CHI '21 Course:
An Introduction to Intelligent User Interfaces
Overview of AI and ML Terms, Concepts and Tools

Albrecht Schmidt, Sven Mayer, Daniel Buschek

Introduction

Albrecht Schmidt

Sven Mayer

Daniel Buschek

Examples & Motivation

Text Suggestions

Google‘s Smart Reply & Smart Compose

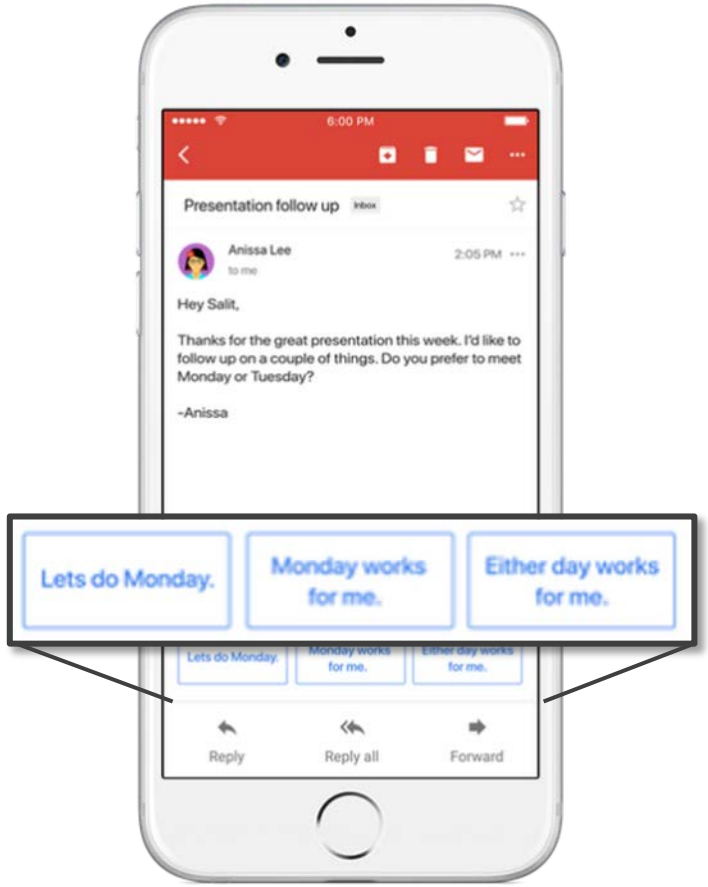

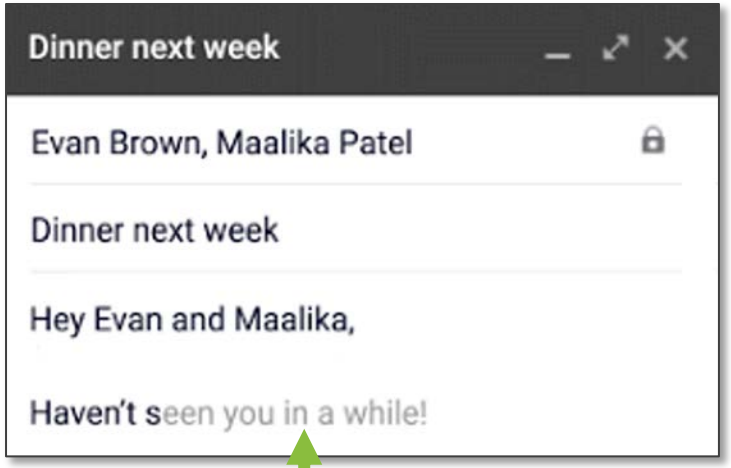

Language model, given email text

https://blog.google/products/gmail/save-time-with-smart-reply-in-gmail/
https://ai.googleblog.com/2018/05/smart-compose-using-neural-networks-to.html

Semantic Image Manipulation

„Smart Portrait Filters“ in Adobe‘s Photoshop

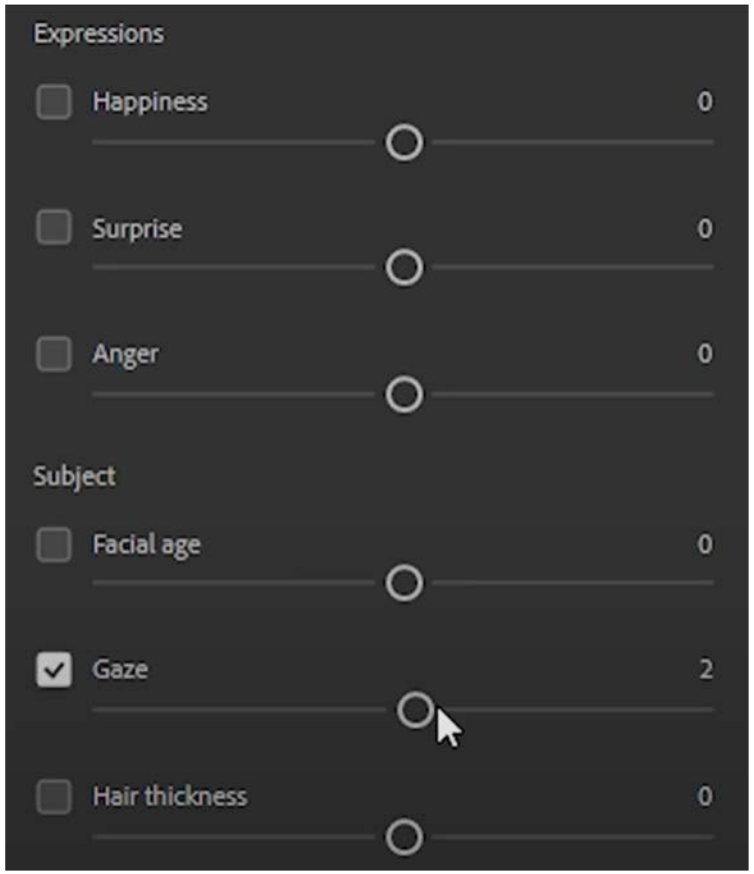

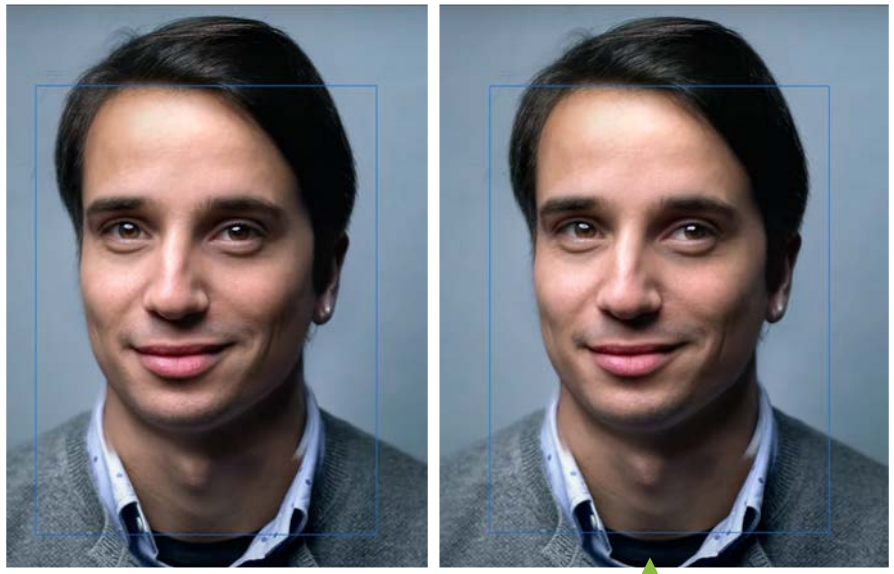

Generative model,
learned from many portraits

https://blog.adobe.com/en/2020/10/20/photoshop-the-worlds-most-advanced-ai-application-for-creatives.html
https://blogs.nvidia.com/blog/2020/10/20/adobe-max-ai/, https://github.com/NVlabs/stylegan2

Recommender Systems

How do recommender systems impact the user experience?

Carlos A. Gomez-Uribe and Neil Hunt. 2015. The Netflix Recommender System: Algorithms, Business Value, and Innovation. ACM Trans. Manage. Inf. Syst. 6, 4, Article 13 (December 2015), 19 pages. DOI: https://doi.org/10.1145/2843948

- Why are recommender systems used?
- How do recommender work?
- What data do recommender systems require?

amazon.com ebay

YouTube Spotify

Text analytics

Where can we use it and how can it improve interaction?

- Answering questions like
  - What is this text about?
  - What did the person communicate?
  - What is the key information in this document?
  - What feelings are communicated?
  - Is this different from what was said before?

- Application areas
  - Social media analytics, e.g. twitter
  - Communication and reading interfaces
  - Customer reviews and feedback
  - Chat bots
  - Text Forensics

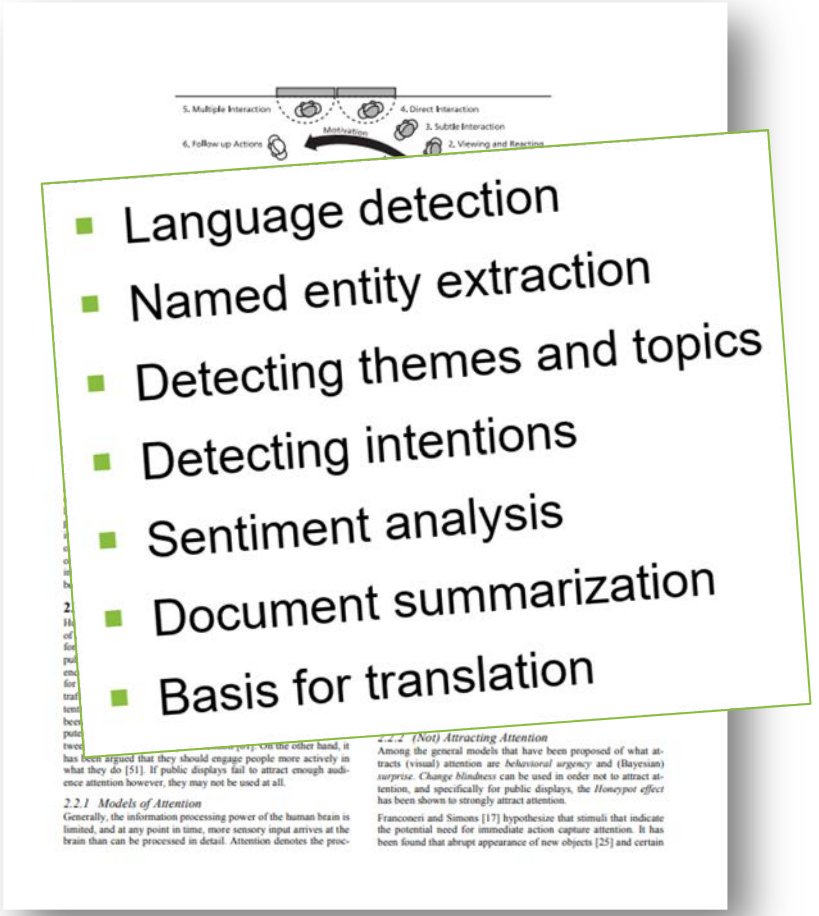

http://www.medien.ifi.lmu.de/pubdb/publications/pub/mueller2010mm/mueller2010mm.pdf

VUI design process

How to design a dialog structure?

- Think of alternatives
  - structure
  - wording
- Try out your dialog
  - wizard of Oz technique!
  - use outside people
- Refine, Revise, Repeat

Image by Gregory Varnum, CC BY-SA 4.0 via Wikimedia Commons
https://commons.wikimedia.org/wiki/File:Amazon_Echo_Dot_-_June_2018_(1952).jpg

A Deceptive UI: redirected Walking

What is real in an intelligent UI?

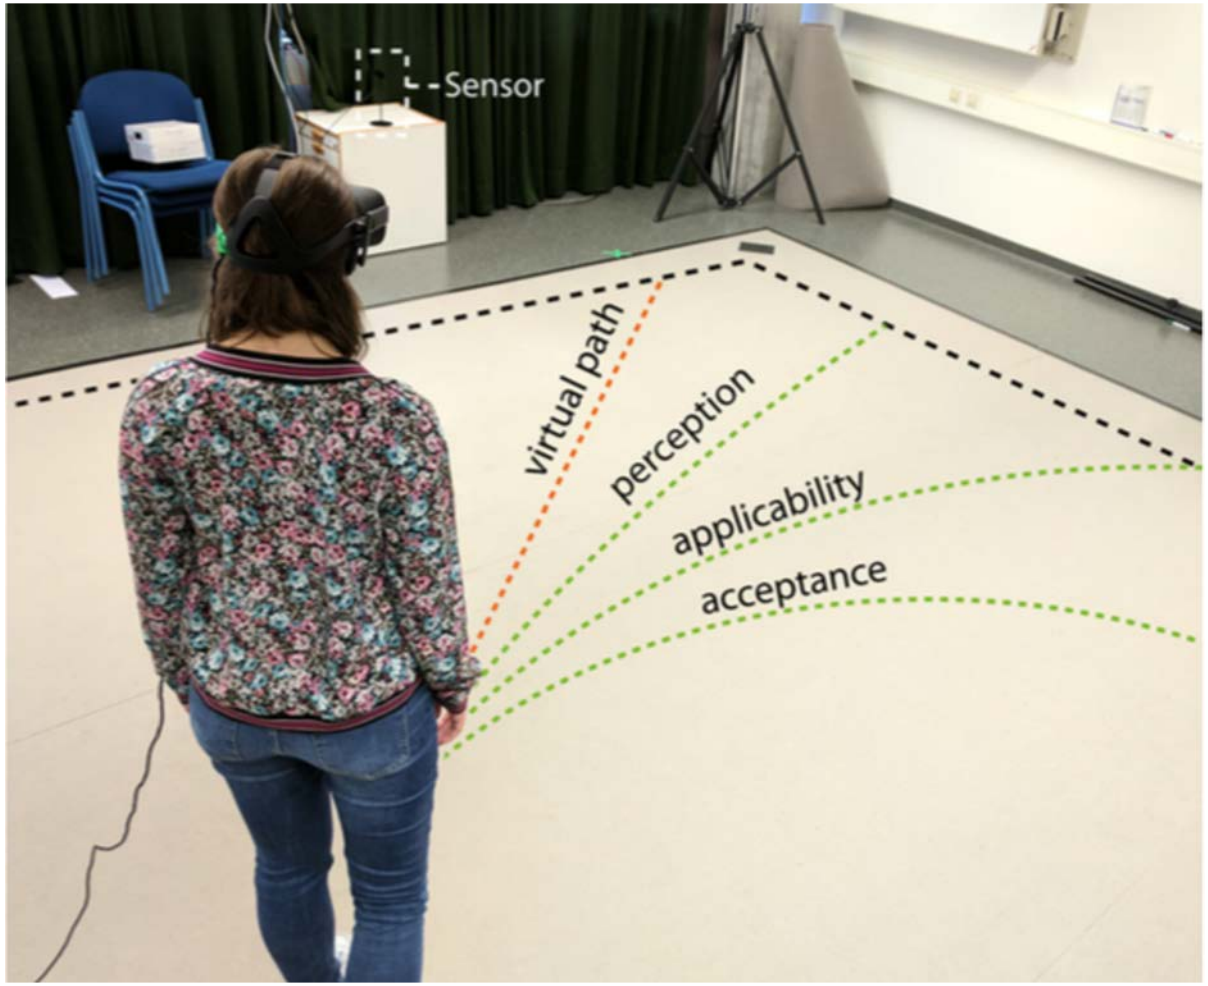

M. Rietzler, J. Gugenheimer, T. Hirzle, M. Deubzer, E. Langbehn and E. Rukzio, "Rethinking Redirected Walking: On the Use of Curvature Gains Beyond Perceptual Limitations and Revisiting Bending Gains," *2018 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, Munich, Germany, 2018, pp. 115-122, doi: 10.1109/ISMAR.2018.00041.

Image from https://ieeexplore.ieee.org/abstract/document/8613757

Facial Recognition

Convenient biometric or overly powerful?

- Unlock your phone
  - Hands-free identification
  - What are the major issues?

- Surveillance
  - Privacy
  - Tricks to „hide“ from facial recognition technology

https://www.bbc.com/news/uk-51237665

(a) Near infrared LED not lit (detection successful)

(b) Near infrared LED lit (detection failed)

http://research.nii.ac.jp/~iechizen/official/research-e.html#research2c

HCI Replacing HHI in Stores

„Just Walk Out“ shopping experience at Amazon Go

- Surveillance-powered shopping
  - Does not use facial recognition
- How does it work?
  - Is it „intelligent“? How so?

Image by SounderBruce, CC BY-SA 4.0 via Wikimedia Commons
https://commons.wikimedia.org/wiki/File:Amazon_Go_in_Seattle,_December_2016.jpg

AI Recruiting

Is an AI a “fairer“ recruiter?

Natural Language Translation

Female historians and male nurses do not exist?

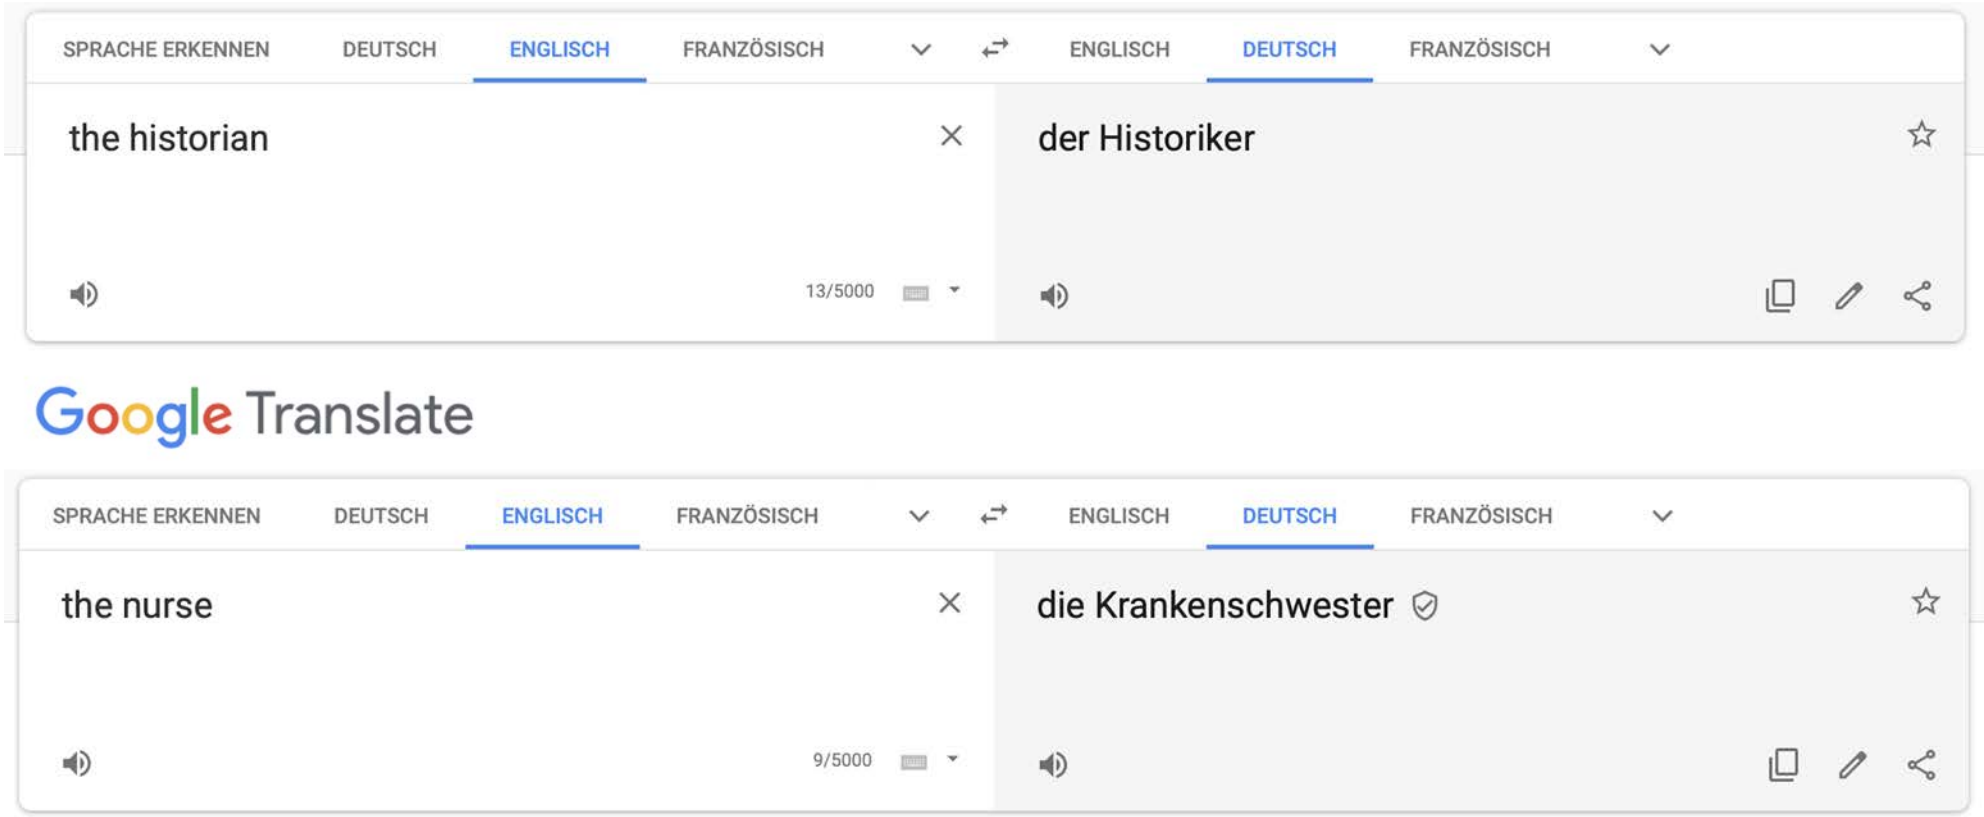

https://translate.google.com
https://algorithmwatch.org/en/story/google-translate-gender-bias/

Intelligent Touch

Why are we so precise with our fingers on a screen?

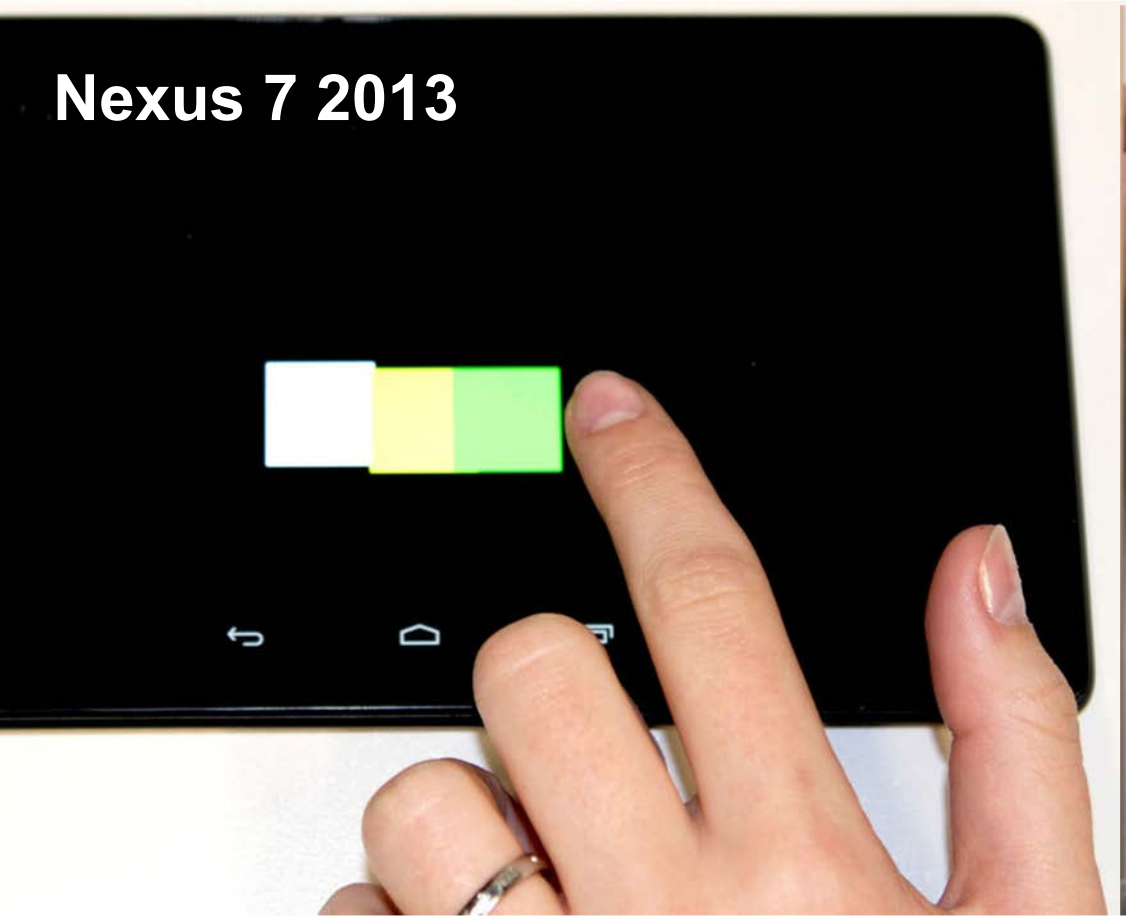

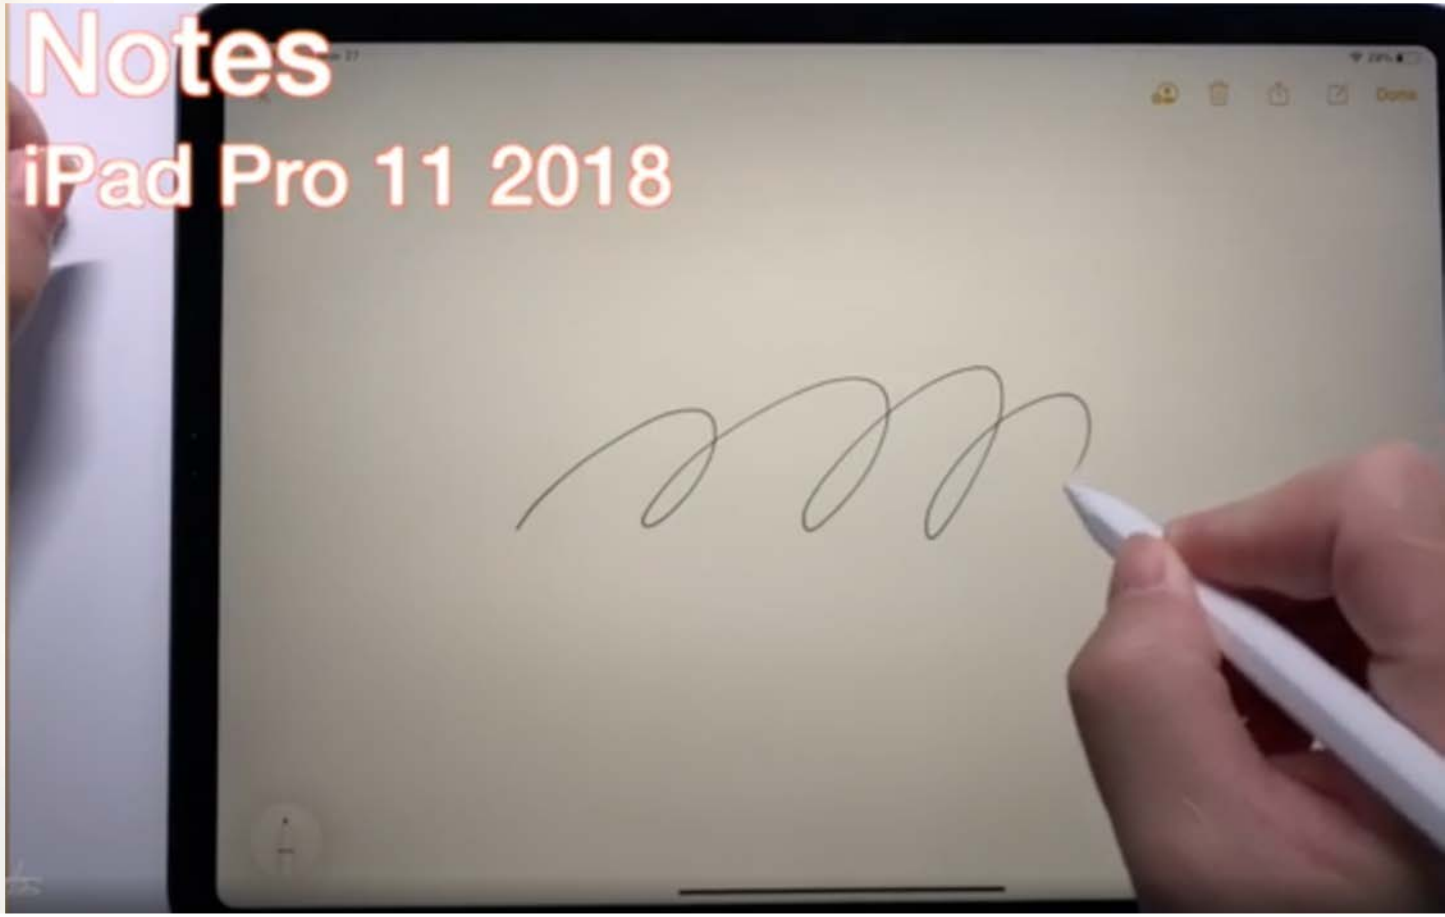

Henze, N., Mayer, S., Le, H.V. and Schwind, V. Improving software-reduced touchscreen latency. *Proc. MobileHCI '17* https://doi.org/10.1145/3098279.3122150

https://www.youtube.com/watch?v=l6Nz8wVUU74

AI for IUIs

HCI perspective

What is considered Artificial Intelligence?

What is considered Artifical Intelligence?

What is Human Intelligence?

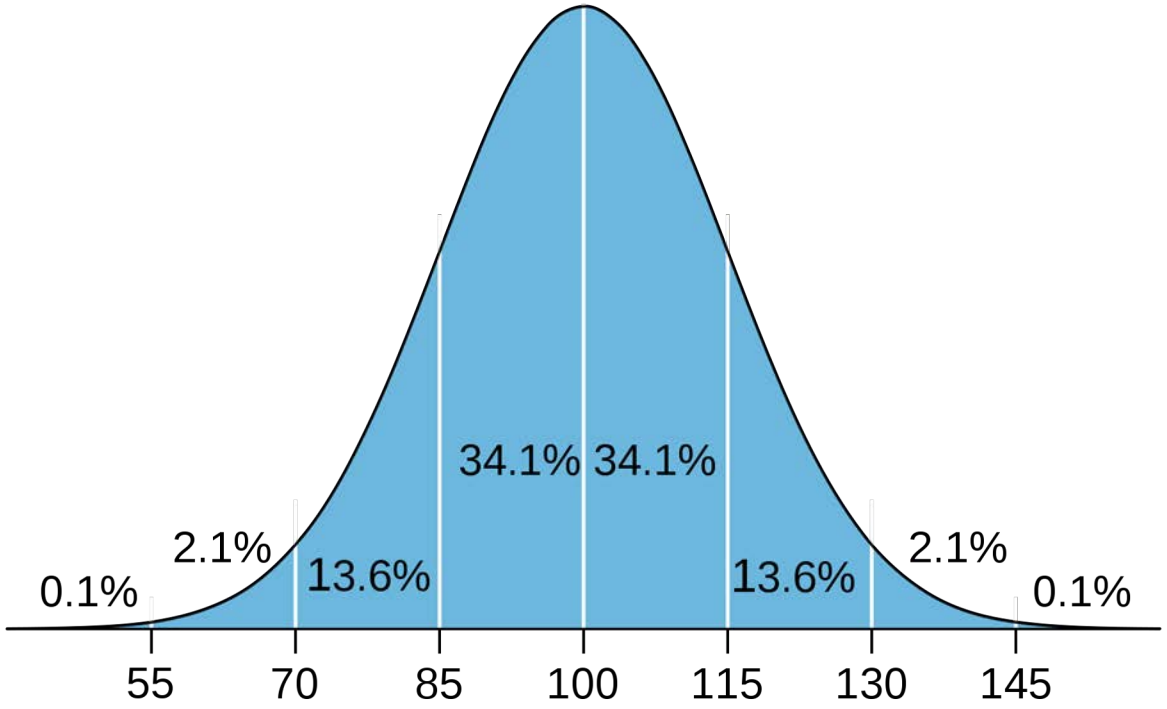

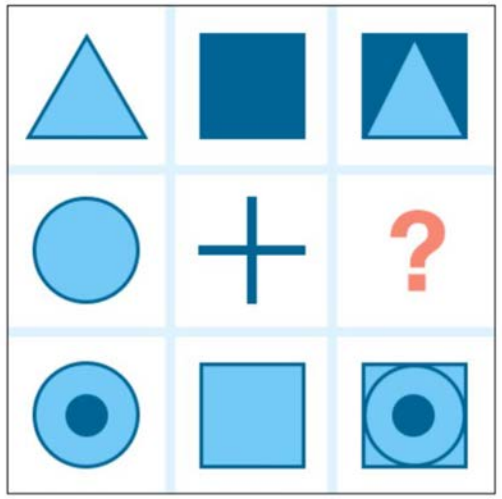

Which figure fits?

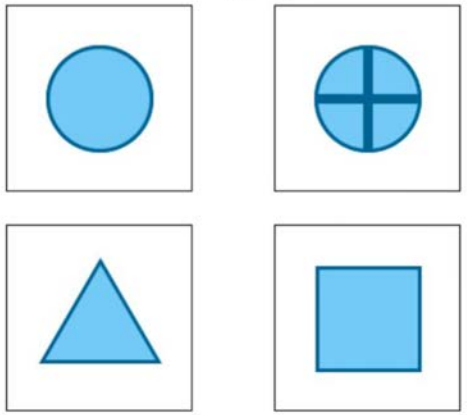

What is in an IQ test?

Many different types… typical questions include:

- Finding Analogies (math and verbal)
  - *Pen and writing, cup and ???*
- Finding / extending Pattern (graphical and math)
  - *45 … 40 … 60 … 55 … 75 … 70 … ???*
- Classification tasks
  - *Make two groups: apple, plate, grape, cake, spoon, knife*
- Making sense spatial and visual representations
- Reasoning and logical
- General knowedge

Is the IQ static or can it be learned?

- Google Search for IQ trainer…

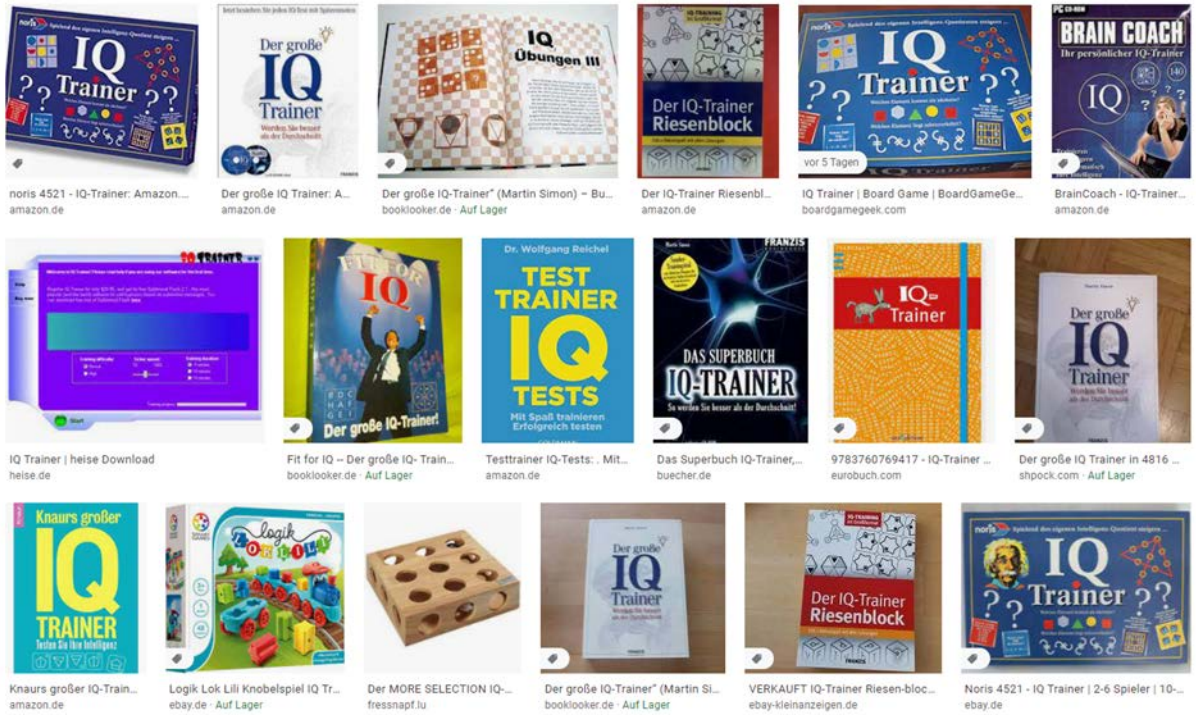

What is…

Mini In-Class Exercise

- Artificial intelligence
- Machine Learning

- Write a definition for one of the term
  - Time limit is 3 minutes
  - No more than 20 words

Artificial Intelligence?

- Artificial Narrow Intelligence (also Weak AI)
  - Solves very specific and well described problems in a specific domain
- Artificial General Intelligence (also Strong AI)
  - An artificial intelligence that has the ability to mimic human intelligence
  - Its behavior cannot be distinguished by observation from a human
- Artificial Super Intelligence
  - An artificial intelligence that surpasses human intelligence

http://www.geminoid.jp/en

AI Knowledge Map: how to classify AI technologies

A sketch of a new AI technology landscape by Francesco Corea

- **Technologies**
  - Logic-based
  - Knowledge-based
  - Probabilistic methods
  - Machine learning
  - Embodied intelligence
  - Search and optimization
- **Domains**
  - Reasoning
  - Knowledge
  - Planning
  - Communication
  - Perception

https://medium.com/@Francesco_AI/ai-knowledge-map-how-to-classify-ai-technologies-6c073b969020

https://medium.com/@Francesco_AI/ai-knowledge-map-how-to-classify-ai-technologies-6c073b969020

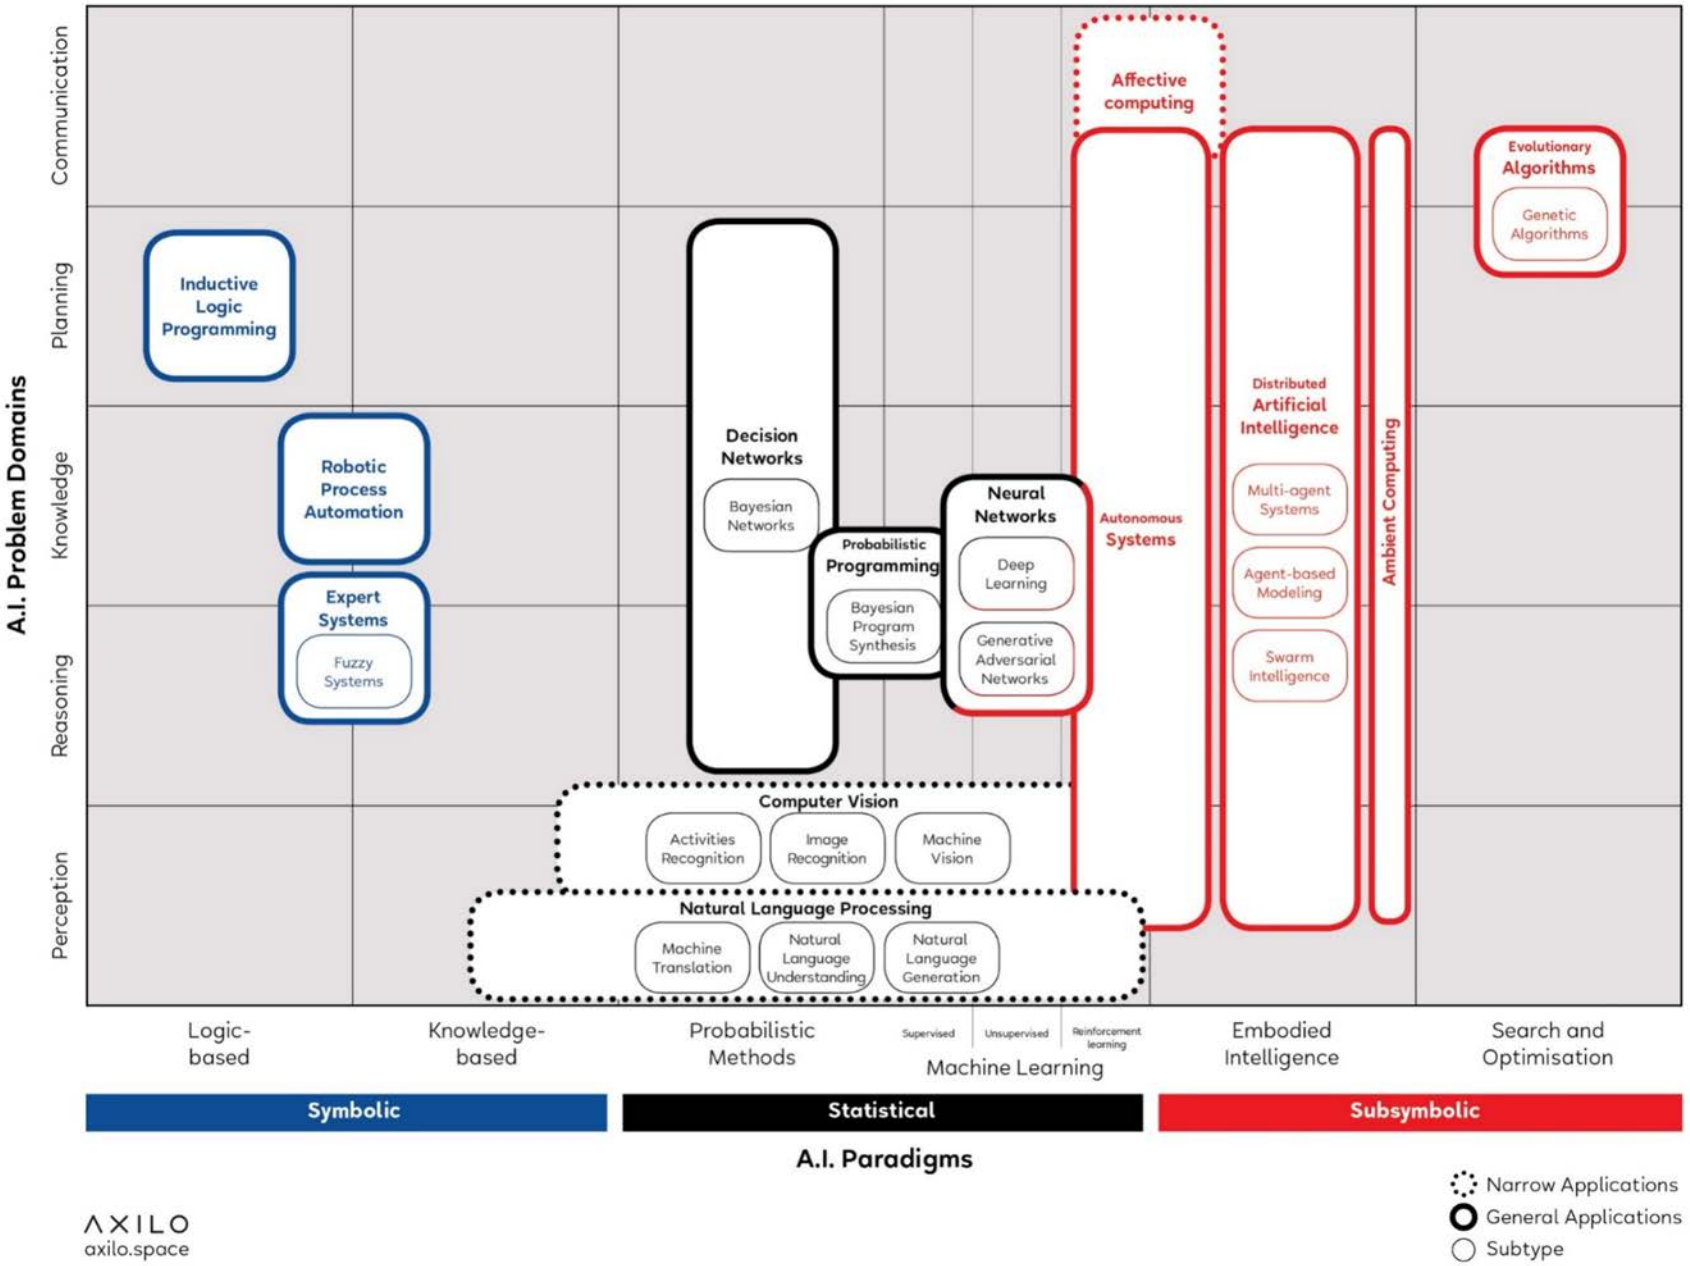

AI Knowledge Map: how to classify AI technologies. A sketch of a new AI technology landscape Francesco Corea

AI Knowledge Map: how to classify AI technologies

A sketch of a new AI technology landscape by Francesco Corea

Technologies

- **Logic-based tools**: tools that are used for knowledge representation and problem-solving
- **Knowledge-based tools**: tools based on ontologies and huge databases of notions, information, and rules
- **Probabilistic methods**: tools that allow agents to act in incomplete information scenarios
- **Machine learning**: tools that allow computers to learn from data
- **Embodied intelligence**: engineering toolbox, which assumes that a body (or at least a partial set of functions such as movement, perception, interaction, and visualization) is required for higher intelligence
- **Search and optimization**: tools that allow intelligently searching through many possible solutions

https://medium.com/@Francesco_AI/ai-knowledge-map-how-to-classify-ai-technologies-6c073b969020

AI Knowledge Map: how to classify AI technologies

A sketch of a new AI technology landscape by Francesco Corea

Domains

- **Reasoning**: the capability to solve problems
- **Knowledge**: the ability to represent and understand the world
- **Planning**: the capability of setting and achieving goals
- **Communication**: the ability to understand language and communicate
- **Perception**: the ability to transform raw sensorial inputs (e.g., images, sounds, etc.) into usable information

https://medium.com/@Francesco_AI/ai-knowledge-map-how-to-classify-ai-technologies-6c073b969020

Supervised vs. Unsupervised Learning

Types of ML Problems

What can we learn just by looking at the data?

Types of ML Problems

What can we learn just by looking at the data?

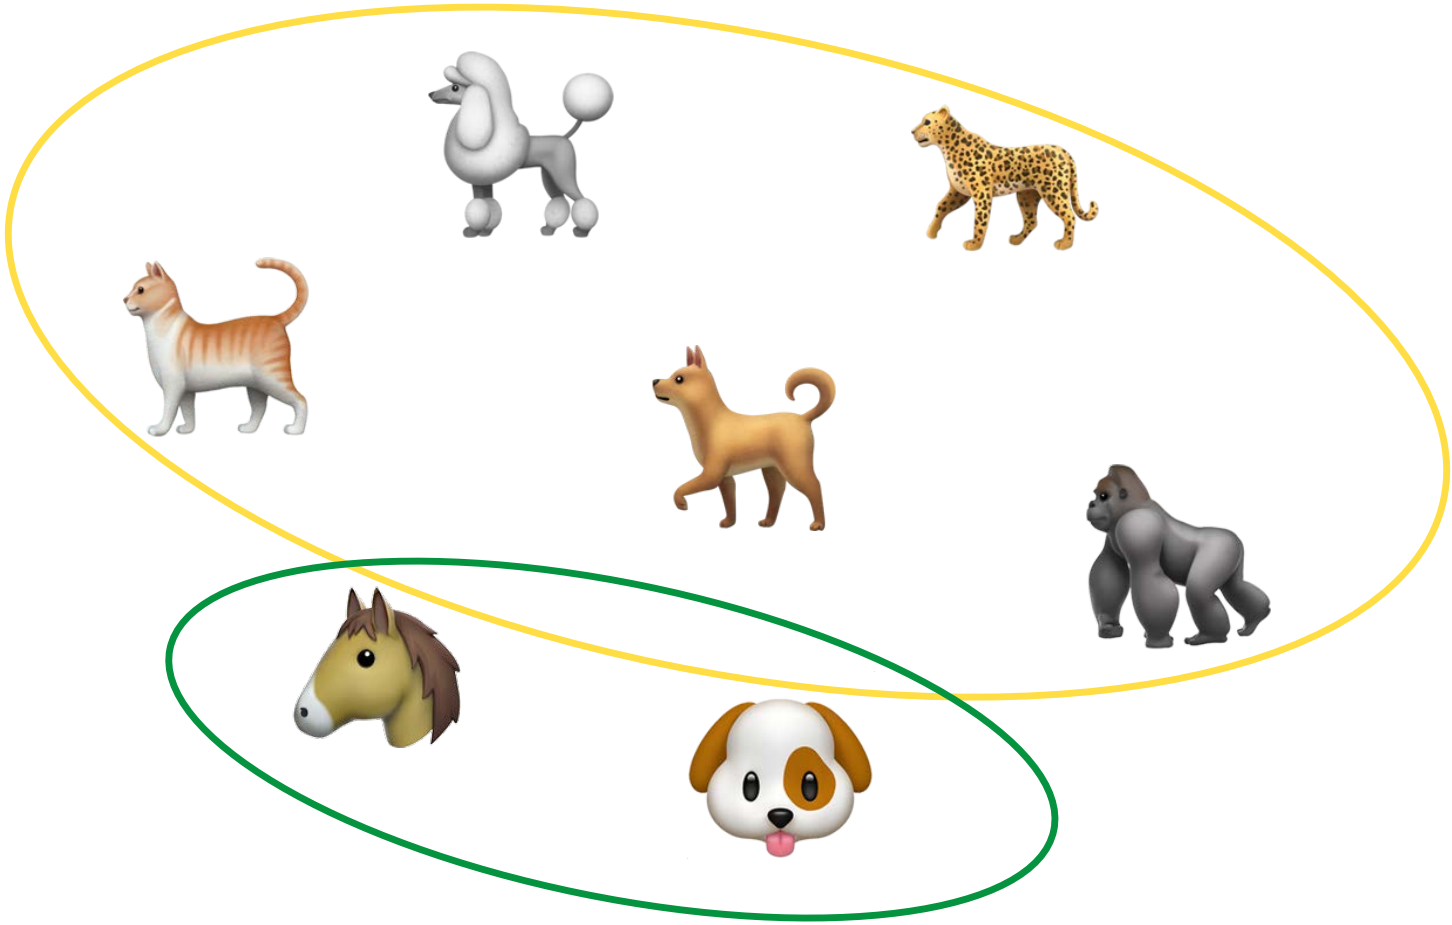

Types of ML Problems

What can we do with binned data?

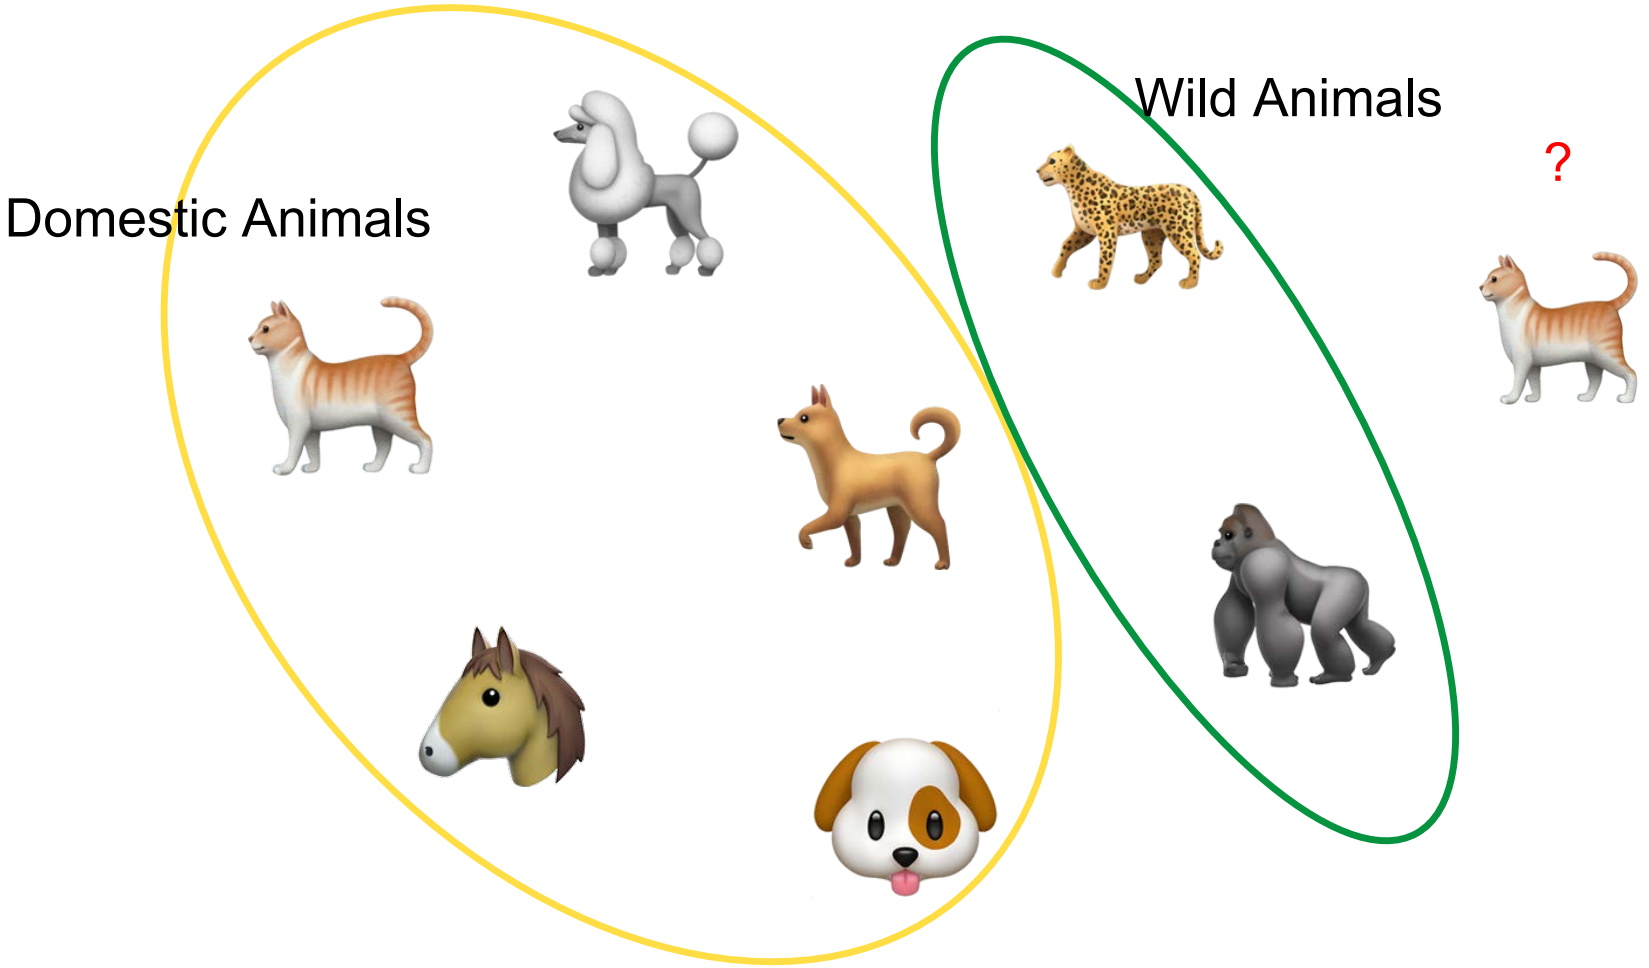

Types of ML Problems

What can we do with binned data?

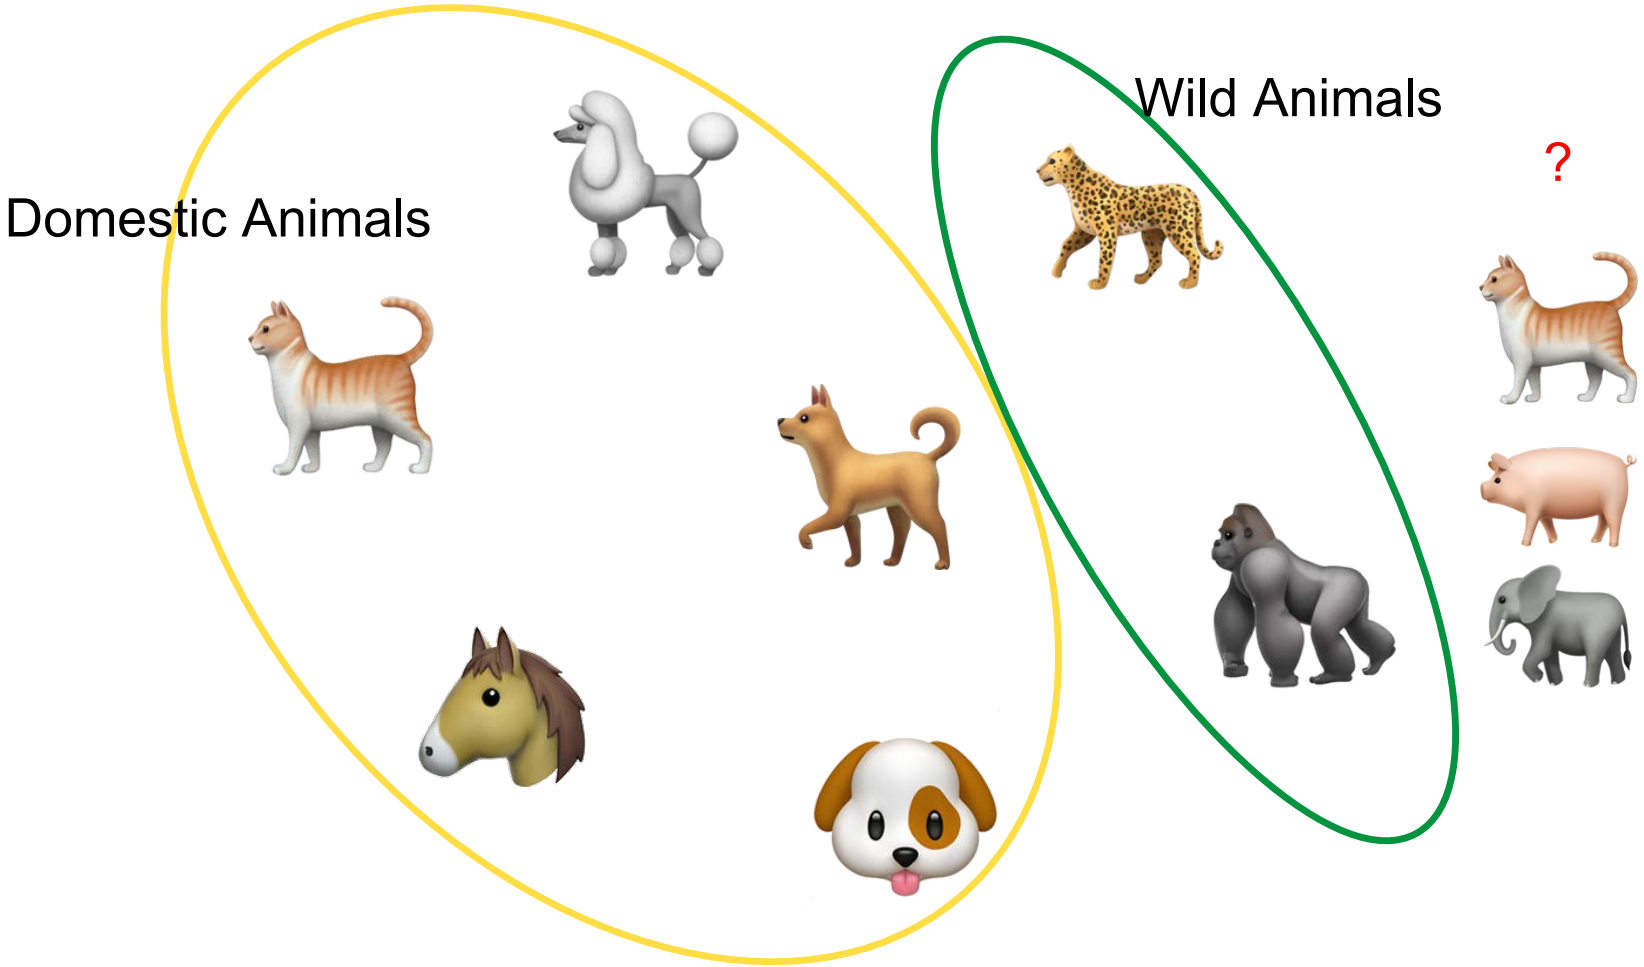

Learning Strategies

Supervised Learning

Unsupervised Learning

Supervised Learning

“Supervised learning is typically done in the context of classification, when we want to map input to output labels, or regression, when we want to map input to a continuous output. Common algorithms in supervised learning include logistic regression, naive bayes, support vector machines, artificial neural networks, and random forests. In both regression and classification, the goal is to find specific relationships or structure in the input data that allow us to effectively produce correct output data. Note that “correct” output is determined entirely from the training data, so while we do have a ground truth that our model will assume is true, it is not to say that data labels are always correct in real-world situations. Noisy, or incorrect, data labels will clearly reduce the effectiveness of your model.”

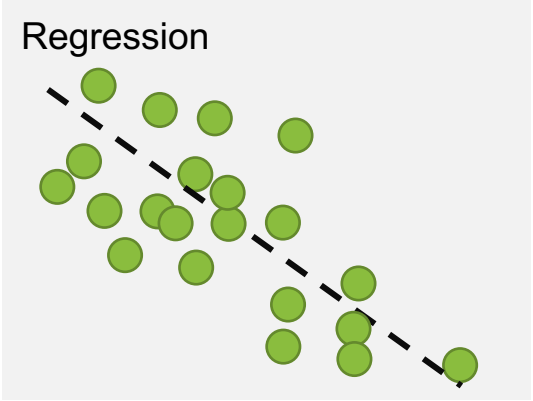

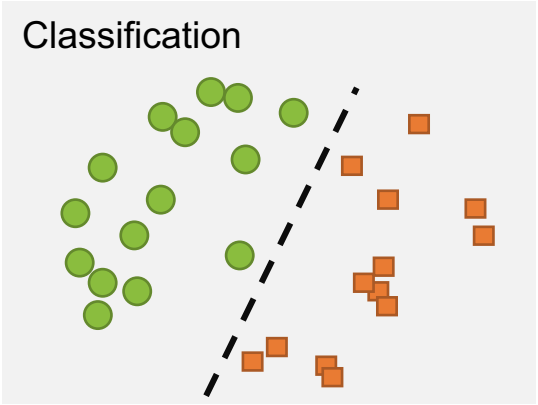

https://towardsdatascience.com/supervised-vs-unsupervised-learning-14f68e32ea8d

Unsupervised Learning

“The most common tasks within unsupervised learning are clustering, representation learning, and density estimation. In all of these cases, we wish to learn the inherent structure of our data without using explicitly-provided labels. Some common algorithms include k-means clustering, principal component analysis, and autoencoders. Since no labels are provided, there is no specific way to compare model performance in most unsupervised learning methods.

Two common use-cases for unsupervised learning are exploratory analysis and dimensionality reduction.”

https://towardsdatascience.com/supervised-vs-unsupervised-learning-14f68e32ea8d

Learning Strategies

	Supervised Learning	Unsupervised Learning
Discrete	Classification or Categorization	Clustering
Continuous	Regression	Dimensionality reduction

Classification

Classification

Supervised Learning

SVM Example

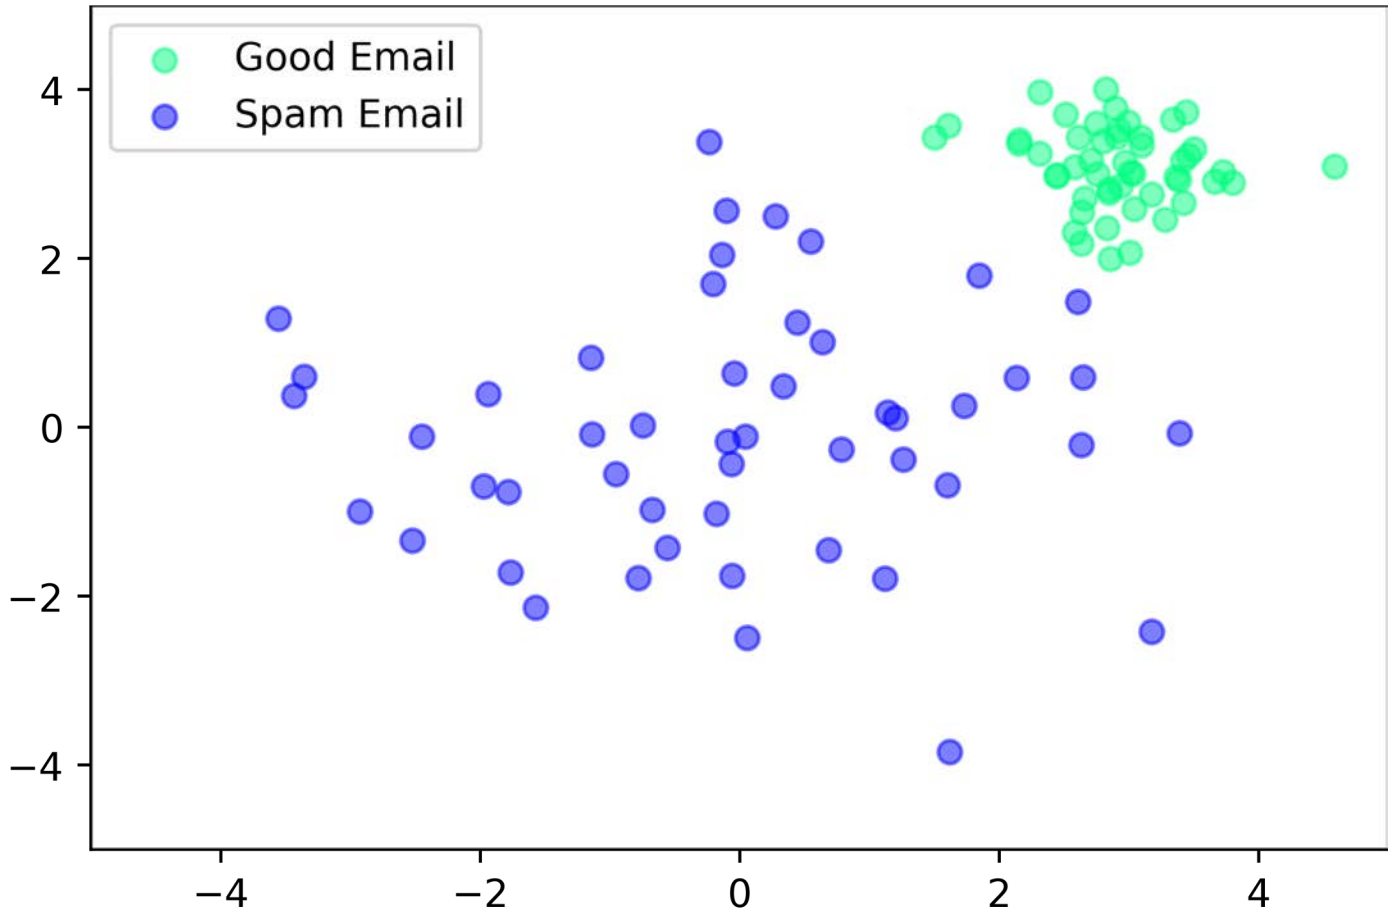

Supervised Learning

SVM Example

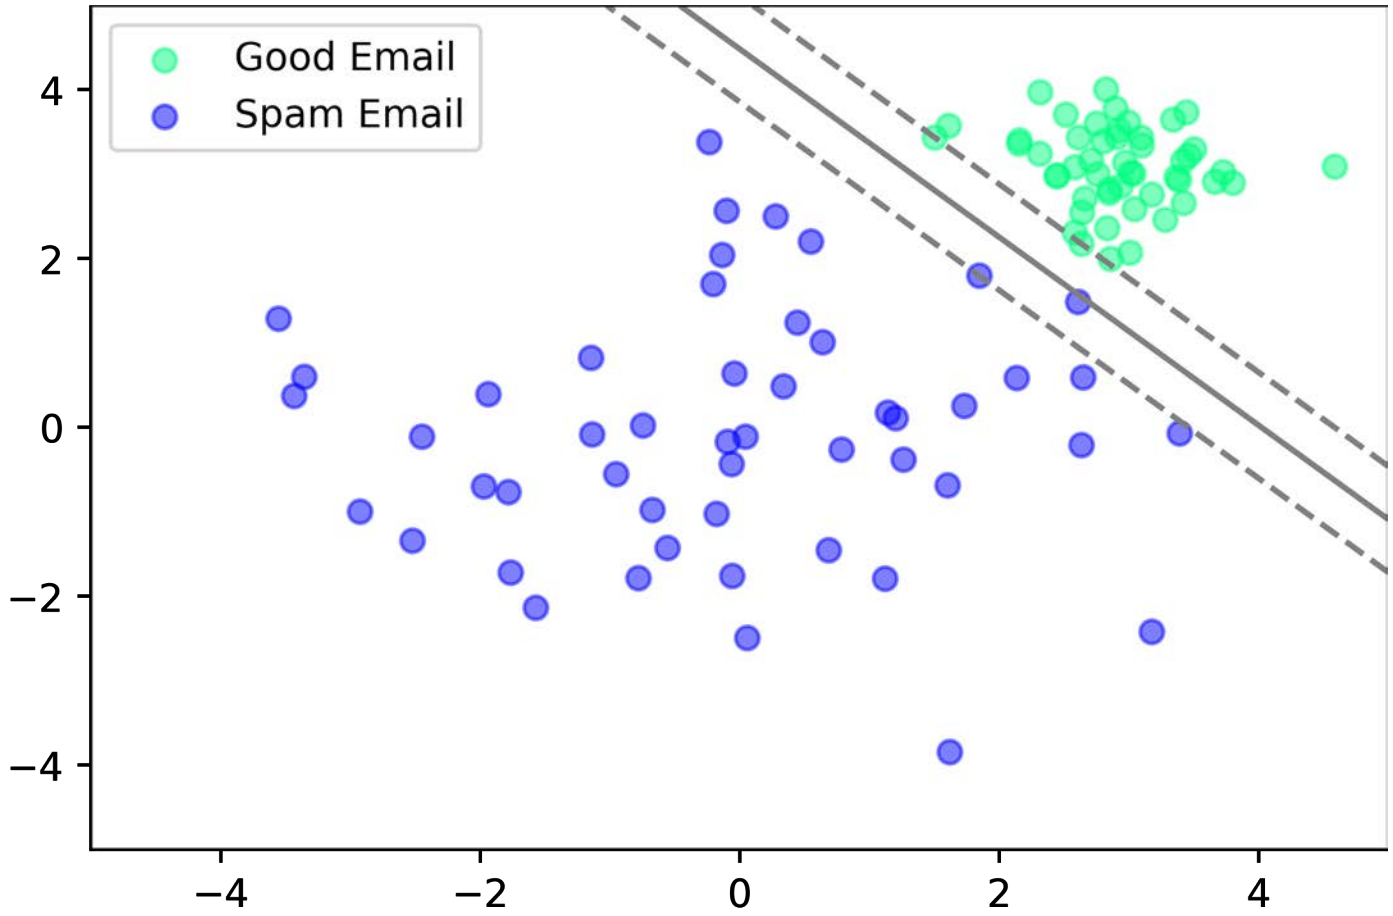

Classification

Example: Support Vector Machines (SVM)

- Determine whether a given data point belongs to a certain class or not
- Supervised method:
  - First training a classifier model on data points for which the class is known (e.g. a set of emails that are labeled as spam or not spam)
  - Then use the model to determine the class of new, unseen data-points
- Find the boundary line that separates two classes
- The boundary line should create a maximum separation between the classes

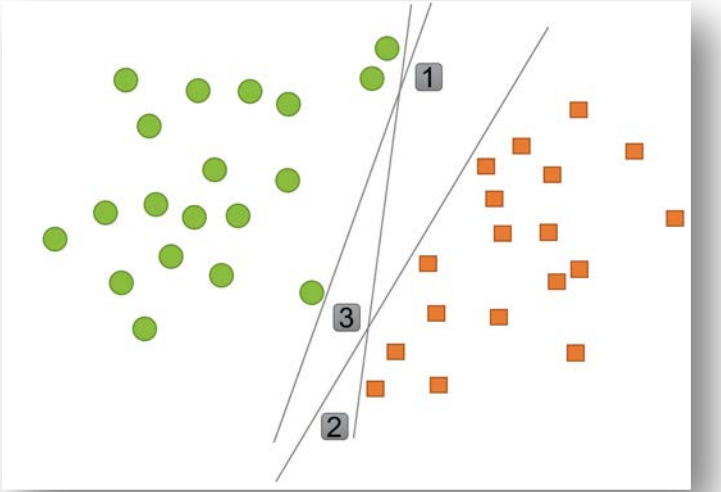

https://www2.deloitte.com/nl/nl/pages/data-analytics/articles/part-2-artificial-intelligence-techniques-explained.html#

Classification

Example: Support Vector Machines (SVM)

- Find the boundary line that separates two classes
- The boundary line should create a maximum separation between the classes

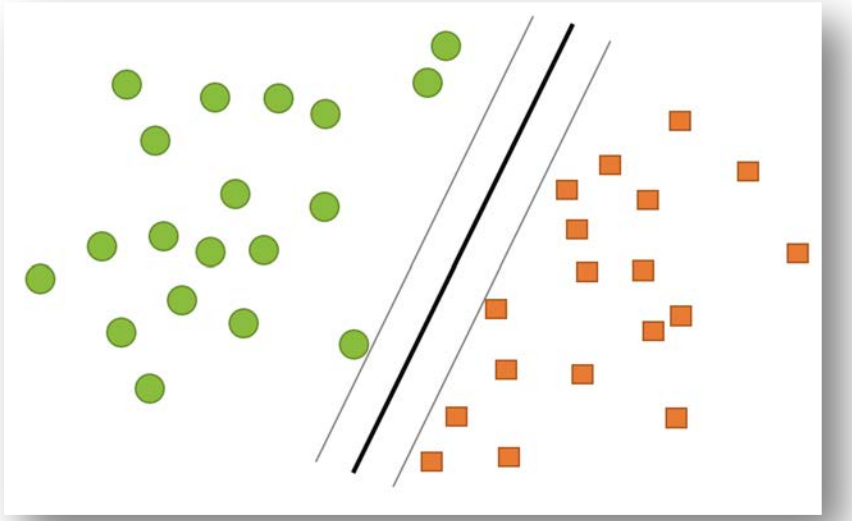

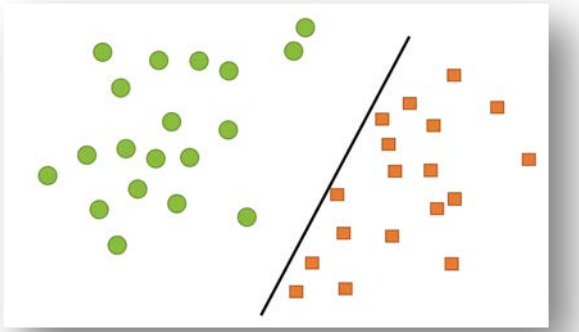

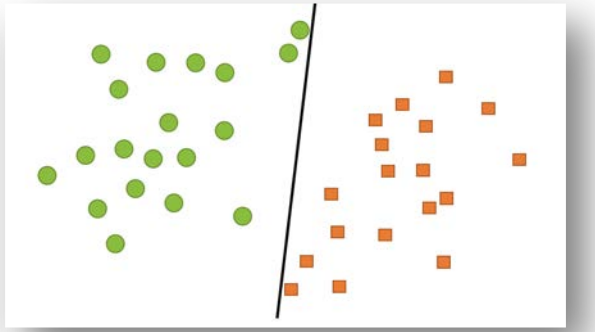

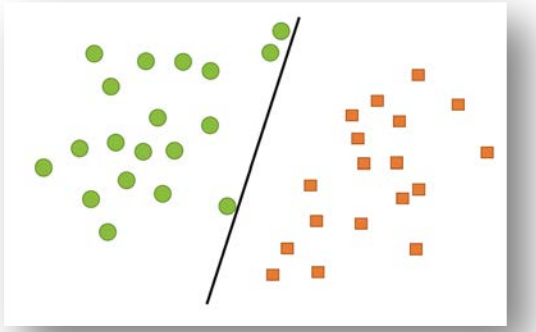

Supervised Learning

SVM Example

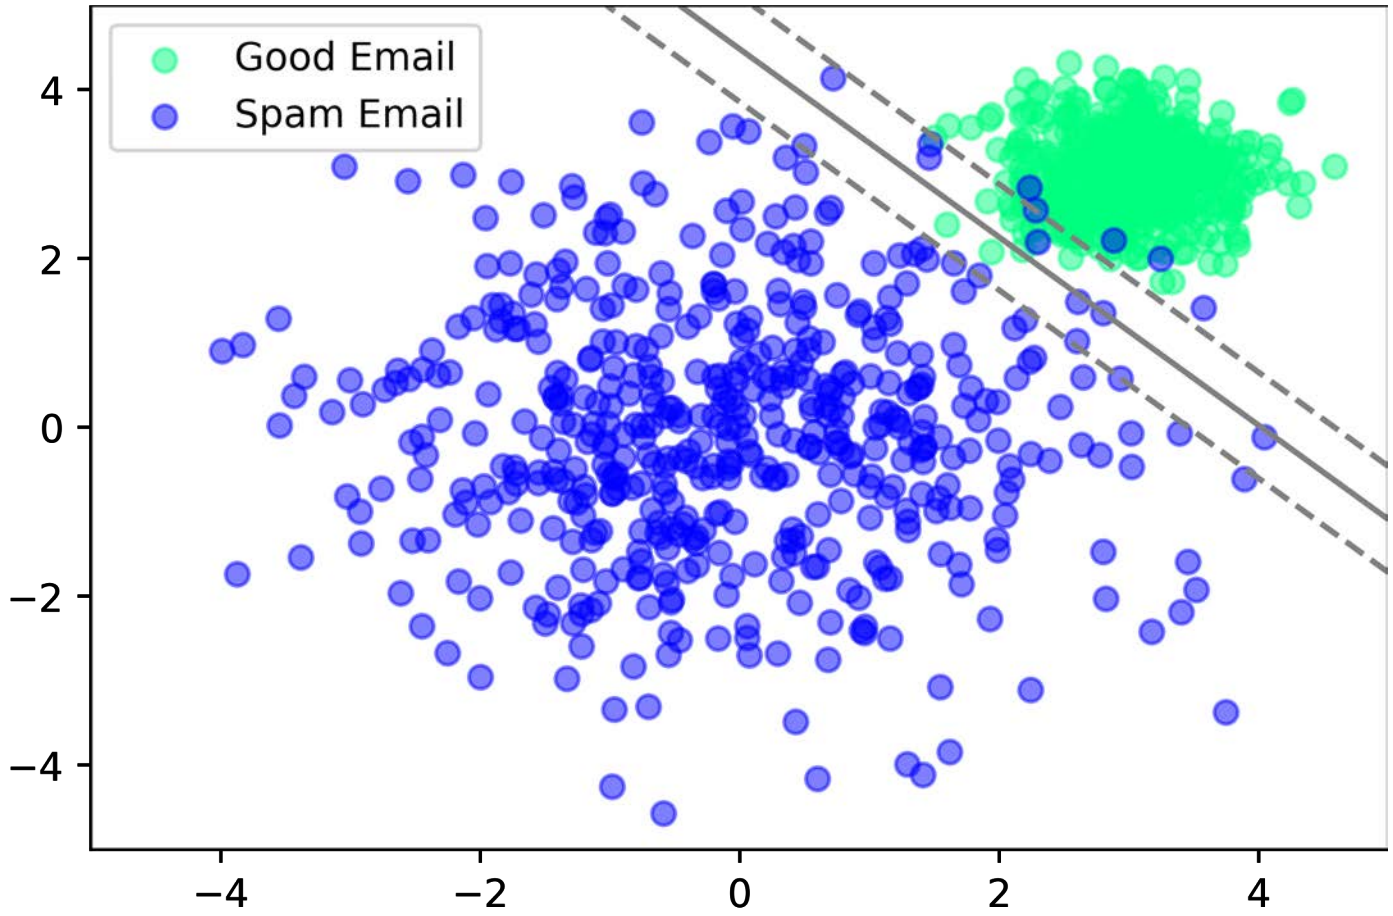

Classification

SVM Example

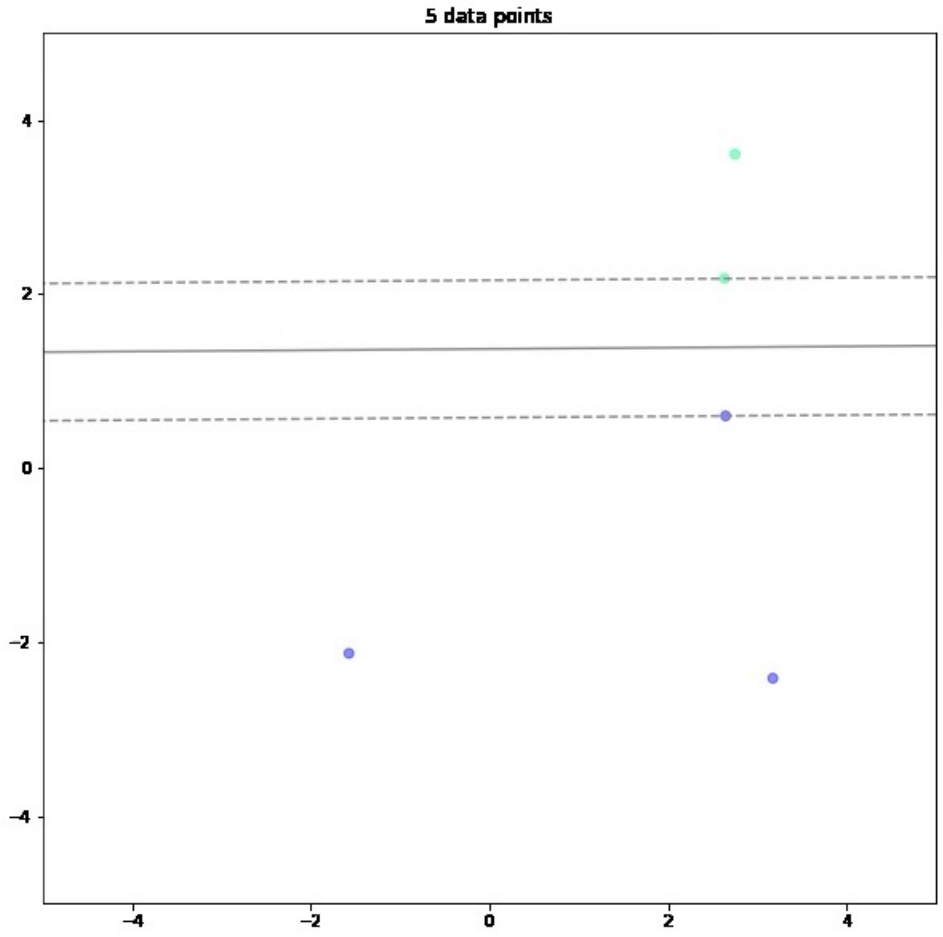

Supervised Learning

Supervised Learning

- Data

- Labels (Classes)
- Continues values

Regression

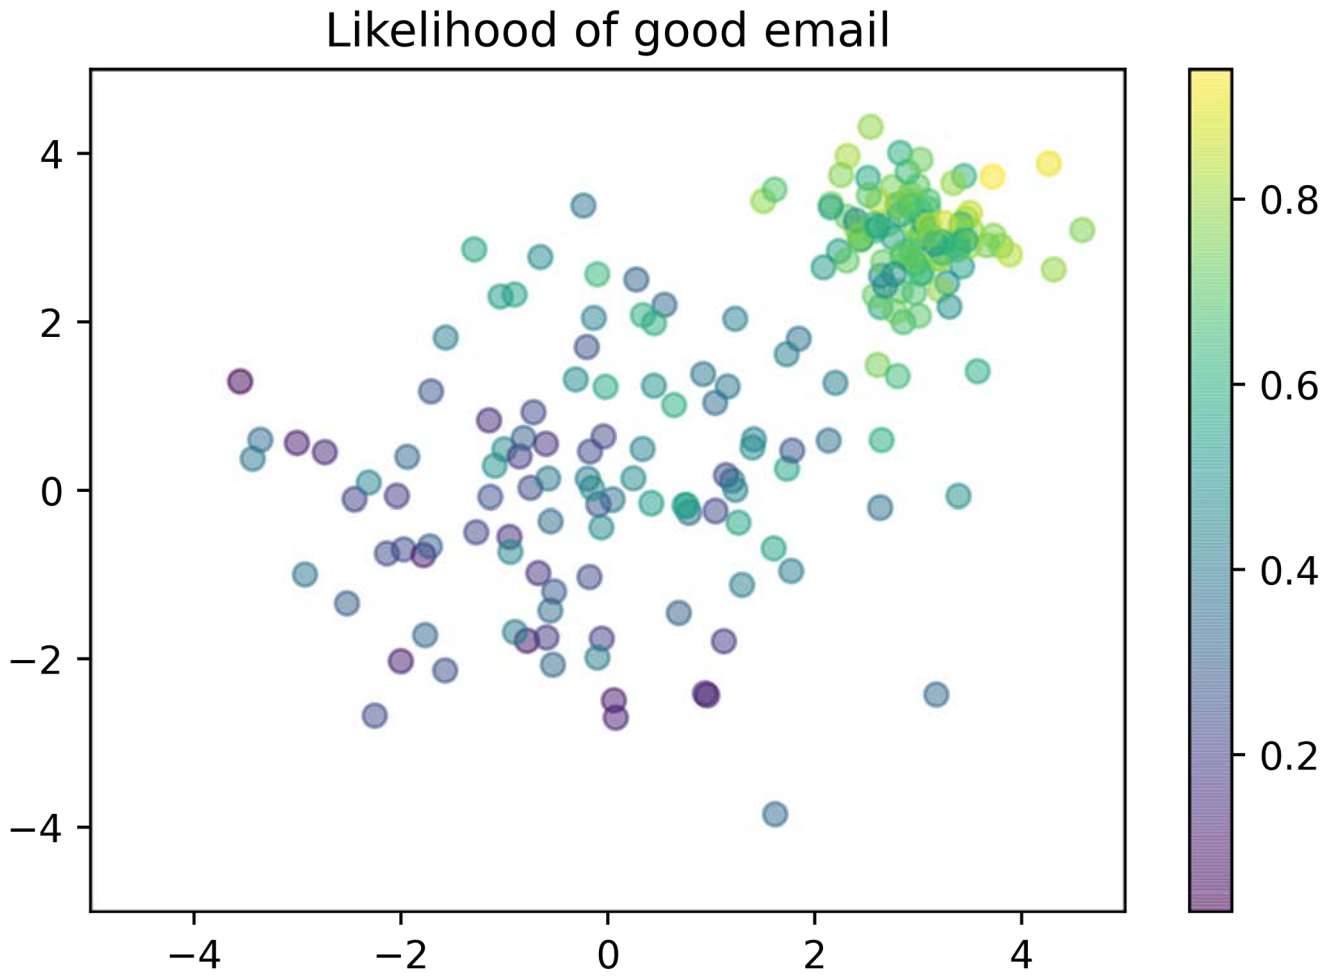

Regression

Linear Support Vector Regression (SVR)

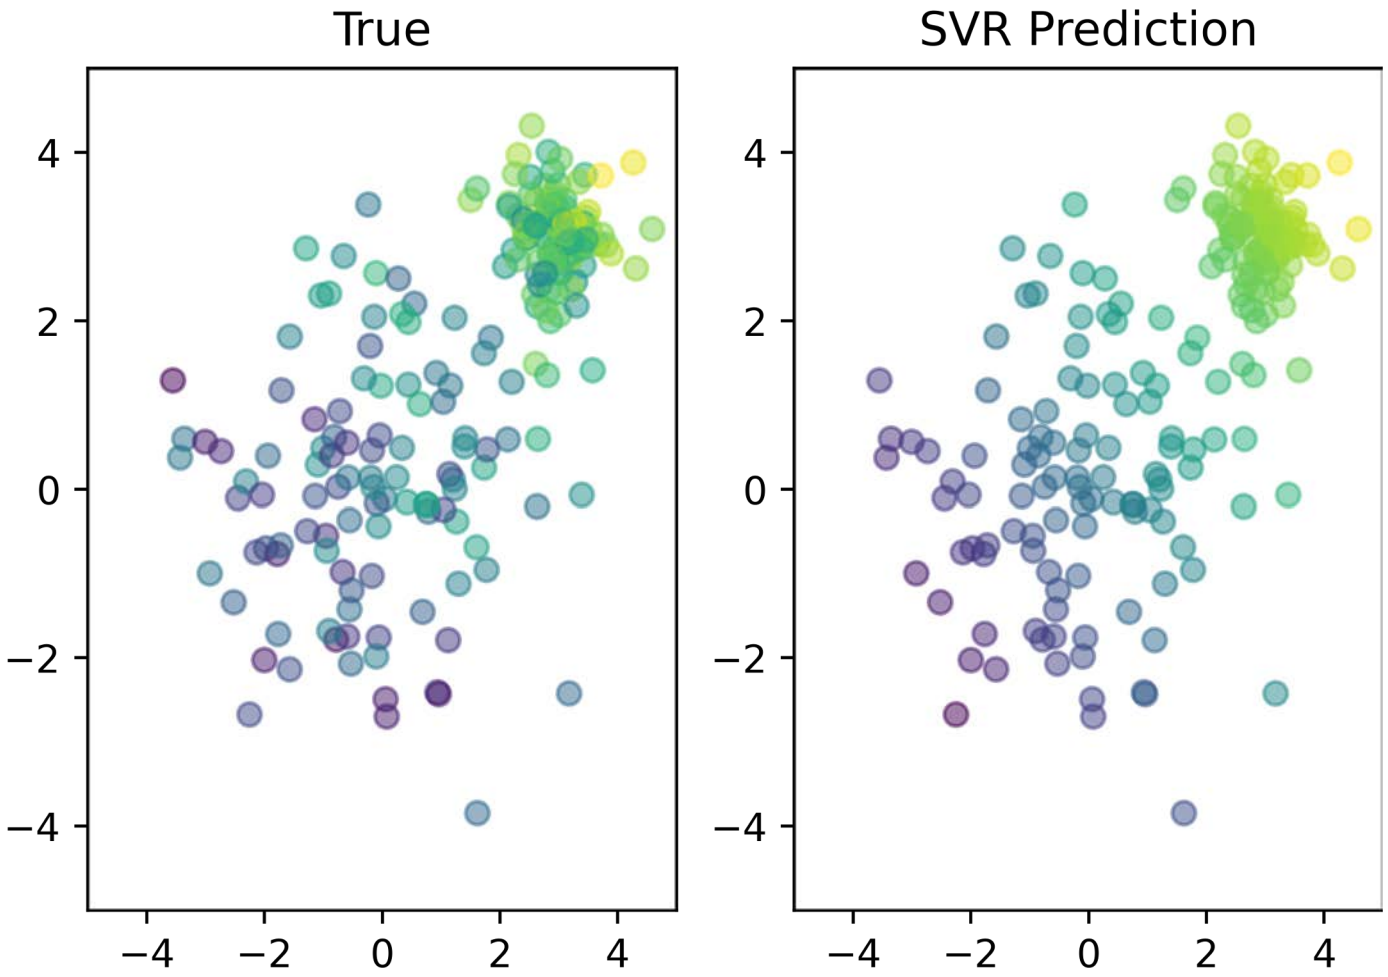

Regression

Linear Support Vector Regression (SVR)

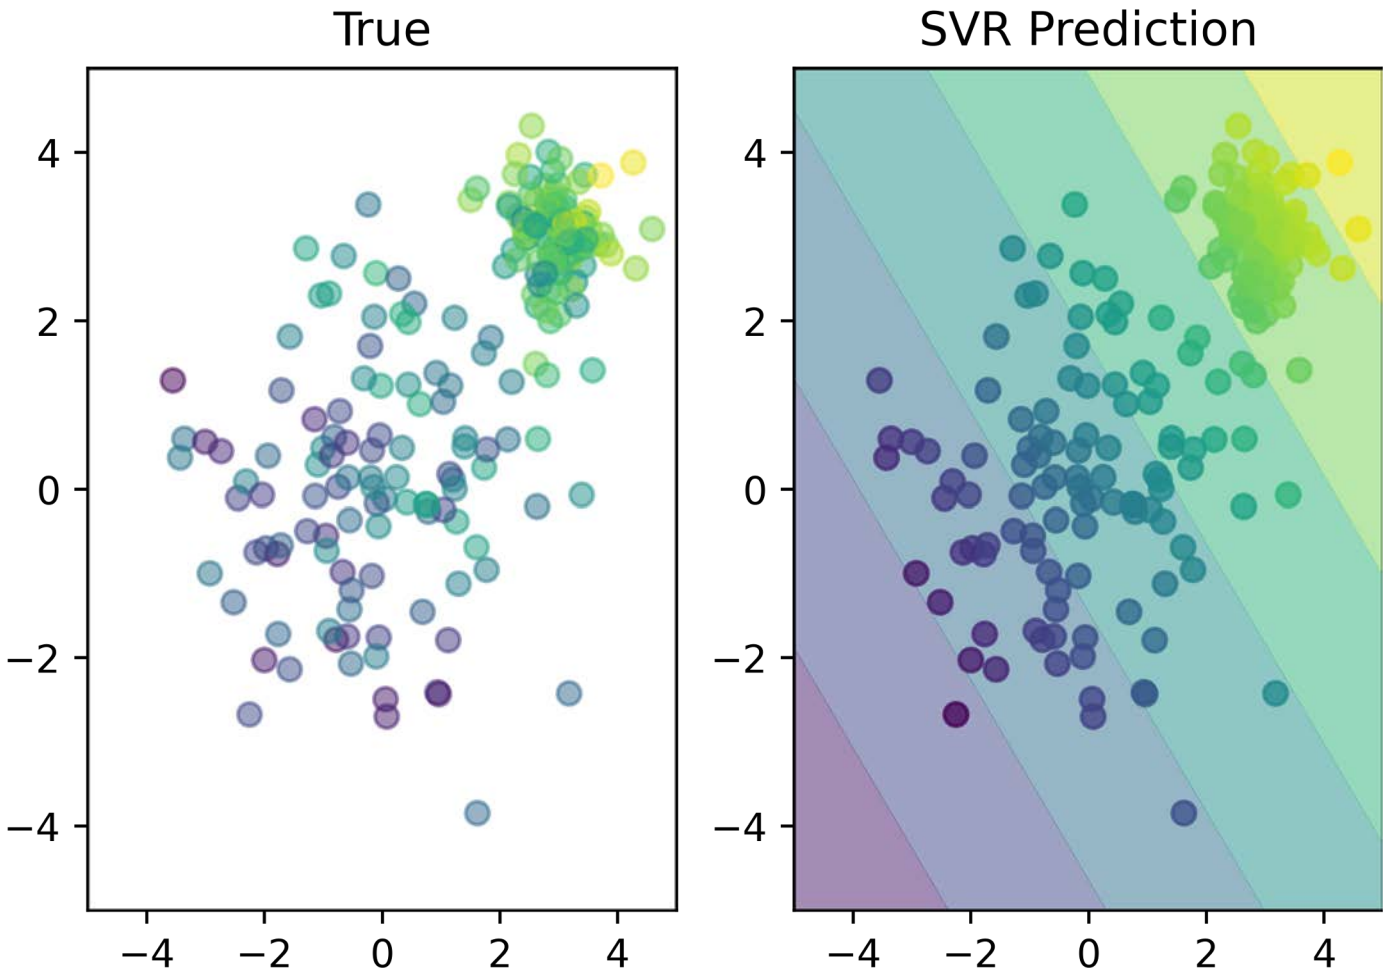

Regression

RBF Support Vector Regression (SVR)

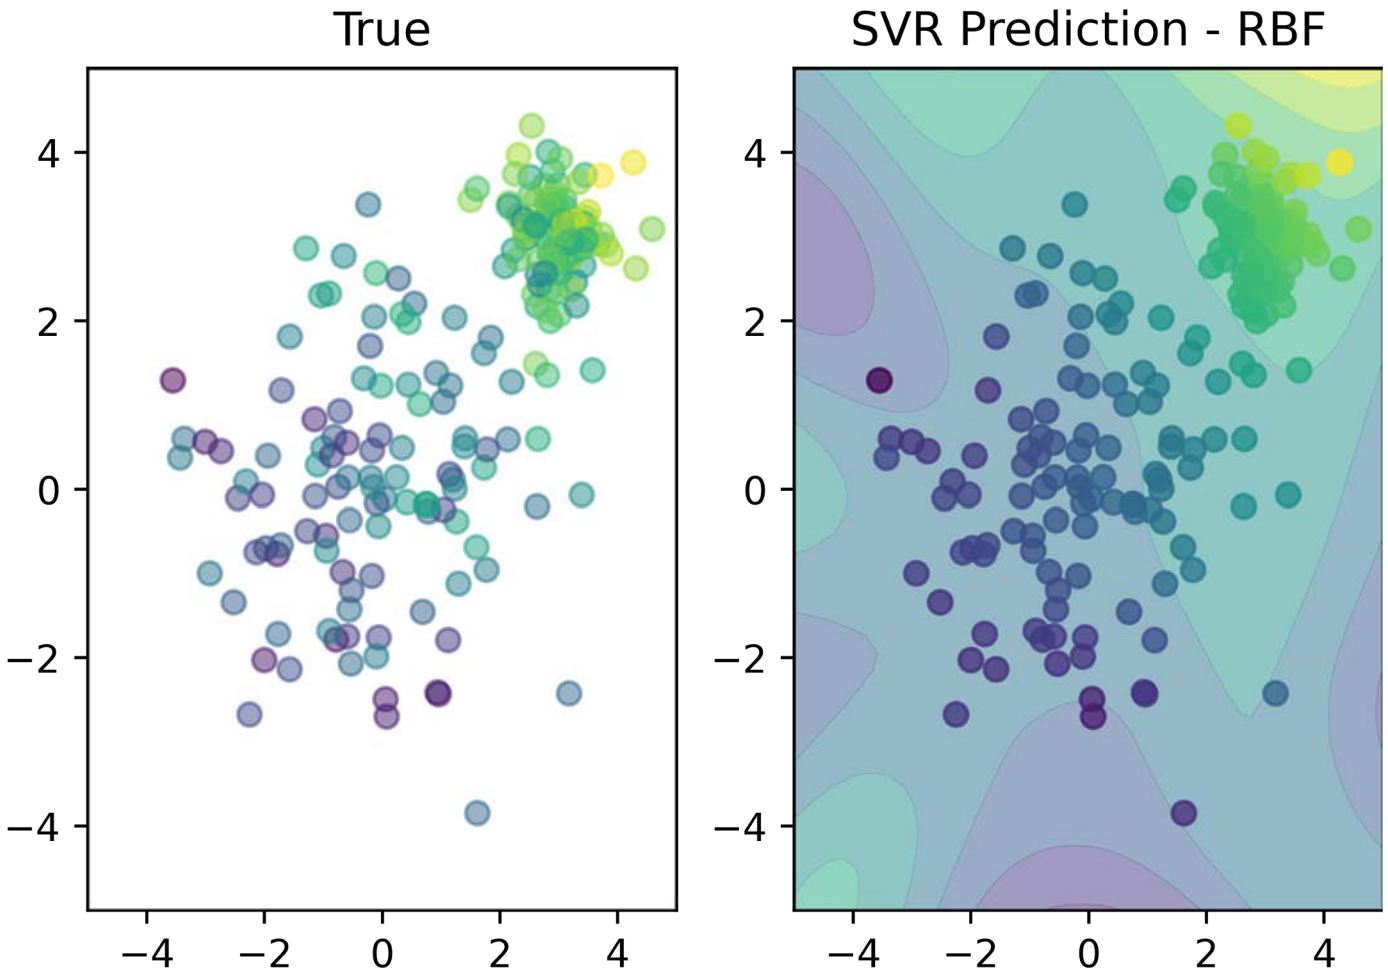

Unsupervised Learning

Unsupervised Learning

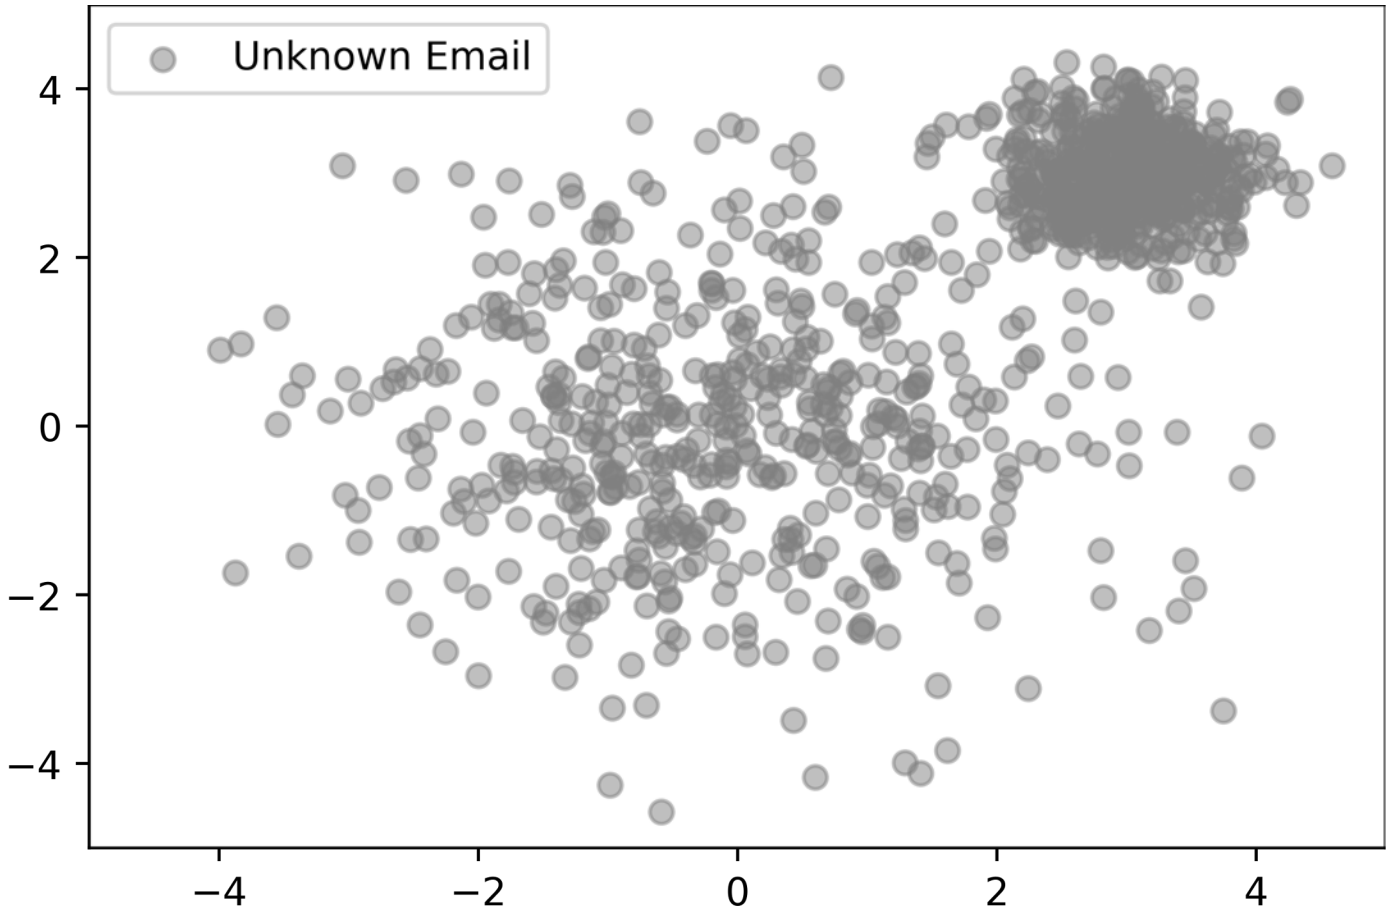

Unsupervised Learning

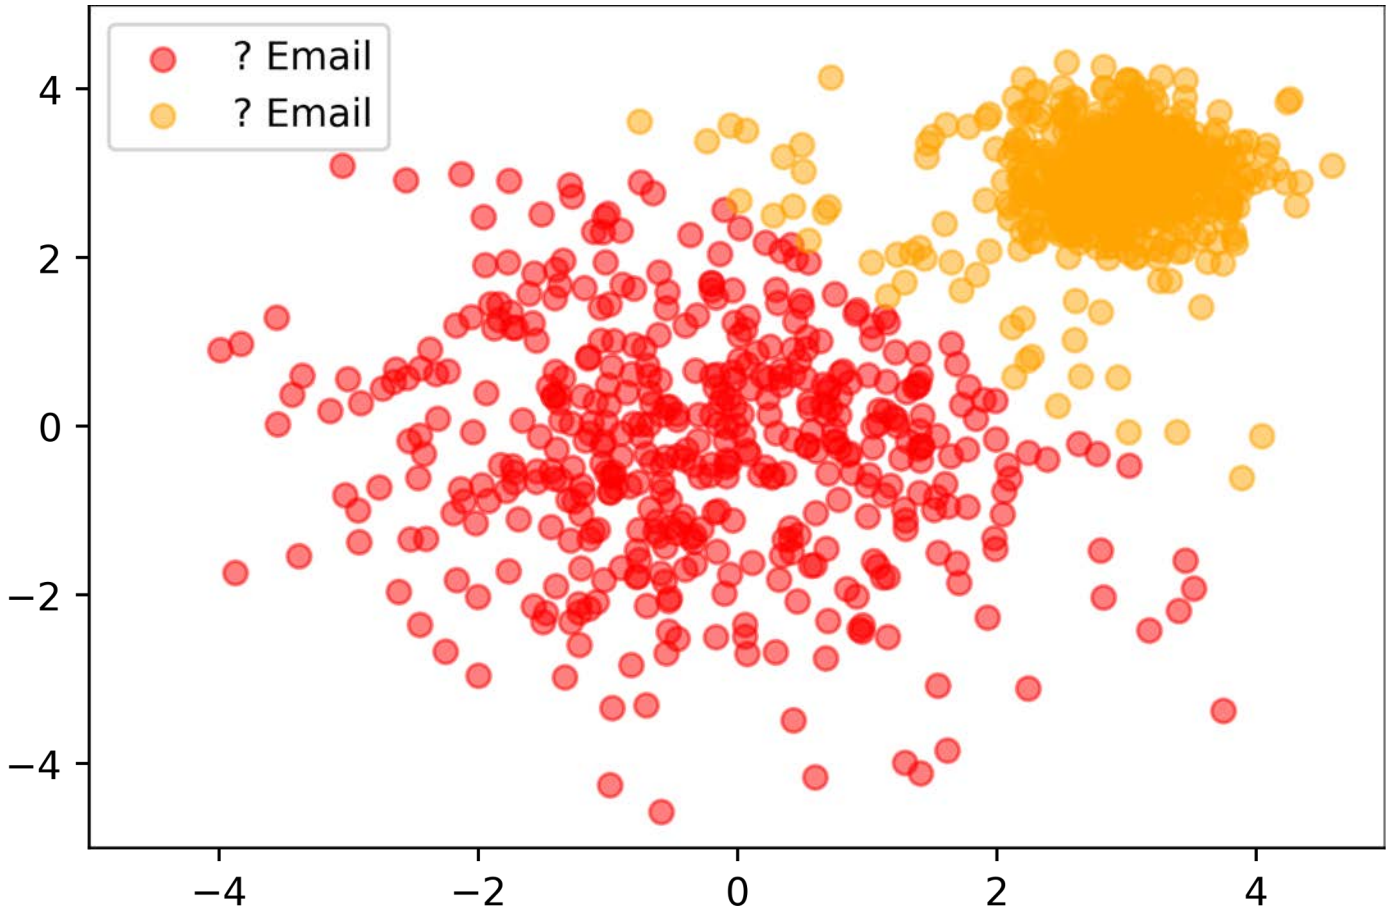

Clustering

K-means clustering

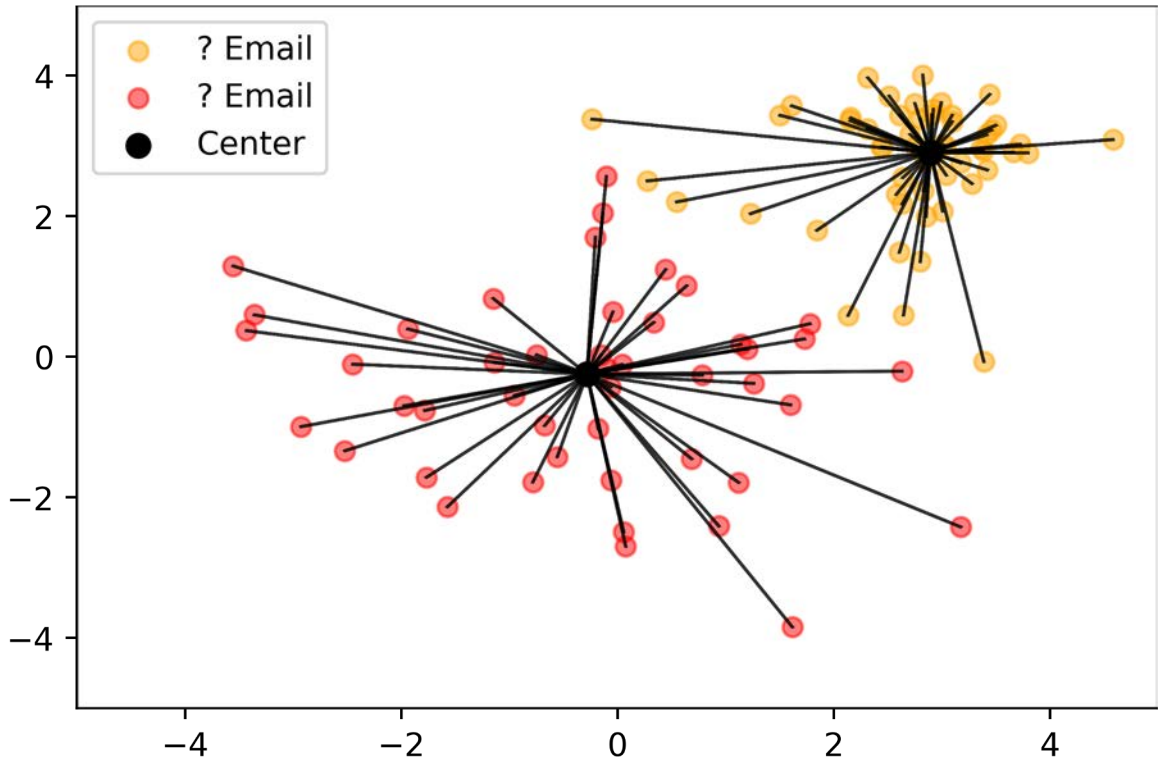

Clustering

K-means clustering

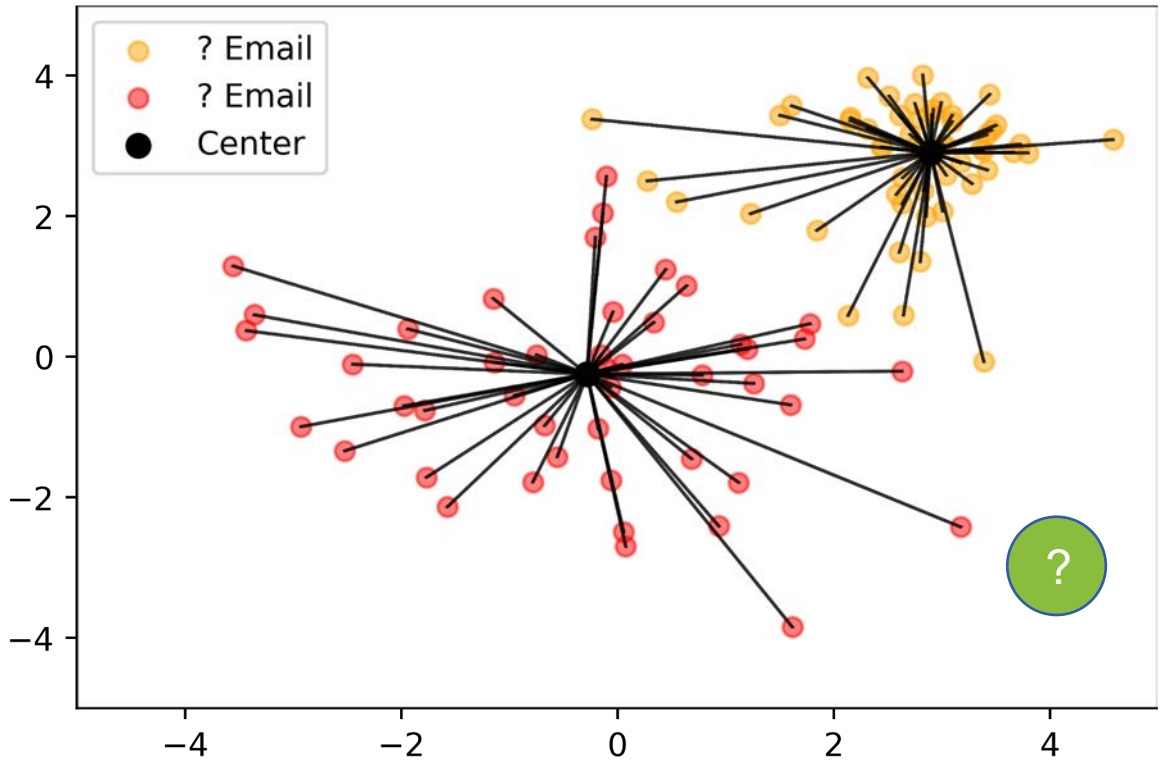

Clustering

Clust

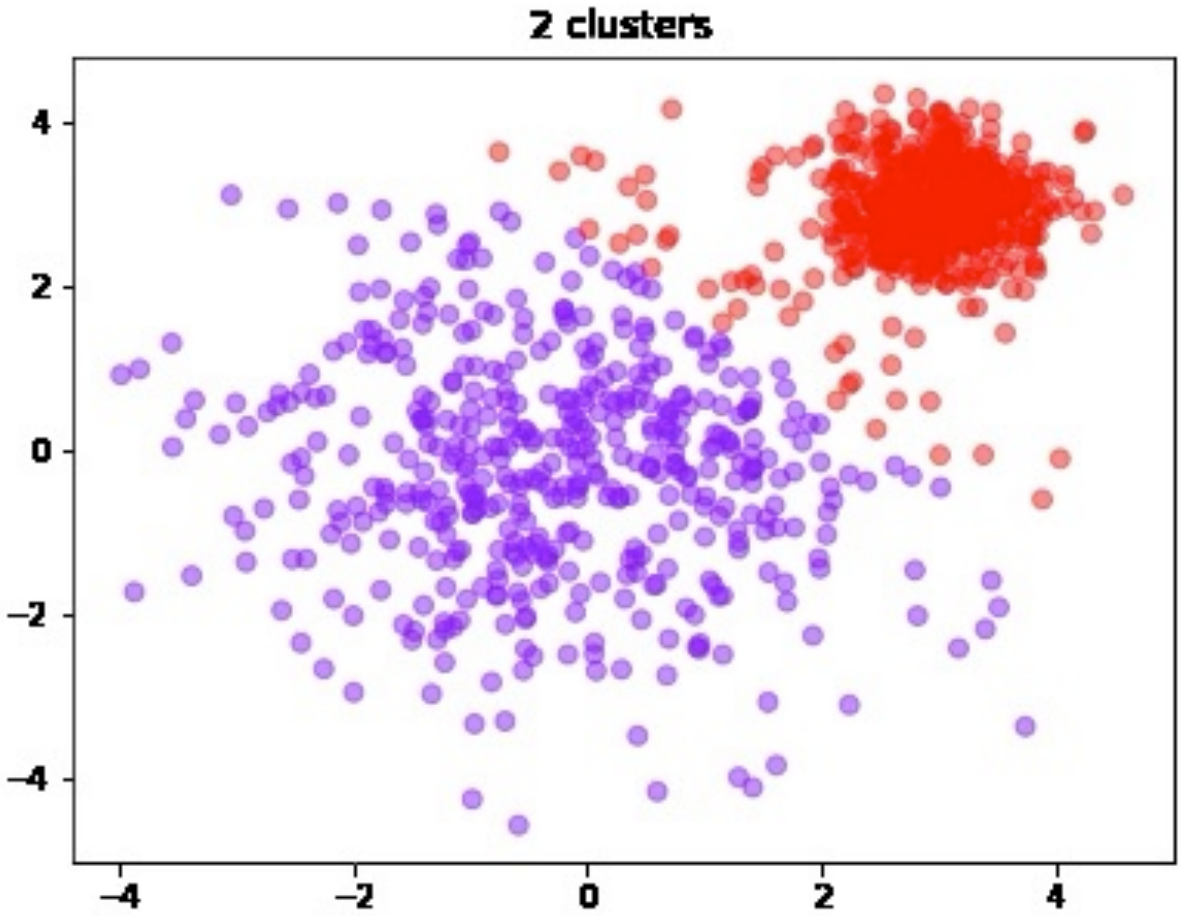

Dimensionality Reduction

Principal Component Analysis (PCA)

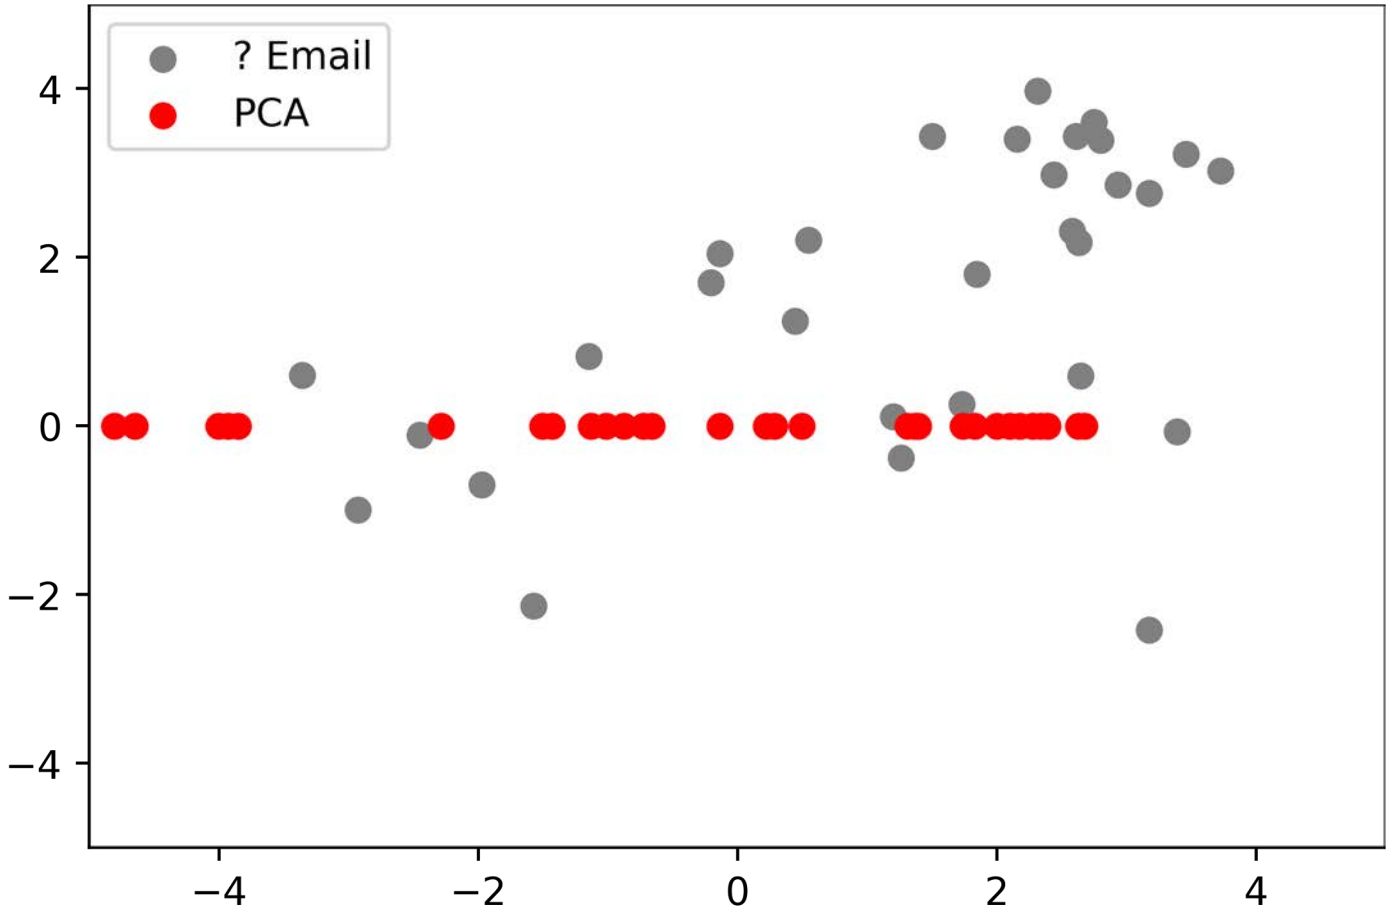

Dimensionality Reduction

Principal Component Analysis (PCA)

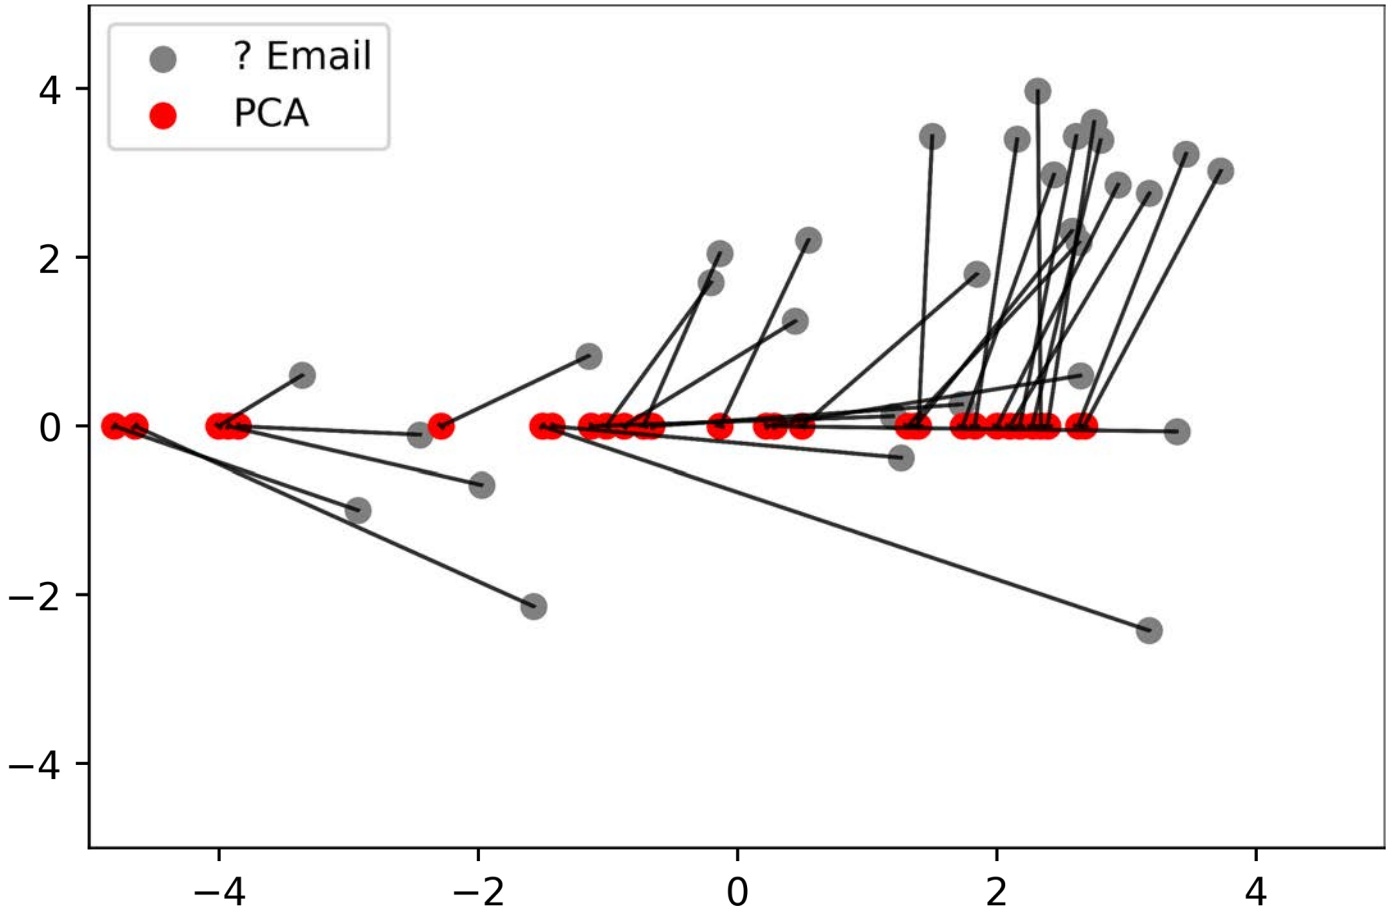

Dimensionality Reduction

Principal Component Analysis (PCA)

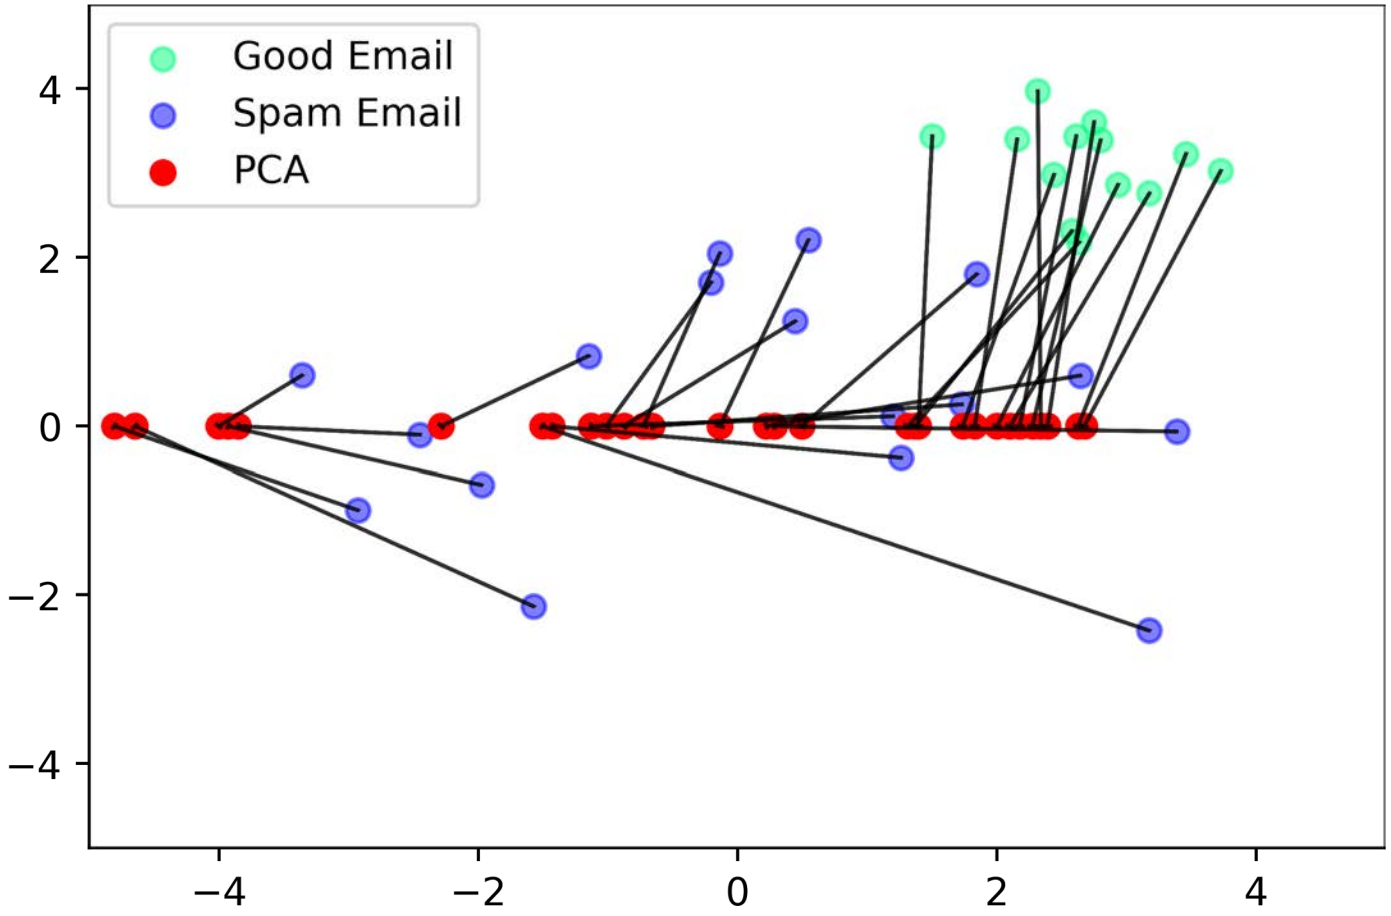

Live Coding Example 2

kMeans in Python

Unsupervised Learning

- Data

Not needed:
- Labels (Classes)
- Continues values

Supervised Learning vs. Unsupervised Learning

	Supervised Learning	Unsupervised Learning
Discrete	Classification or Categorization	Clustering
Continuous	Regression	Dimensionality reduction

https://towardsdatascience.com/supervised-vs-unsupervised-learning-14f68e32ea8d

Implementing IUIs

Practical advice

Implementing IUIs

Libraries, Python, web services, …

- Build your own, e.g.:
  - Web-based frontend
  - Python backend
    (e.g. using ML libraries)

Pro:
Full flexibility,
integrate own models or
models from others etc.

Con:
More development work,
computational costs

- Pre-built, e.g.:
  - External APIs, web-services
  - Devkits, e.g.
    https://developers.google.com/ml-kit

Pro:
Faster prototyping

Con:
…if it fits your needs;
API costs

Implementing IUIs

ML models

- HCI & user-centred work is often iterative, uses prototyping
- Training e.g. state-of-the-art text or image generator from scratch is costly (time + computation)

- Prototyping with pretrained models
  - NLP e.g.: https://huggingface.co/
  - Other models e.g.: https://www.tensorflow.org/resources/models-datasets
  - Many (ML/AI) papers come with code/model releases (e.g. https://paperswithcode.com/)

- Also: „Buying a time machine“ for IUI research
  - Models don‘t need to run on your target device

cf. Hudson and Mankoff 2014

Some Resources

Frameworks, pretrained models, etc.

- TensorFlow Hub, Model Garden, TensorFlow.js models: https://www.tensorflow.org/resources/models-datasets
- p5.js & ml5.js: https://p5js.org/ & https://ml5js.org/
- Keras: https://keras.io/
- PyTorch: https://pytorch.org/
- Fast AI: https://docs.fast.ai/
- MLKit: https://developers.google.com/ml-kit
- Scikit-learn: https://scikit-learn.org/stable/index.html
- Hugging Face (NLP): https://huggingface.co/
- Teachable Machine: https://teachablemachine.withgoogle.com/
- Tensorflow lite, js: https://www.tensorflow.org/lite/, https://www.tensorflow.org/js
- Conversational UIs: https://rasa.com/docs/rasa/, https://developers.google.com/learn/pathways/chatbots-dialogflow

Cloud Services
Example: Microsoft Cognitive Services

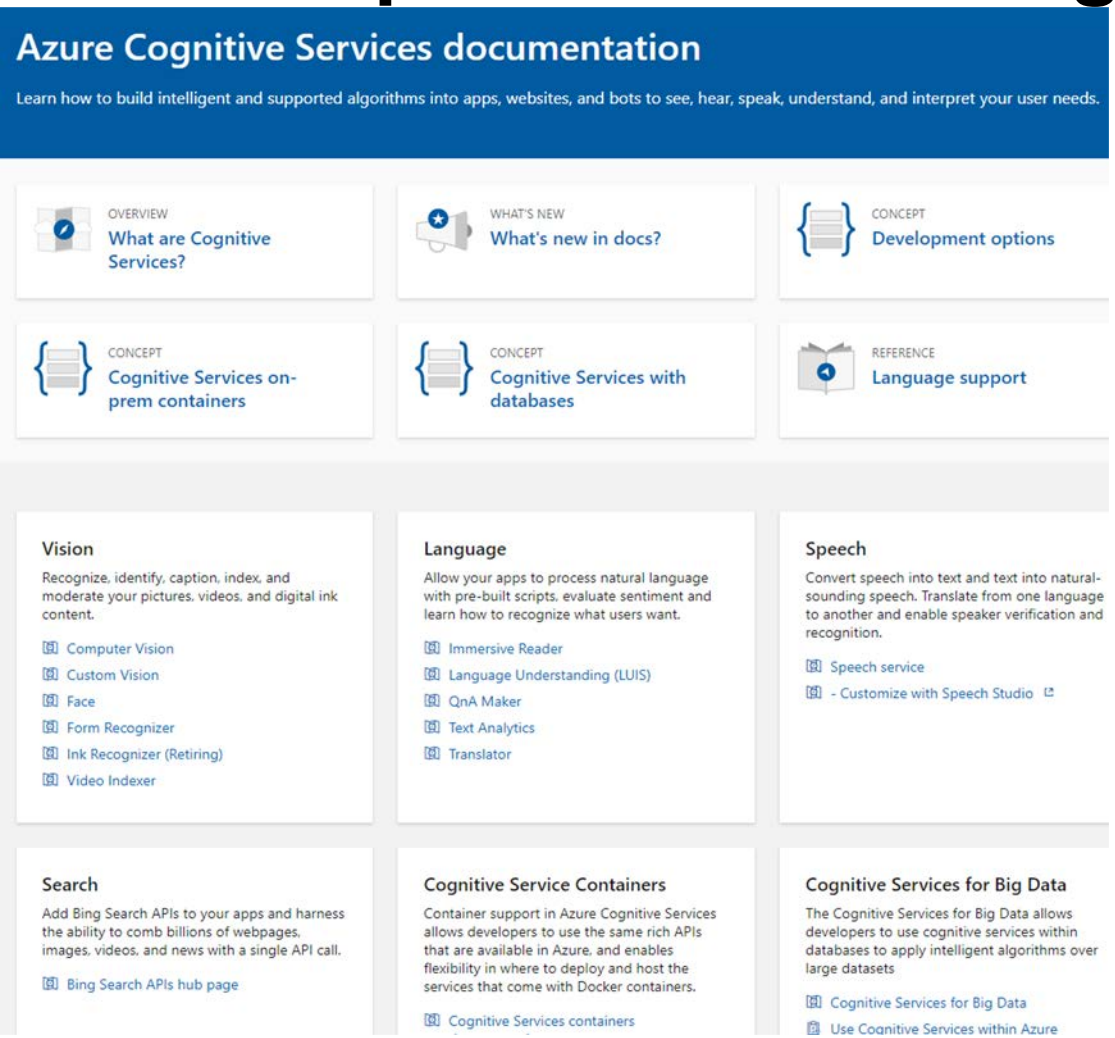

https://docs.microsoft.com/en-us/azure/cognitive-services/

Cloud Services from:

- IBM
- Google
- Amazon
- Microsoft
- … and many others

Why do people use them?

What is the risk?

Examples Text Analytics

Identification of the Language

- Can tell what language the text is, e.g. English, German, Spanish,…
- Relevant for understanding and translation
- Example (Online) APIs:
  - https://console.bluemix.net/apidocs/language-translator
  - https://docs.microsoft.com/en-us/azure/cognitive-services/translator/
  - https://cloud.google.com/translate/docs/basic/detecting-language
  - https://pypi.org/project/langdetect/

To detect the language of the text:

```
>>> from langdetect import detect
>>> detect("War doesn't show who's right, just who's left.")
'en'
>>> detect("Ein, zwei, drei, vier")
'de'
```

To find out the probabilities for the top languages:

```
>>> from langdetect import detect_langs
>>> detect_langs("Otec matka syn.")
[sk:0.572770823327, pl:0.292872522702, cs:0.134356653968]
```

How to Install Python

- Python via Anaconda www.anaconda.com
- Python with pip https://www.python.org/downloads/
- Python in the web via Google Colab https://colab.research.google.com

License