Machine Learning
Lecture 1

1.

Machine learning is everywhere

- Search, content recommendation, image/scene analysis, machine translation, dialogue systems, automated assistants, game playing, sciences (biology, chemistry, etc), …

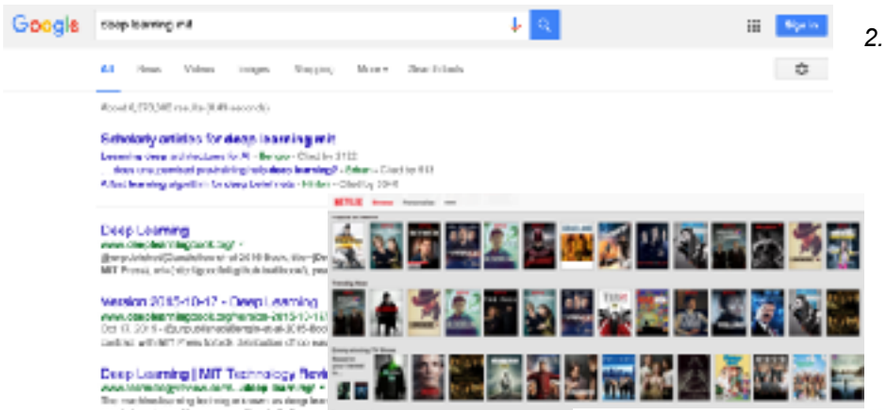

2.

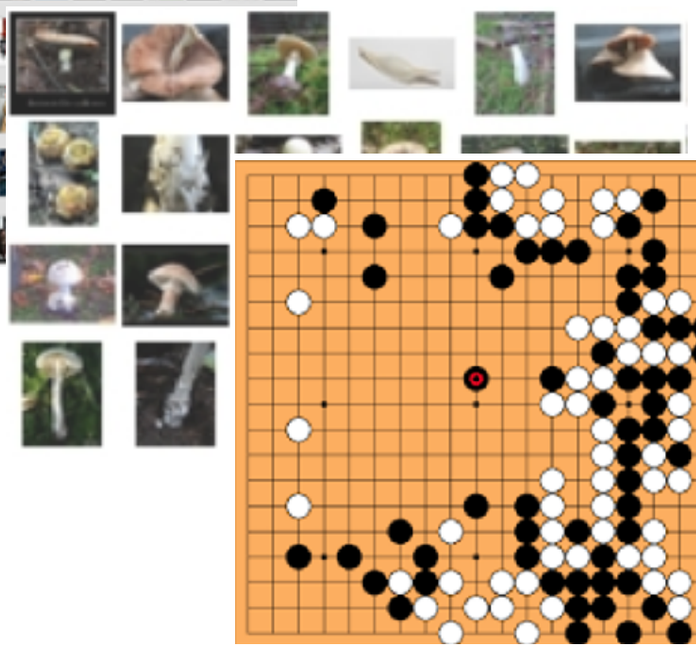

3.

4.

5.

2,4. Google and the Google logo are registered trademarks of Google LLC

3. Screenshot of Netflix (n.d.), (c) Netflix, Inc.

5. © Senseis Library

Machine learning: what is it?

- A brief definition

  Machine learning as a discipline aims to design, understand and apply computer programs that learn from experience (i.e., data) for the purpose of modeling, prediction, or control

Prediction problems

- About future events

- Also collision avoidance, monitoring, medical risk, etc.

Prediction problems

- About properties we don’t yet know

6. © Black Bear Pictures

would I like this movie?

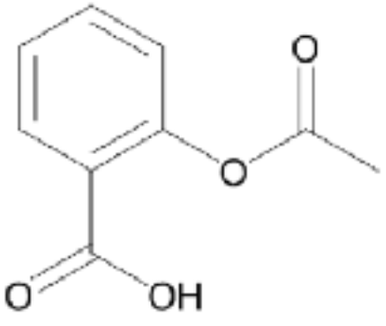

soluble in water?

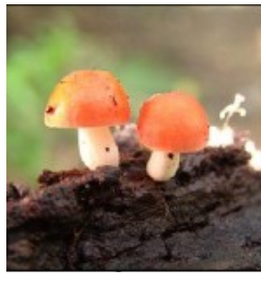

7. , © Neural Information Processing Systems Foundation, Inc.

what is the image about?

“ML is very cool”

what is it in Spanish?

Example: supervised learning

- It is easier to express tasks in terms of examples of what you want (rather than how to solve them)
- E.g., image classification (1K categories)

Image	**Category**
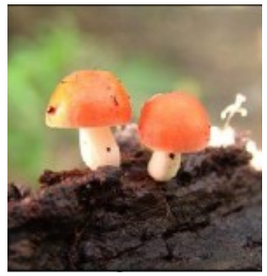	mushroom
. © Neural Information Processing Systems Foundation, Inc.	
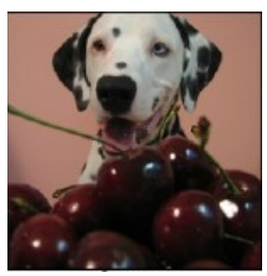	cherry
8. © Geoffrey Hinton, University of Toronto.	
...	...

Example: supervised learning

- It is easier to express tasks in terms of examples of what you want (rather than how to solve them)
- E.g., image classification (1K categories)

Image	**Category**
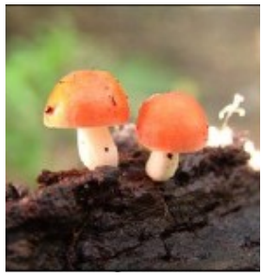	mushroom
© Neural Information Processing Systems Foundation, Inc	
	cherry
© Geoffrey Hinton, University of Toronto.	
...	...

- Rather than specify the solution directly (hard), we automate the process of finding one based on examples

Example: supervised learning

- It is easier to express tasks in terms of examples of what you want (rather than how to solve them)
- E.g., image classification (1K categories)

Image		**Category**
$h\left(\right.$ 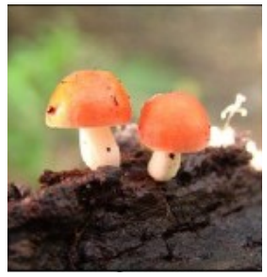 $; \theta \left. \right)$	$=$	mushroom
		cherry
...		...

- Rather than specify the solution directly (hard), we automate the process of finding one based on examples

Example: supervised learning

- It is easier to express tasks in terms of examples of what you want (rather than how to solve them)
- No limit to what you can learn to predict…

English		**Spanish**
$h($ Is it real? ; $\theta)$	$=$	¿Es real?
Will it continue?		¿Continuará?
For how long?		¿Por cuanto tiempo?
…		…

- Already in production for some language pairs (Google)

A concrete example

- Learning to predict preferences from just a little data…

11.© 2020 Sony Pictures Digital Productions Inc.

10. © Twentieth Century Fox.

9. © 2020 The Studios At Paramount.

12.© Disney

© Black Bear Pictures

A concrete example

- Learning to predict preferences from just a little data…

11.© 2020 Sony Pictures Digital Productions Inc.

-1

10. © Twentieth Century Fox.

-1

9. © 2020 The Studios At Paramount.

+1

12.© Disney

+1

© Black Bear Pictures

?, ?, ….

A concrete example

- Learning to predict preferences from just a little data…

11.© 2020 Sony Pictures Digital Productions Inc.

10. © Twentieth Century Fox.

9. © 2020 The Studios At Paramount.

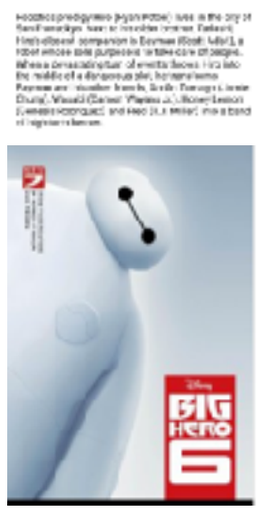

12.© Disney

© Black Bear Pictures

-1 -1 +1 +1 ?, ?, ….

comedy? top lead?

$$x^{(1)} = [1\ 0\ 0\ 1\ 1\ \ldots\ 0]^T \quad \text{(feature vector)}$$

action? romance? Spielberg?

Supervised learning

- Learning to predict preferences from just a little data…

11.© 2020 Sony Pictures Digital Productions Inc.

$x^{(1)}$

-1

10. © Twentieth Century Fox.

$x^{(2)}$

-1

9. © 2020 The Studios At Paramount.

$x^{(3)}$

+1

12.© Disney

$x^{(4)}$

+1

Training set

© Black Bear Pictures

$x^{(5)}, x^{(6)}, \ldots$

?, ?, ...

Test set

Supervised learning

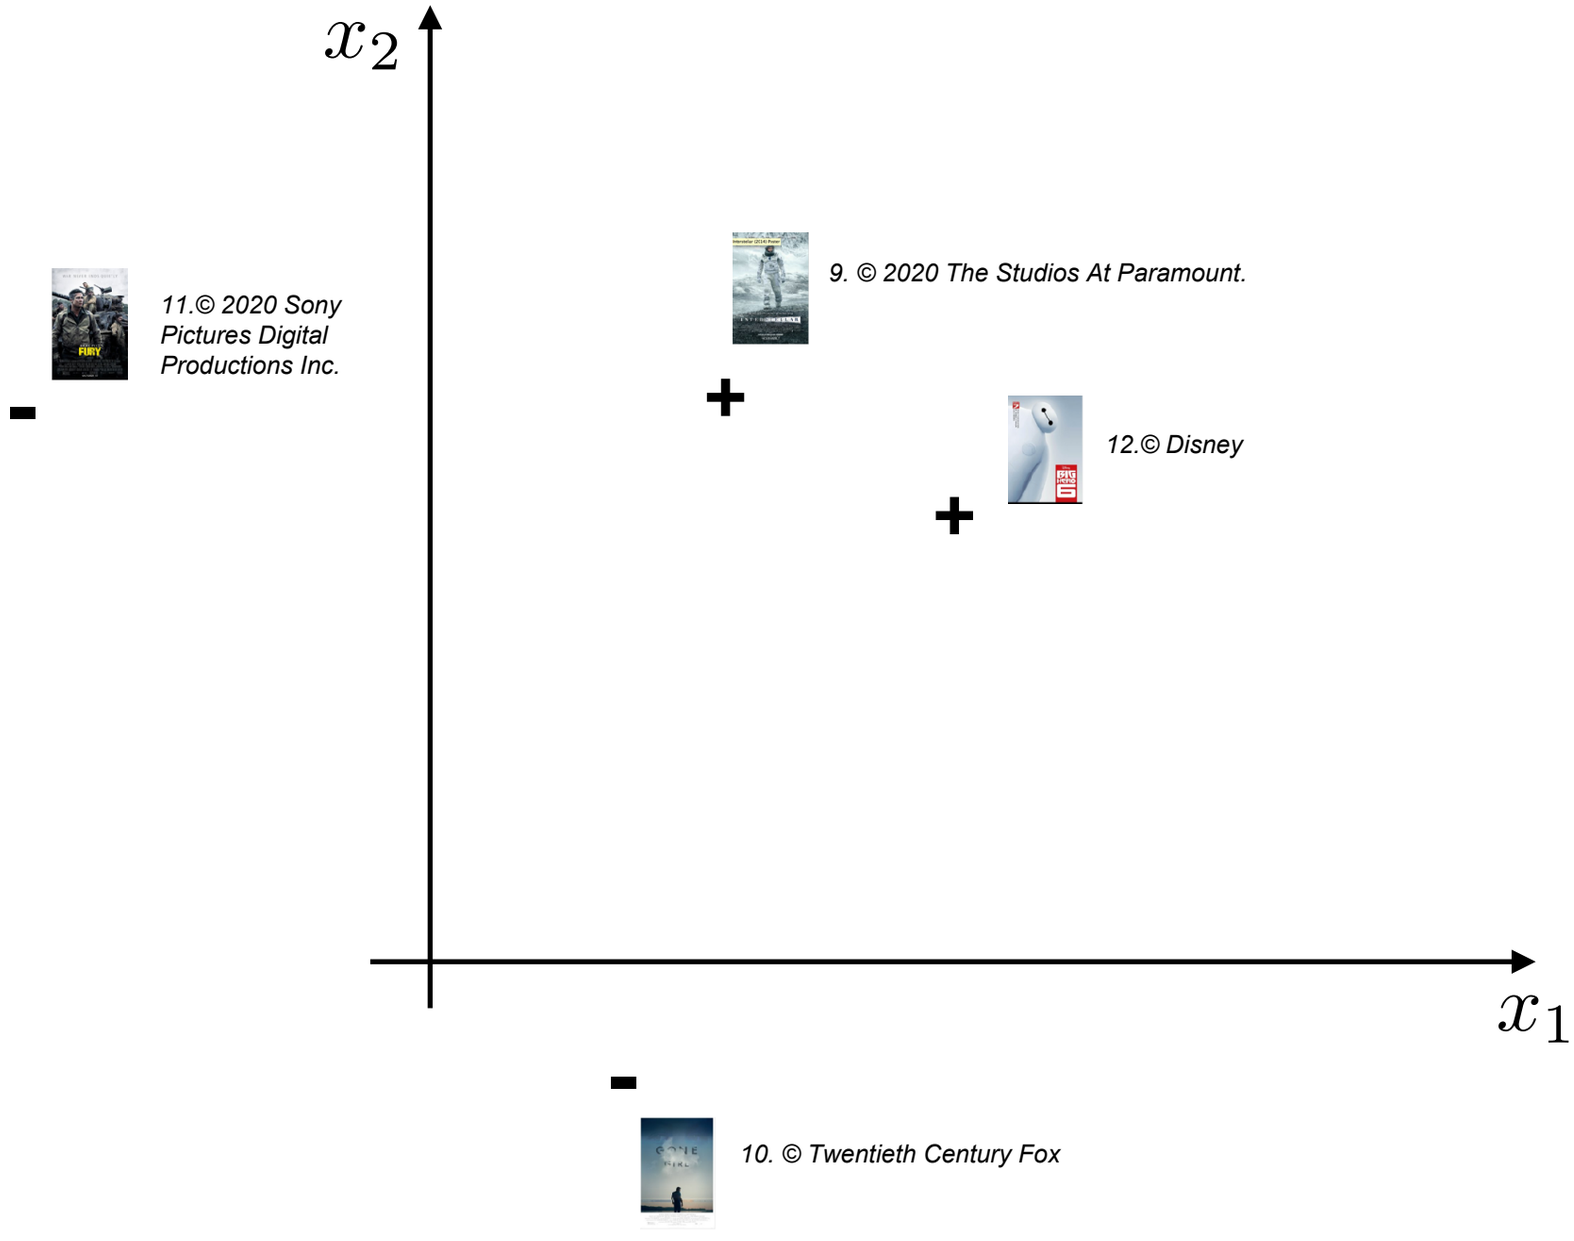

11.© 2020 Sony Pictures Digital Productions Inc.

9. © 2020 The Studios At Paramount.

12.© Disney

10. © Twentieth Century Fox

Supervised learning: training set

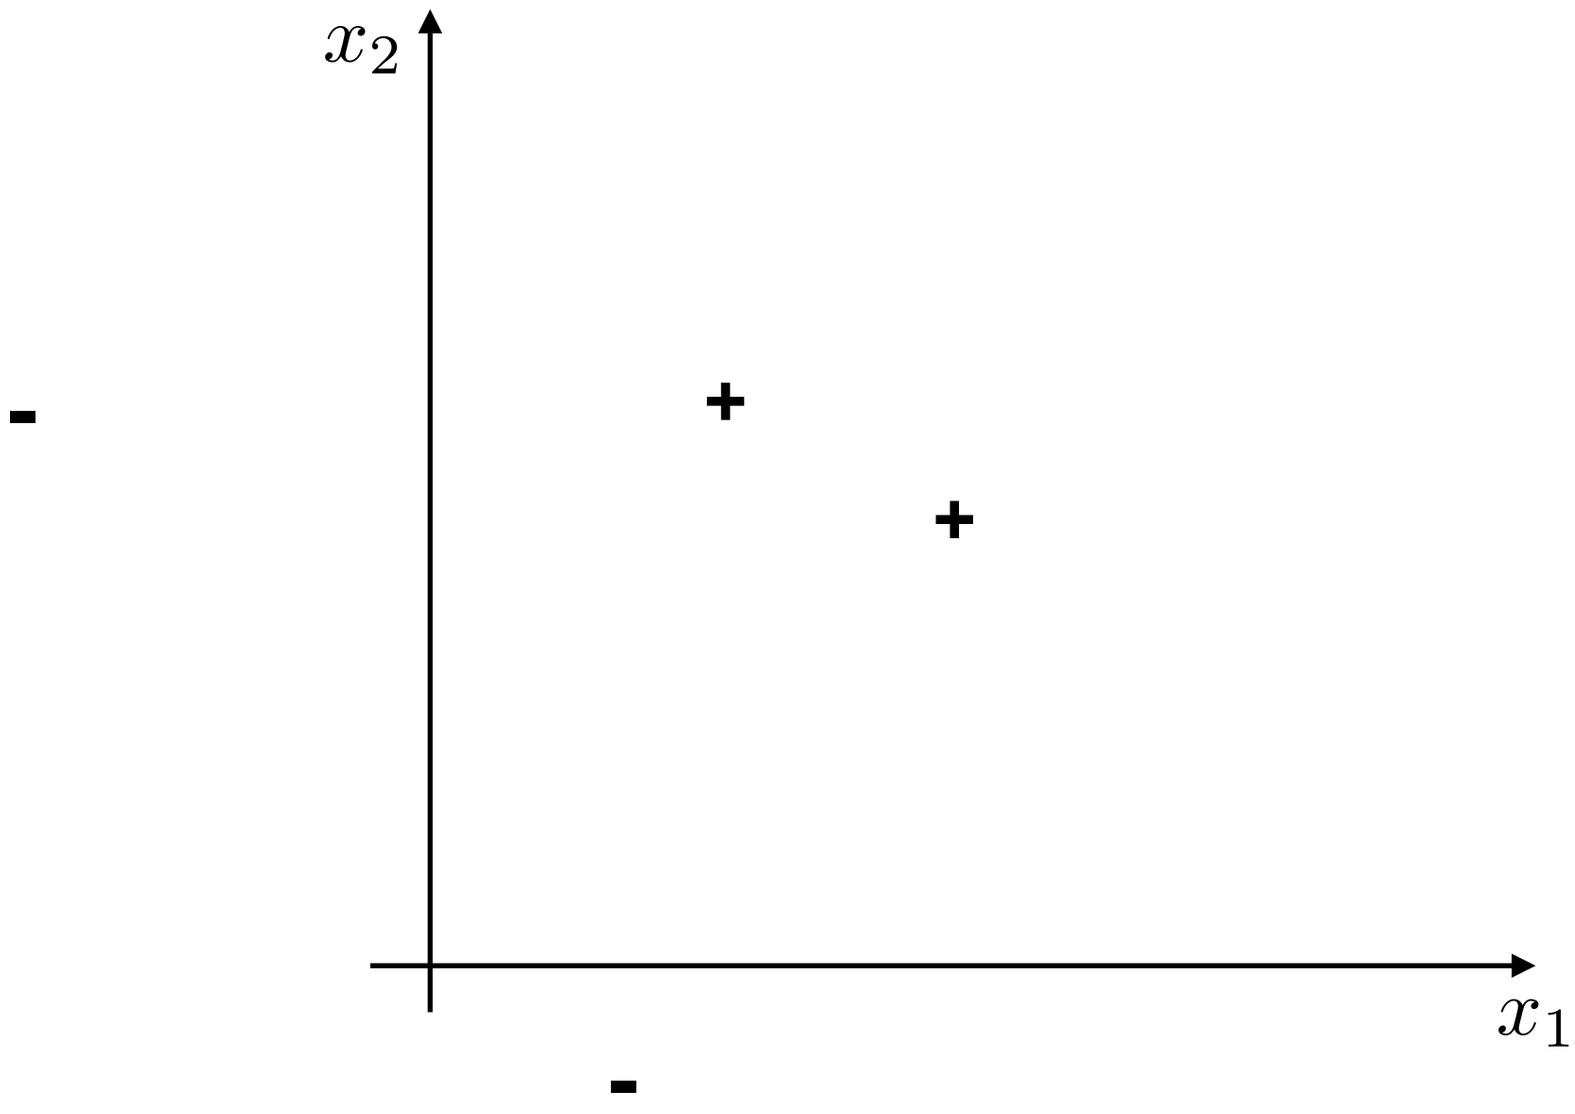

Supervised learning: test set

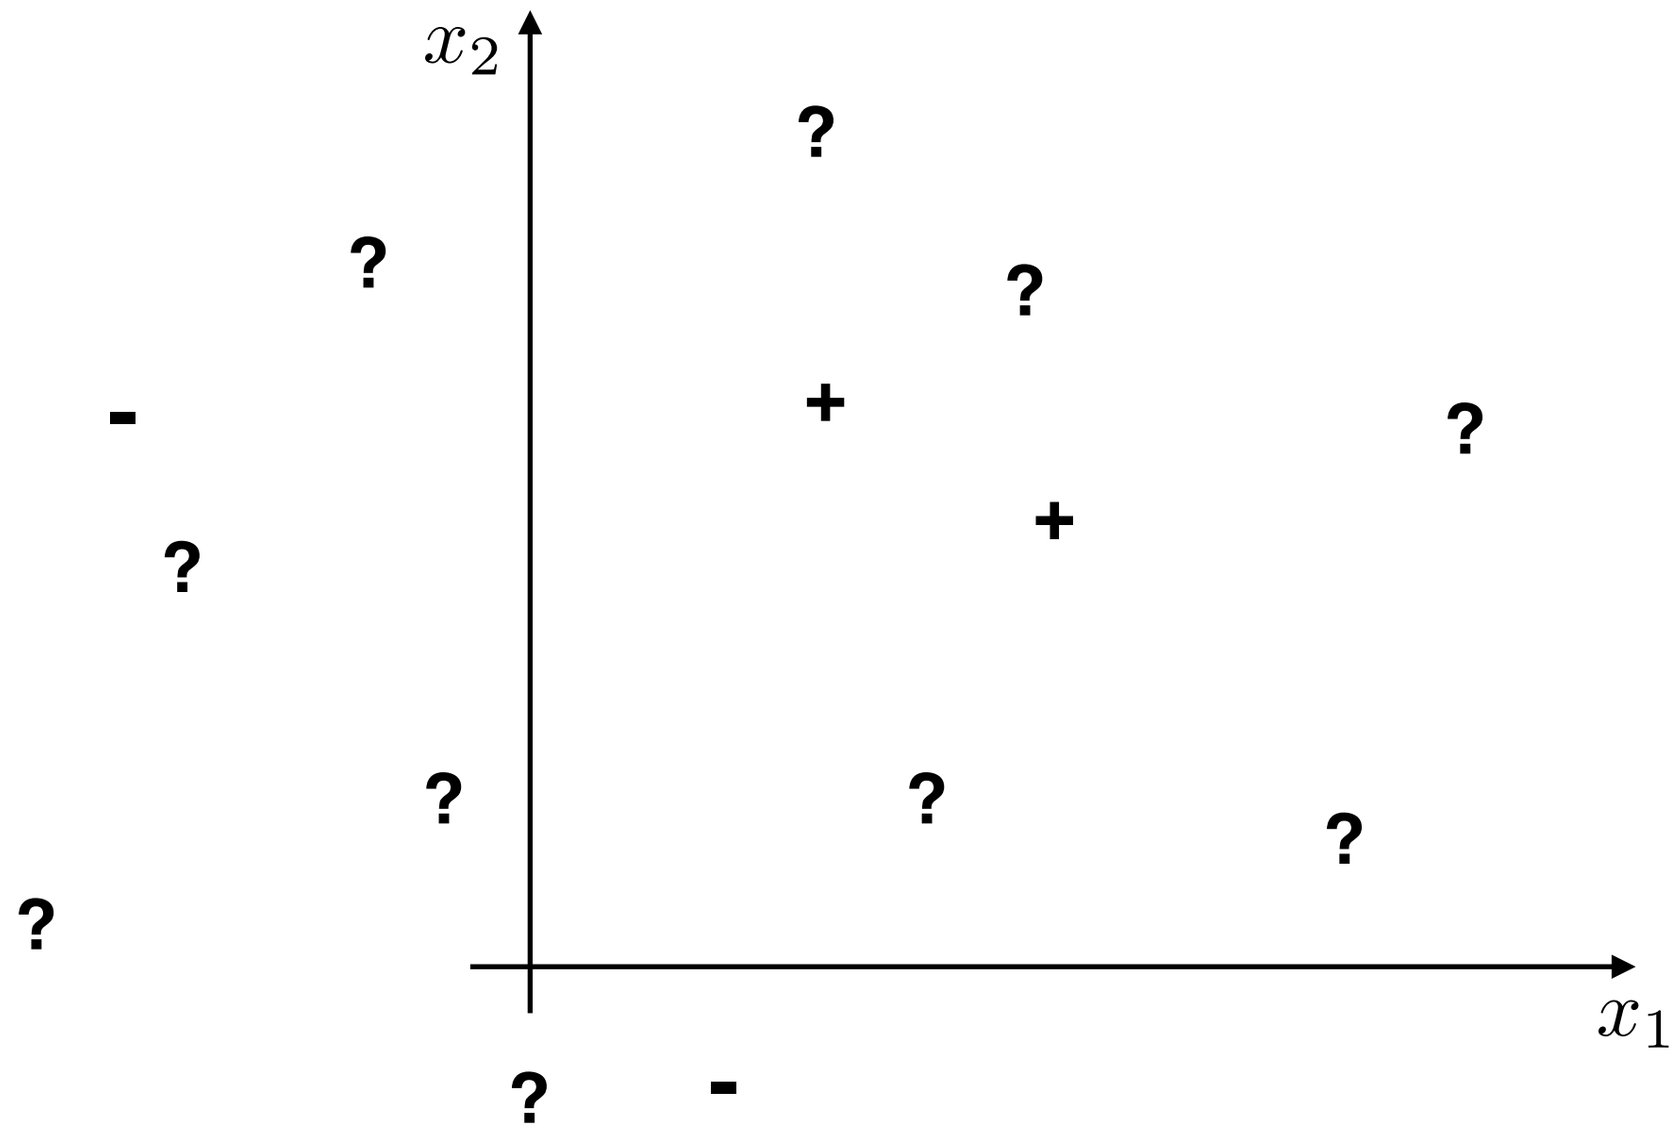

Supervised learning: training set

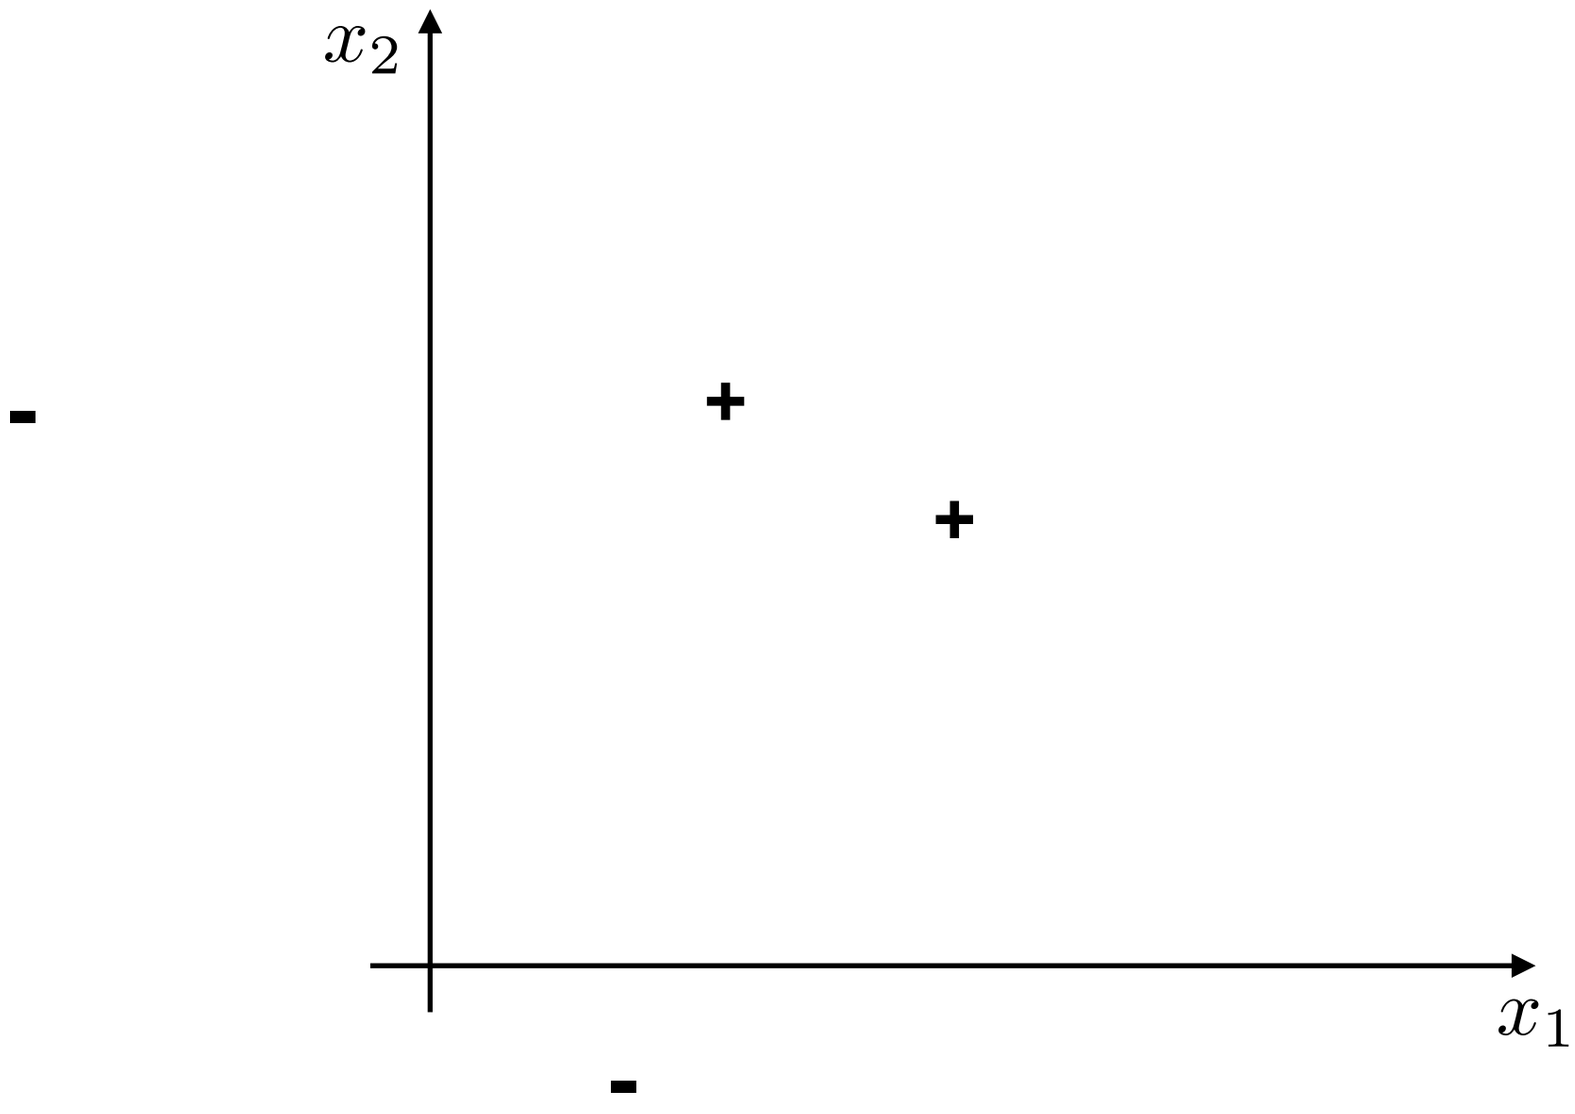

Supervised learning: classifier

$$h : \mathcal{X} \to \{-1, 1\}$$

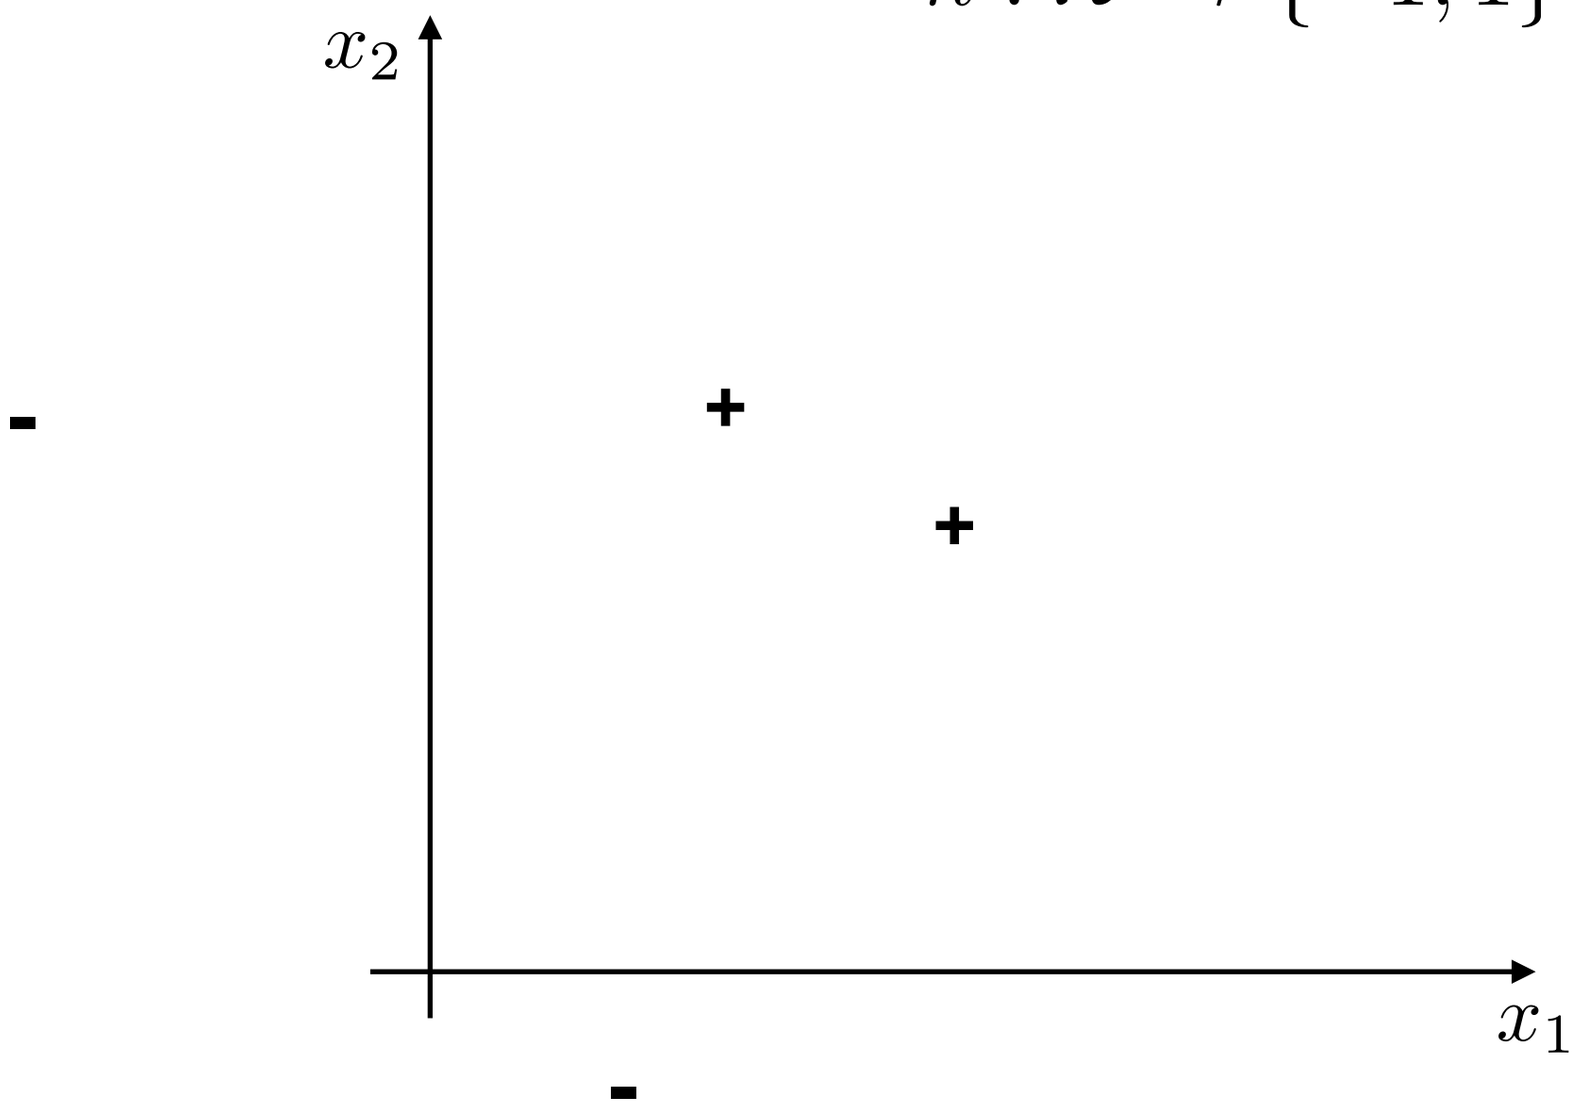

Supervised learning: classifier

$h(x) = -1$ $h(x) = +1$

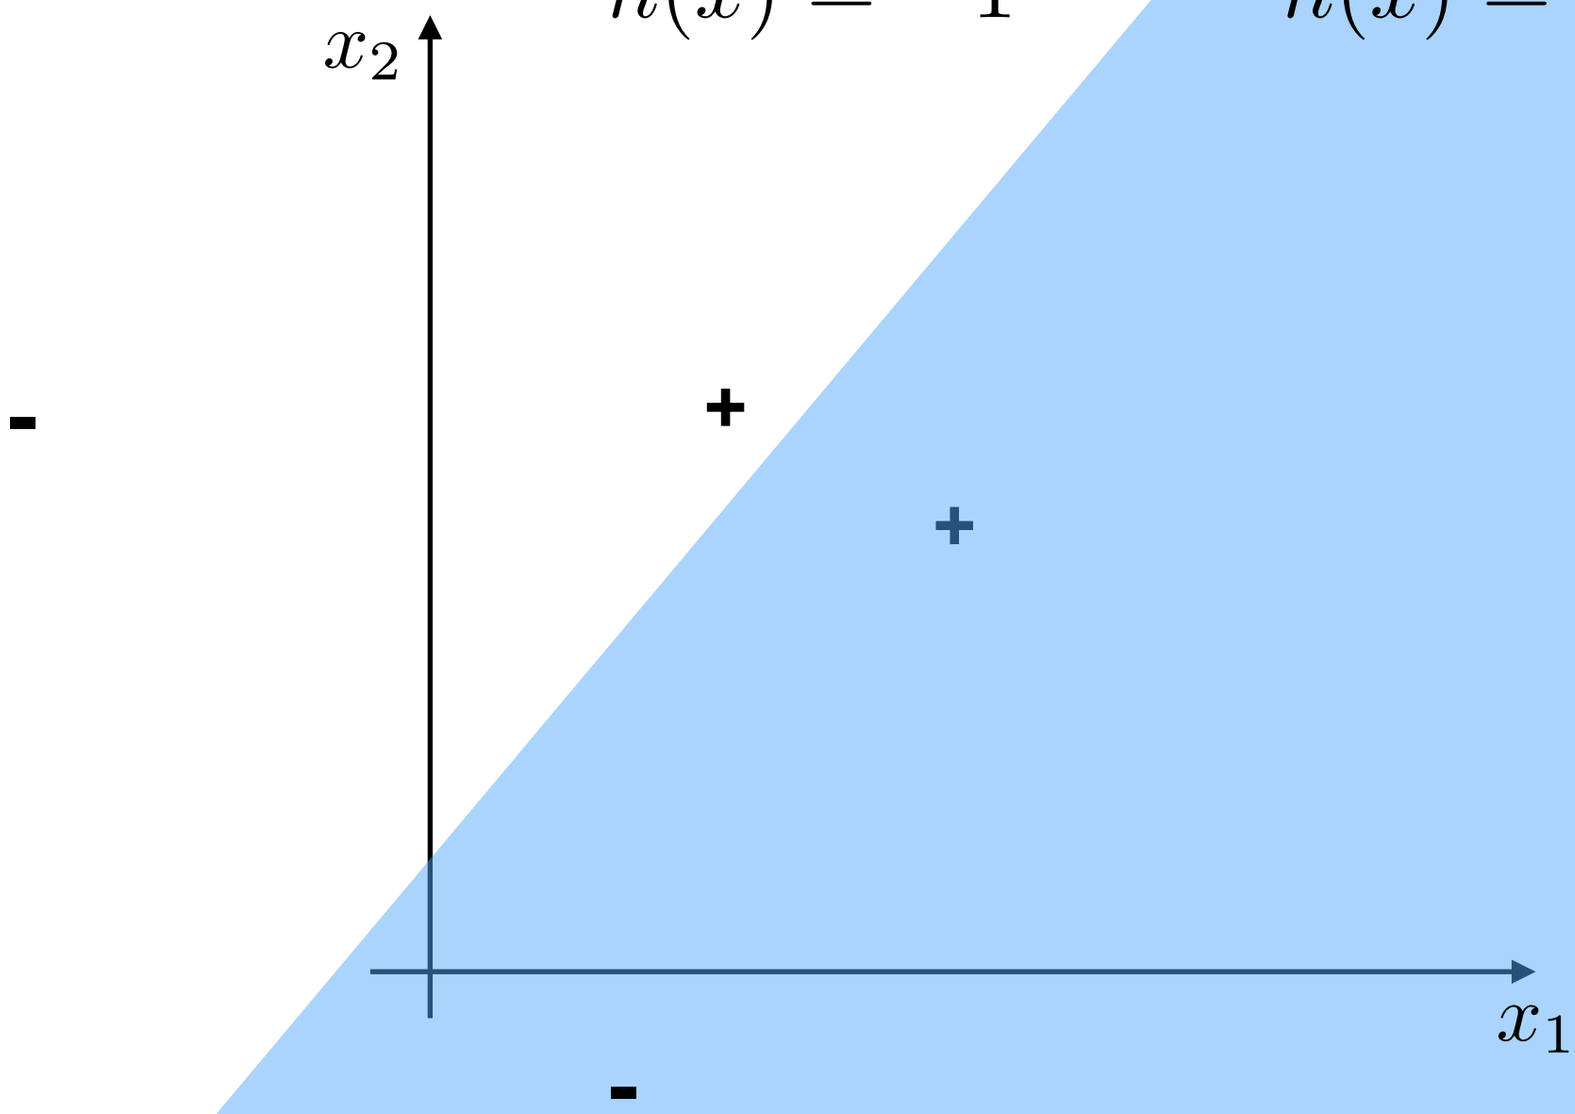

Supervised learning: classifier

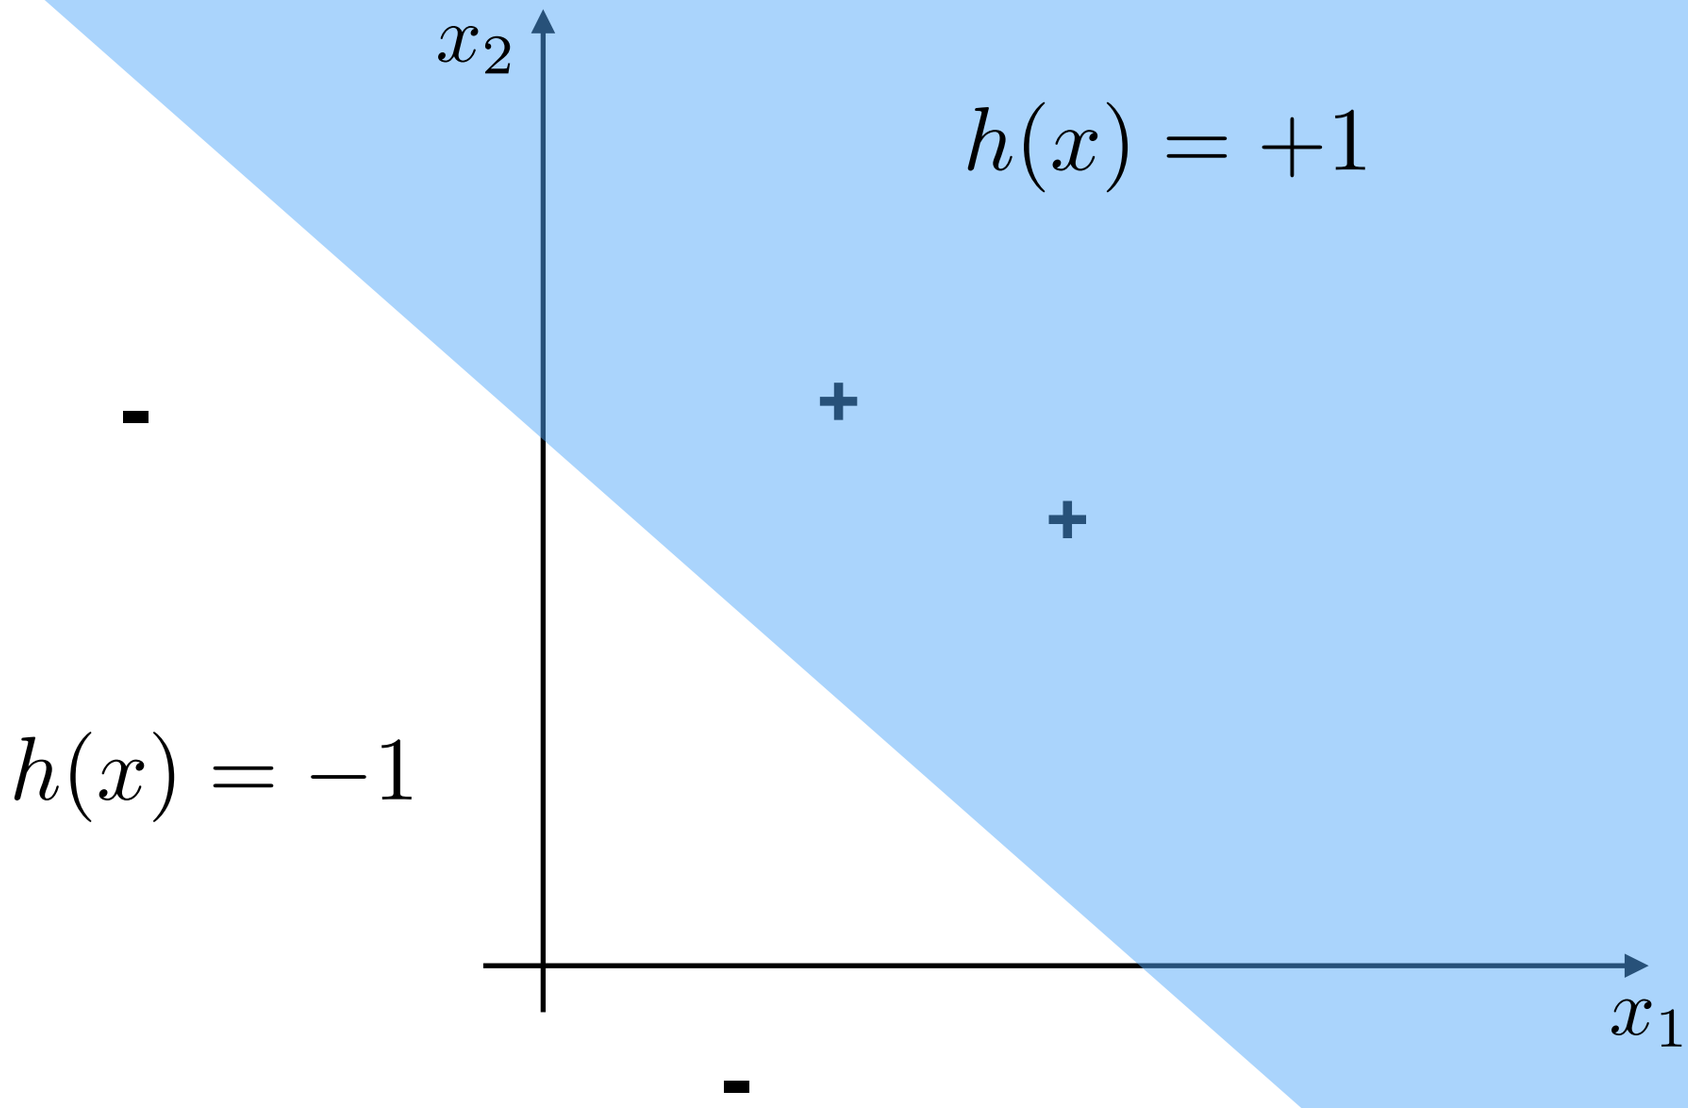

Supervised learning: classifier

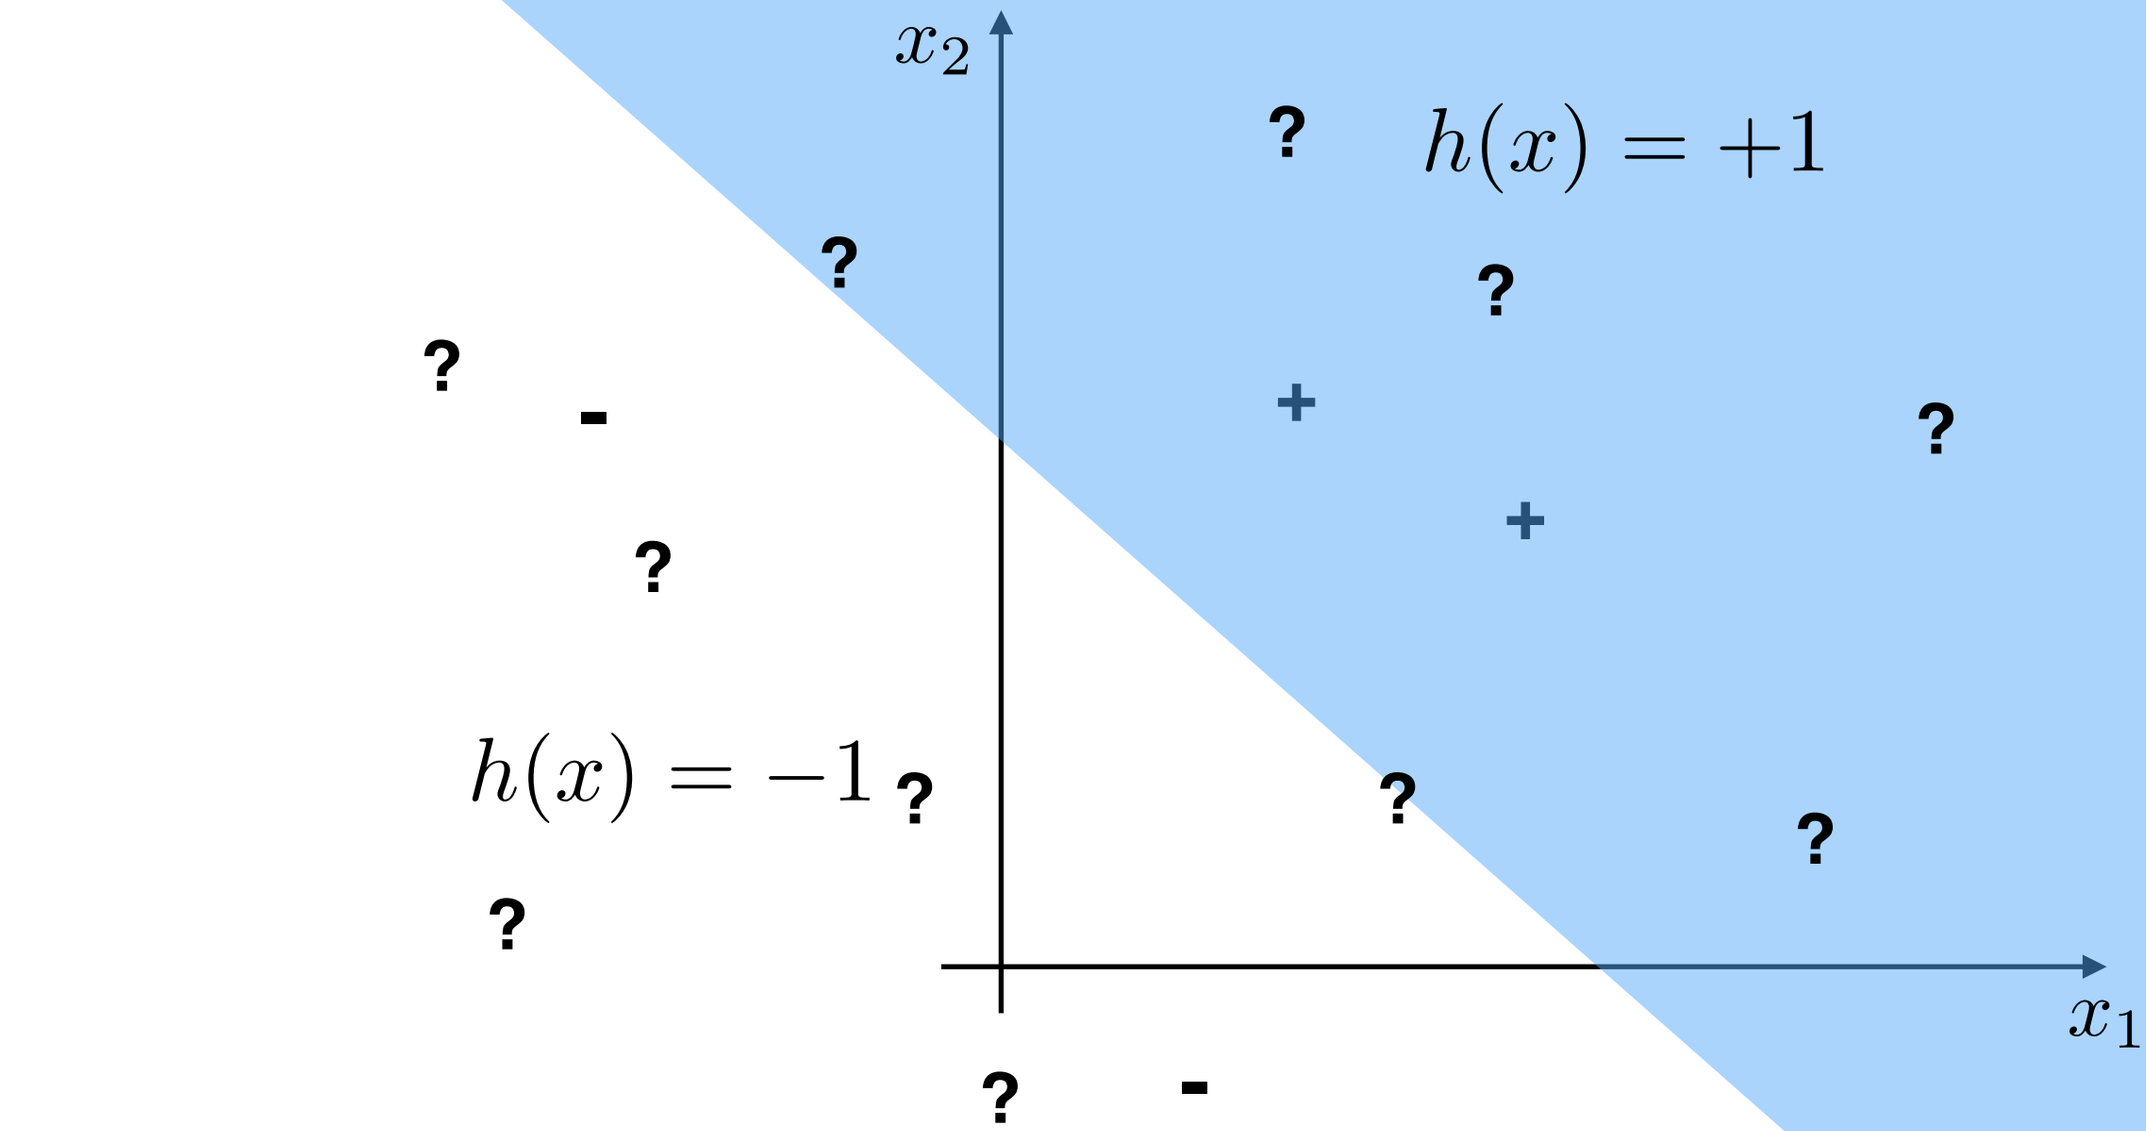

Supervised learning: classifier

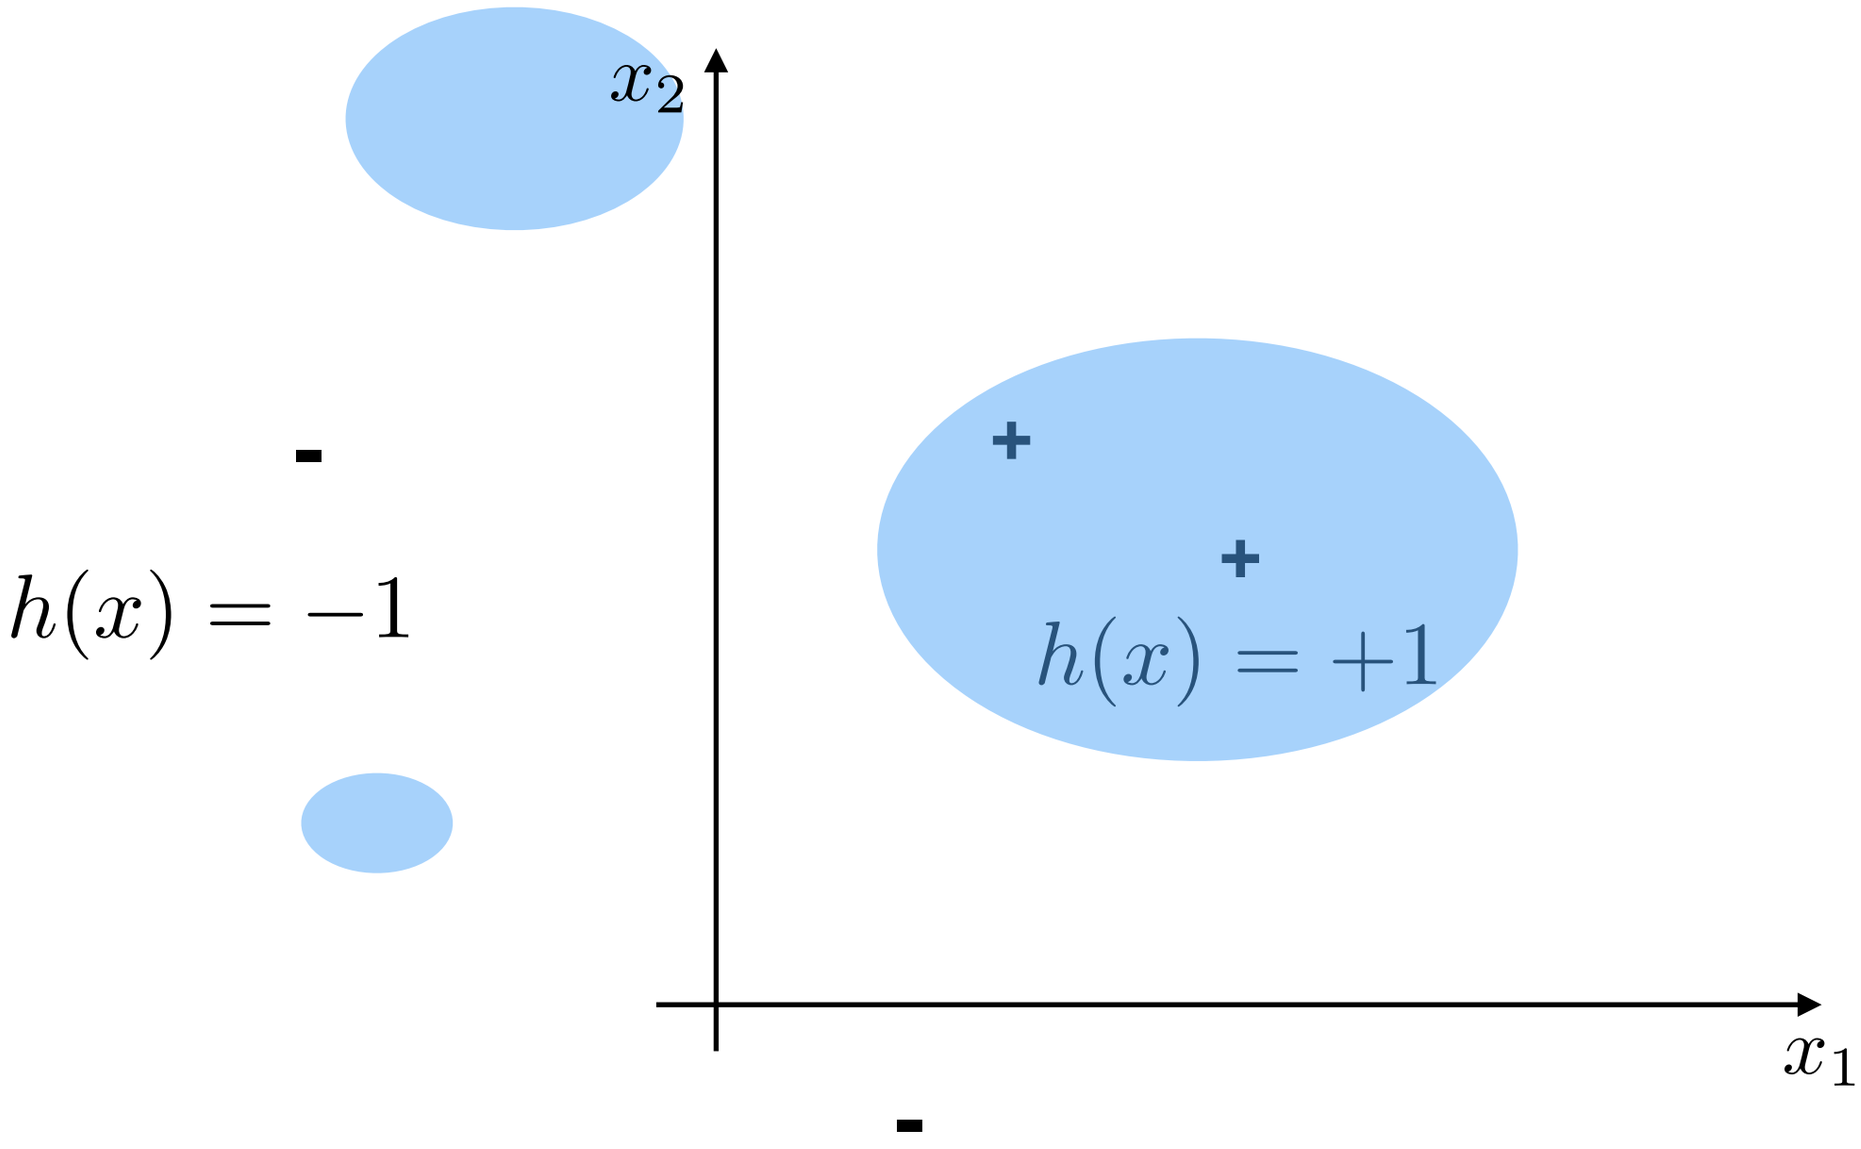

Supervised learning: classifier

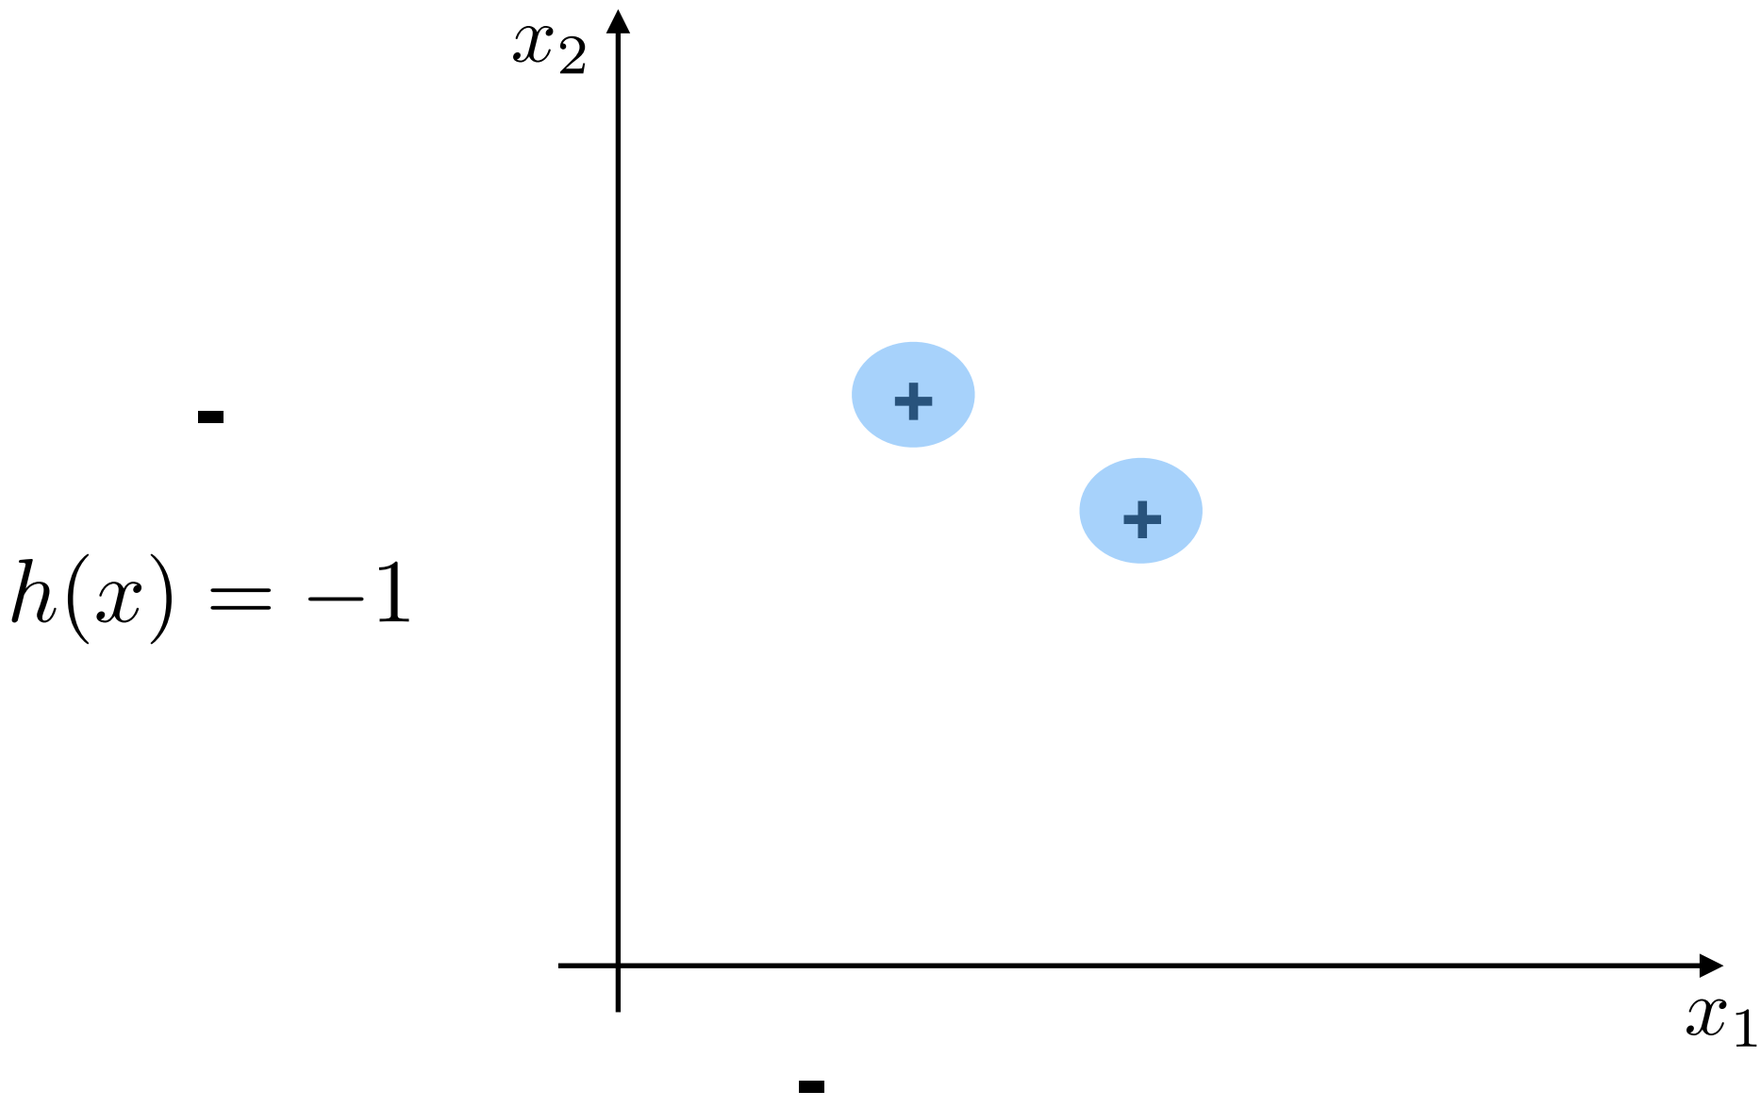

Supervised learning: generalization

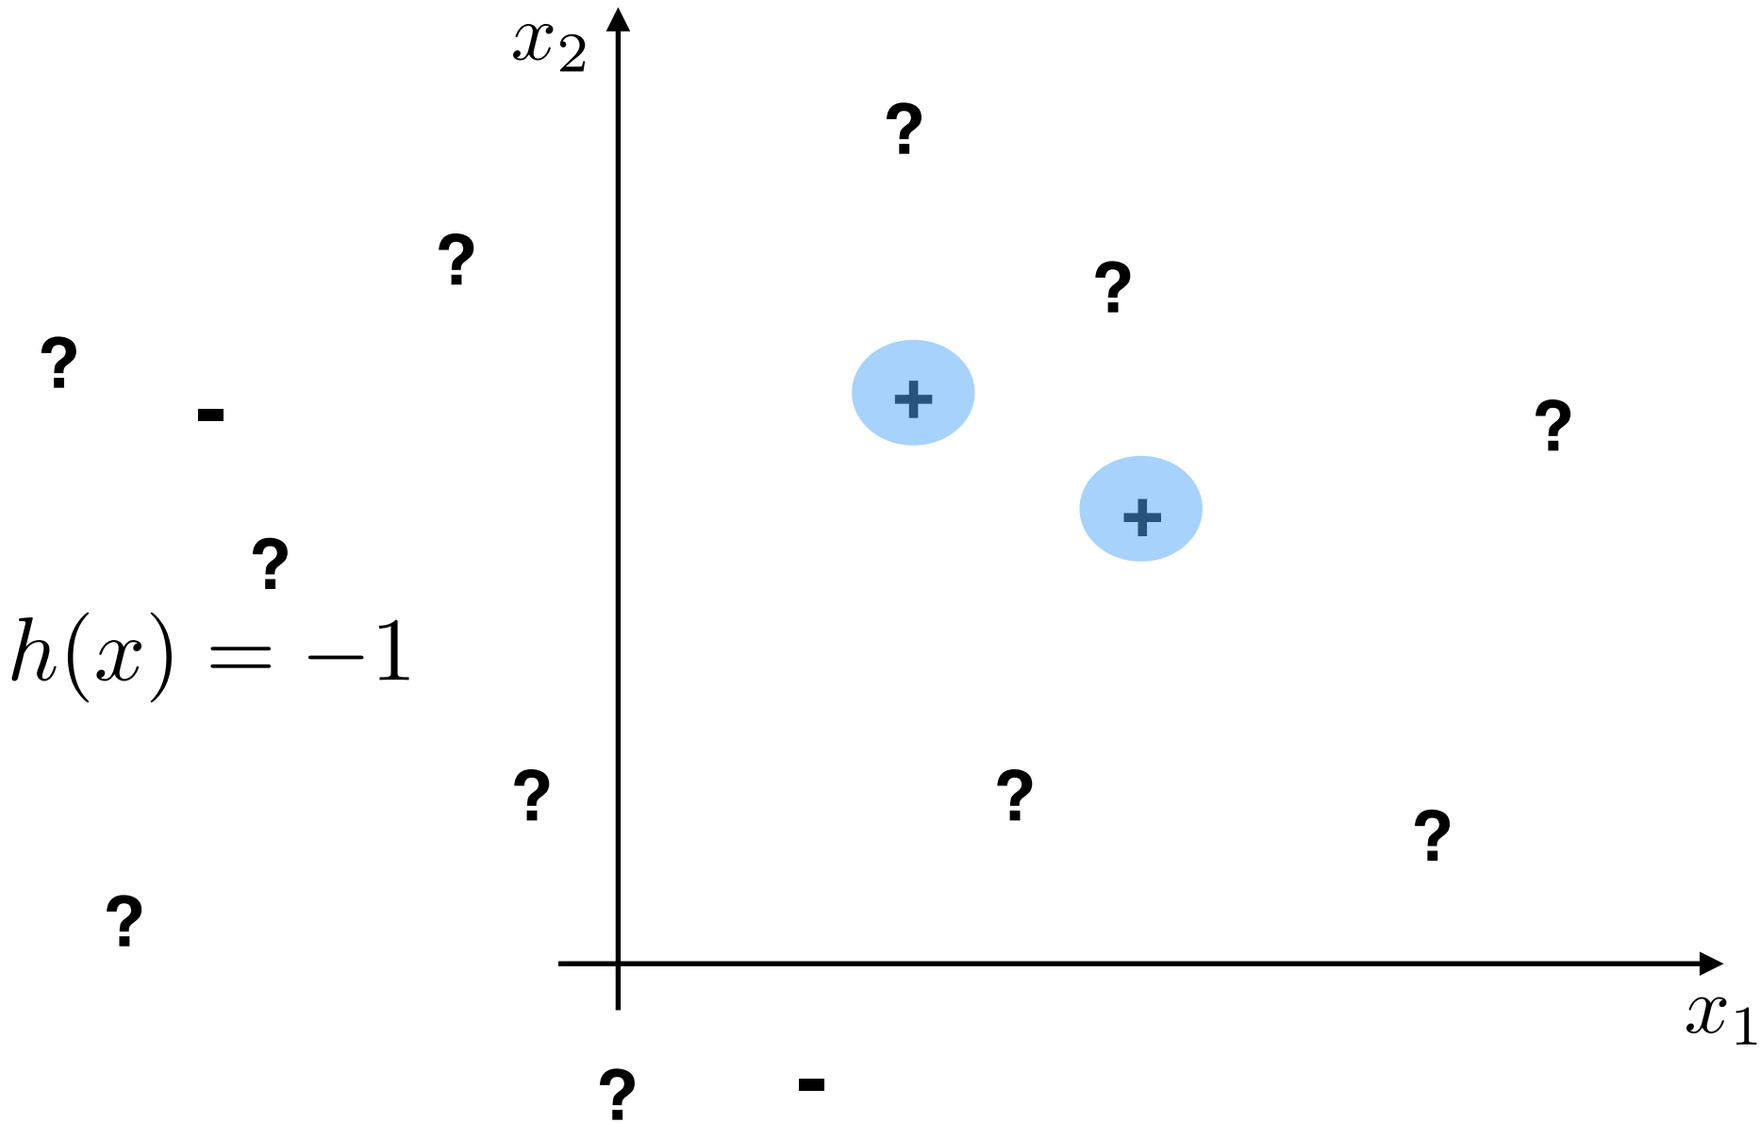

Supervised learning +

- Multi-way classification (e.g., three-way classification)

$$h\left(\quad\right) = \text{politics} \qquad h : \mathcal{X} \to \{\text{politics}, \text{sports}, \text{other}\}$$

13. BBC news and the BBC logo are registered trademarks of Google LLC.

- Regression

$$h\left(\quad\right) = \$1{,}349{,}000 \qquad h : \mathcal{X} \to \mathbb{R}$$

14. © Robert Paul Properties, Cambridge.

- Structured prediction

$$h\left(\quad\right) = \begin{array}{l}\text{A group of people}\\ \text{shopping at an}\\ \text{outdoor market}\end{array} \qquad h : \mathcal{X} \to \{\text{English sentences}\}$$

15. © Computer Vision & Pattern Recognition Lab, 2017.

Types of machine learning

- Supervised learning
  - prediction based on examples of correct behavior
- Unsupervised learning
  - no explicit target, only data, goal to model/discover
- Semi-supervised learning
  - supplement limited annotations with unsupervised learning
- Active learning
  - learn to query the examples actually needed for learning
- Transfer learning
  - how to apply what you have learned from A to B
- Reinforcement learning
  - learning to act, not just predict; goal to optimize the consequences of actions
- Etc.

Key things to understand

- Posing supervised machine learning problems
- Supervised classification
- The role of training/test sets
- A classifier
- A set of classifiers
- Errors, generalization

Attribution List - Lecture 1 - 6.86x Maxchine Learning

1.
Location: Throughout course
Object Source / URL: own work
Citation/Attribution: Instructor content (c) Tommi Jaakkola. Used with permission.

2.
Unit 1 Lecture 1: Introduction to Machine Learning
Screenshot of a Google search of "deep learning"
Slide: # 2
Object Source / URL: Screenshot of a Google search of "deep learning"
Citation/Attribution: Google and the Google logo are registered trademarks of Google LLC.

3.
Unit 1 Lecture 1: Introduction to Machine Learning
Screenshot of a Google search of "Netflix"
Slides: #2
Object Source / URL: Screenshot of a Google search of "Netflix"
Citation/Attribution: (c) Netflix, Inc.

4.
Screenshot of a Google search of "mushrooms"
Slides: #2
Object Source / URL: Screenshot of a Google search of "mushrooms"
Citation/Attribution: Google and the Google logo are registered trademarks of Google LLC.

5.
Unit 1 Lecture 1: Introduction to Machine Learning
Image of “Ear-reddening game move”
Slides: #2
Object Source / URL: https://senseis.xmp.net/?EarReddeningMove
Citation/Attribution: © Senseis Library

6.
Unit 1 Lecture 1: Introduction to Machine Learning
Movie poster thumbnail for the film "Imitation Game"
Slides: 1 at #5, #6, #7, #8. Lecture 12 at #3, #4, #5
Object Source / URL: https://www.imdb.com/title/tt2084970/
Citation/Attribution: © Black Bear Pictures

7.
Unit 1 Lecture 1: Introduction to Machine Learning
Thumbnail of two mushrooms
Slides: # 5, #10, #11, #12, #13, #14
Object Source / URL: https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf
Citation/Attribution: Image from "ImageNet Classification with Deep Convolutional Neural Networks" by Alex Krizhevsky et all, © Neural Information Processing Systems Foundation, Inc.

Attribution List - Lecture 1 - 6.86x Maxchine Learning

8.
Unit 1 Lecture 1: Introduction to Machine Learning
Thumbnail of one Dalmatian dog
Slides: # 5, #10, #11, #12, #13, #14
Object Source / URL: https://www.slideserve.com/rufin/an-overview-of-deep-learning
Citation/Attribution: “An Overview of Deep Learning” © Geoffrey Hinton, University of Toronto.

9.
Unit 1 Lecture 1: Introduction to Machine Learning
Movie poster thumbnail for the film "Interstellar"
Slides: #10, #11, #12, #13, #14
Object Source / URL: https://www.imdb.com/title/tt0816692/
Citation/Attribution: © 2020 The Studios At Paramount

10.
Unit 1 Lecture 1: Introduction to Machine Learning
Movie poster thumbnail for the film "Gone Girl"
Slides: #10, #11, #12, #13, #14
Object Source / URL: https://www.imdb.com/title/tt2267998/?ref_=fn_al_tt_1
Citation/Attribution: © Twentieth Century Fox

11.
Unit 1 Lecture 1: Introduction to Machine Learning
Movie poster thumbnail for the film "Fury"
Slides: #10, #11, #12, #13, #14
Object Source / URL: https://www.imdb.com/title/tt2713180/?ref_=fn_al_tt_1
Citation/Attribution: © 2020 Sony Pictures Digital Productions Inc.

12.
Unit 1 Lecture 1: Introduction to Machine Learning
Movie poster thumbnail for "Big Hero 6"
Slides: #10, #11, #12, #13, #14
Object Source / URL: https://www.imdb.com/title/tt2245084/?ref_=nv_sr_srsg_0
Citation/Attribution: © Disney

13.
Unit 1 Lecture 1: Introduction to Machine Learning
Screenshot of BBC news page
Slides: #25
Object Source / URL: http://news.bbc.co.uk/2/hi/uk_news/7616996.stm
Citation/Attribution: BBC news and the BBC logo are registered trademarks of Google LLC.

14.
Unit 1 Lecture 1: Introduction to Machine Learning
Thumbnail of a living room
Slides: #25
Object Source / URL: https://www.robertpaul.com/real-estate/150-cambridge-st-410-cambridge-ma-02141/71952983/28587689
Citation/Attribution: © Robert Paul Properties, Cambridge.

15.
Unit 1 Lecture 1: Introduction to Machine Learning
Thumbnail of a group of people shopping at an outdoor market
Slides: #25
Object Source / URL: http://166.104.231.121/ysmoon/mip2017/lecture_note/제9장.pdf -- page 2
Citation/Attribution: © Computer Vision & Pattern Recognition Lab, 2017.