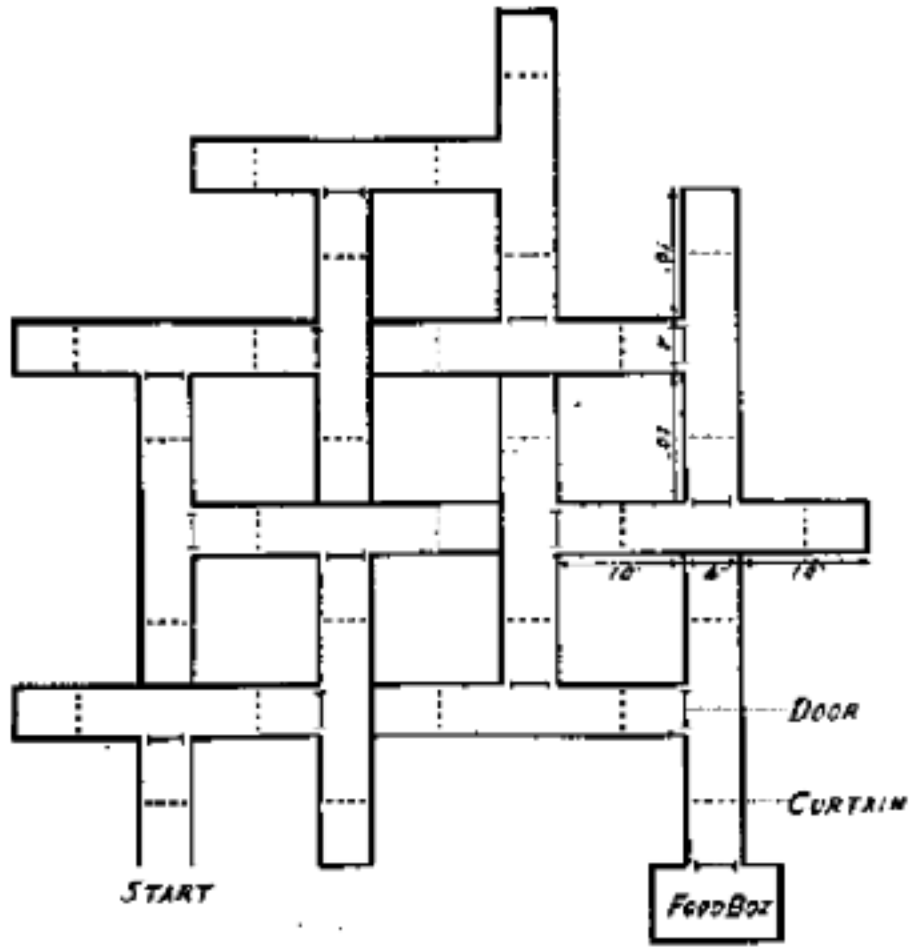

Plan of maze
14 Unit T-Alley Maze

FIG. 1

(From M. H. Elliott, The effect of change of reward on the maze performance of rats. *Univ. Calif. Publ. Psychol.*, 1928, 4, p. 20.)

Tolman (1948): Different interpretations

“All students agree as to the facts. They disagree, however on theory and explanation”

- **S-R school**: learning consists of strenthening/weakening of S-R connections (like a telephone exchange)
  - subgroup a) more frequent responses are strengthened (Law of Exercise)
  - subgroup b) more rewarded responses are strengthened (Law of Effect)
- **S-S school**: in the course of learning, “*a field map of the environment gets established”*
  - Sampling of stimuli is not passive, but active and selective during learning w.r.t. to a goal or purpose
  - Stimuli are not just routed to associations, but used to construct some new map-like representation that captures the relational structure of the environment
  - The nature of these map-like representations (strip-like vs. broad) have consequences for generalization

Experiments

1. Latent Learning

2. Vicarious trial and error

3. Searching for the stimulus

4. Hypotheses

5. Spatial orientation

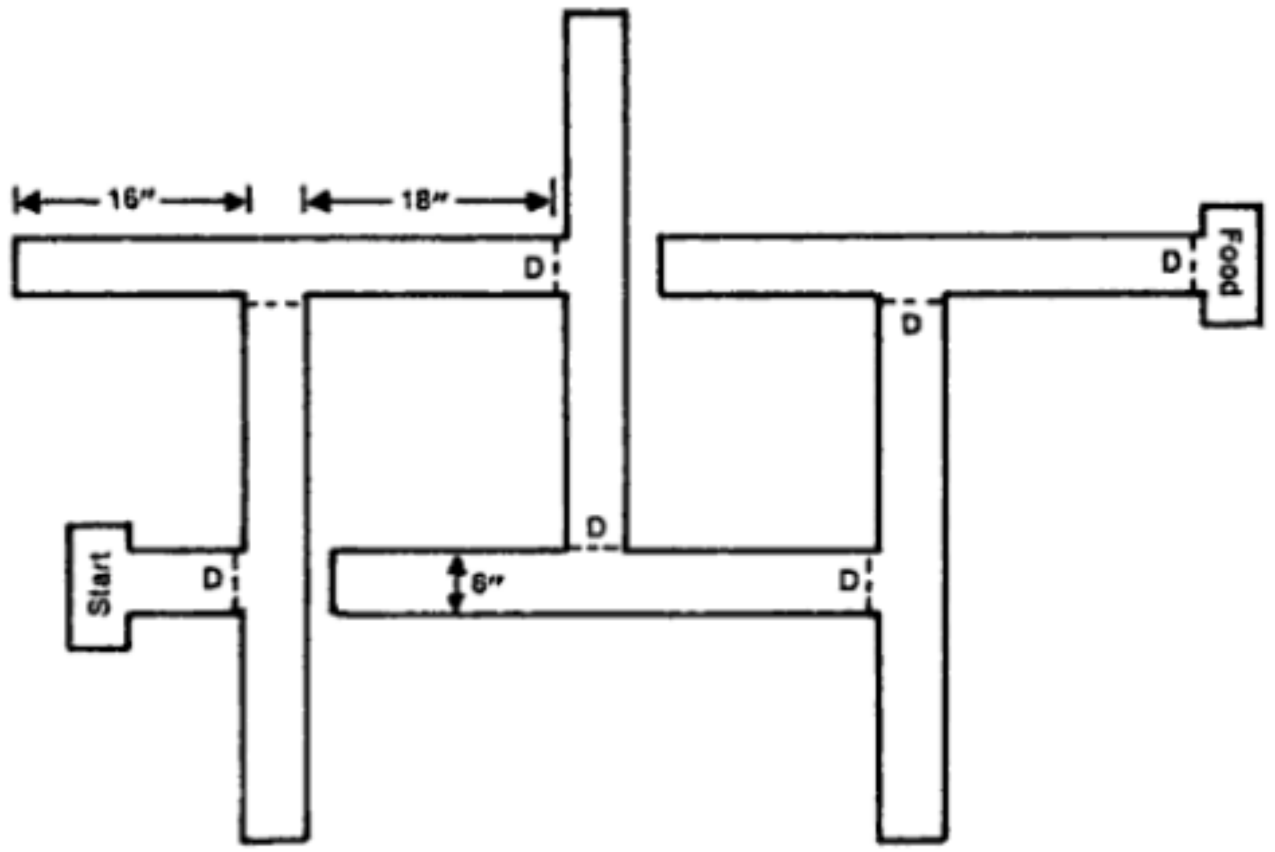

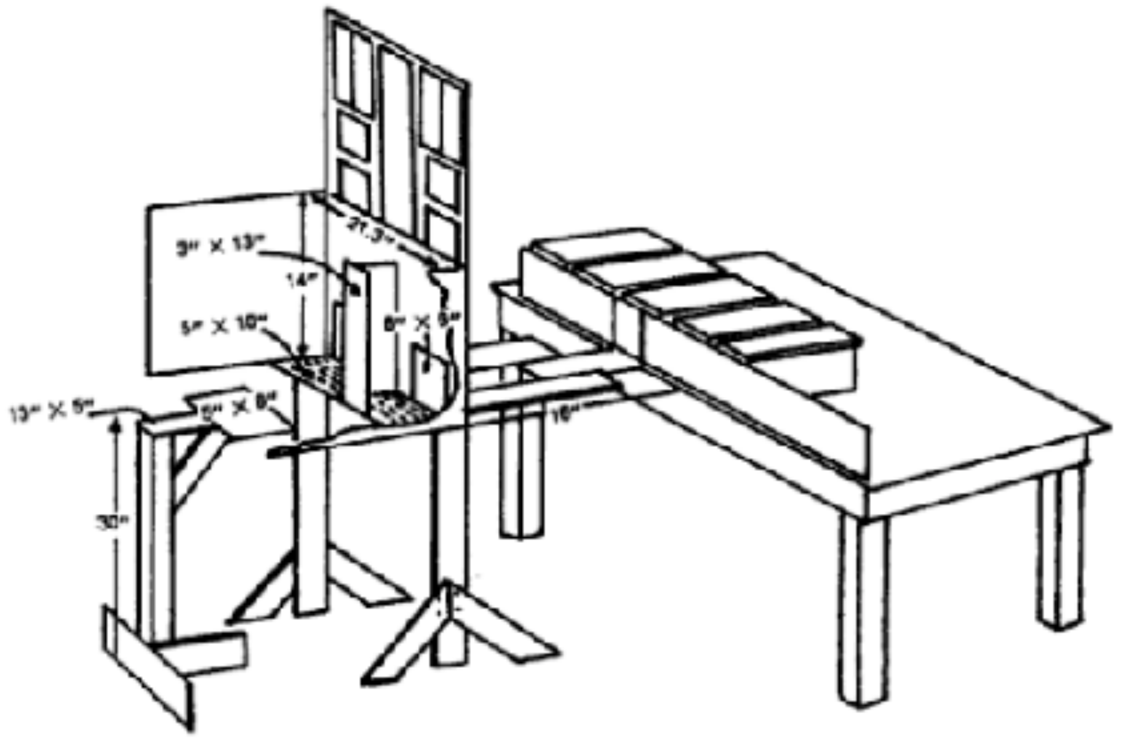

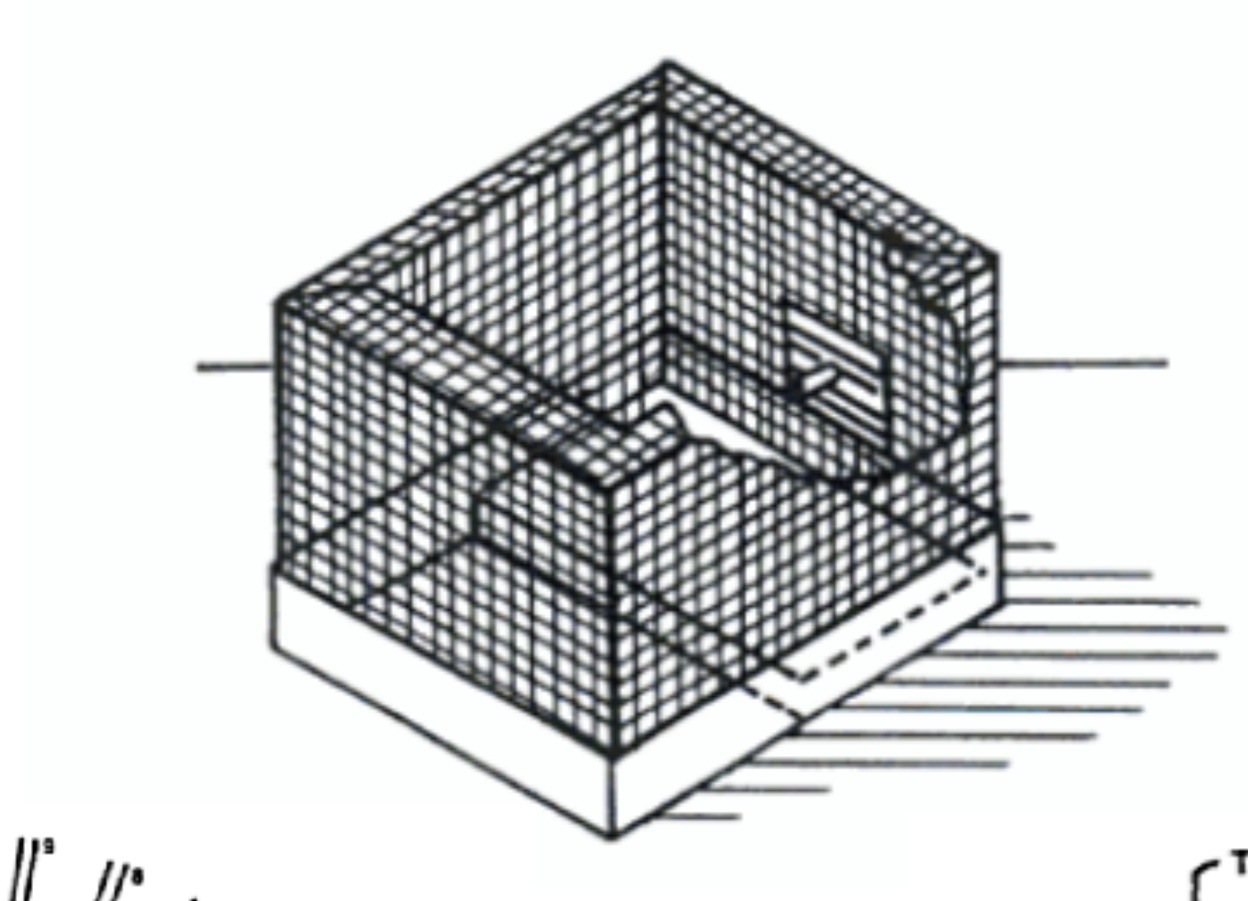

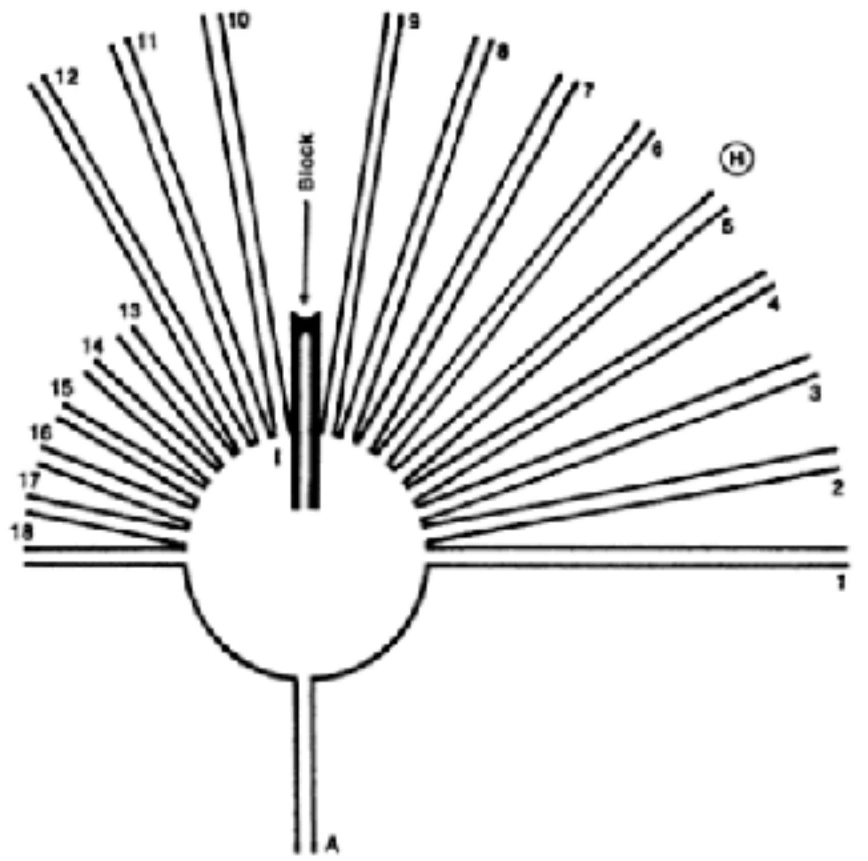

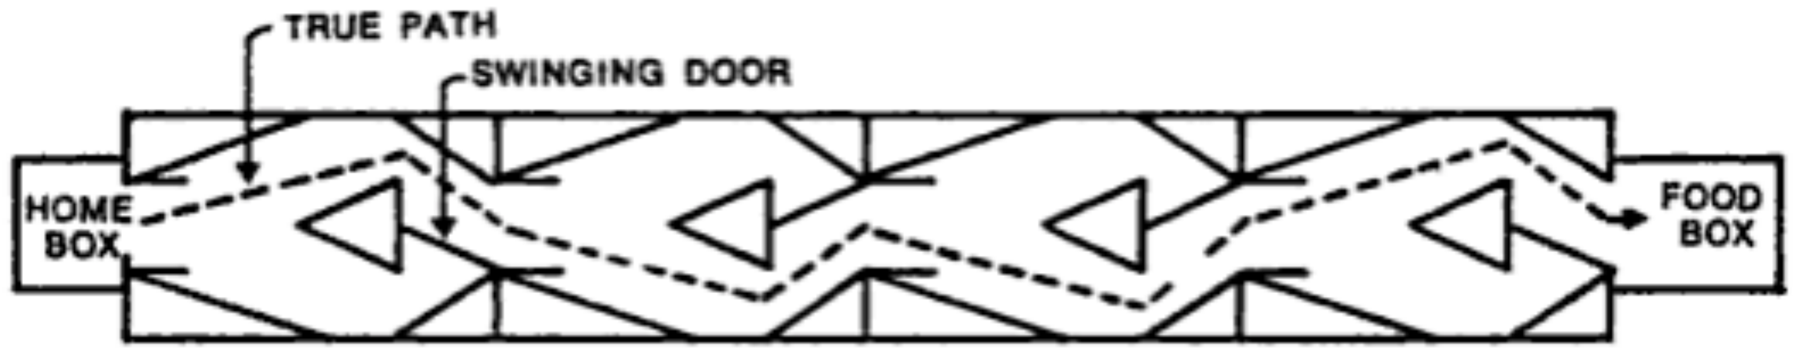

Experiments

1. Latent Learning

2. Vicarious trial and error

3. Searching for the stimulus

4. Hypotheses

5. Spatial orientation

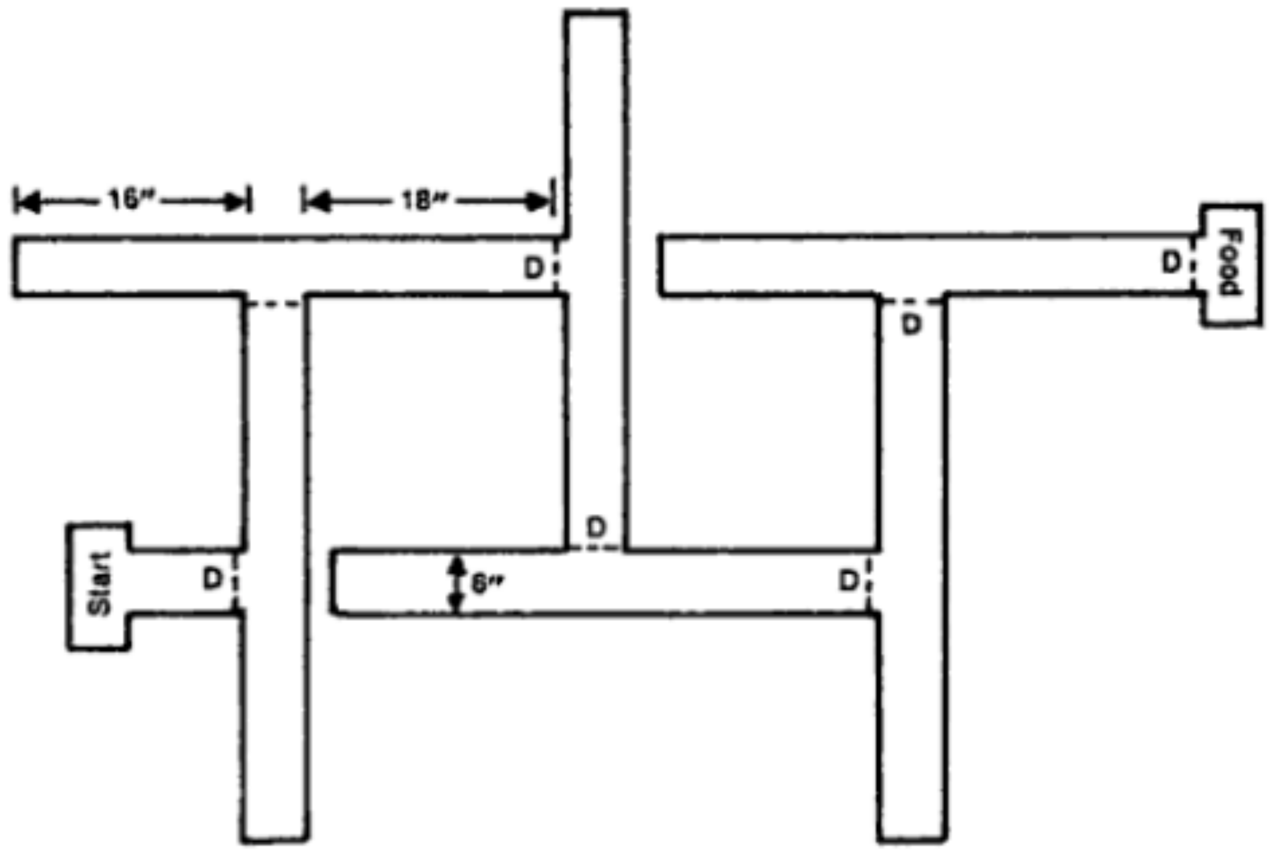

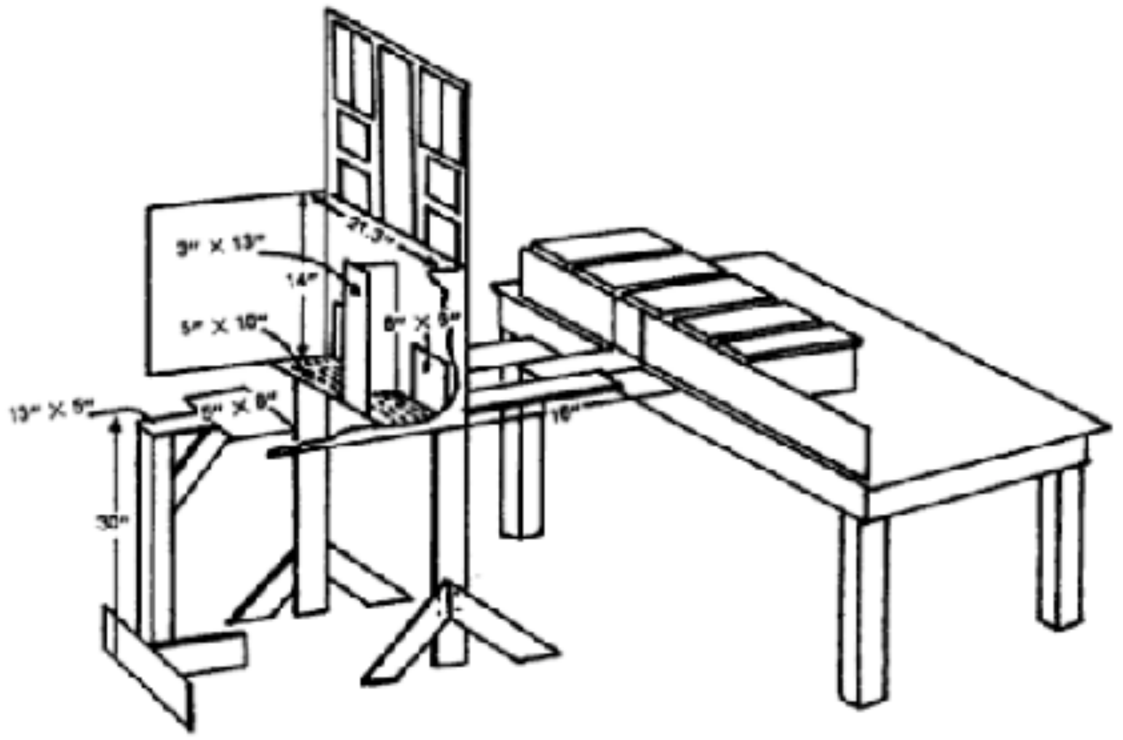

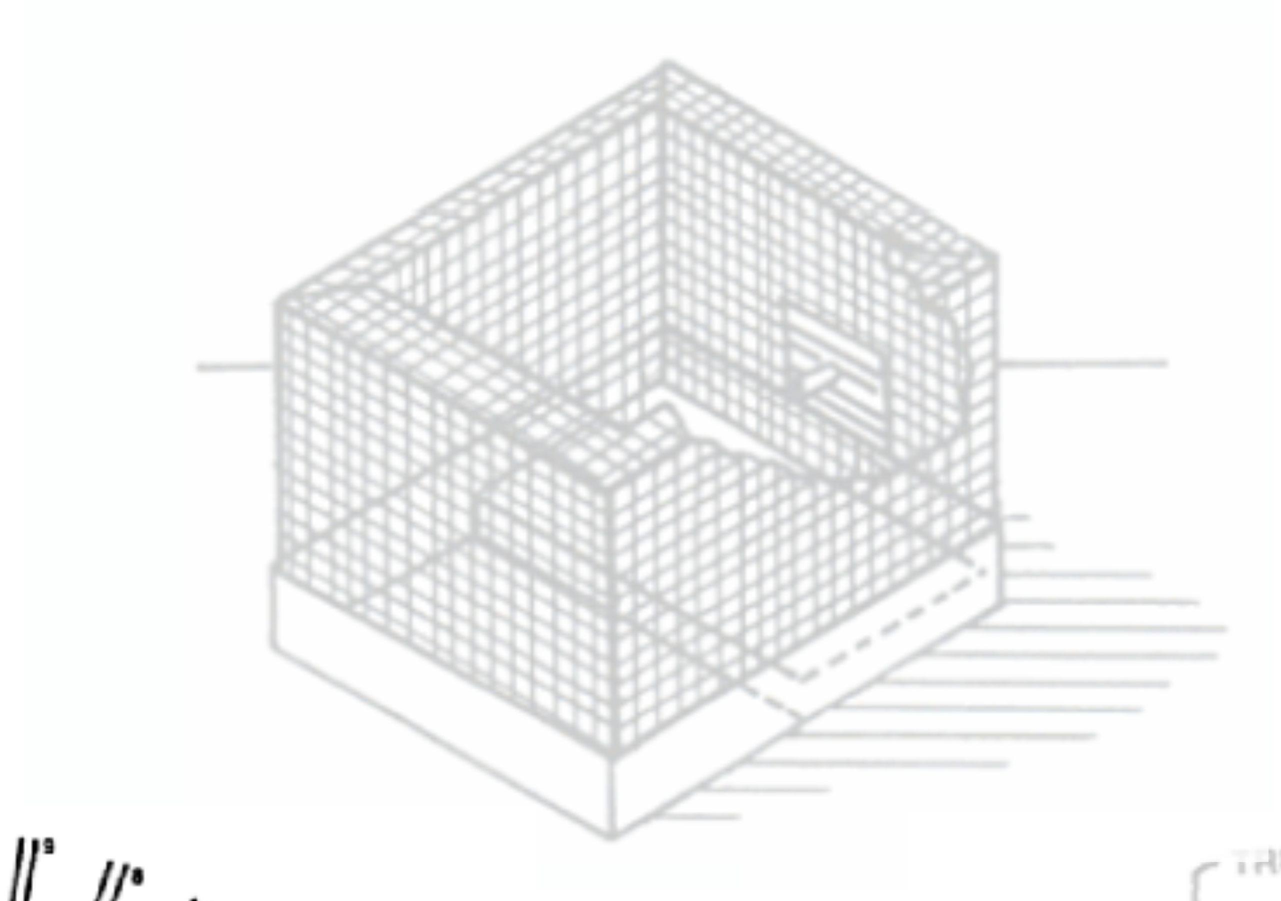

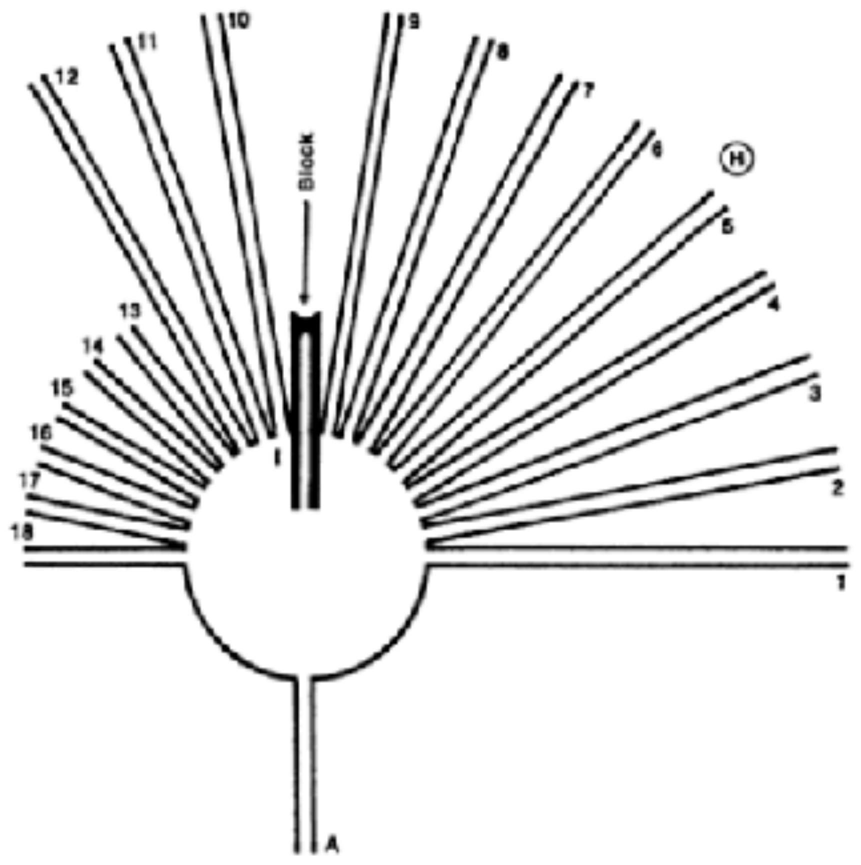

Latent Learning

- Blodgett (1929) Maze navigation task
    - **Group 1** [Control]: one trial a day with food in the goal box at the end
    - **Group 2** [Late food] No food in the maze for days 1-6, then food provided at the end on day 7
    - **Group 3** [Early food] … food added on day 3
- Learning curves dropped dramatically when food was added
    - This suggests latent learning prior to reward
    - “They had been building up a ‘map’”
    - Once the reward was added, they could use the map rather than starting from scratch

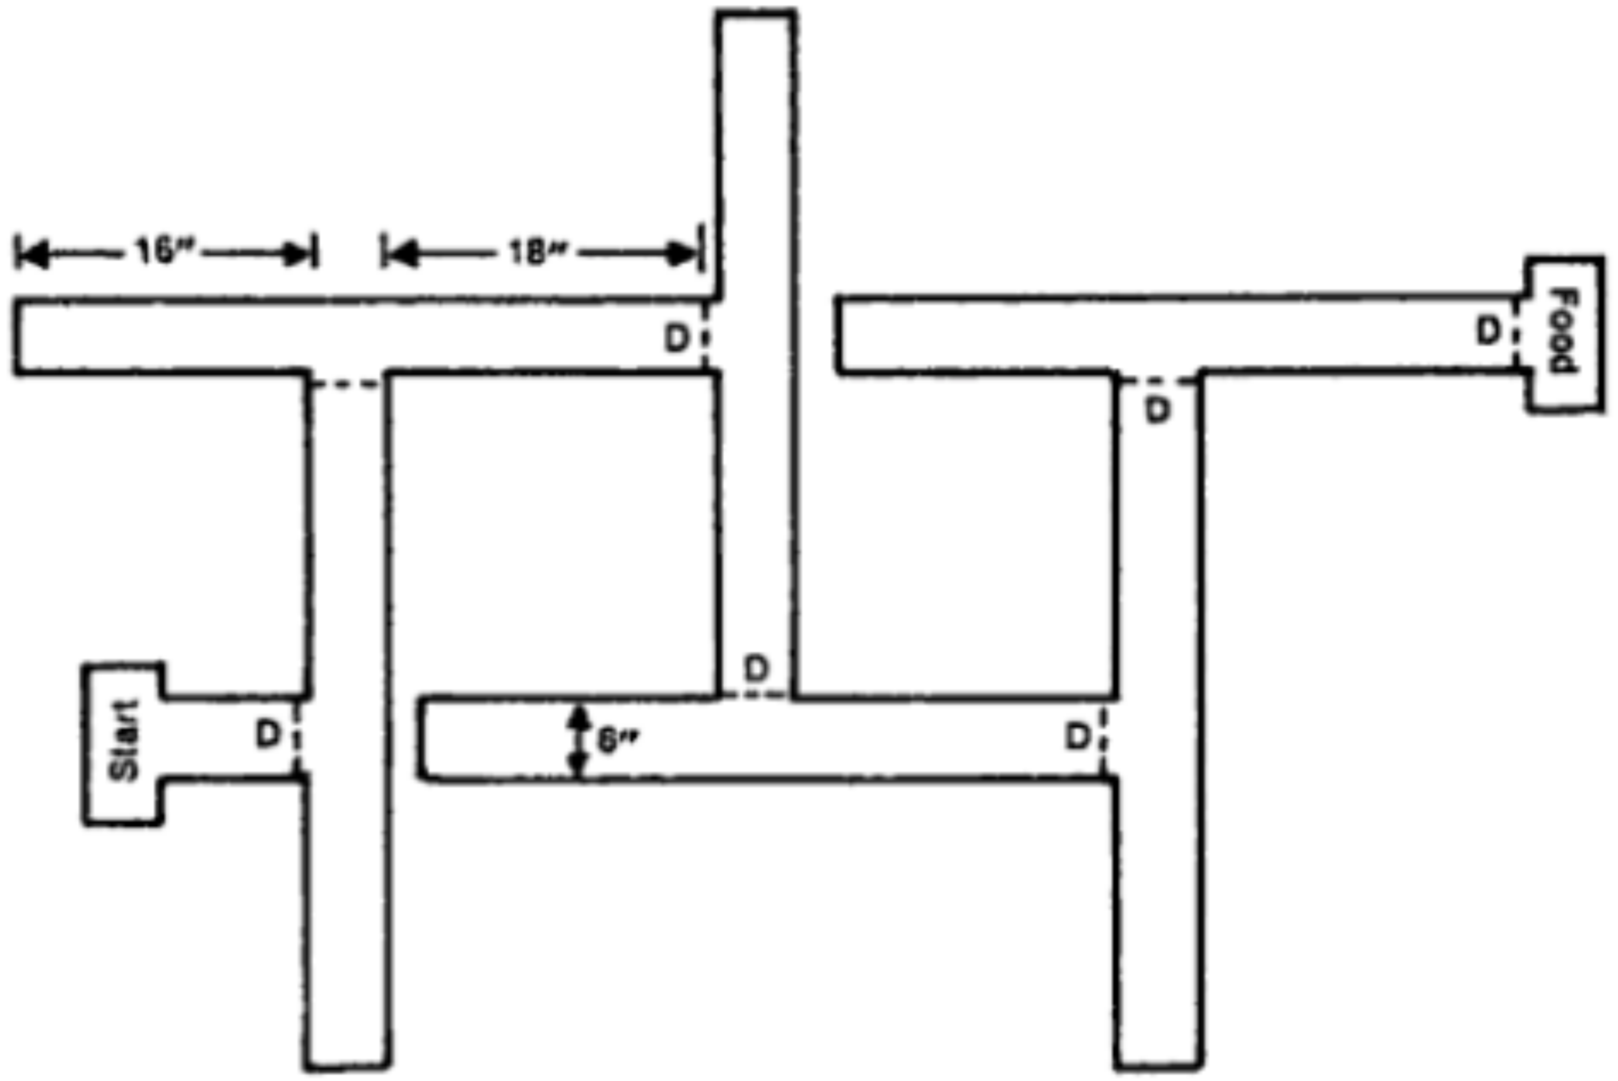

(lower is better)

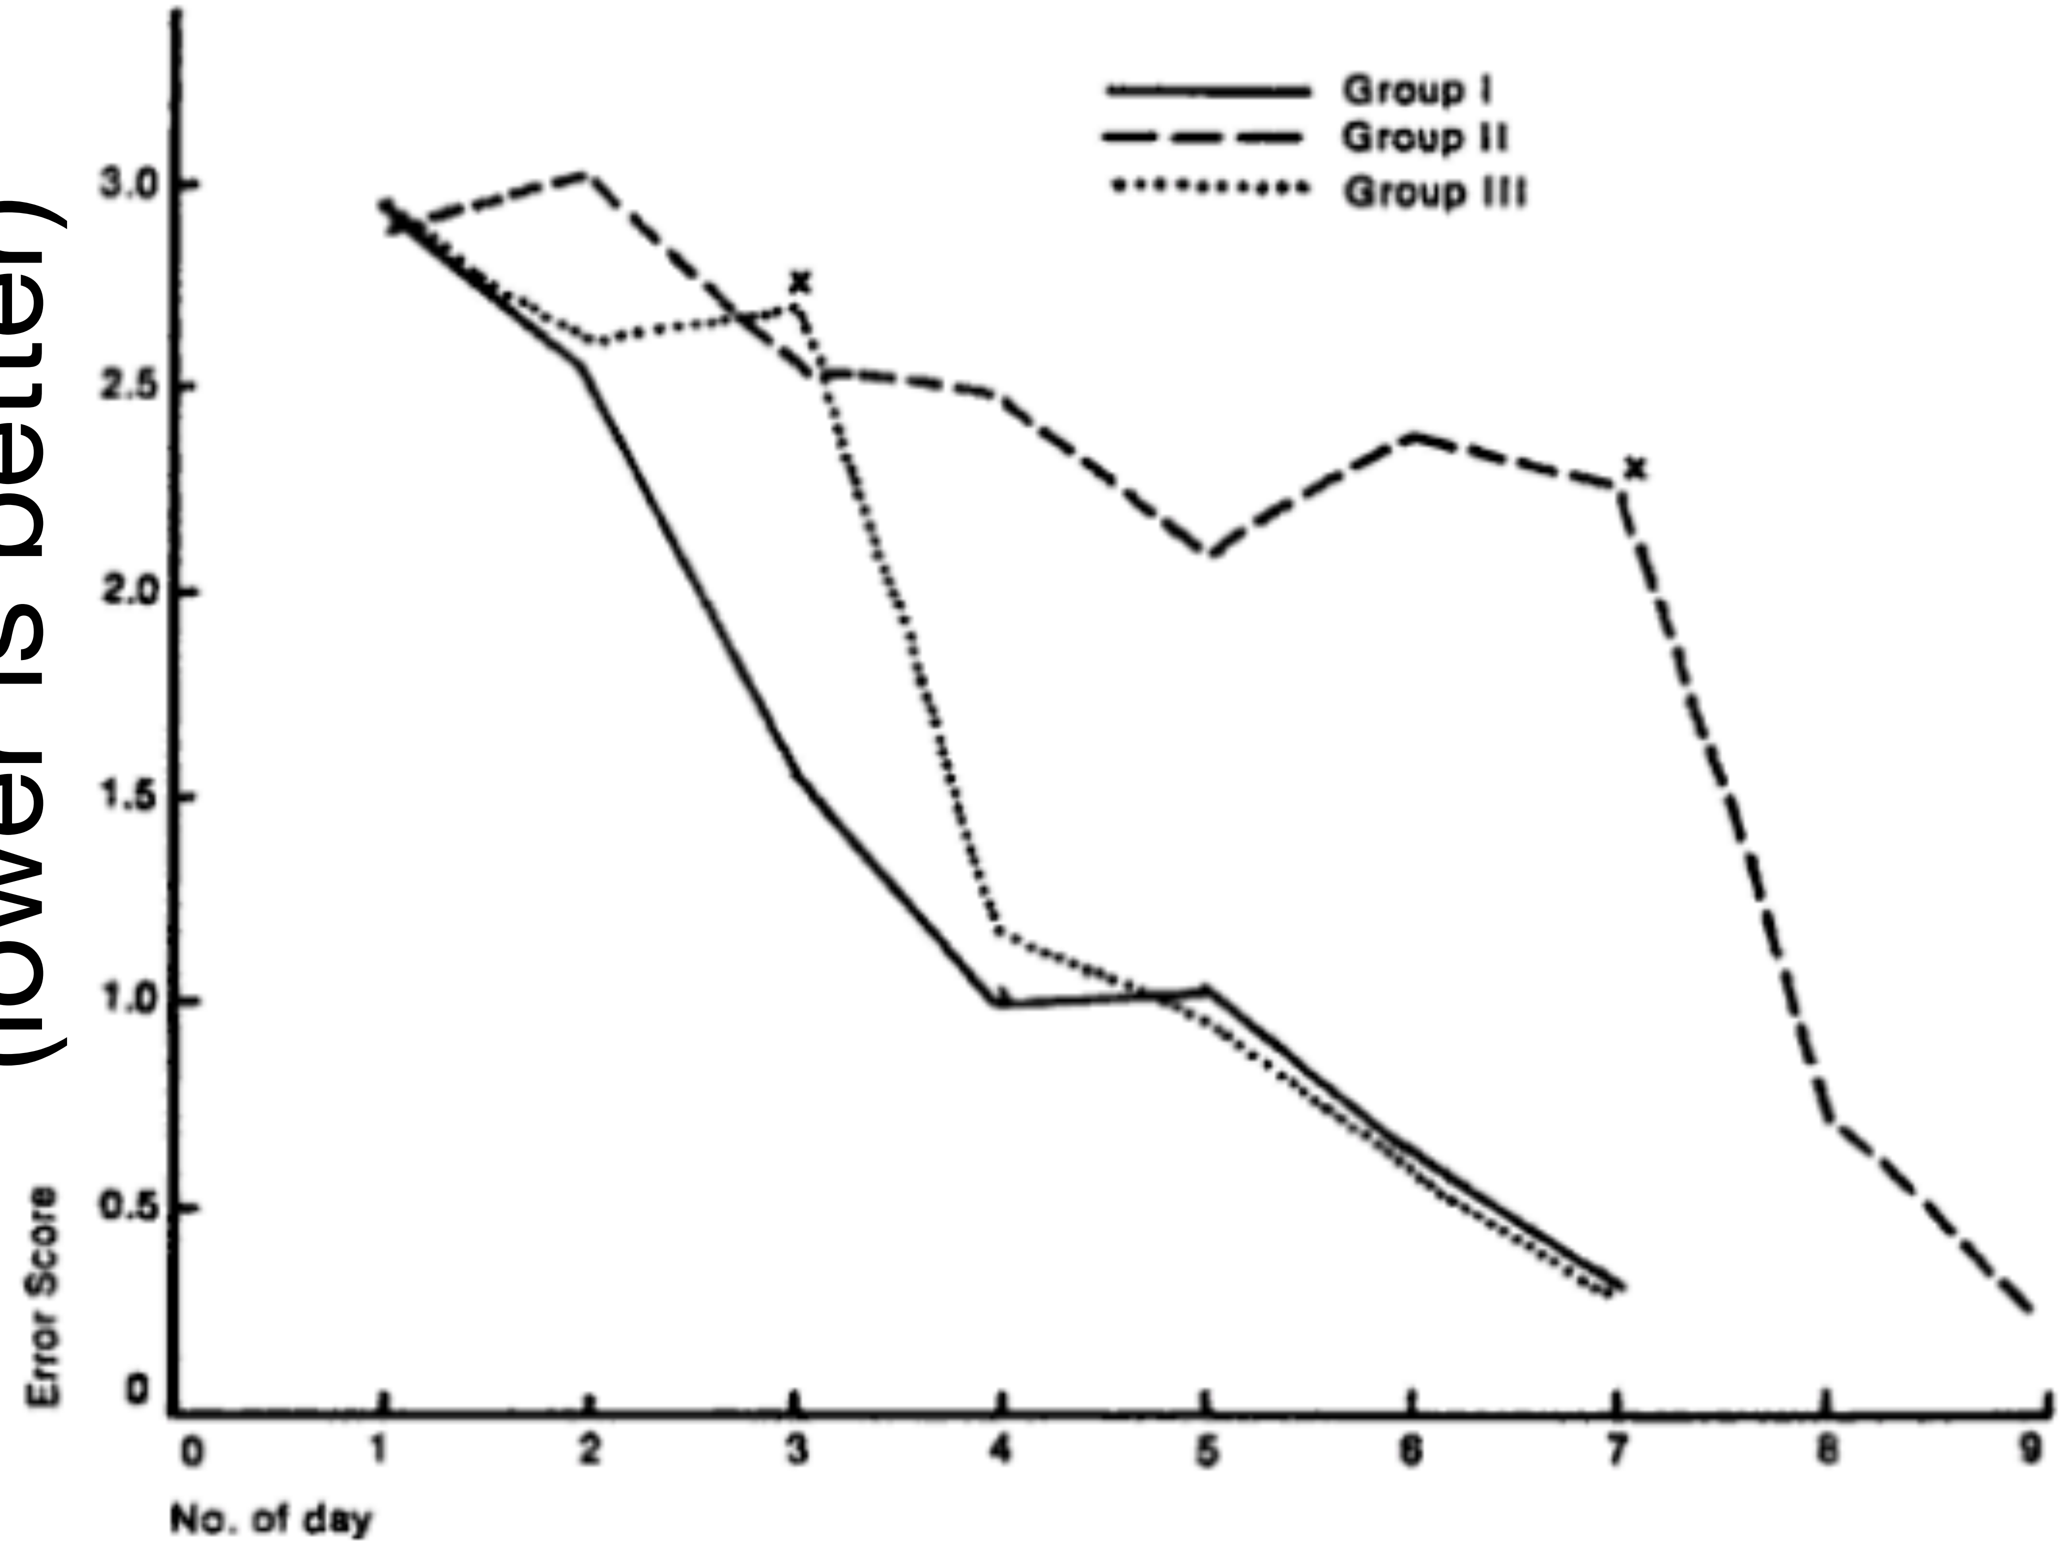

Latent Learning

- Blodgett (1929) Maze navigation task
  - **Group 1** [Control]: one trial a day with food in the goal box at the end
  - **Group 2** [Late food] No food in the maze for days 1-6, then food provided at the end on day 7
  - **Group 3** [Early food] … food added on day 3
- Learning curves dropped dramatically when food was added
  - This suggests latent learning prior to reward
  - “They had been building up a ‘map’”
  - Once the reward was added, they could use the map rather than starting from scratch

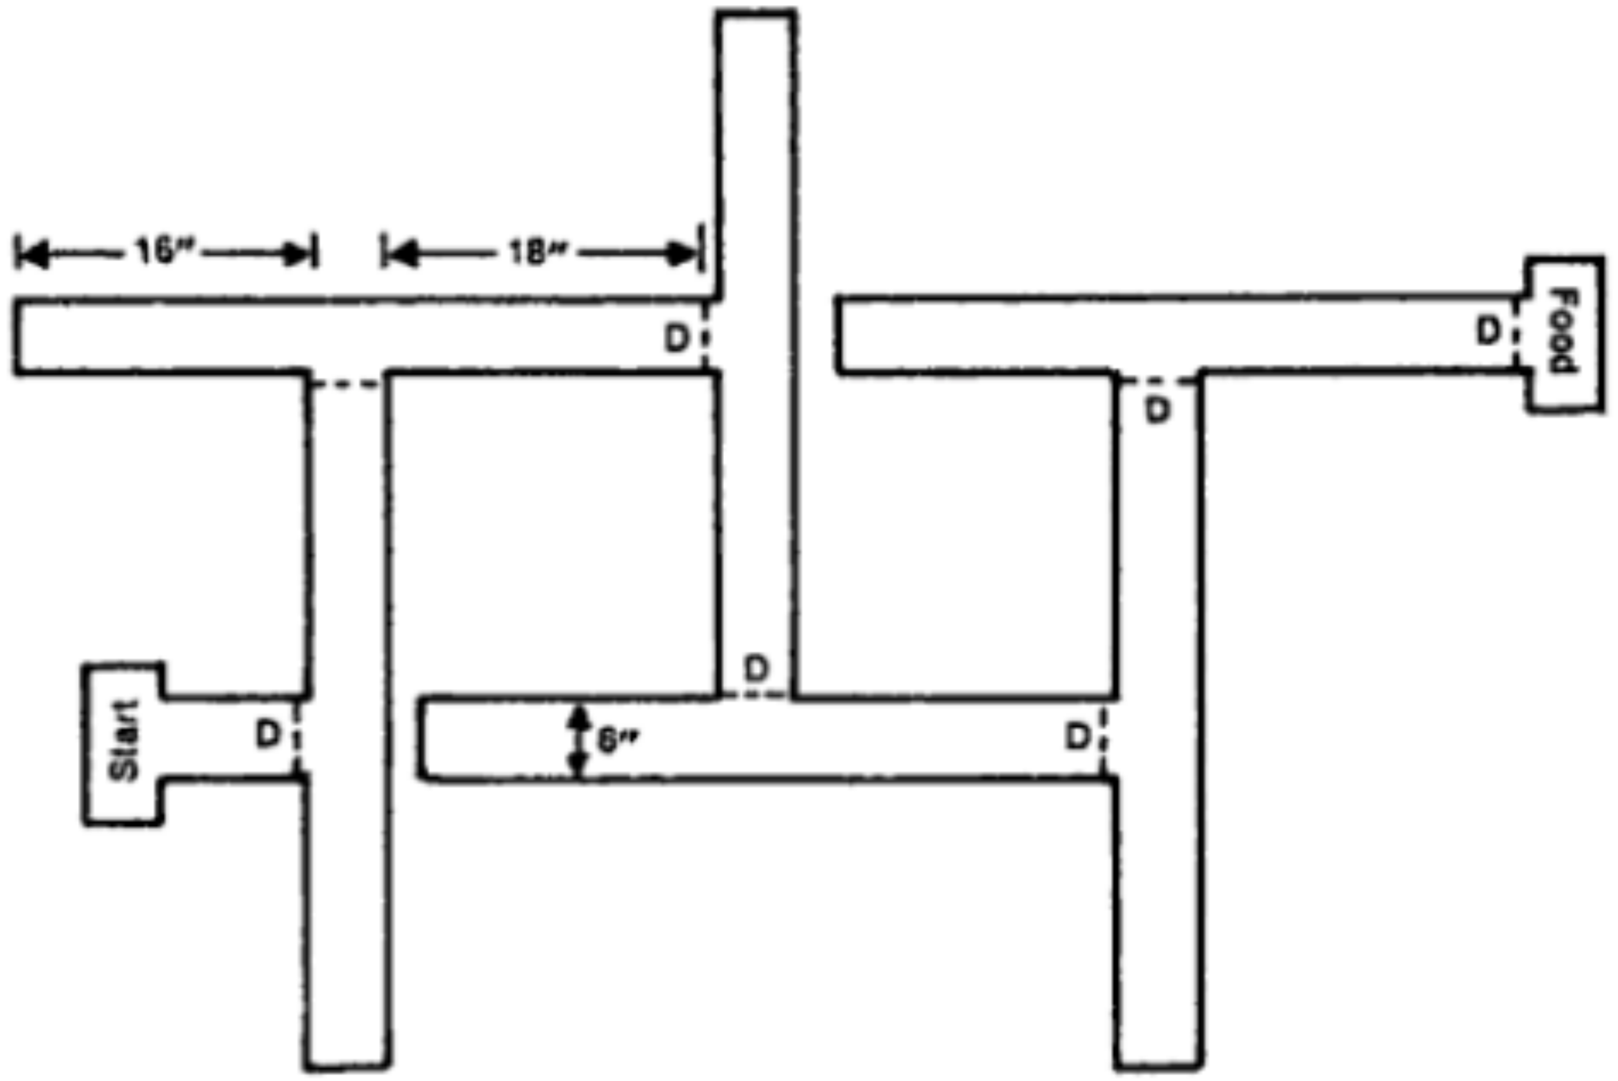

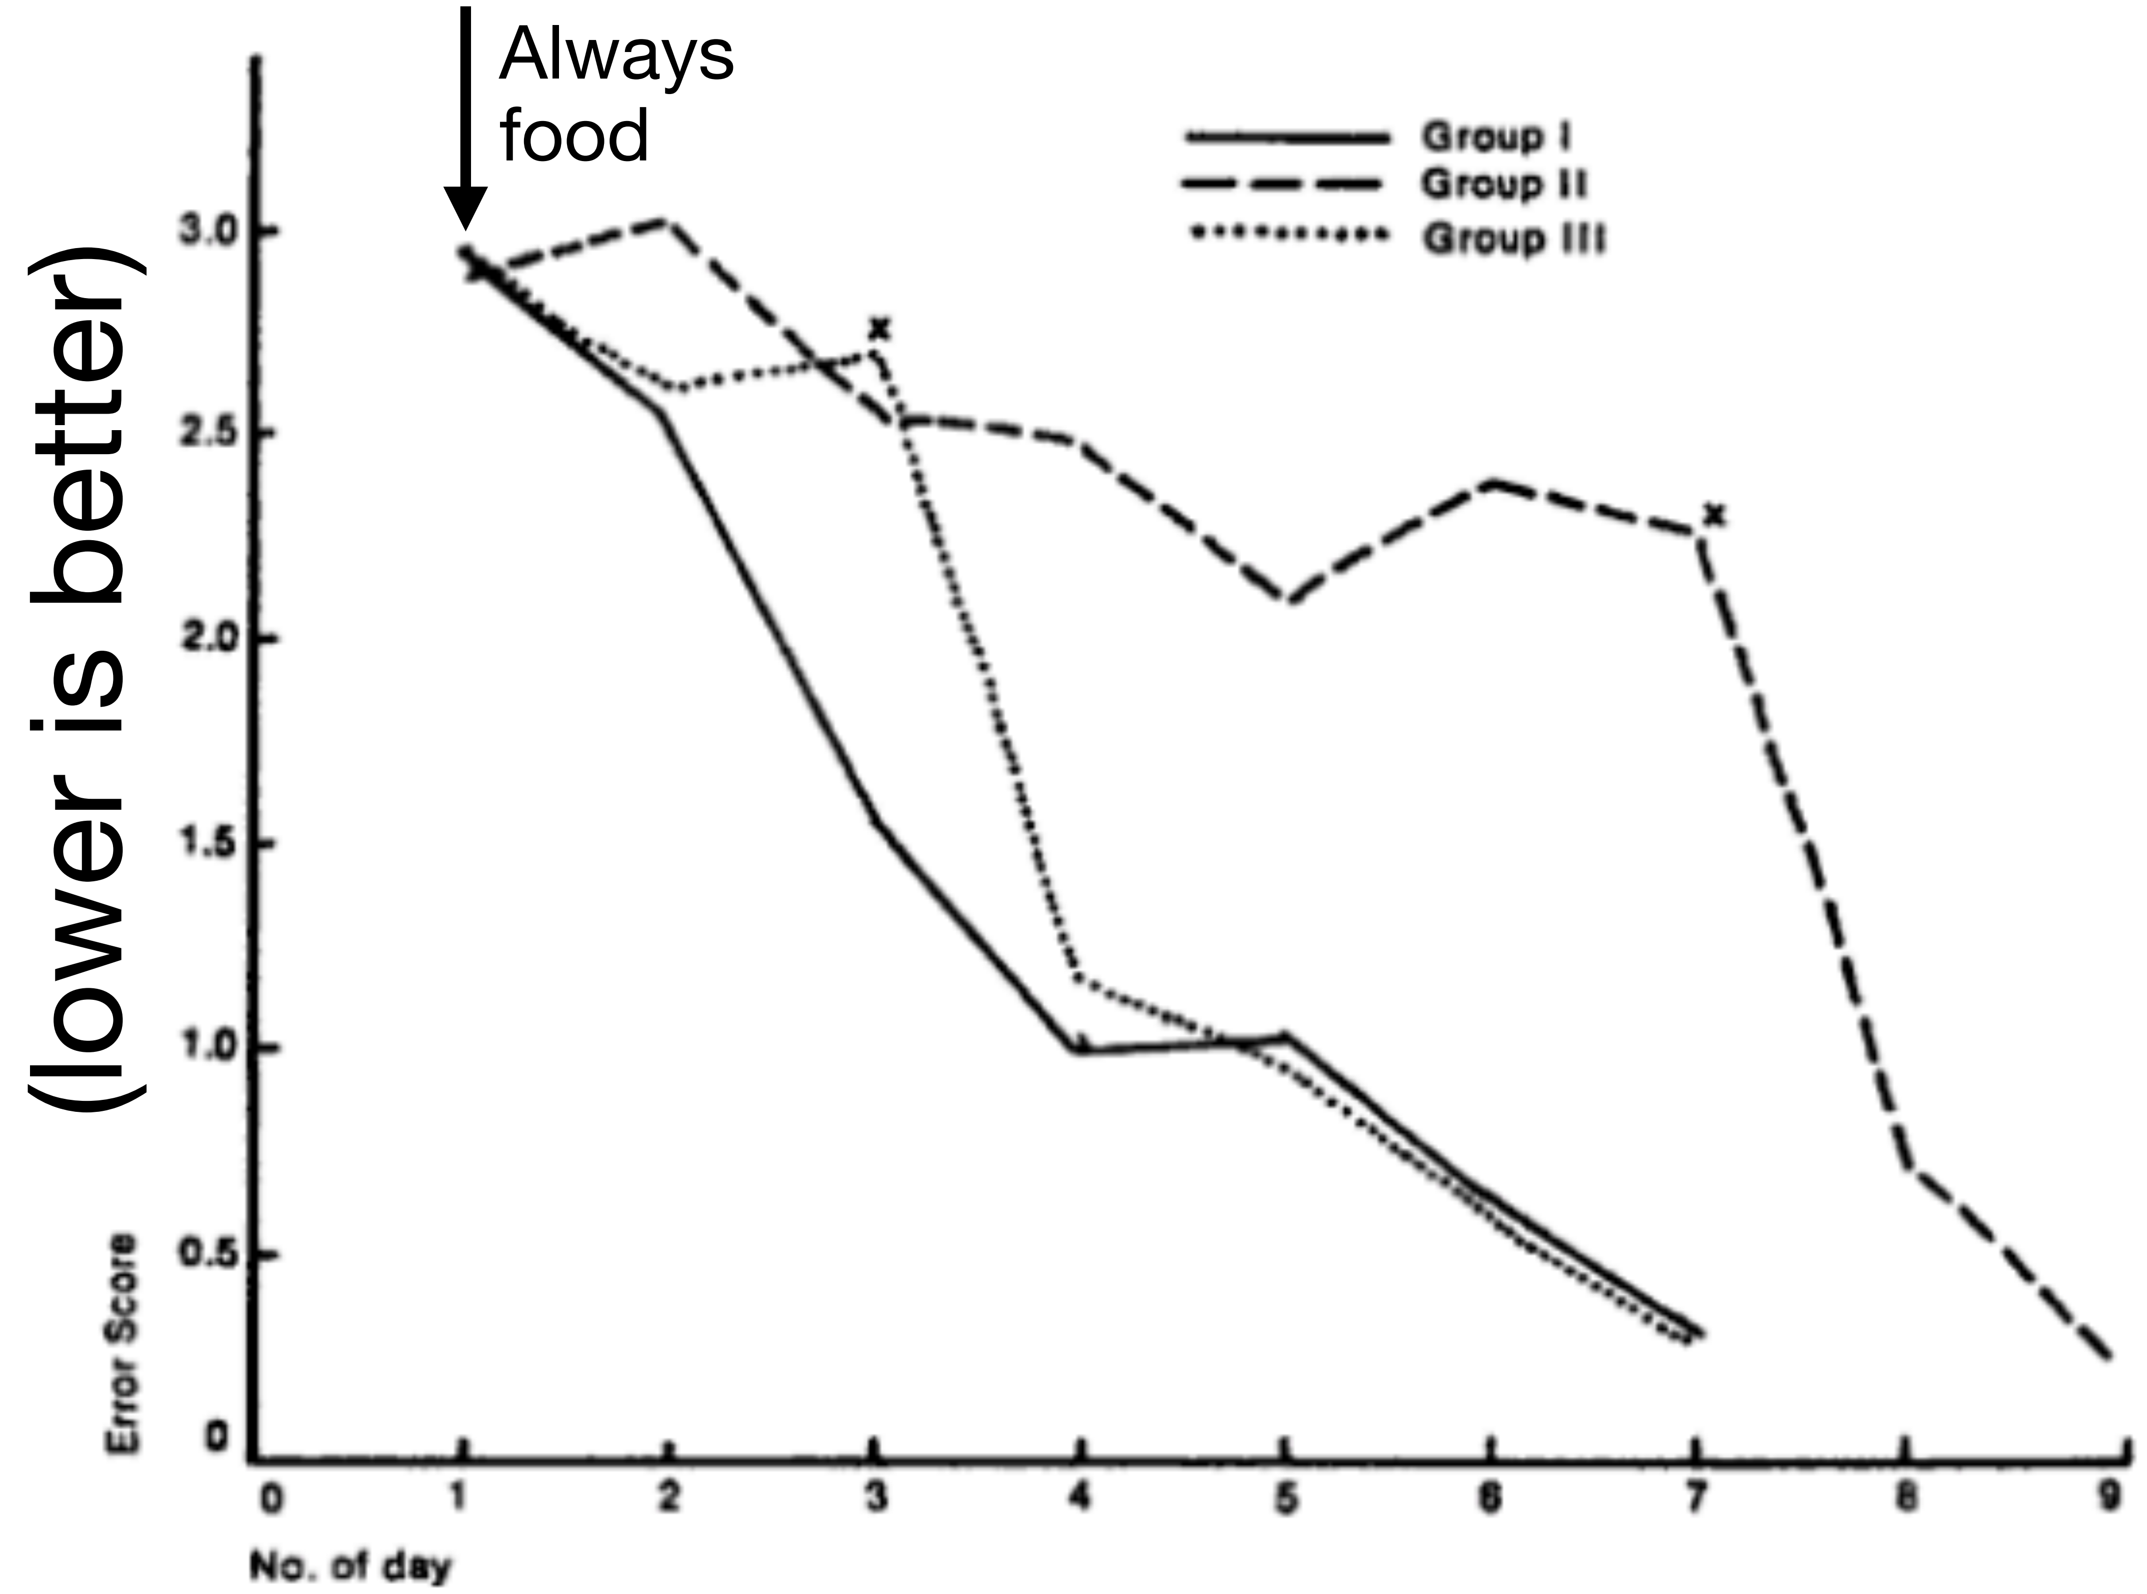

Latent Learning

- Blodgett (1929) Maze navigation task
  - **Group 1** [Control]: one trial a day with food in the goal box at the end
  - **Group 2** [Late food] No food in the maze for days 1-6, then food provided at the end on day 7
  - **Group 3** [Early food] … food added on day 3
- Learning curves dropped dramatically when food was added
  - This suggests latent learning prior to reward
  - "They had been building up a 'map'"
  - Once the reward was added, they could use the map rather than starting from scratch

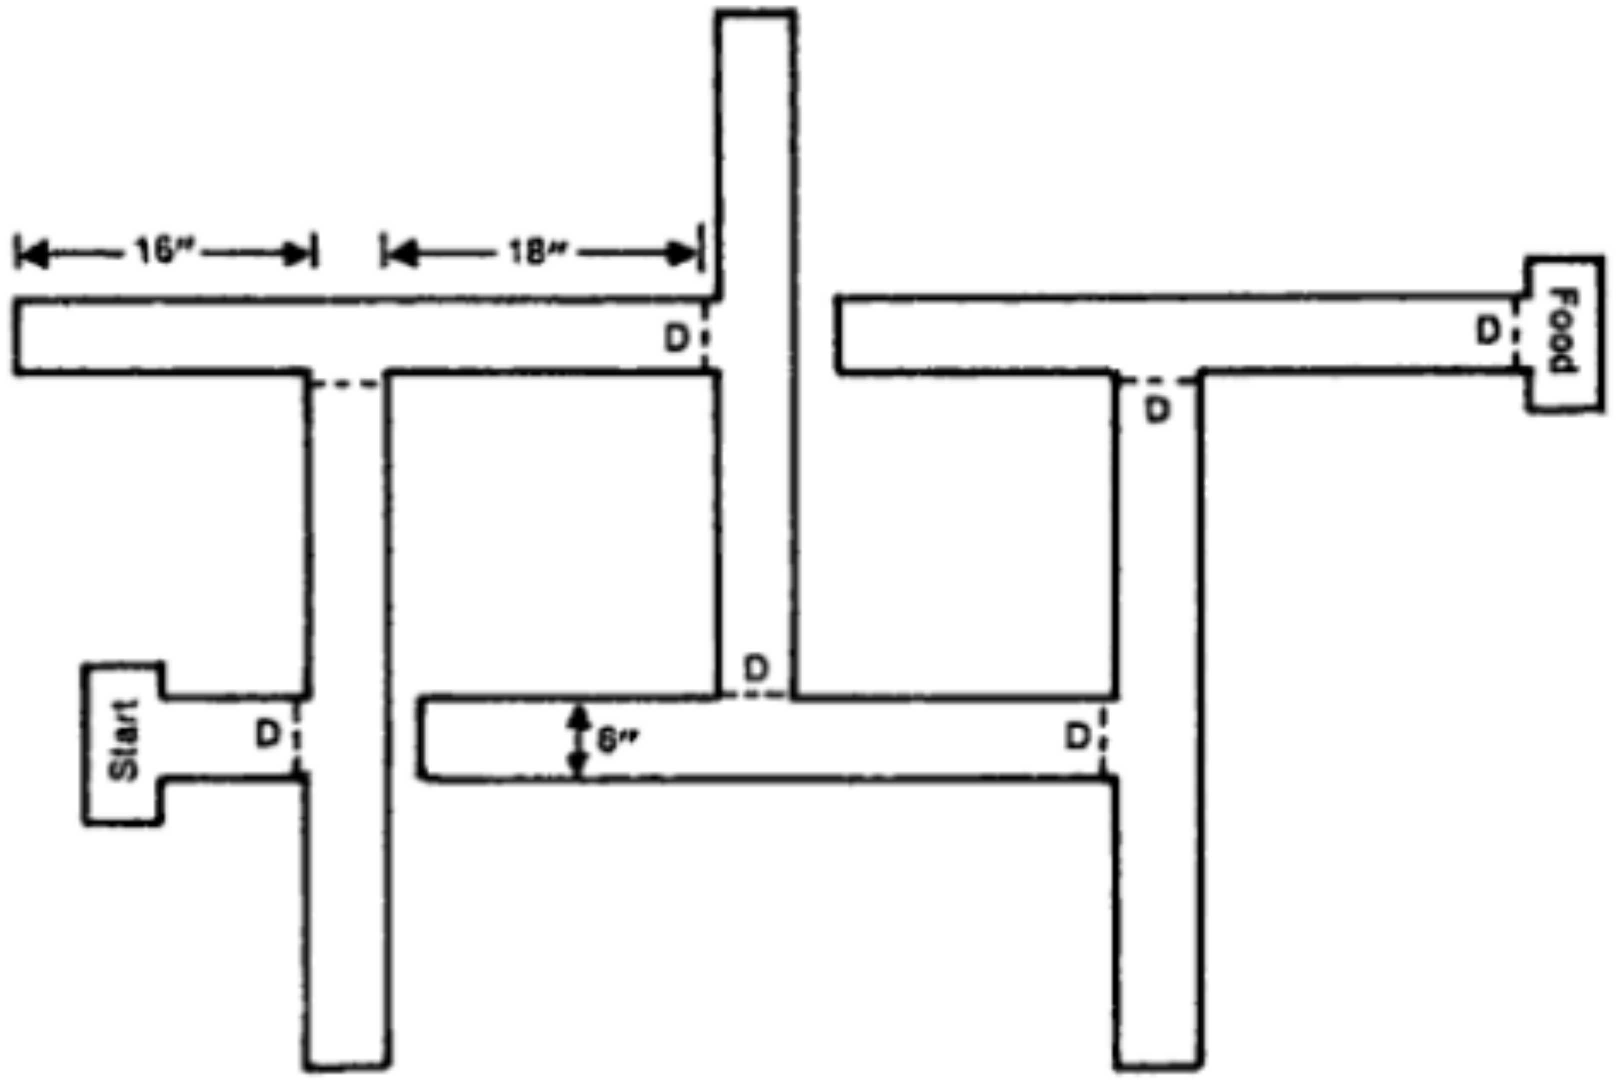

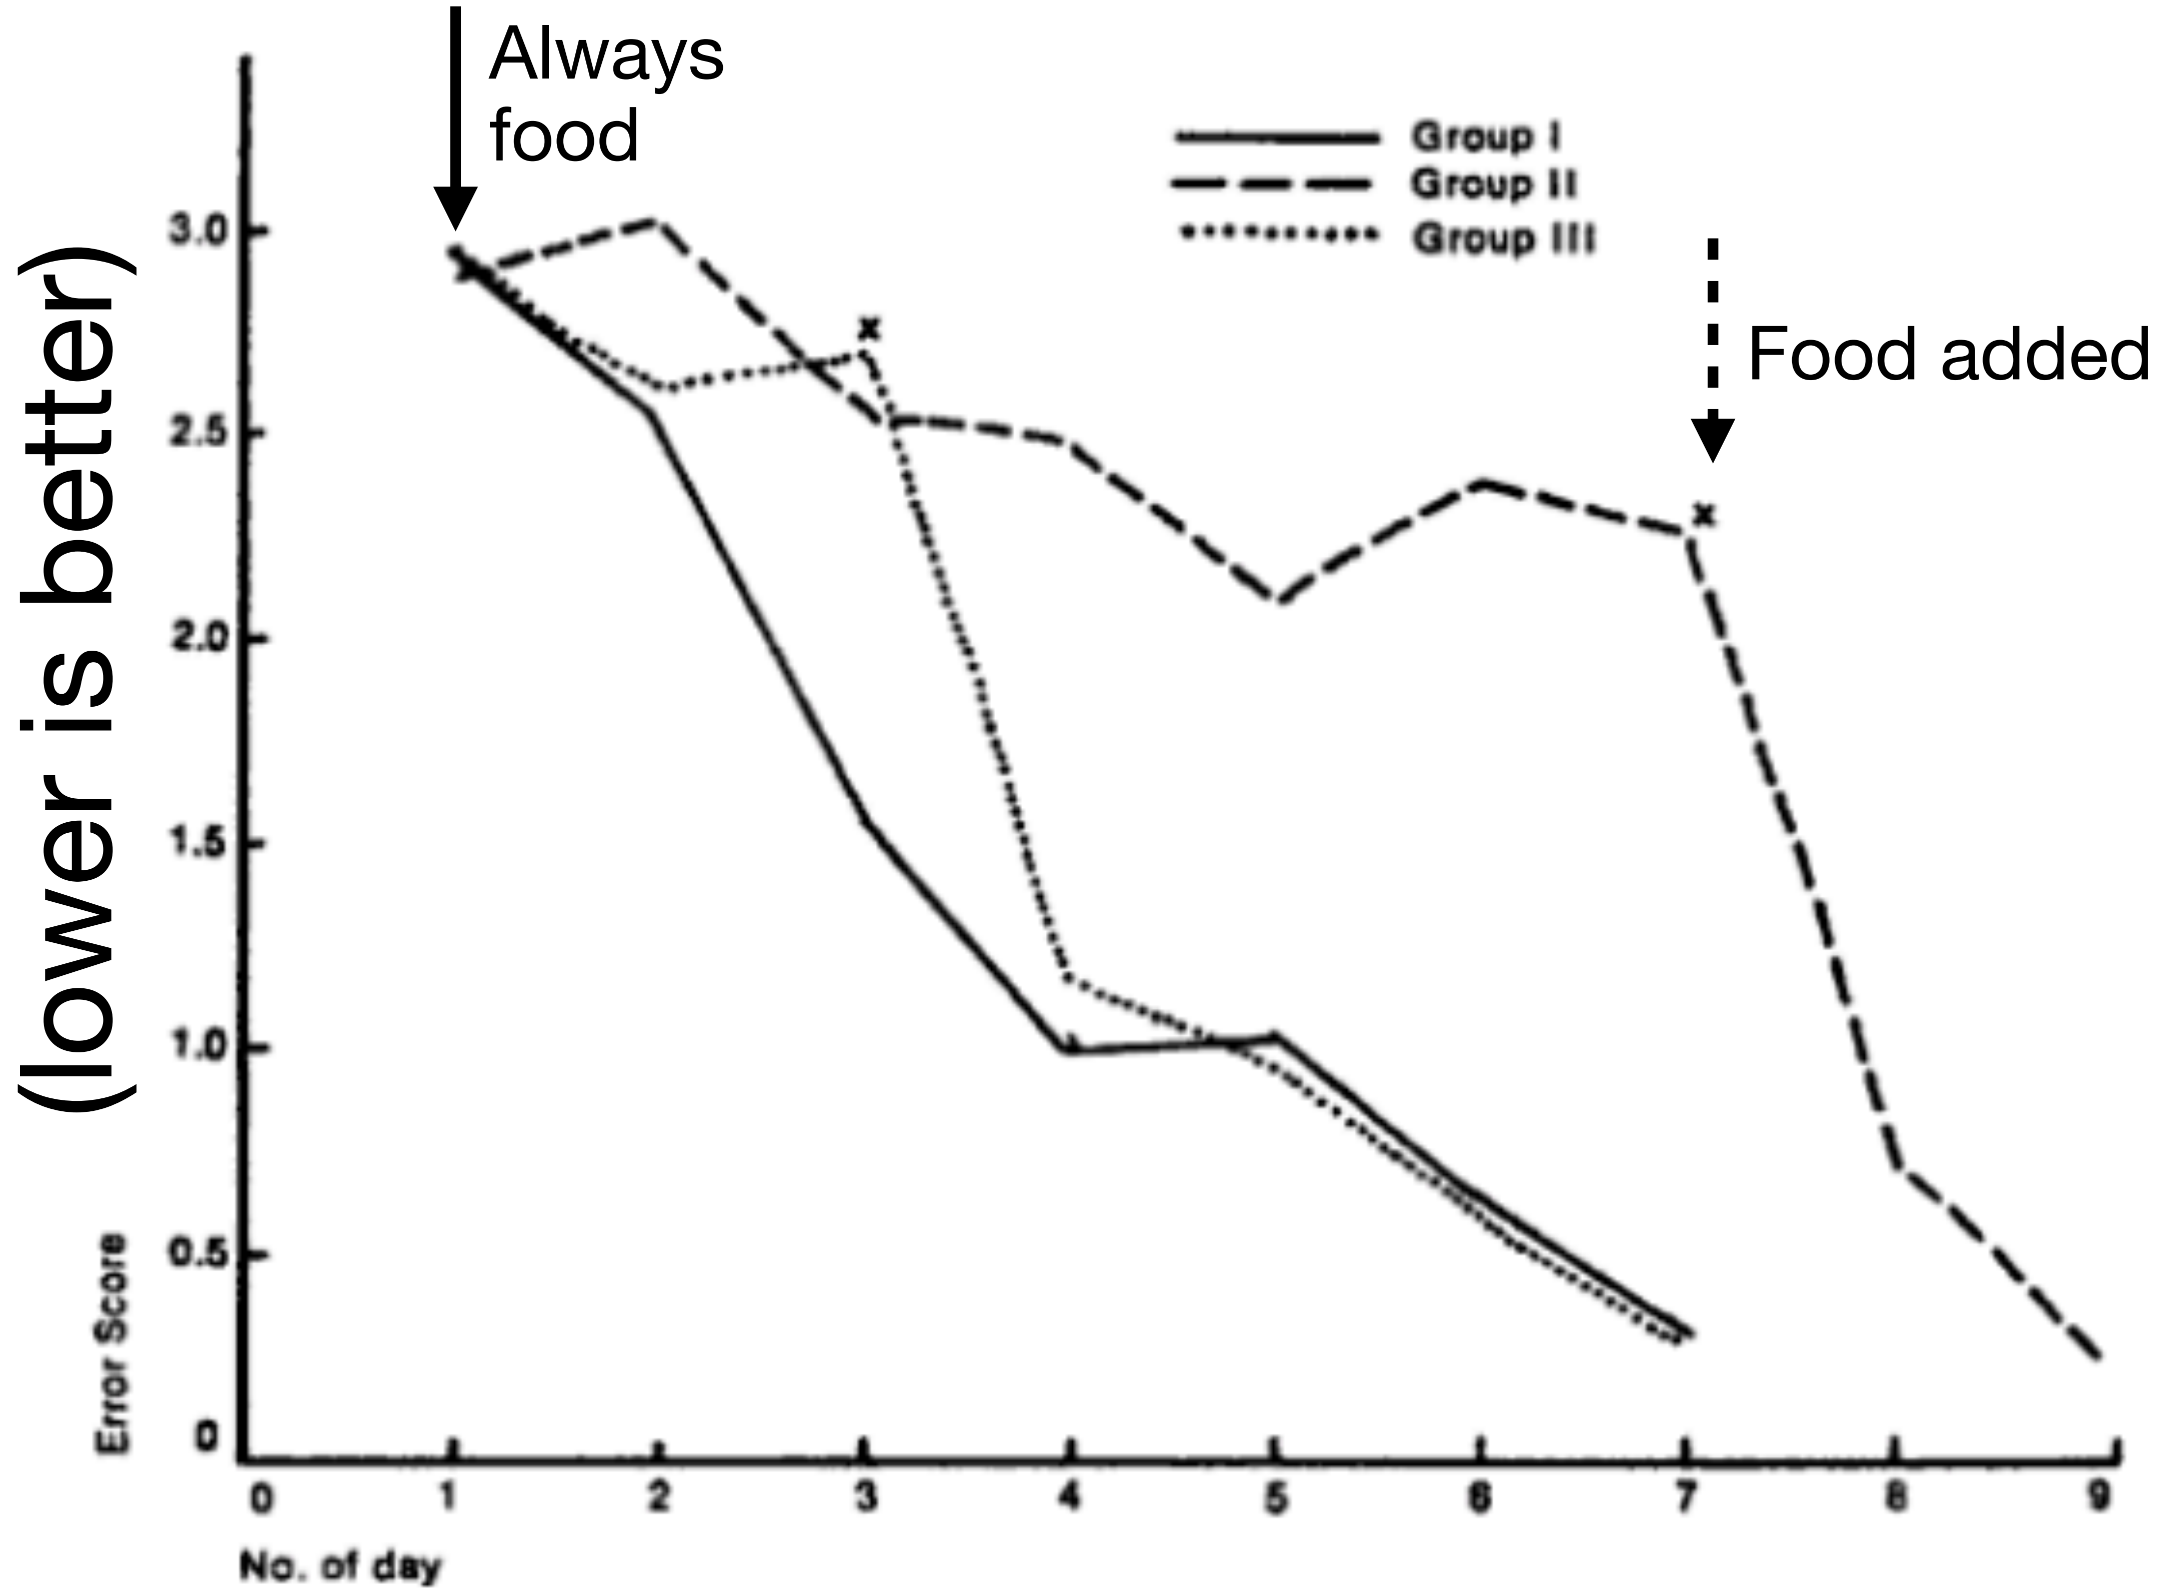

Latent Learning

- Blodgett (1929) Maze navigation task
  - **Group 1** [Control]: one trial a day with food in the goal box at the end
  - **Group 2** [Late food] No food in the maze for days 1-6, then food provided at the end on day 7
  - **Group 3** [Early food] … food added on day 3
- Learning curves dropped dramatically when food was added
  - This suggests latent learning prior to reward
  - “They had been building up a ‘map’”
  - Once the reward was added, they could use the map rather than starting from scratch

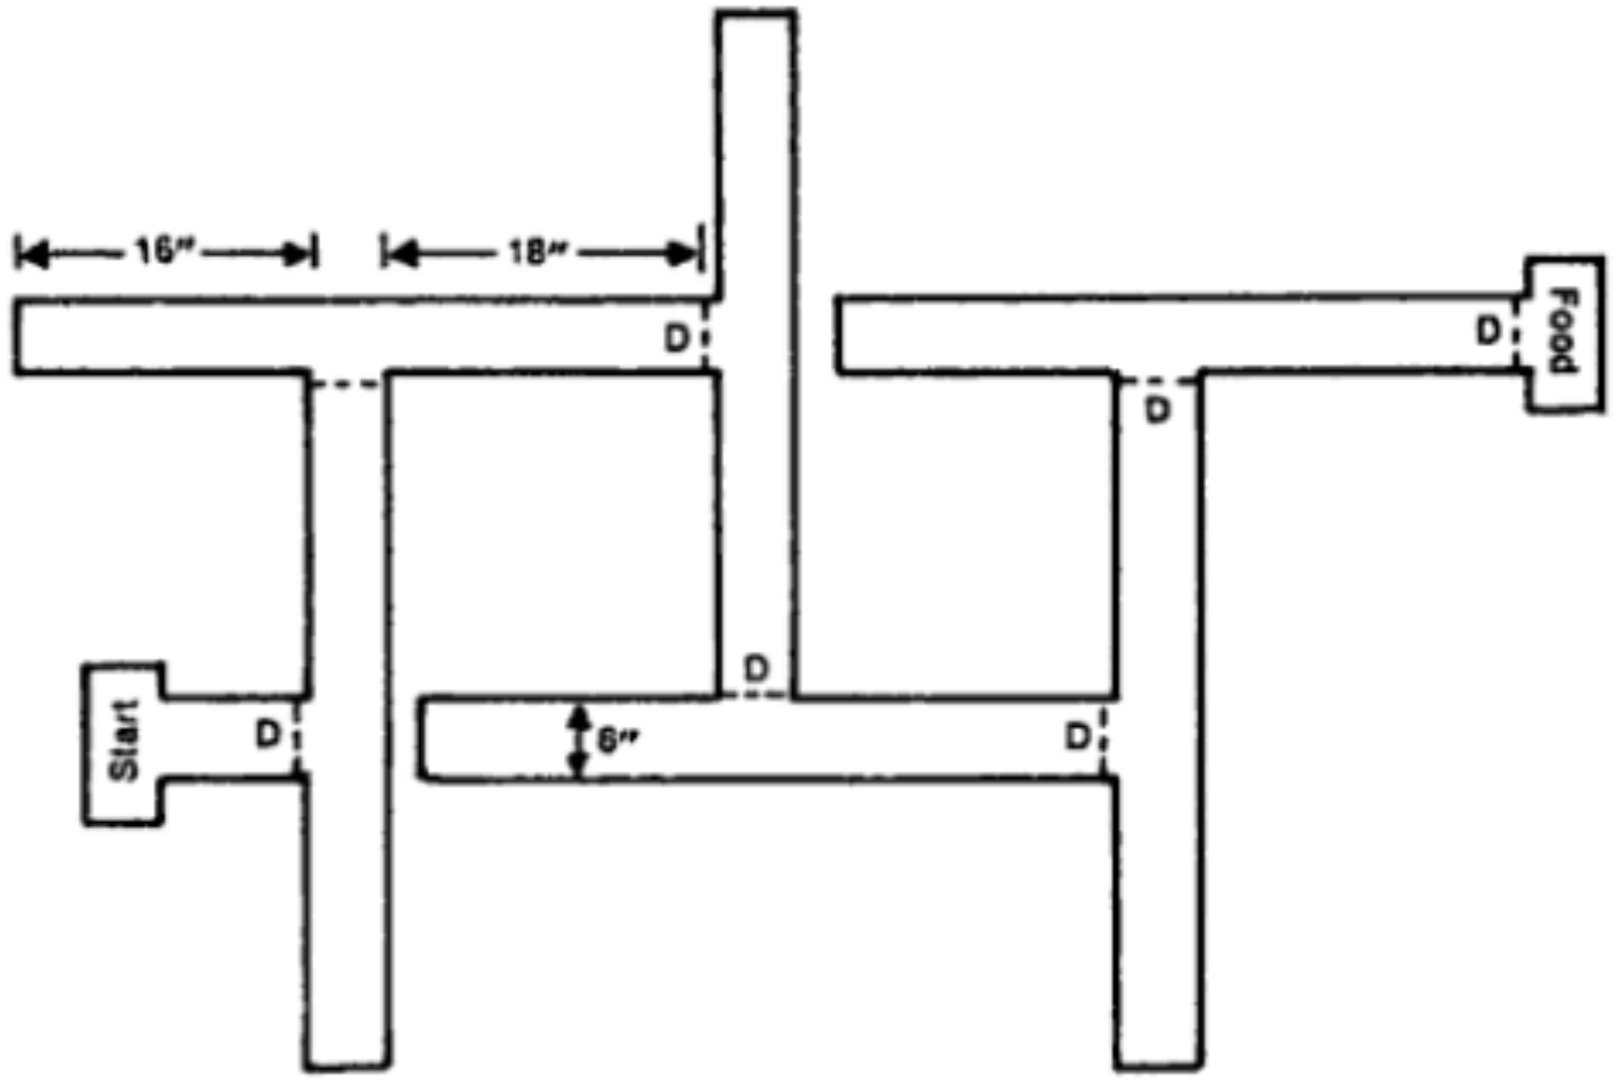

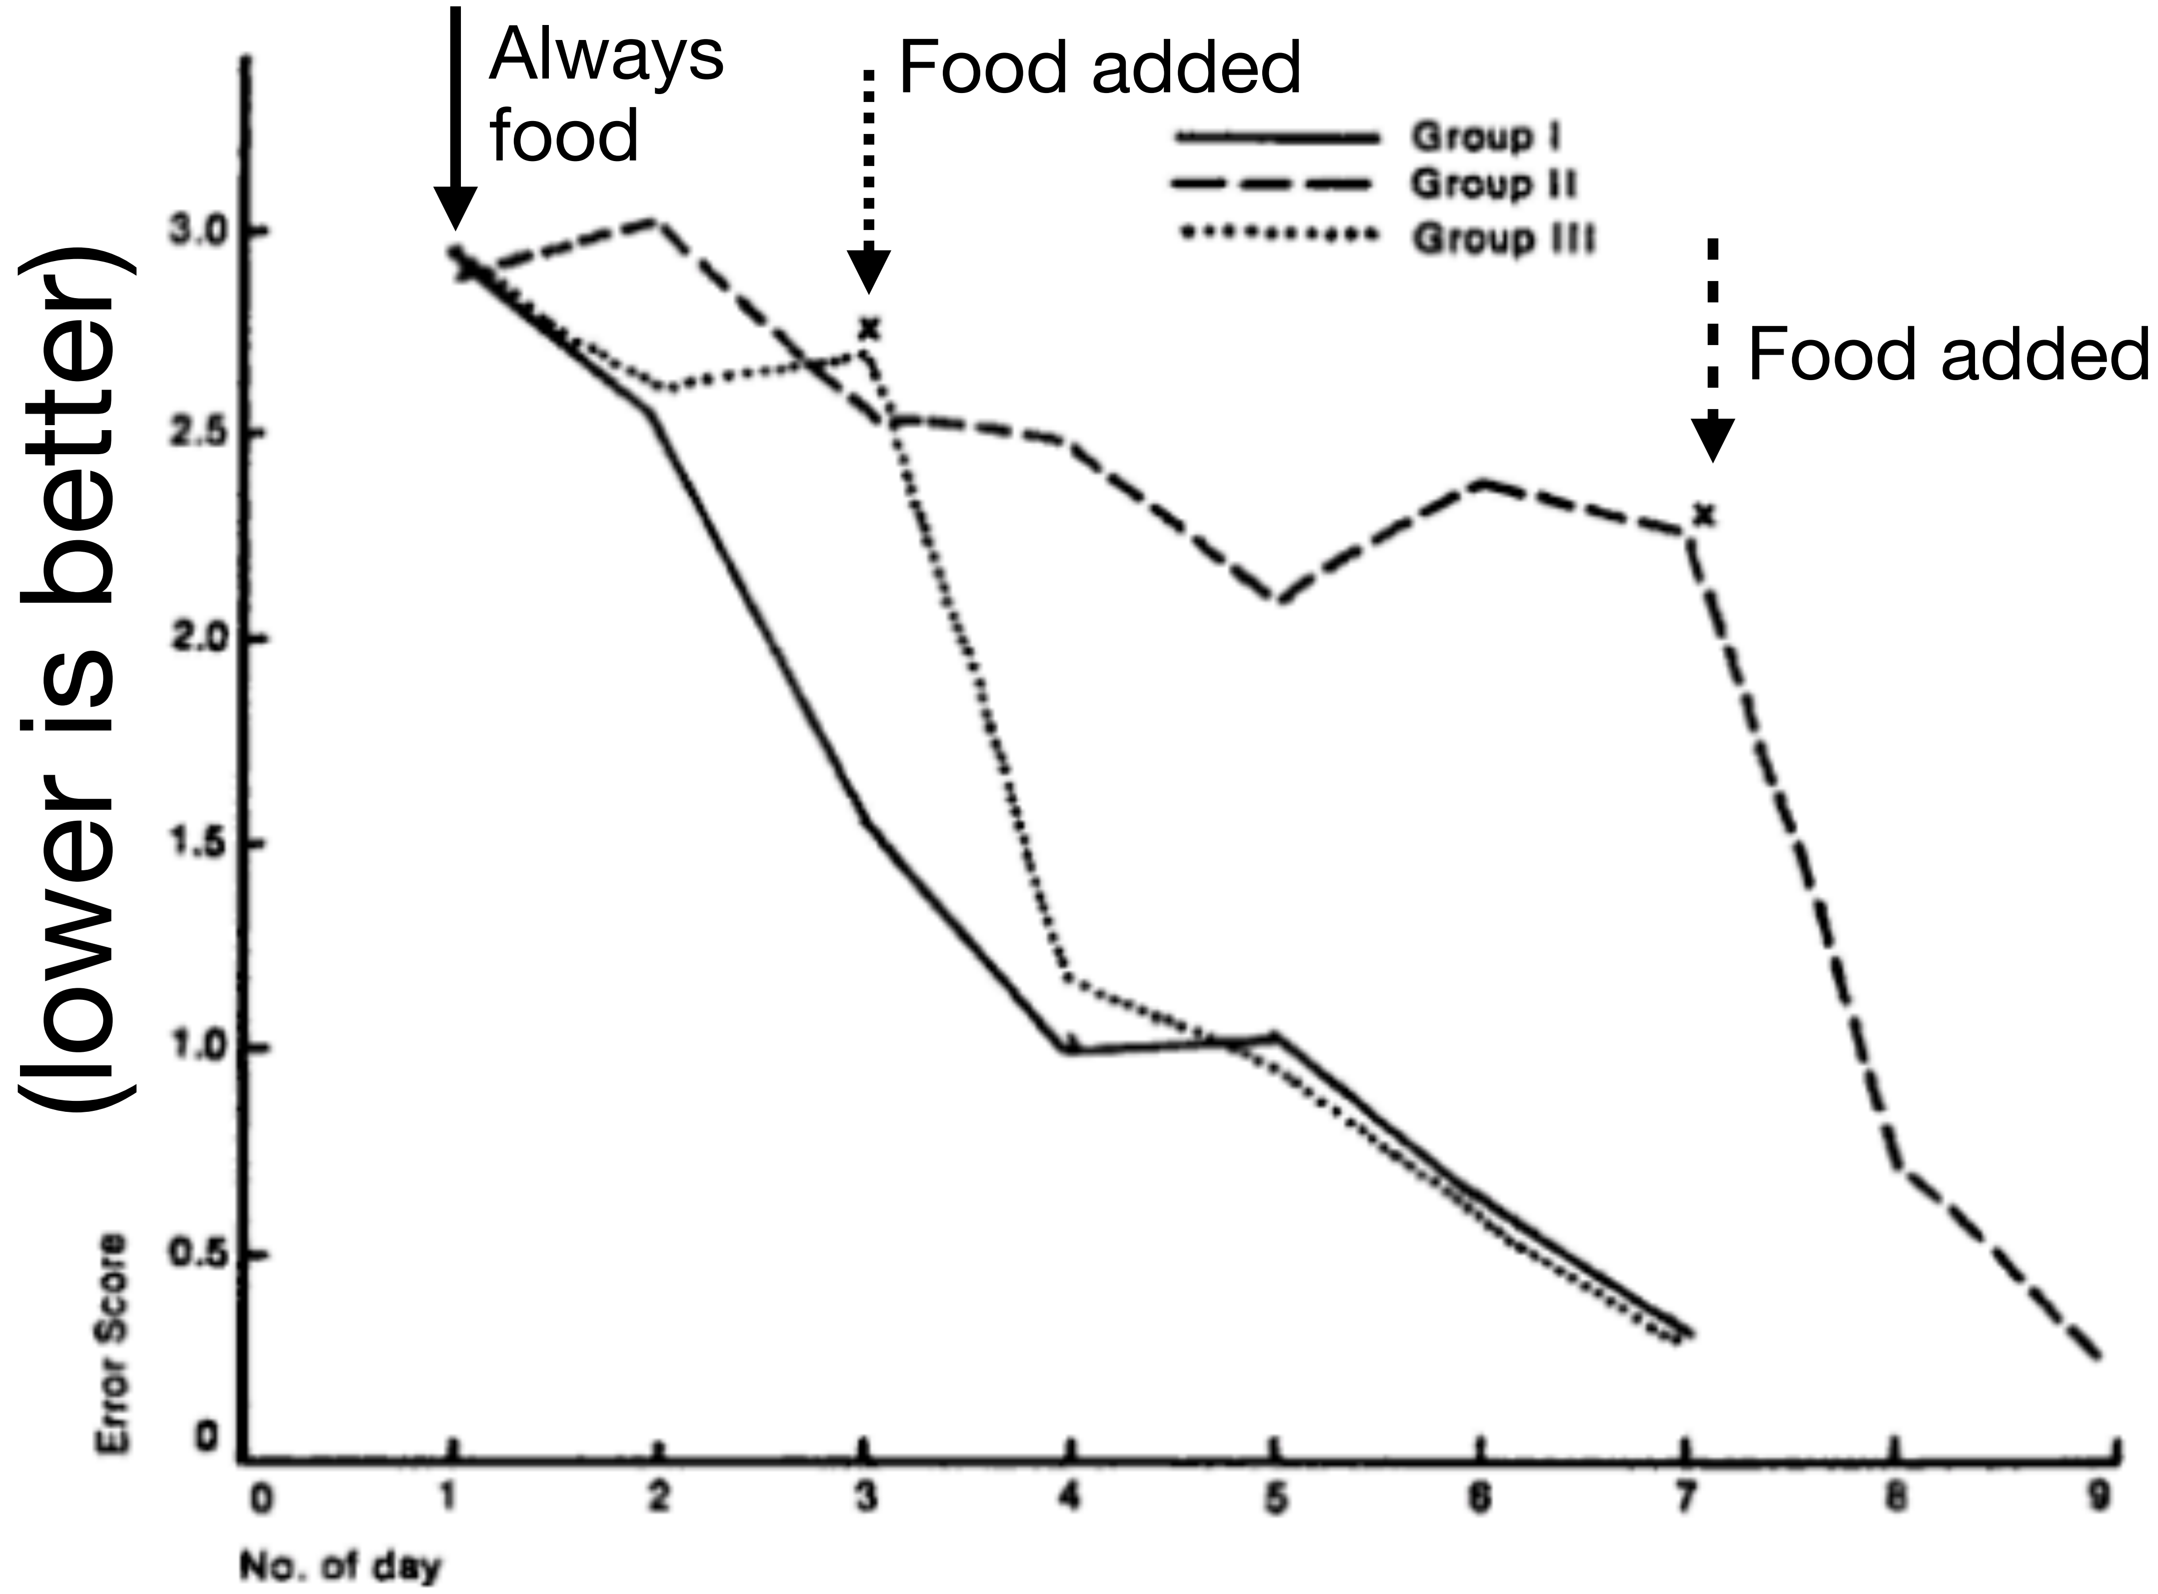

Latent Learning

- Replicates with more complex environment (Tolman & Honzik, 1930)
- **Always reward** better than **no reward**
- Adding **reward later** produces the same dramatic drop in error

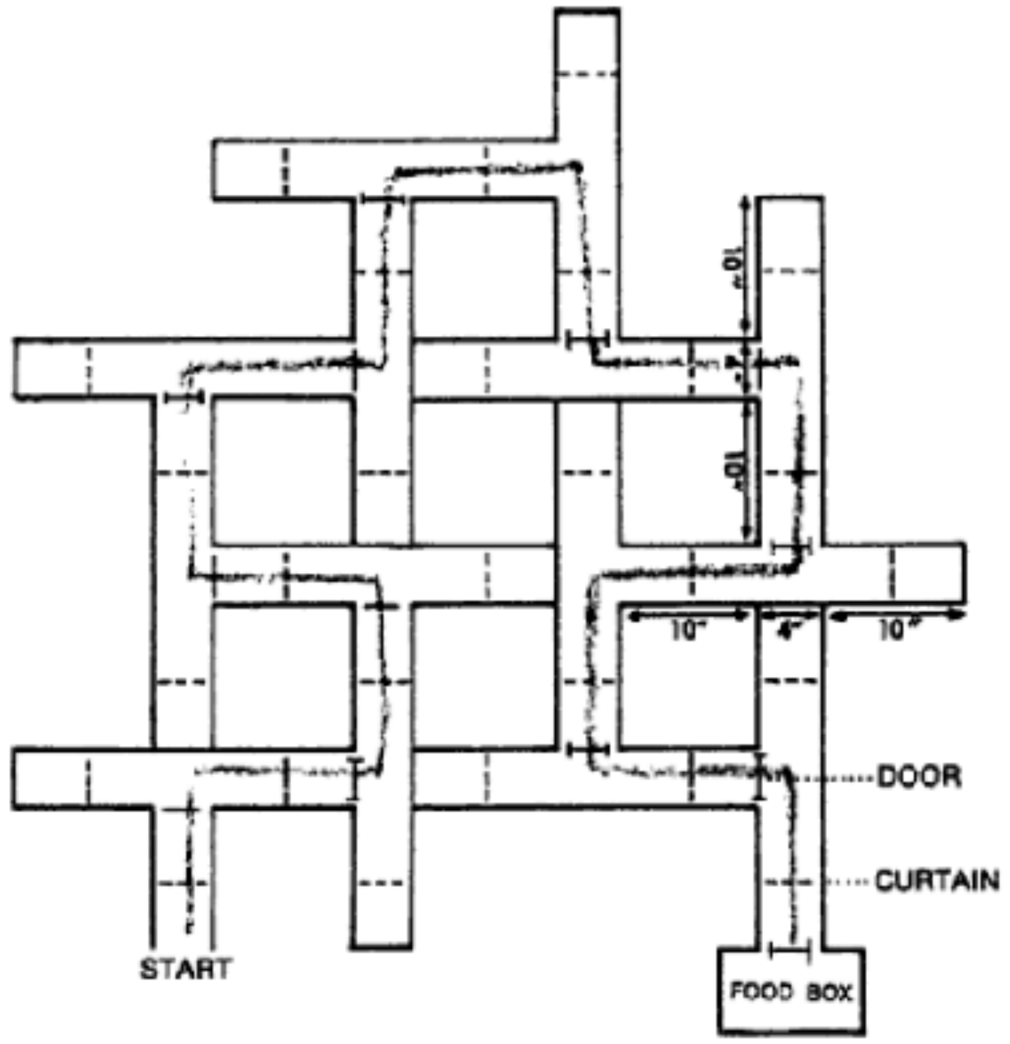

Tolman & Honzik (1930)

Latent Learning

- Replicates with more complex environment (Tolman & Honzik, 1930)
- **Always reward** better than **no reward**
- Adding **reward later** produces the same dramatic drop in error

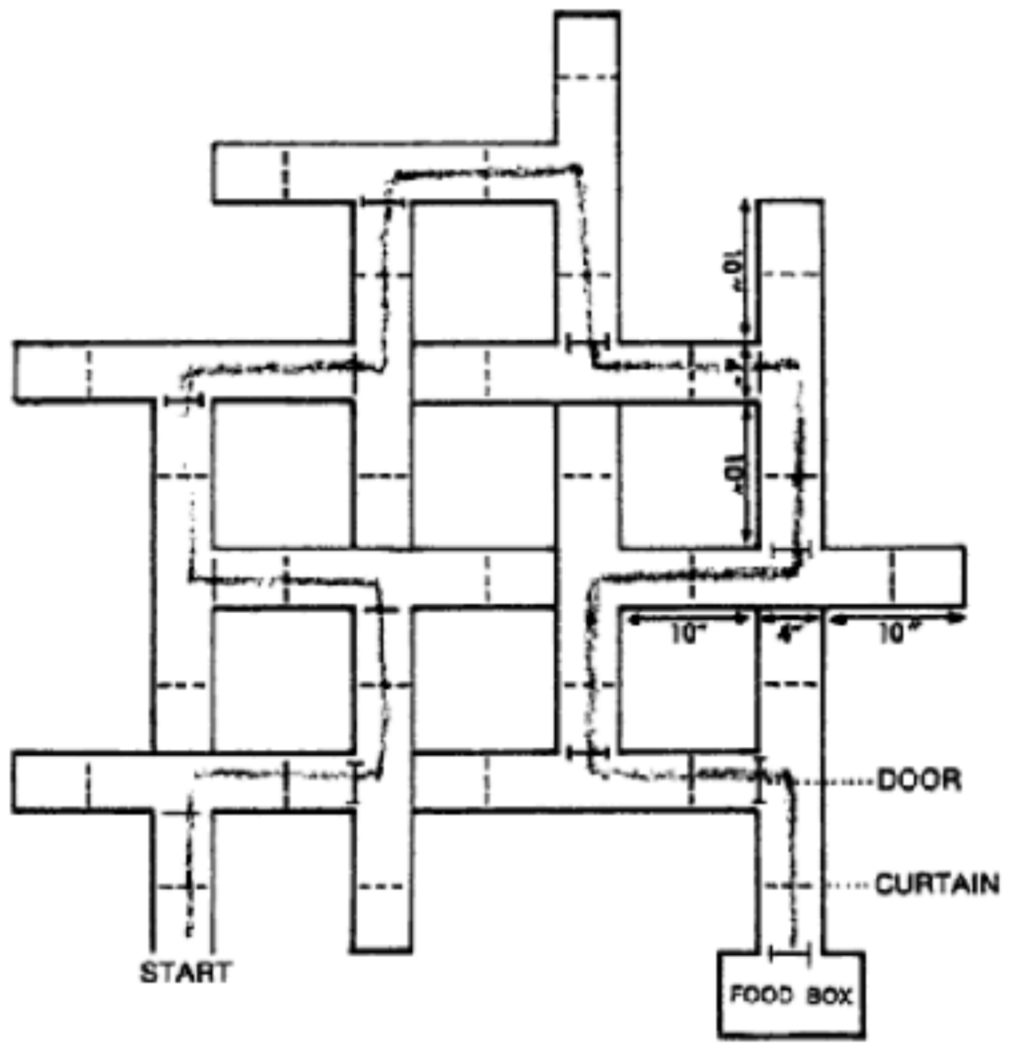

Tolman & Honzik (1930)

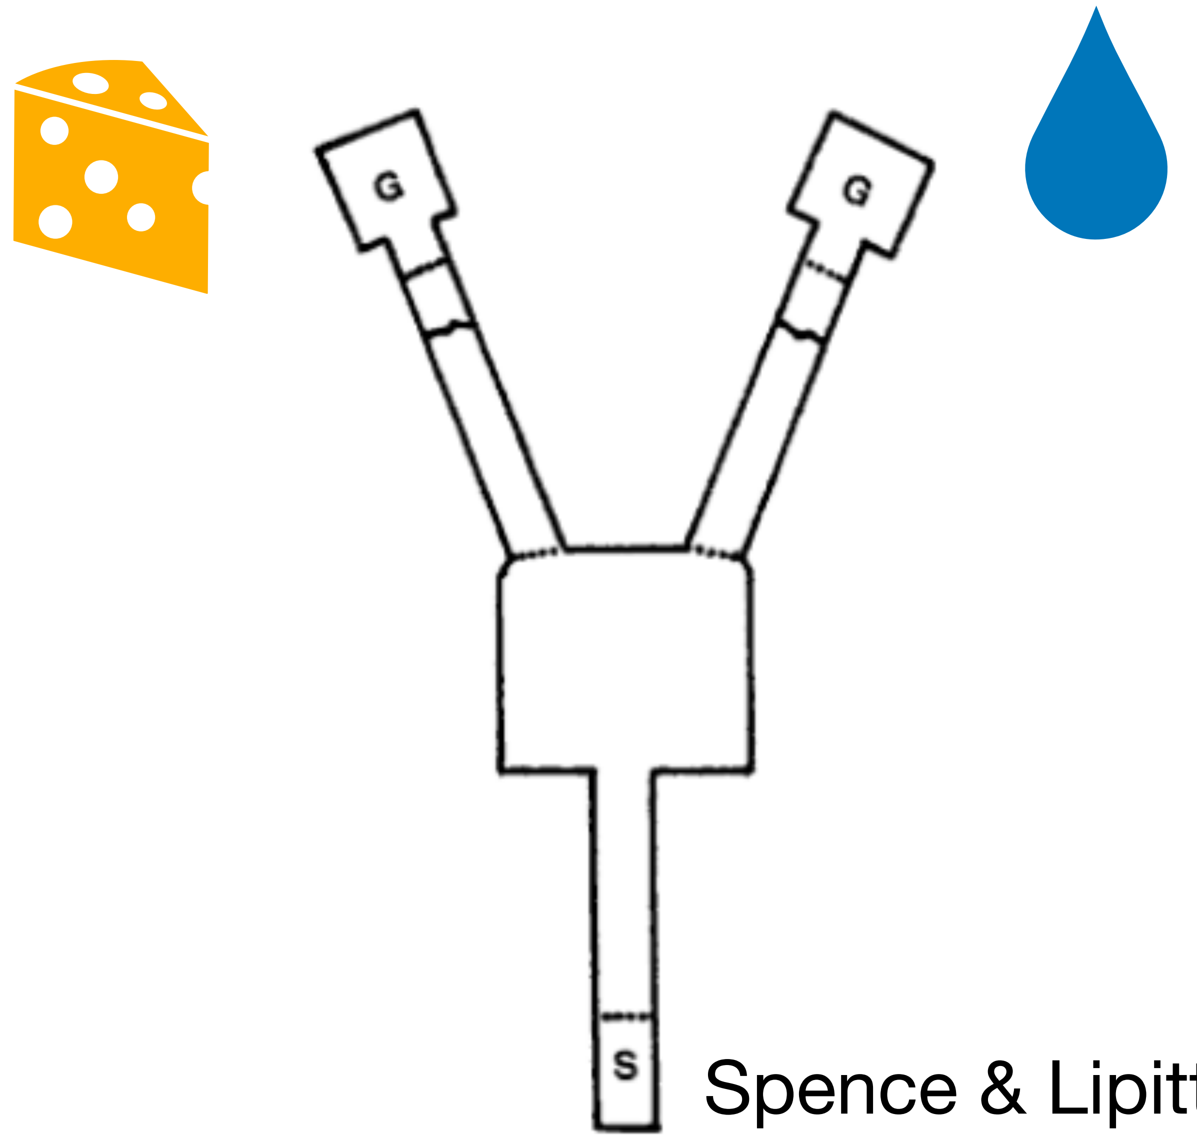

Spence & Lipitt (1946)

- **Y-maze** with separate food + water rewards
- Rats exposed to maze while satiated (no **hunger** + no **thirst**)
  - One group reintroduced when **hungry** goes left towards
  - Another group reintroduced when **thirsty** goes right towards

Vicarious Trial and Error (VTE)

- Animal put on jumping stand, facing two doors (l vs. r) with different visual properties (e.g., horizontal vs. vertical stripes)
    - One door is correct, the other incorrect
    - location is randomly swapped but visual features are predictive
    - If the animal jumps towards the correct door, it opens and reveals food on a platform behind… and if incorrect ….
- Tolman (1939) added landing platforms infront of the doors
    - When the choice was easy (black vs. white stimuli), the animals learned quicker and did more VTEing than for hard problems
    - After learning had been established, VTEs went down
    - Better learners also did more VTEing (Geier, LEvin & Tolman, 1941)

Muenzinger (1938)

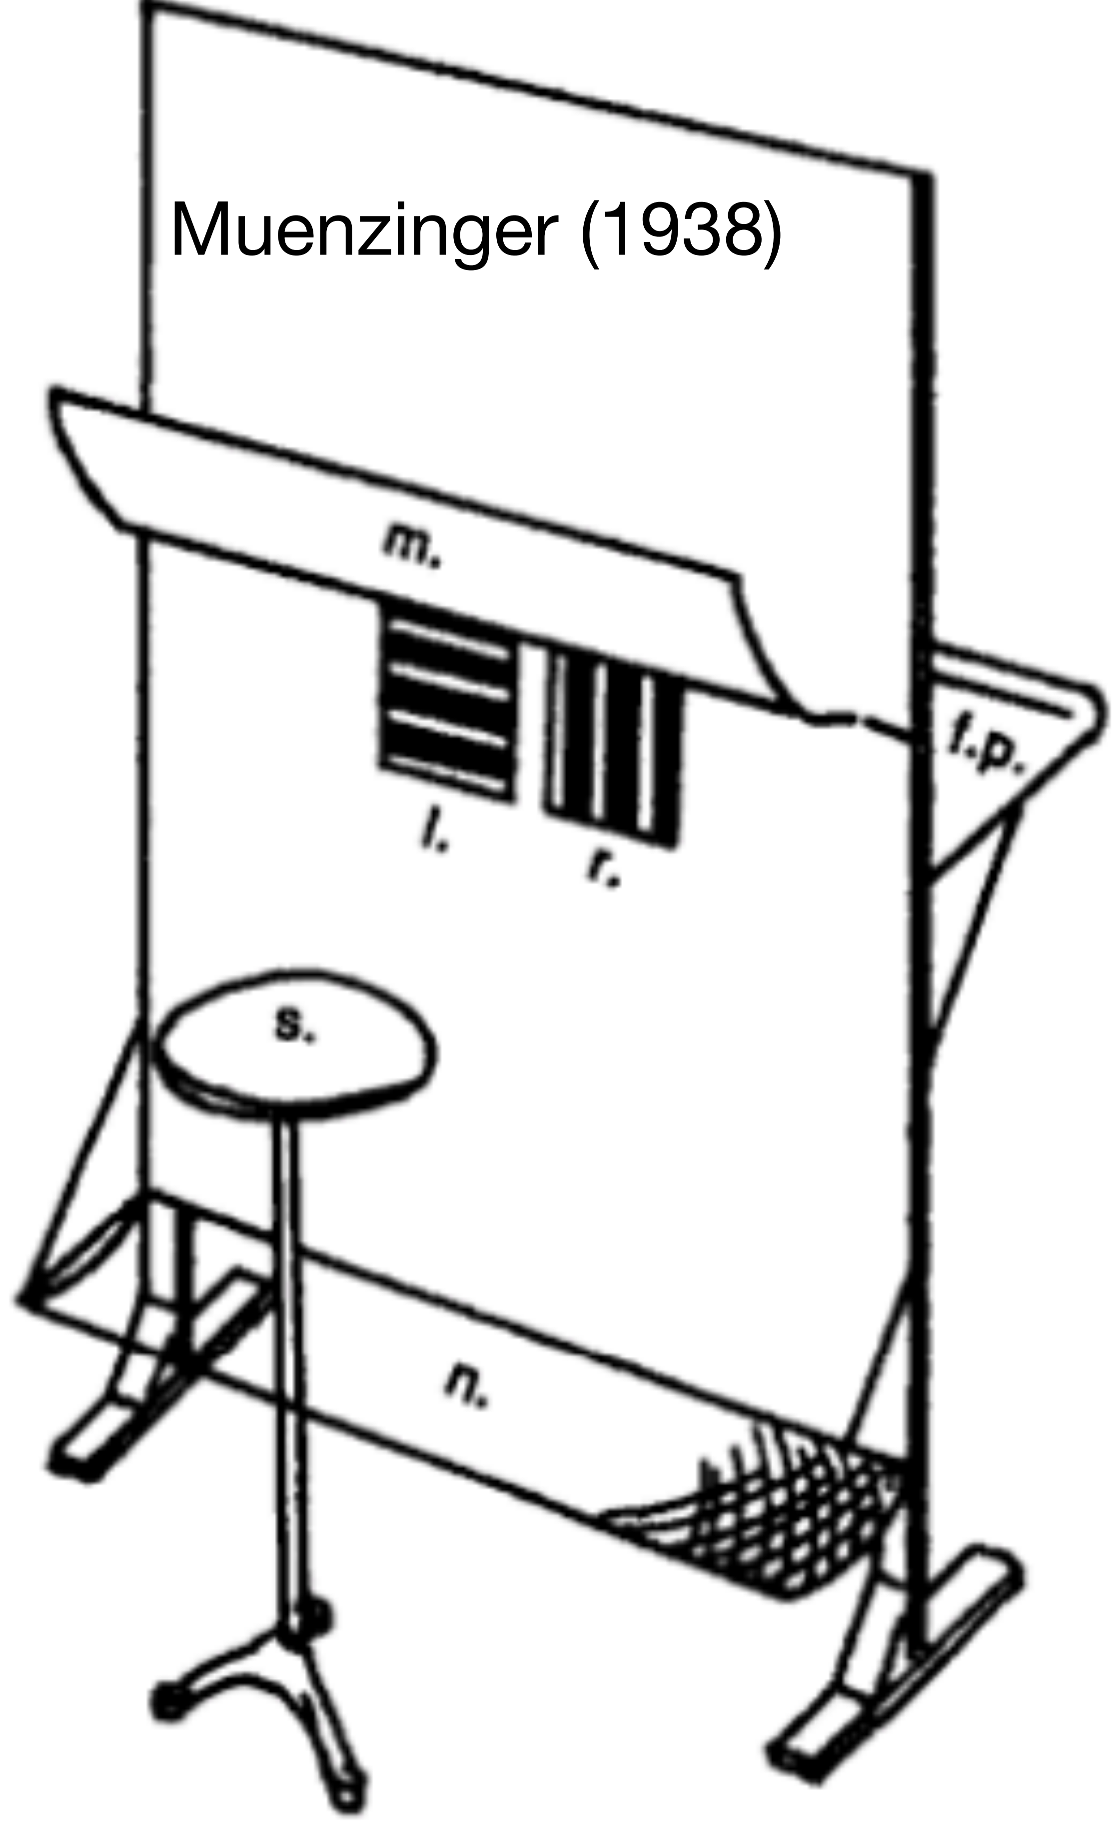

Vicarious Trial and Error (VTE)

- Animal put on jumping stand, facing two doors (l vs. r) with different visual properties (e.g., horizontal vs. vertical stripes)
    - One door is correct, the other incorrect
    - location is randomly swapped but visual features are predictive
    - If the animal jumps towards the correct door, it opens and reveals food on a platform behind… and if incorrect ….
- Tolman (1939) added landing platforms infront of the doors
    - When the choice was easy (black vs. white stimuli), the animals learned quicker and did more VTEing than for hard problems
    - After learning had been established, VTEs went down
    - Better learners also did more VTEing (Geier, LEvin & Tolman, 1941)

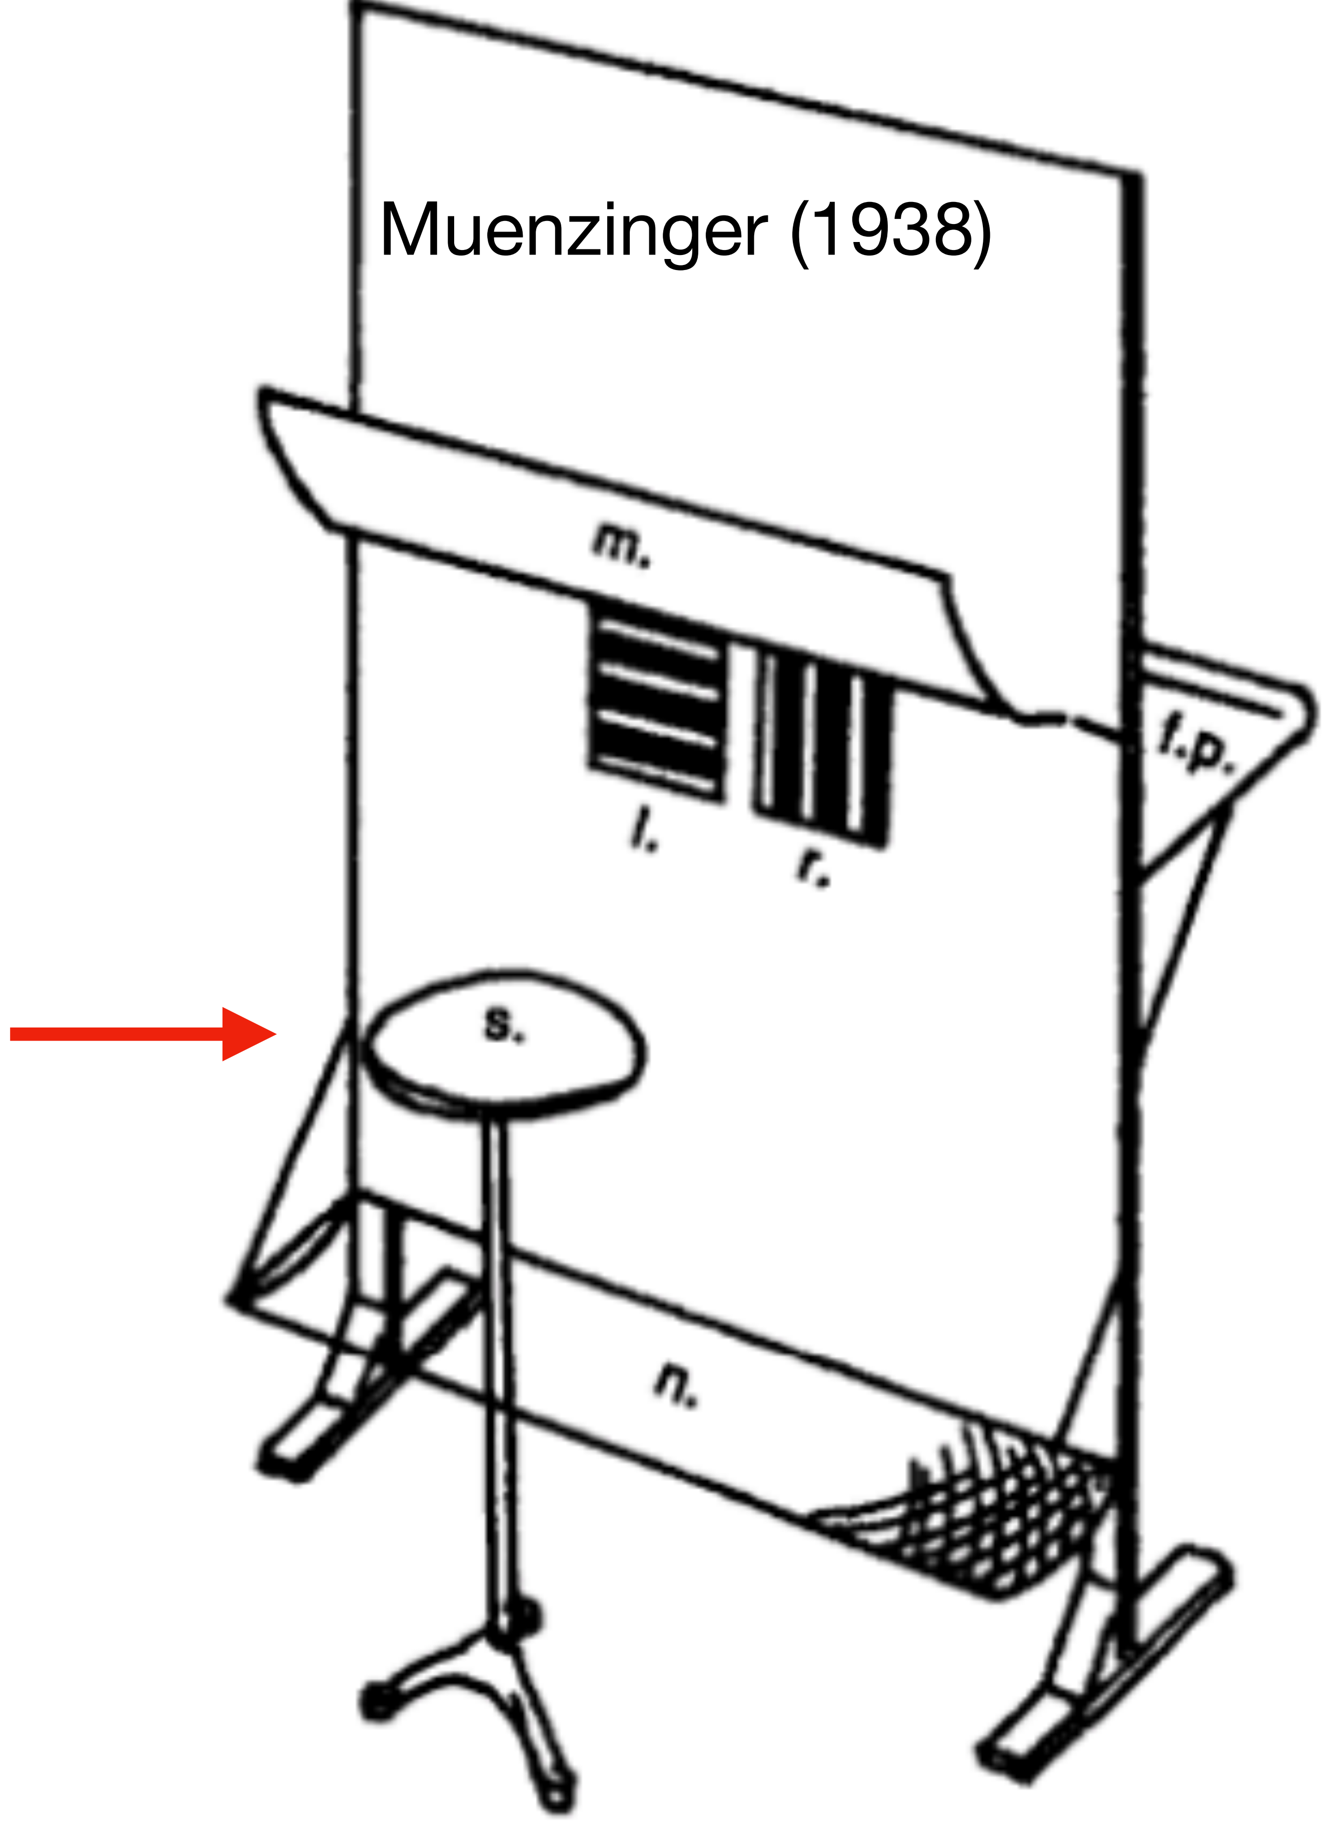

Vicarious Trial and Error (VTE)

- Animal put on jumping stand, facing two doors (l vs. r) with different visual properties (e.g., horizontal vs. vertical stripes)
  - One door is correct, the other incorrect
  - location is randomly swapped but visual features are predictive
  - If the animal jumps towards the correct door, it opens and reveals food on a platform behind… and if incorrect ….
- Tolman (1939) added landing platforms infront of the doors
  - When the choice was easy (black vs. white stimuli), the animals learned quicker and did more VTEing than for hard problems
  - After learning had been established, VTEs went down
  - Better learners also did more VTEing (Geier, LEvin & Tolman, 1941)

Muenzinger (1938)

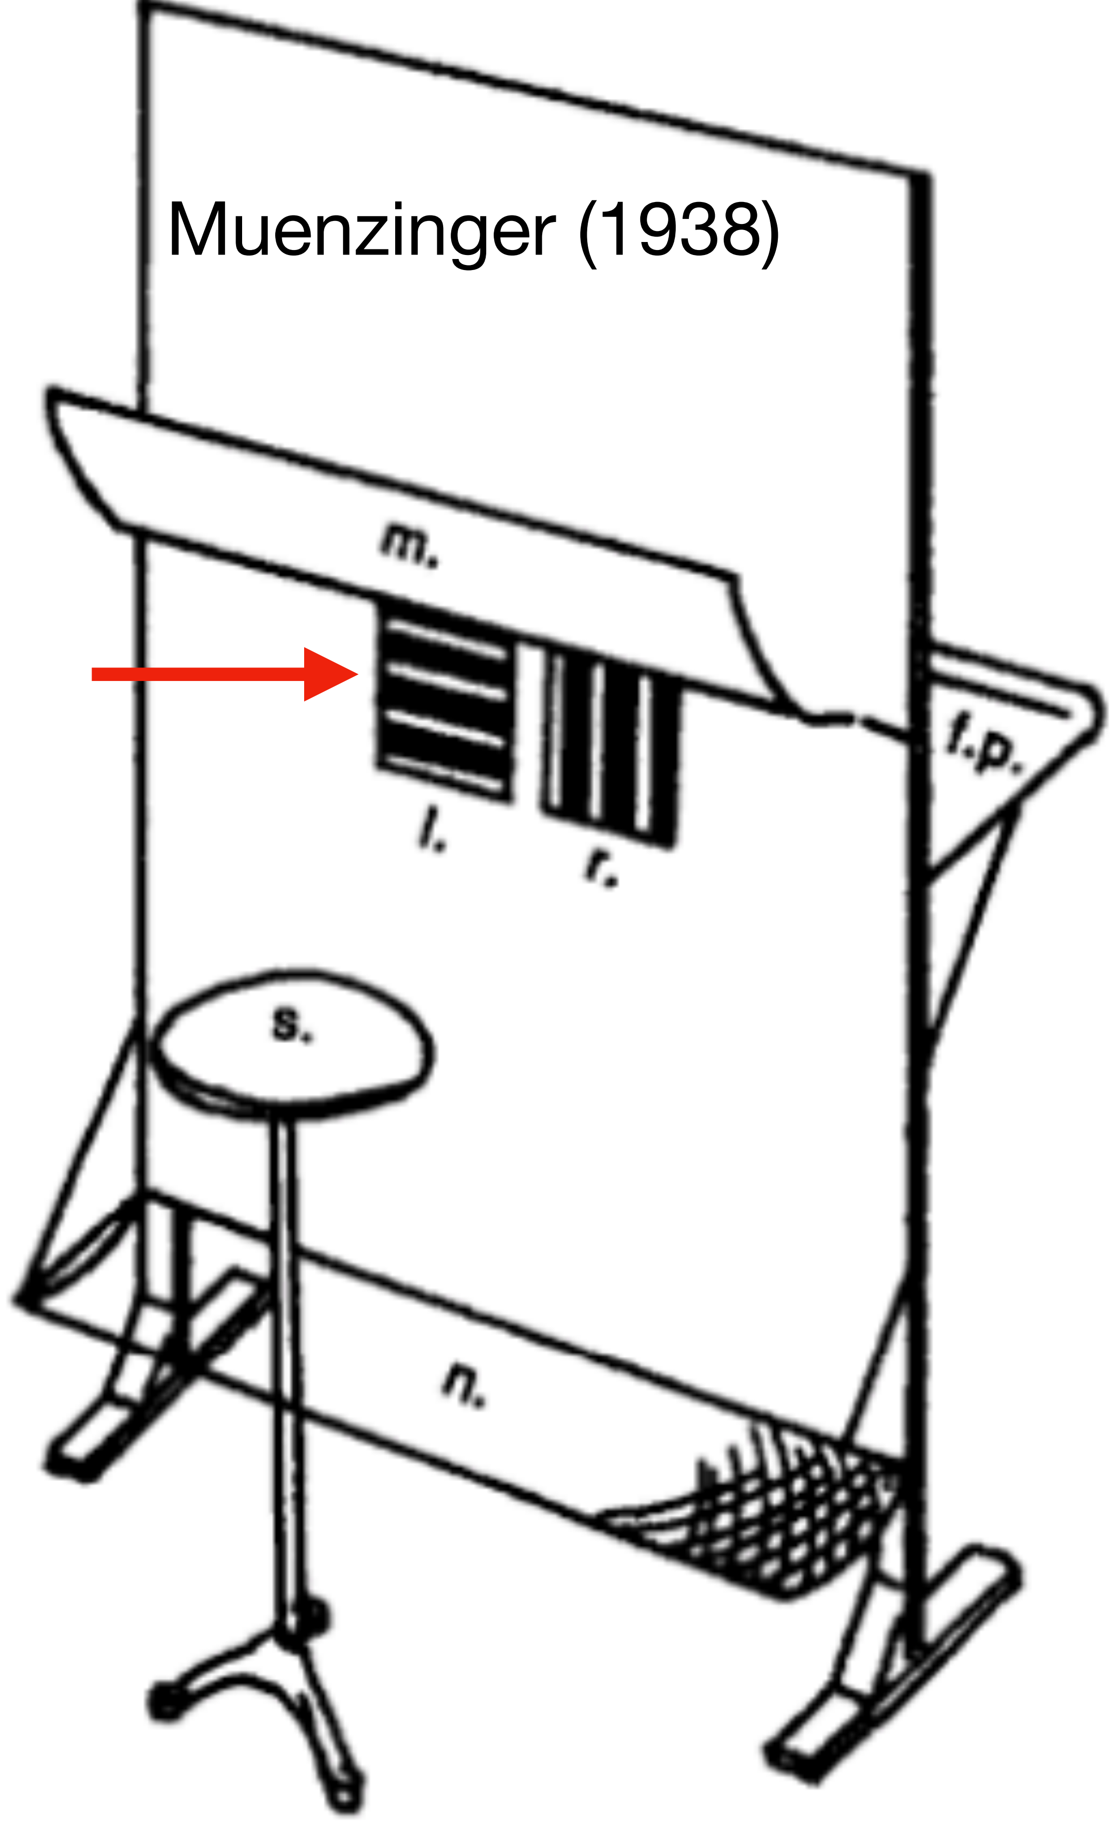

Vicarious Trial and Error (VTE)

- Animal put on jumping stand, facing two doors (l vs. r) with different visual properties (e.g., horizontal vs. vertical stripes)
  - One door is correct, the other incorrect
  - location is randomly swapped but visual features are predictive
  - If the animal jumps towards the correct door, it opens and reveals food on a platform behind… and if incorrect ….
- Tolman (1939) added landing platforms infront of the doors
  - When the choice was easy (black vs. white stimuli), the animals learned quicker and did more VTEing than for hard problems
  - After learning had been established, VTEs went down
  - Better learners also did more VTEing (Geier, LEvin & Tolman, 1941)

Muenzinger (1938)

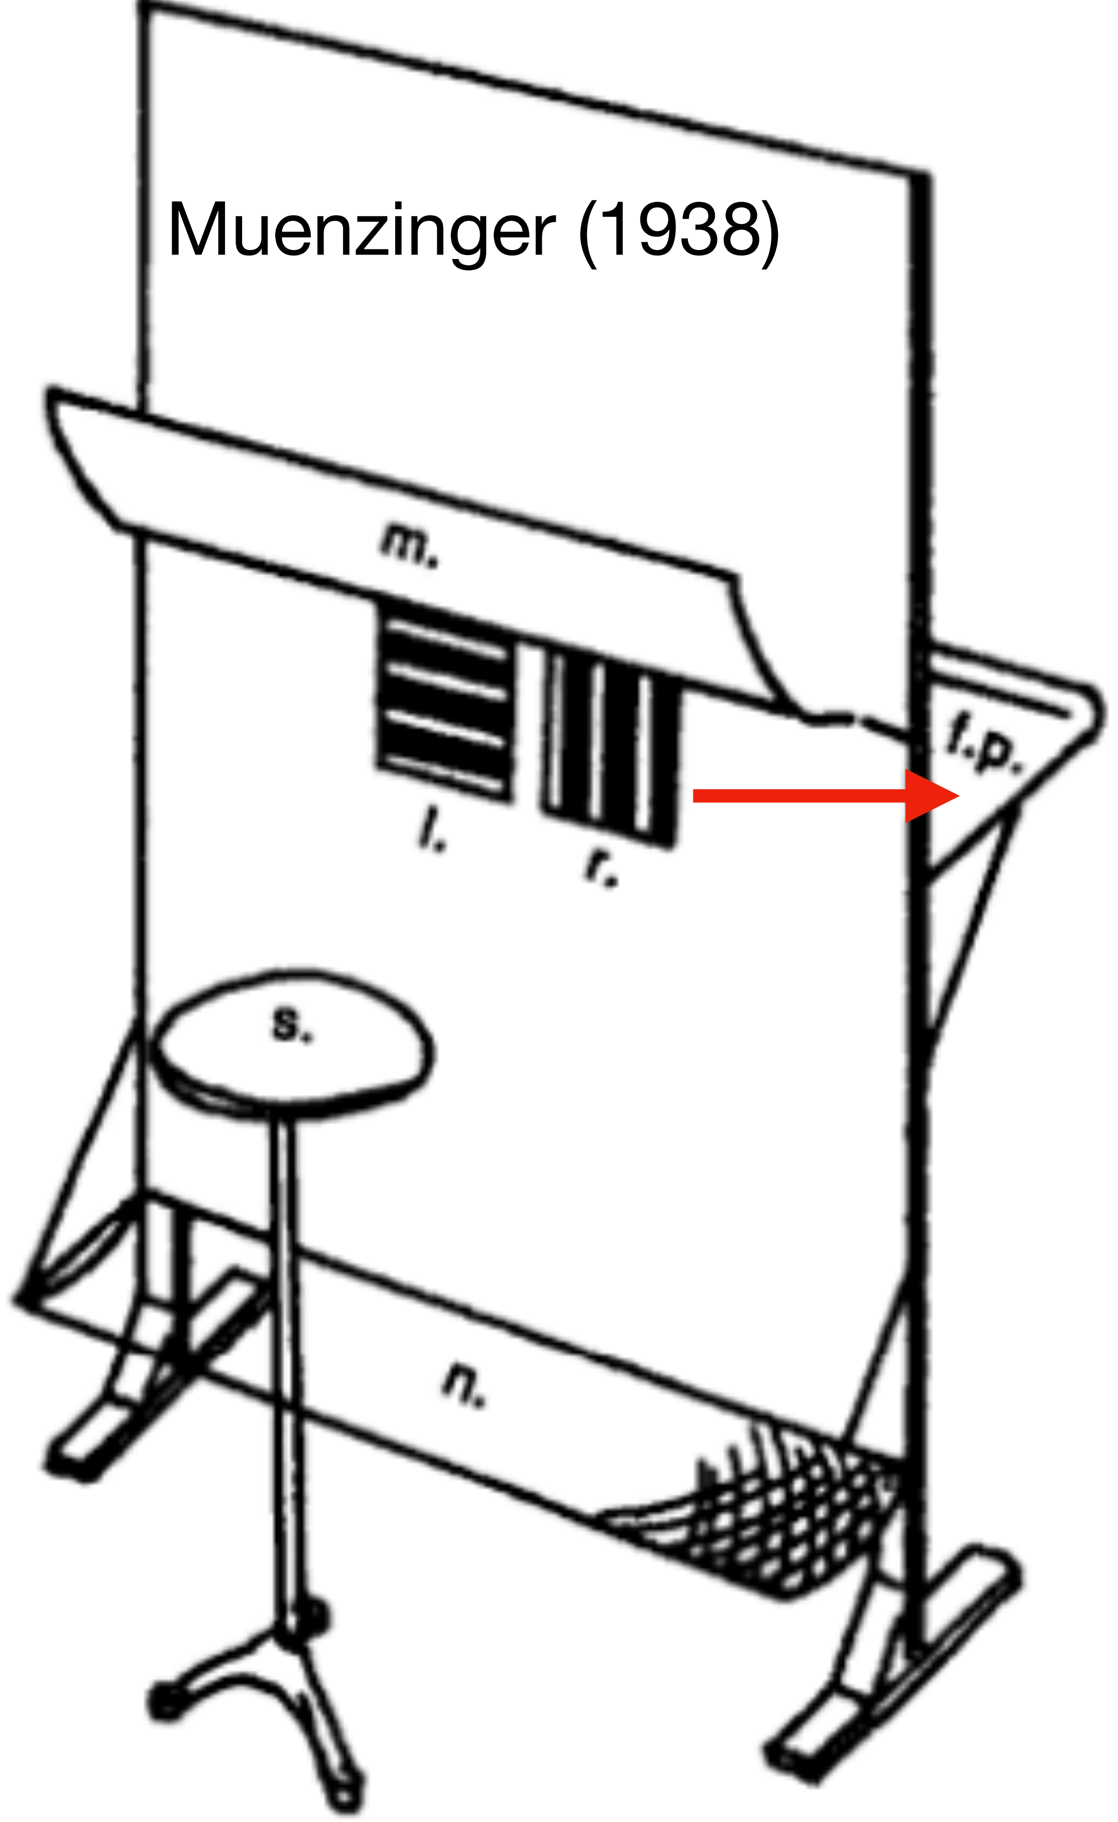

Vicarious Trial and Error (VTE)

- Animal put on jumping stand, facing two doors (l vs. r) with different visual properties (e.g., horizontal vs. vertical stripes)
  - One door is correct, the other incorrect
  - location is randomly swapped but visual features are predictive
  - If the animal jumps towards the correct door, it opens and reveals food on a platform behind… and if incorrect ….
- Tolman (1939) added landing platforms infront of the doors
  - When the choice was easy (black vs. white stimuli), the animals learned quicker and did more VTEing than for hard problems
  - After learning had been established, VTEs went down
  - Better learners also did more VTEing (Geier, LEvin & Tolman, 1941)

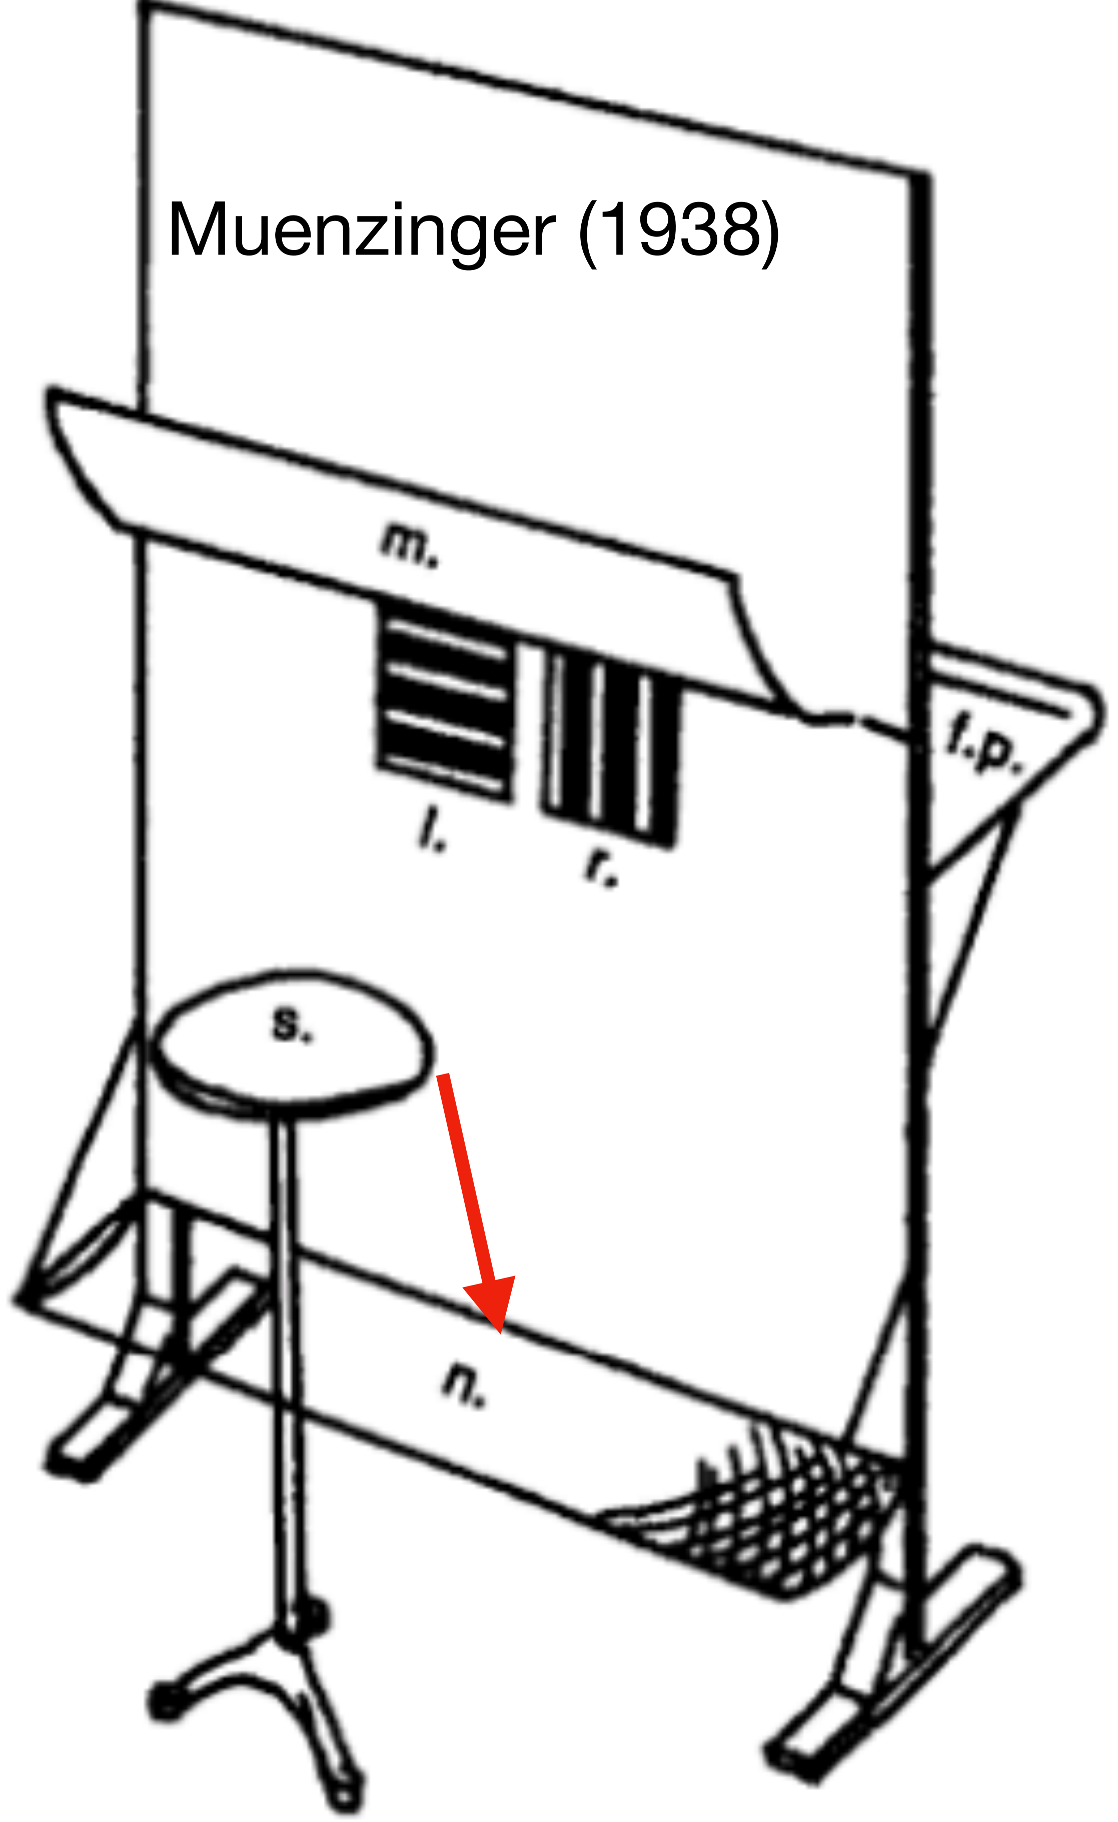

Vicarious Trial and Error (VTE)

- Animal put on jumping stand, facing two doors (l vs. r) with different visual properties (e.g., horizontal vs. vertical stripes)
    - One door is correct, the other incorrect
    - location is randomly swapped but visual features are predictive
    - If the animal jumps towards the correct door, it opens and reveals food on a platform behind… and if incorrect ….
- Tolman (1939) added landing platforms infront of the doors
    - When the choice was easy (black vs. white stimuli), the animals learned quicker and did more VTEing than for hard problems
    - After learning had been established, VTEs went down
    - Better learners also did more VTEing (Geier, LEvin & Tolman, 1941)

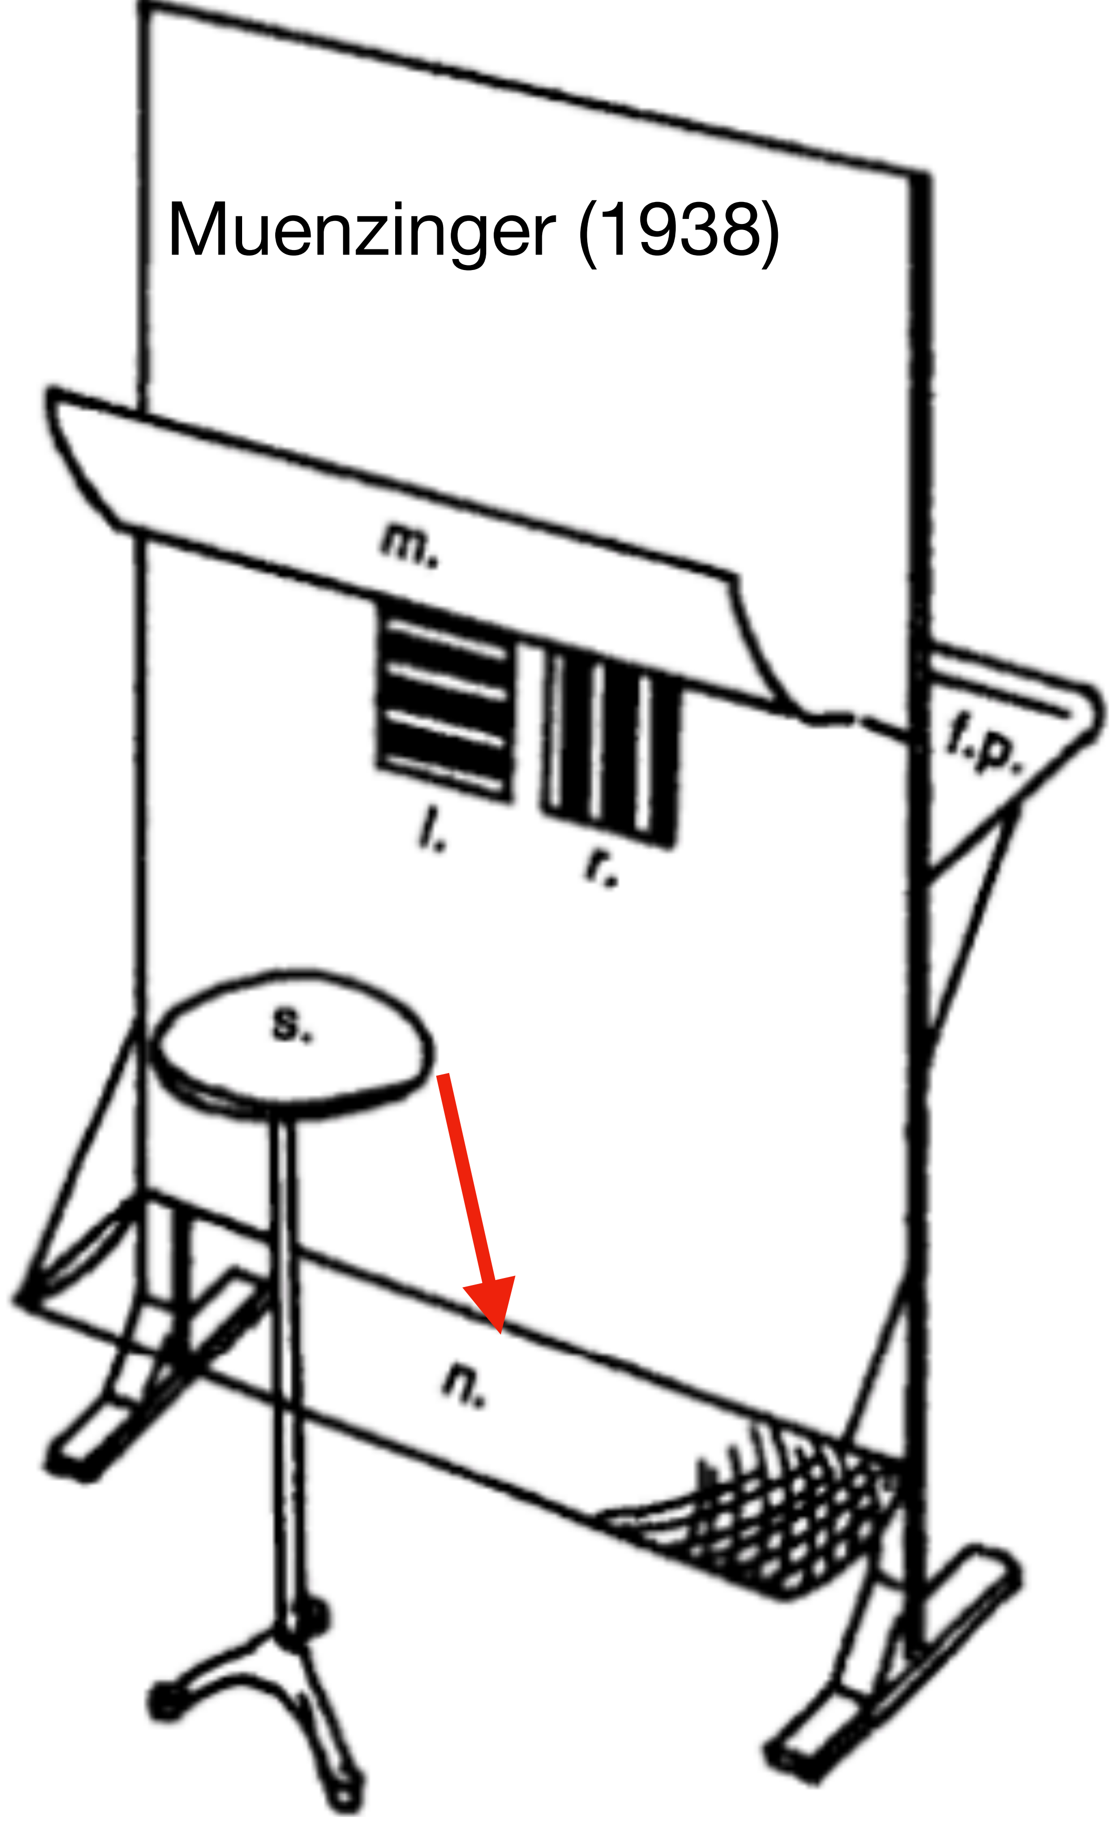

Muenzinger (1938)

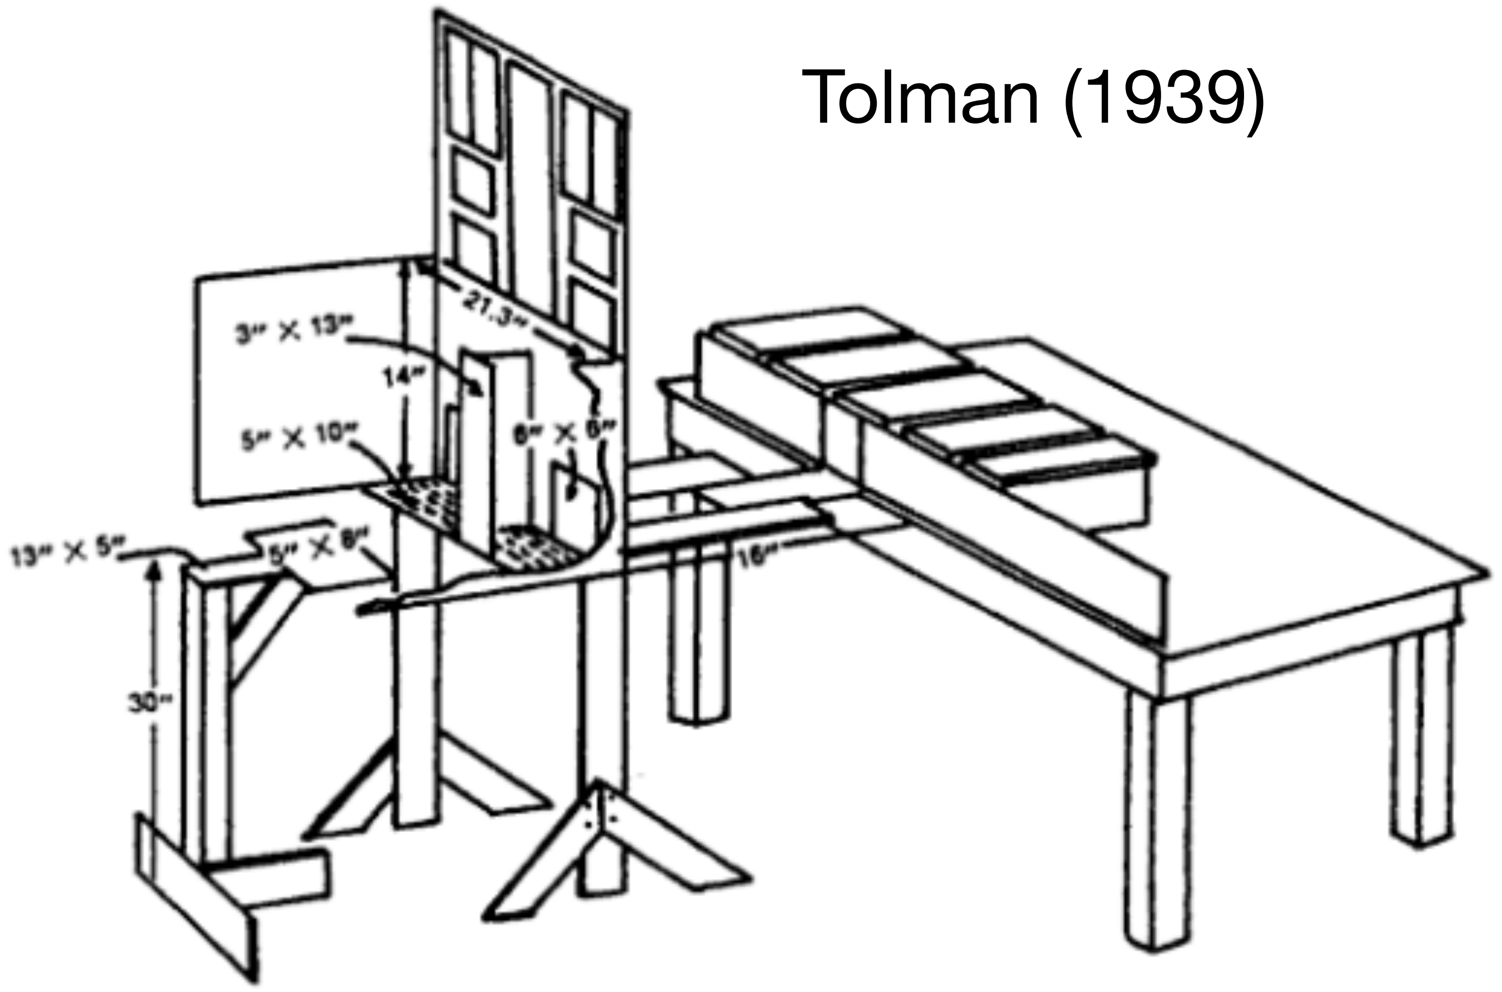
Tolman (1939)

Vicarious Trial and Error (VTE)

- Animal put on jumping stand, facing two doors (l vs. r) with different visual properties (e.g., horizontal vs. vertical stripes)
    - One door is correct, the other incorrect
    - location is randomly swapped but visual features are predictive
    - If the animal jumps towards the correct door, it opens and reveals food on a platform behind… and if incorrect ….
- Tolman (1939) added landing platforms infront of the doors
    - When the choice was easy (black vs. white stimuli), the animals learned quicker and did more VTEing than for hard problems
    - After learning had been established, VTEs went down
    - Better learners also did more VTEing (Geier, LEvin & Tolman, 1941)

Muenzinger (1938)

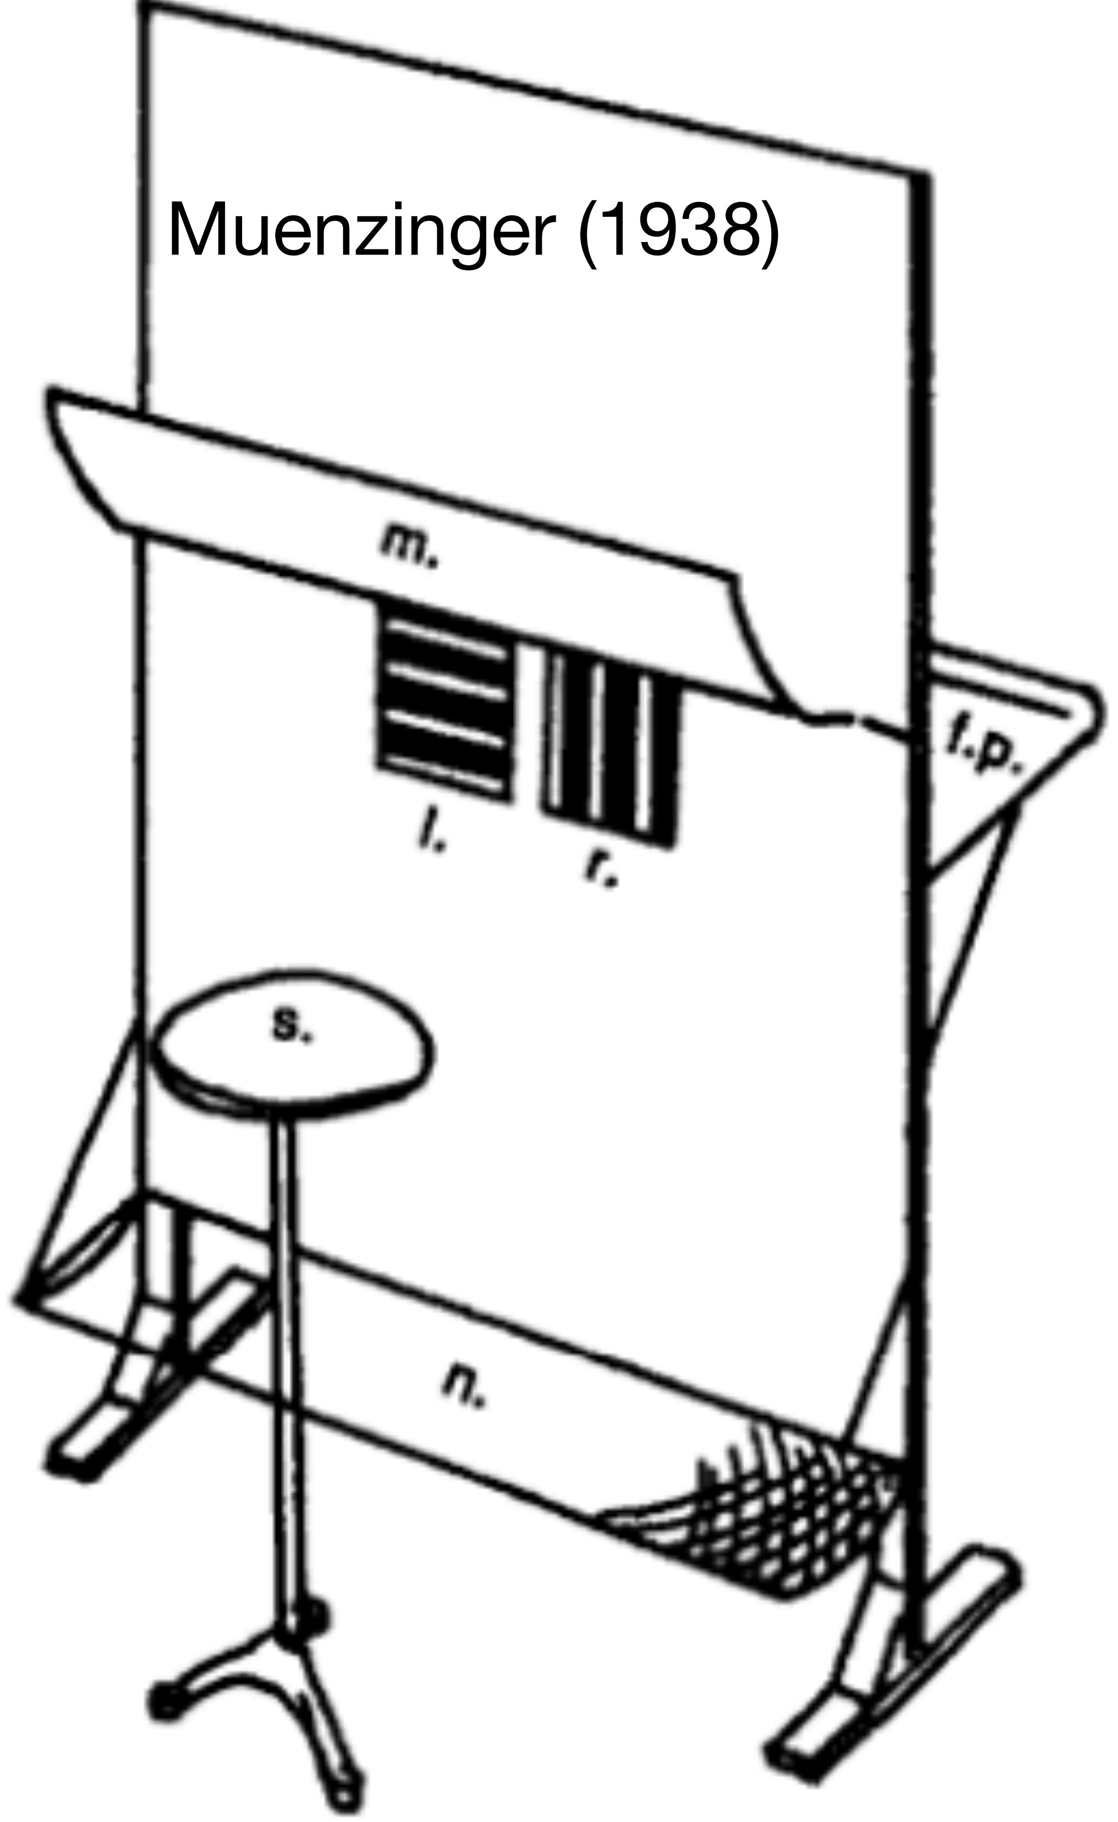

Tolman (1939)

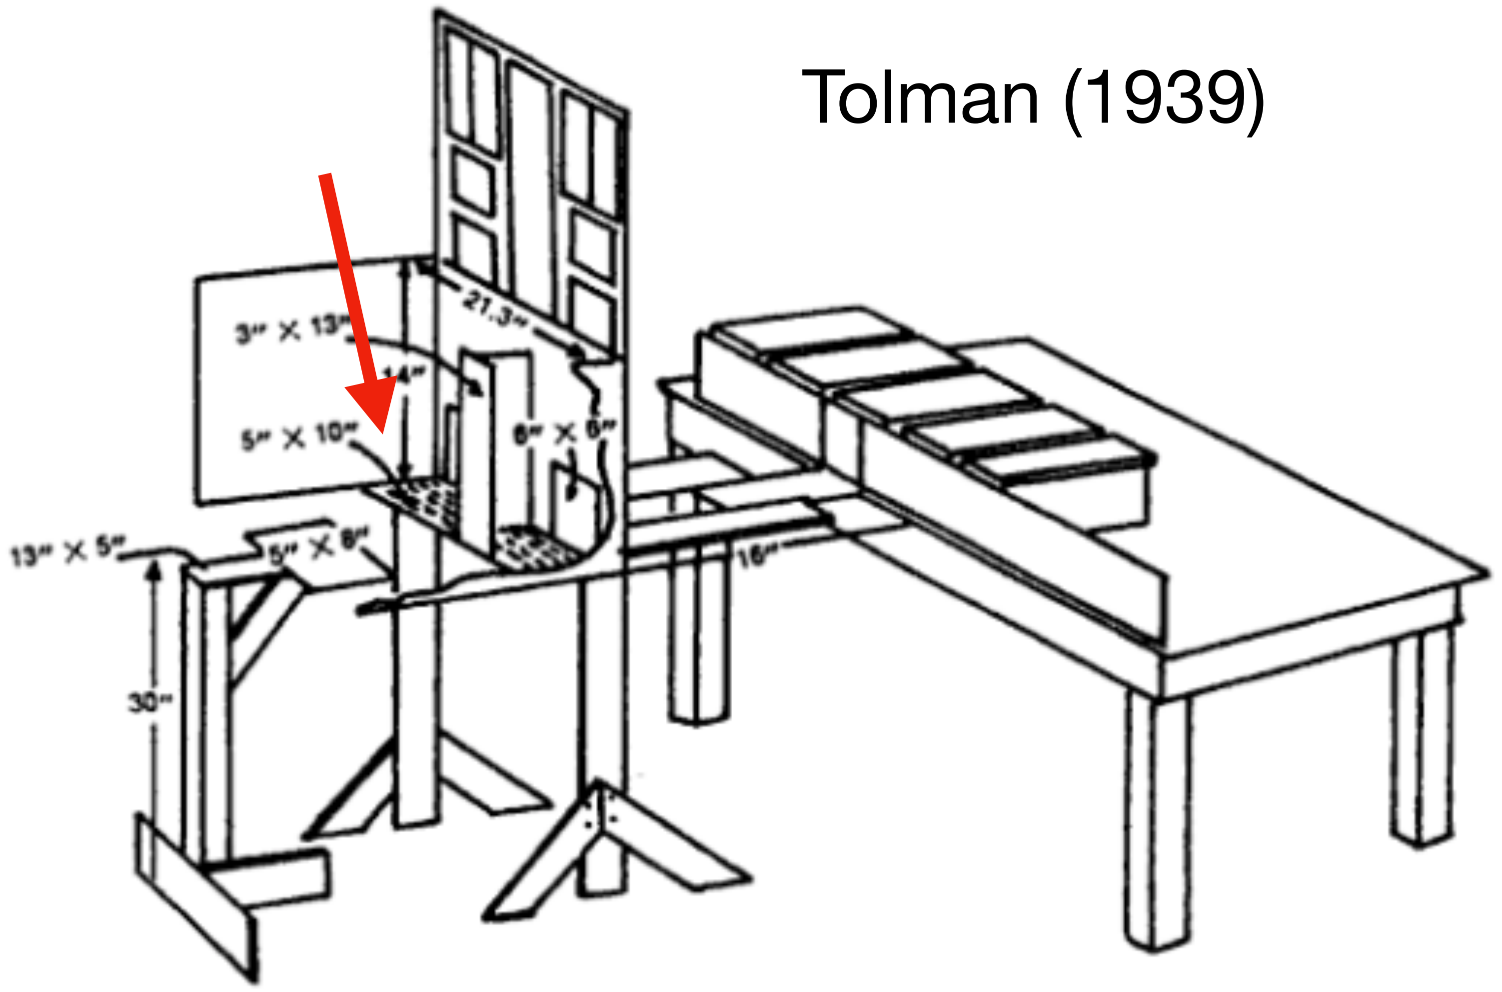

Vicarious Trial and Error (VTE)

- Animal put on jumping stand, facing two doors (l vs. r) with different visual properties (e.g., horizontal vs. vertical stripes)
  - One door is correct, the other incorrect
  - location is randomly swapped but visual features are predictive
  - If the animal jumps towards the correct door, it opens and reveals food on a platform behind… and if incorrect ….
- Tolman (1939) added landing platforms infront of the doors
  - When the choice was easy (black vs. white stimuli), the animals learned quicker and did more VTEing than for hard problems
  - After learning had been established, VTEs went down
  - Better learners also did more VTEing (Geier, LEvin & Tolman, 1941)

Vicarious trial and error (VTE): hesitating, looking-back-and-forth behavior observed in rats when confronted with a choice

Muenzinger (1938)

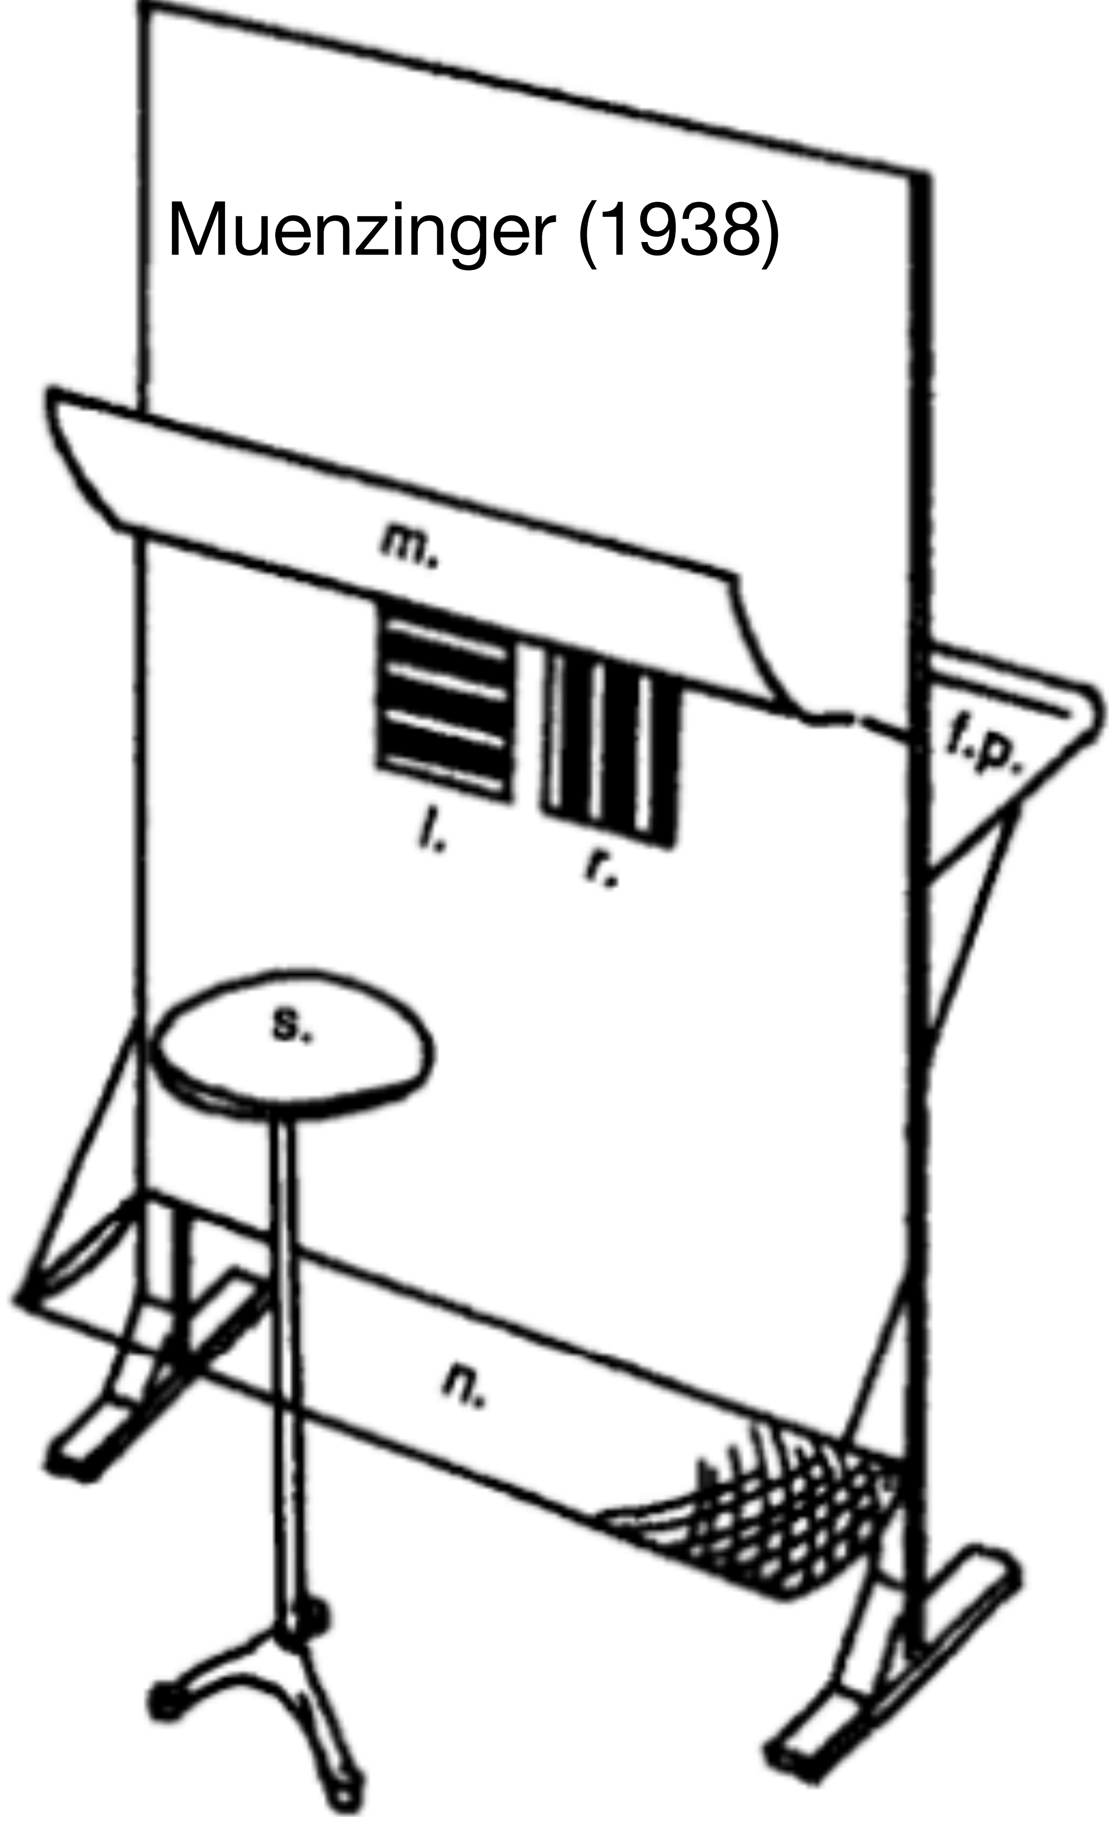

Tolman (1939)

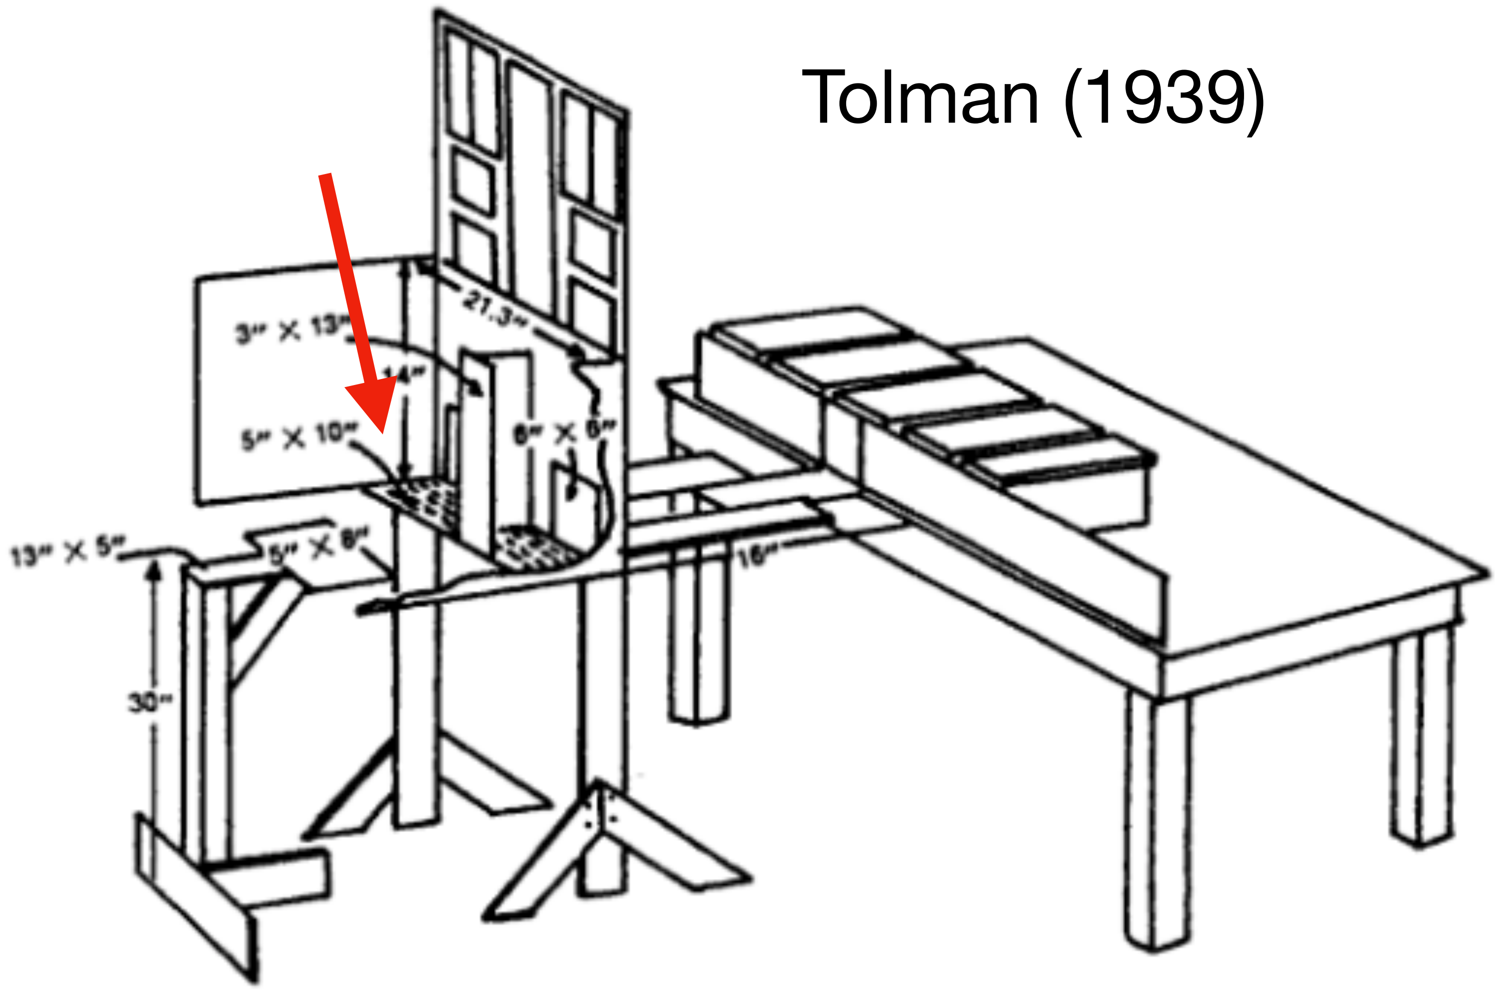


AVIOLI·P


AVIOLI·P

Vicarious Trial and Error (VTE)

- VTEs coincide with the start of learning, and fade away afterwards
- Not just passive association of stimuli, but active selecting and comparison of stimuli

Learning Curves

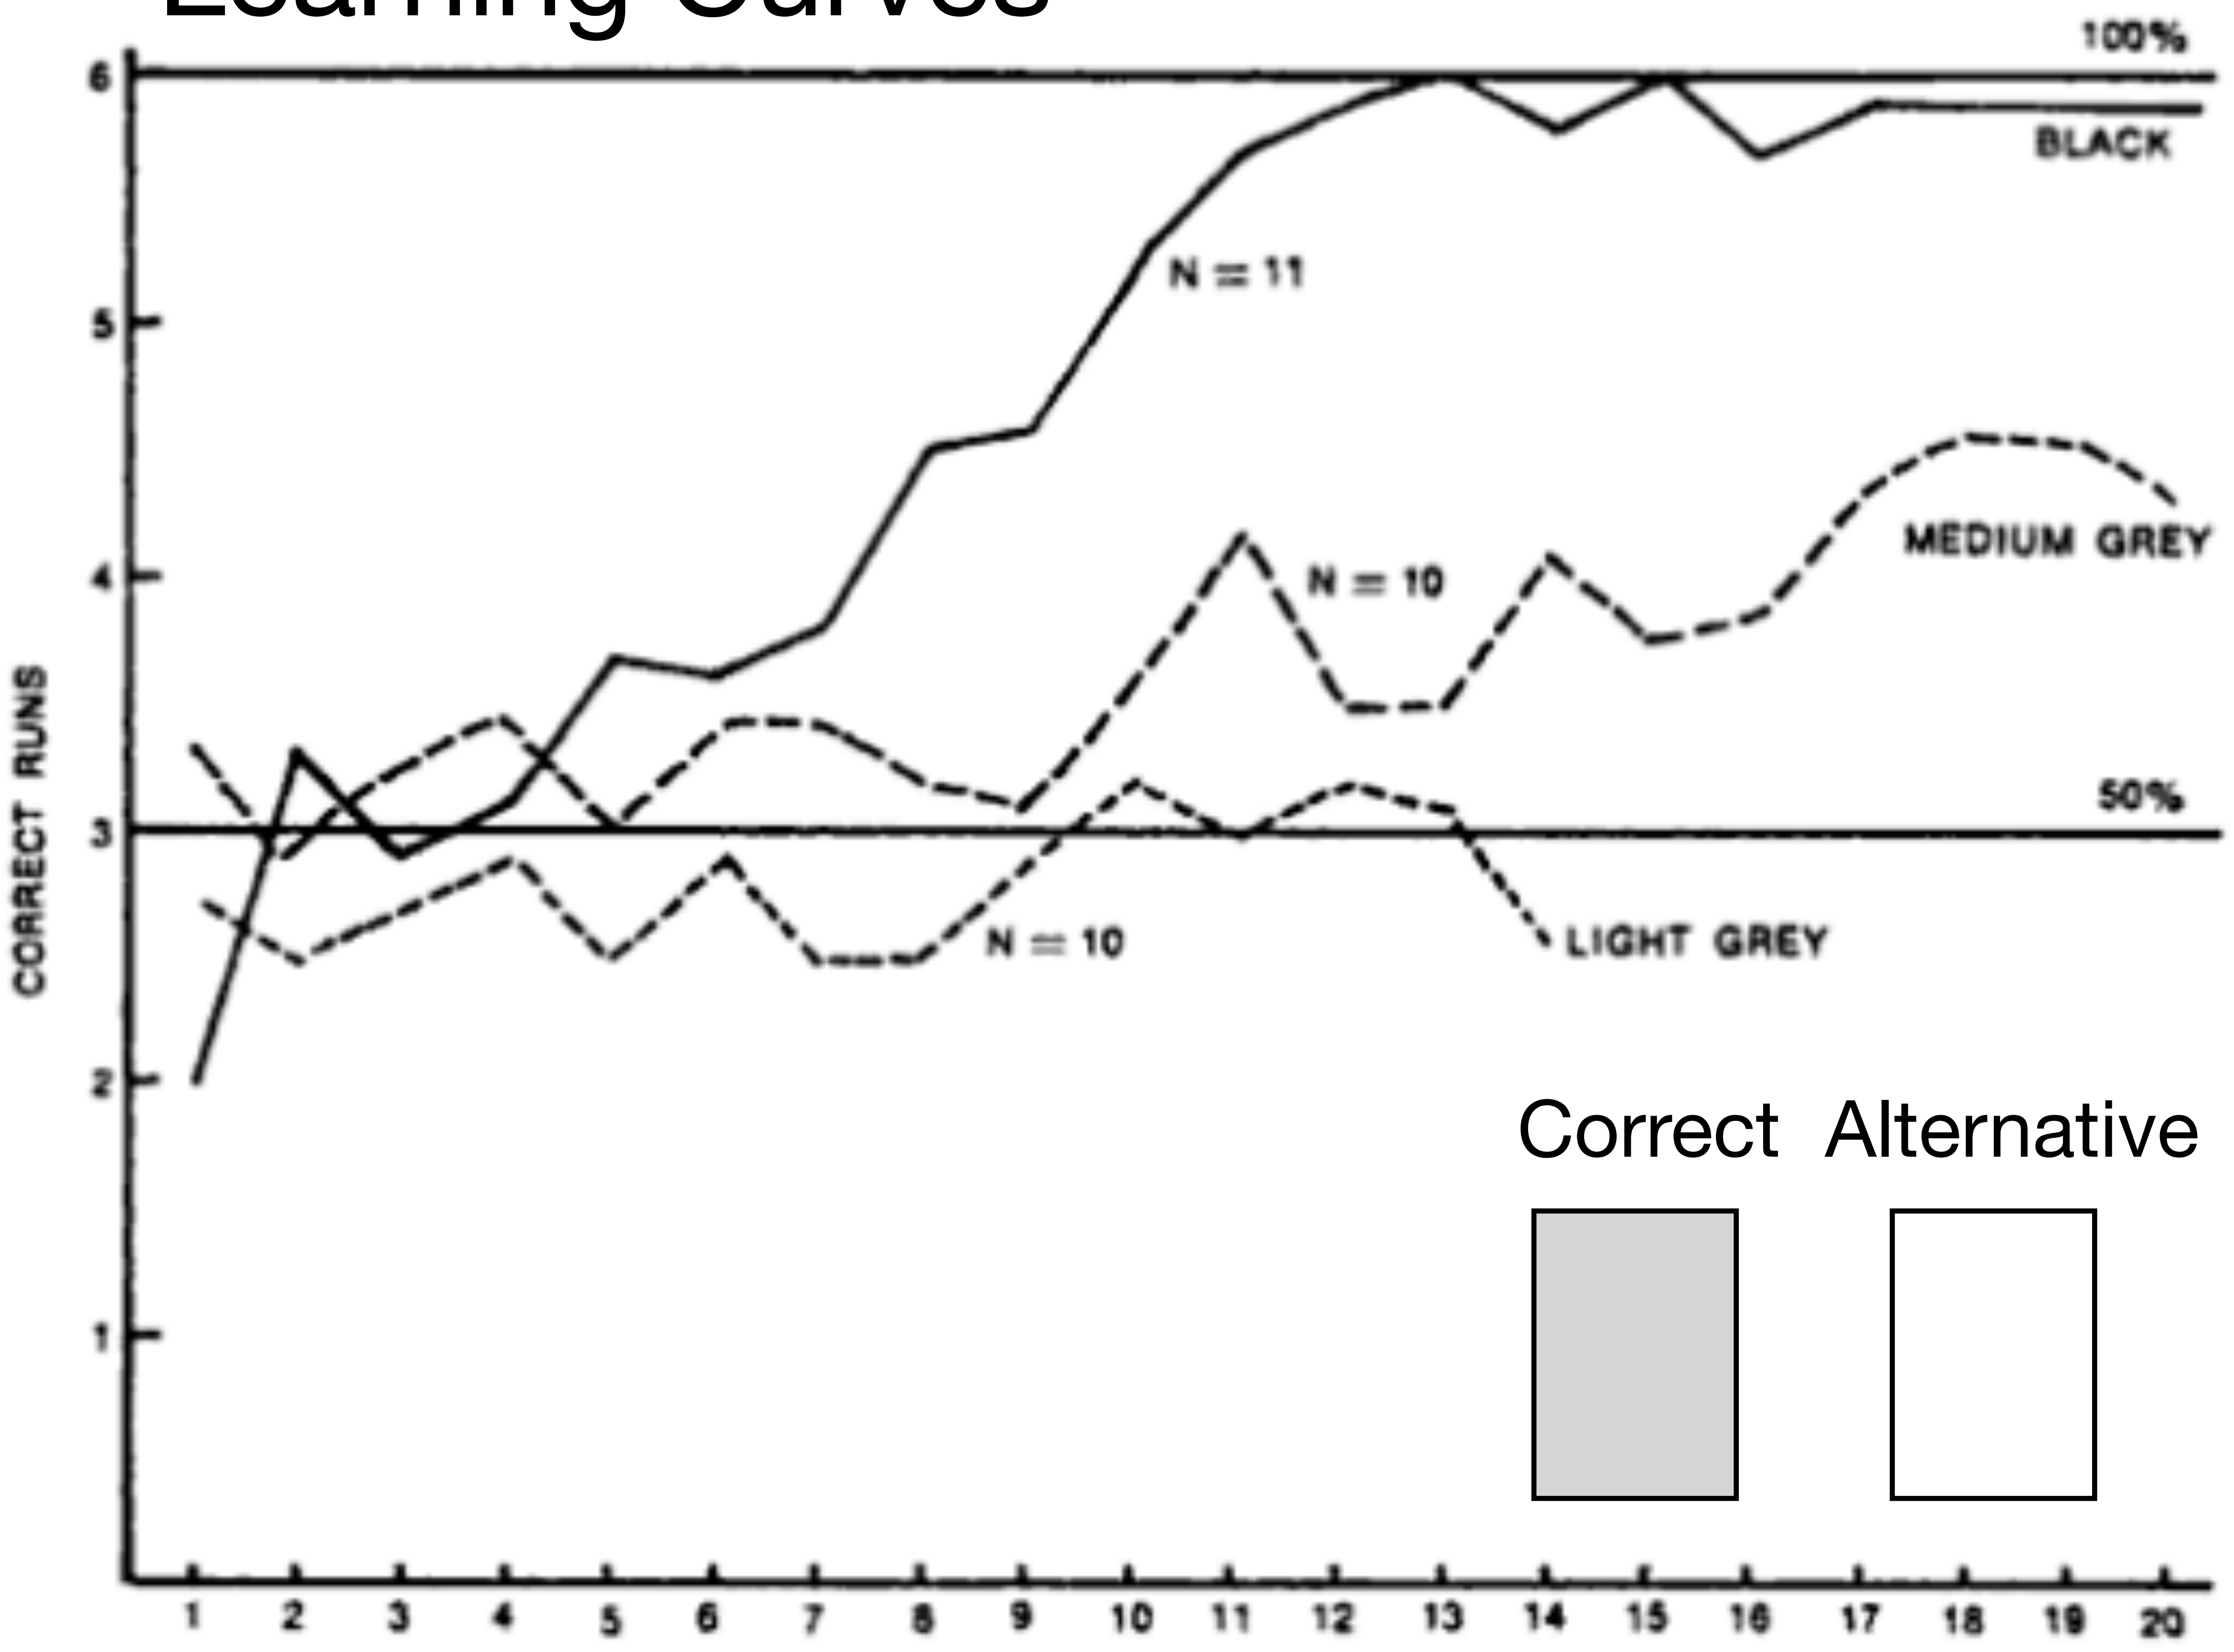

VTE rate

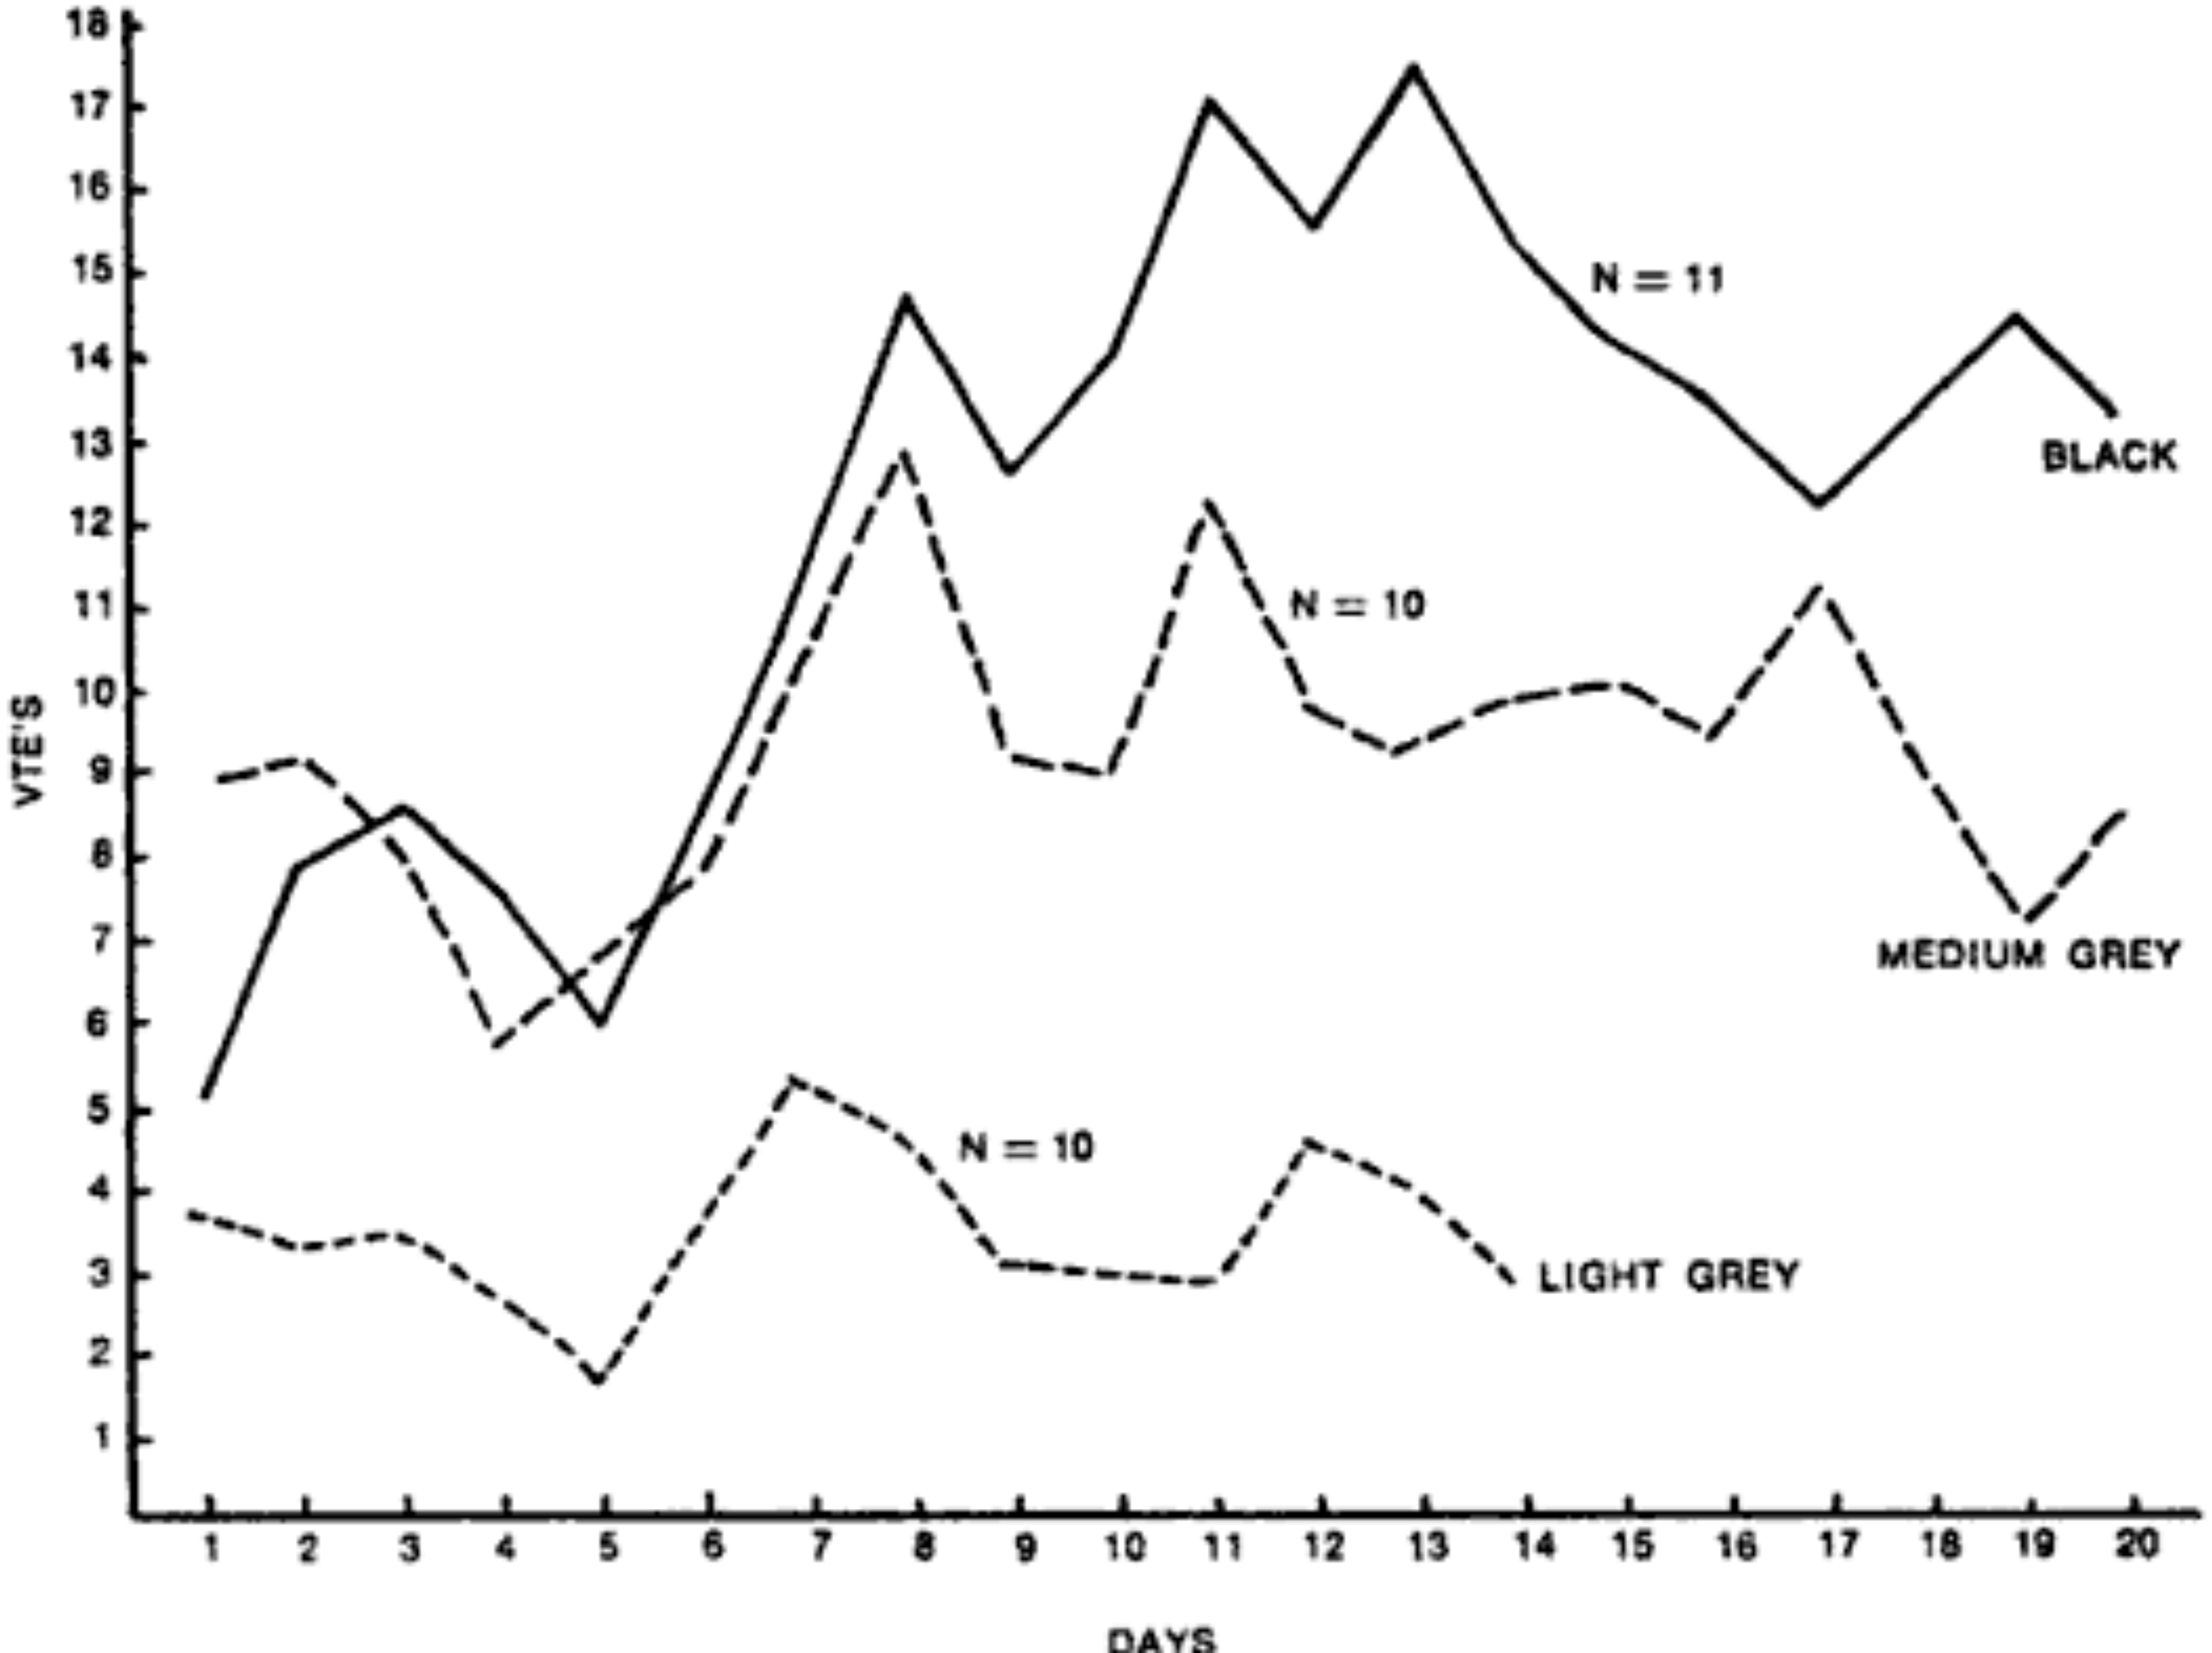

Vicarious Trial and Error (VTE)

- VTEs coincide with the start of learning, and fade away afterwards
- Not just passive association of stimuli, but active selecting and comparison of stimuli

Learning Curves

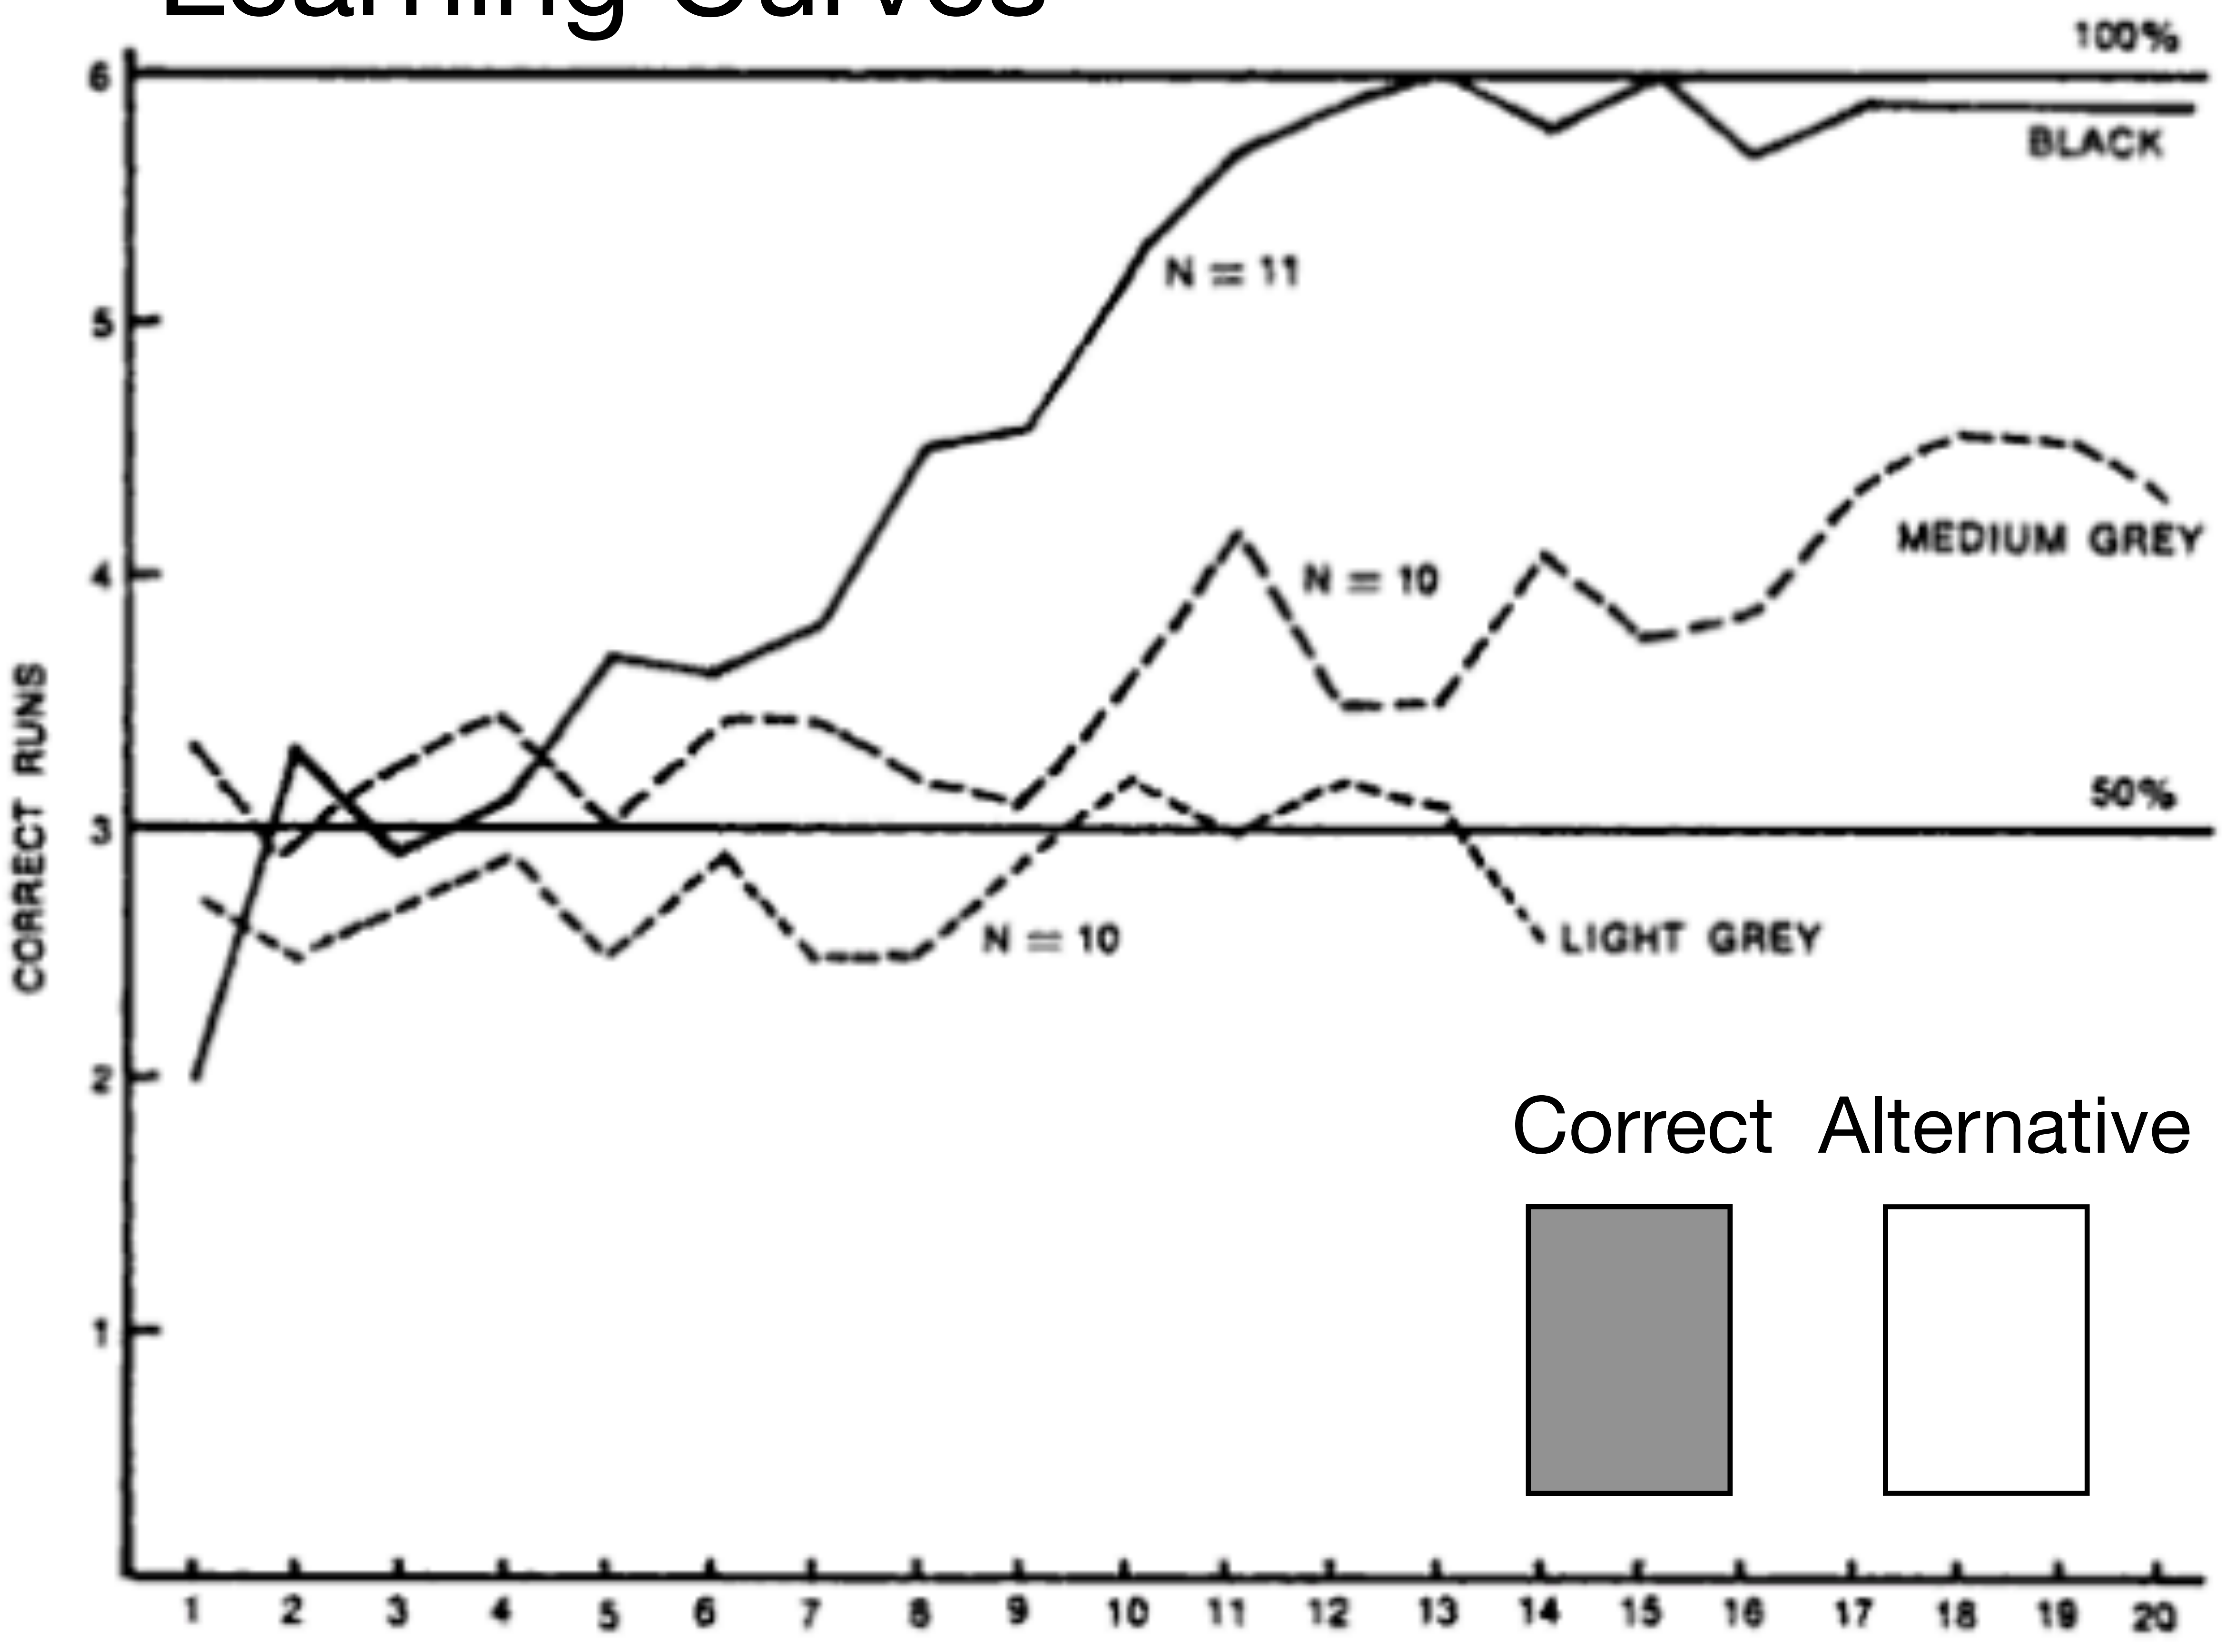

VTE rate

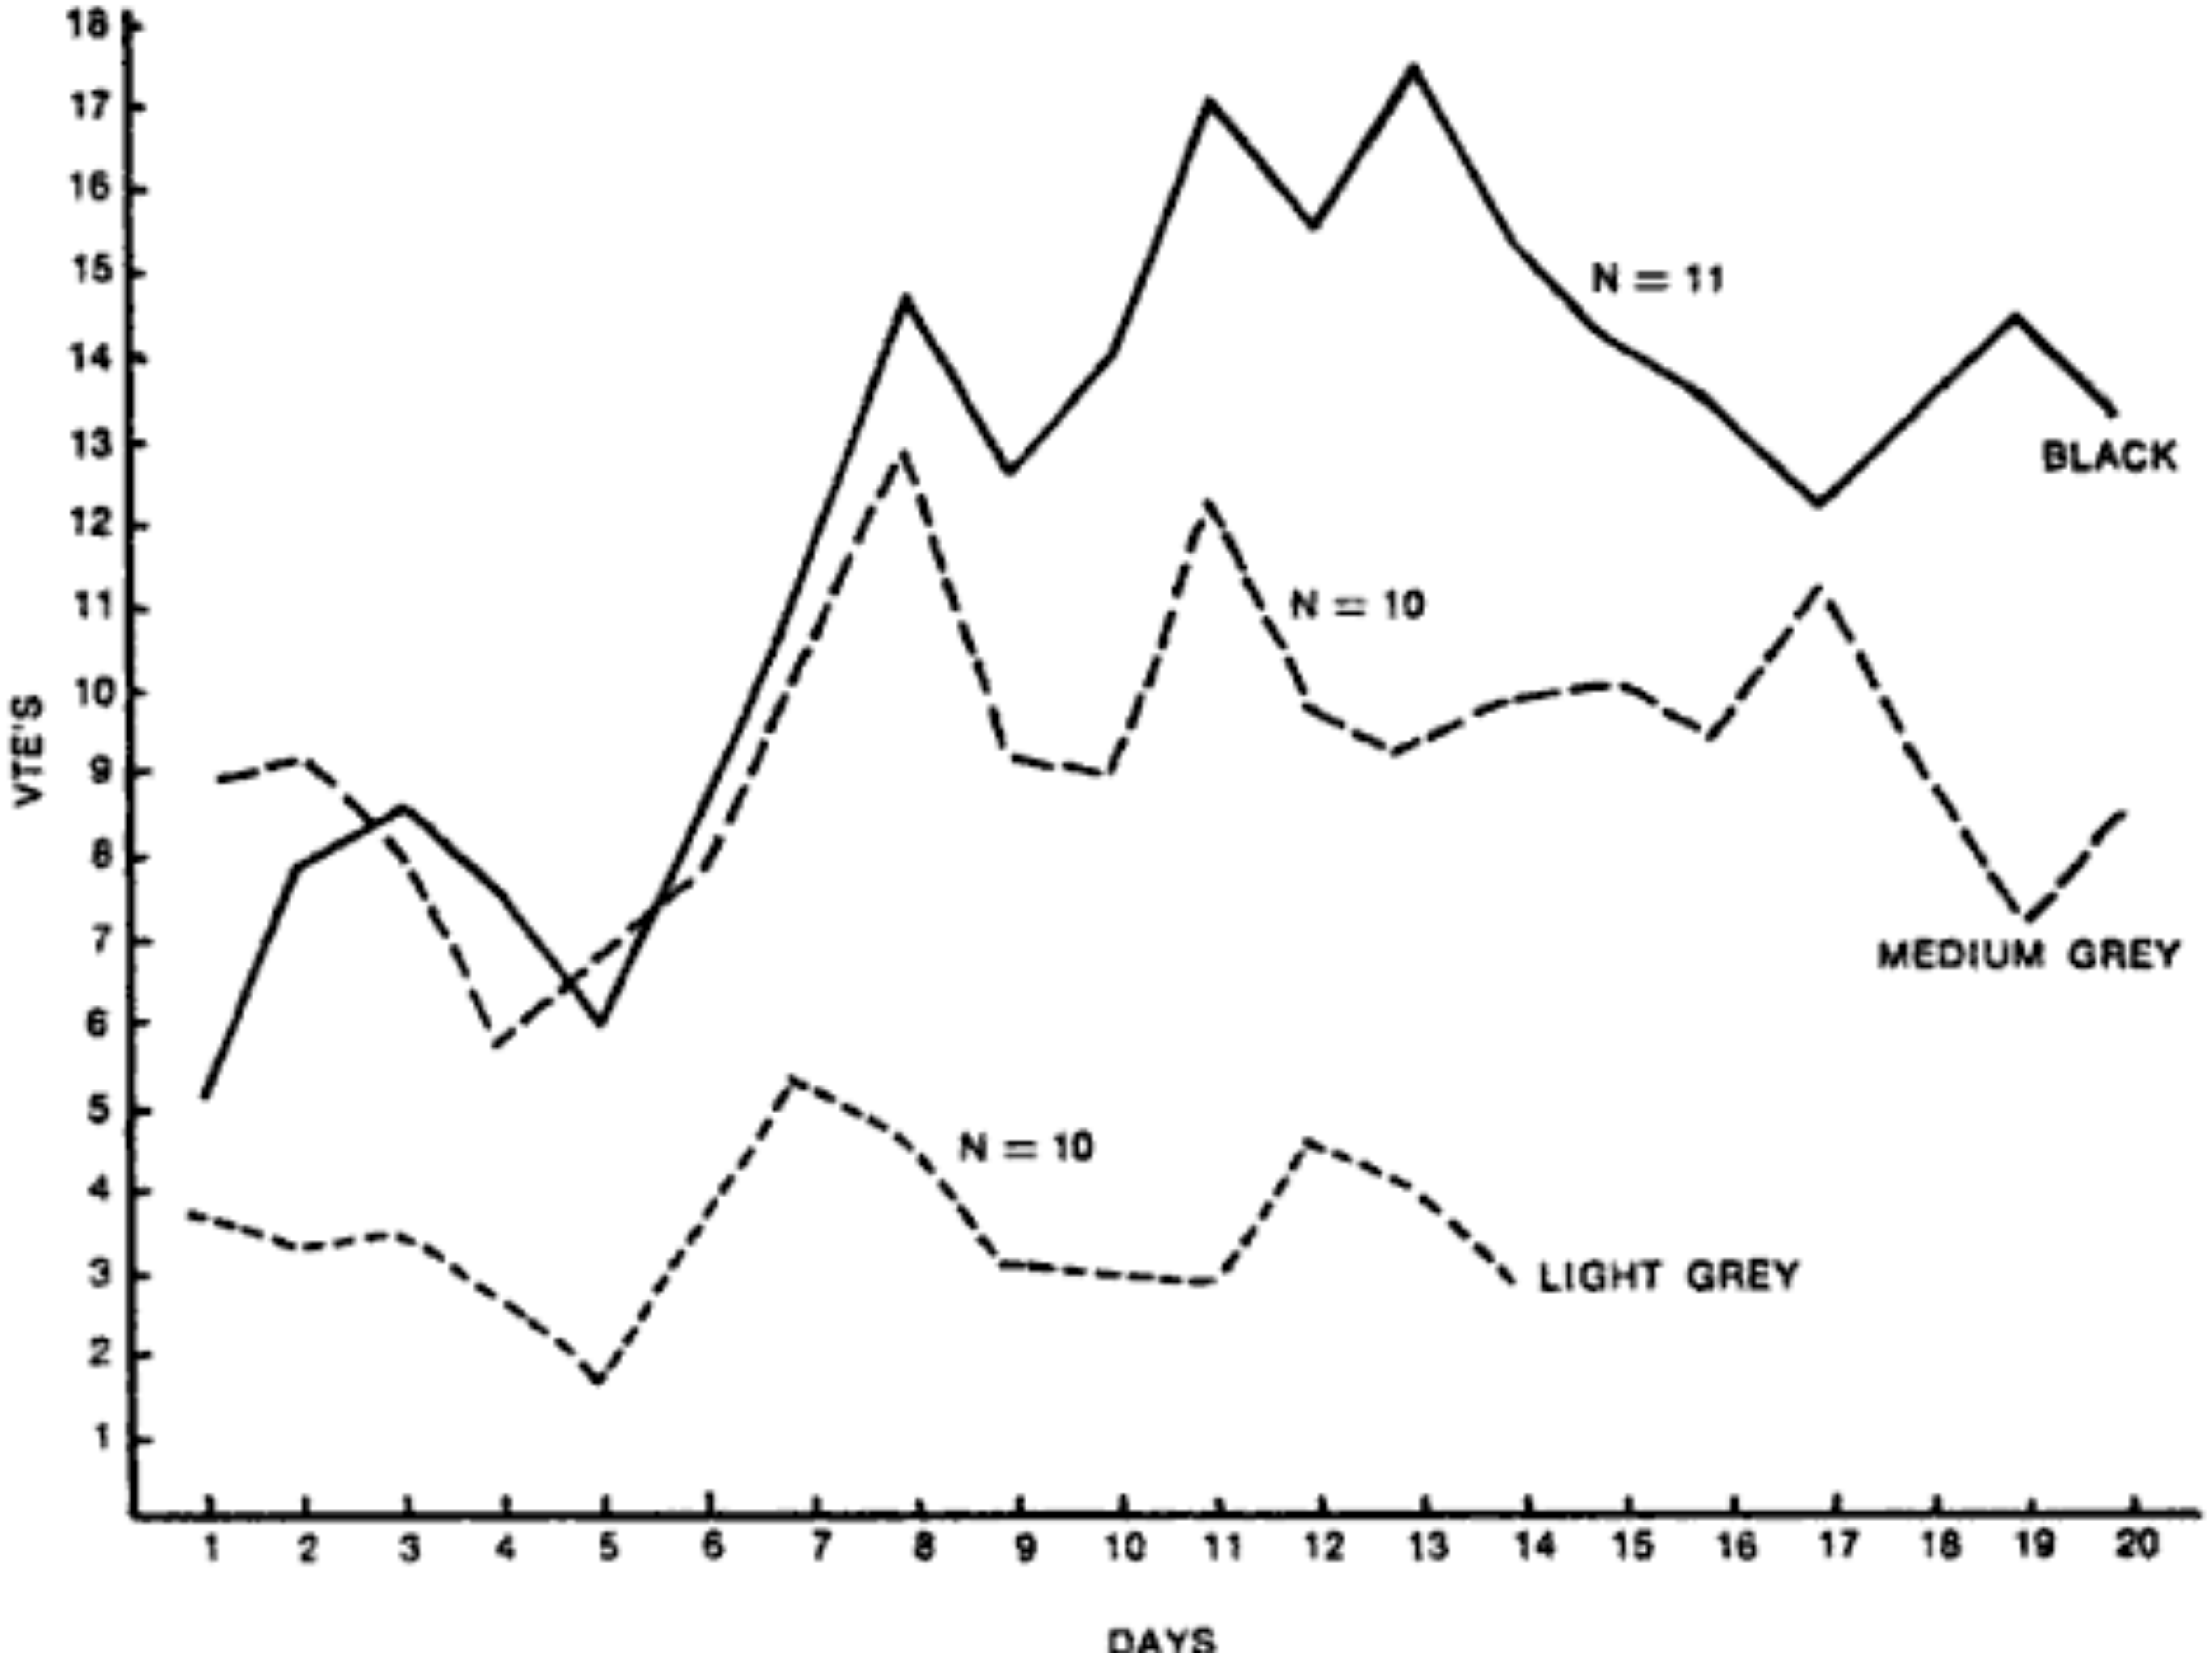

Vicarious Trial and Error (VTE)

- VTEs coincide with the start of learning, and fade away afterwards
- Not just passive association of stimuli, but active selecting and comparison of stimuli

Learning Curves

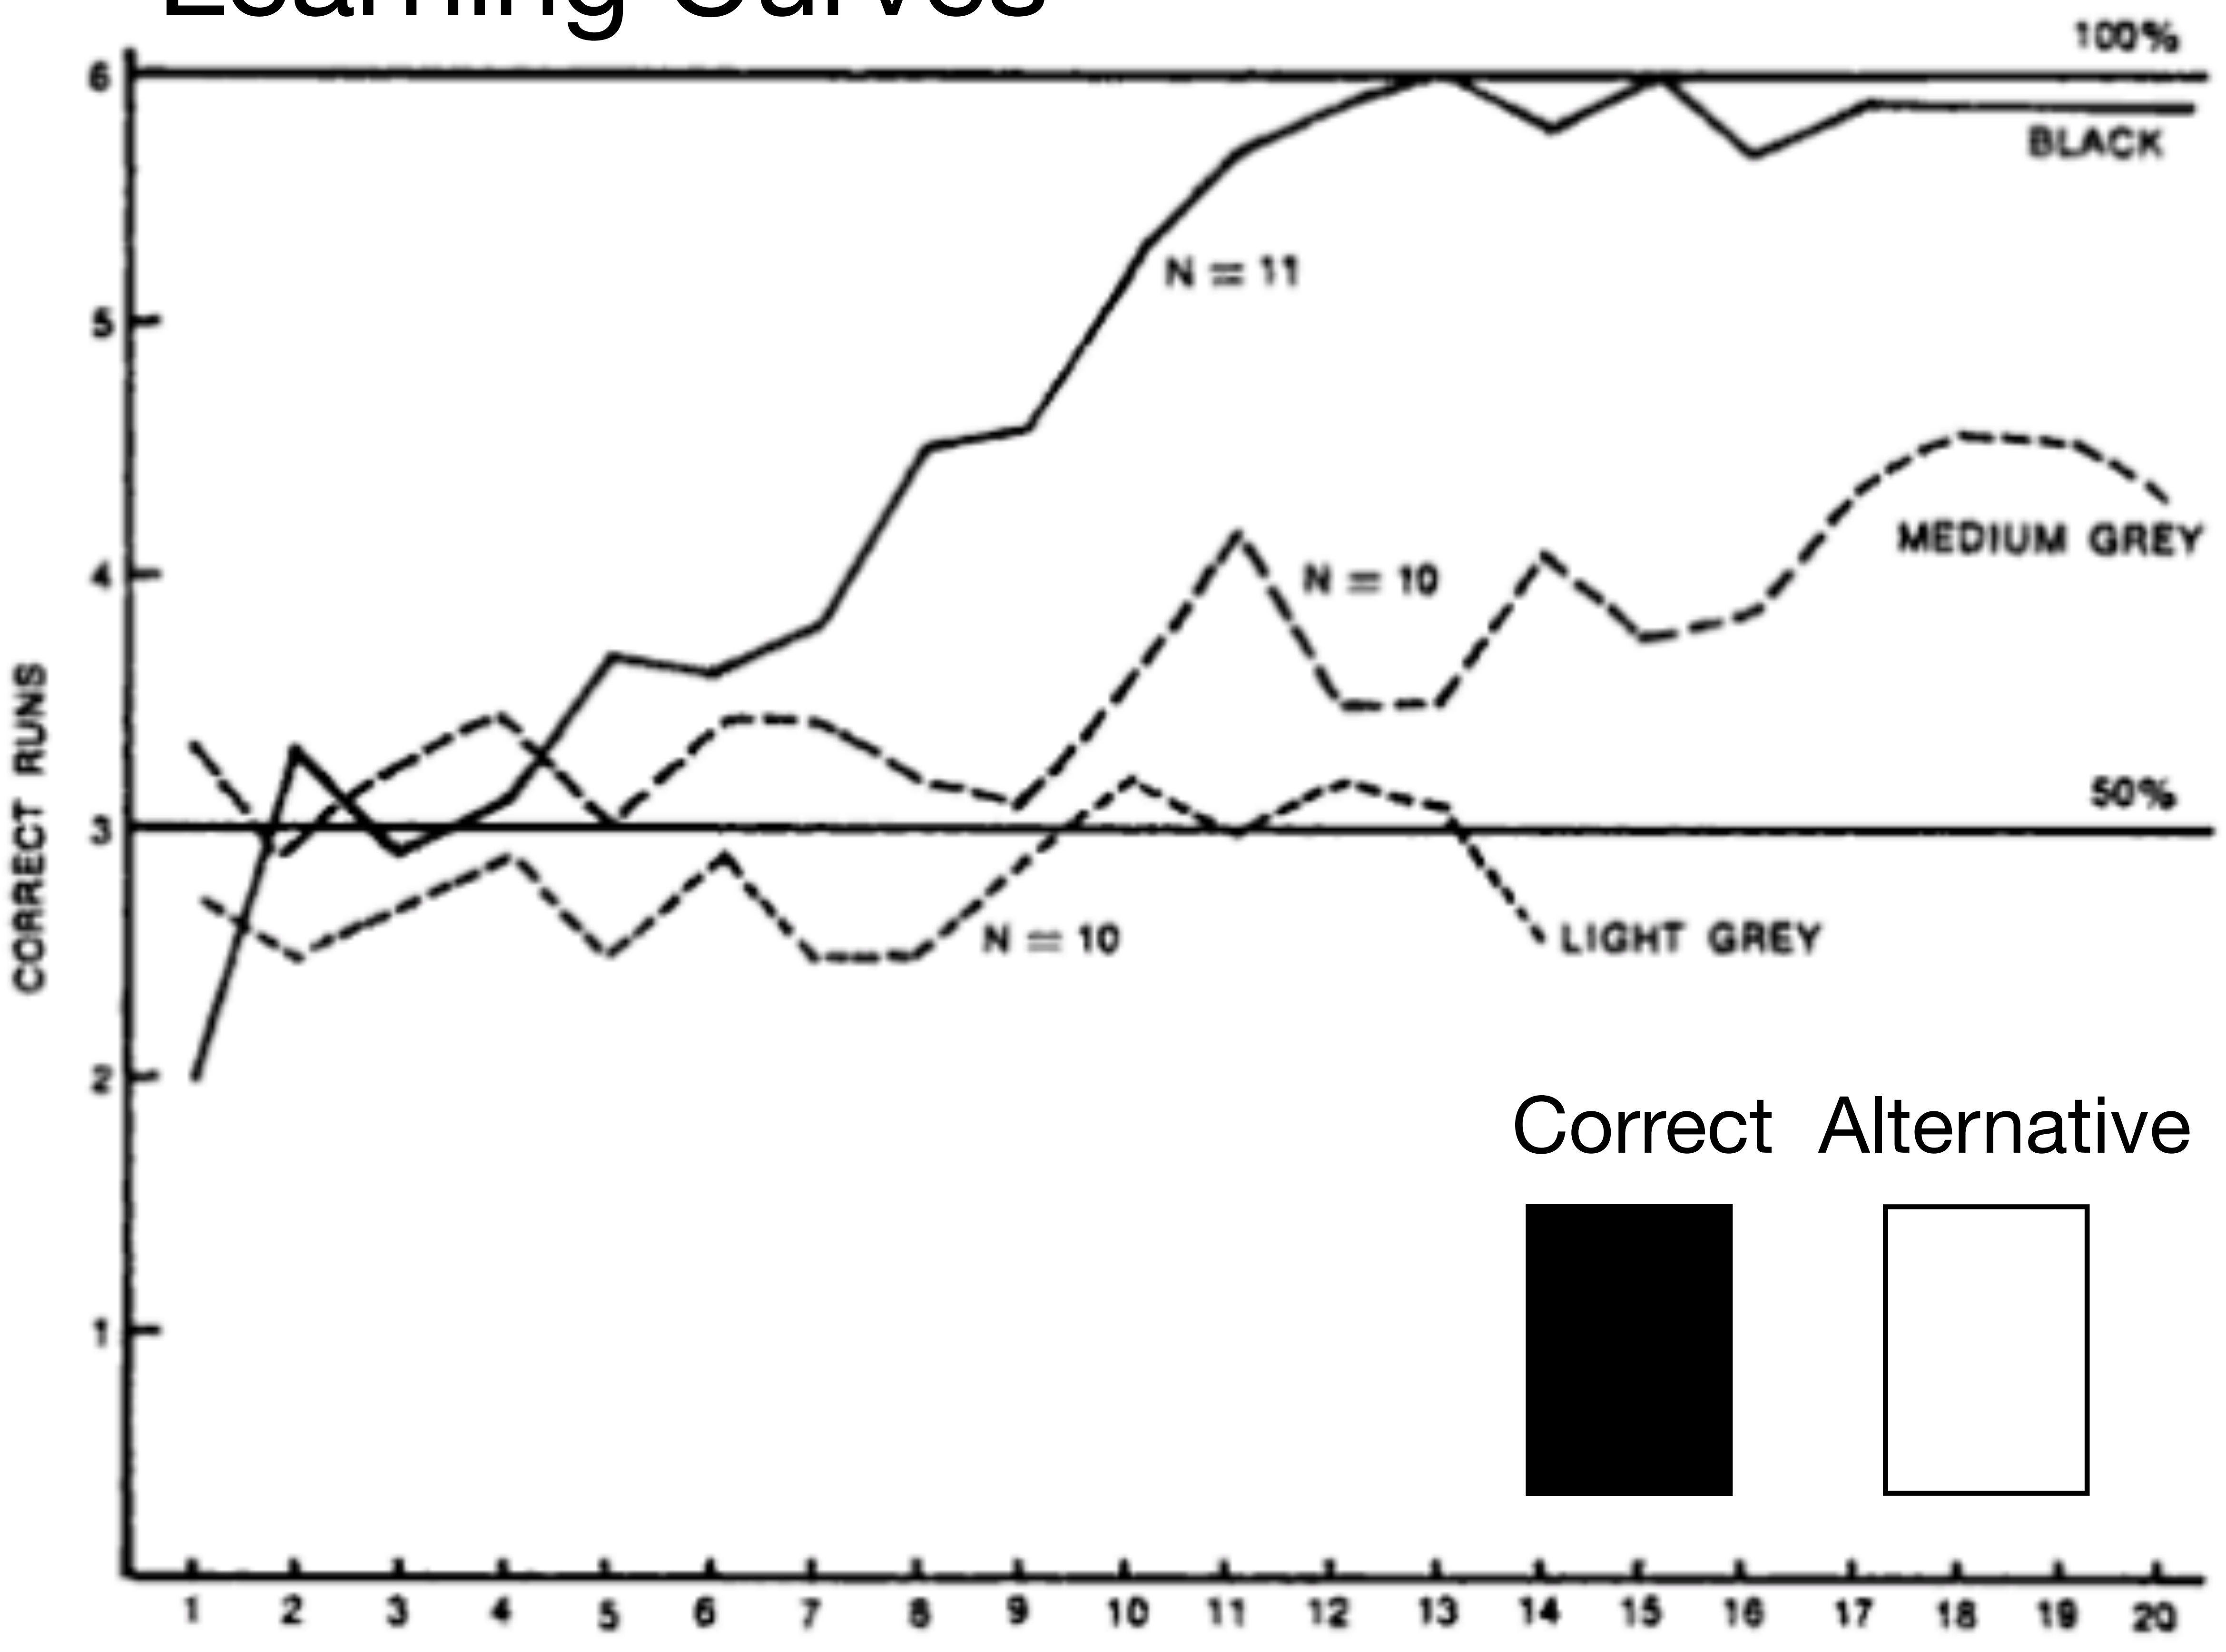

VTE rate

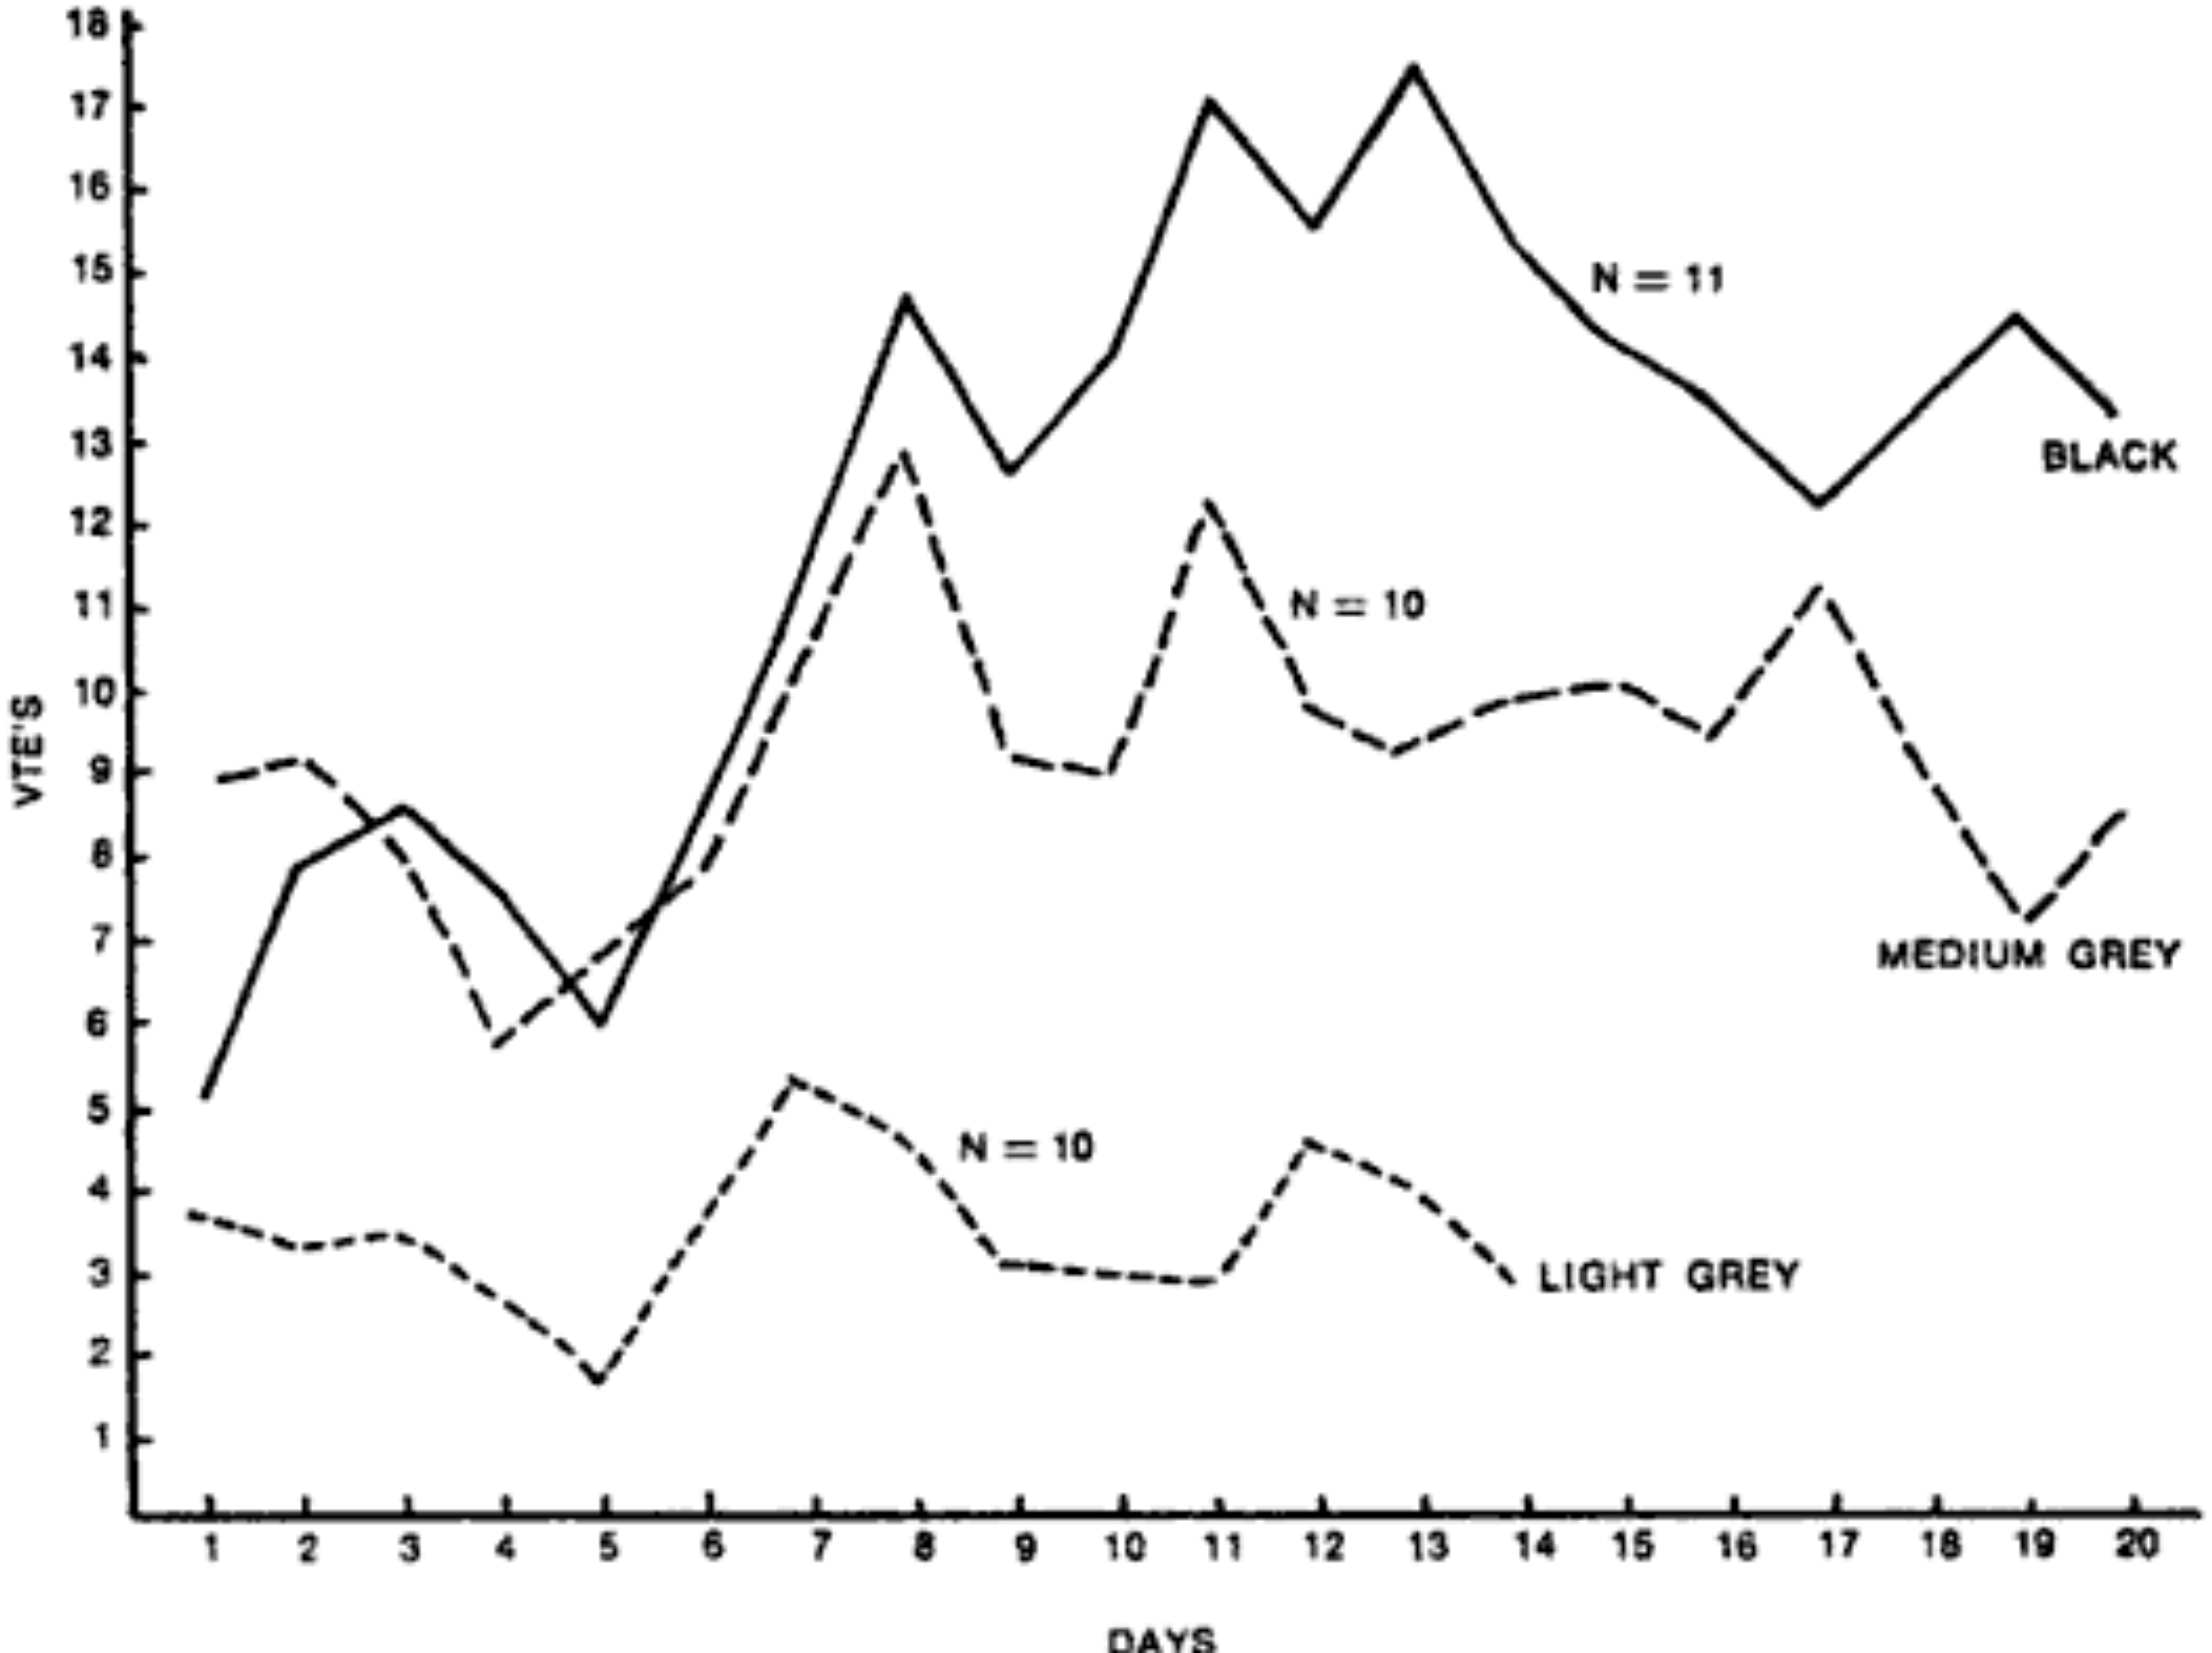

Spatial Orientation

- 3 trials of alley maze task, where H was a light shining from G-F
- Afterwards, rats transferred to sun-burst maze
  - Initially tried the C-D move, but found it blocked
  - Returned to circle and prefered the radiating path in the same direction as the original food location

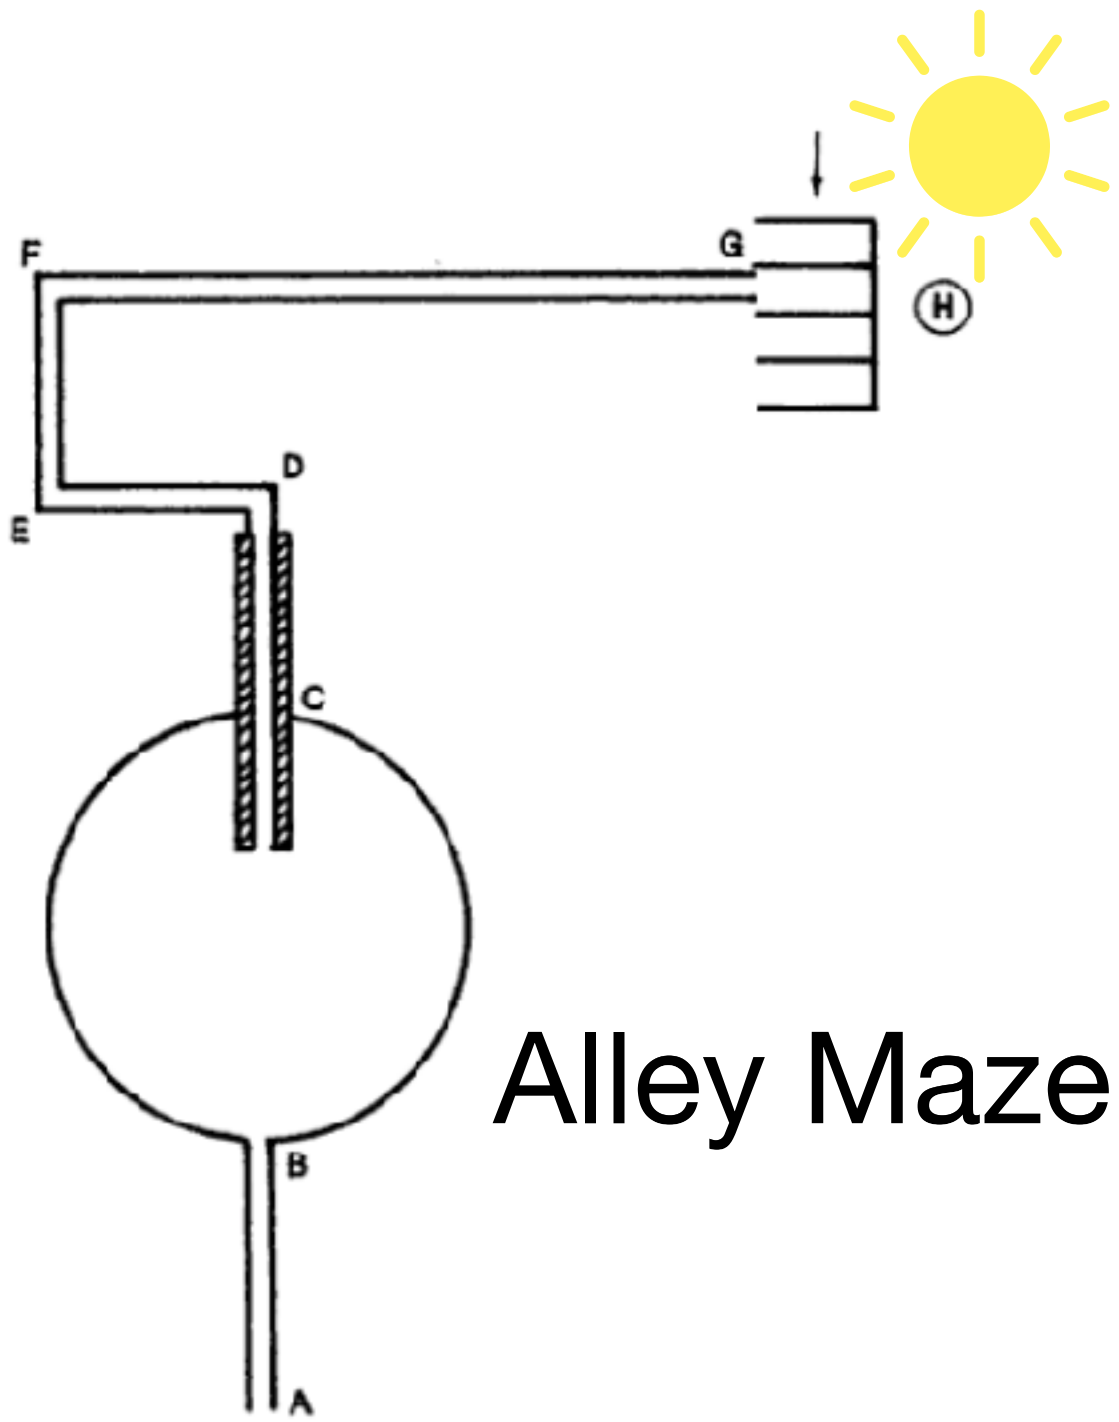

Alley Maze

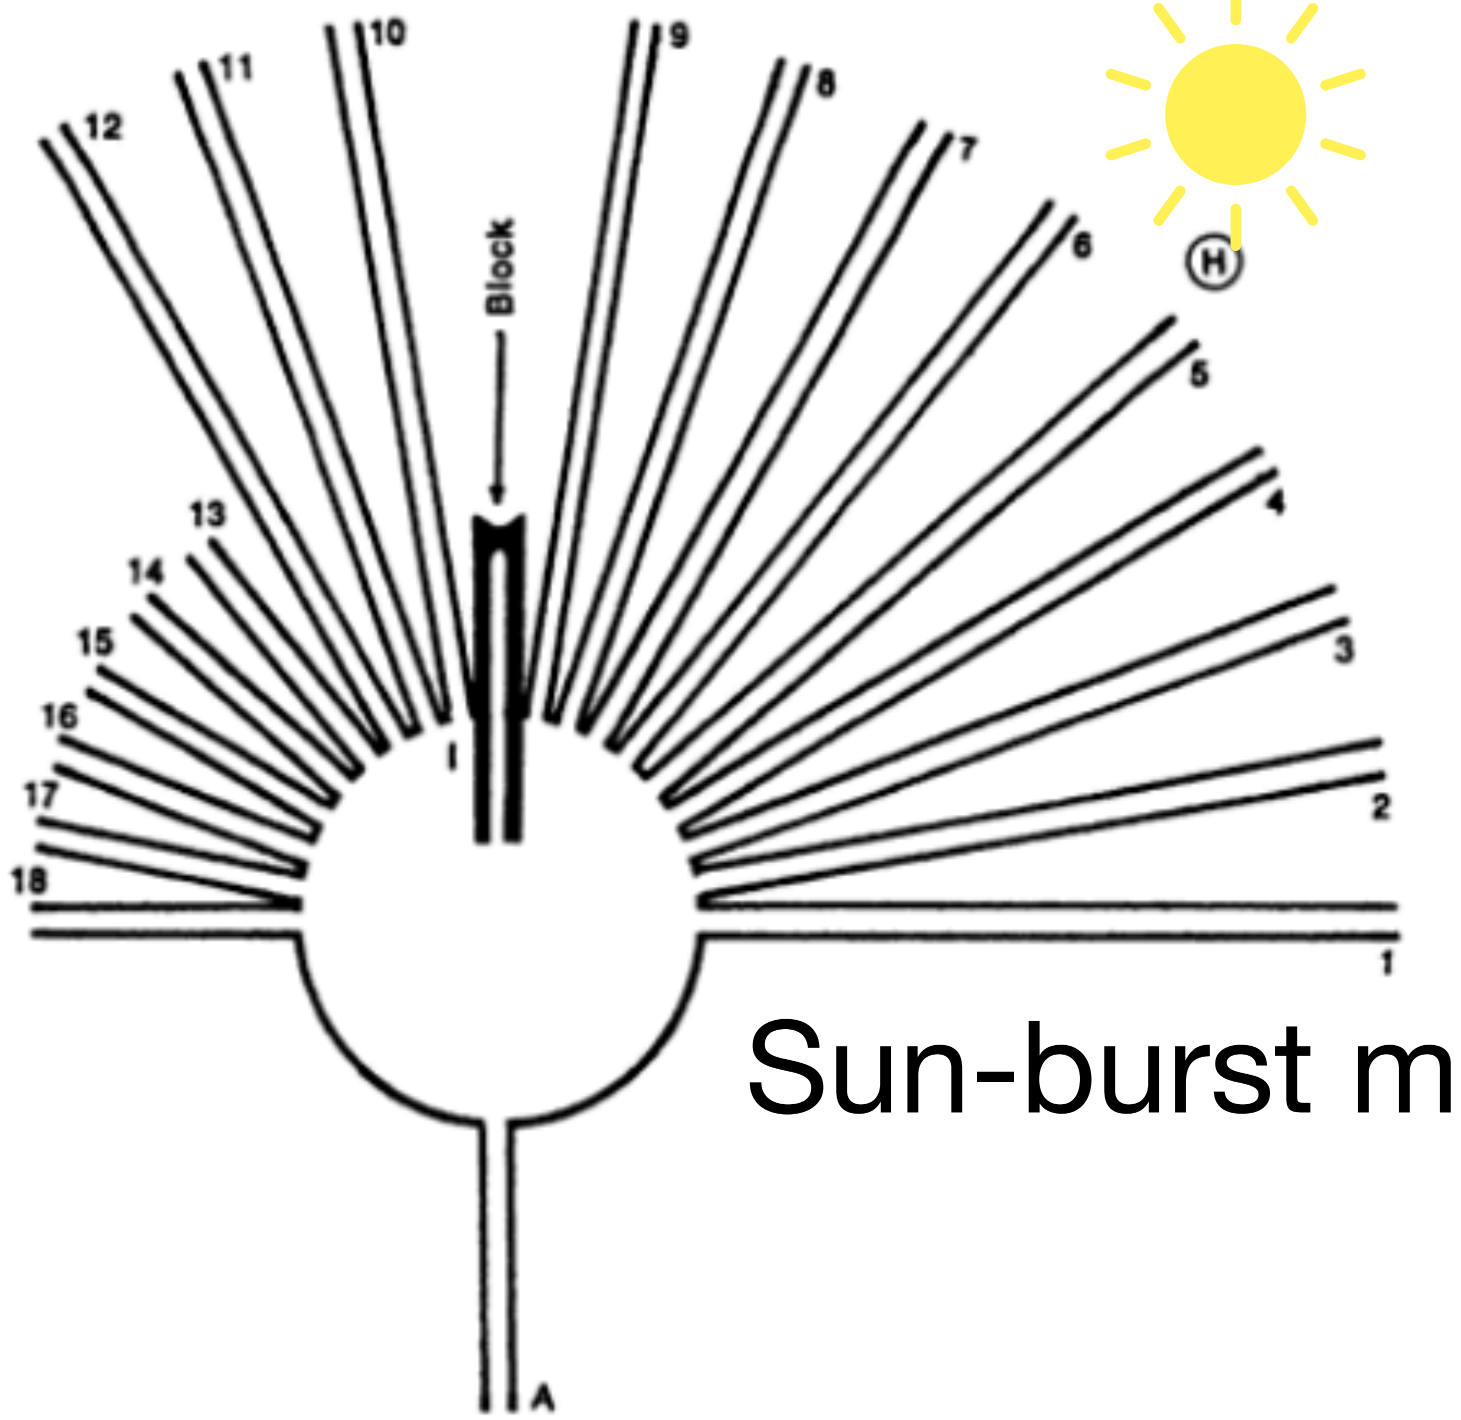

Sun-burst maze

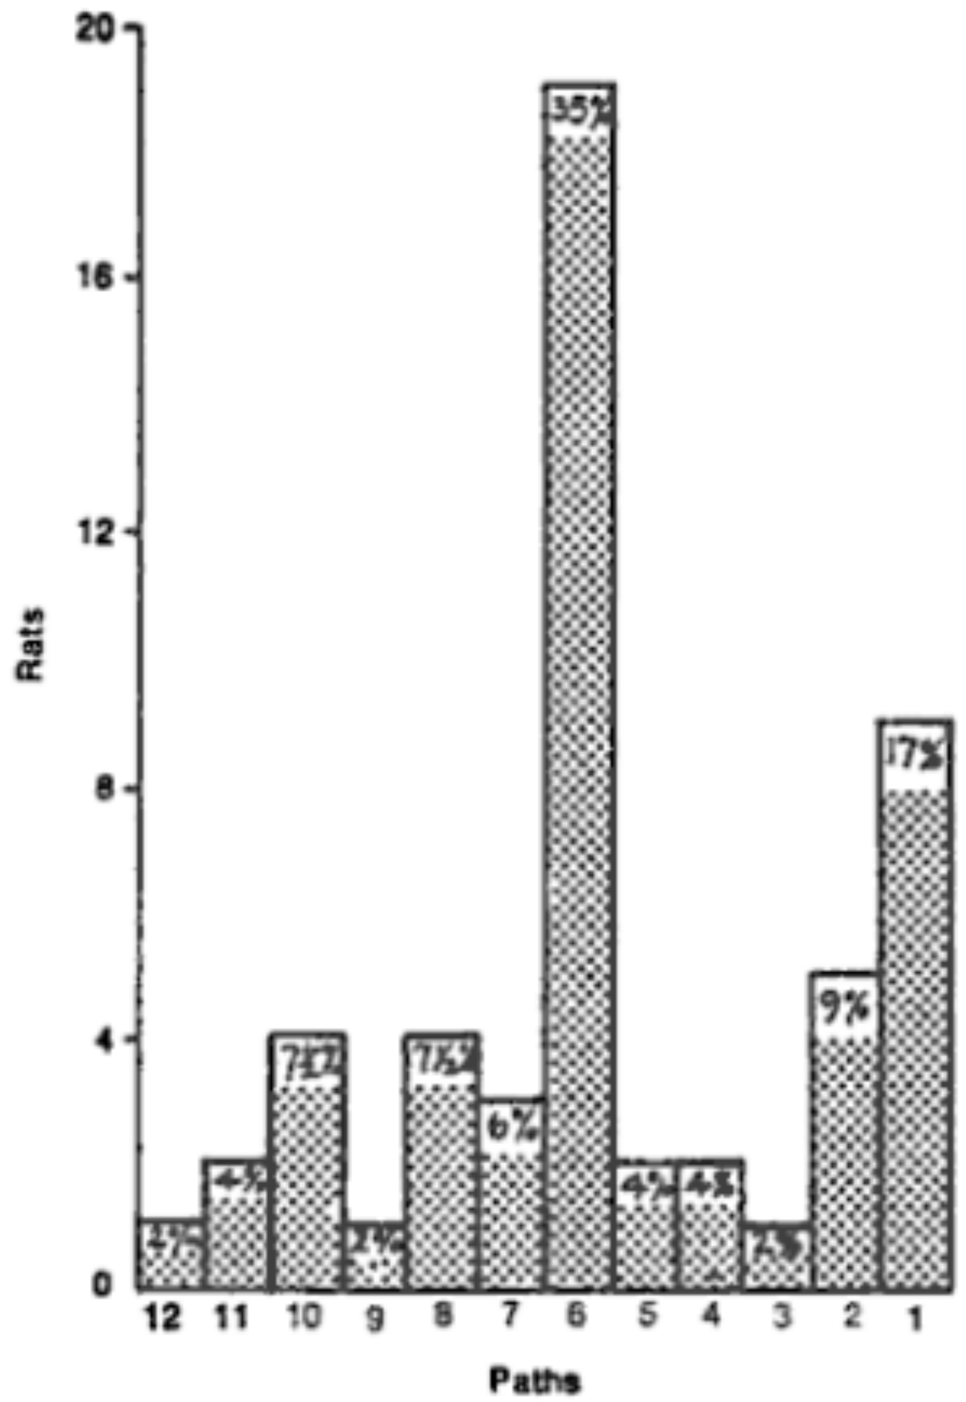

Tolman, Ritchie, & Kalish (1946)

Spatial Orientation

- 3 trials of alley maze task, where H was a light shining from G-F
- Afterwards, rats transferred to sun-burst maze
    - Initially tried the C-D move, but found it blocked
    - Returned to circle and prefered the radiating path in the same direction as the original food location

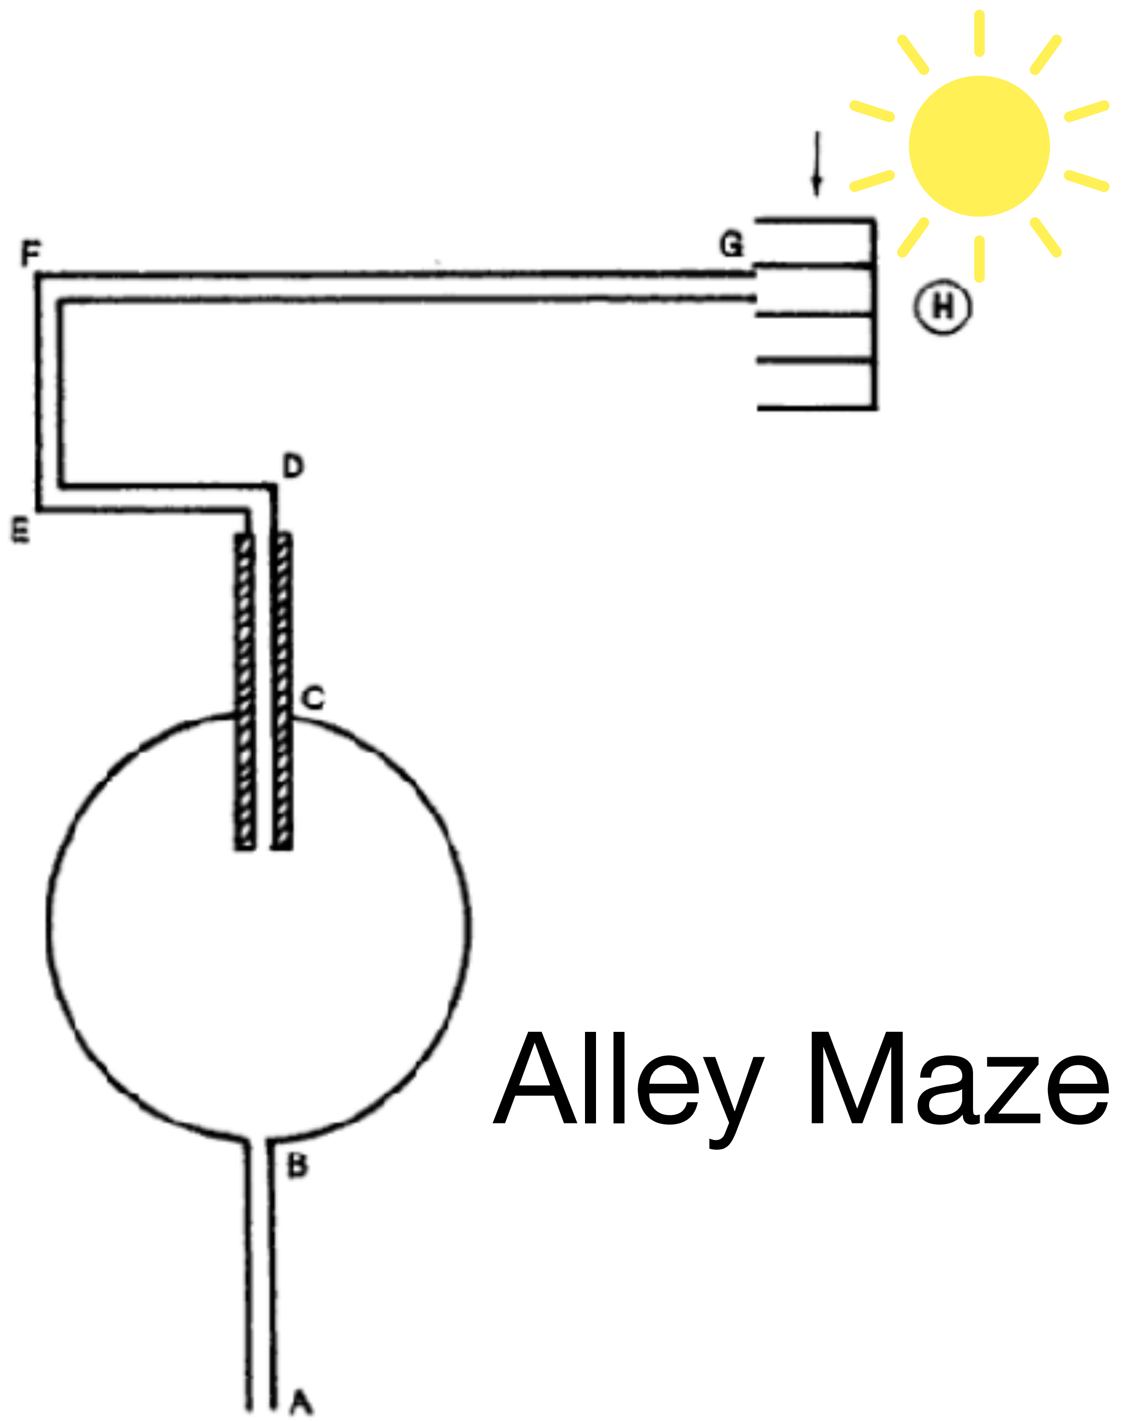

Alley Maze

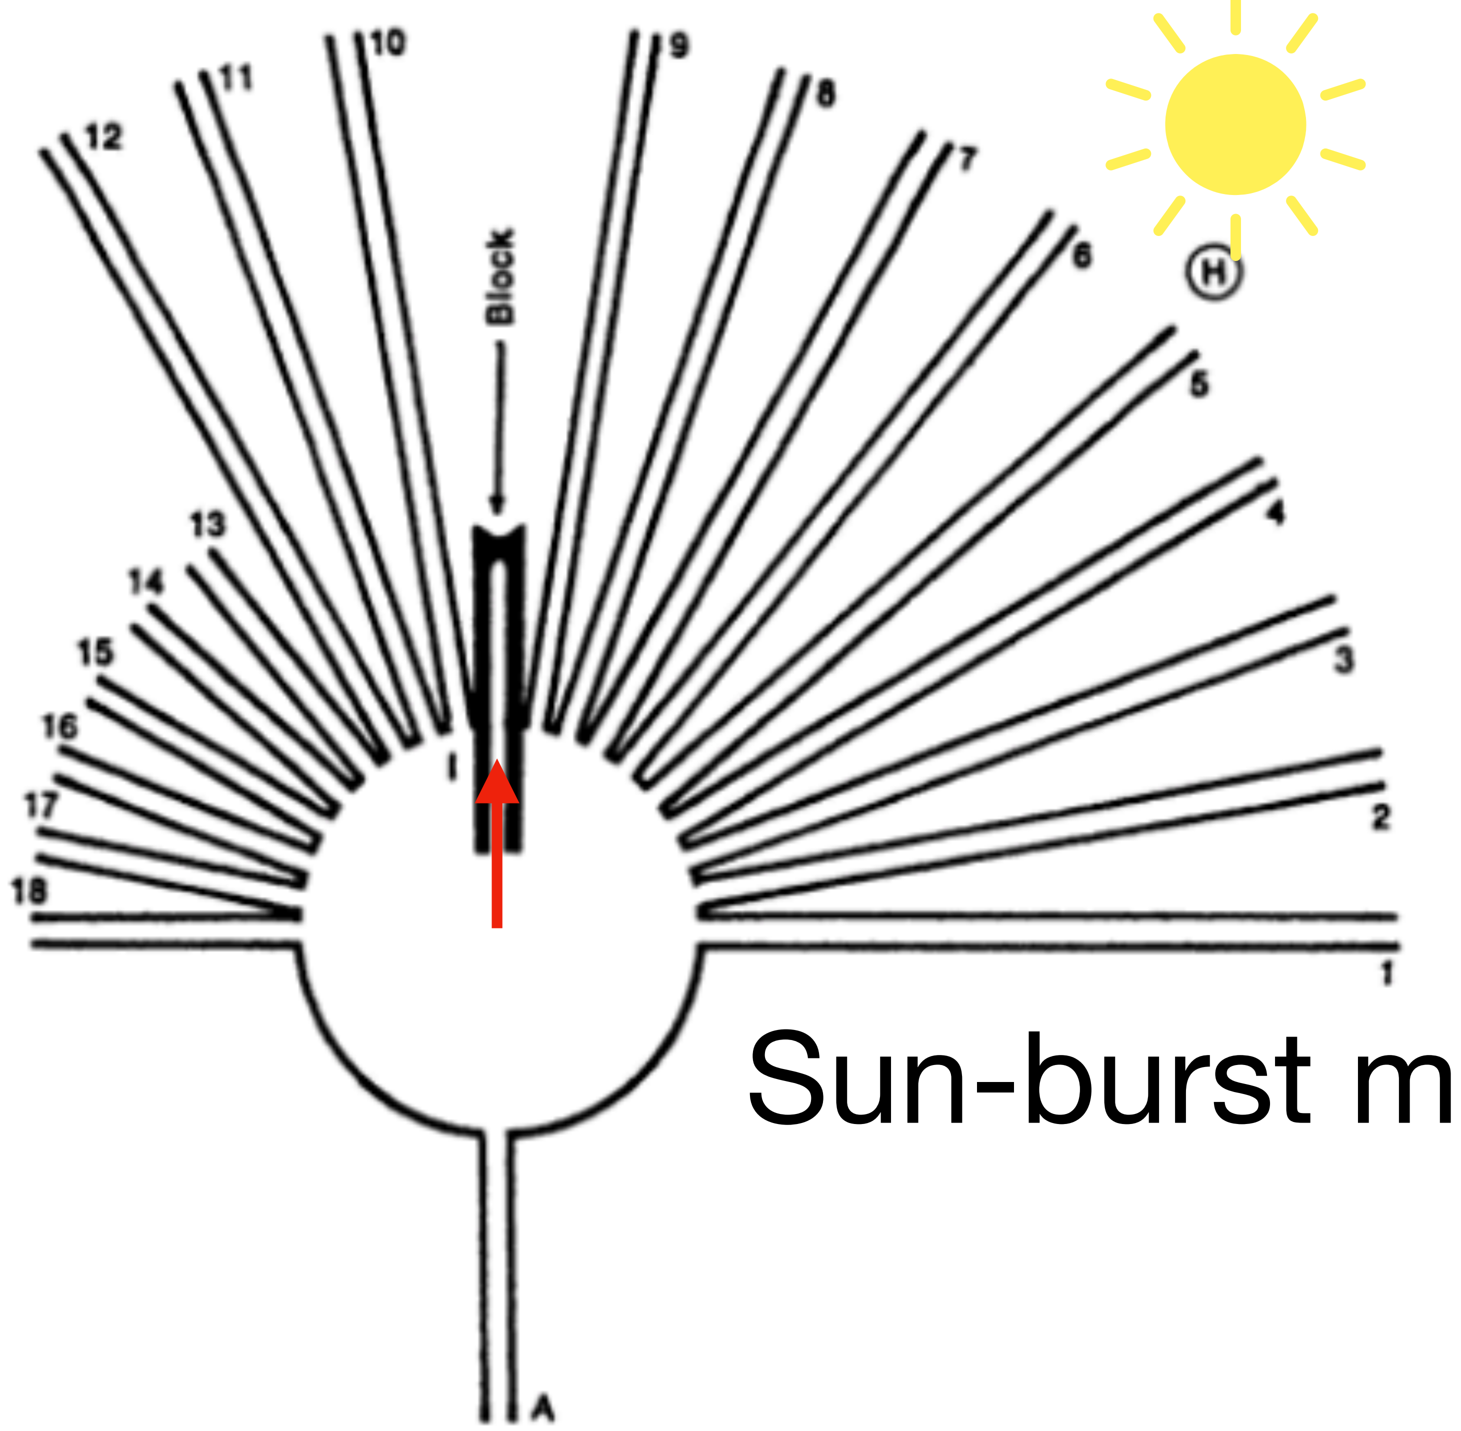

Sun-burst maze

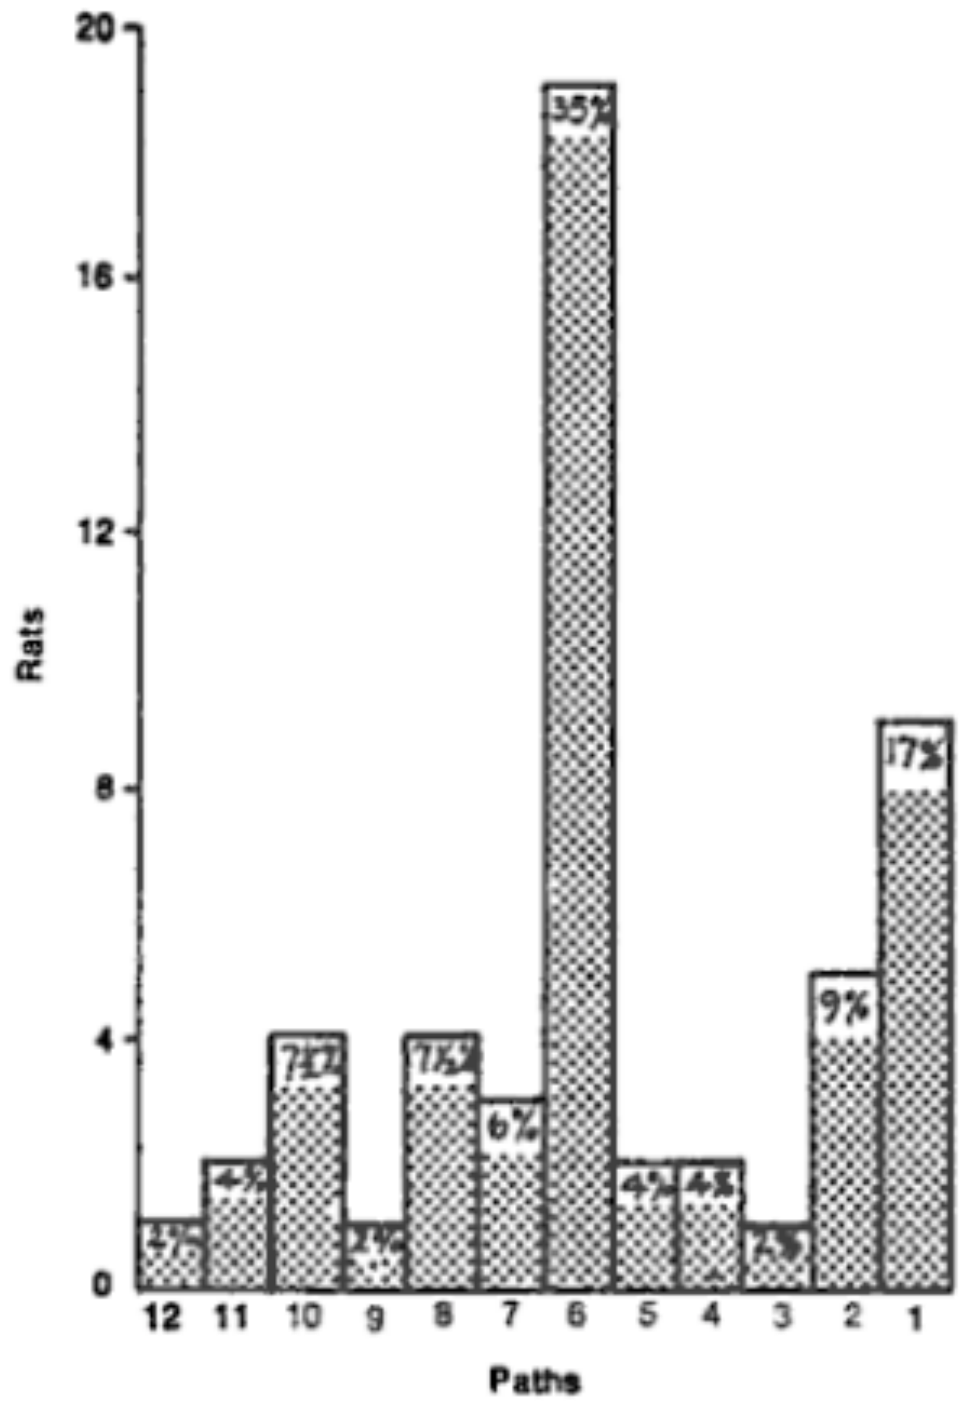

Spatial Orientation

- 3 trials of alley maze task, where H was a light shining from G-F
- Afterwards, rats transferred to sun-burst maze
  - Initially tried the C-D move, but found it blocked
  - Returned to circle and prefered the radiating path in the same direction as the original food location

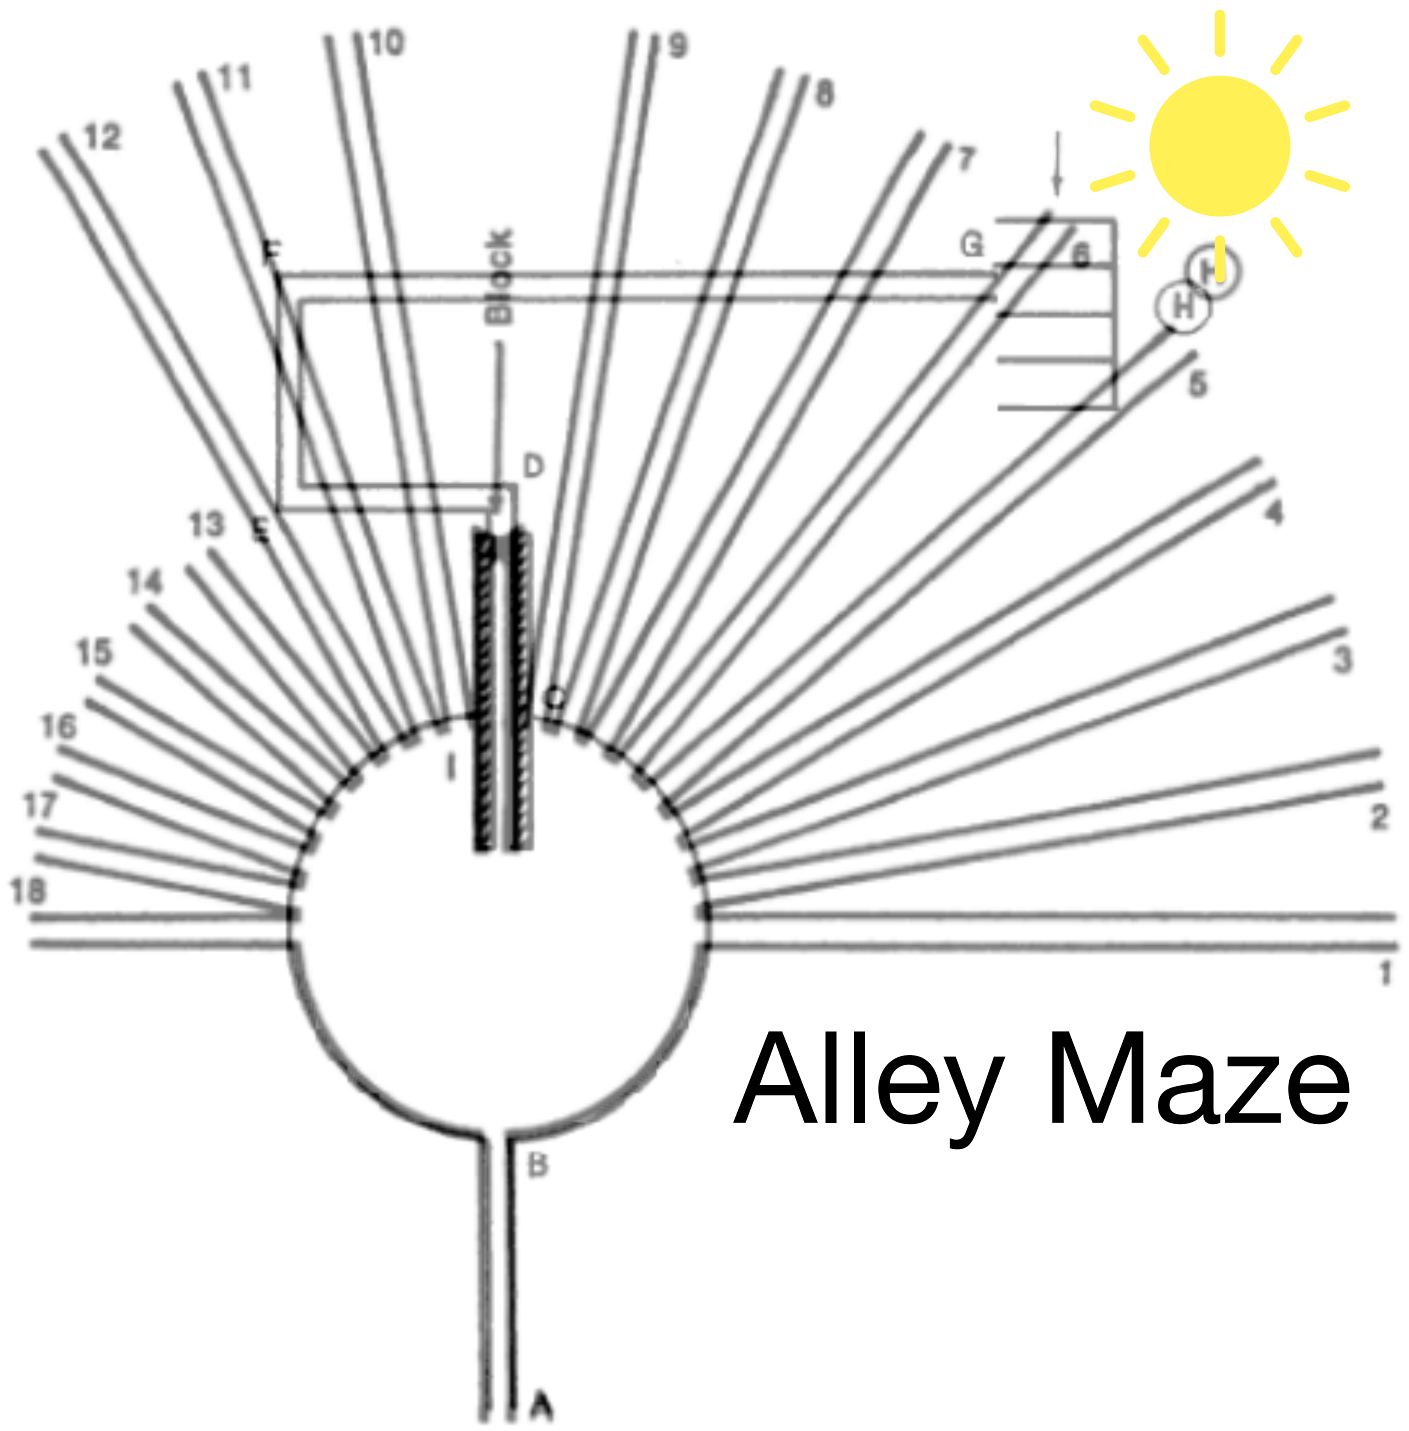

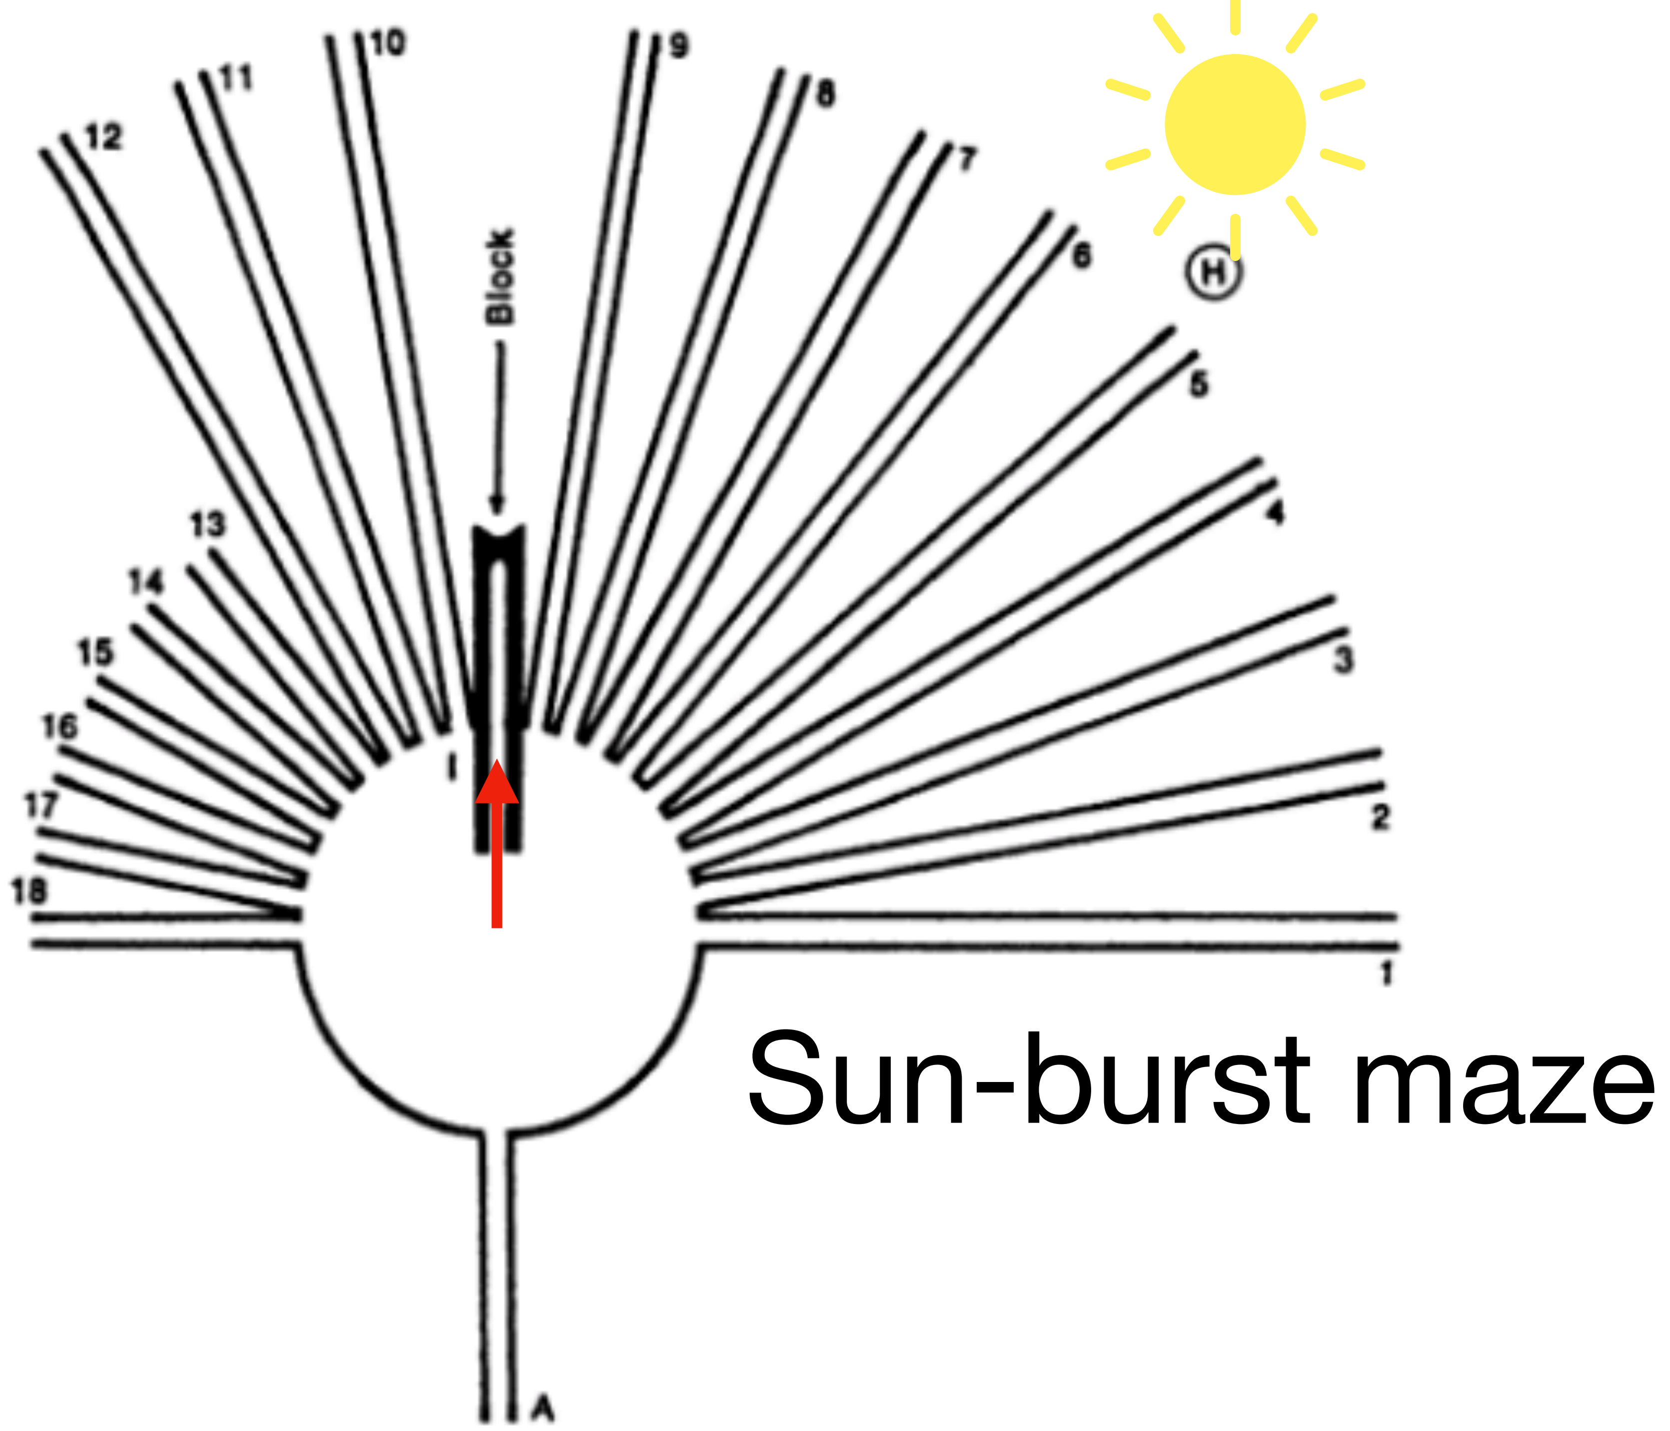

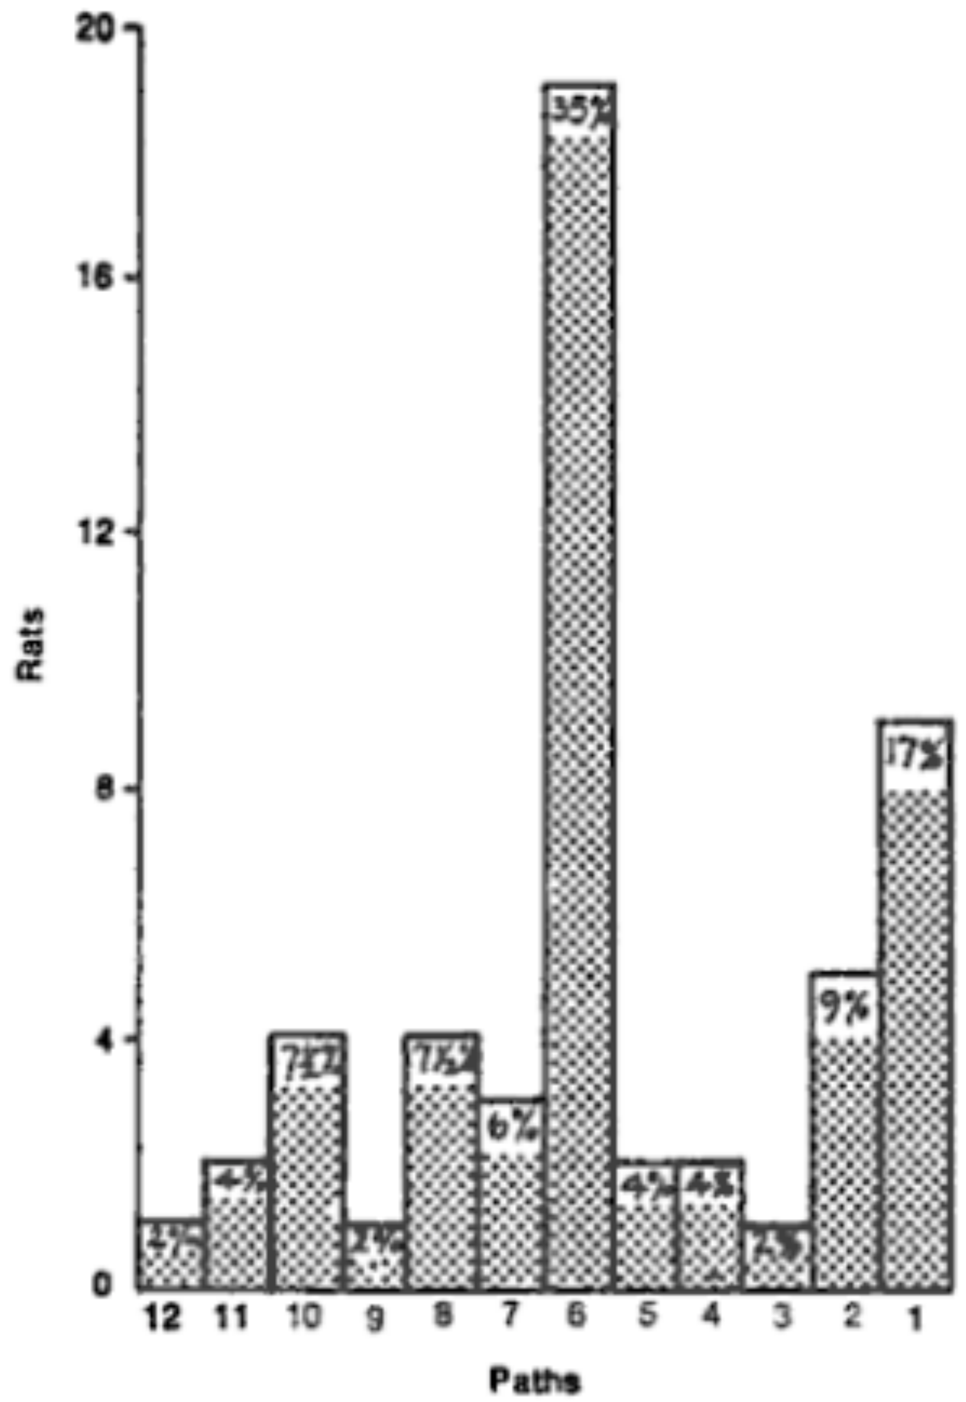

Tolman, Ritchie, & Kalish (1946)

Spatial Orientation

- 3 trials of alley maze task, where H was a light shining from G-F
- Afterwards, rats transferred to sun-burst maze
  - Initially tried the C-D move, but found it blocked
  - Returned to circle and prefered the radiating path in the same direction as the original food location

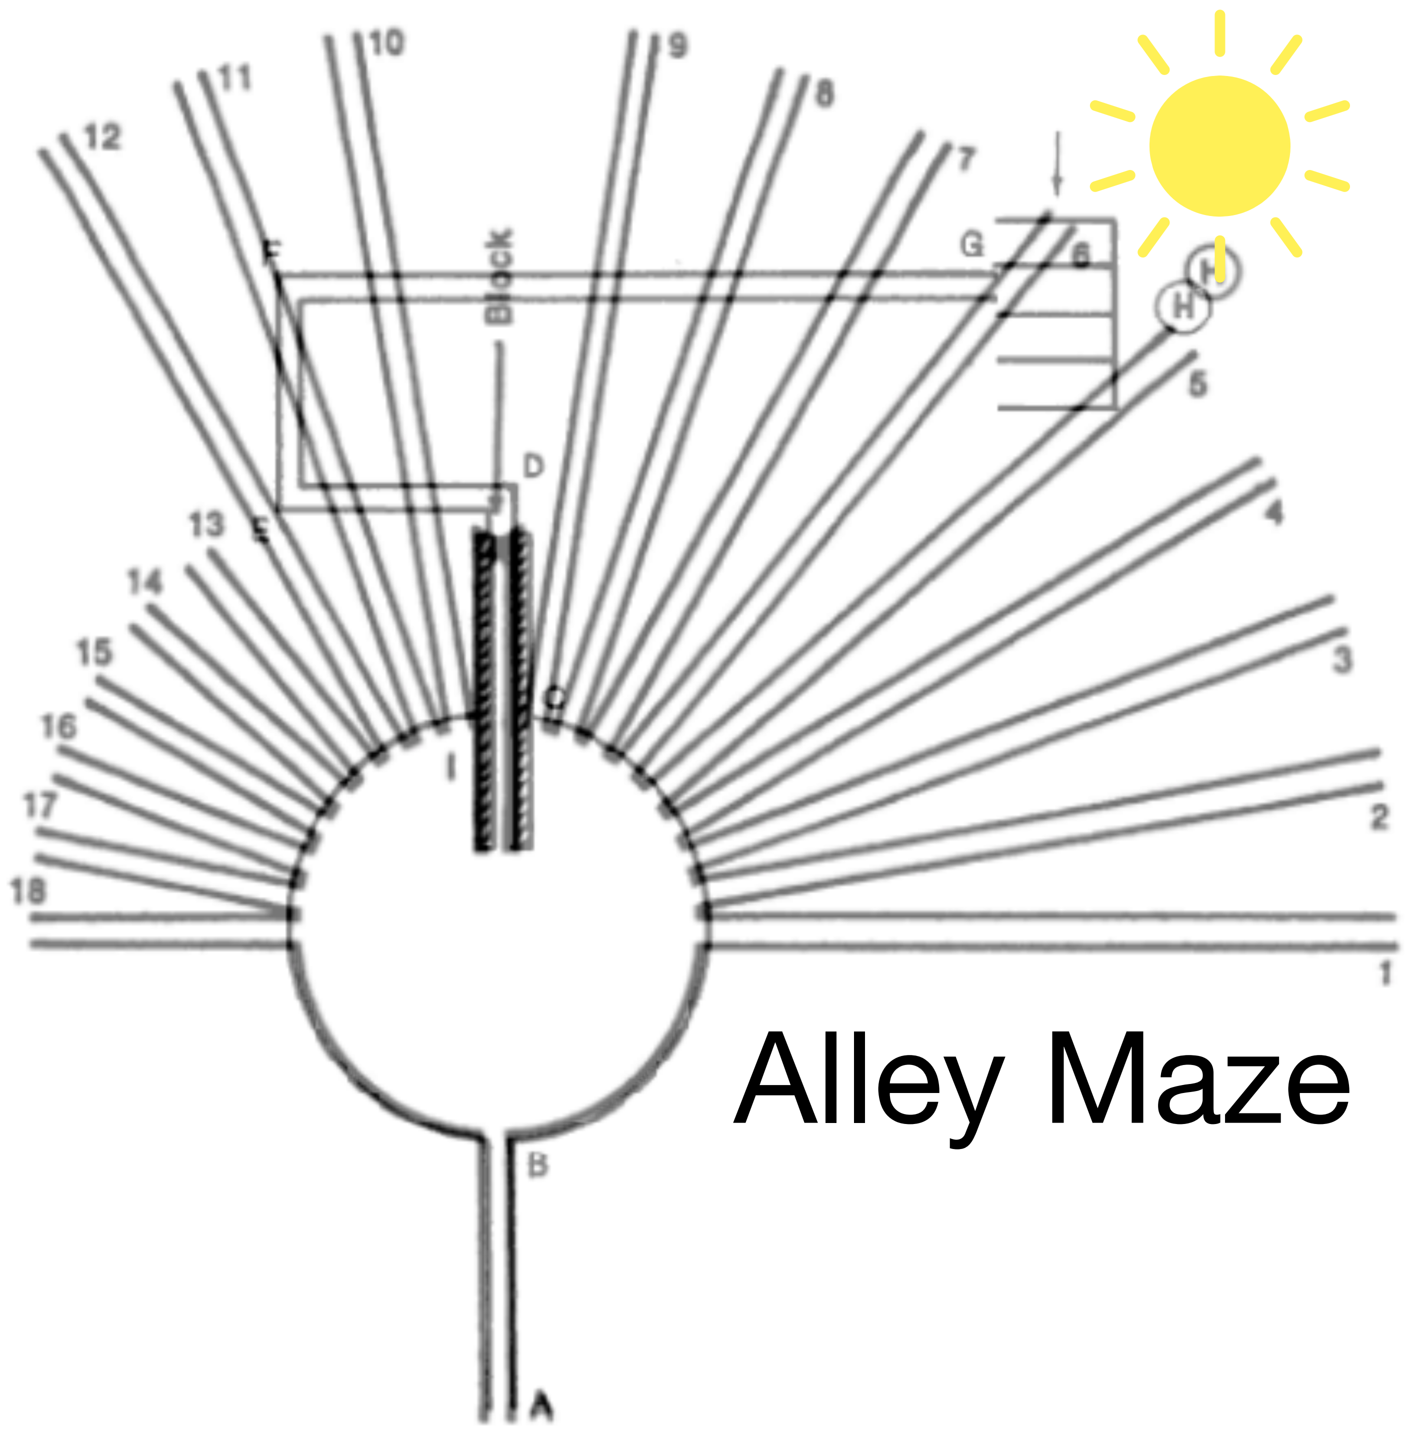

Alley Maze

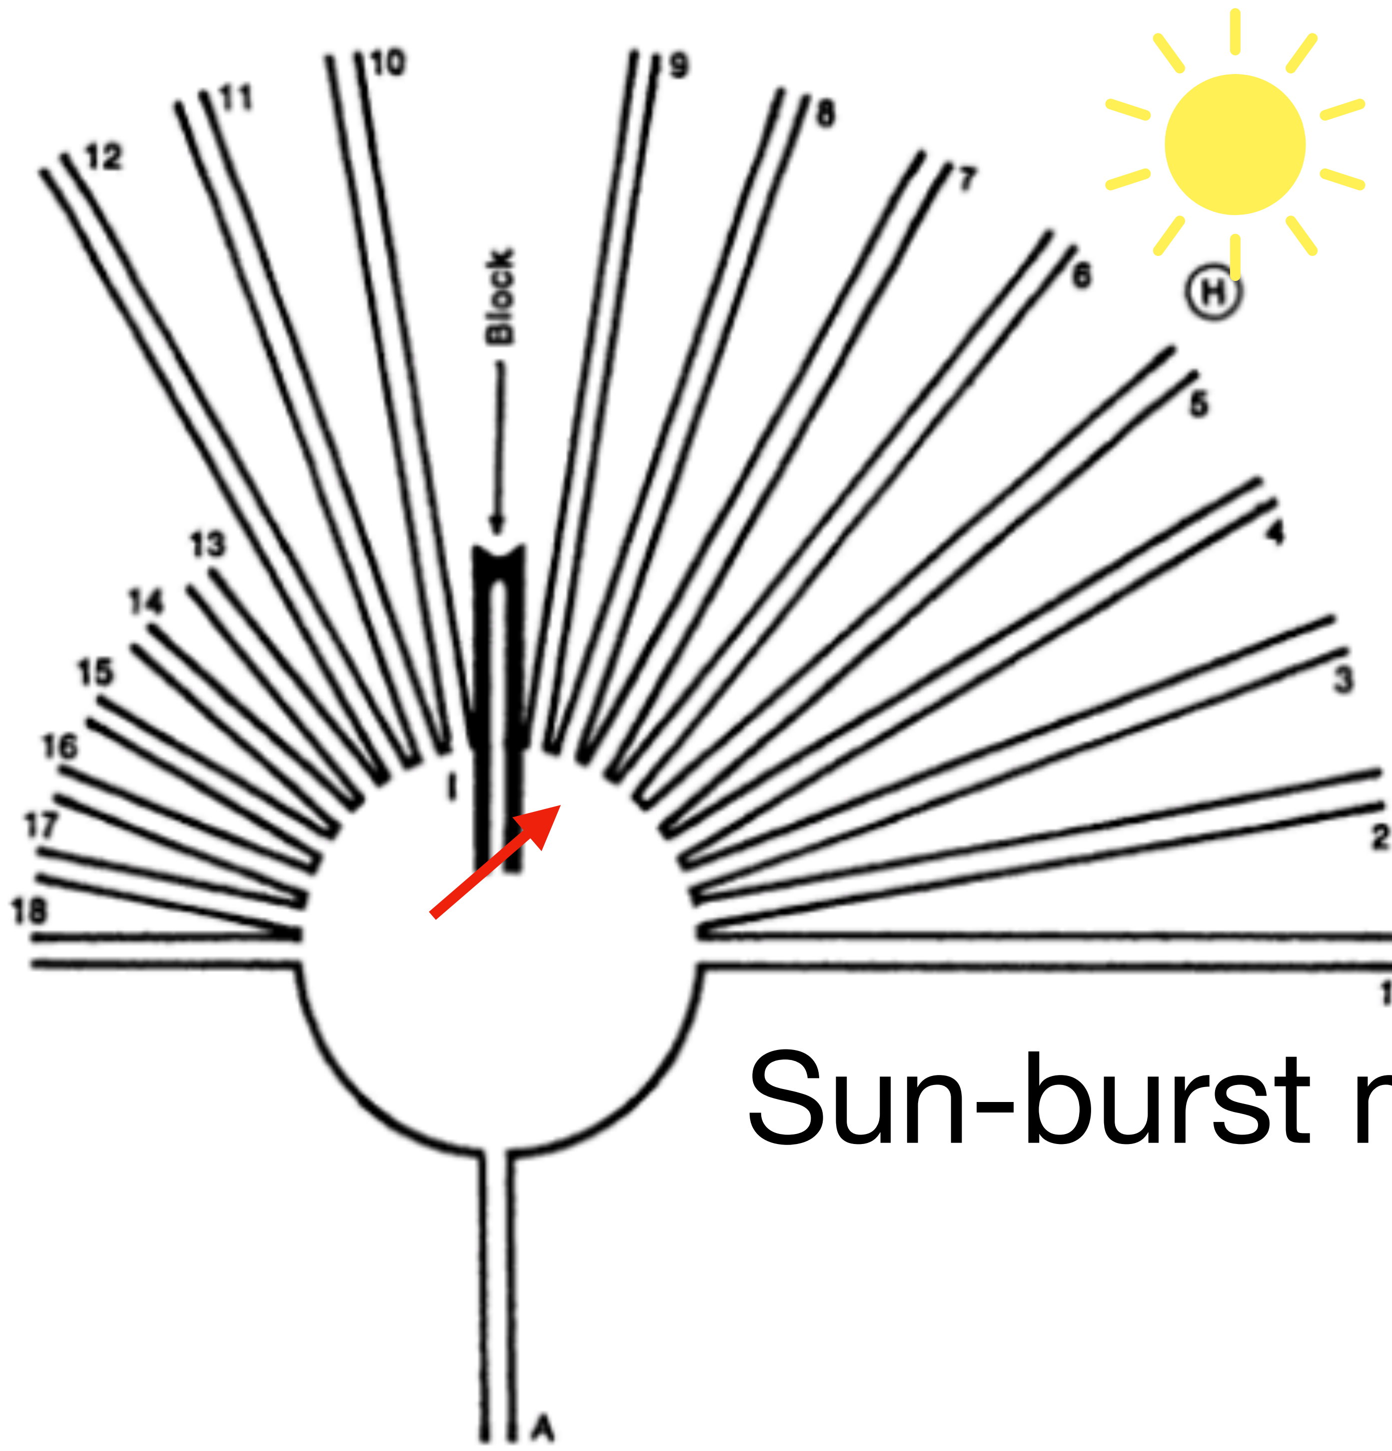

Sun-burst maze

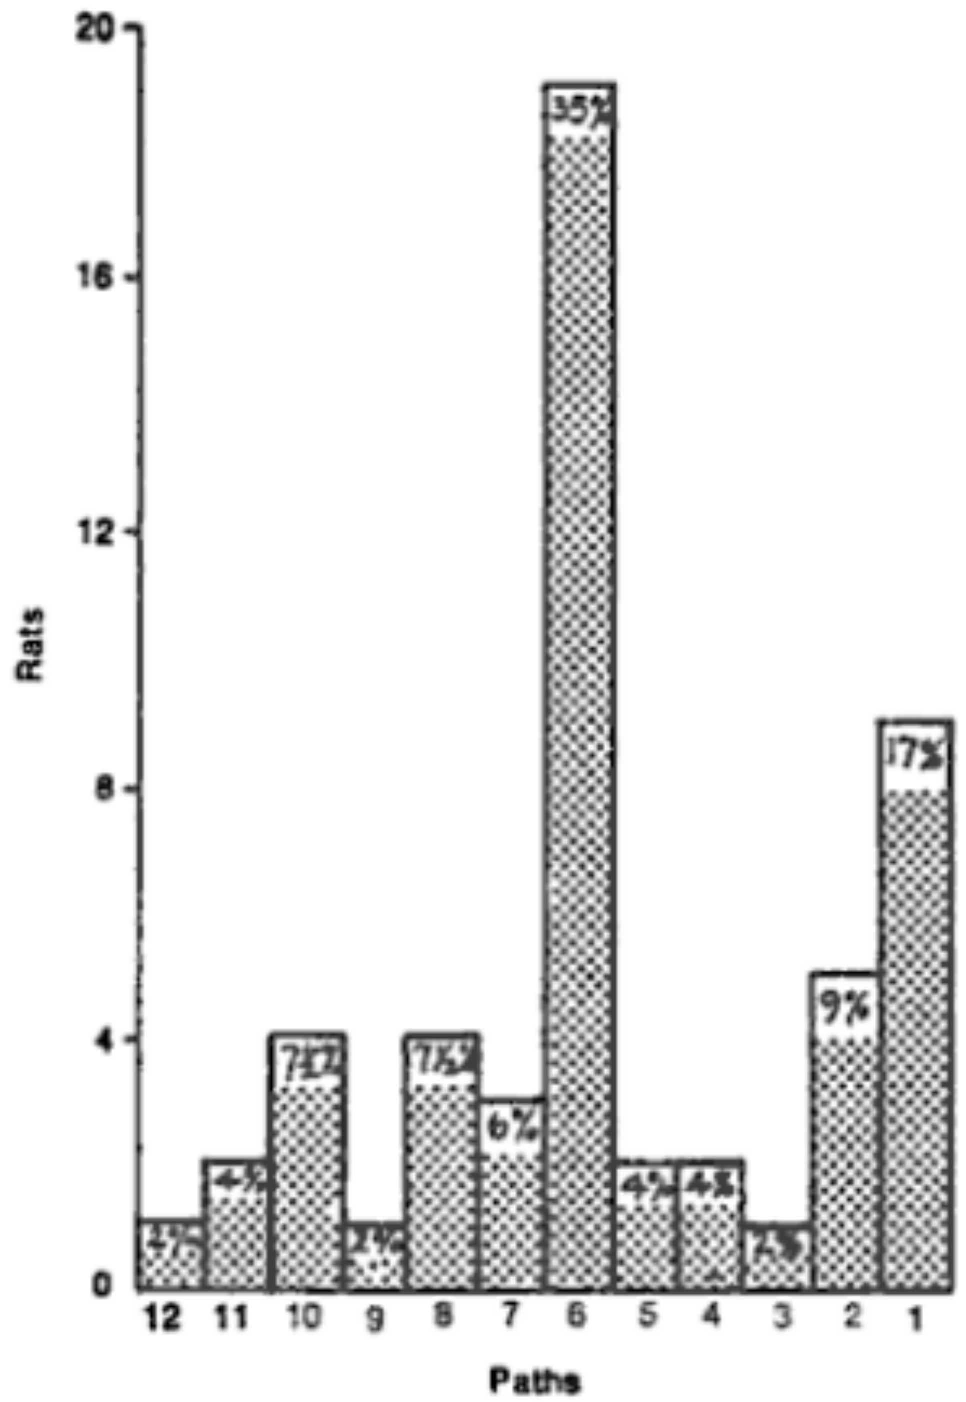

Cognitive Maps shape generalization

- The nature of the maps we learn shape how we generalize
  - *"the narrower and more strip-like the original map, the less will it carry over successfully to the new problem; whereas, the wider and the more comprehensive it was, the more adequately it will serve in the new set-up"*
- What conditions favor learning a narrow strip-map vs. a broad comprehensive map?
  - narrow maps induced by :

    1) damaged brains

    2) impoverished environments

    3) overdose of repetition

    4) too strongly motivational/frustrating conditions

Maladaptive psychopathologies

- **Regression** to childlike behavior

  "take an example, the overprotected middle-aged woman […] who, after losing her husband, regressed […] into dressing in too youthful a fashion and into competing for their beaux and then finally into behaving like a child requiring continuous care […]"

- **Fixation** on various addictive behaviors

  "If rats are too strongly motivated in their original learning, they find it very difficult to relearn when the original path is no longer correct"

- **Displacement** of agression towards outgroups

  - *"The individual comes no longer to distinguish the true locus of the cause of his frustration"*
  - *"The poor Southern whites, who take it out on the Negroes, are displacing their aggressions from the landlords"*
  - *"the southern economic system, the northern capitalists, or wherever the true cause of their frustration may lie, [displace their frustration] onto a mere convenient outgroup*
  - *[physicists vs. humanities, psychologists vs. all other depts., university vs. secondary school, americans vs. russians]…*
  - *"nothing more than such irrational displacements of our aggressions onto outgroups"*

What is the solution?

“We must, in short, subject our children and ourselves … to the optimal conditions of moderate motivation and of an absence of unnecessary frustrations…. I cannot predict whether or not we will be able, or be allowed, to do this; but I ***can*** *say that, only insofar as we* ***are*** *able and* ***are*** *allowed, have we cause for hope.*

What is the solution?

*"We must, in short, subject our children and ourselves … to the optimal conditions of moderate motivation and of an absence of unnecessary frustrations…. I cannot predict whether or not we will be able, or be allowed, to do this; but I **can** say that, only insofar as we **are** able and **are** allowed, have we cause for hope.*

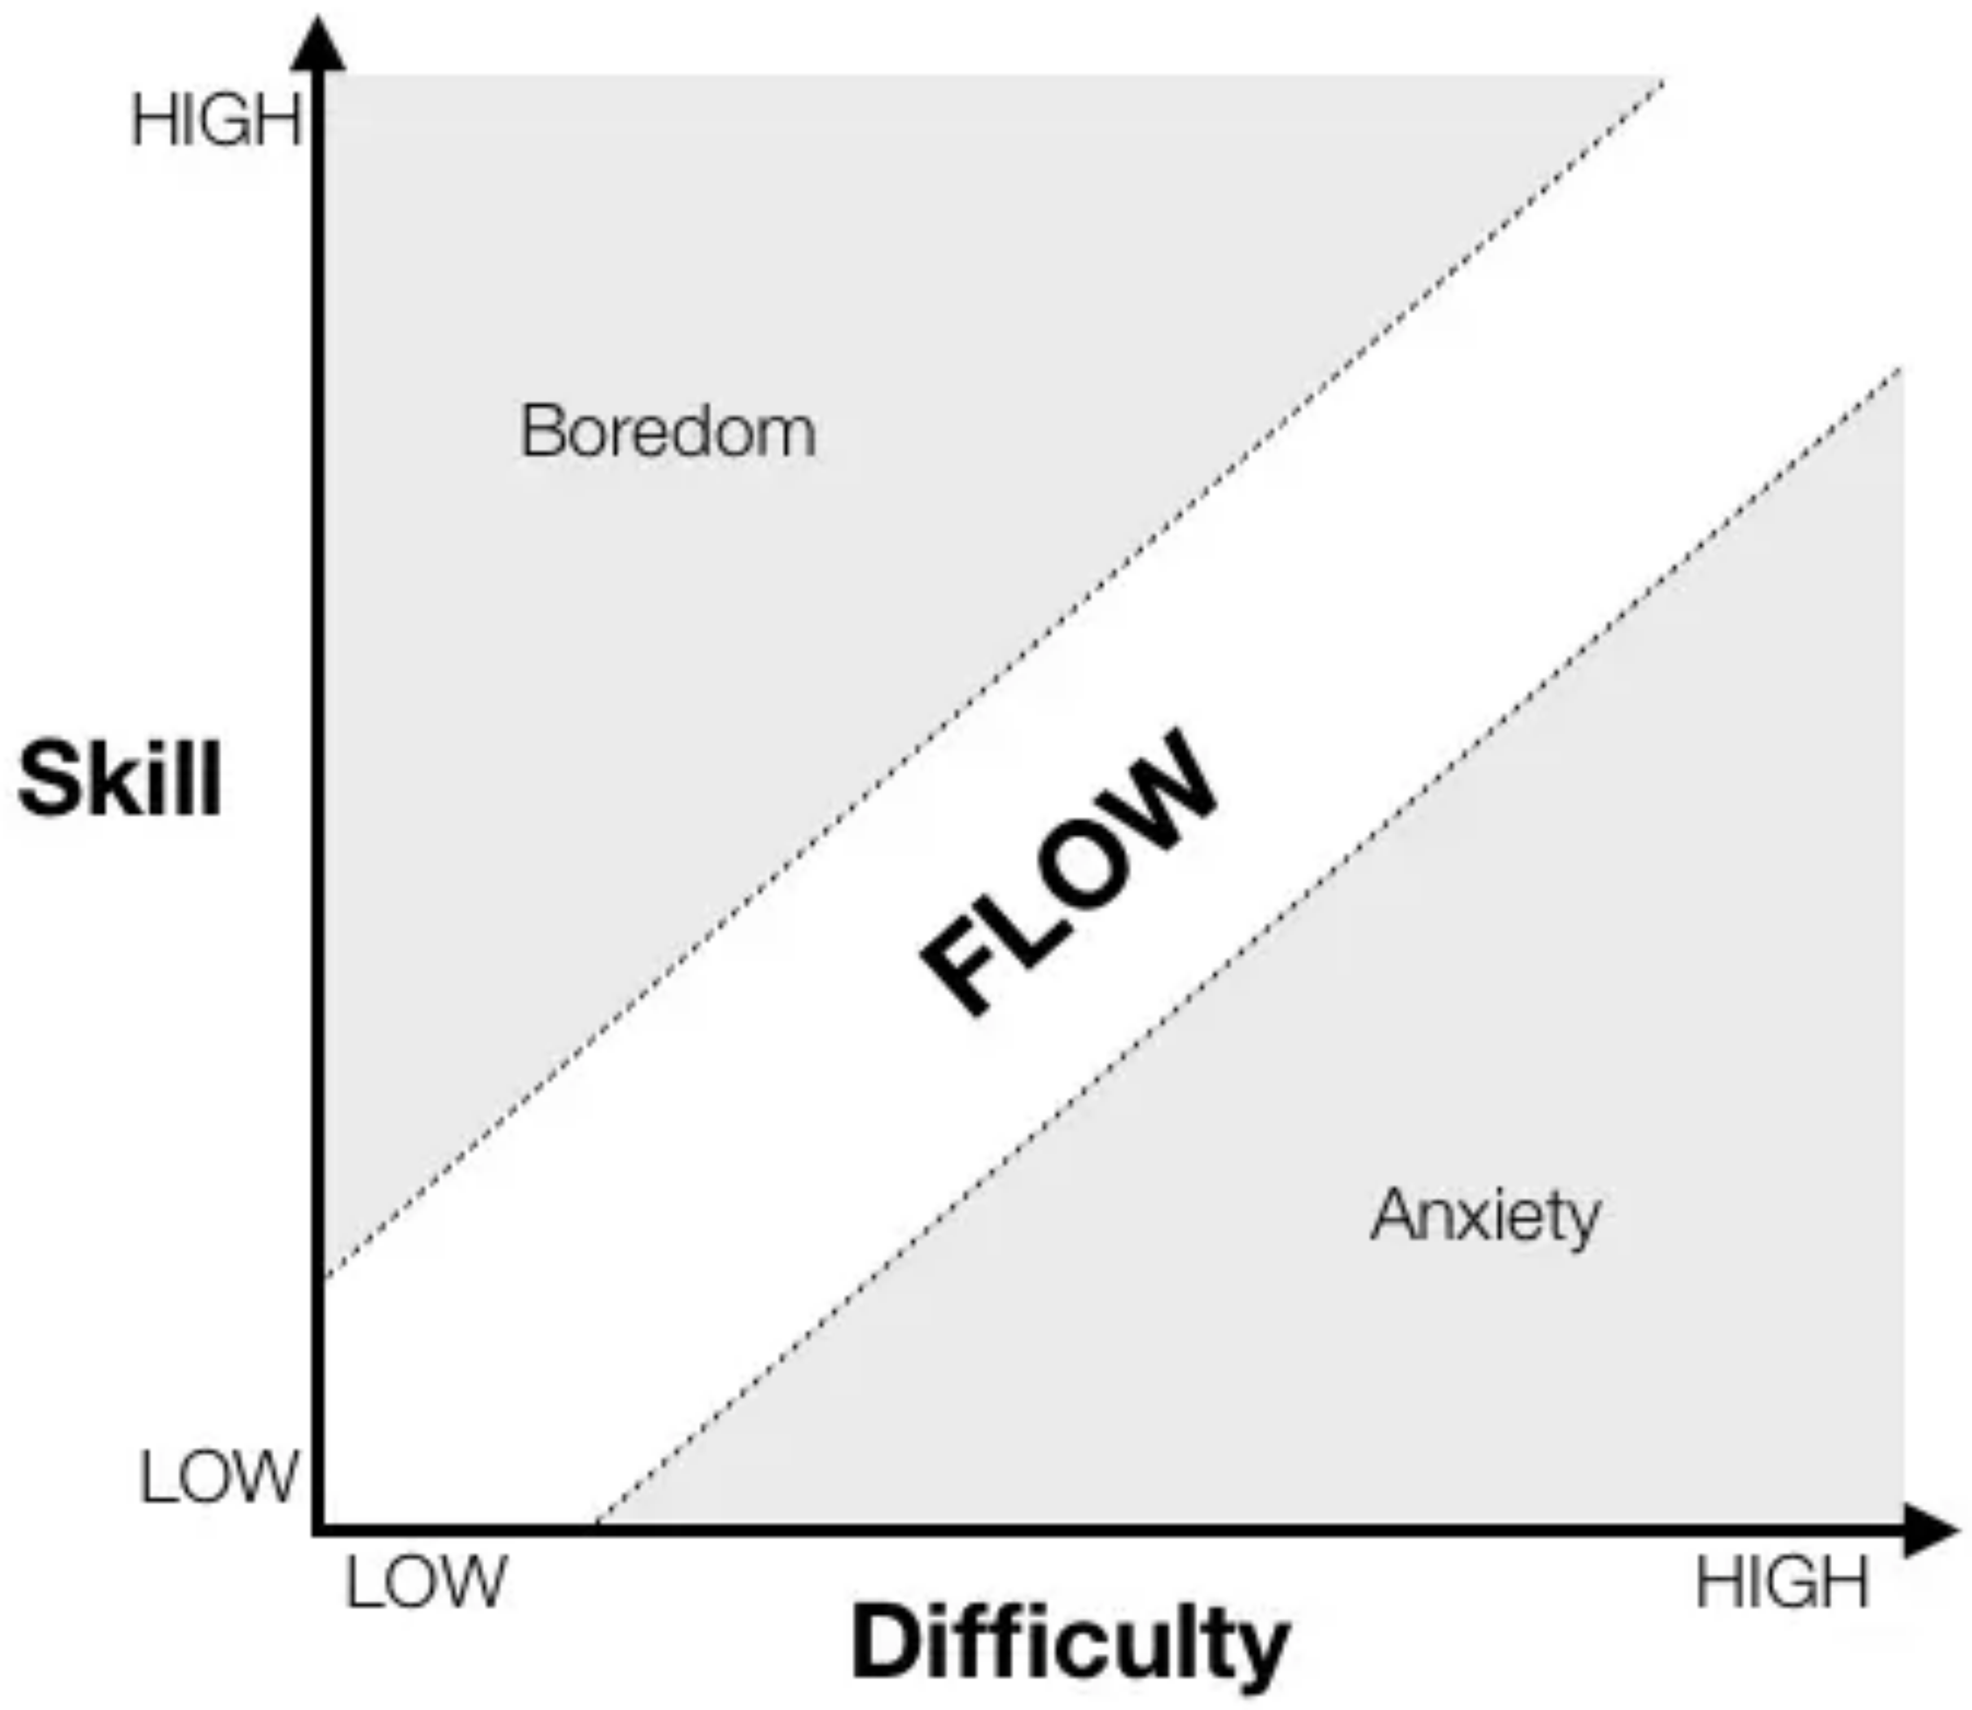

Csikszentmihalyi (1990)

Cognitive Maps in the Brain

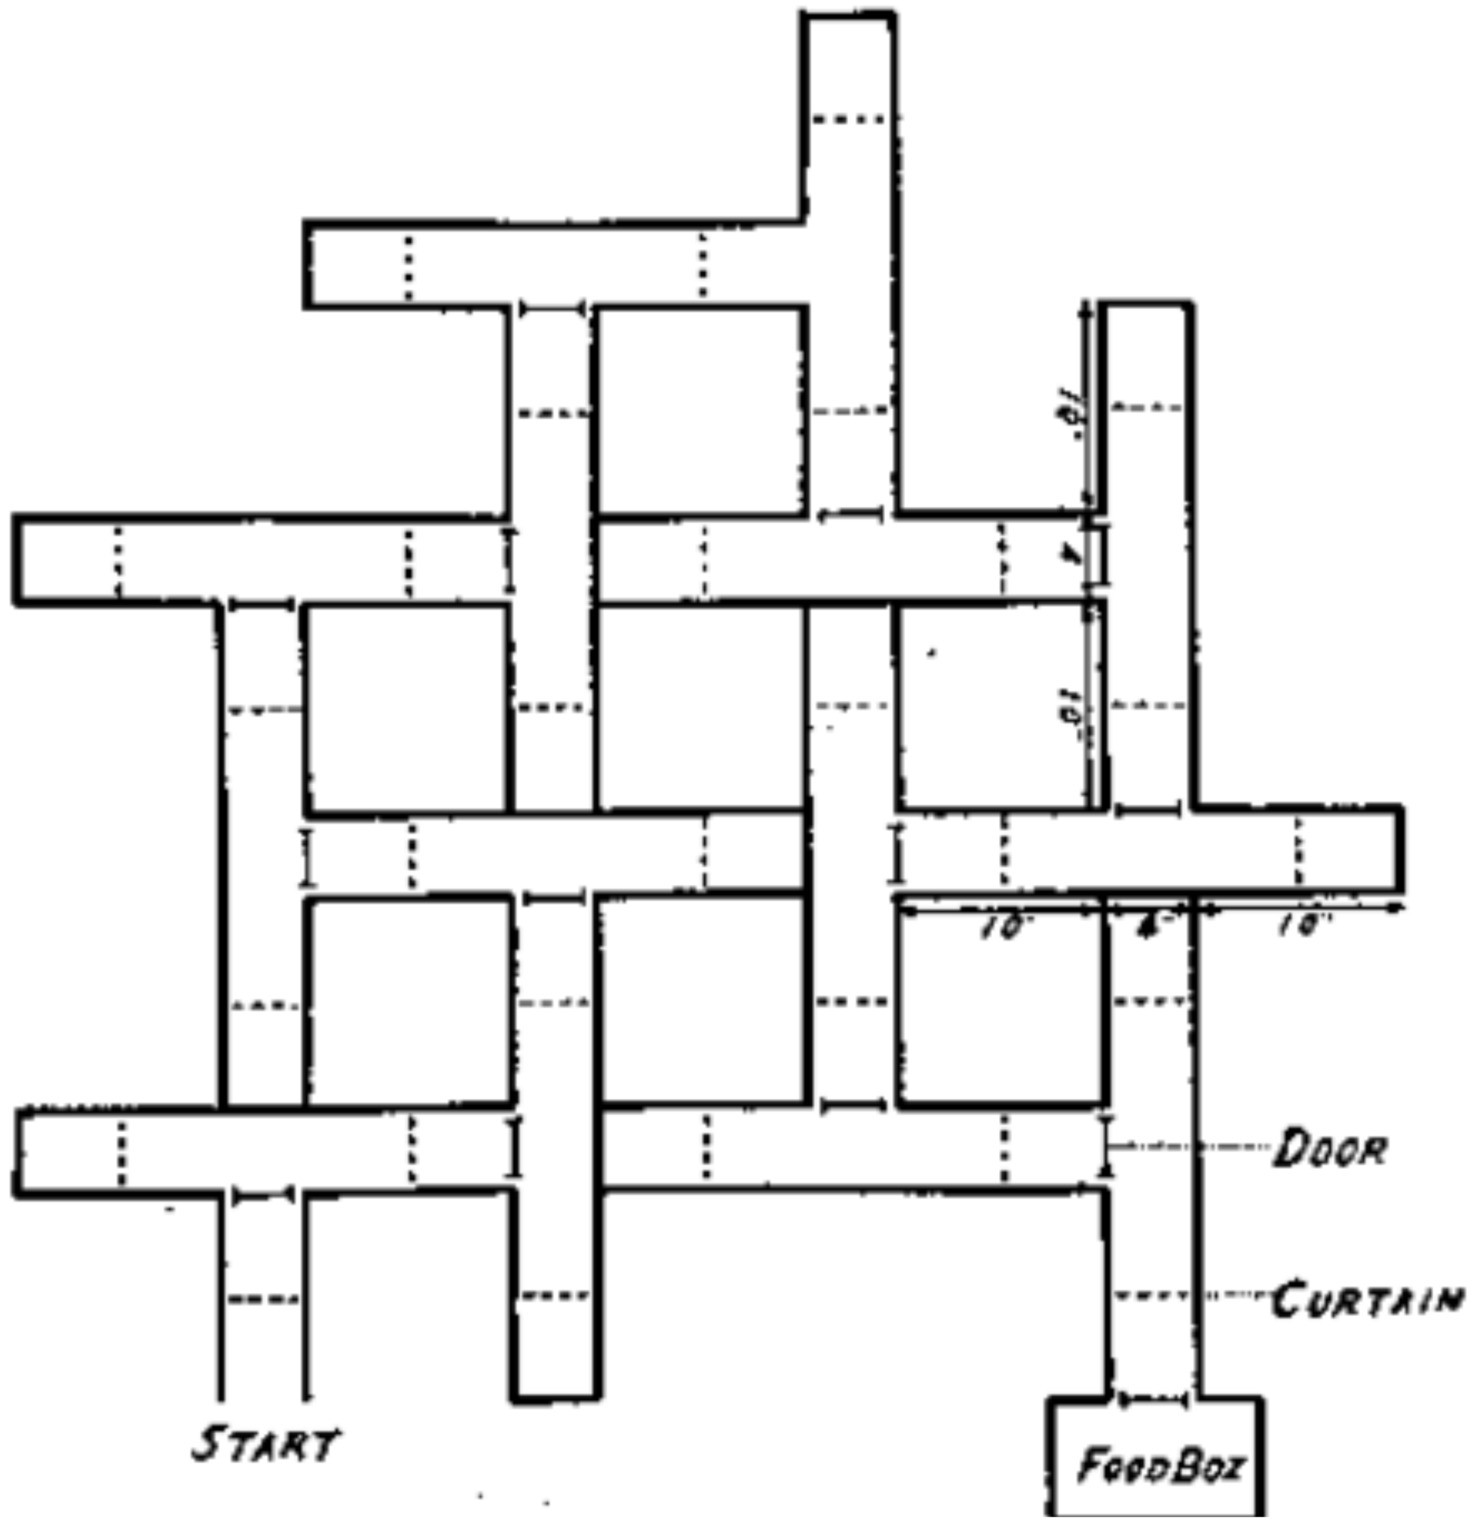

Place cells in the hippocampus represent location in an environment

(O'keefe & Nadel 1978)

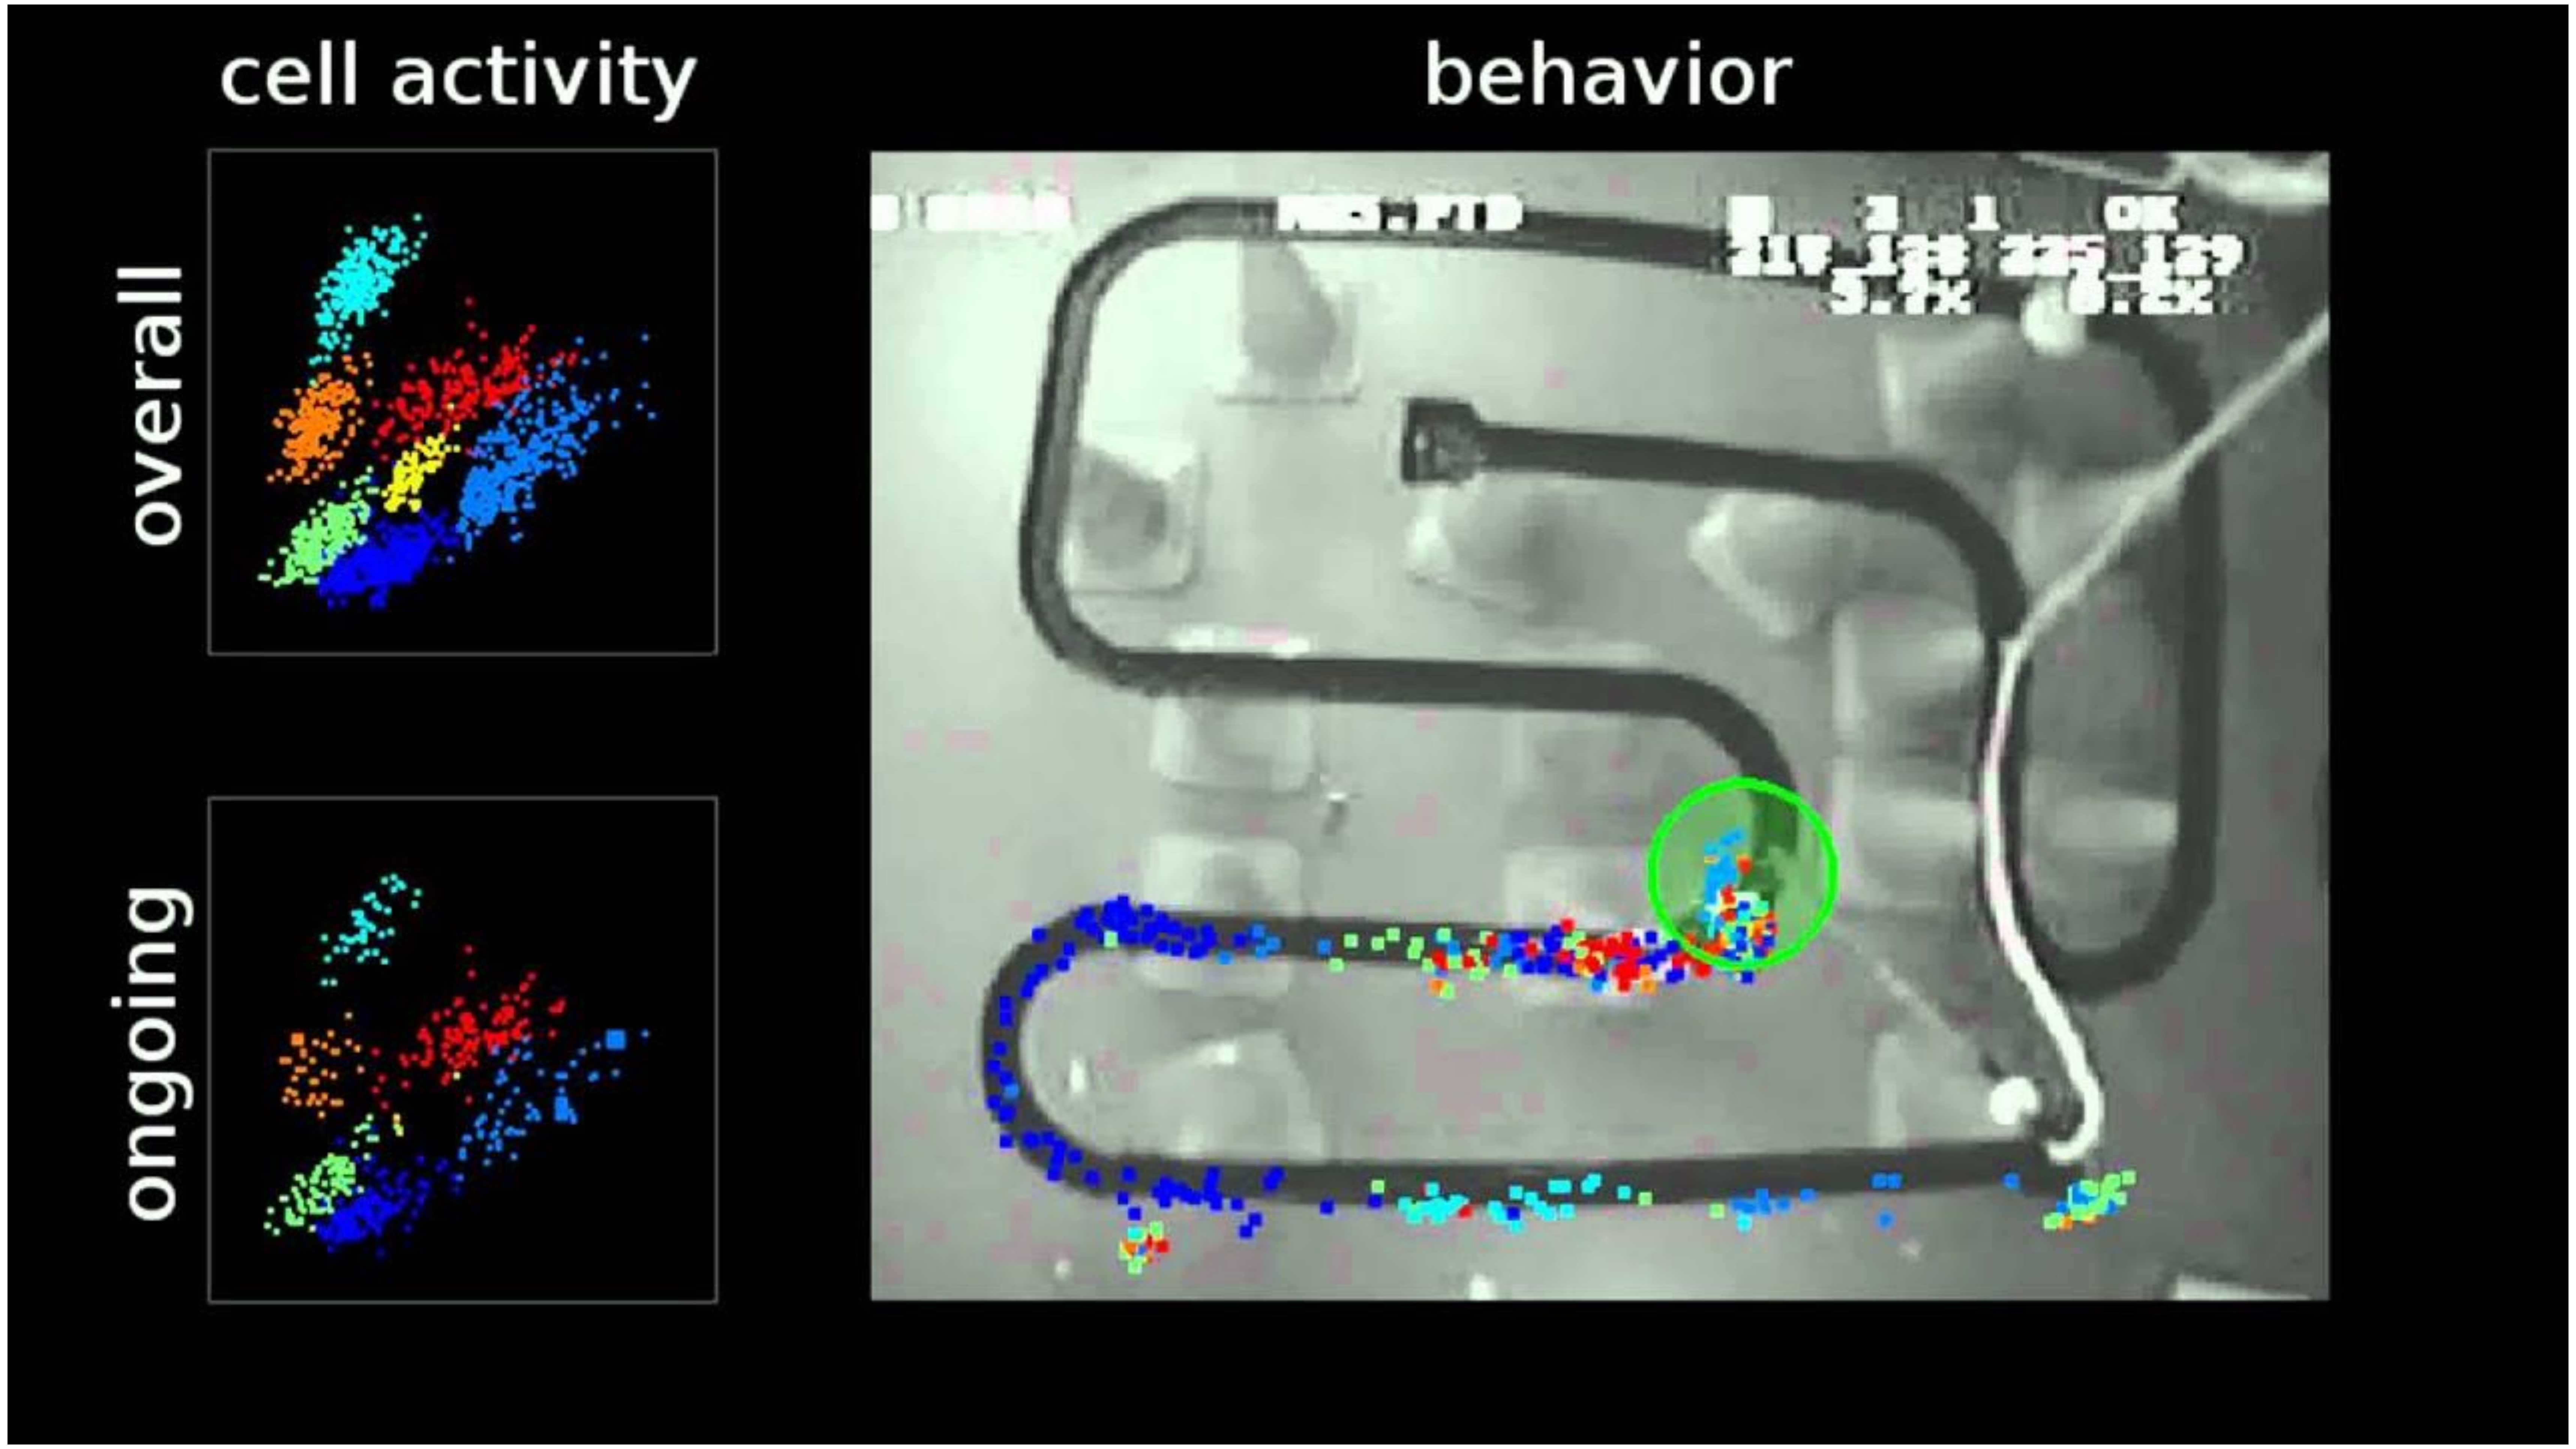

John O'Keefe
Nobel Prize in Physiology or Medicine 2014

Wilson Lab (MIT)

Place cells in the hippocampus represent location in an environment

(O'keefe & Nadel 1978)

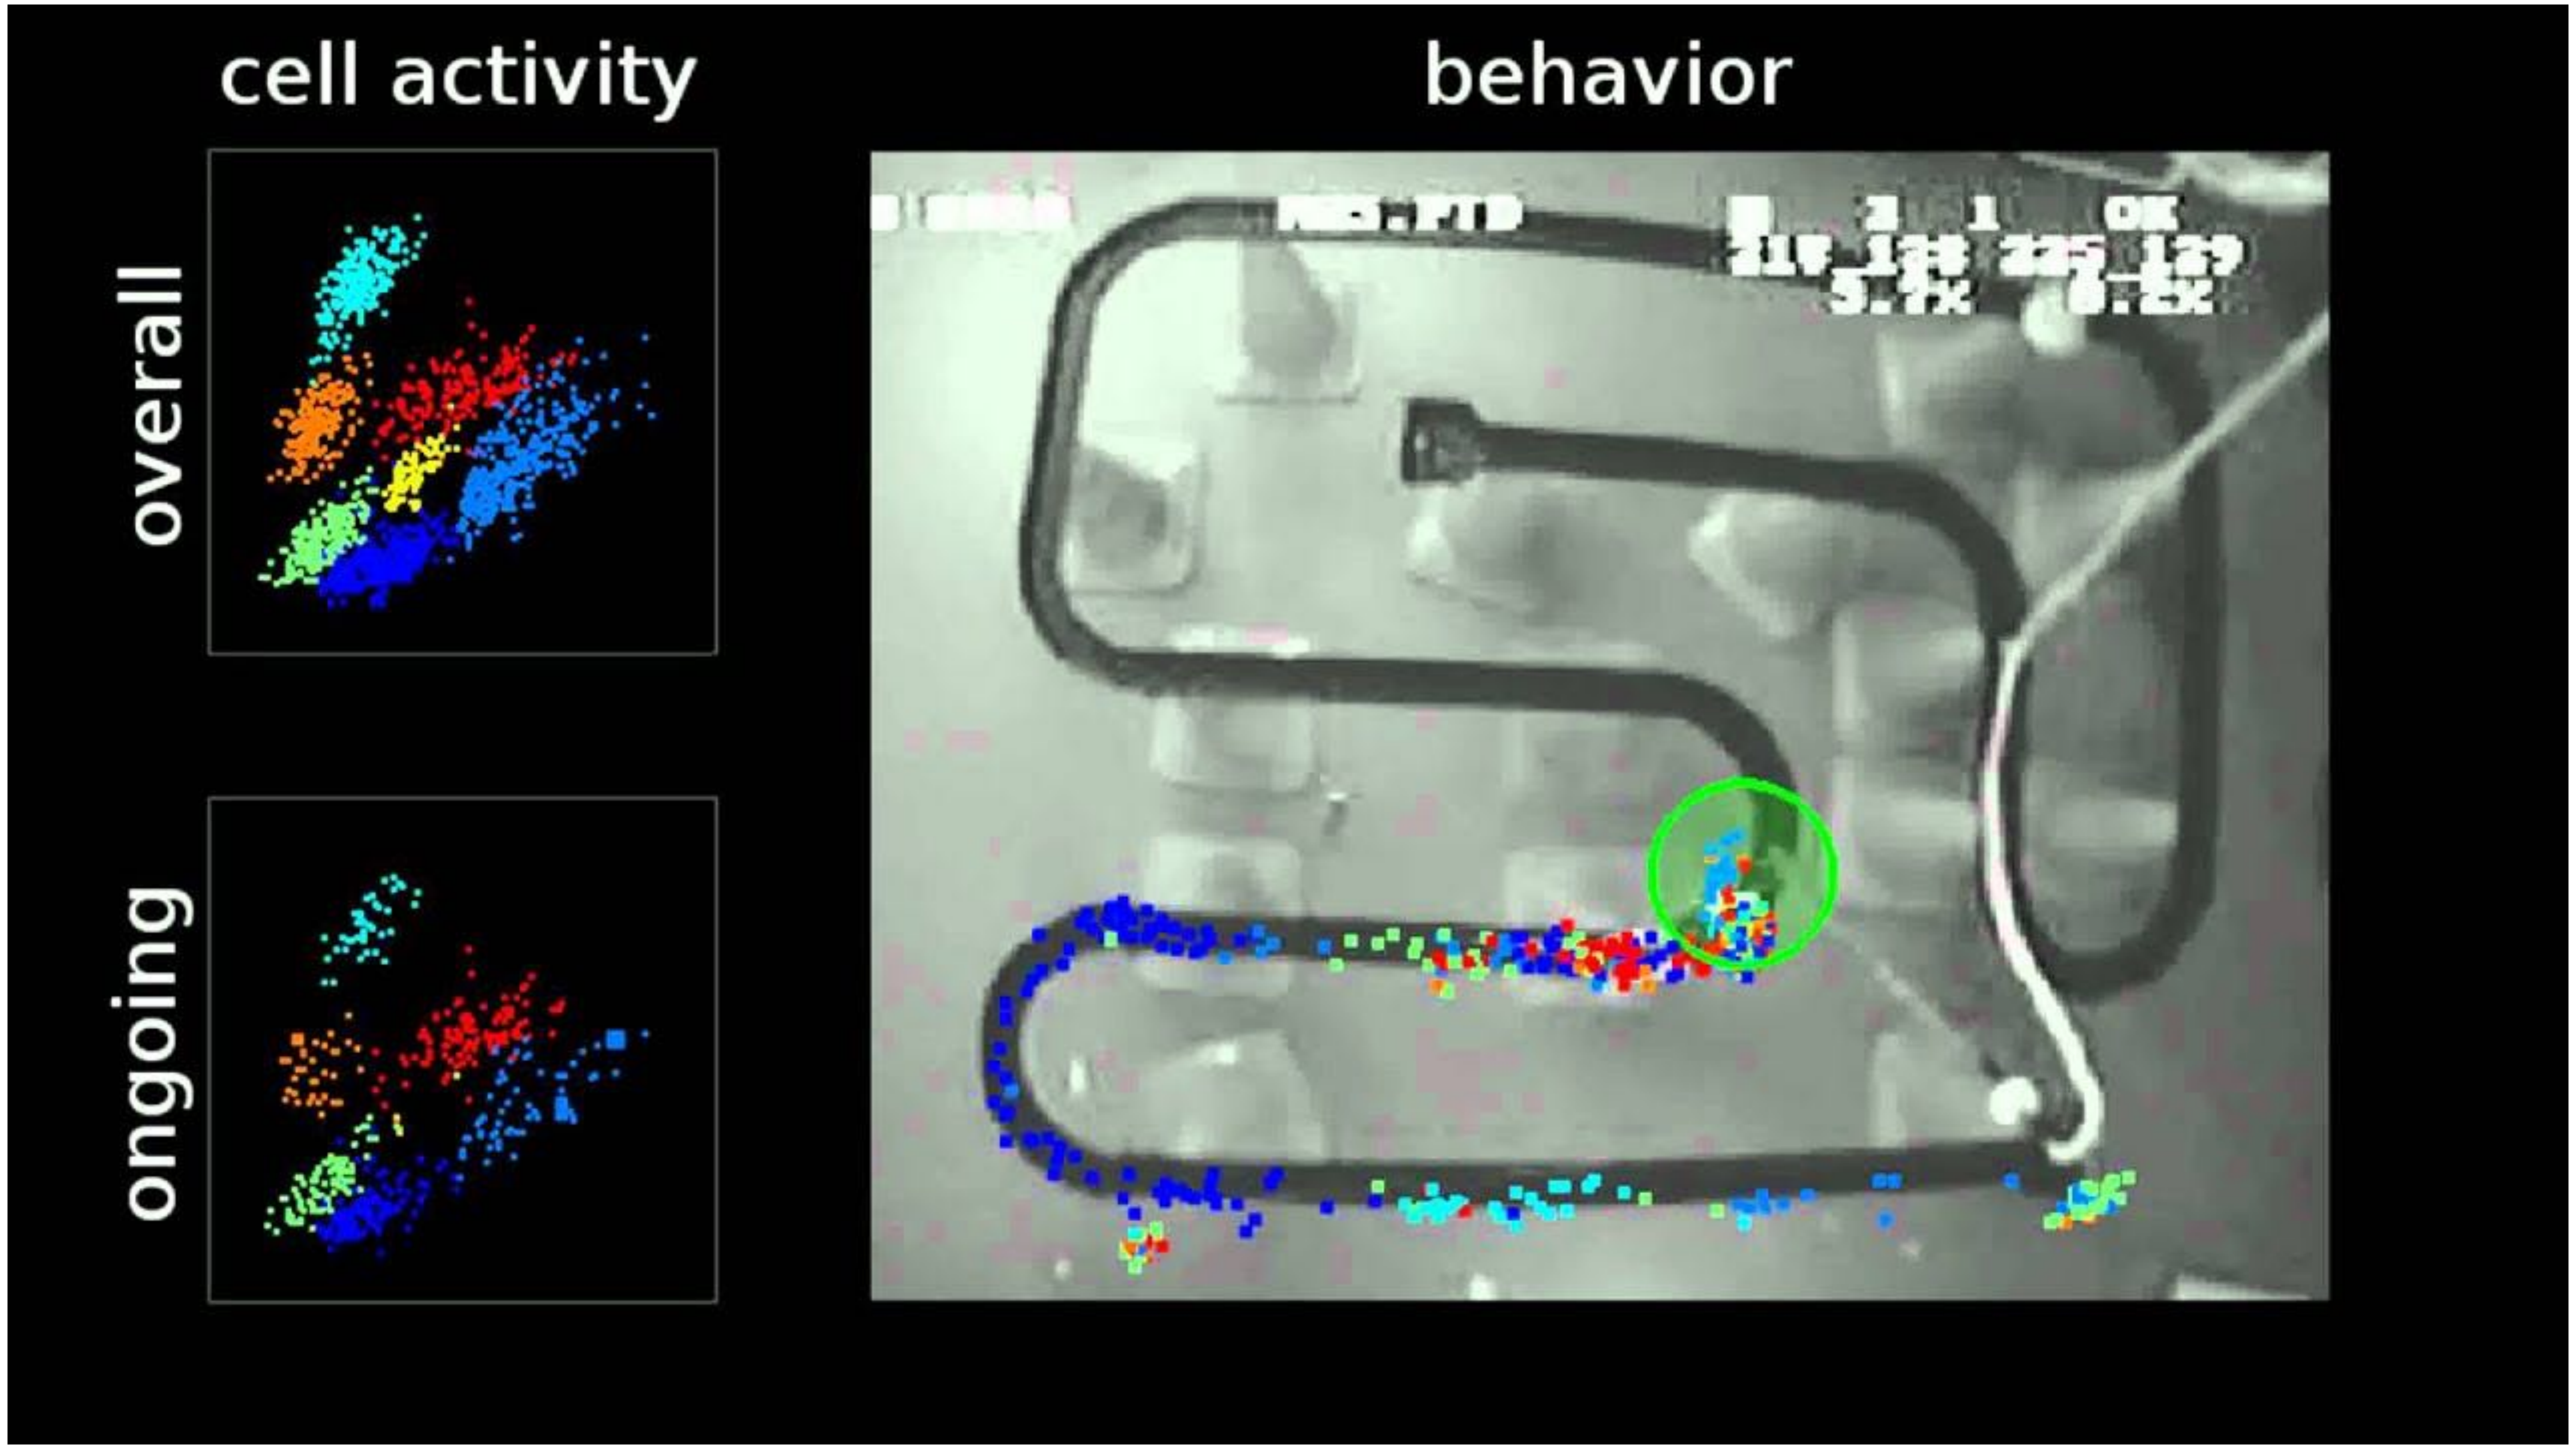

John O'Keefe
Nobel Prize in Physiology or Medicine 2014

Wilson Lab (MIT)

Grid cells in the Entorhinal Cortex provide a coordinate system

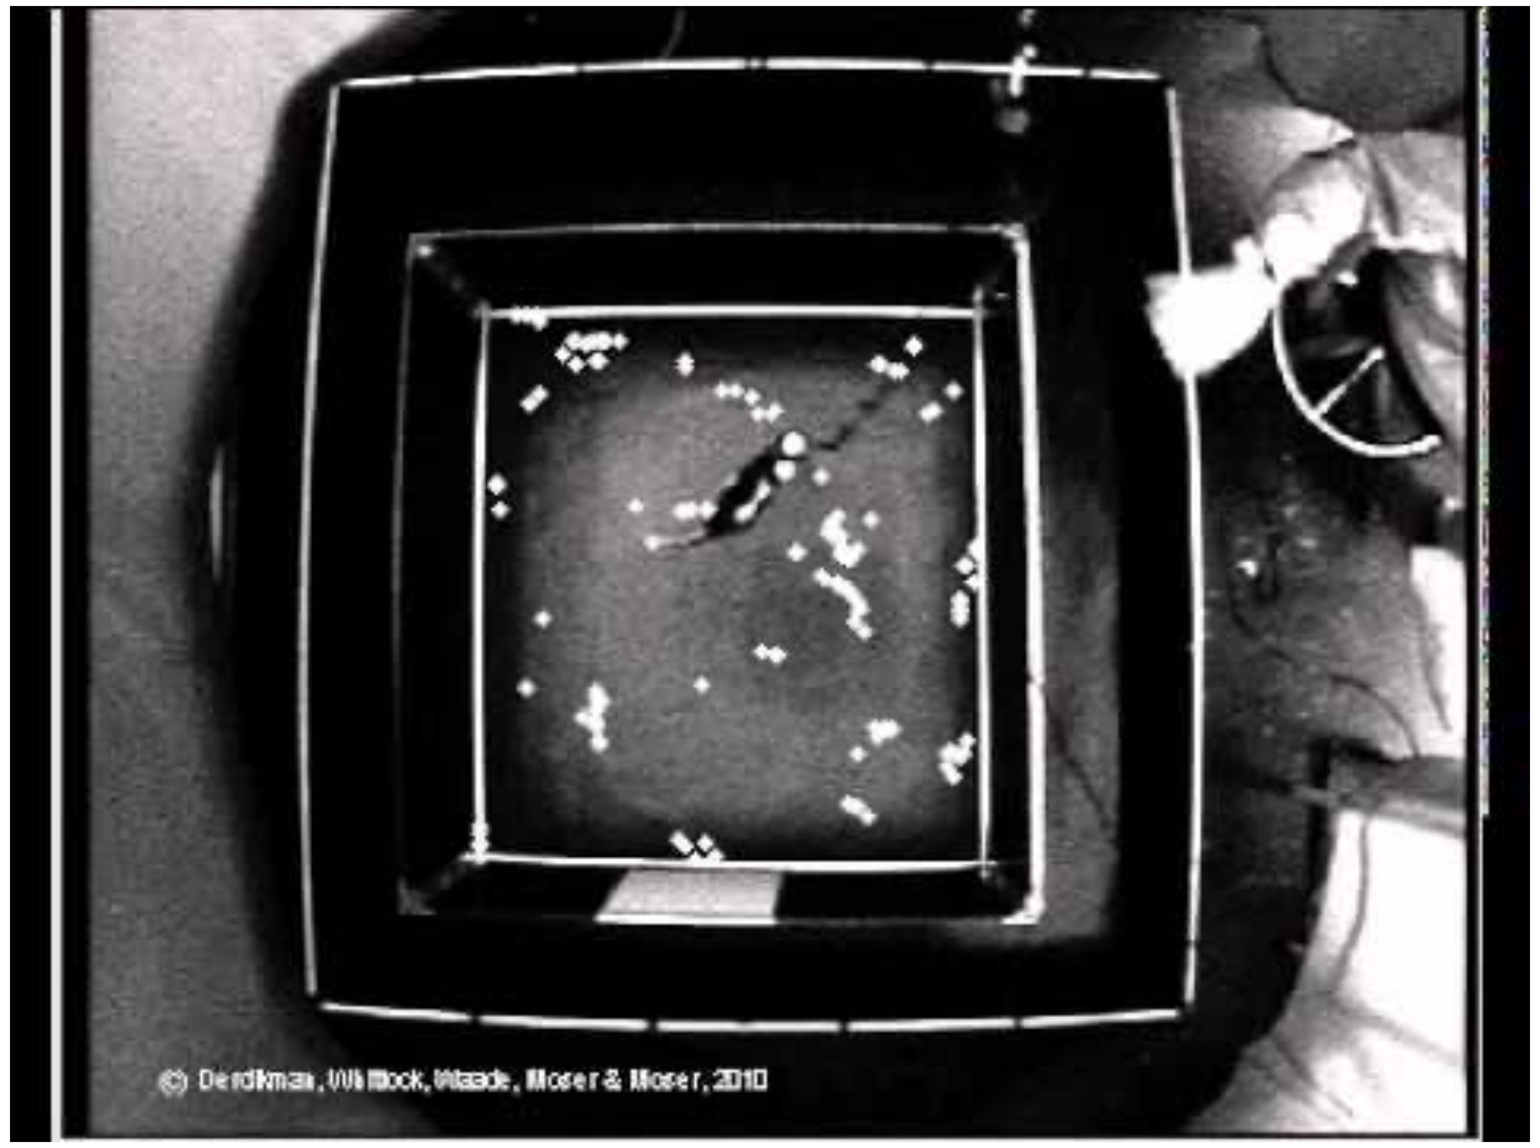

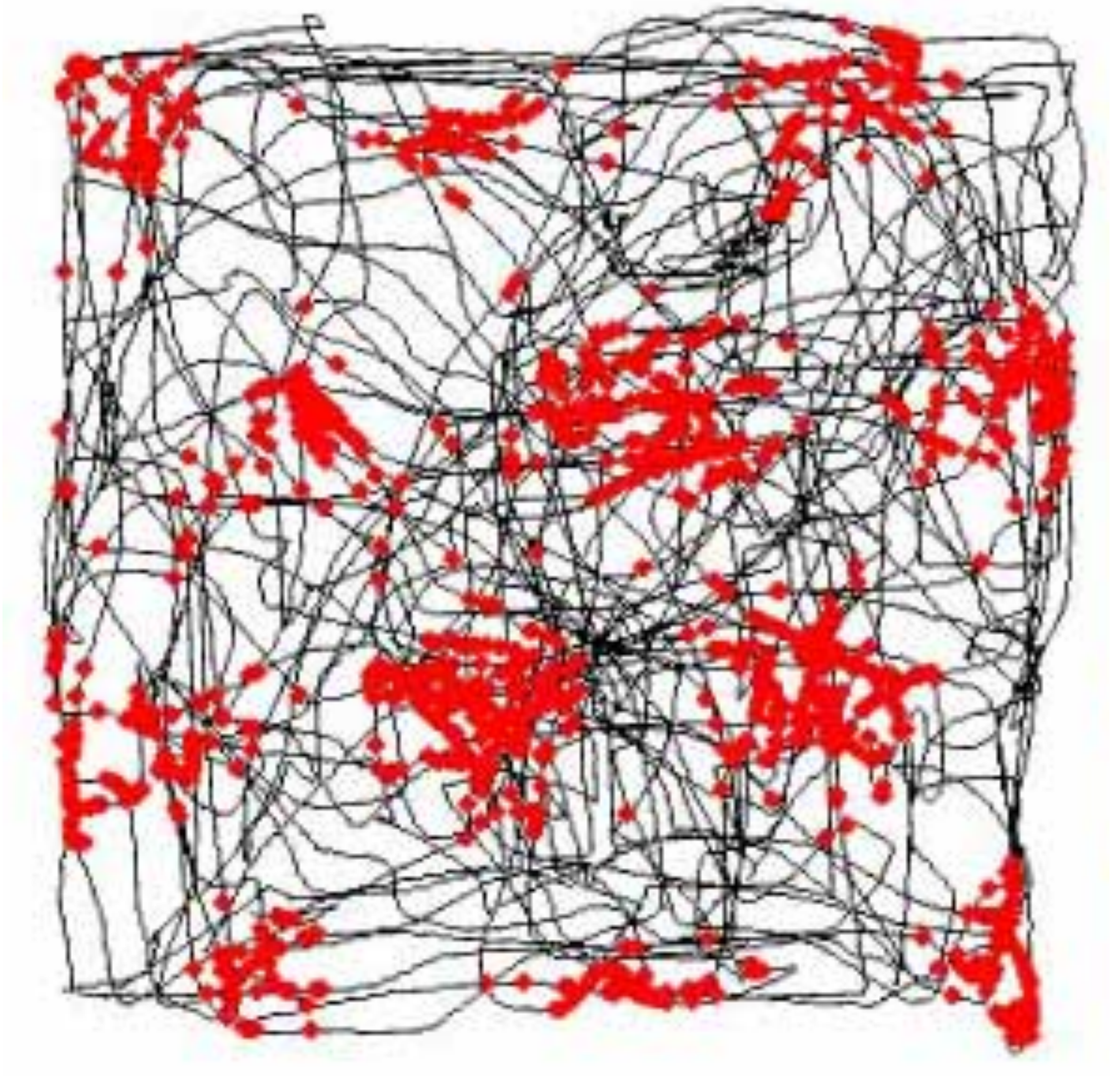

Trajectory

Peaks

Edvard and Maj-Britt Moser
Nobel Prize in Physiology or Medicine 2014

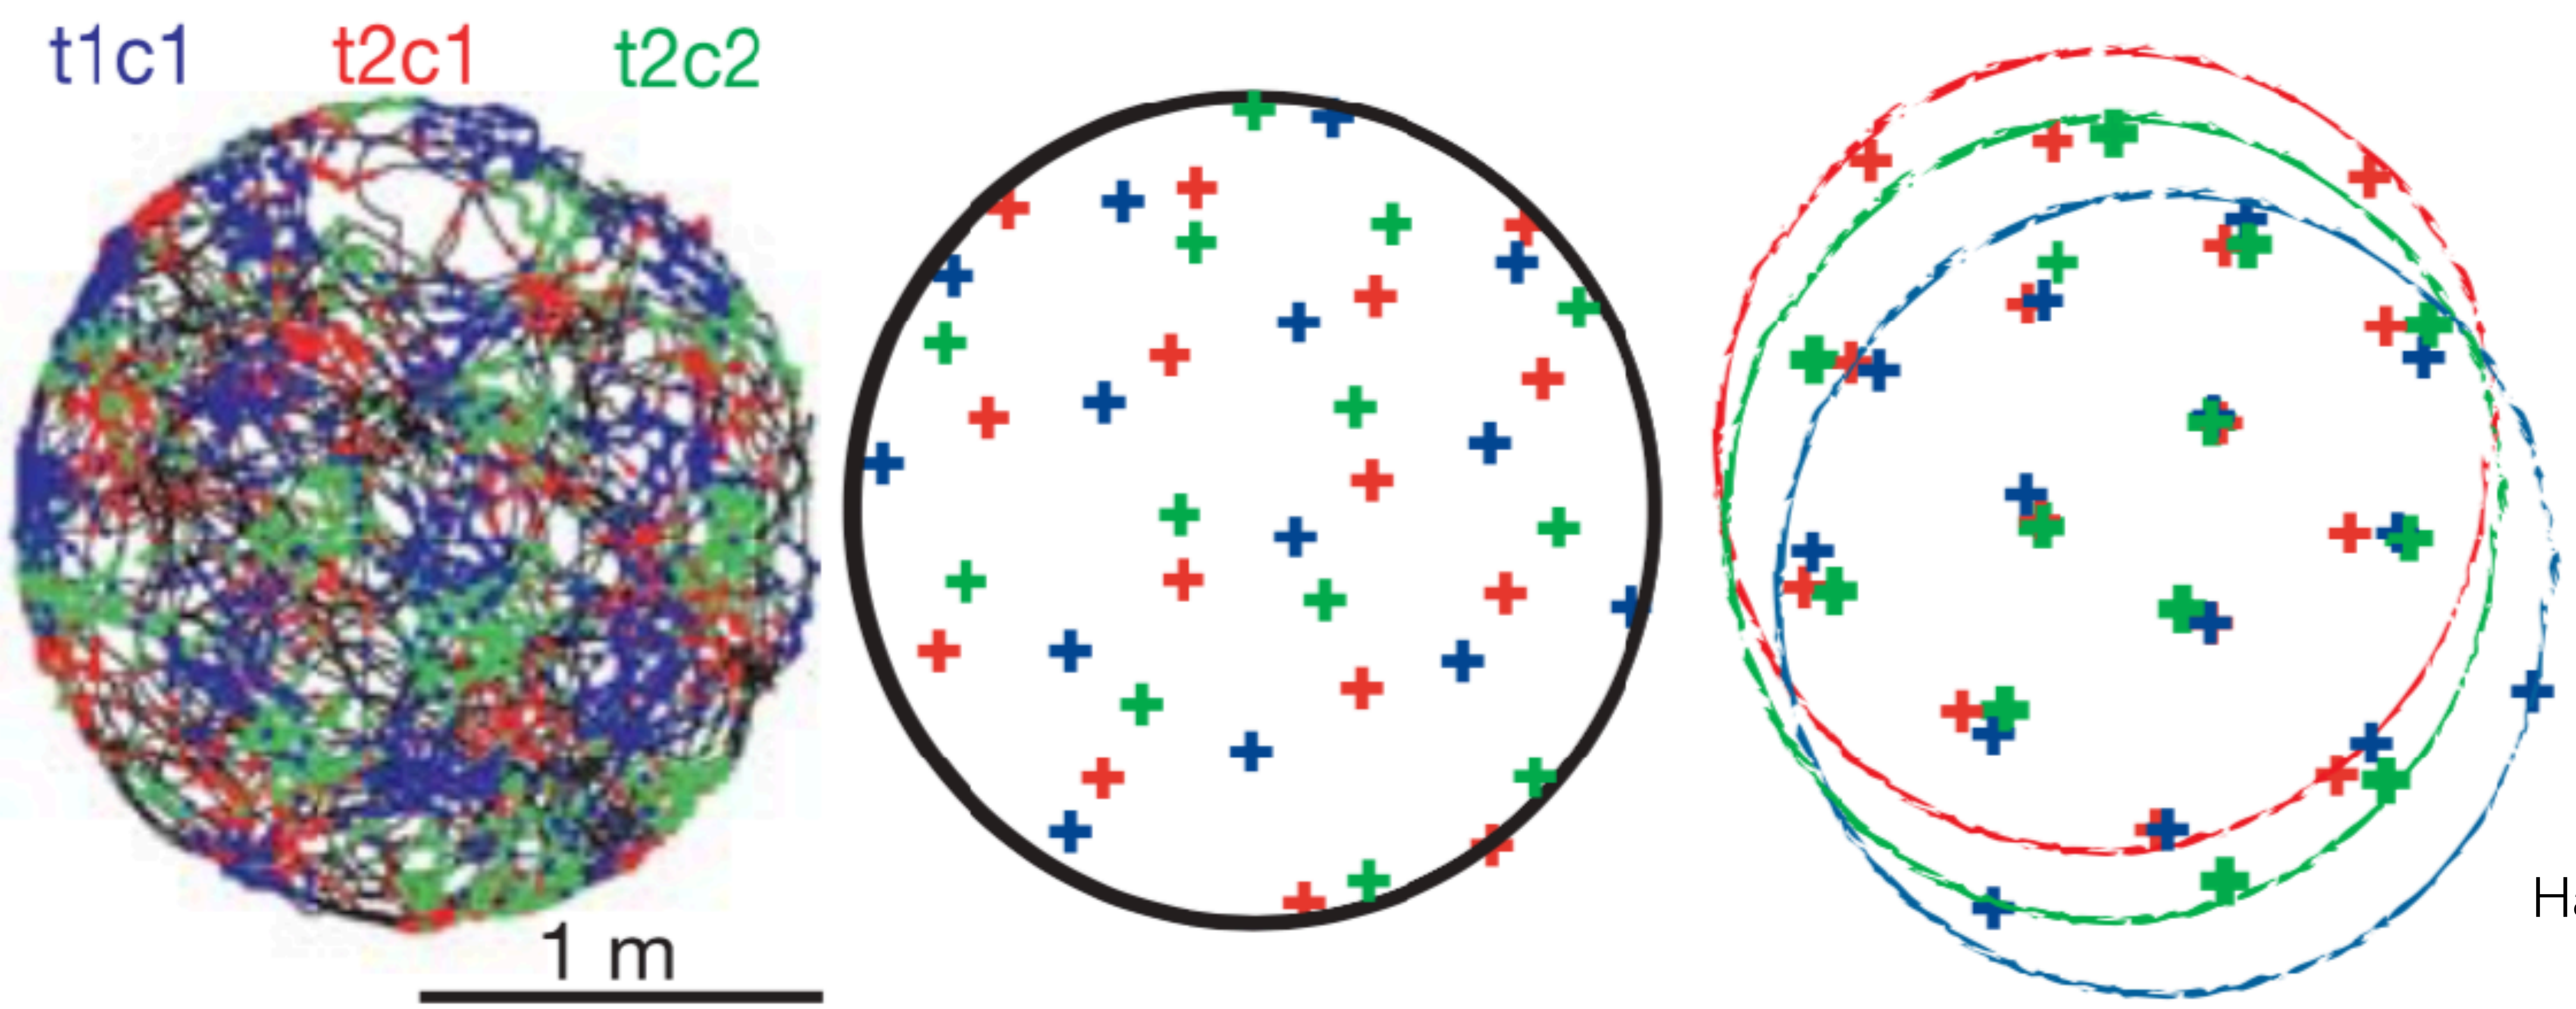

+ Peak

Hafting *et al* (Nature, 2005)

Grid cells in the Entorhinal Cortex provide a coordinate system

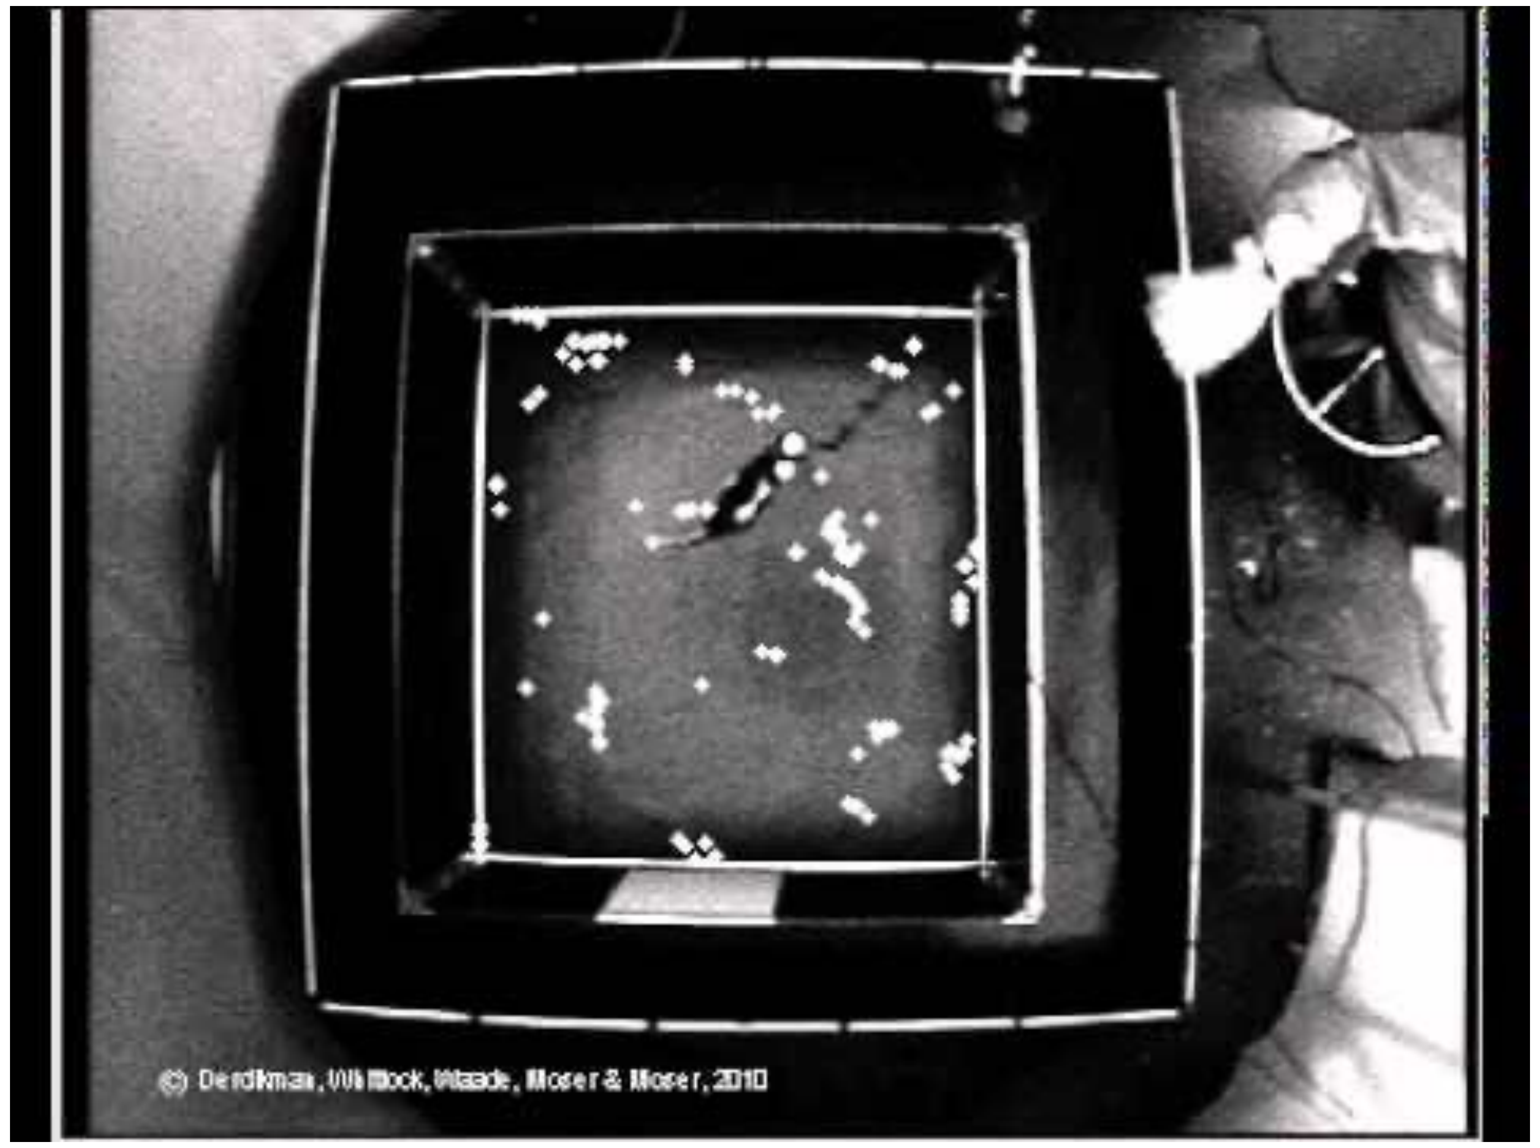

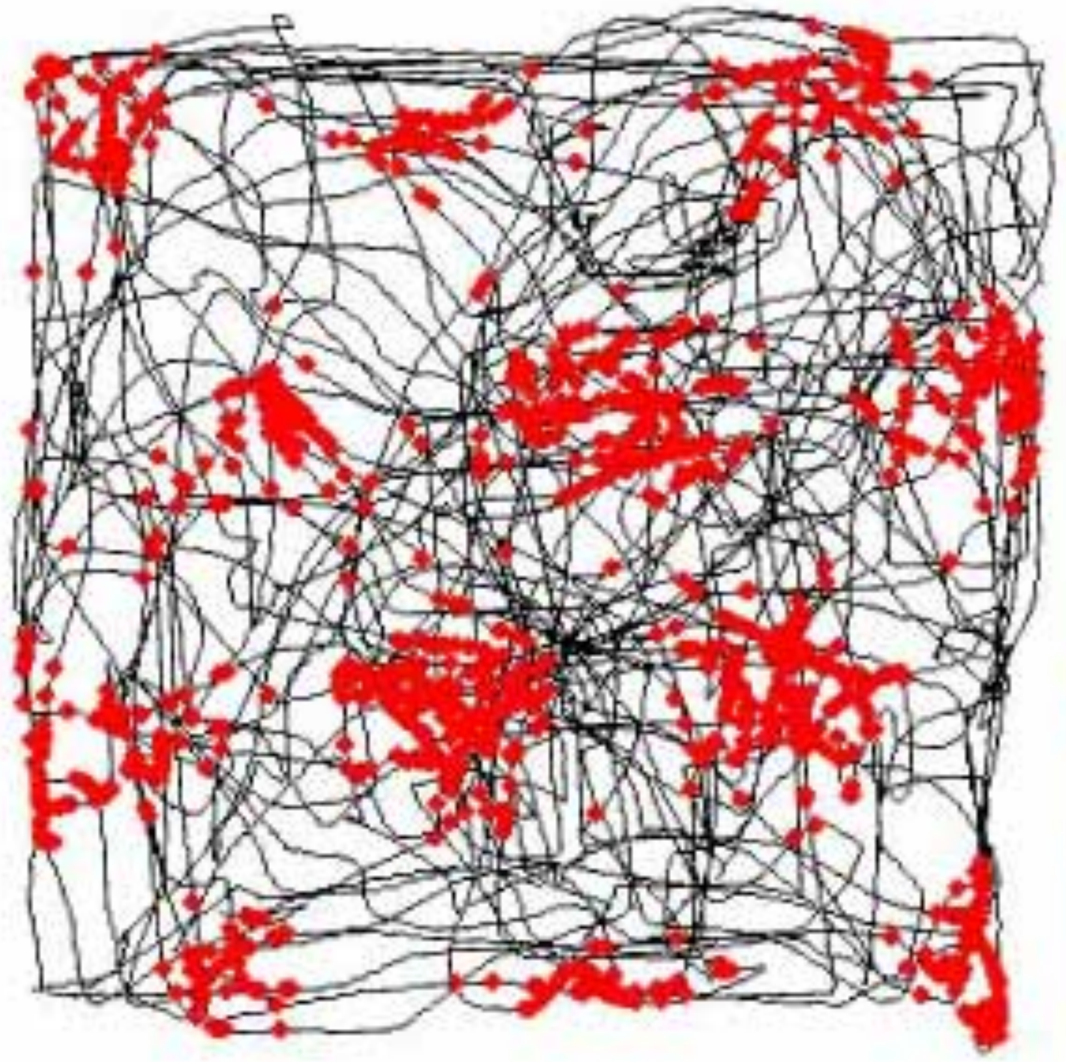

Trajectory
Peaks

Edvard and Maj-Britt Moser
Nobel Prize in Physiology or Medicine 2014

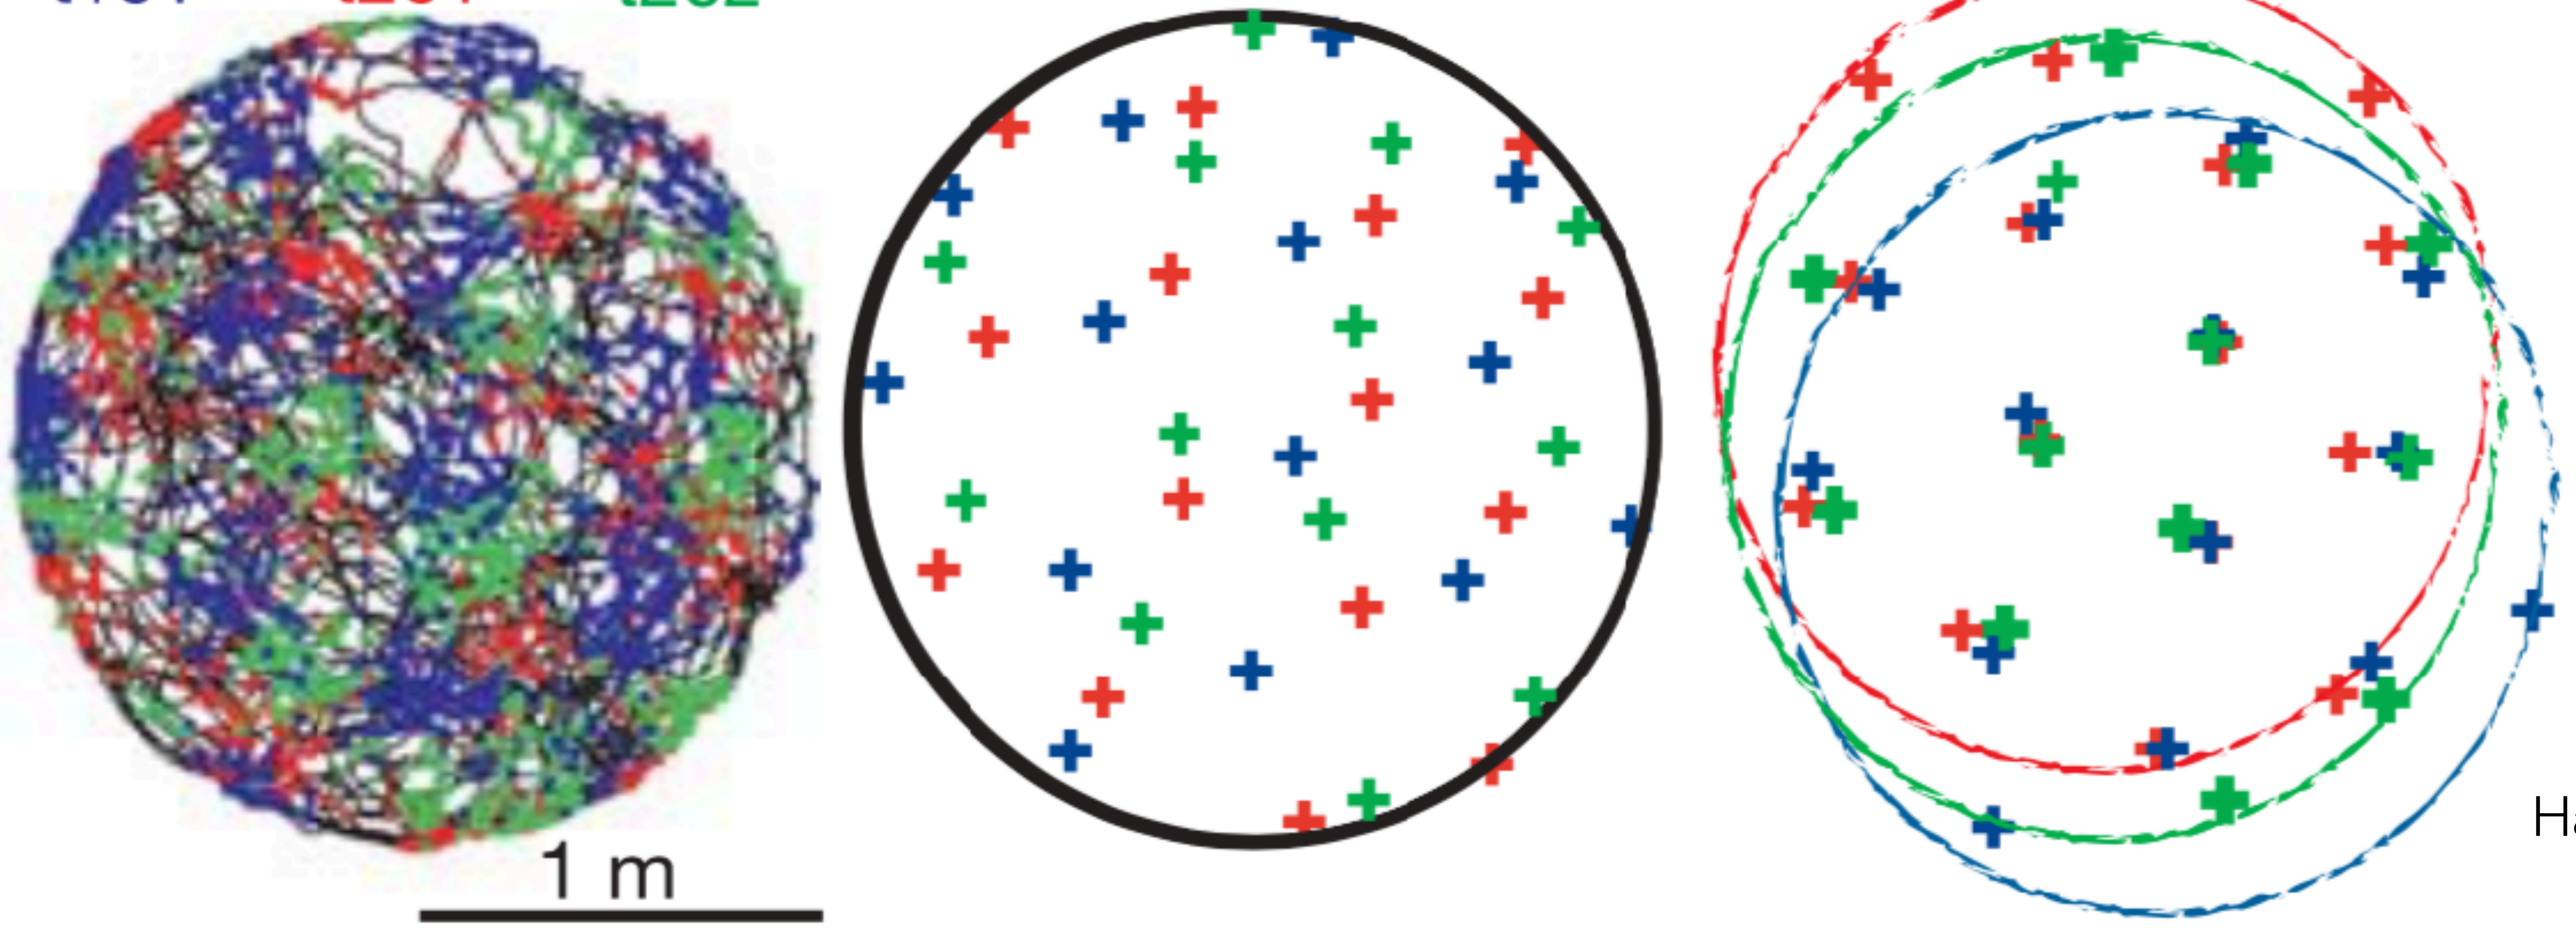

Hafting *et al* (Nature, 2005)

Grid cells in the Entorhinal Cortex provide a coordinate system

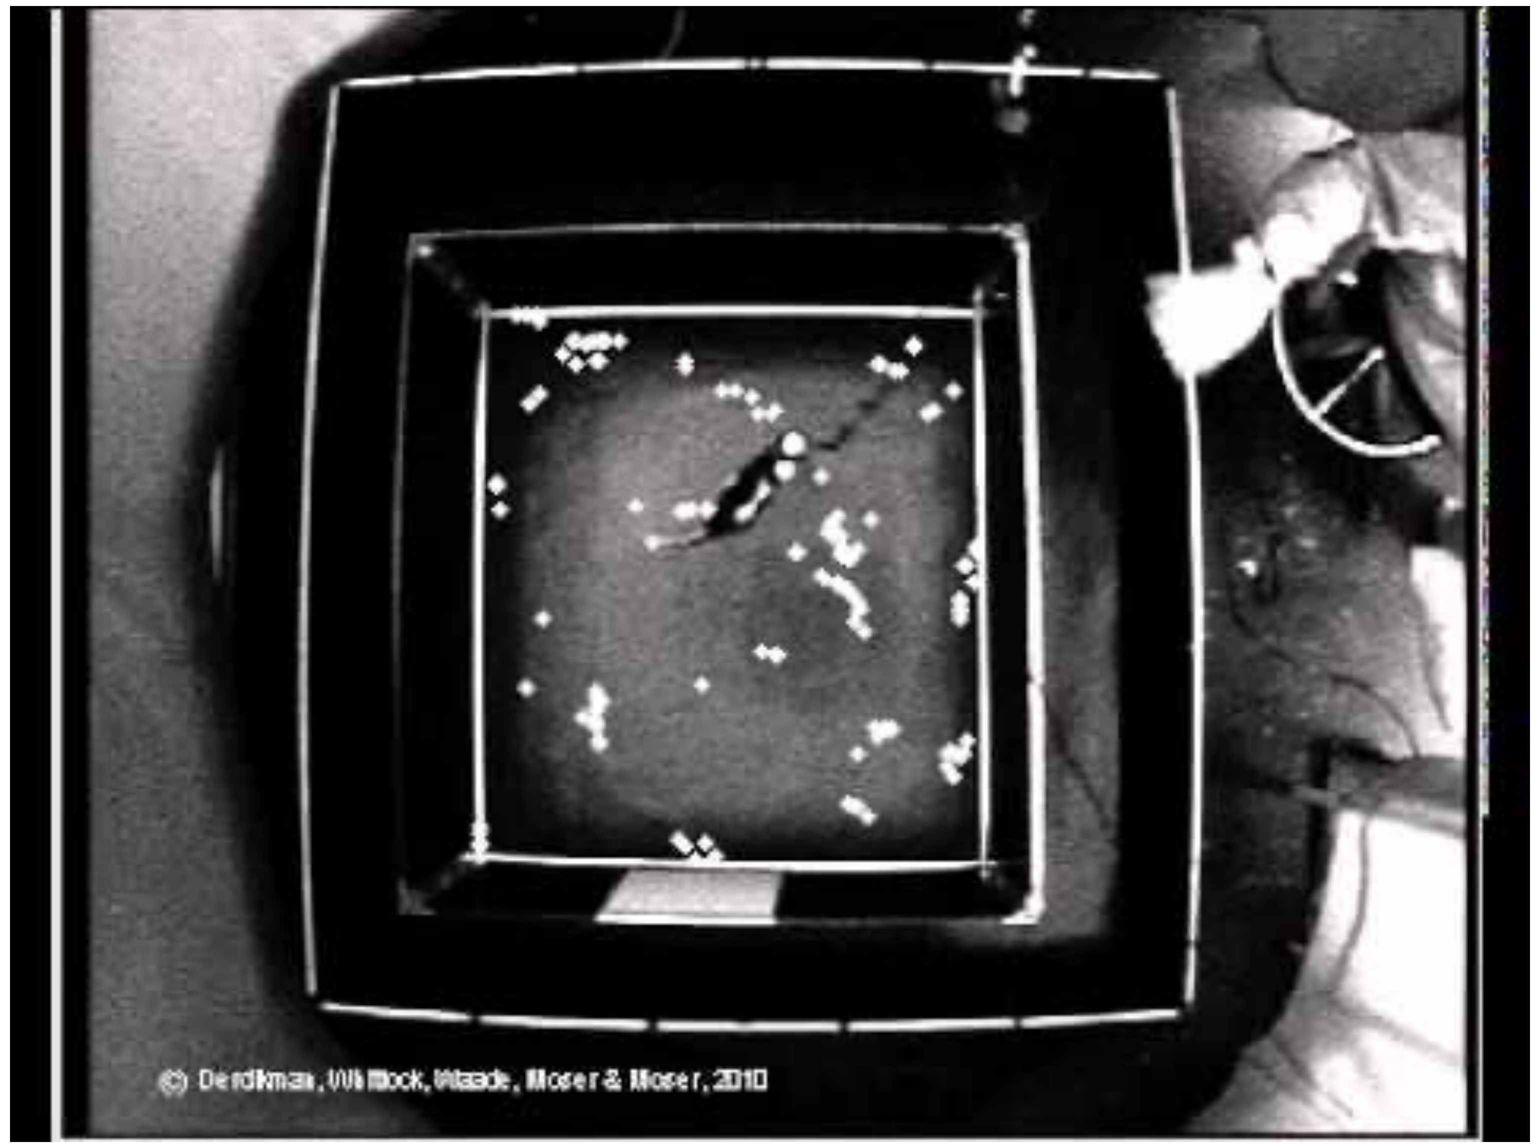

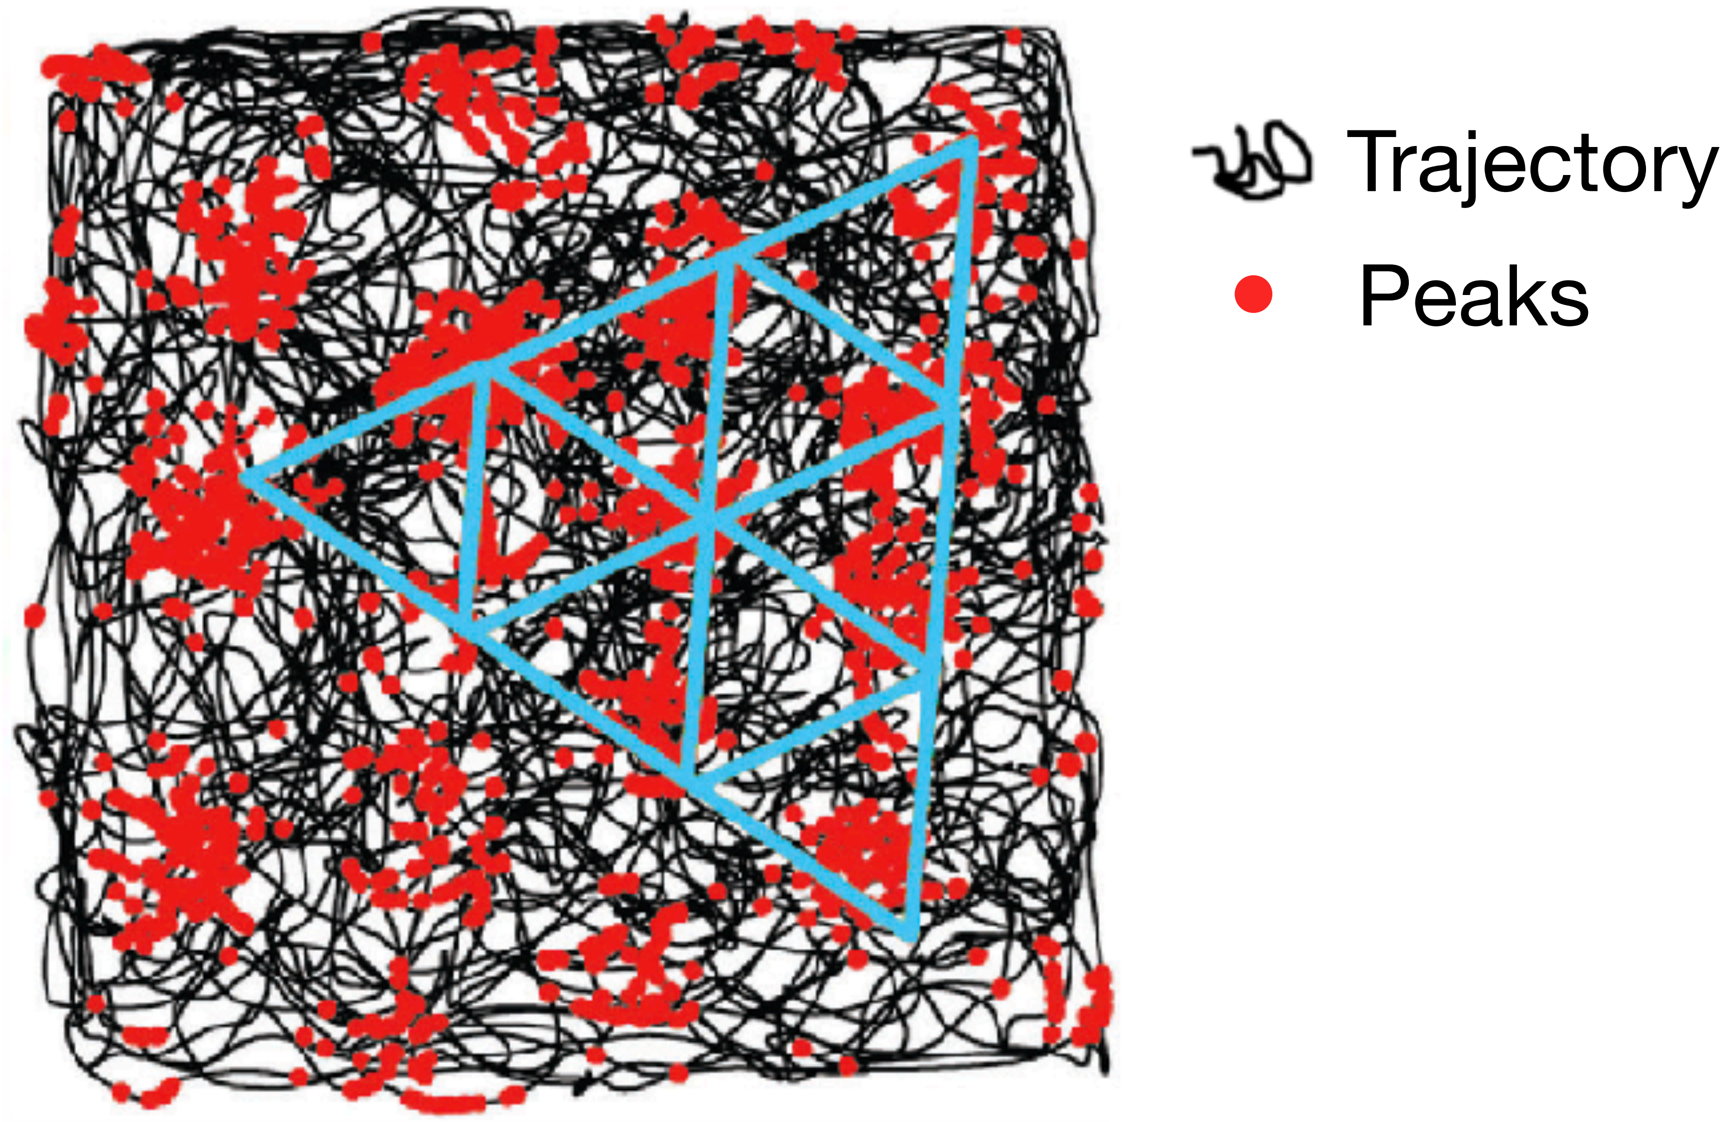

Edvard and Maj-Britt Moser
Nobel Prize in Physiology or Medicine 2014

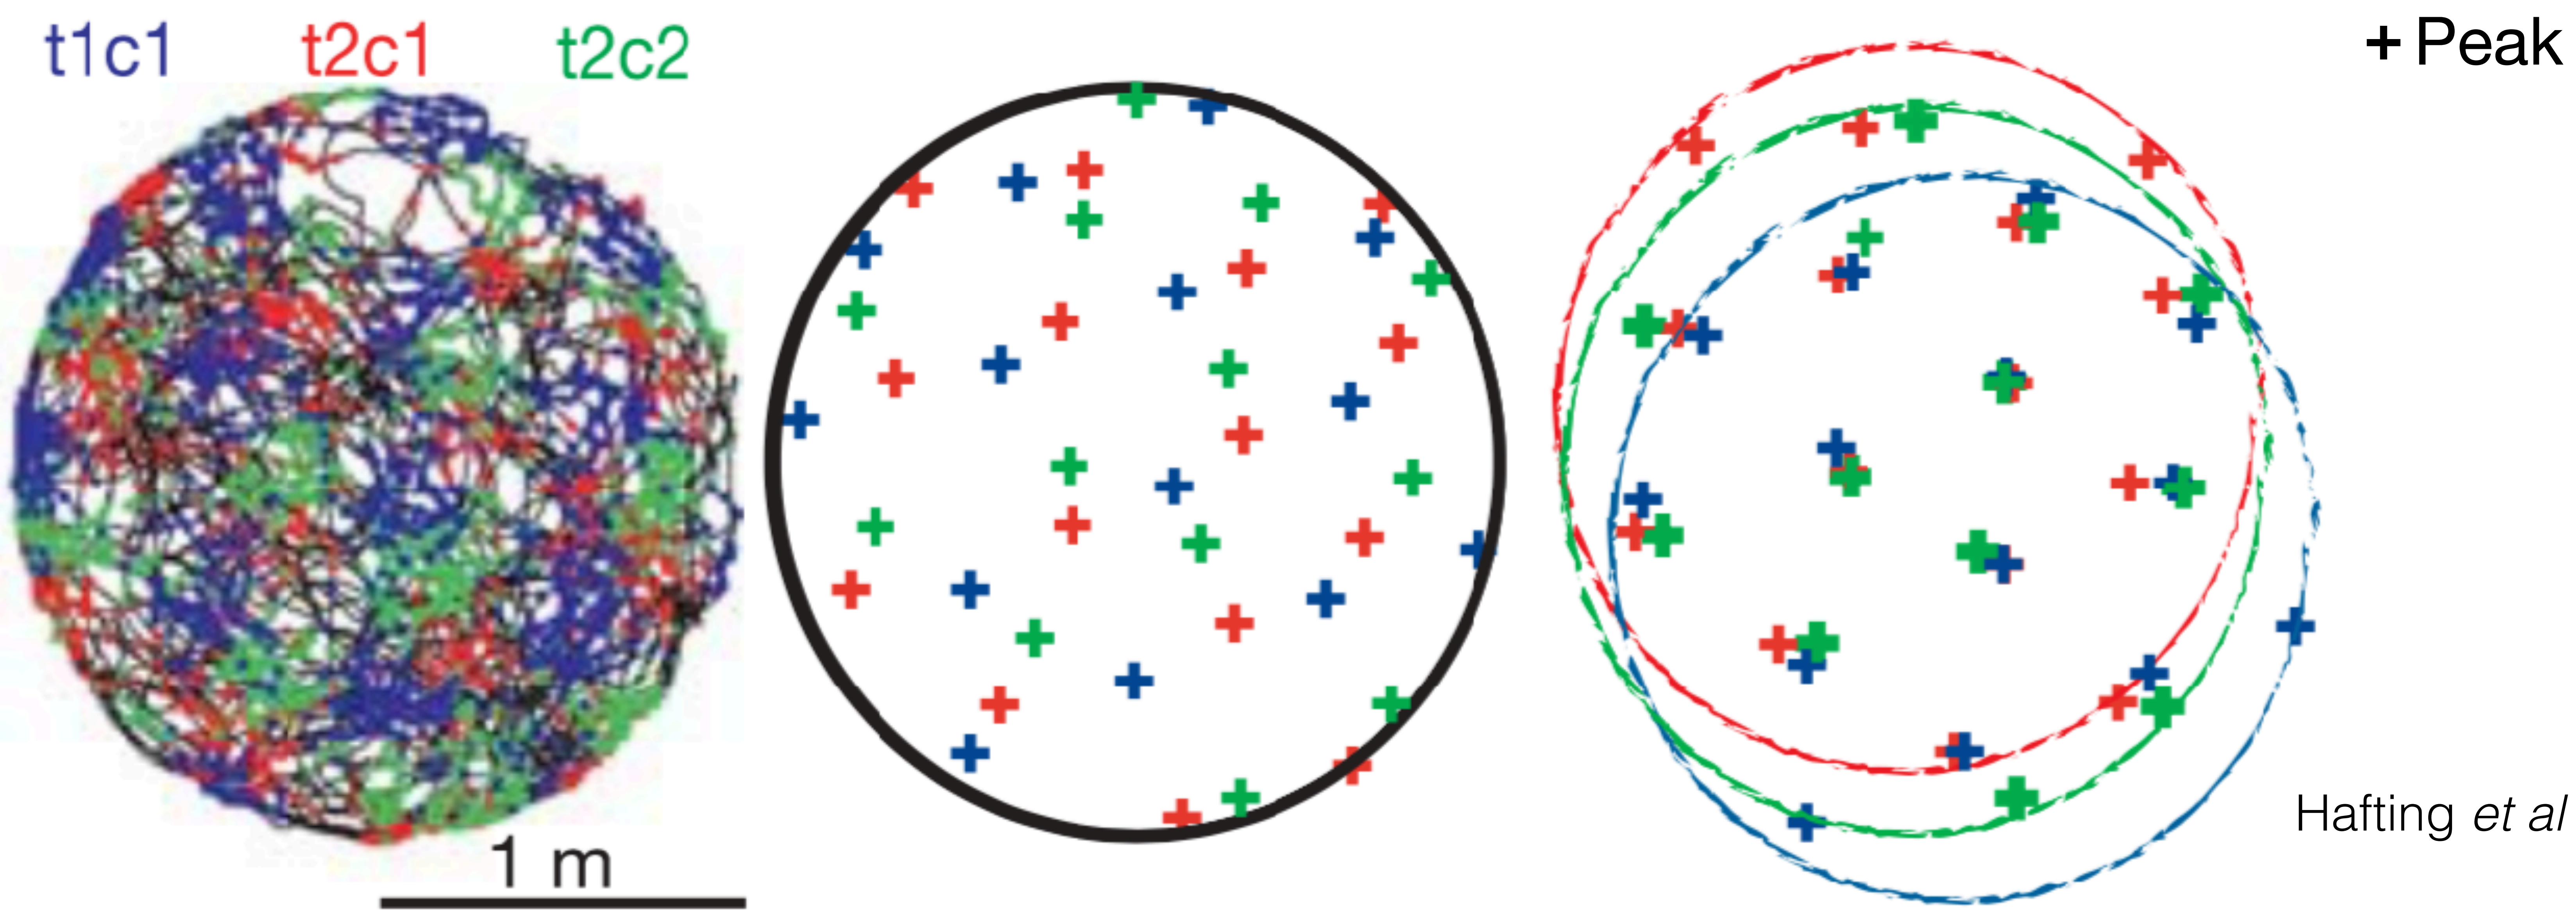

Hafting *et al* (Nature, 2005)

"Hippocampal Zoo"

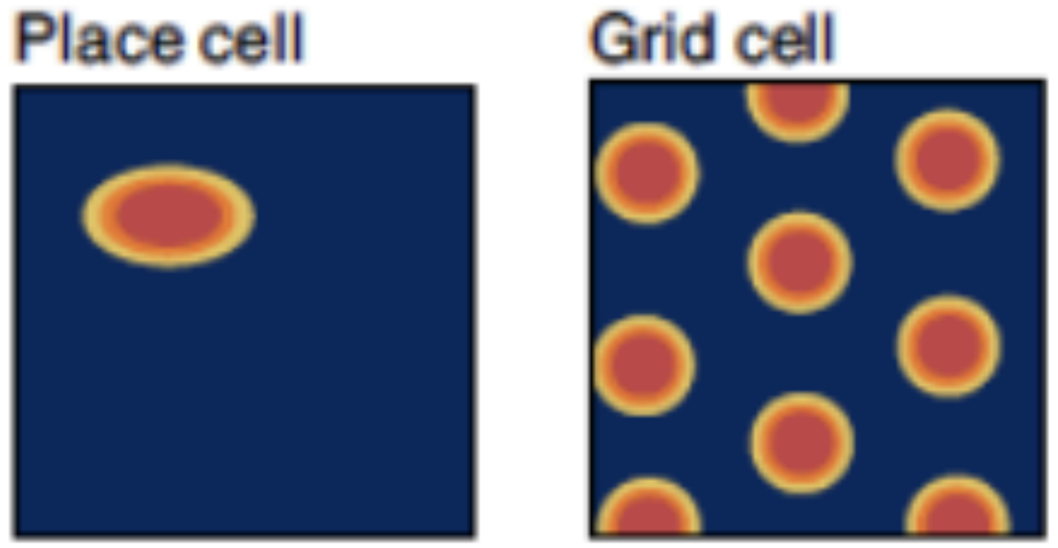

Whittington et al,. (2022)

Behrens et al., (2018)

“Hippocampal Zoo”

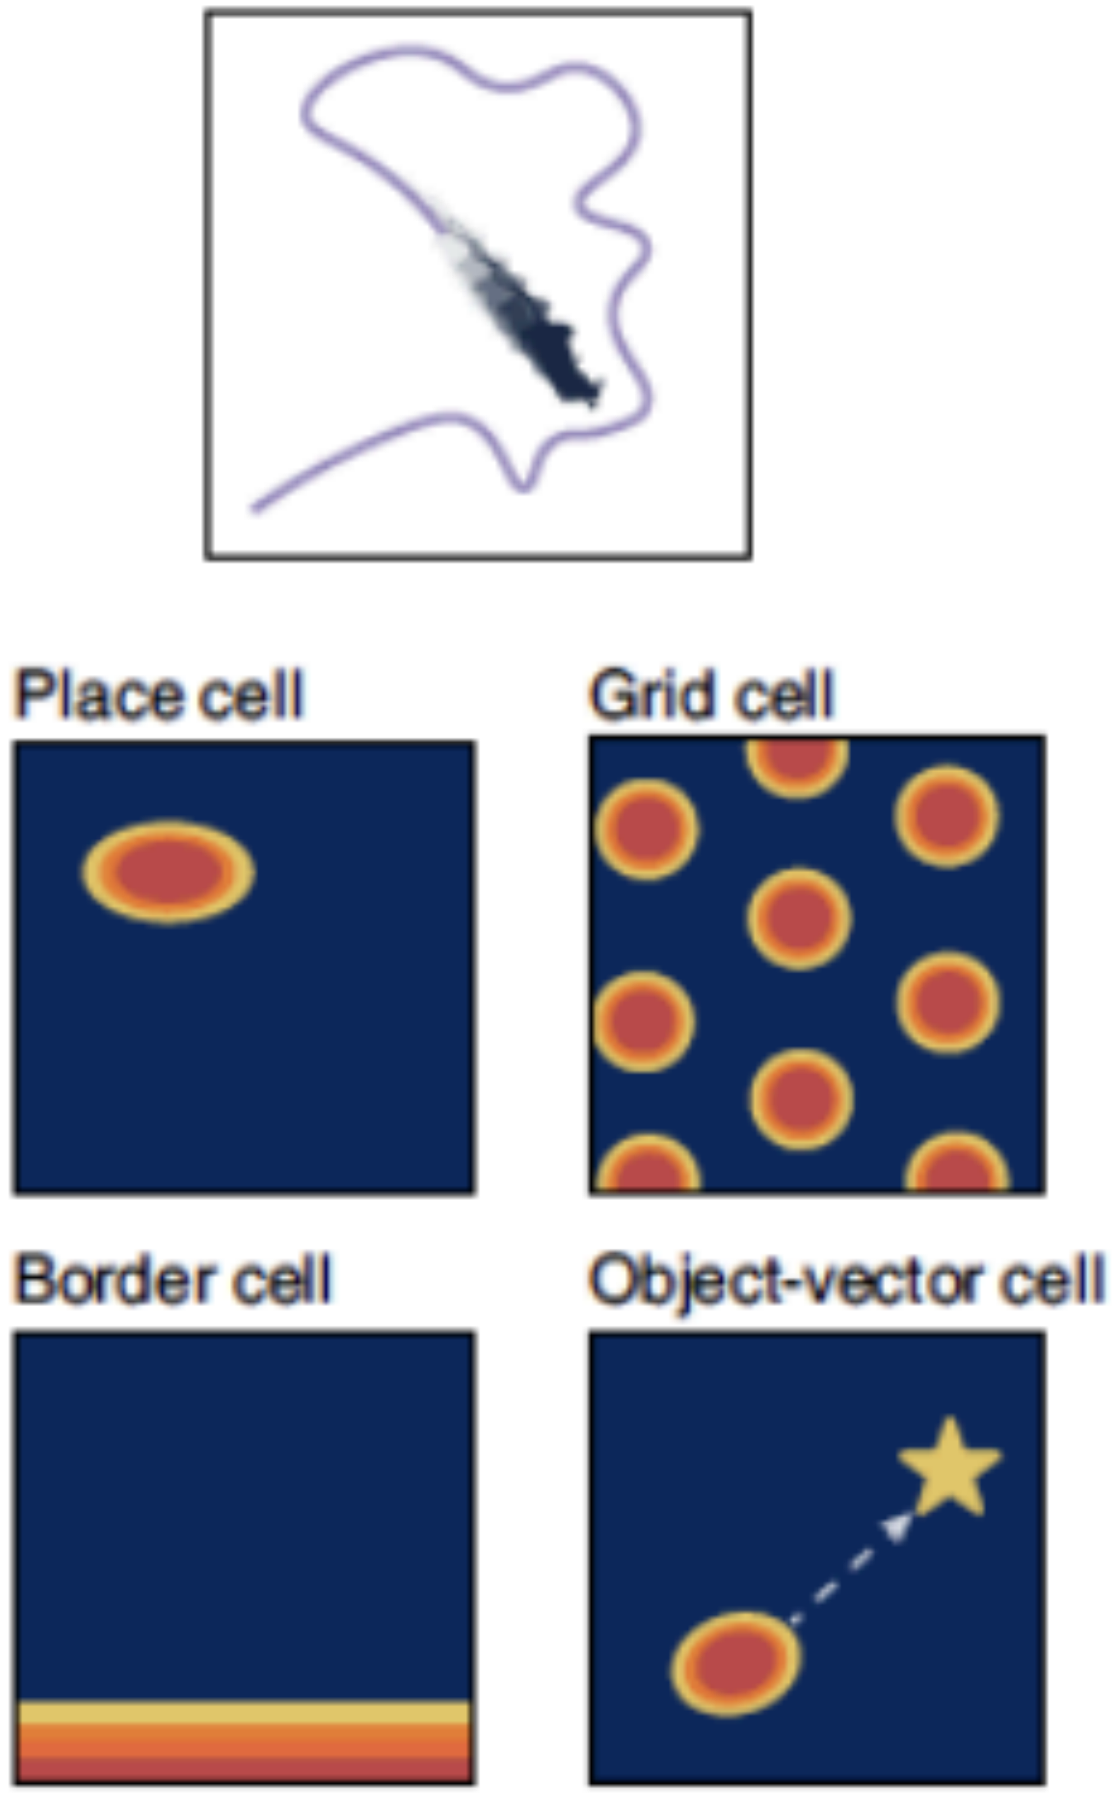

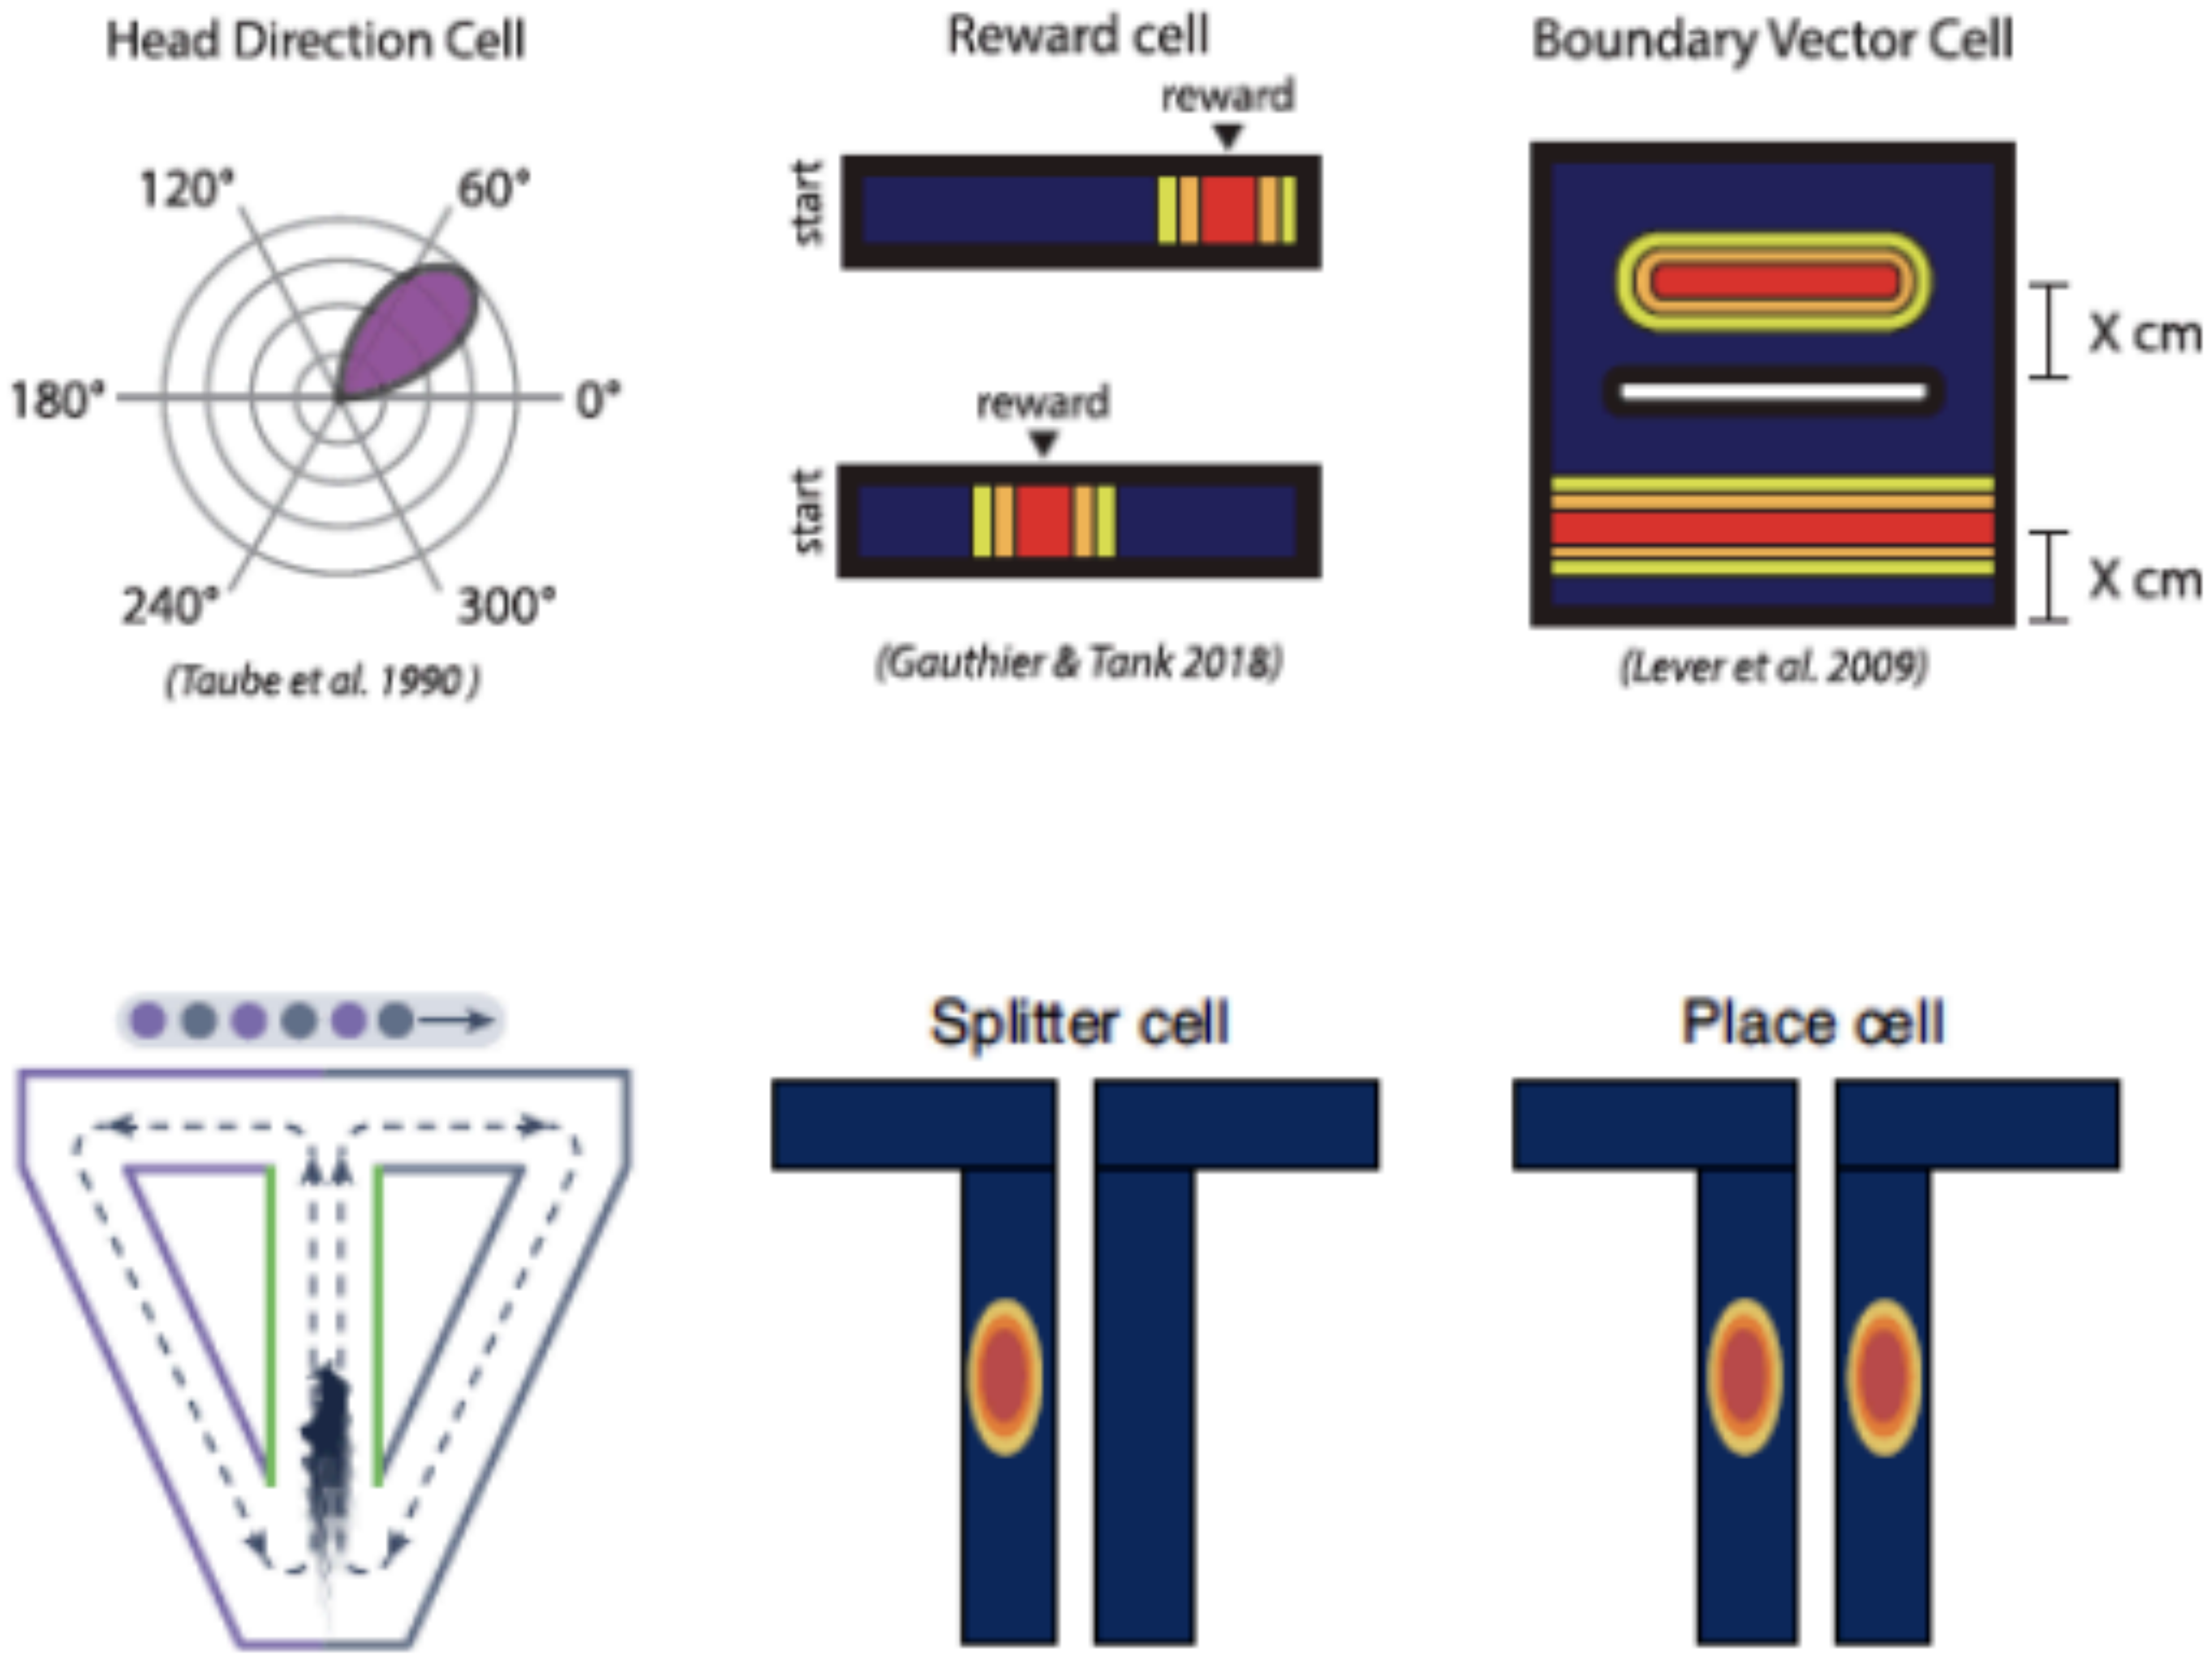

Whittington et al,. (2022)

Behrens et al., (2018)

Tools for navigation

- The **Hippocampus** represents spatial distance between landmarks (Morgan et al., 2011) nand between events (Nielson et al., 2015)

- The **Entorhinal Cortex (EC)** encodes the direction of travel (Doeller et al., 2015)
    - Participants moved in a VR environment
    - When direction aligned with one of the 3 axes of their grid cells, we observe stronger BOLD activation in the EC
    - These angles are remarkably robust, and are preserved (in the same environment) when participants return to the scanner days or weeks later

Tools for navigation

- The **Hippocampus** represents spatial distance between landmarks (Morgan et al., 2011) nand between events (Nielson et al., 2015)

- The **Entorhinal Cortex (EC)** encodes the direction of travel (Doeller et al., 2015)

  - Participants moved in a VR environment

  - When direction aligned with one of the 3 axes of their grid cells, we observe stronger BOLD activation in the EC

  - These angles are remarkably robust, and are preserved (in the same environment) when participants return to the scanner days or weeks later

Morgan et al., (*JNeuro*, 2011)

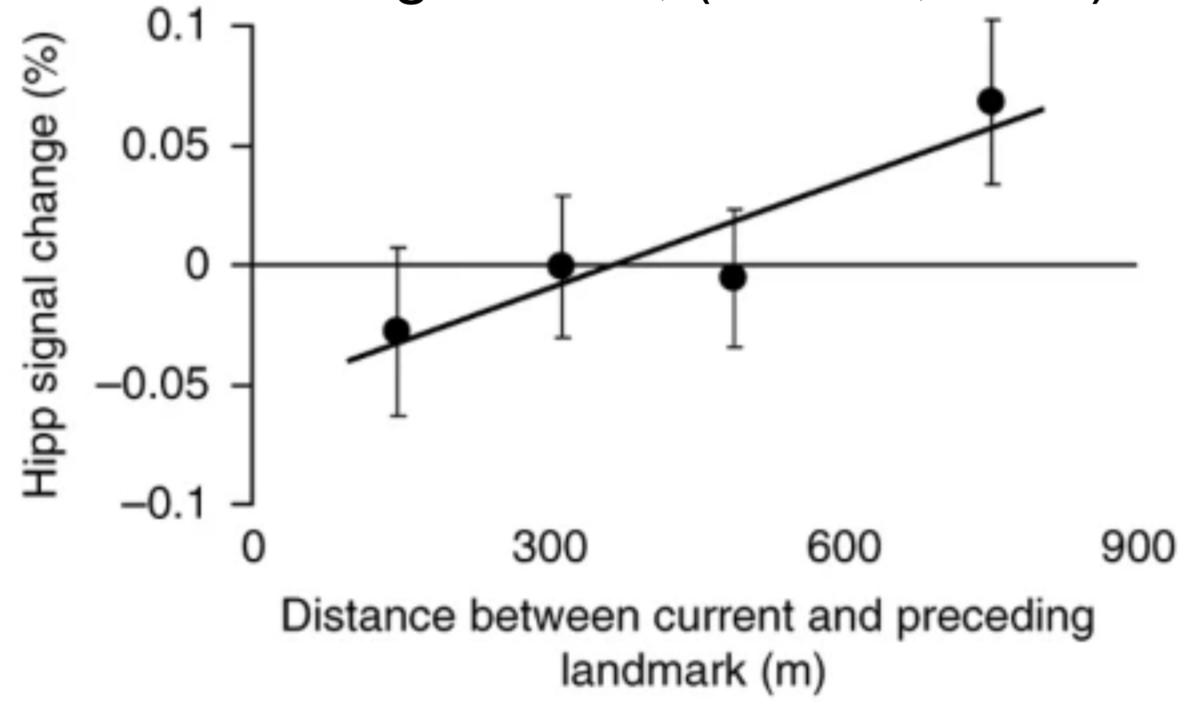

Tools for navigation

- The **Hippocampus** represents spatial distance between landmarks (Morgan et al., 2011) nand between events (Nielson et al., 2015)

- The **Entorhinal Cortex (EC)** encodes the direction of travel (Doeller et al., 2015)

  - Participants moved in a VR environment

  - When direction aligned with one of the 3 axes of their grid cells, we observe stronger BOLD activation in the EC

  - These angles are remarkably robust, and are preserved (in the same environment) when participants return to the scanner days or weeks later

Morgan et al., (*JNeuro*, 2011)

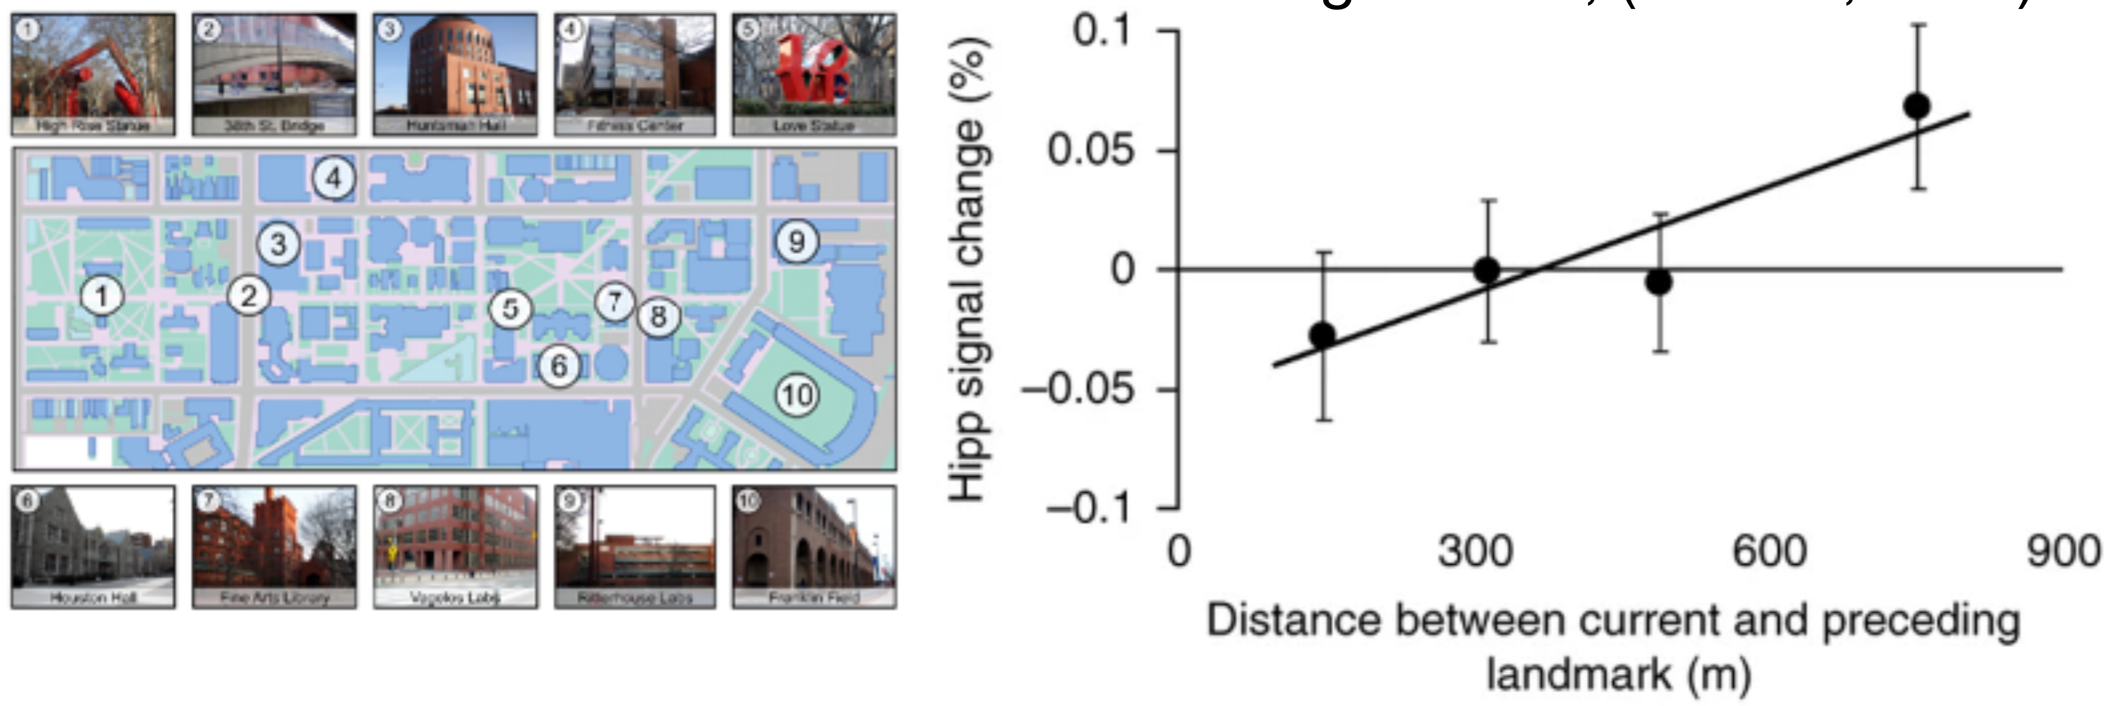

Nielson et al., (*PNAS*, 2015)

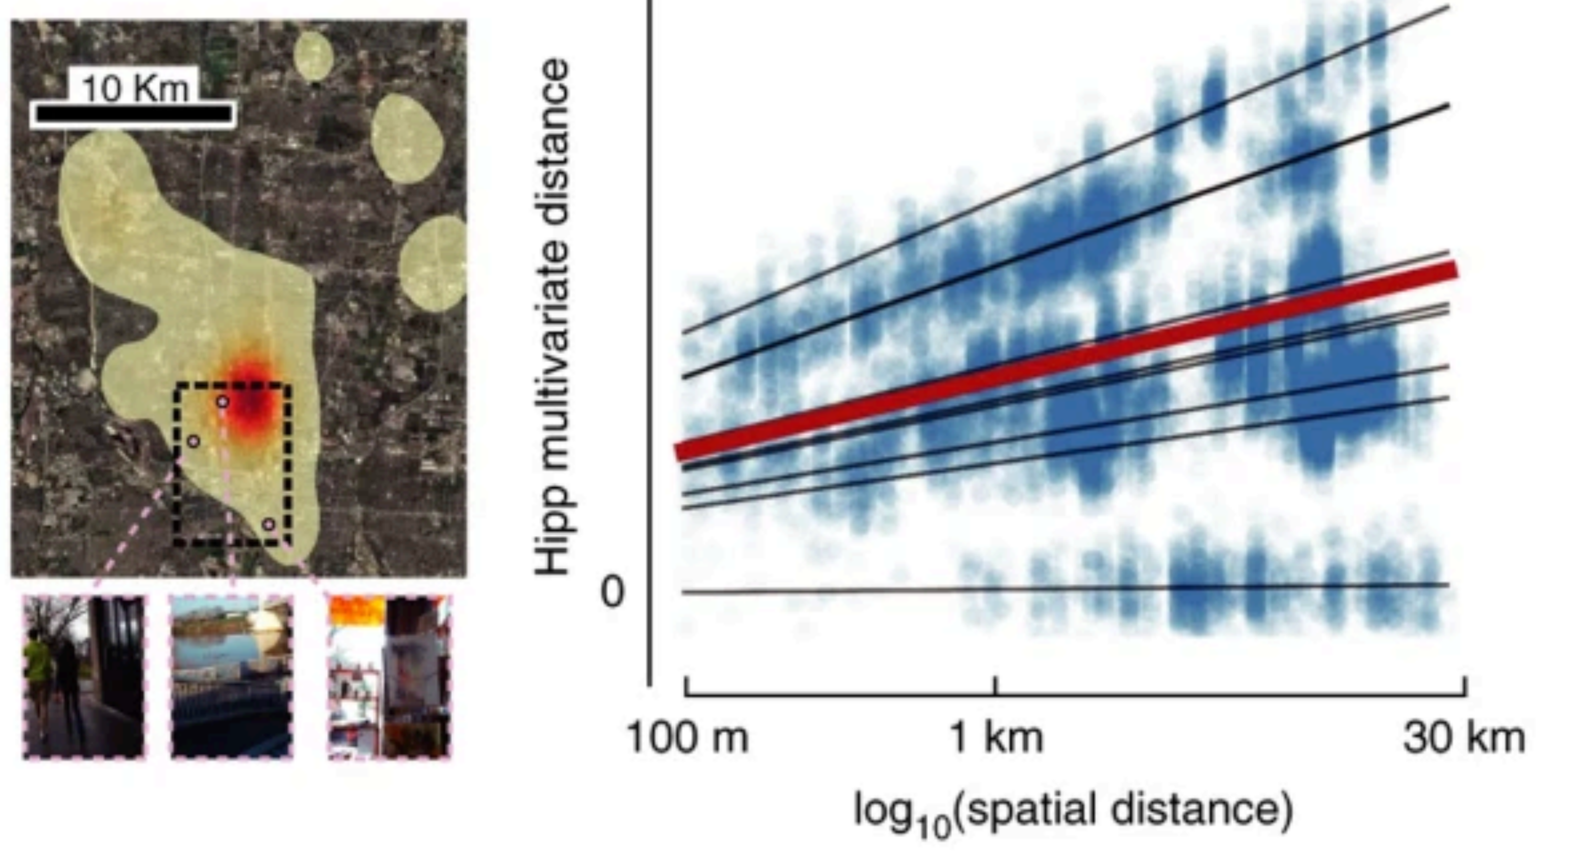

Tools for navigation

- The **Hippocampus** represents spatial distance between landmarks (Morgan et al., 2011) nand between events (Nielson et al., 2015)
- The **Entorhinal Cortex (EC)** encodes the direction of travel (Doeller et al., 2015)
  - Participants moved in a VR environment
  - When direction aligned with one of the 3 axes of their grid cells, we observe stronger BOLD activation in the EC
  - These angles are remarkably robust, and are preserved (in the same environment) when participants return to the scanner days or weeks later

Morgan et al., (*JNeuro*, 2011)

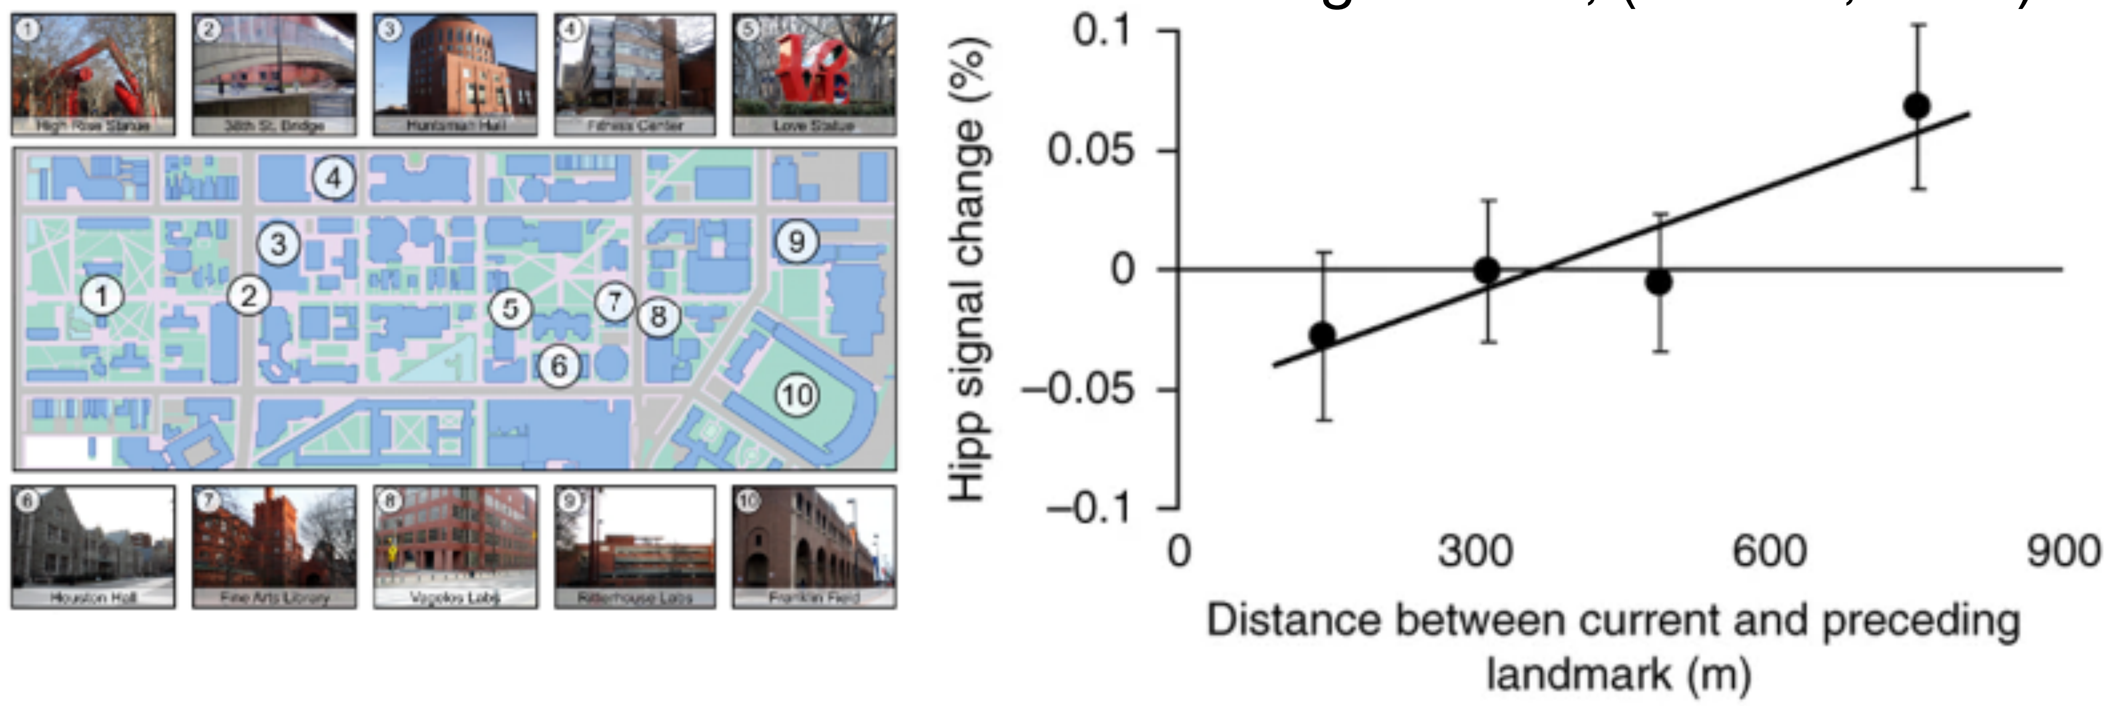

Nielson et al., (*PNAS*, 2015)

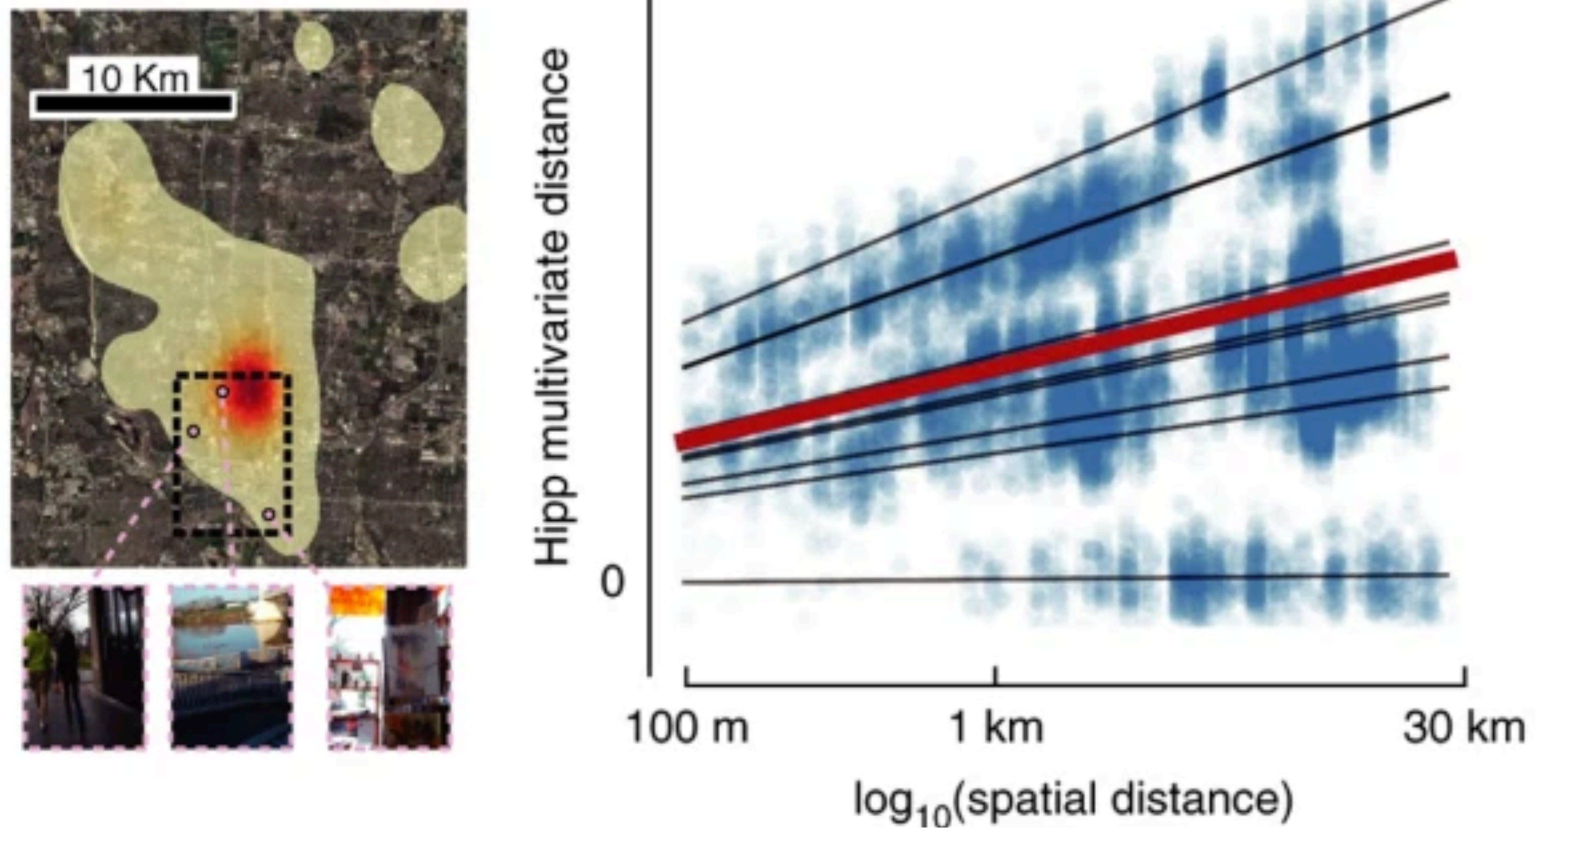

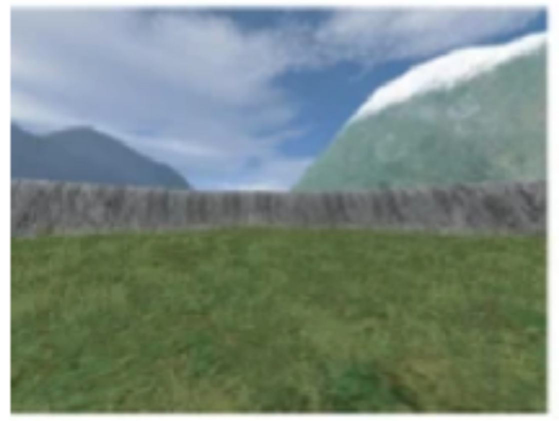

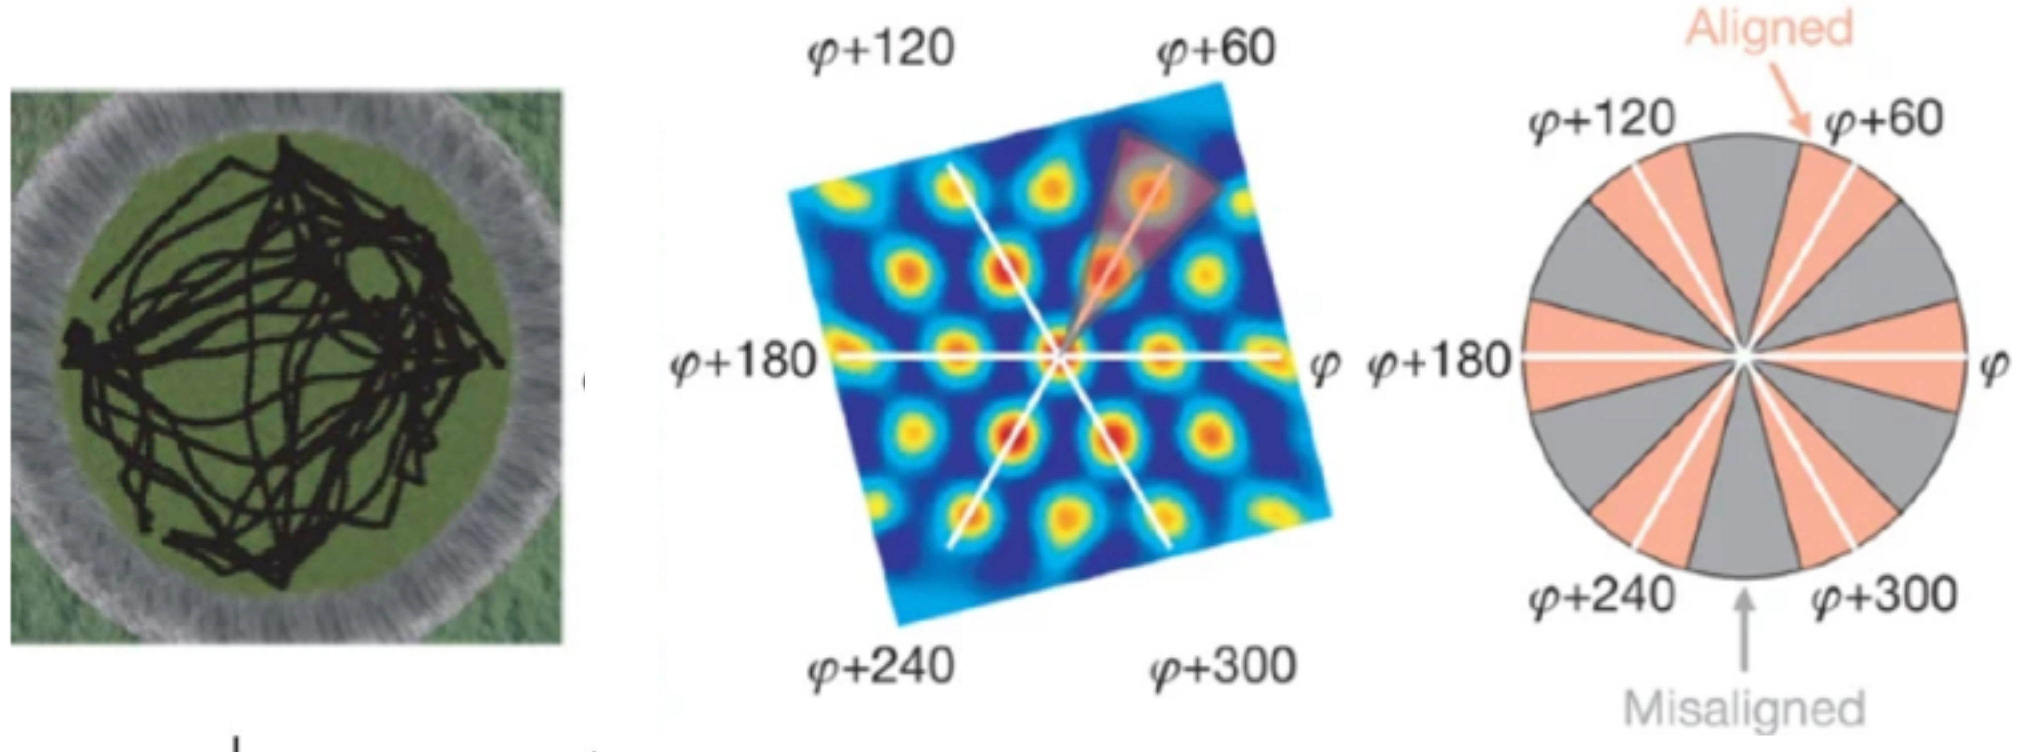

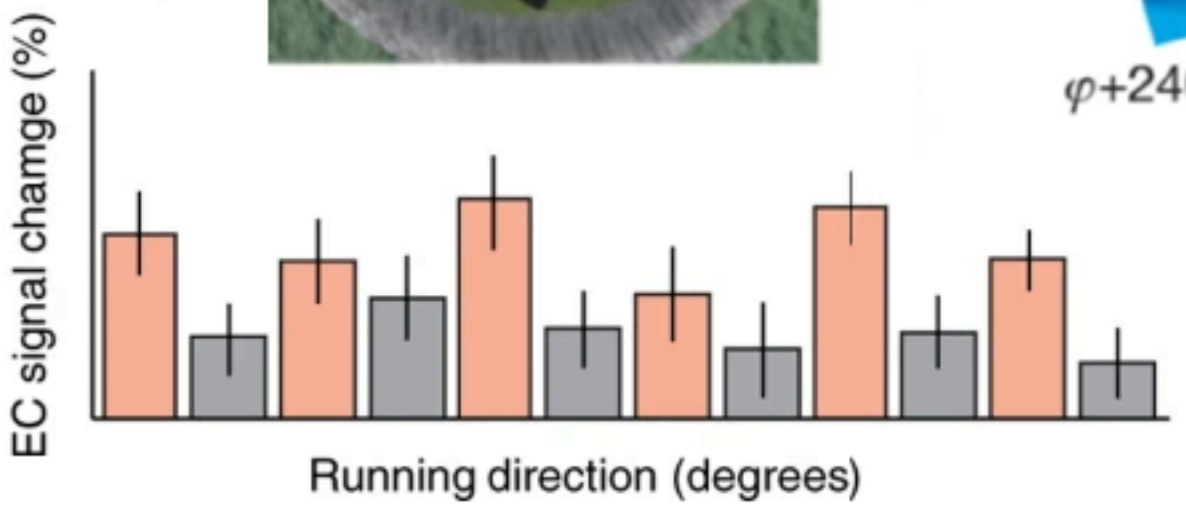

Doeller et al., (*Nature*, 2015)

Not just naïve distance, but based on the structure of the environment

A = goal

- As in Tolman's experiments, the brain represents distance in the environment based on the transition structure
- Not just "as the crow flies" but a structure-informed distance metric

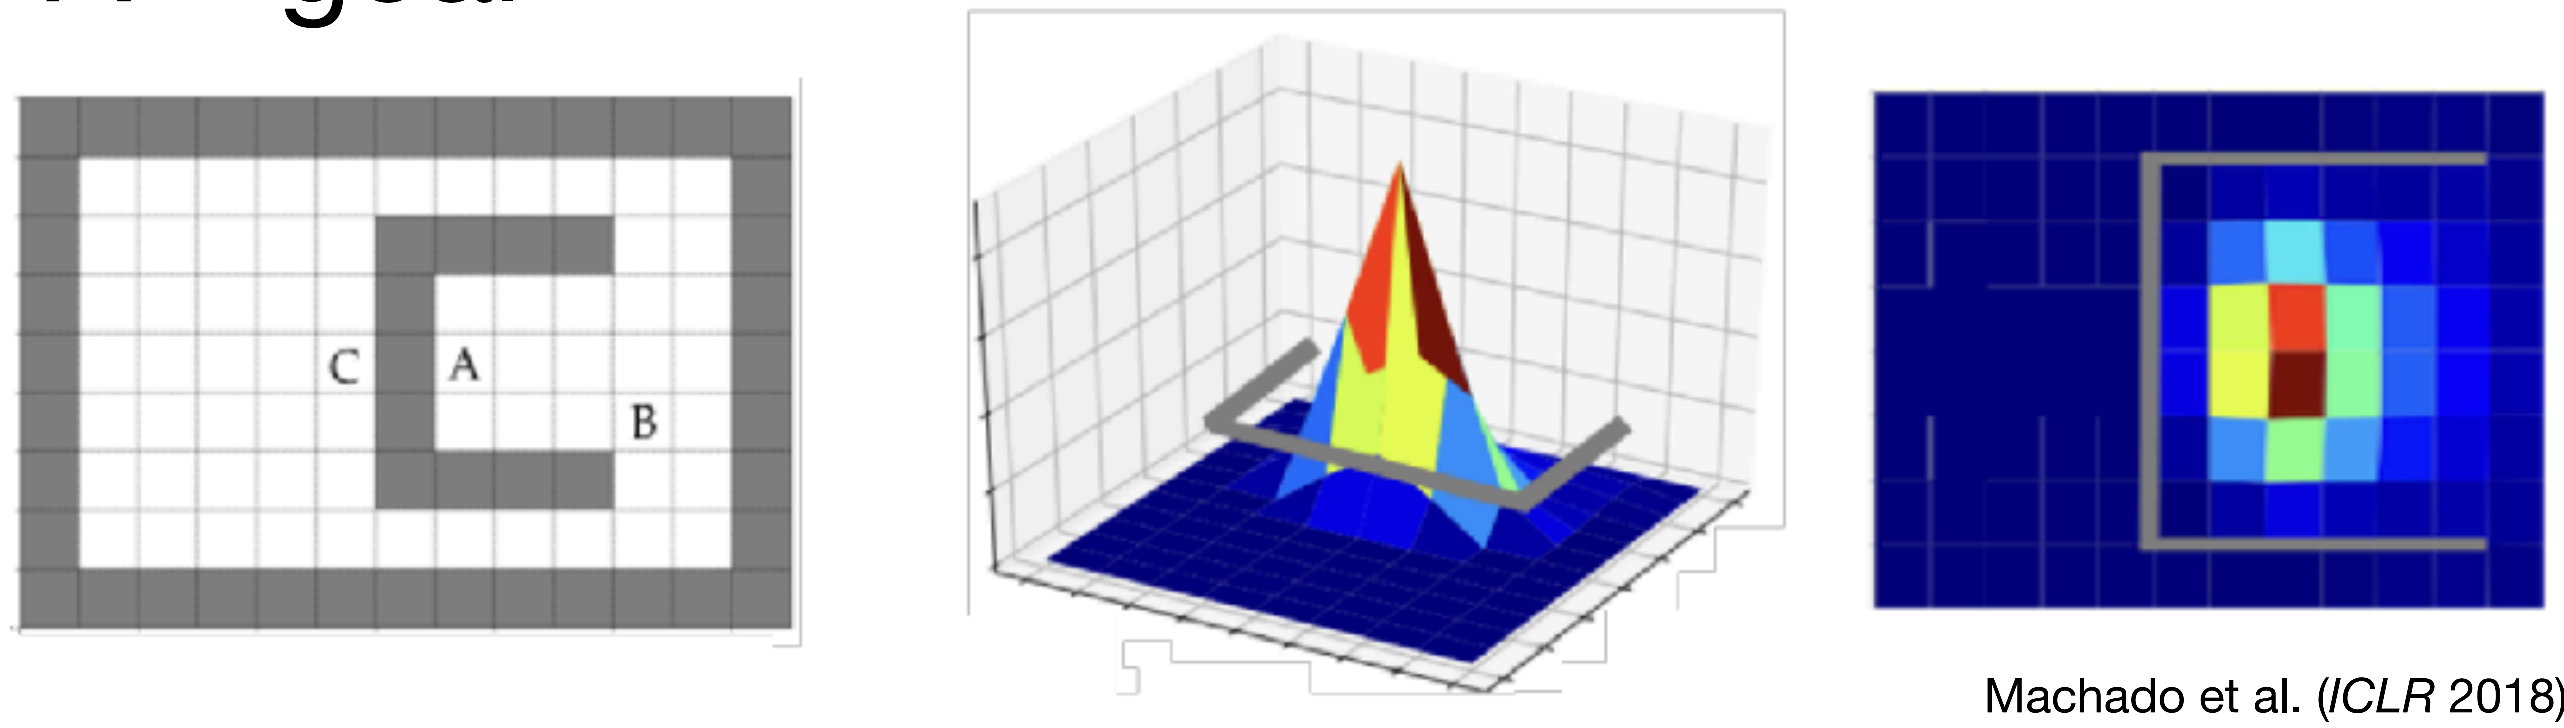

Machado et al. (*ICLR* 2018)

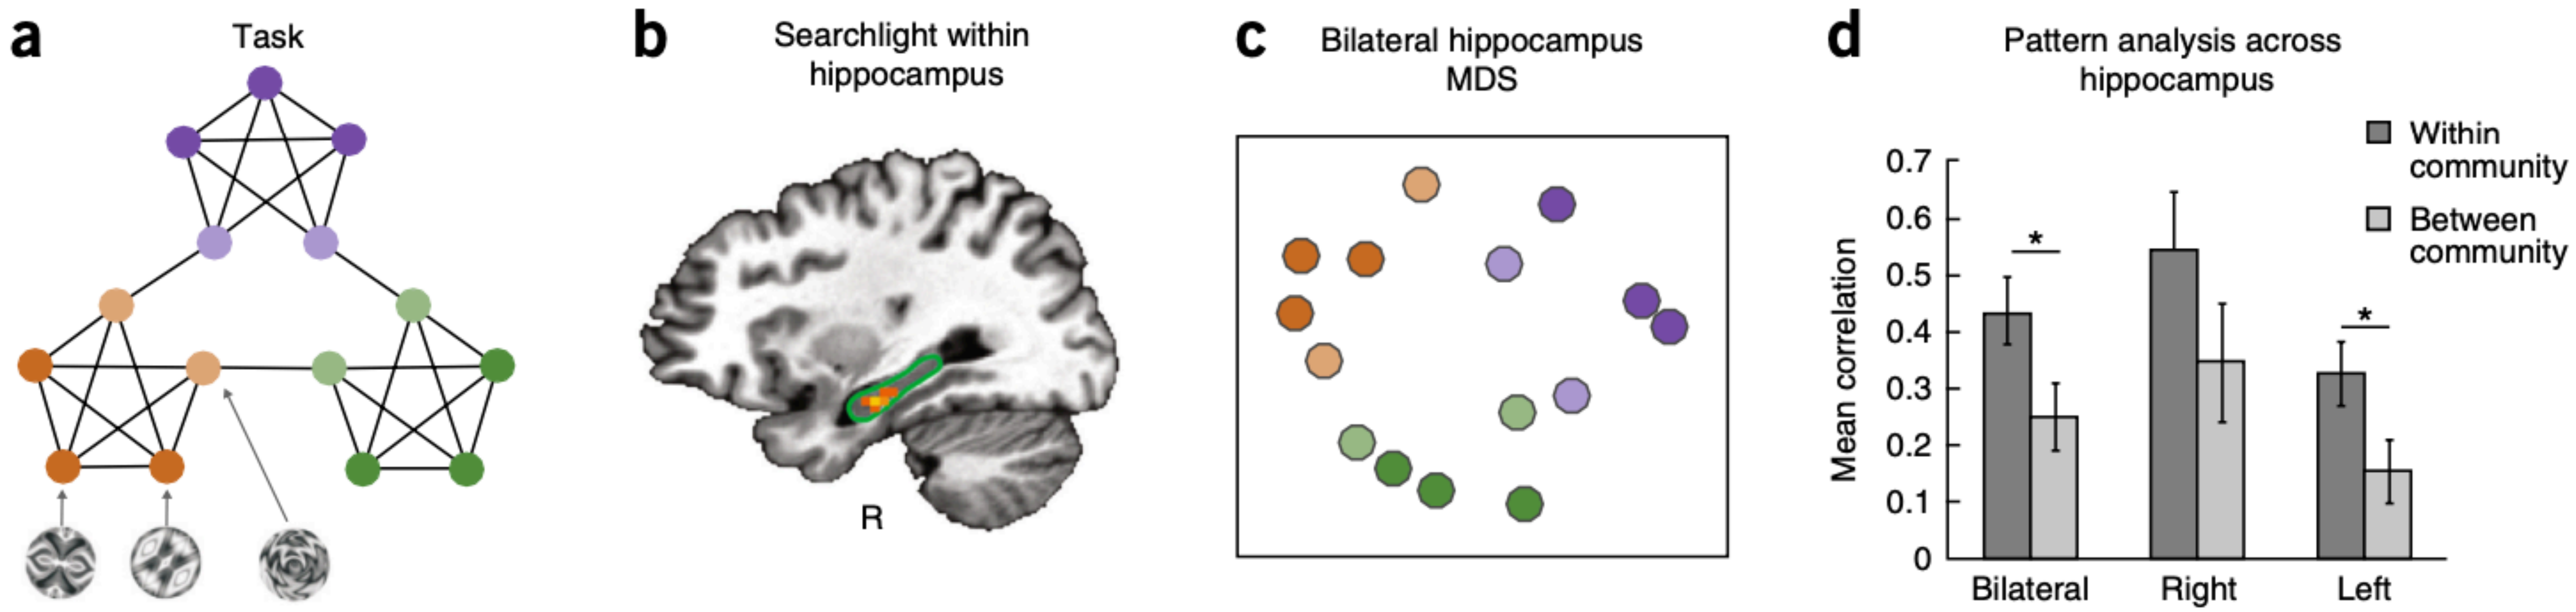

Schapiro et al. (*Hippocampus* 2013)

Not just spatial, but also conceptual navigation

Abstract features

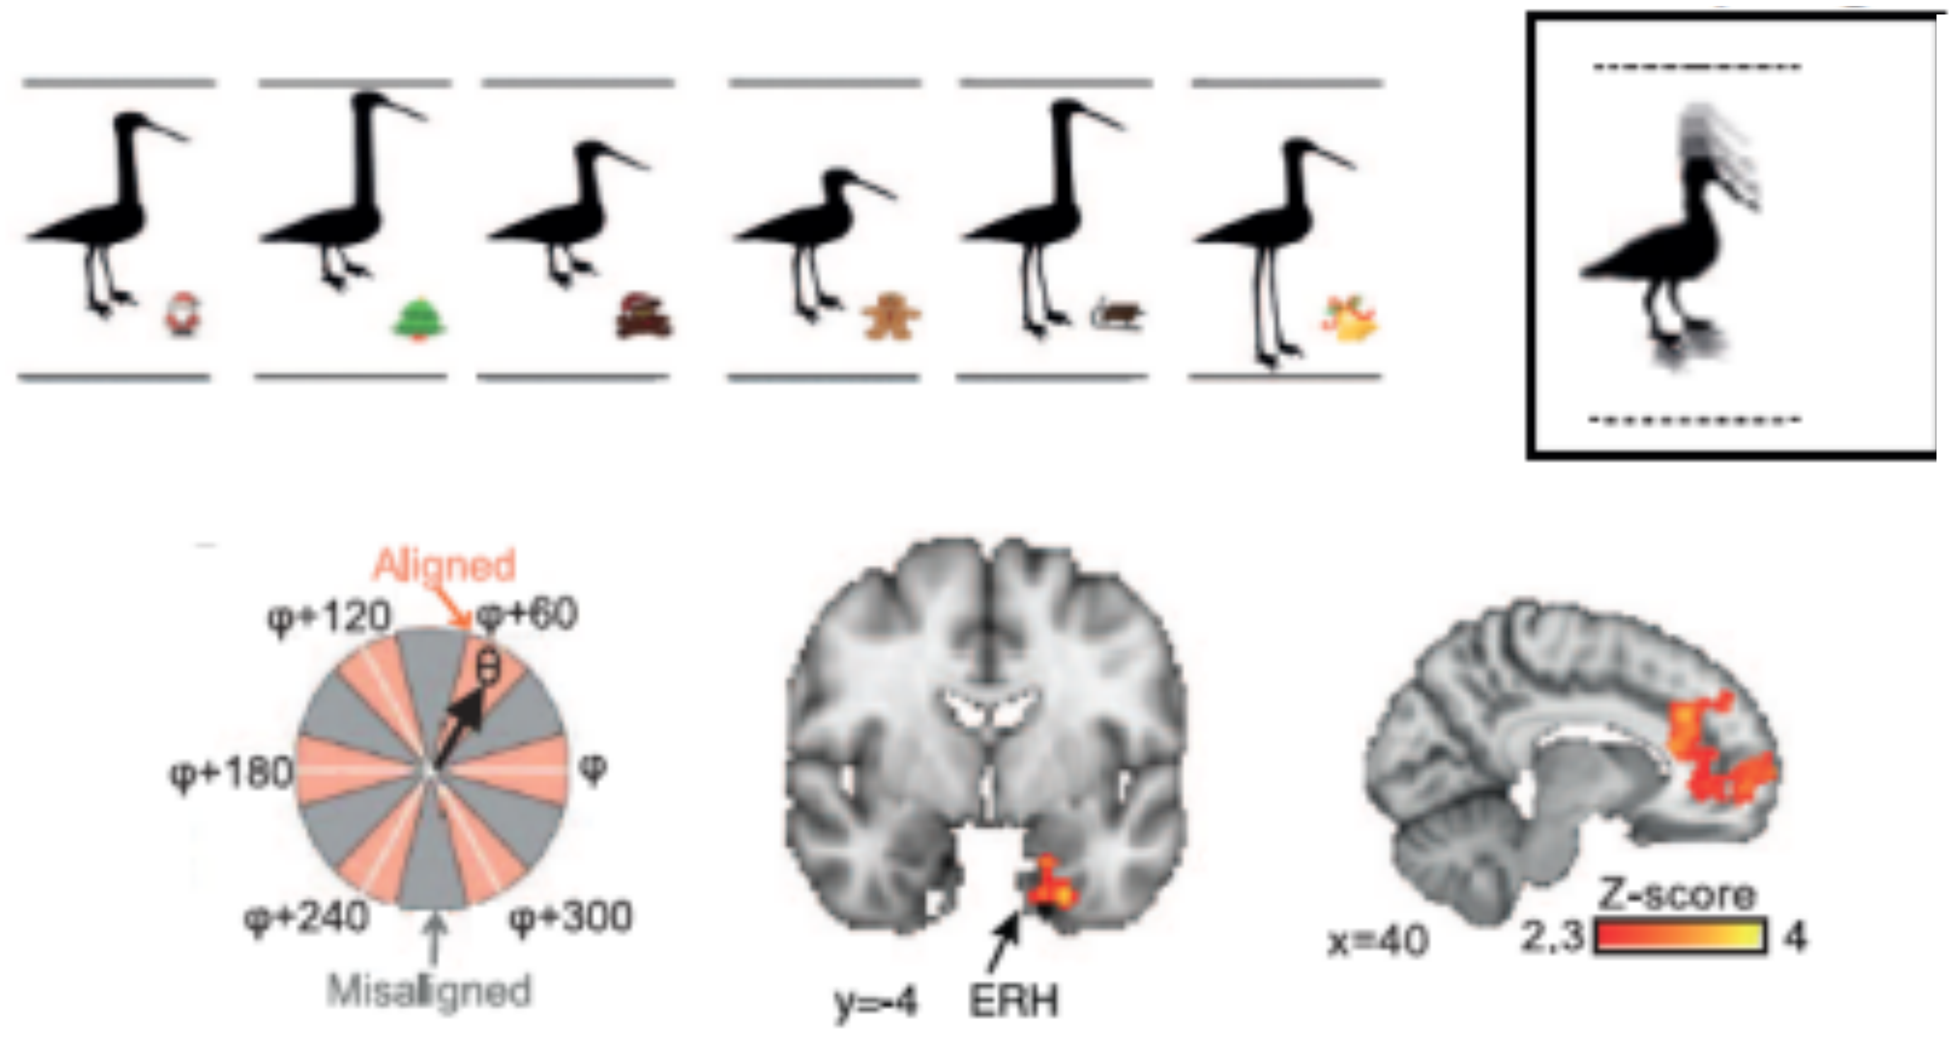

Relational structure

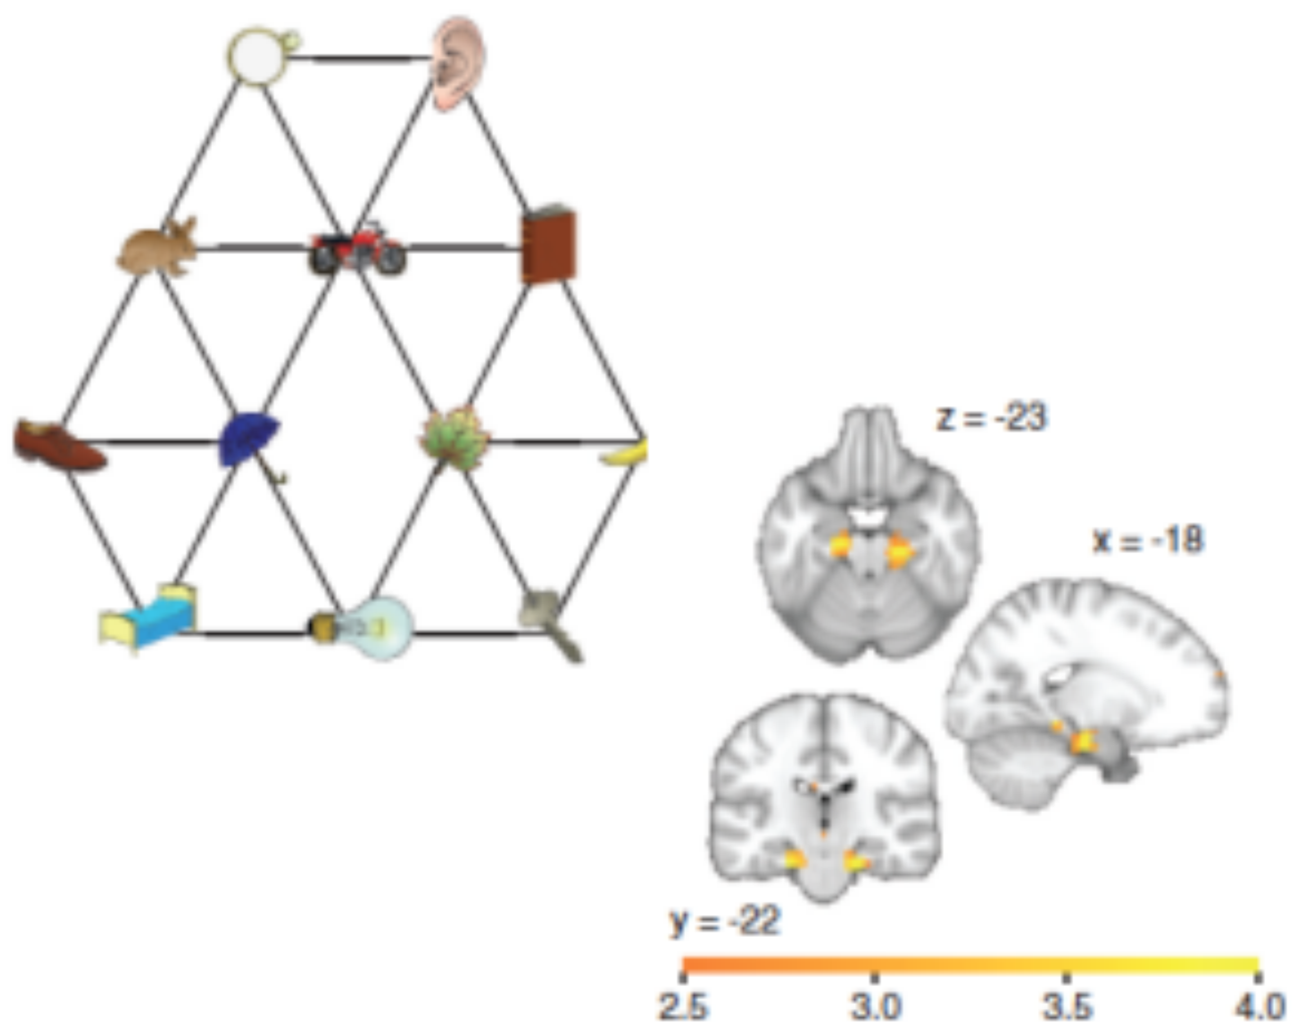

Social Hierarchies

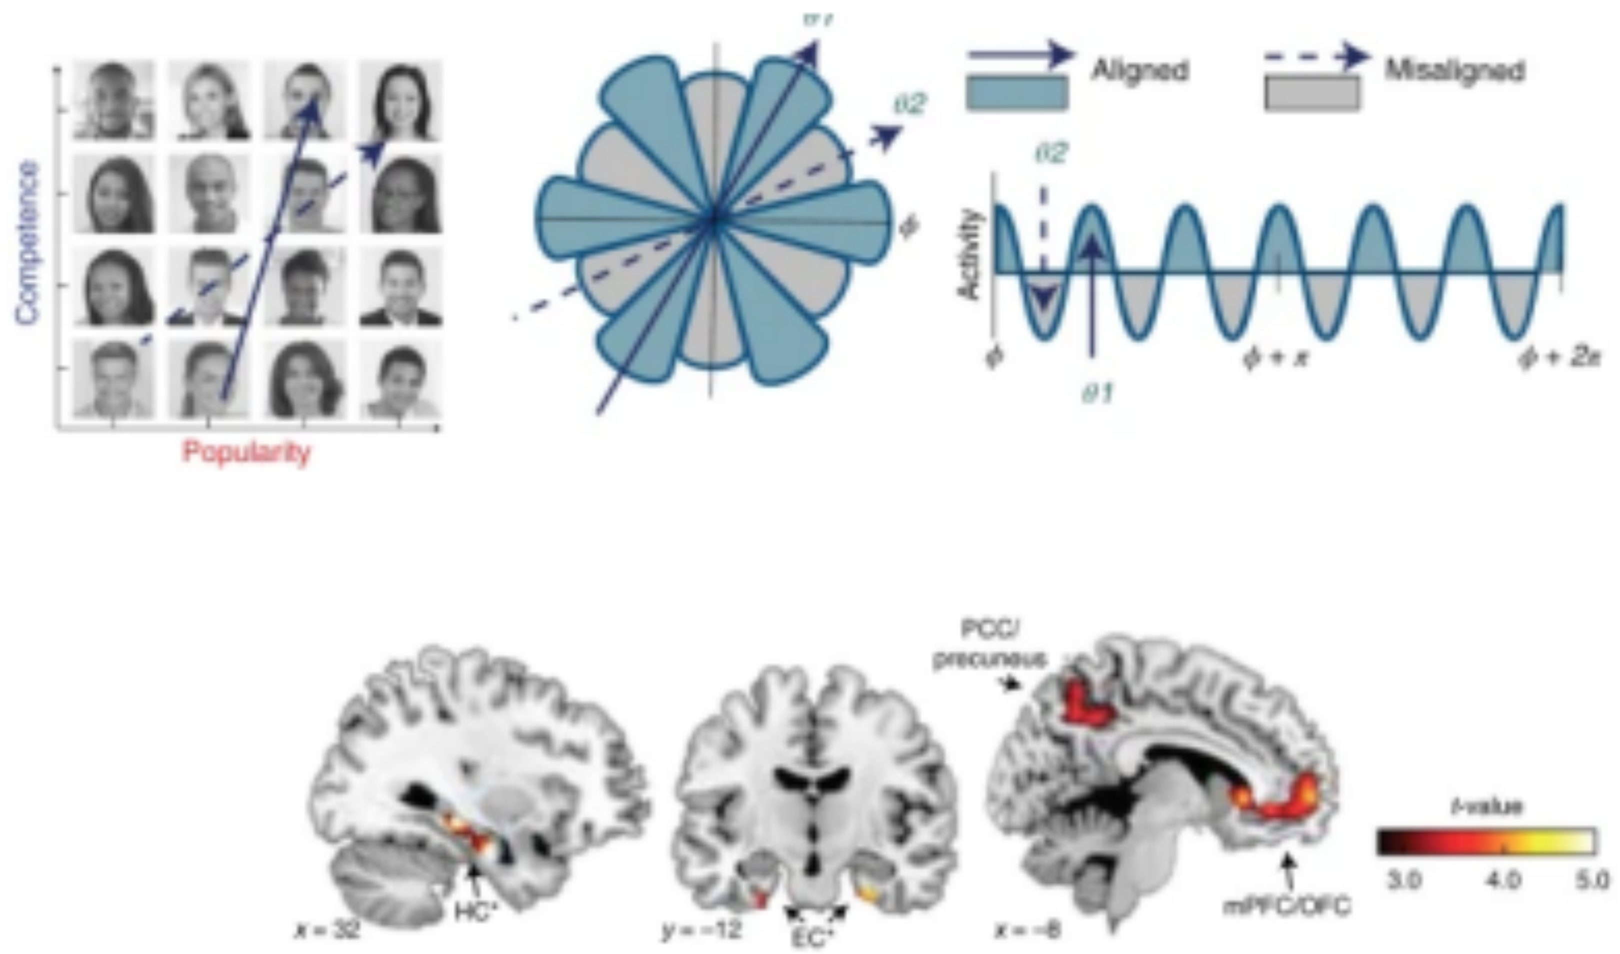

Constantinescu et al., (*Nature* 2016)

Garvert et al., (*eLife* 2017)

Park et al., (*NatNeuro* 2021)

Do we always need a representation of the environment?

An ant, viewed as a behaving system, is quite simple. The apparent complexity of its behavior over time is largely a reflection of the complexity of the environment in which it finds itself. I should like to explore this hypothesis with the word "man" substituted for "ant."

- Herbert Simon (1970)

Herbert Simon

Grandfather of AI and proponent of Bounded Rationality

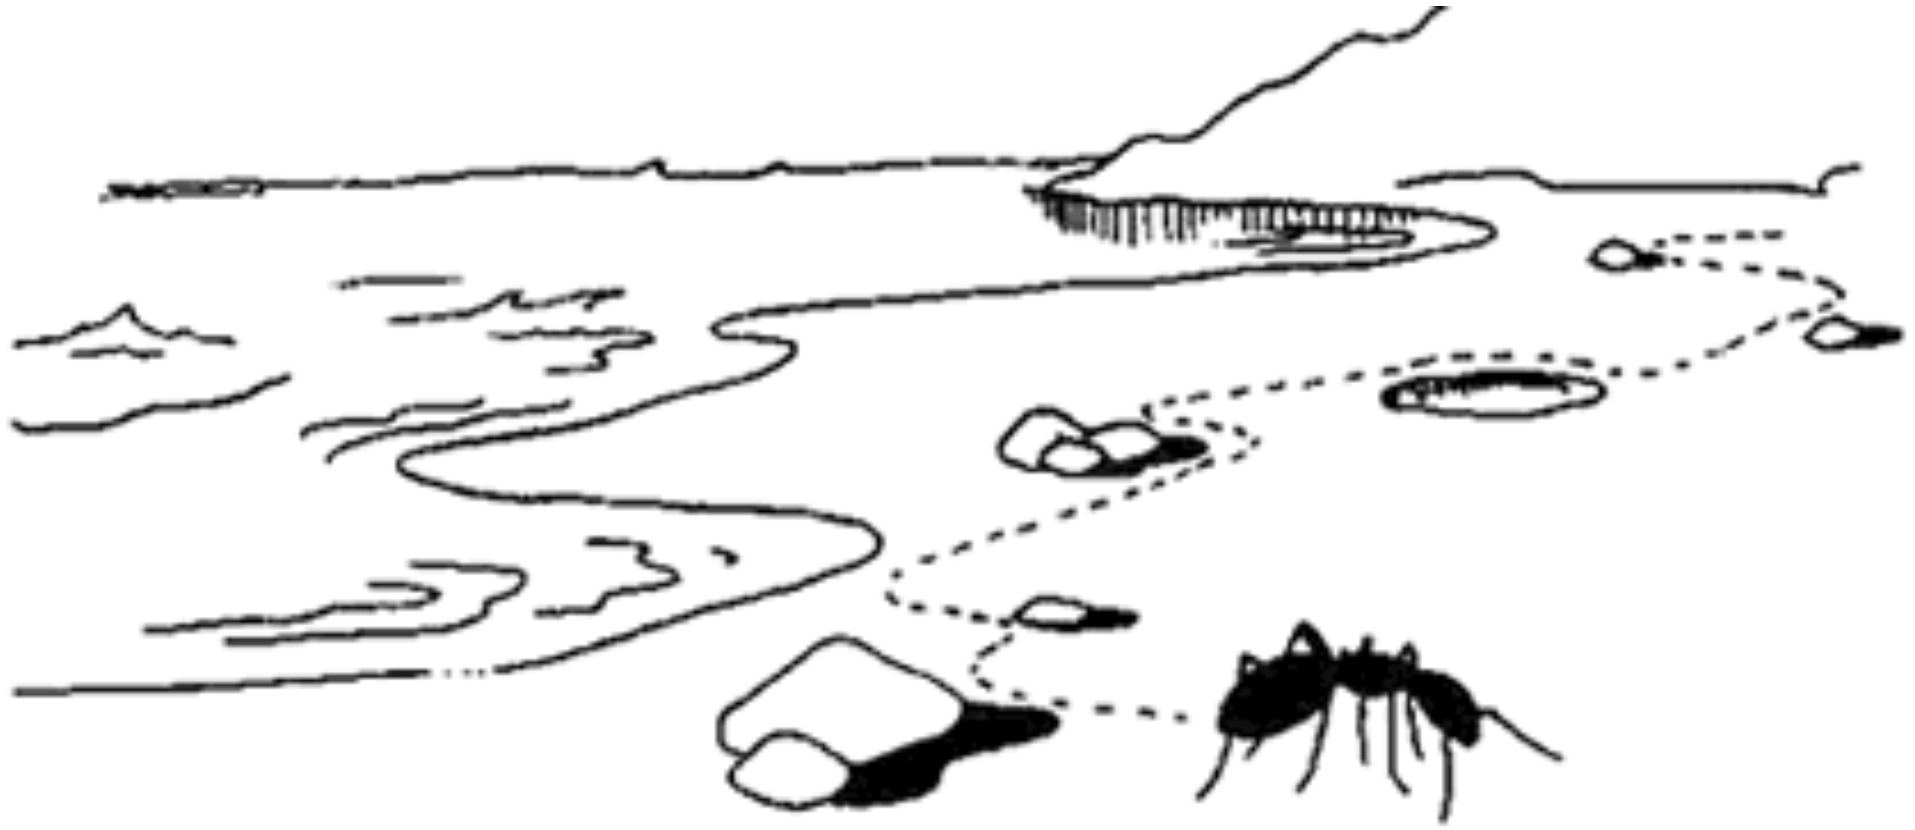

Cognitive Maps: Summary

- Learning is more than just a telephone switchboard of Stimulus-Response (S-R) associations
- We learn a **map-like representation** of the environment, allowing us to rapidly generalize and plan efficiently
  - Tolman refers to this as **S-S learning**
- Neural evidence for a cognitive map in the brain
  - **Place cells** in the Hippocampus encode location and distances
  - **Grid cells** in the Entorhinal Cortex provide a coordinate system and encode direction of travel
  - + a whole zoo of other specialized cells in the hippocampal-entorhinal system
- Cognitive maps are sensitive to **transition structure** and used in abstract, **conceptual** contexts as well

General principles

- **Symbolic AI**: Learning as infering rules and manipulating symbols
  - In contrast to *subsymbolic* AI (i.e., neural networks), which learn by updating associating weights
  - For symbolic AI, learning corresponds to search over hypotheses, but current solutions are intractable/inefficient in most interesting settings
    - How do people manage to learn symbolic rules/programs efficiently?
- **Cognitive maps**: Learning as inferring a representation of the structure of the environment
  - Not just S-R relationships but also S-S latent learning
  - Do we always need a representation of the environment?
- Both lines of research capture mechanisms for learning **structure**
  - Structure as the relationships between different symbolic concepts
  - Structure as the relationship between stimuli in the environment
- Is there a common basis for both forms of learning? Or are they complementary systems?

Next week: Introduction to Reinforcement Learning

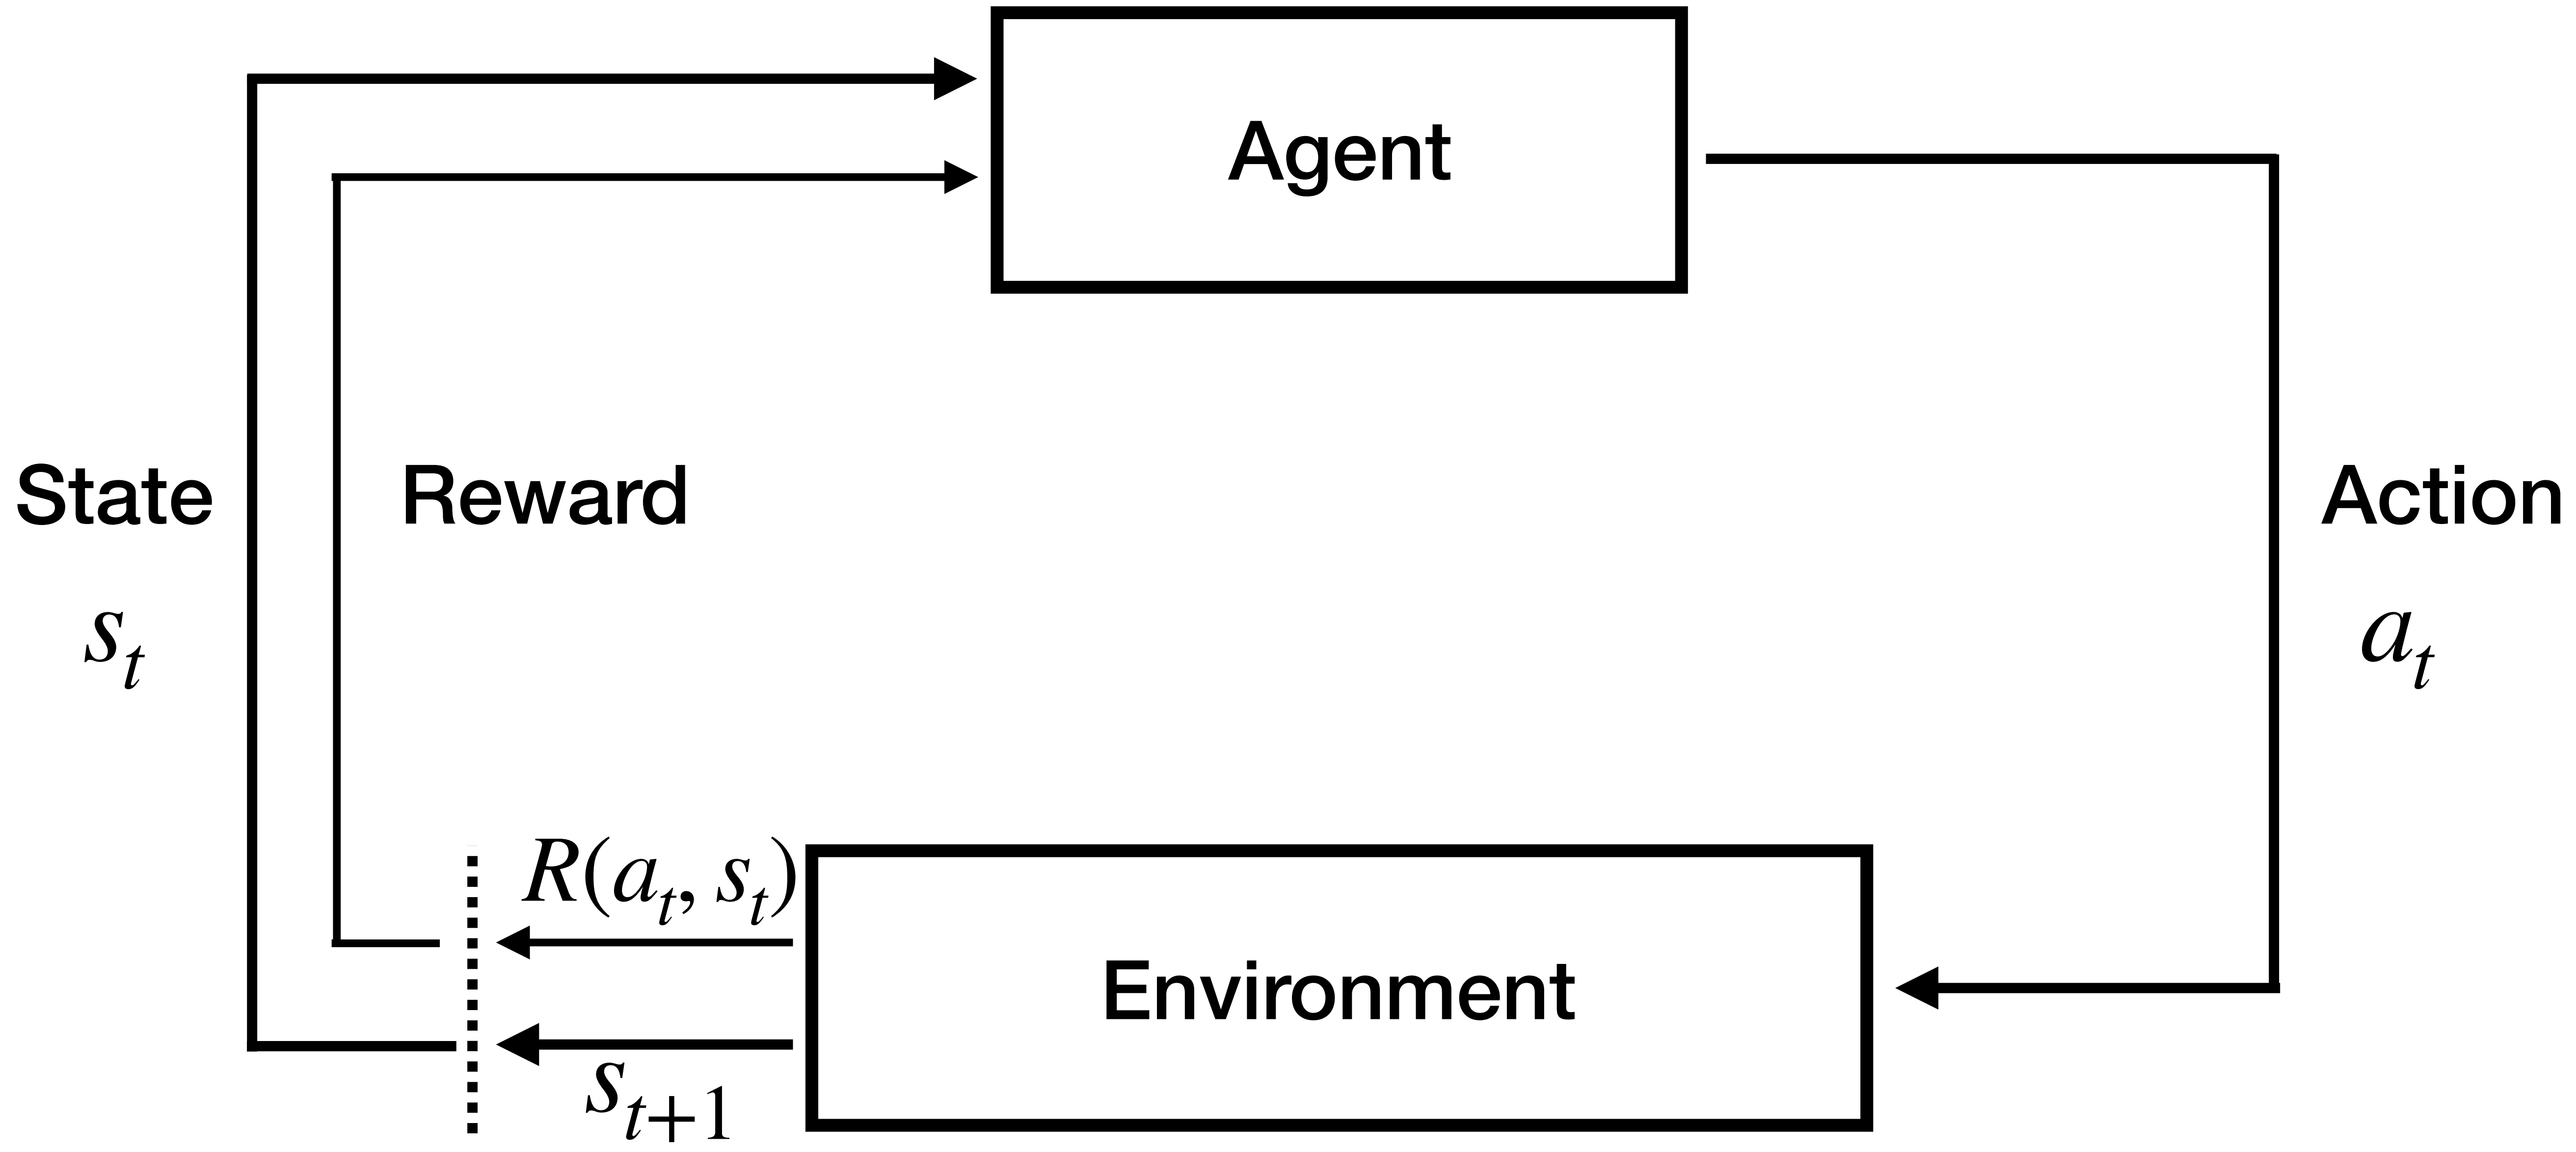

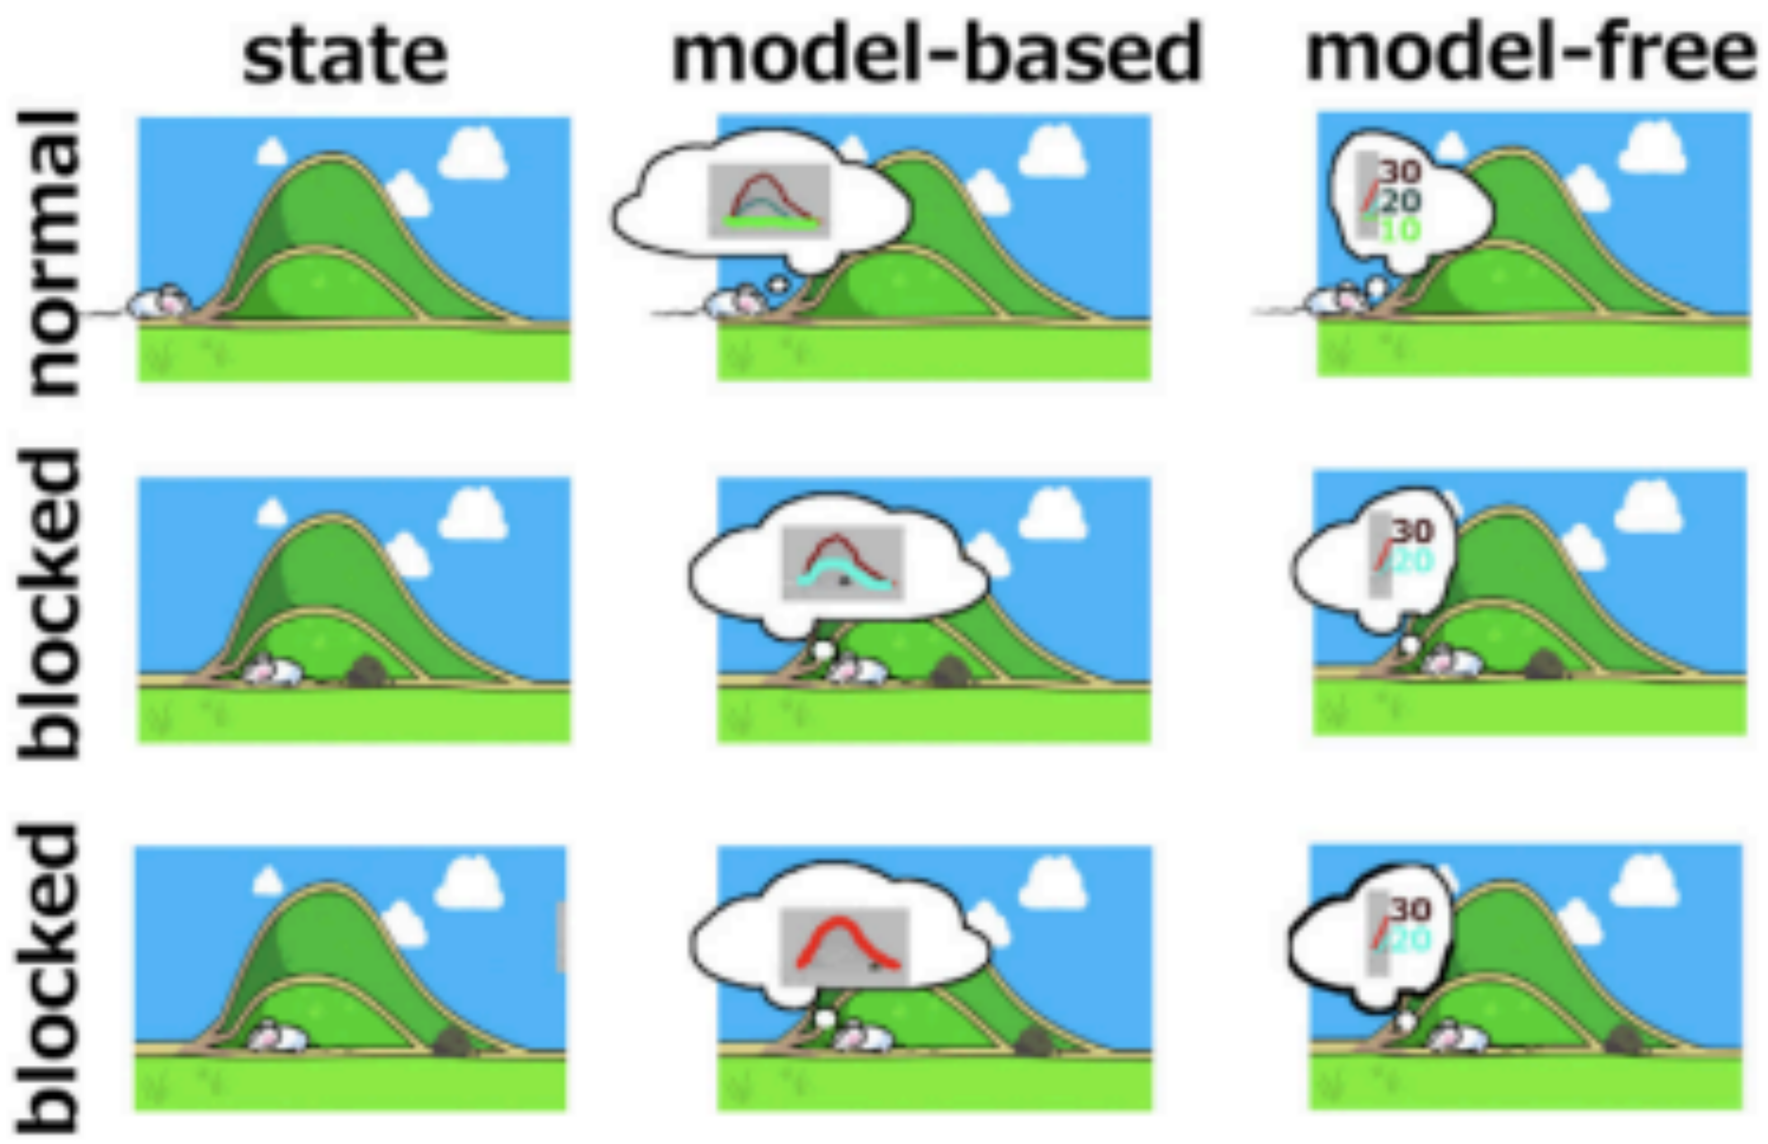

Model-based and **model-free decision making** in a cartoon of a maze invented by Tolman and Honzik (1930)