General Principles of Human and Machine Learning

Tutorial 12: General Principles

Dr. Charley Wu

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

Exam

- Combination of multiple choice and short answer questions
  - No complex calculations are needed 
  - No need to memorize formulas or dates
  - Focus on understanding the main theoretical ideas and how they connect across fields
  - Bring pens/pencils
- First taking: Friday, Feb 21st, 13:00 -15:00
  - Hörsaal 1, F119 (SAND 6/7)
- Second taking: Friday April 11th, 12:00 –14:00
  - Ground floor lecture room, AI building (Maria-von-Linden-Str. 6)

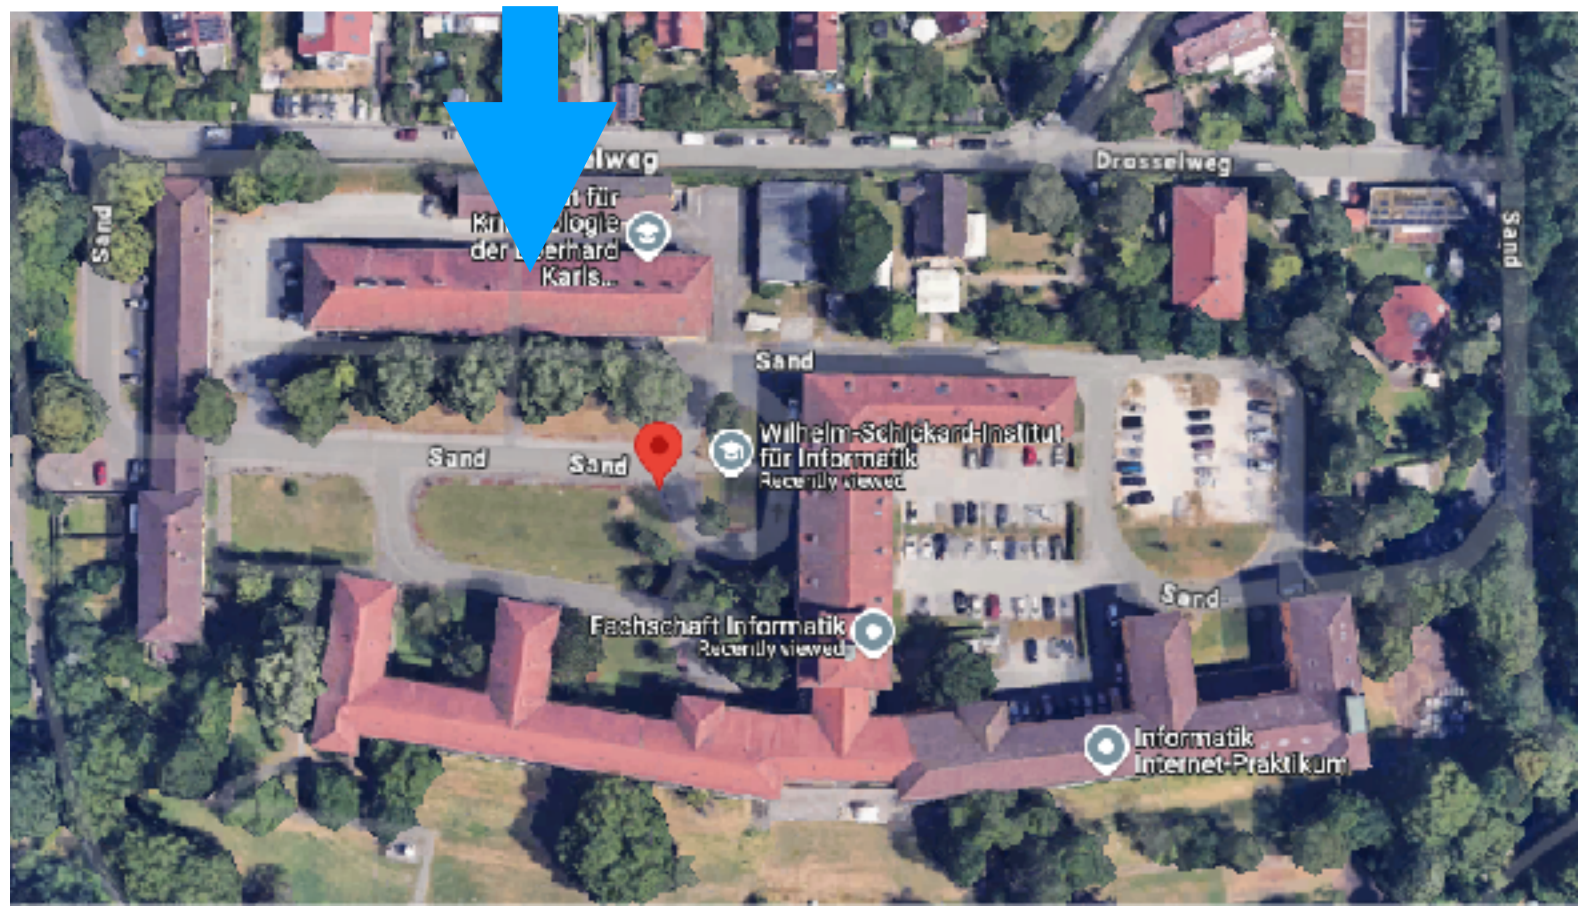

Tutorial Exercise

- Review of some challenging questions from Quiz #4

- Exam preparation
  - Last chance to enter some candidate exam questions in the google doc
    - Short answer question format
    - You are incentivized to bring plausible questions that would be sufficiently challenging, thought provoking, and feasible
    - Good questions will be included on the exam

[Google doc]

1. Name one advantage and one disadvantage for each approach to concept learning:

   a. Rule-based approaches

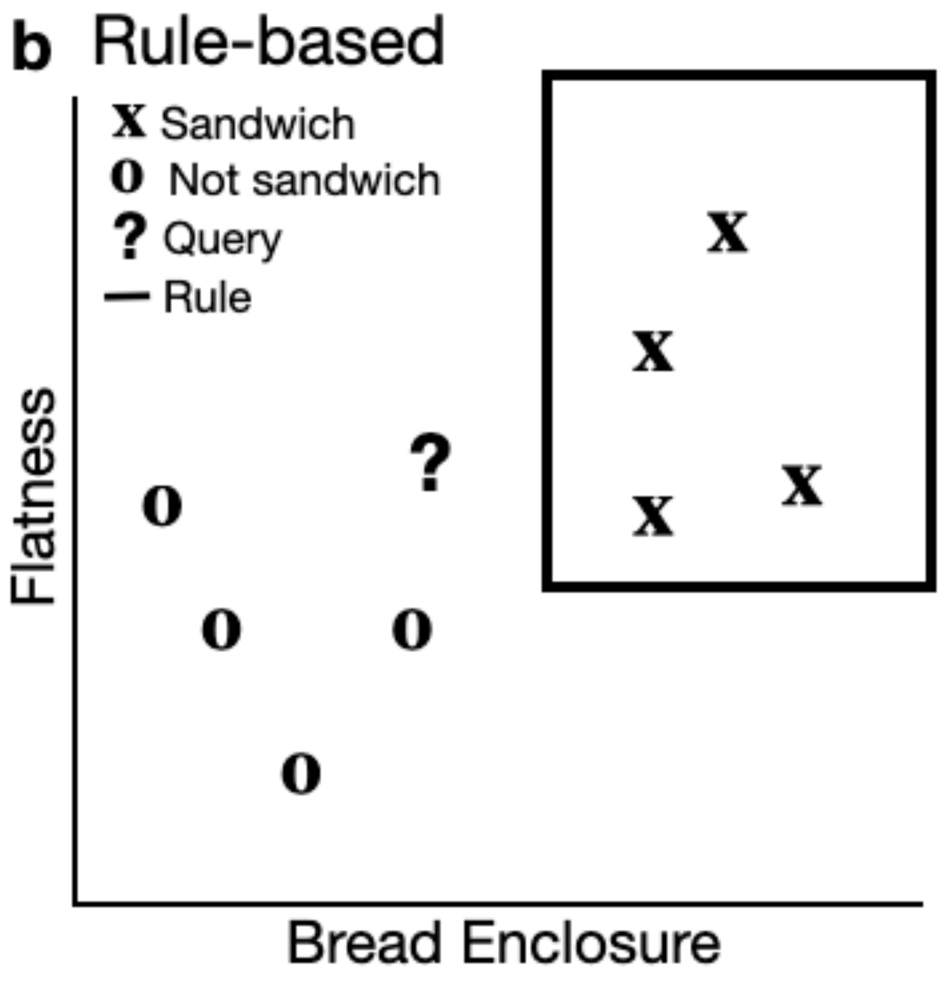

   b. Similarity-based approaches

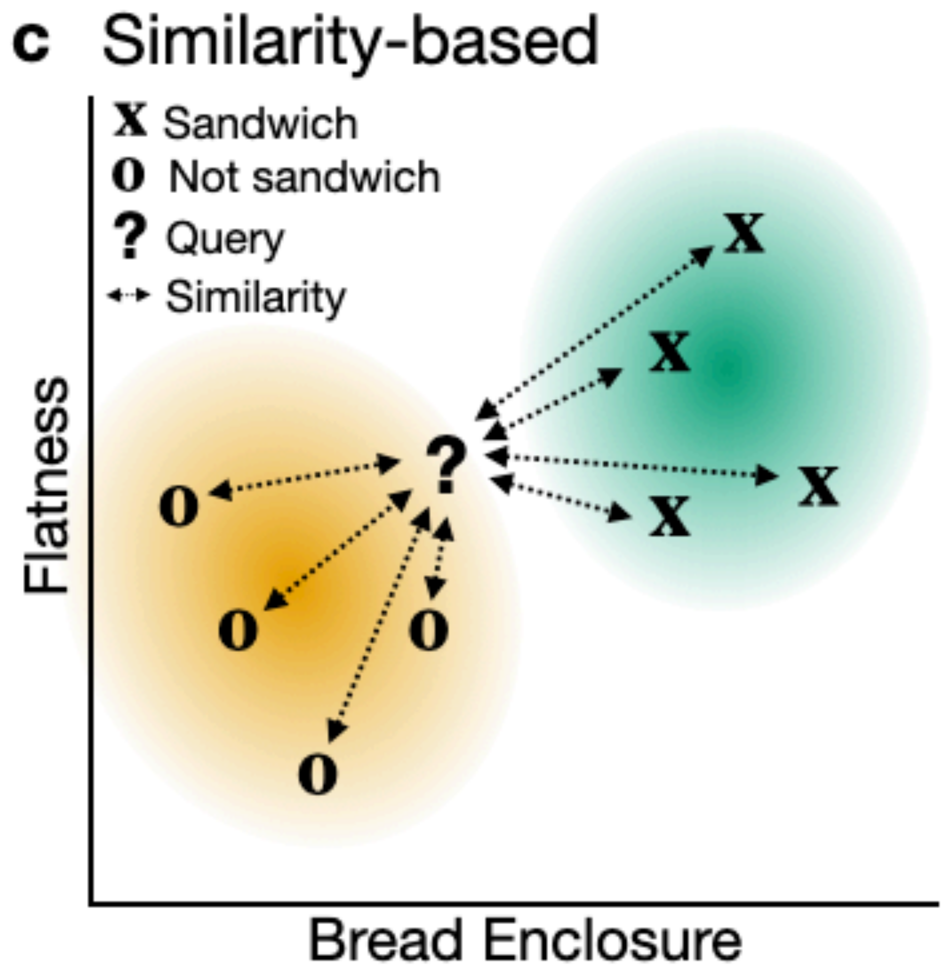

1. Name one advantage and one disadvantage for each approach to concept learning:

a. Rule-based approaches

Advantages. Facilitates rapid generalization and allows compositional rules to create infinitely productive systems. (1 mark)

Disadvantages. Rigidity and inflexibility; struggles with exceptions or fuzzy categories (e.g., open-faced sandwiches). (1 mark)

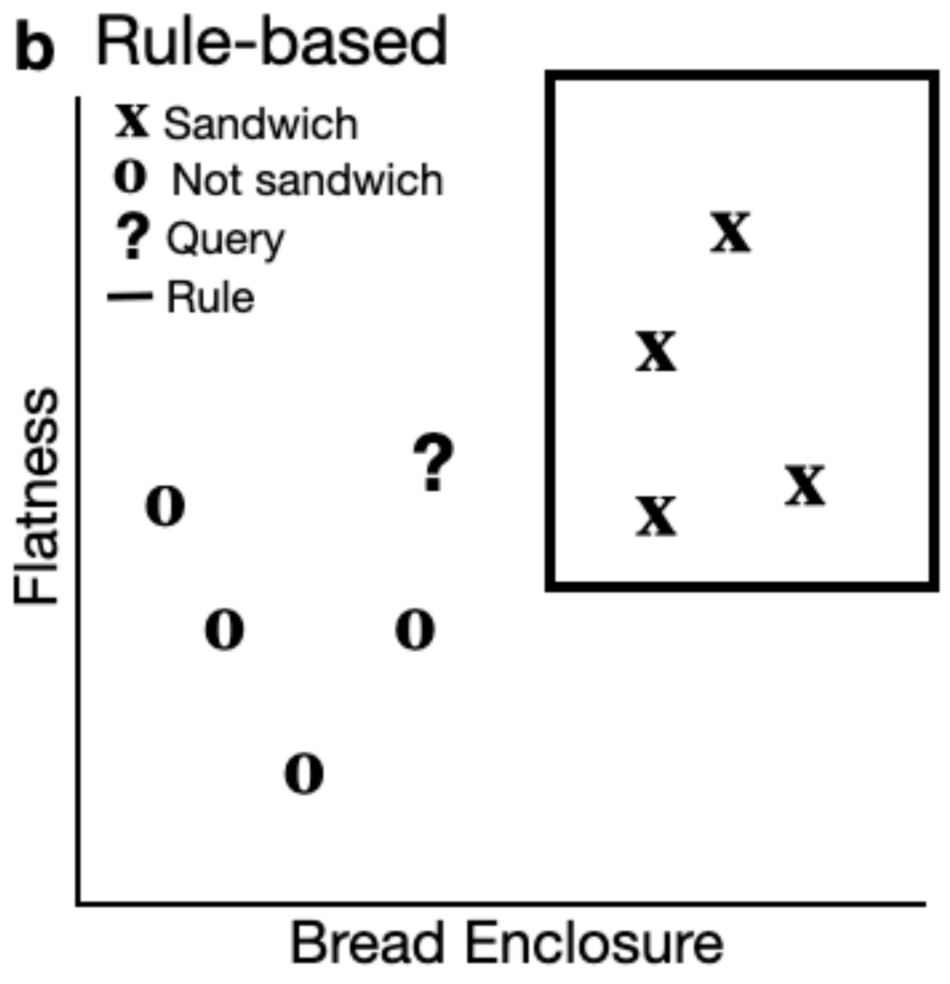

b. Similarity-based approaches

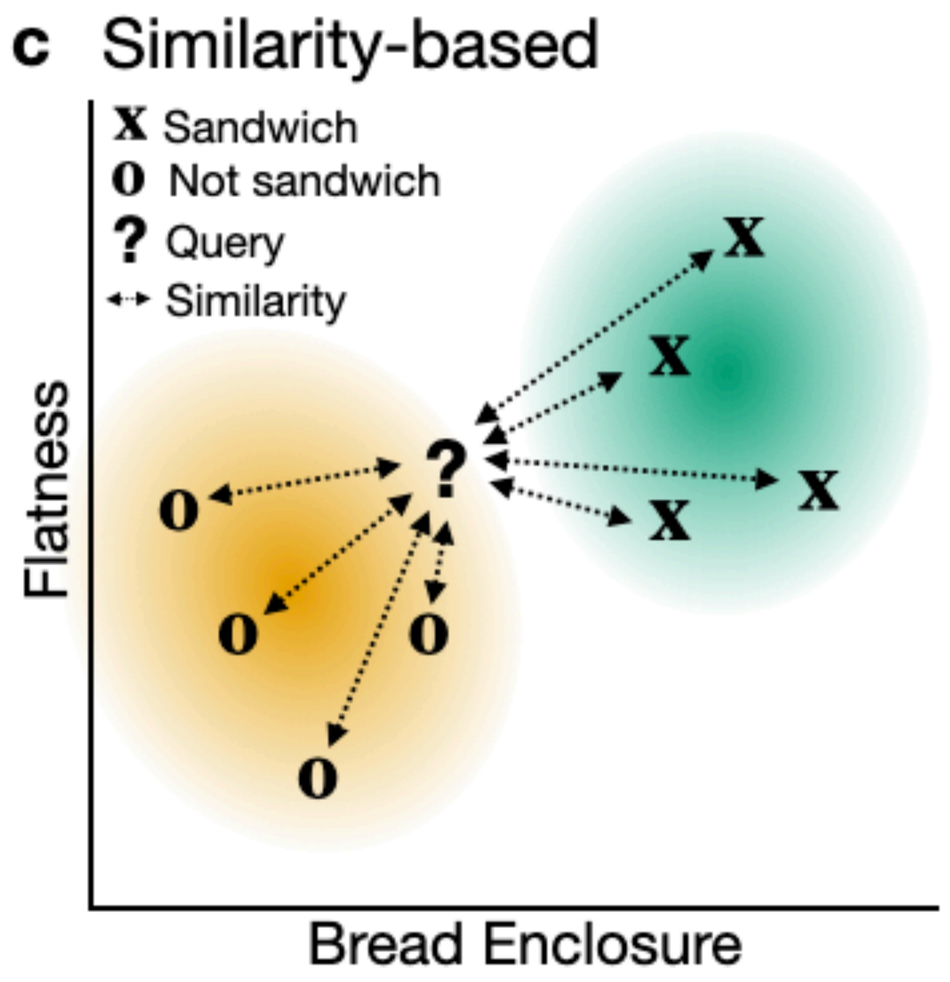

1. Name one advantage and one disadvantage for each approach to concept learning:

a. Rule-based approaches

Advantages. Facilitates rapid generalization and allows compositional rules to create infinitely productive systems. (1 mark)

Disadvantages. Rigidity and inflexibility; struggles with exceptions or fuzzy categories (e.g., open-faced sandwiches). (1 mark)

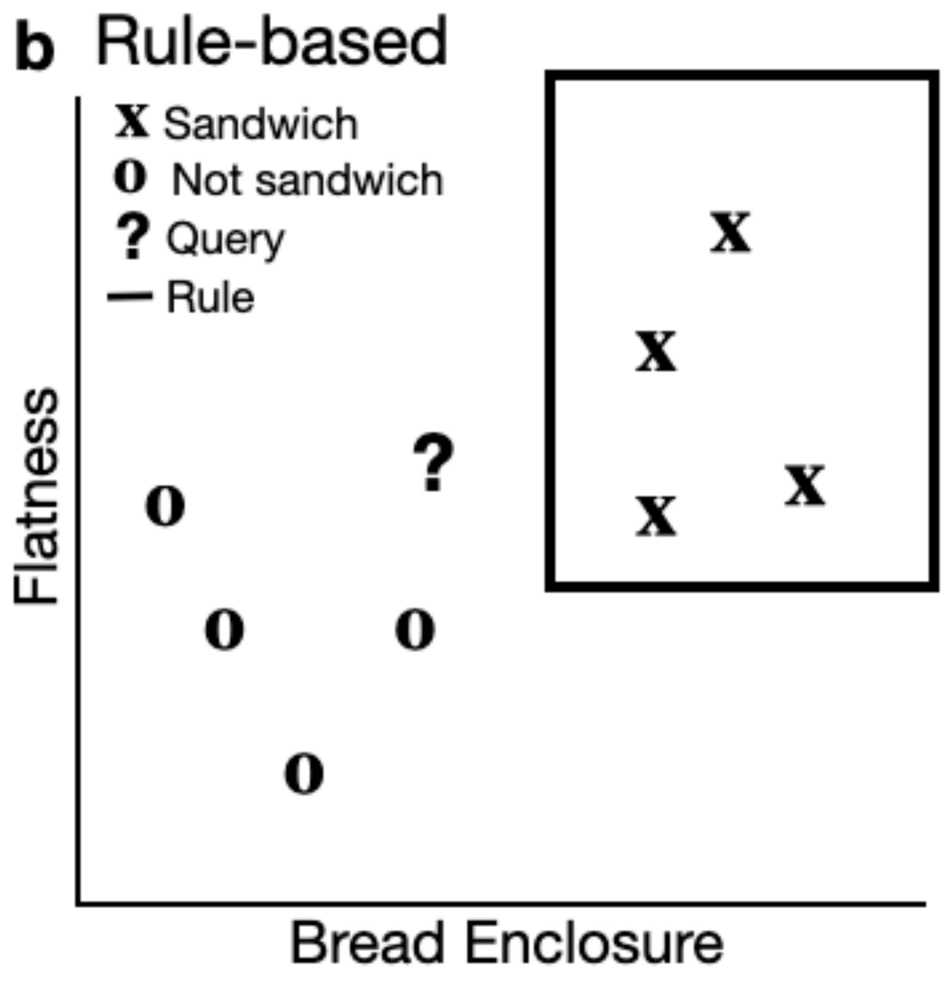

b. Similarity-based approaches

Advantages. Can make on-the-fly generalizations by comparison to past stimuli, intuitive that stimuli with similar features are more likely to belong to the same category, simple method for evaluating class membership based on distance in feature space

Disadvantages: Fails to capture more structured concept representations, difficulties in choosing the right similarity metric, (when using metric similarity) fails to account for violations of symmetry and triangle inequality axioms

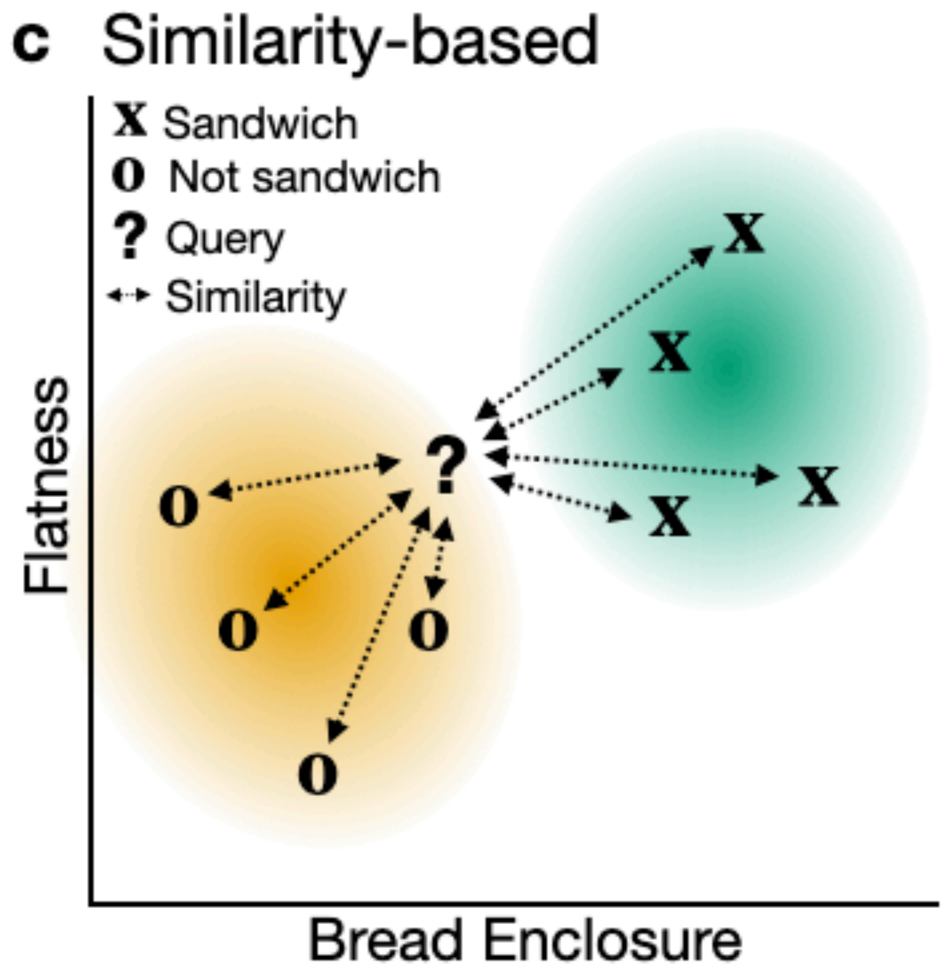

2. How do exemplar and prototype theories explain category representation differently? In what situation might each be more useful?:

a. Exemplar

b. Prototype

2. How do exemplar and prototype theories explain category representation differently? In what situation might each be more useful?:

a. Exemplar

Definition: Categories are represented by specific stored examples. Useful for fine-grained or novel categorization based on unique experiences. (1 mark)

Examples of when most useful: Identifying a peacock as a bird based on a specific zoo visit; Recognizing a painting as an artist's work due to familiarity with a prior piece; Diagnosing a rare disease from memory of a similar past case. In general, when the category boundaries are always clearcut (i.e., there are outliers), and comparison to other category members is more helpful than simply aggregating over all members. (1 mark for examples; the “in general” part is just to clarify the logic)

b. Prototype

2. How do exemplar and prototype theories explain category representation differently? In what situation might each be more useful?:

a. Exemplar

Definition: Categories are represented by specific stored examples. Useful for fine-grained or novel categorization based on unique experiences. (1 mark)

Examples of when most useful: Identifying a peacock as a bird based on a specific zoo visit; Recognizing a painting as an artist's work due to familiarity with a prior piece; Diagnosing a rare disease from memory of a similar past case. In general, when the category boundaries are always clearcut (i.e., there are outliers), and comparison to other category members is more helpful than simply aggregating over all members. (1 mark for examples; the “in general” part is just to clarify the logic)

b. Prototype

Definition: Categories are represented by an abstract prototype summarizing the most typical features. Useful for generalization across typical category members. (1 mark)

Examples of when most useful: Categorizing a robin as a bird because it looks like a typical bird; Classifying a cushioned, four-legged object as a chair based on typical features; Recognizing a Labrador as a dog because it aligns with the average dog prototype. In *general*, when the memory limitations make it difficult to retain a large number of past exemplars or not much information is lost when aggregating over past category members (1 mark)

3. According to Shepard, why does generalization occur?

3. According to Shepard, why does generalization occur?

- Shepard (1987) believed that representations about categories or natural kinds correspond to a *consequential region* in psychological space

- **Generalization arises from uncertainty about the extent of these regions**

- As representational distance between stimuli **x** and **x'** increases (i.e., become less **similar**), they are less likely to belong to the same region, and thus produce less similar outcomes

- This produces the smooth gradient of generalization

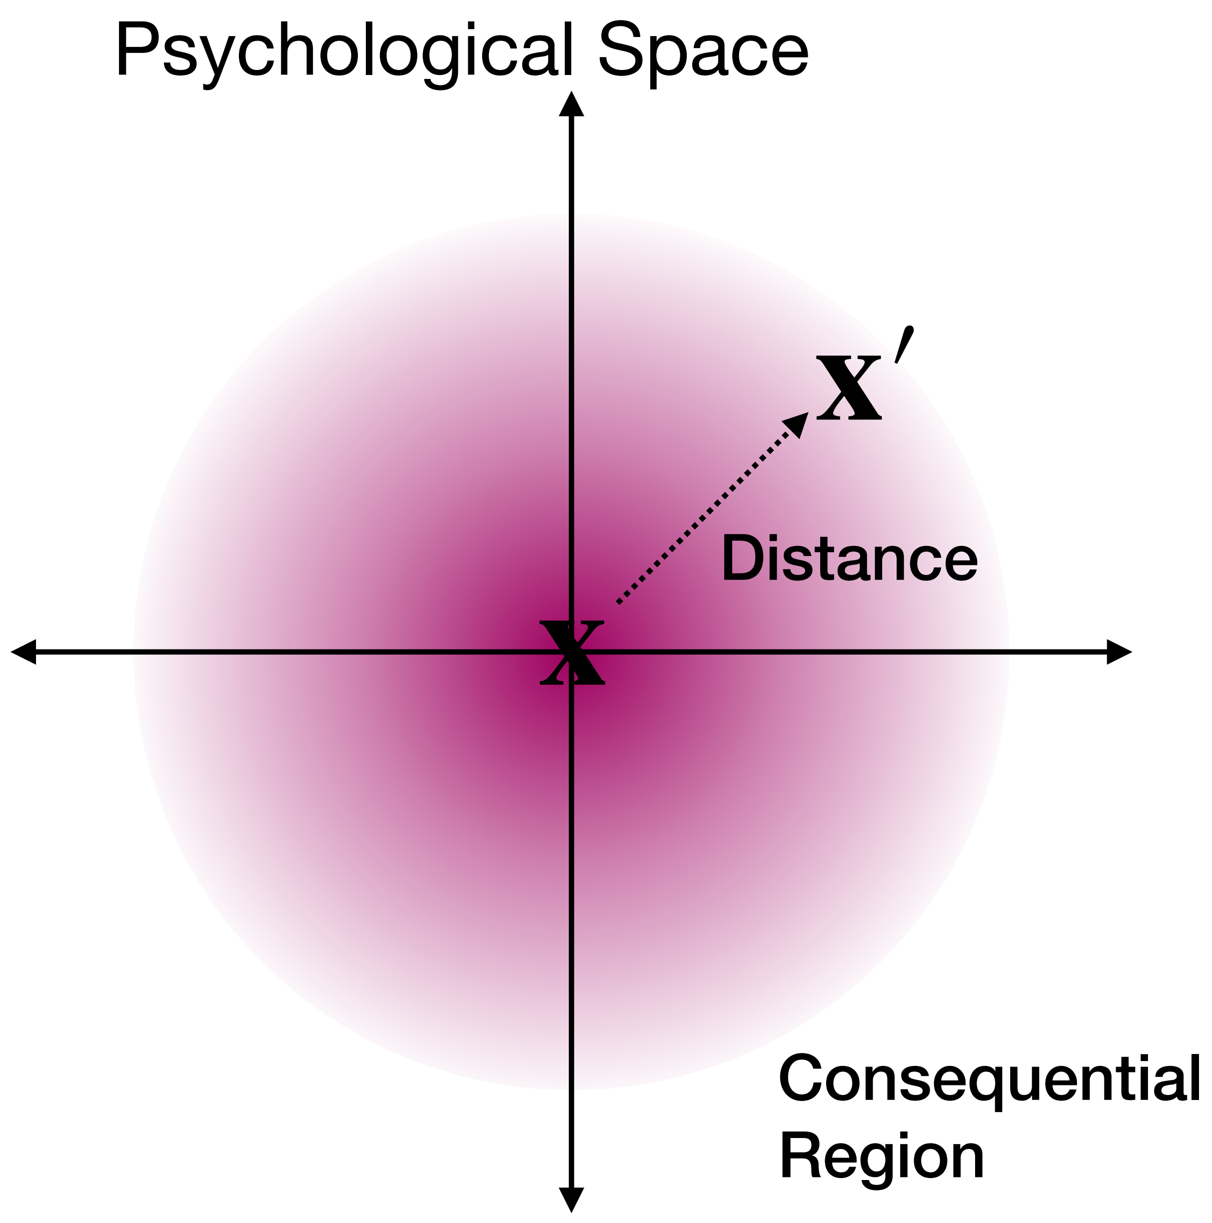

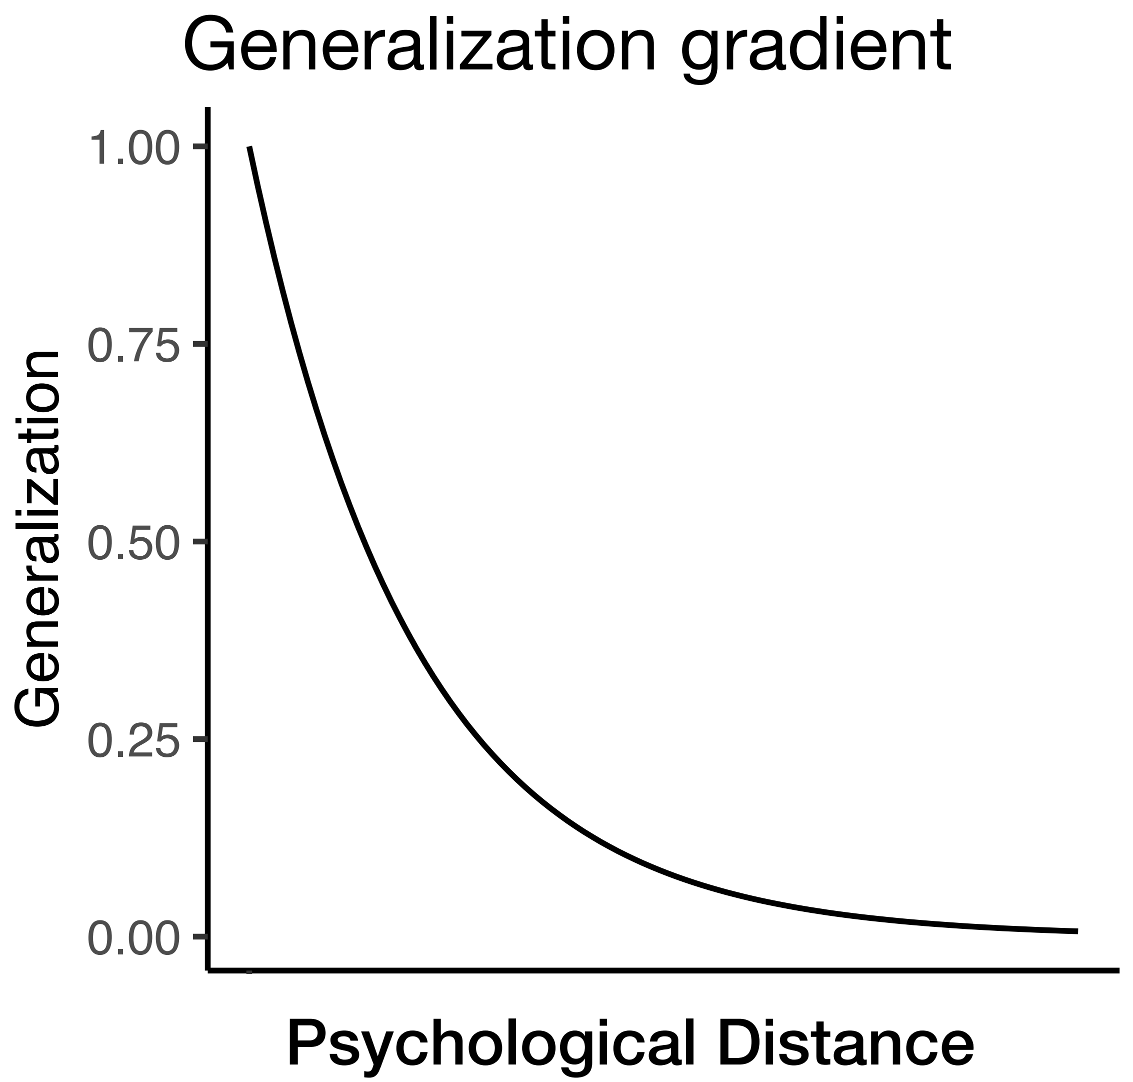

4. In Bayesian concept learning, what is the size principle?

4. In Bayesian concept learning, what is the size principle?

Likelihood:

x'es generated randomly

$$p(x|h) = \begin{cases} 1 & \text{if } x \in h \\ 0 & \text{otherwise} \end{cases} \quad \text{[weak sampling].}$$

x'es generated to be positive examples

$$p(x|h) = \begin{cases} \frac{1}{|h|} & \text{if } x \in h \\ 0 & \text{otherwise} \end{cases} \quad \text{[strong sampling],}$$

Bayesian size principle: under *strong sampling*, smaller h'es (consistent with the data) are more likely

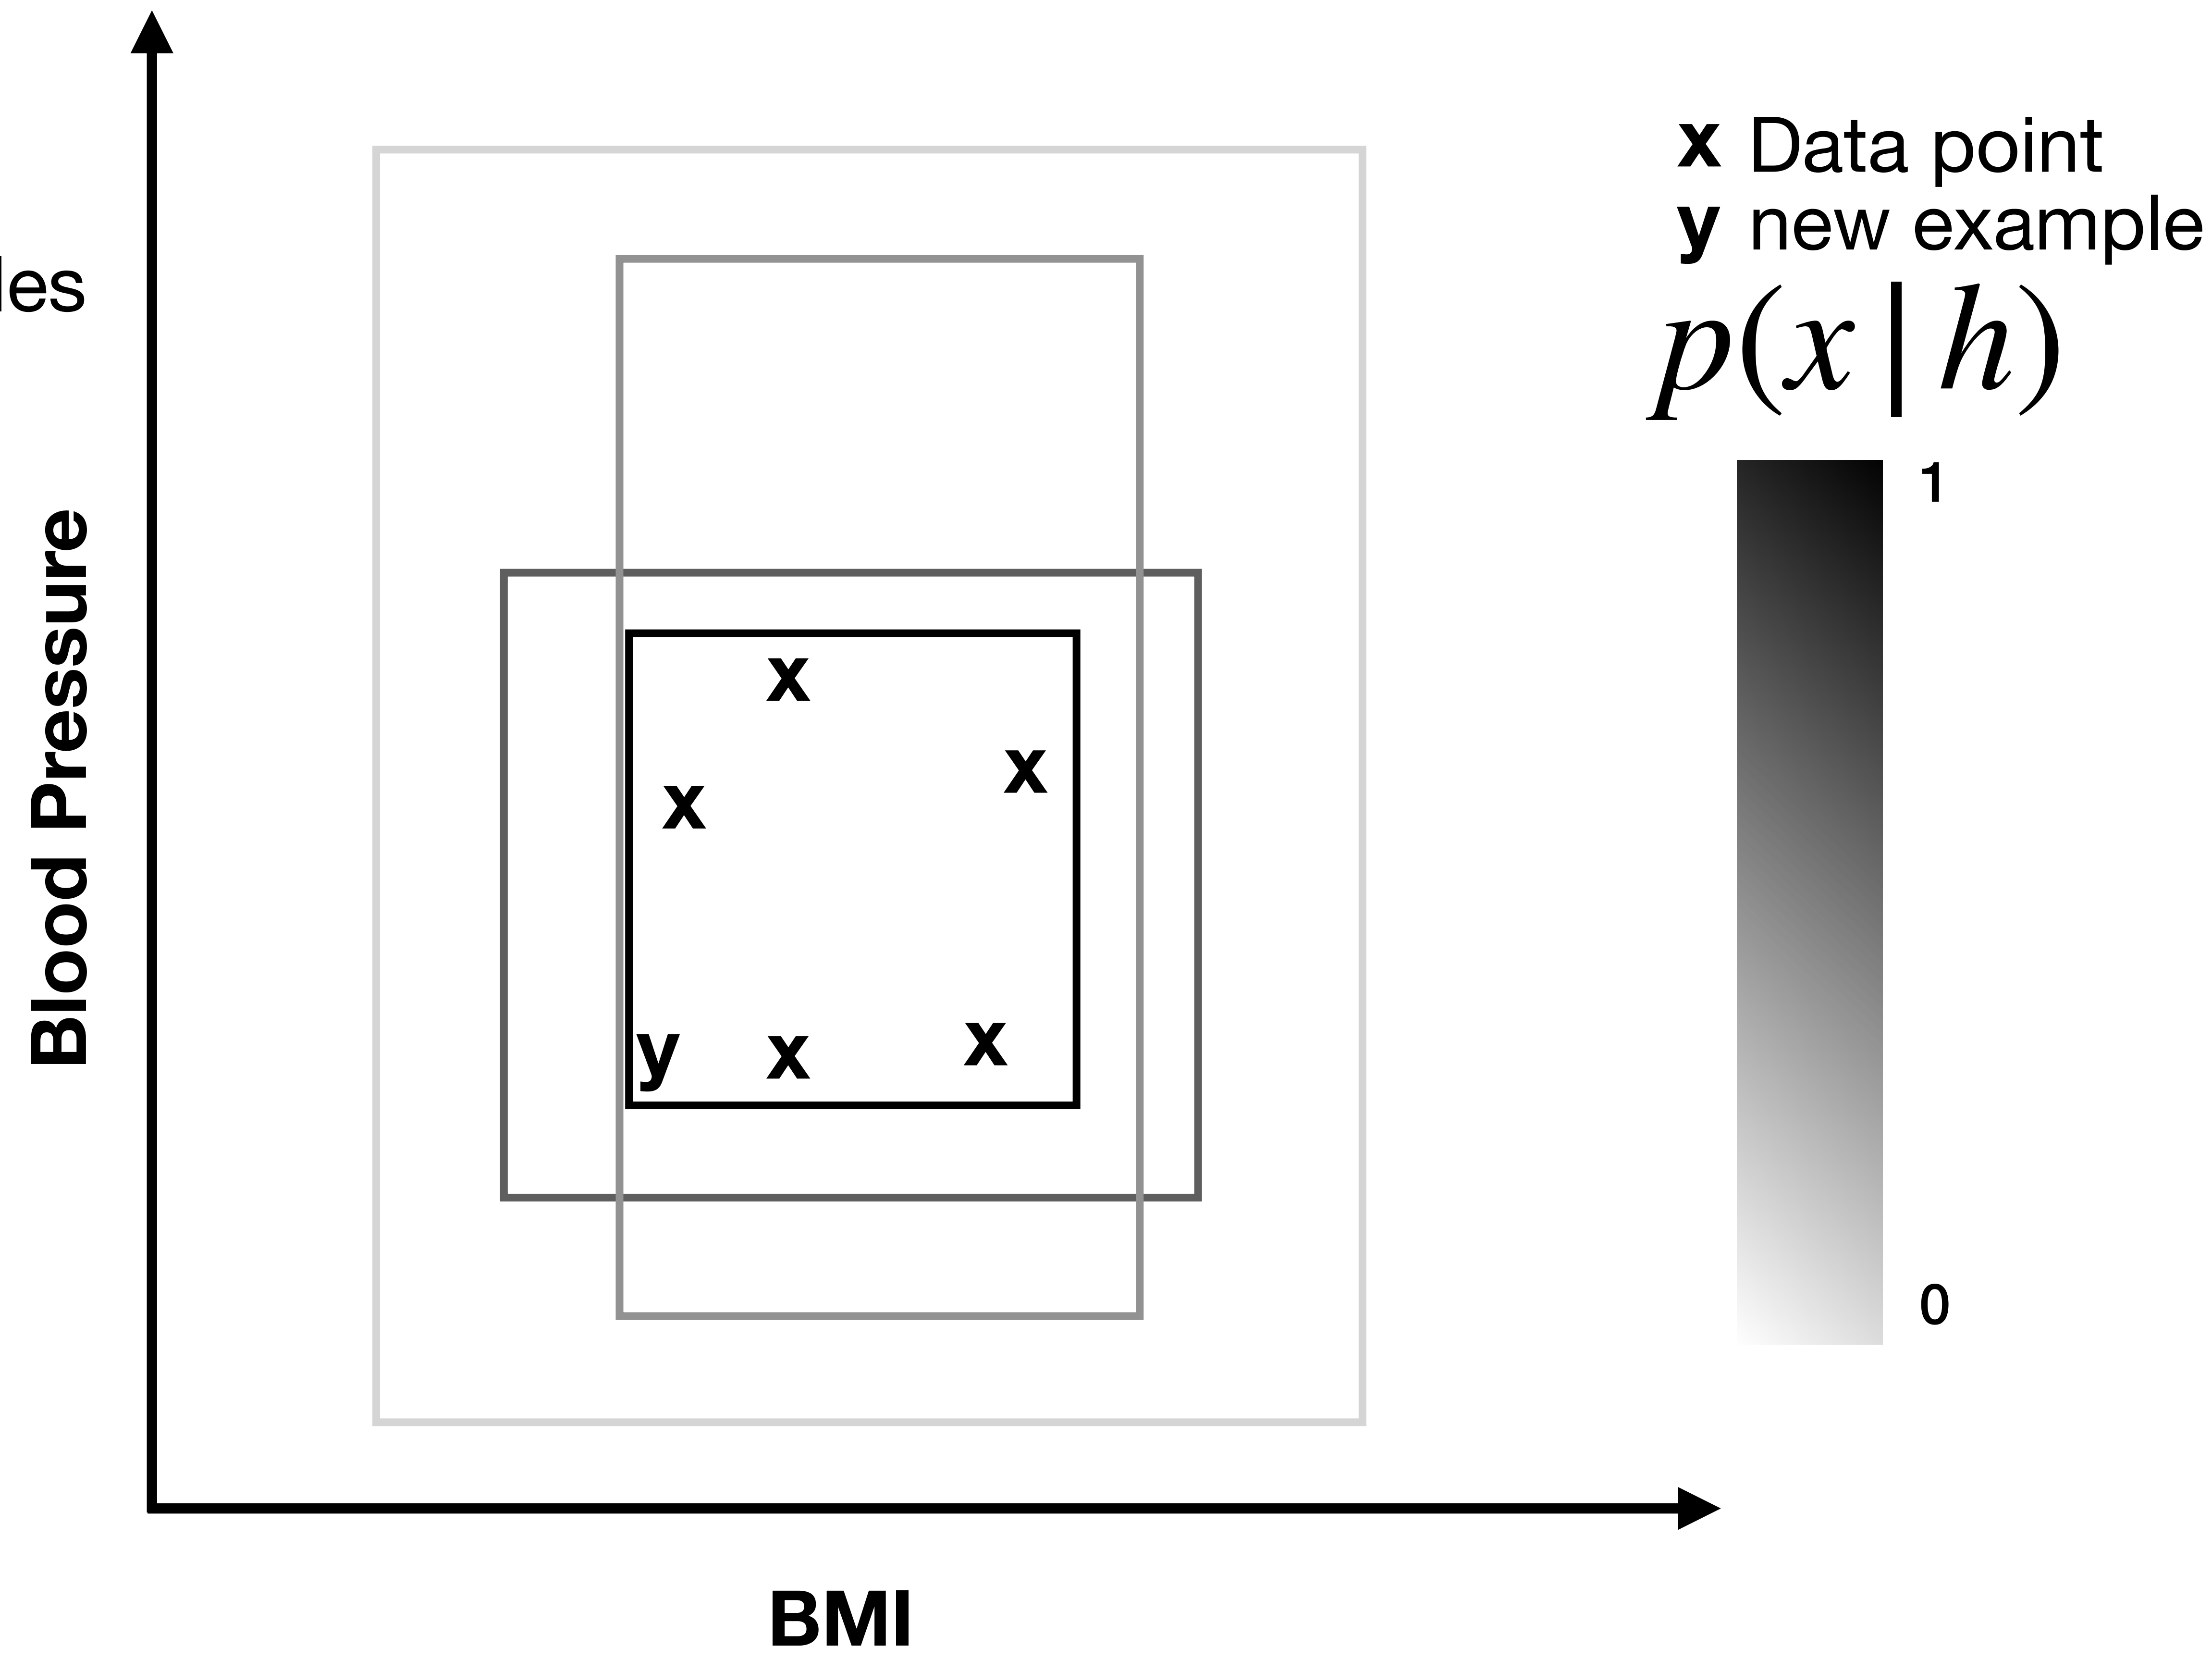

Tenenbaum (*NIPS* 1999)
Tenenbaum & Griffiths (*BBS* 2001)

4. In Bayesian concept learning, what is the size principle?

- To summarize….
- The probability of y being in the same category of x is based on summing over all hypotheses consistent with the data

$$p(y \in C|x) = \sum_{h: y \in h} p(h|x).$$

- Where narrower hypotheses are favored under strong sampling

$$p(h|x) = \frac{p(x|h)p(h)}{p(x)}$$

$$p(x|h) = \begin{cases} \frac{1}{|h|} & \text{if } x \in h \\ 0 & \text{otherwise} \end{cases} \qquad \text{[strong sampling]},$$

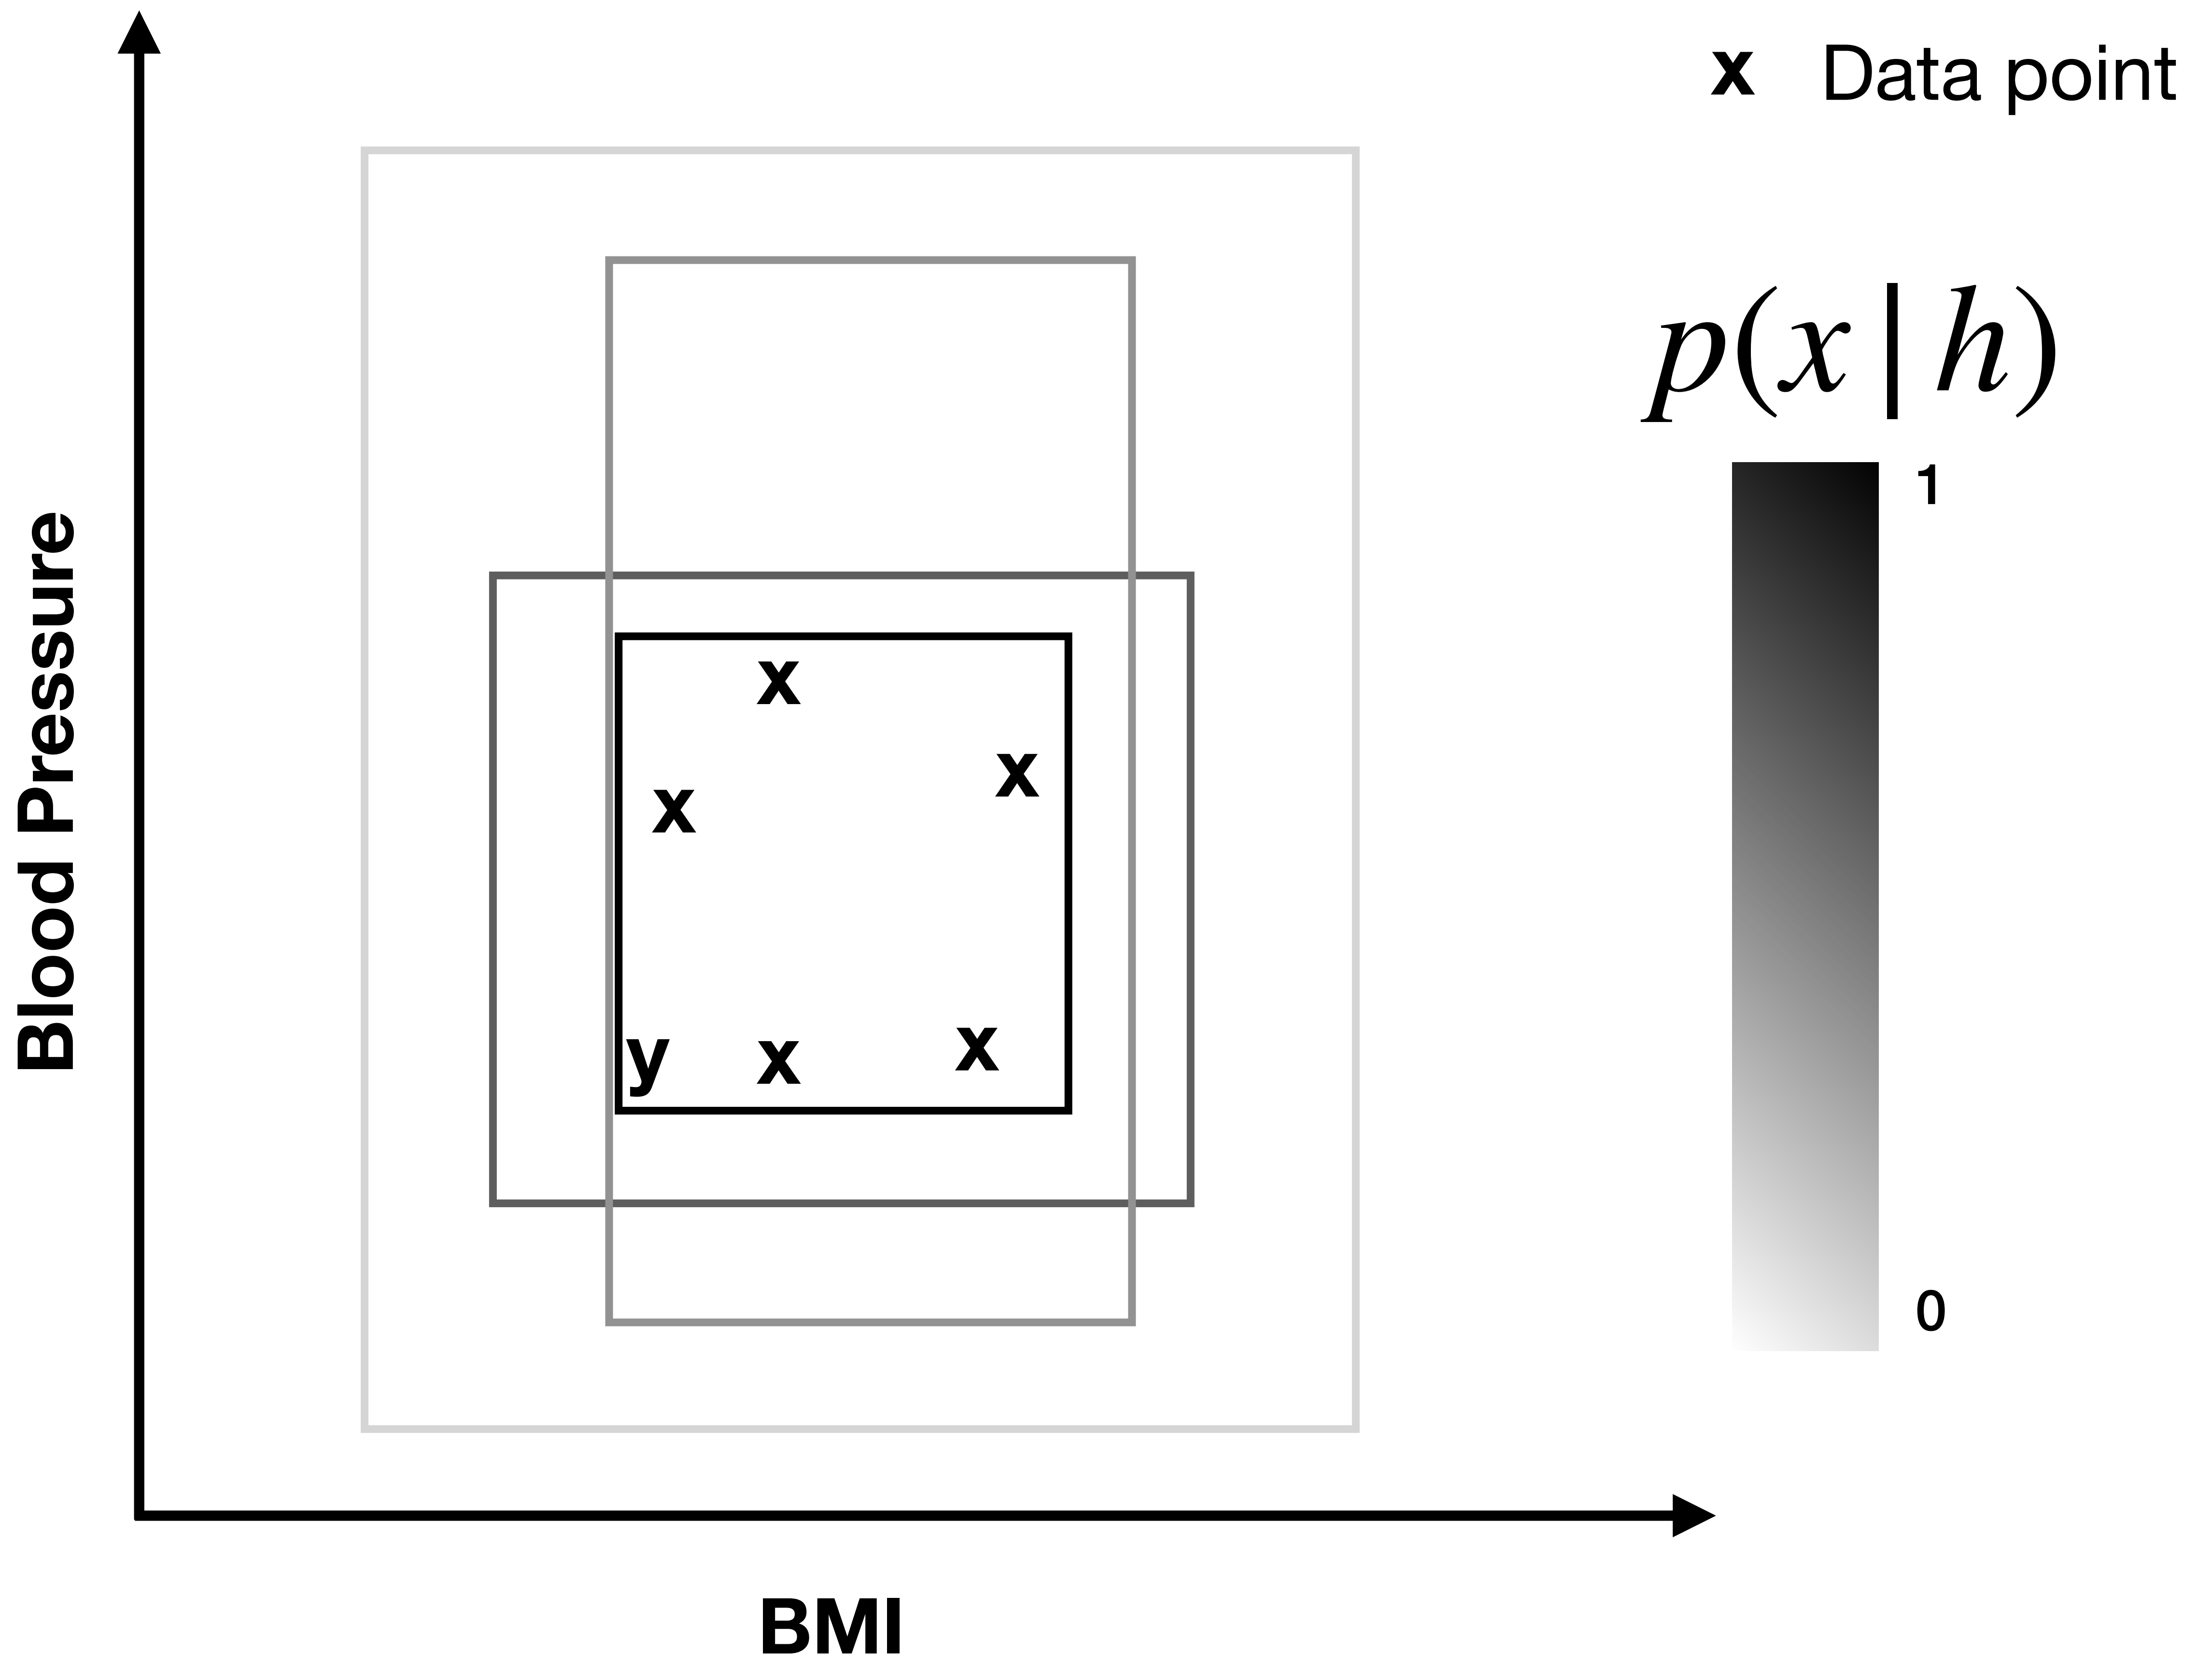

Tenenbaum (*NIPS* 1999)
Tenenbaum & Griffiths (*BBS* 2001)

Bayesian Concept Learning Extends Shepard's Law of Generalization to Multiple Examples

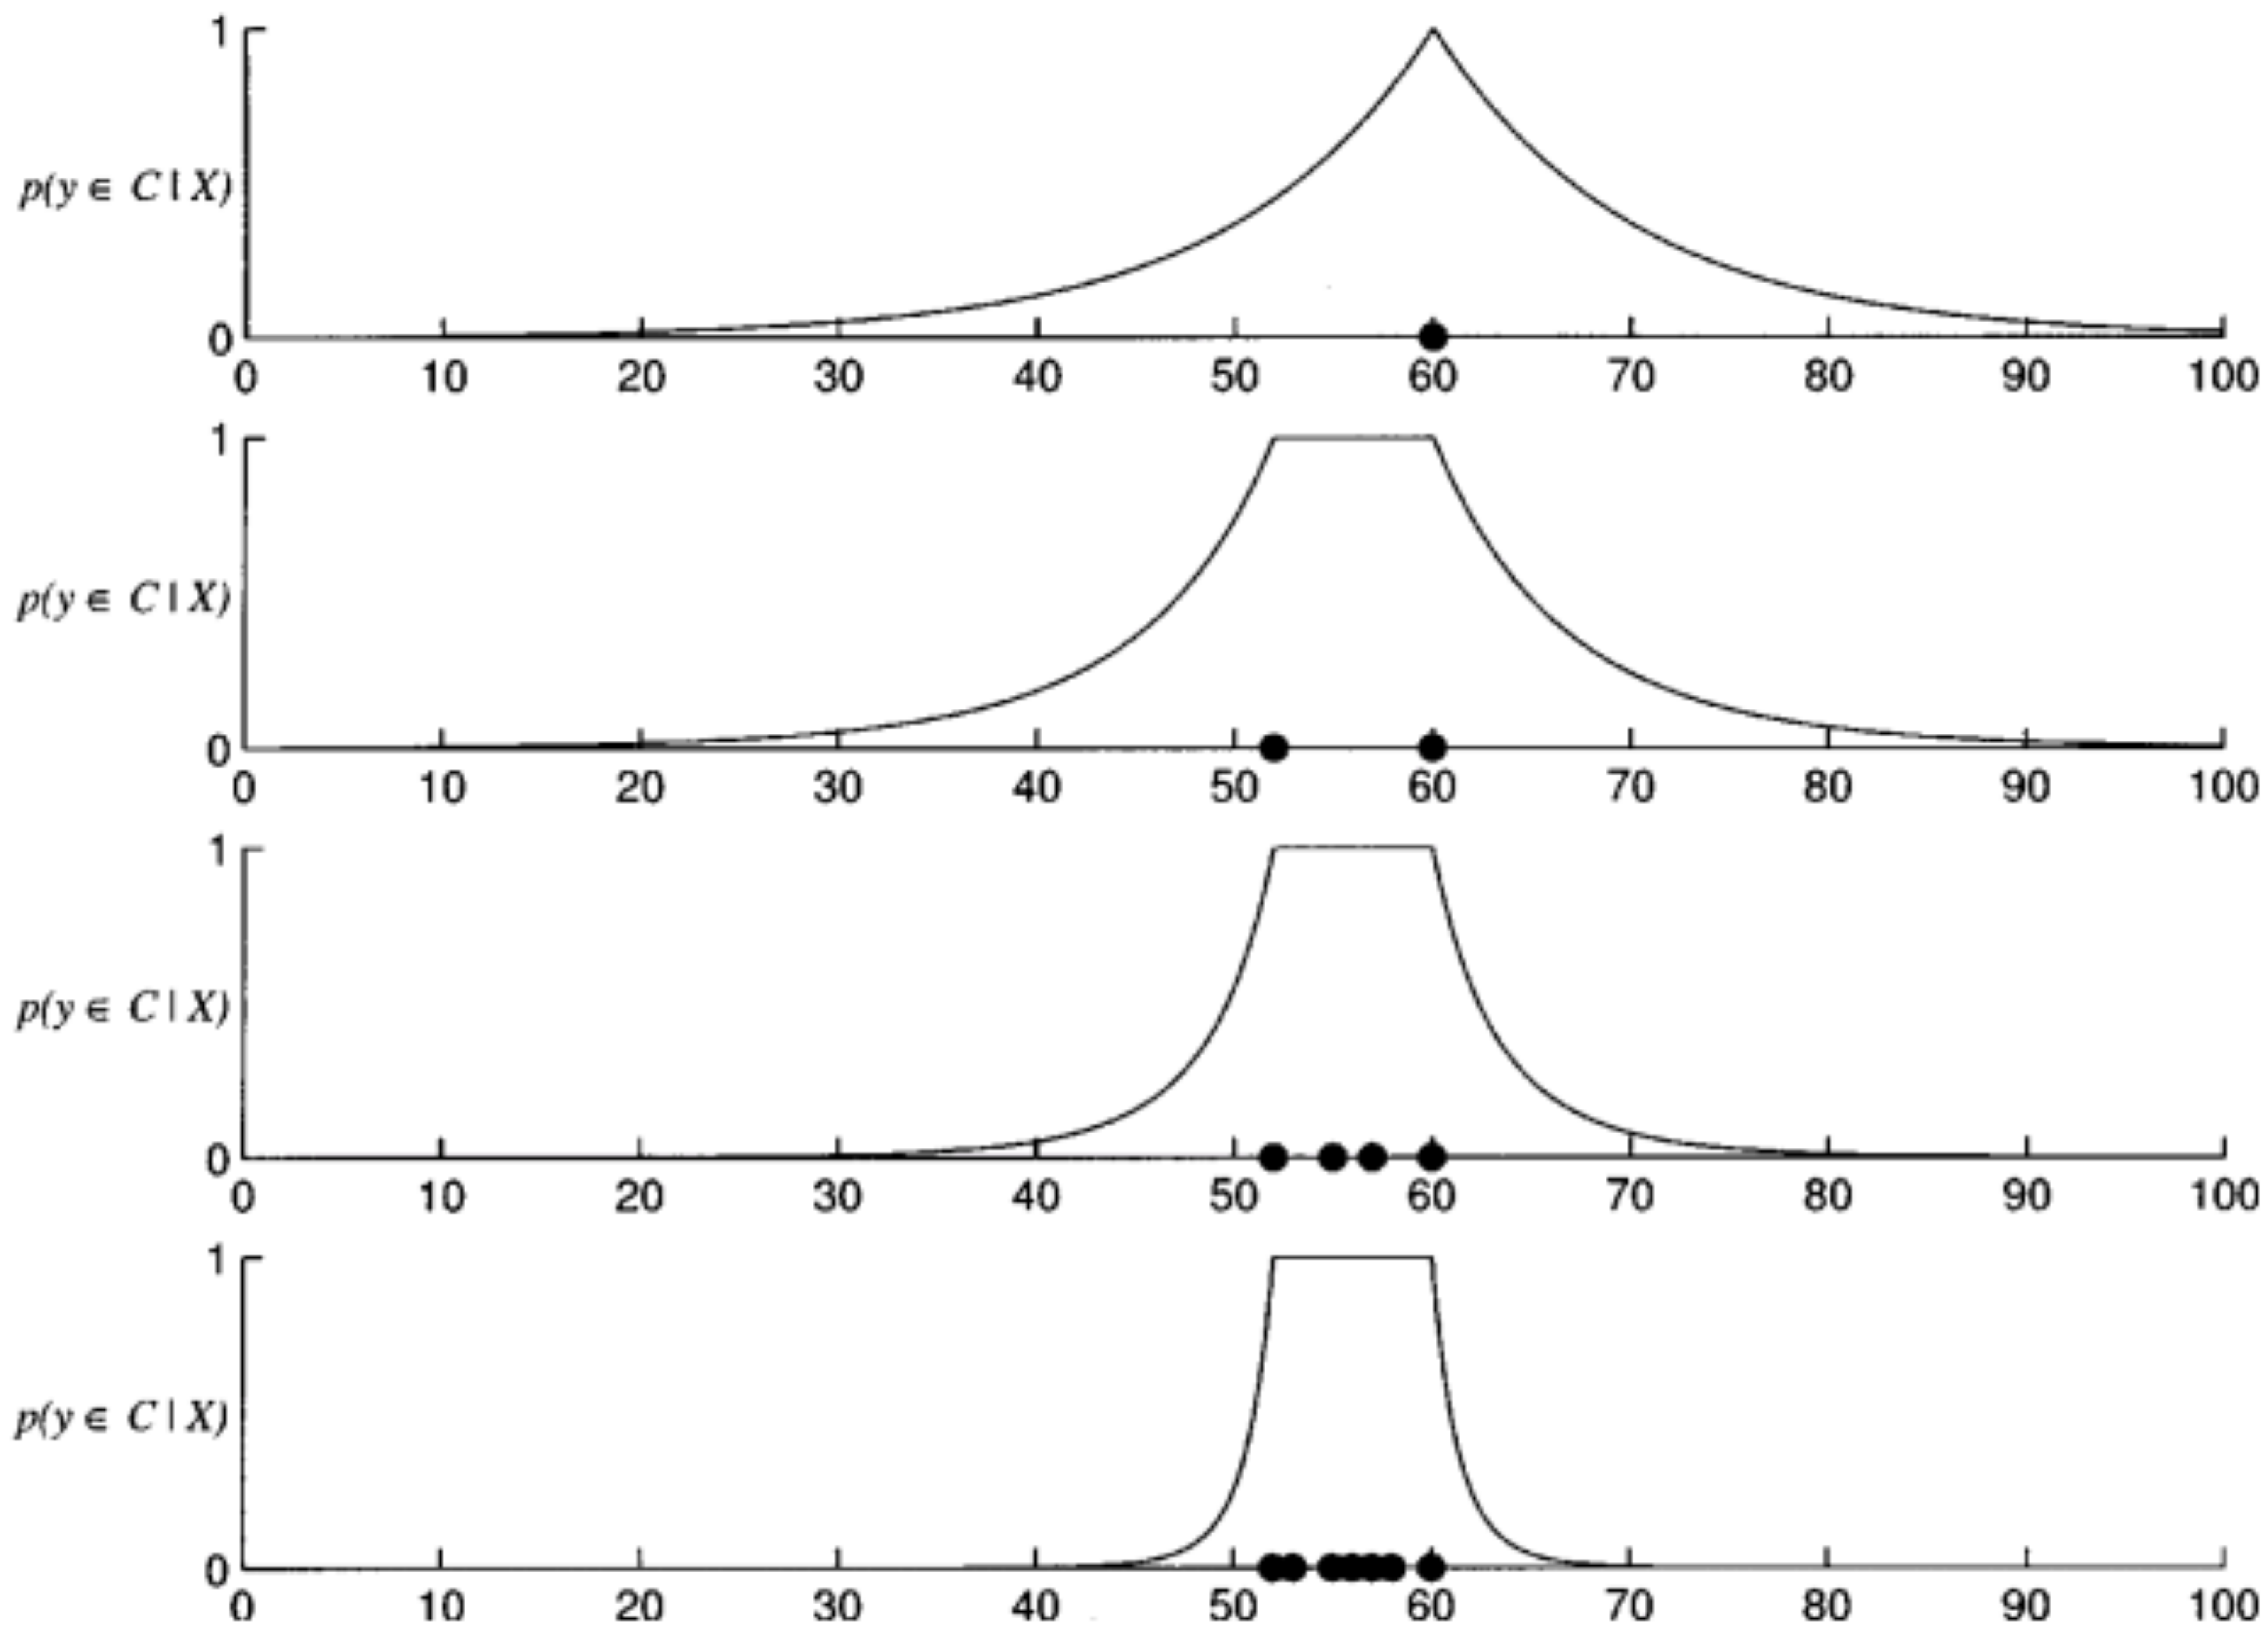

Figure 3. The effect of the number of examples on Bayesian generalization (under the assumptions of strong sampling and an Erlang prior, $\mu = 10$). Filled circles indicate examples. The first curve is the gradient of generalization with a single example, for the purpose of comparison. The remaining graphs show that the range of generalization decreases as a function of the number of examples.

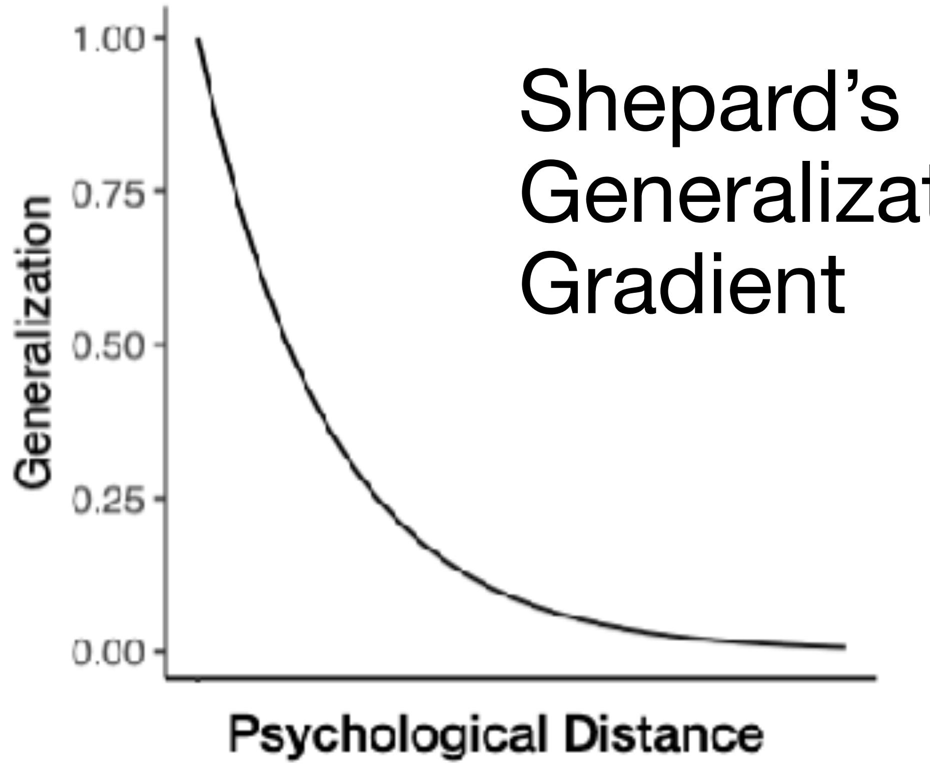

Shepard's Generalization Gradient

Bayesian Concept Learning Extends Shepard's Law of Generalization to Multiple Examples

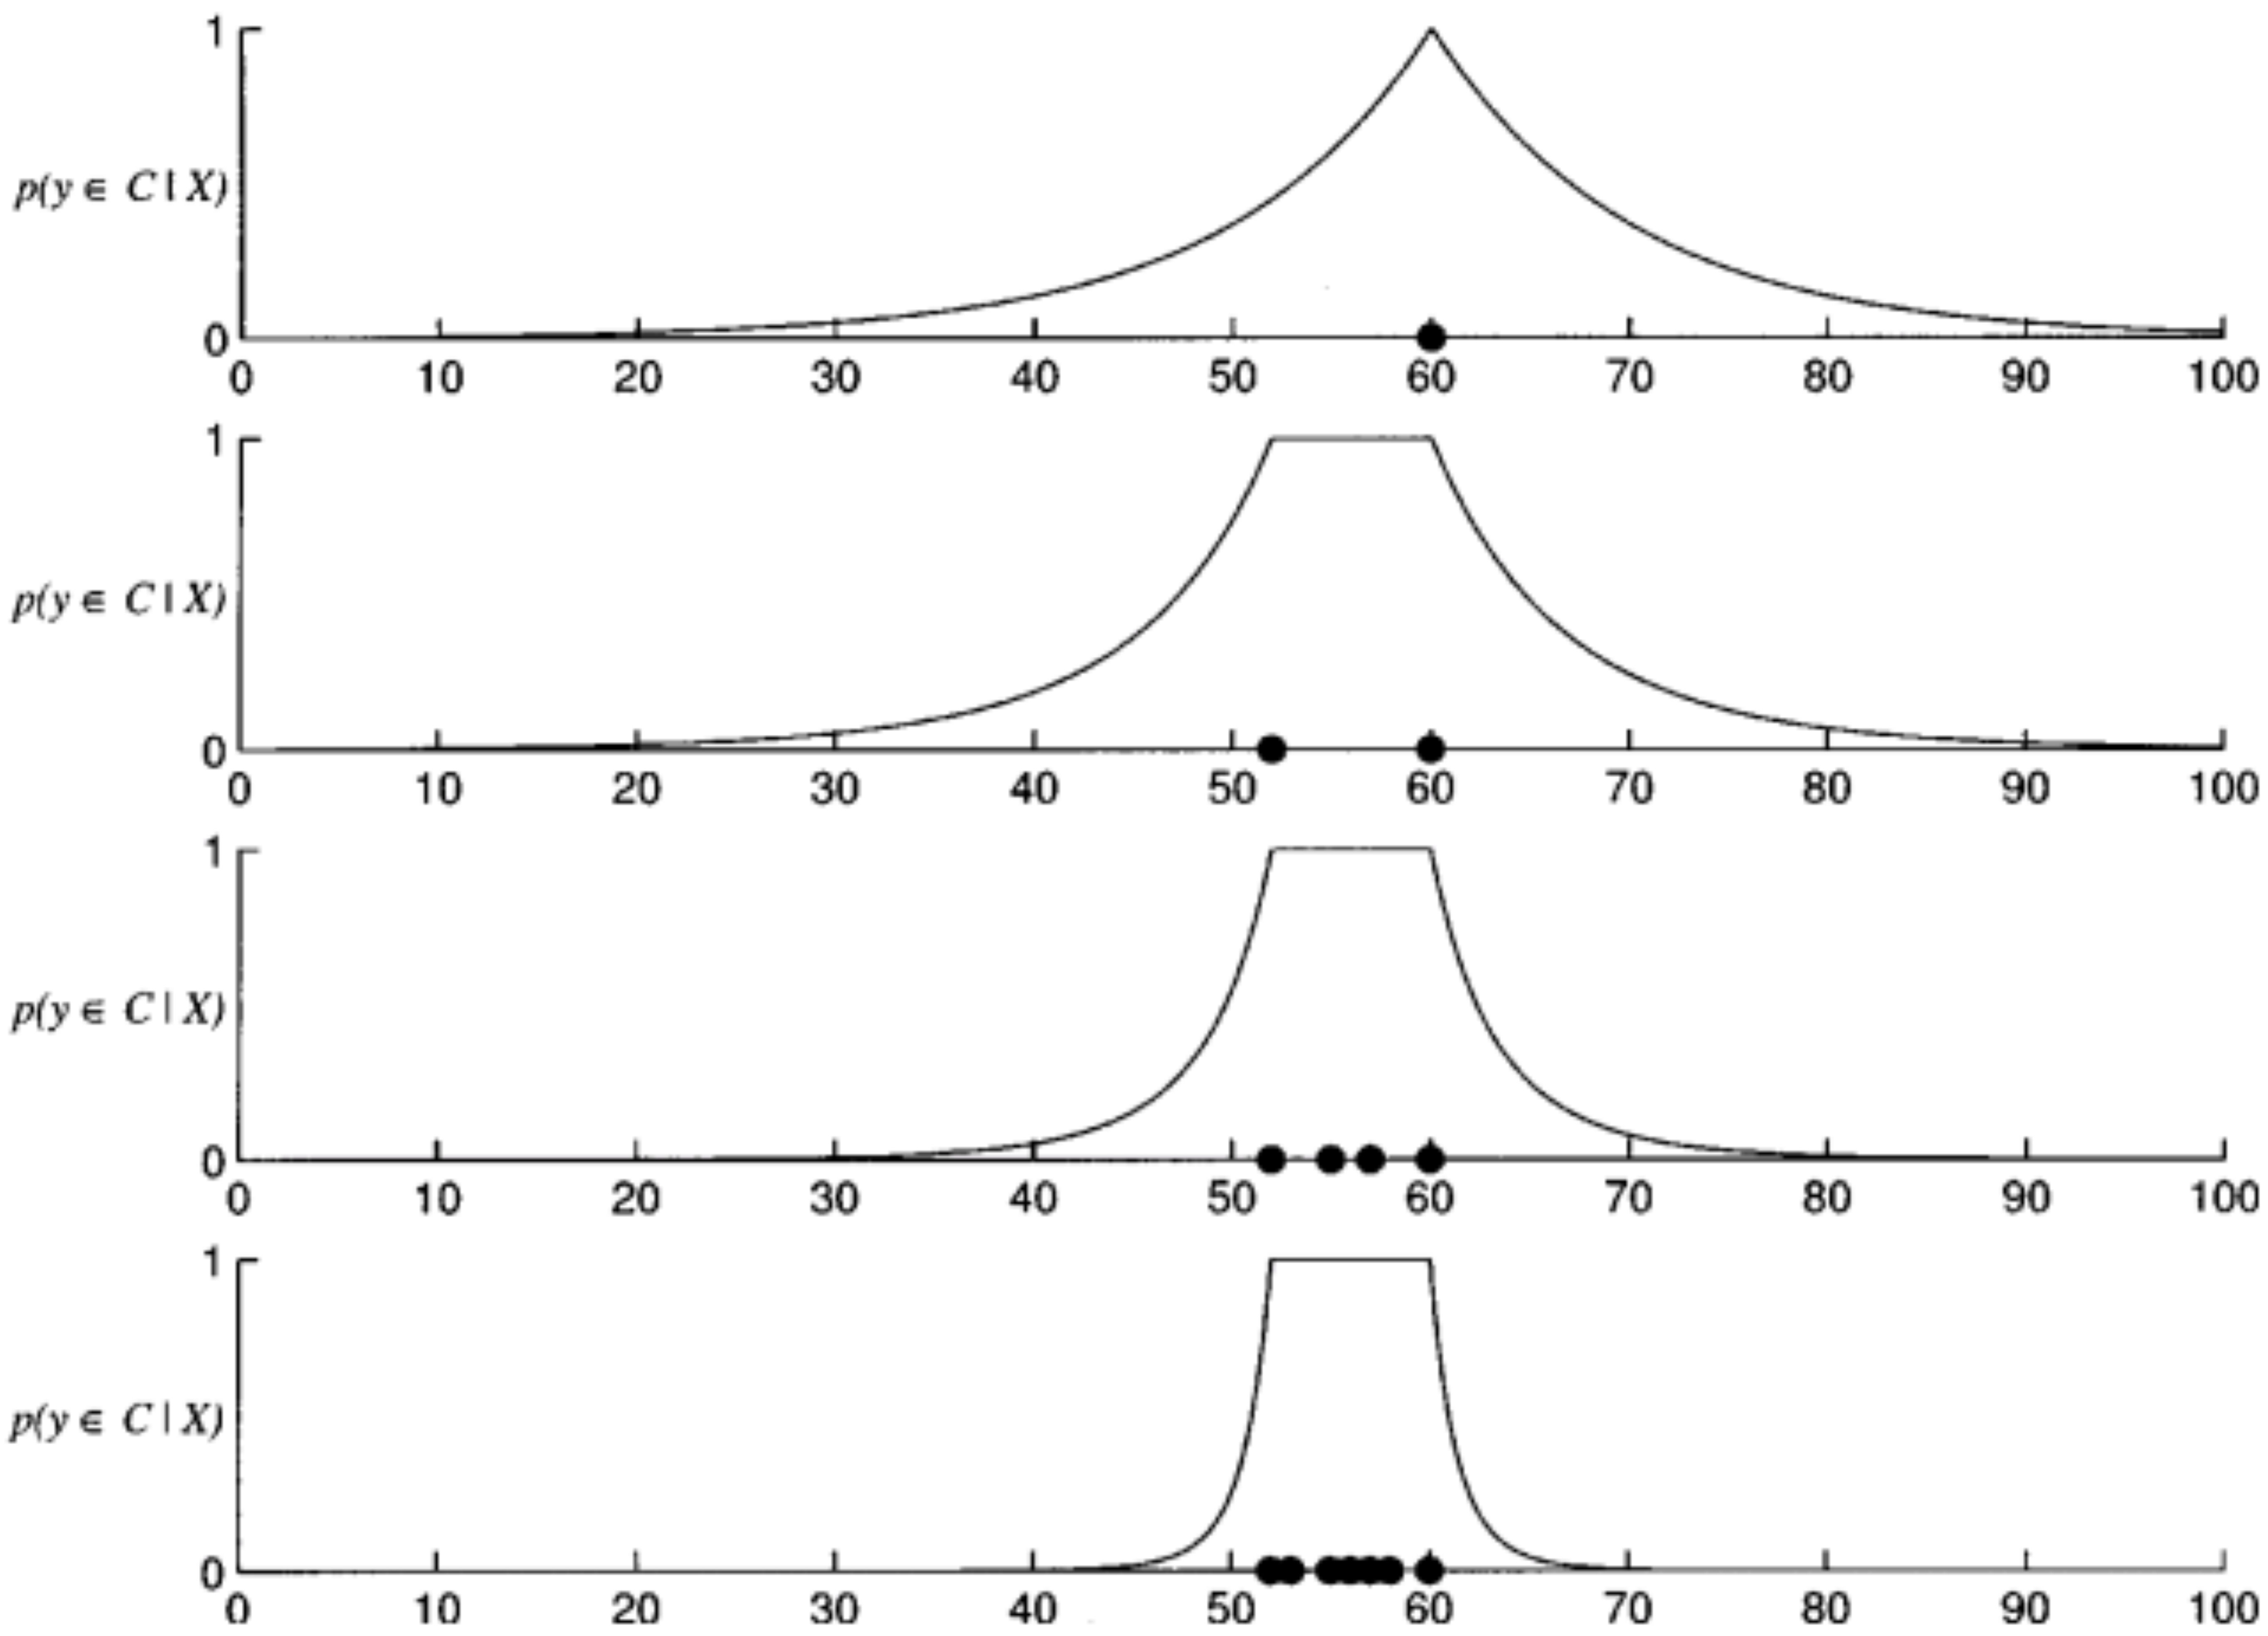

Figure 3. The effect of the number of examples on Bayesian generalization (under the assumptions of strong sampling and an Erlang prior, $\mu = 10$). Filled circles indicate examples. The first curve is the gradient of generalization with a single example, for the purpose of comparison. The remaining graphs show that the range of generalization decreases as a function of the number of examples.

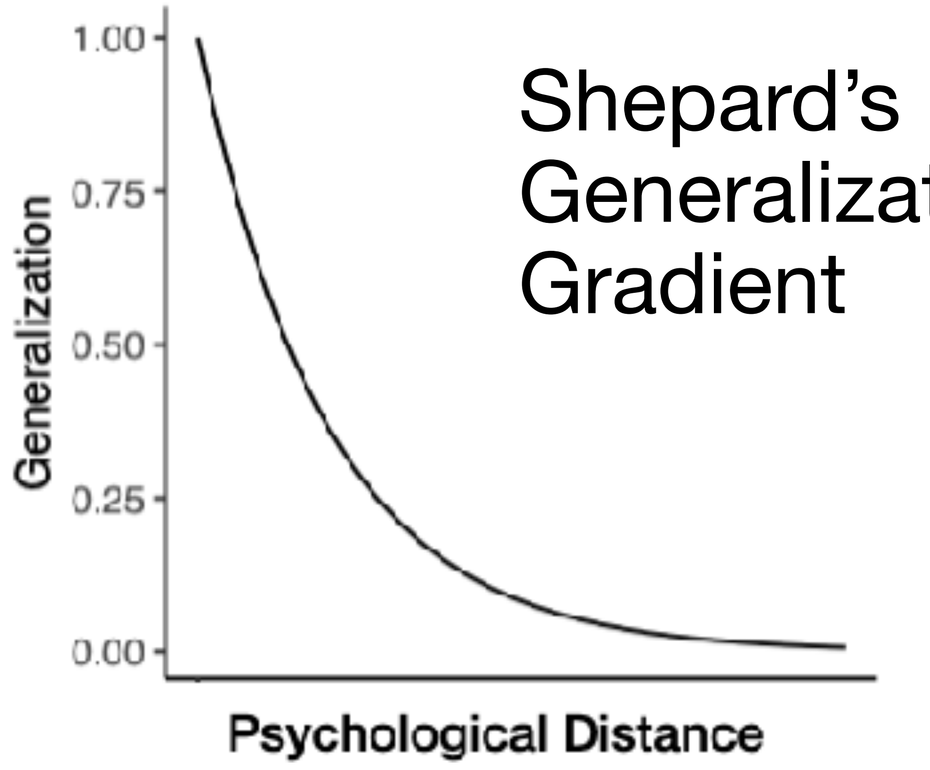

Range of generalization decreases with more examples

more examples = less uncertainty about the extent of consequential region

10. How is function learning useful in reinforcement learning settings?

- Function learning is the basis for **value approximation** in RL

- *Value approximation*: Learning an (implicit) value function mapping states to value expectations
  $V(s') = f(s')$

- This allows for more flexible generalization and adaptive behavior in settings with vast or continuous state spaces

Silver et al., (*Nature* 2016)

10. How is function learning useful in reinforcement learning settings?

- Function learning is the basis for **value approximation** in RL

- *Value approximation*: Learning an (implicit) value function mapping states to value expectations
  $$V(s') = f(s')$$

- This allows for more flexible generalization and adaptive behavior in settings with vast or continuous state spaces

Value network

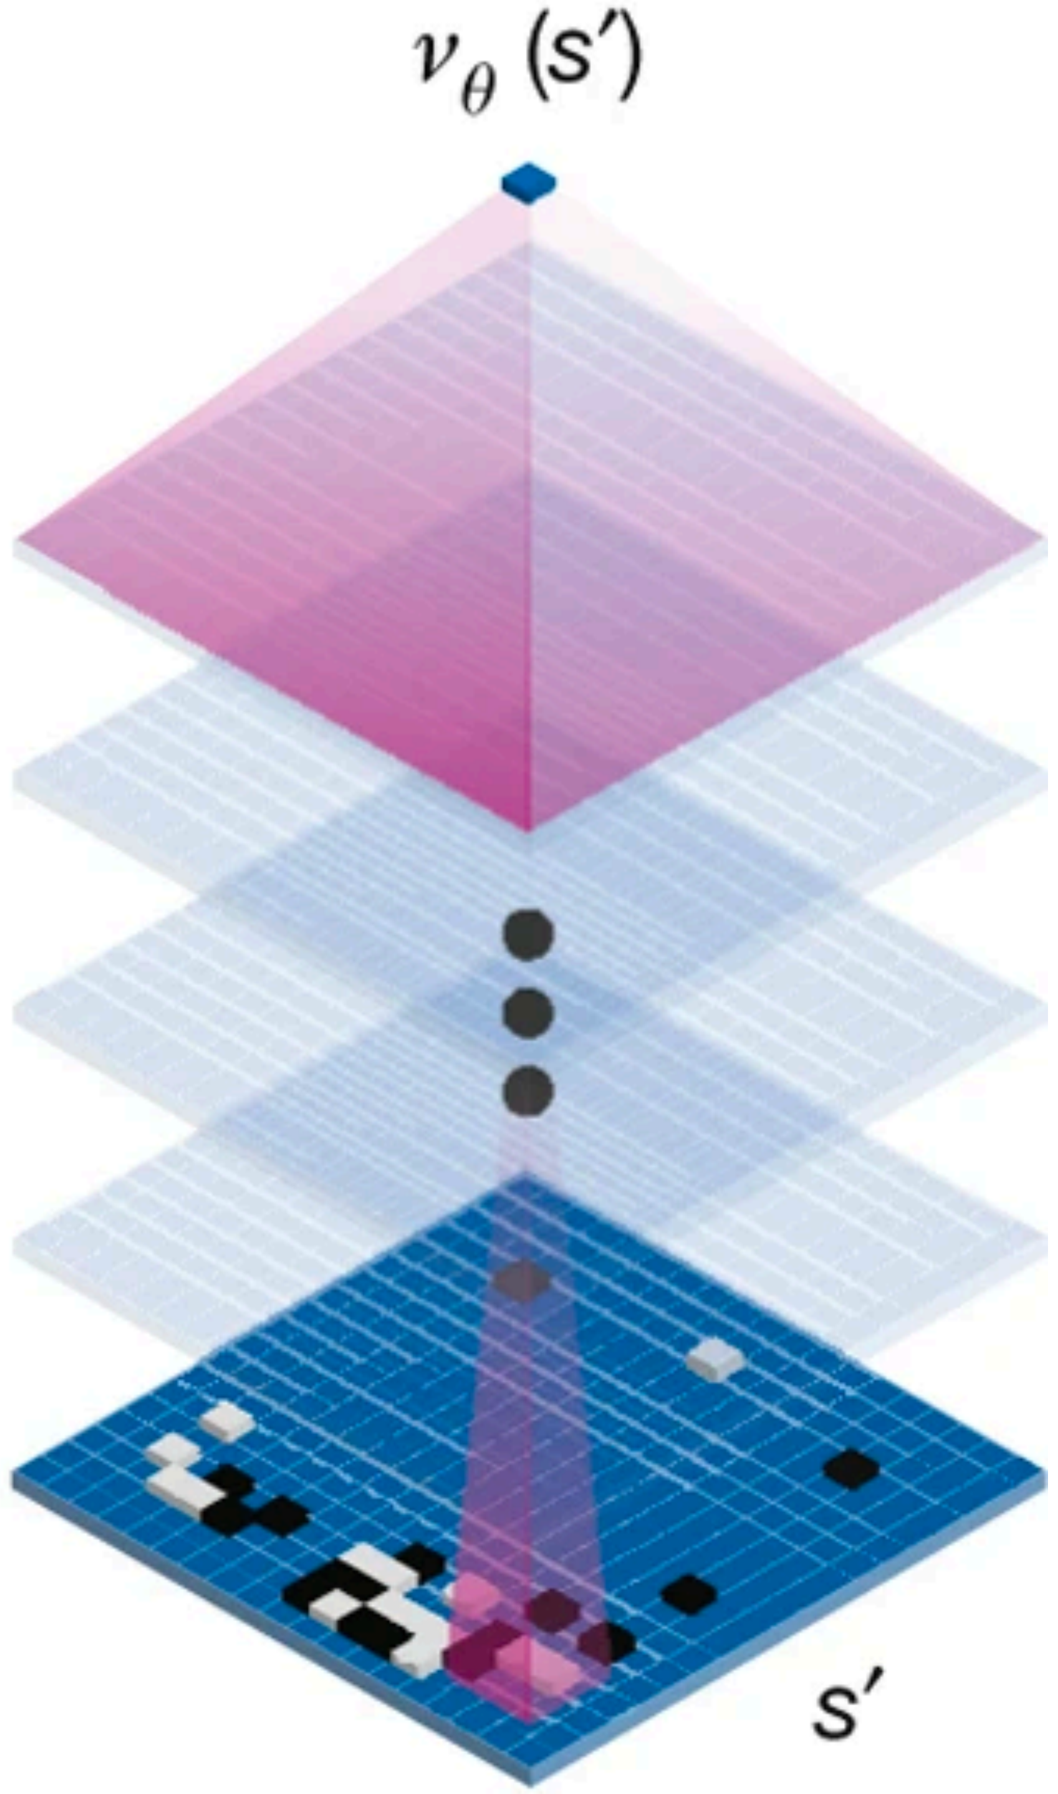

Silver et al., (*Nature* 2016)

Exam prep

Google doc