

**DINFK**

# Introduction to Machine Learning

Lecture 1

Fanny Yang, Andreas Krause


**ETH** zürich



# Before we start ... eduApp installed?


Go to WebApp [eduapp.ethz.ch](https://eduapp.ethz.ch) (mobile app doesn't work anymore) and log in with ETH credentials

**ETH** zürich EduApp

 Dashboard

## COURSE ACTIVITIES

 Schedule

 Course Clicker

## Upcoming Lectures

TODAY

**V: Introduction to Machine Learning**


14:15–16:00, ETA F 5

Which departments are you from?

Why do you want to take  
Intro to Machine Learning?

# What is the goal of machine learning?

Goal: develop machines to perform complex tasks. Example tasks:



“Who can take this course?”

“Students who know risk assessment from what eigenvalues are” physical measurements

Complete blood count		Serum chemistry			
WBC	5400/ $\mu$ L	TP	6.7 g/dL	Na	137 mEq/L
RBC	443 $\times 10^3$ / $\mu$ L	Alb	4.0 g/dL	K	4.4 mEq/L
Hb	12.6 g/dL	T-Bil	0.6 mg/dL	Cl	102 mEq/L
Ht	35.50%	D-Bil	0.1 mg/dL	Ca	9.1 mg/dL
Plt	11.7 $\times 10^3$ / $\mu$ L	BUN	18 mg/dL	CSP	0.55 mg/dL
		Cr	0.90 mg/dL	TC	112 mg/dL
Blood coagulation test		LDH	209 IU/L	HDL-C	16 mg/dL
PT (INR)	1.1	CK	70 IU/L	LDL-C	15 mg/dL
PT	82.80%	AST	21 IU/L	TG	96.3 mg/dL
APTT	39.6 s	ALT	19 IU/L	SIL-2R	281 U/ml



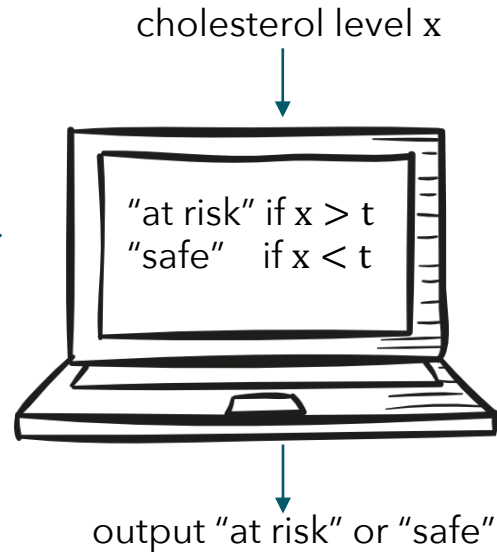
classify image



anomaly or not?

# Classification: The traditional pipeline

Classical approach: Machine performs tasks following specific rules **defined by humans**



# Classification: What would you do now?

Possible LLM era approach: Given **example data**, you **could** feed that and your own cholesterol level  $x$  to an LLM to ask it to predict your risk



If you let them **explain how they did it**, they'll tell you that under the hood they followed some **machine learning pipeline...**

Given examples (145, 0), (172, 0), (181, 0), (195, 0), (207, 1), (218, 1), (225, 1), (237, 1), (249, 1), (260, 1), predict the label for 200

The predicted label for a cholesterol level of **200 mg/dL** is **0 (Low Risk)** based on the trained logistic regression model.

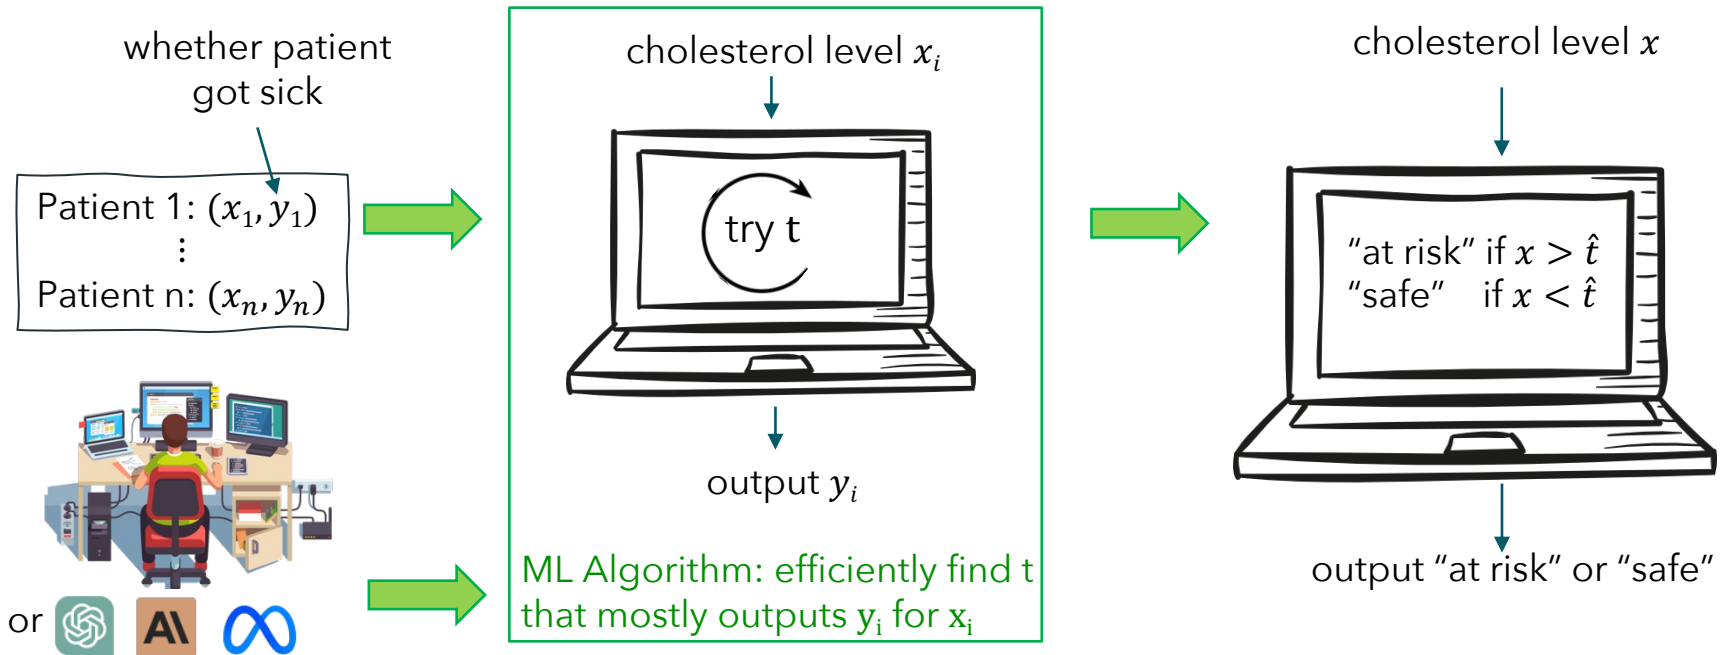
explain how you got there

### Summary of How I Predicted the Label for Cholesterol Level 200

- **Dataset Analysis:** Cholesterol levels below 200 mg/dL are mostly labeled **0 (Low Risk)**, and those above are labeled **1 (High Risk)**.
- **Method Choice:** Used **logistic regression**, which models the probability of high risk using a **sigmoid curve**.
- **Model Training:**
  - **X\_train:** Cholesterol levels.
  - **y\_train:** Corresponding labels (0 or 1).
  - **Fit logistic regression model** to find the decision boundary.
- **Prediction Step:** Passed **200 mg/dL** to the trained model.

# Classification: The Machine Learning pipeline

Machine Learning (ML) approach: Machines learn rules using example data



# If an LLM can design the algorithm, what's our role?

why do i have to take a course like "introduction to machine learning" if llms can automatically solve machine learning problems for me? Give a one-paragraph bullet-point response.

- **LLMs don't replace understanding—they amplify it:** you still need to frame the right question, choose success metrics, spot data leakage, and recognize when the model is confidently wrong.
- **Most "ML problems" are really data + evaluation problems:** cleaning data, defining labels, handling imbalance, setting baselines, and designing a sound train/validation/test split are where projects succeed or fail—and LLMs can't reliably infer your ground truth or business constraints.
- **You're accountable for correctness, safety, and compliance:** issues like bias, privacy, security, interpretability, and regulatory requirements (e.g., in hiring, finance, health) require principled reasoning, not just plausible code.
- **Costs and constraints matter:** understanding compute, latency, memory, deployment, monitoring, drift, and failure modes helps you ship something reliable instead of a demo that breaks in production.
- **Knowing fundamentals lets you verify and improve LLM output:** you'll catch subtle bugs (leakage, wrong loss, bad cross-validation), tune models intelligently, and decide when simpler methods beat fancy ones.
- **You'll face cases where LLMs are weak:** novel domains, scarce/noisy labels, shifting distributions, weird edge cases, and strict performance guarantees often need domain expertise and ML judgment.
- **The course teaches durable concepts:** math/statistics intuition, generalization, overfitting, optimization, and evaluation transfer across tools—while specific LLM capabilities and APIs change fast.

*For trustworthiness, safety and reliability:*

*We need to make sure it works as **we** need it!*

*However fast automation will progress,*

- *even if it tells us how it performed a task, **we need to understand** whether this was a reasonable approach*
- *It might **not explain "truthfully"** how it arrived at the result (sometimes hard to check)*

# Simplified diagram of machine learning



The focus of our course:  
Machine learning concepts  
and methods

# Today

- Who is in this class and what is machine learning
- Learning objectives, prerequisites, other courses
- Logistics
  - Lectures, Tutorials
  - Projects, Moodle
- More examples of machine learning problems
  - Supervised learning
  - Unsupervised learning

# Scope and objectives of the lecture

By the end of the course you should

1. have **familiarity** with some common machine learning methods, i.e. be able to
  - **explain** basic machine learning concepts
  - **apply** basic machine learning methods on real problems
  - **validate/evaluate** the output of a learning method
2. have **reliable intuition** for why and when which algorithms work well, i.e. be able to
  - **choose** a small set of promising methods for specific problems and datasets
  - **justify** the intuition mathematically and **explain the why and when**



# ...ultimately for your future career

2. have intuition for why and when which algorithms work how “well” and be able to
  - choose a small set of promising methods for specific problems and datasets
  - justify the intuition mathematically and explain the why and when



## mathematical approach is important for reliable and transferable intuition:

- intuition can be misleading, may often guide you wrongly when conditions change!
- mathematical derivations may characterize precisely when intuitions are valid
- formal characterization helps to transfer intuition reliably to new problems



3. ultimately in the future, you want to be able to tackle new real–world problems
  - develop new methods for new challenges in a principled manner
  - evaluate new algorithmic developments/LLM designed methods in the future



# Who is this course designed for?

No prior knowledge about what machine learning is required, BUT you need

- a good grasp of basic concepts in linear algebra, analysis and probability
  - the Preliminaries chapter in the [lecture notes](#) should only serve as a reminder of the core math concepts, if you haven't heard of most concepts there, the course might be difficult to follow
  - other reference: e.g. Part I of "Mathematics for Machine Learning" by Deisenroth, Faisal, Ong at <https://mml-book.com/>
- basic programming experience (we will work in Python)
  - dedicated tutorial Monday next week 23.2. during Q&A session!
  - links to more online tutorials on the course website

# Intro to ML: foundation for other courses at ETH

ML Methods for specific applied domains (courses at D-INFK)

- Privacy Enhancing Technologies (Fall)
- Natural Language Processing (Fall)
- Computer Vision (Fall)
- Large Language Models (Spring)
- Machine Perception (Spring)
- Machine Learning for Health Care (Spring)
- Machine Learning for Genomics (Spring) ...

# Intro to ML: foundation for other courses at ETH

## More on ML Methods & Theory at D-INFK

- Advanced Machine Learning (Fall): Continuation and advanced topics
- Deep Learning (Fall): Deep neural networks and their applications
- Probabilistic Artificial Intelligence (Fall): Learning-based decision making under uncertainty
- Reliable and trustworthy artificial intelligence (Fall): (certifiable) Robustness of ML methods
- Guarantees for Machine Learning (Fall): Statistical guarantees for ML methods (sample complexity)
- Computational Intelligence Lab (Spring): Matrix Factorization, Recommender Systems, projects
- Foundations of Reinforcement Learning (Spring): Mathematics of Reinforcement Learning


# Today

- Who is in this class and what is machine learning
- Learning objectives, prerequisites, other courses
- **Logistics**
  - Lectures, Tutorials
  - Projects, Moodle
- **More examples of machine learning problems**
  - Supervised learning
  - Unsupervised learning

# People

- Instructors

- Andreas Krause
- Fanny Yang



- Teaching assistants

- Head TA: Marco Bagatella
- Other Heads: Malte Franke (projects), Frederike Lübeck (exam), Cynthia Chen (tutorials)
- Other TAs: Nicolas Menet, Riccardo De Santi, Logan Numerow, Weronika Ormaniec, Tiago Pimentel, Junhyung Park, Koren Nodens, Yarden As, Scott Sussex, Yunke Ao, Patrik Okanovic, Nuria Armengol Urpi, Shubh Goel, Tamas Nemes, Piersilvio de Bartolomeis, Lars Lorch, Rajesh Sharma, Barna Pasztor, Jonas Hübotter, Bruce Lee, Javier Abad, Sarah Martinson, Federico di Gennaro, Zifan Lyu, Juan Montes, Mengyao Fan, Anuj Srivastava, Balázs Szekér, Jinlei Zhou, Philip Pawlowsky, Gil Kur, Mert Albaba, Leo Schmidt-Traub, Tobias Wegel, Alex Shevchenko, Klejdi Sevdari, Anja Sjöström, Martin Wertich, Joudaki Amir

# Most important resource

Course website with all course info: <https://las.inf.ethz.ch/teaching/introml-s26>

- includes slides, lecture notes, homeworks & solutions, project details, recording links, book recommendations, access instructions to all of those and FAQ
- for slides/tutorial/homework material, use username: ml; password: predict
- additional access instructions can be retrieved when you're in the ETH VPN

# Lecture logistics

- Lectures are recorded (no livestream)
  - unannotated slides uploaded before lecture
  - annotated ones uploaded after the lecture
- Ask questions during lectures
  - in-person: simply raise your hand, or ask during the break, or after lecture
  - or write in [eduApp course channel](#) (TA will monitor and read aloud or answer most questions)
- Other forms of in-person participation
  - will use clicker questions on eduApp (download if you haven't yet)
  - will sometimes encourage peer discussions

# Attendance & Waitlist

**800 registered, 550+ on waitlist** but in-person space is limited. **This week** for all that are registered:

- Even student IDs can attend in **ETA F5**, odd student IDs can go to **ETA E1**
- In the second week all of you should be able to find a spot in ETA F5
- Depending on enrollment and exam constraints, we'll **admit more waitlisted students** during the first weeks of the course. Usually, everybody who wanted to take the course for credit, could
- **Please de-register if you don't plan to take the course for credit to make space for your peers!**

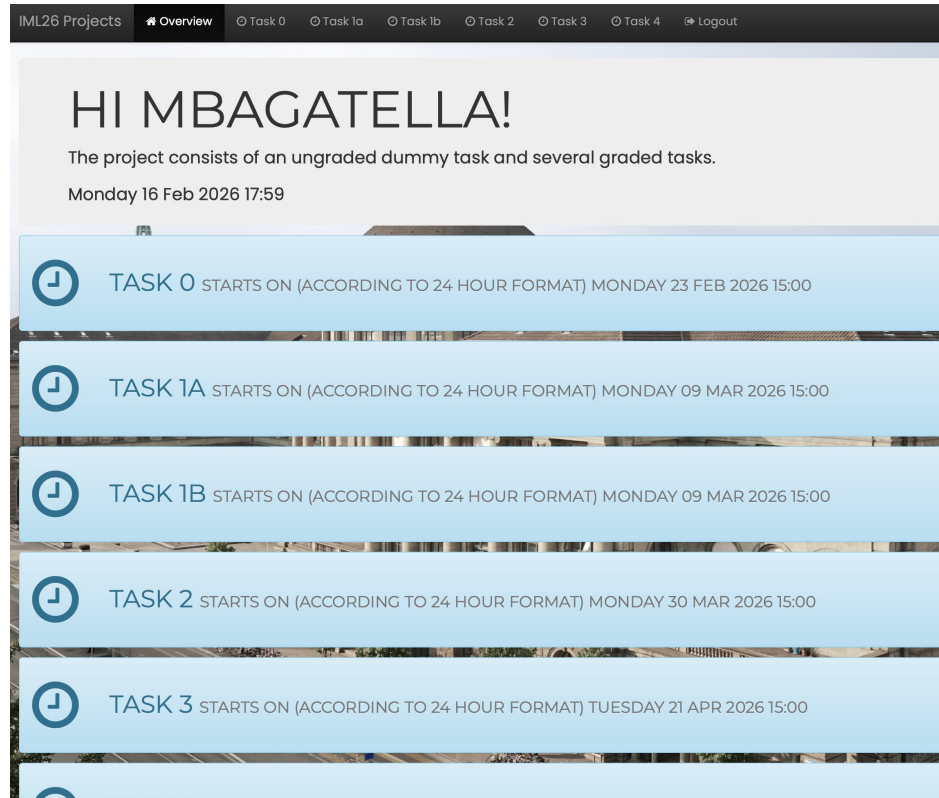
# Tutorials & Homeworks

- Take homework seriously if you want to pass the exam.  
Released [bi-weekly on Tuesdays](#), starting in two weeks, discussed in the Friday tutorial of the same week
- [Tutorials: Fridays 2pm-4pm](#), in ETA F5+ETF E1, recorded
  - Starting 20.2., this week: math recap (instructions on how to use the math recap notes and walking through some basic linear algebra, calculus and probability)
  - Other weeks: Recap of lecture material OR homework solutions chosen via student feedback
- Solutions are also on the website after the tutorials

# Projects

- In the 4 course projects, you will apply basic learning methods to [make predictions on real data](#)
- To do now:
  - Team up in groups of **(up to)** three students (there is a moodle topic for this)
  - More instructions are posted on the website
- “Graded” only as pass/fail, best project teams will be awarded a “best project” certificate
- 3 of 4 projects [must be passed to register for the exams](#)
- Solutions will be presented in the [Q&A sessions](#) after the deadlines (recorded)

# Project server <https://project.las.ethz.ch>



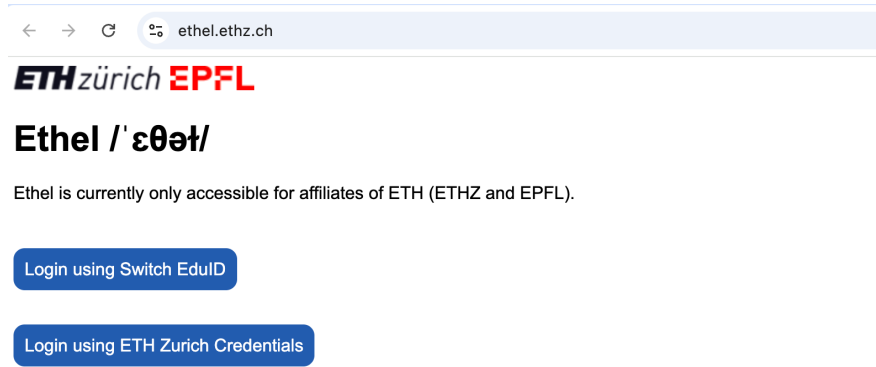
The screenshot displays the 'IML26 Projects' web application. The navigation bar at the top includes 'Overview', 'Task 0', 'Task 1a', 'Task 1b', 'Task 2', 'Task 3', 'Task 4', and 'Logout'. The main content area shows a personalized greeting 'HI MBAGATELLA!' and a description: 'The project consists of an ungraded dummy task and several graded tasks.' The current date and time are 'Monday 16 Feb 2026 17:59'. Below this, a list of tasks is shown, each with a clock icon and a start time:

- TASK 0 STARTS ON (ACCORDING TO 24 HOUR FORMAT) MONDAY 23 FEB 2026 15:00
- TASK 1A STARTS ON (ACCORDING TO 24 HOUR FORMAT) MONDAY 09 MAR 2026 15:00
- TASK 1B STARTS ON (ACCORDING TO 24 HOUR FORMAT) MONDAY 09 MAR 2026 15:00
- TASK 2 STARTS ON (ACCORDING TO 24 HOUR FORMAT) MONDAY 30 MAR 2026 15:00
- TASK 3 STARTS ON (ACCORDING TO 24 HOUR FORMAT) TUESDAY 21 APR 2026 15:00

# I still have a question ...

- In-person: Ask during and after lectures, tutorials (recorded)
- 1-on1: Dedicated Q&A session (office hours, optional) - recorded
  - will be held on **Mondays, 5.15-6 pm** virtually on Zoom: <https://ethz.zoom.us/j/62324294937>  
(TA from the project team will also be there and present project-related information)
- Forum: Moodle will be monitored by TAs
  - you are all encouraged to engage and answer!  
in fact trying to explain is the best way to check if you really understood
- Email: [introml26@inf.ethz.ch](mailto:introml26@inf.ethz.ch) (for logistics), [introml26-projects@inf.ethz.ch](mailto:introml26-projects@inf.ethz.ch) (for projects)

# ML for Intro to ML - Ethel chatbot



Log in with ETH  
credentials!

- A chatbot specialized with material from IML (still in the process of getting updated with latest dates and last year's material...)
- *Disclaimer: not guaranteed to give correct responses - use at your own risk!  
In particular, Ethel answer cannot be quoted to request points in the exam*
- New this year: upon informed consent, anonymous data collection to improve teaching - you will be able to give your consent some time in the upcoming weeks

what does subscript of expectation mean?

The subscript in an expectation, such as

$$\mathbb{E}_X[f(X, Y)],$$

indicates that the expectation is taken with respect to the random variable  $X$  only—that is, you average over the marginal distribution of  $X$  while treating  $Y$  as fixed (or as a parameter).

In more detail:

- When we write  $\mathbb{E}[f(X, Y)]$  (without a subscript), we mean the average of  $f(X, Y)$  taken over the joint distribution of  $X$  and  $Y$ .
- When we write  $\mathbb{E}_X[f(X, Y)]$ , the subscript  $X$  tells us to perform the integration (or summation) over  $X$  only, according to its marginal probability density (or mass) function  $p_X(x)$ , treating  $Y$  as a constant. This results in a function of  $Y$ .

For example, if  $X$  is continuous, then

$$\mathbb{E}_X[f(X, Y)] = \int f(x, Y) p_X(x) dx.$$

Similarly, a subscript in a variance expression, like

$$\text{Var}_X(f(X, Y)),$$

denotes that the variance is computed with respect to the distribution of  $X$  only.

This notation is particularly useful when dealing with conditional expectations. For instance, when you see

$$\mathbb{E}[f(X, Y)|Y = y],$$

this means the expectation is taken over the conditional distribution of  $X$  given  $Y = y$ .

In summary, the subscript in the expectation tells you which variable's distribution is being used to average the function, leaving other variables either fixed or as parameters.



# All links at one glance

- Course website with all course info: <https://las.inf.ethz.ch/teaching/introml-s26>
- Q&A session: <https://ethz.zoom.us/j/62324294937> (password on website)
- Moodle: <https://moodle-app2.let.ethz.ch/course/view.php?id=27418>
- eduApp: <https://eduapp.ethz.ch>
- Project: <https://project.las.ethz.ch/>

## Other FAQ

- Distance exams?  
Possible (as exception), but need to officially request with study administration.
- Project counts for this summer + winter's session exams.  
If you write the exam later, you have to redo the projects.
- Lecture notes for all lectures? - Mostly complete, one chapter on LLMs missing.

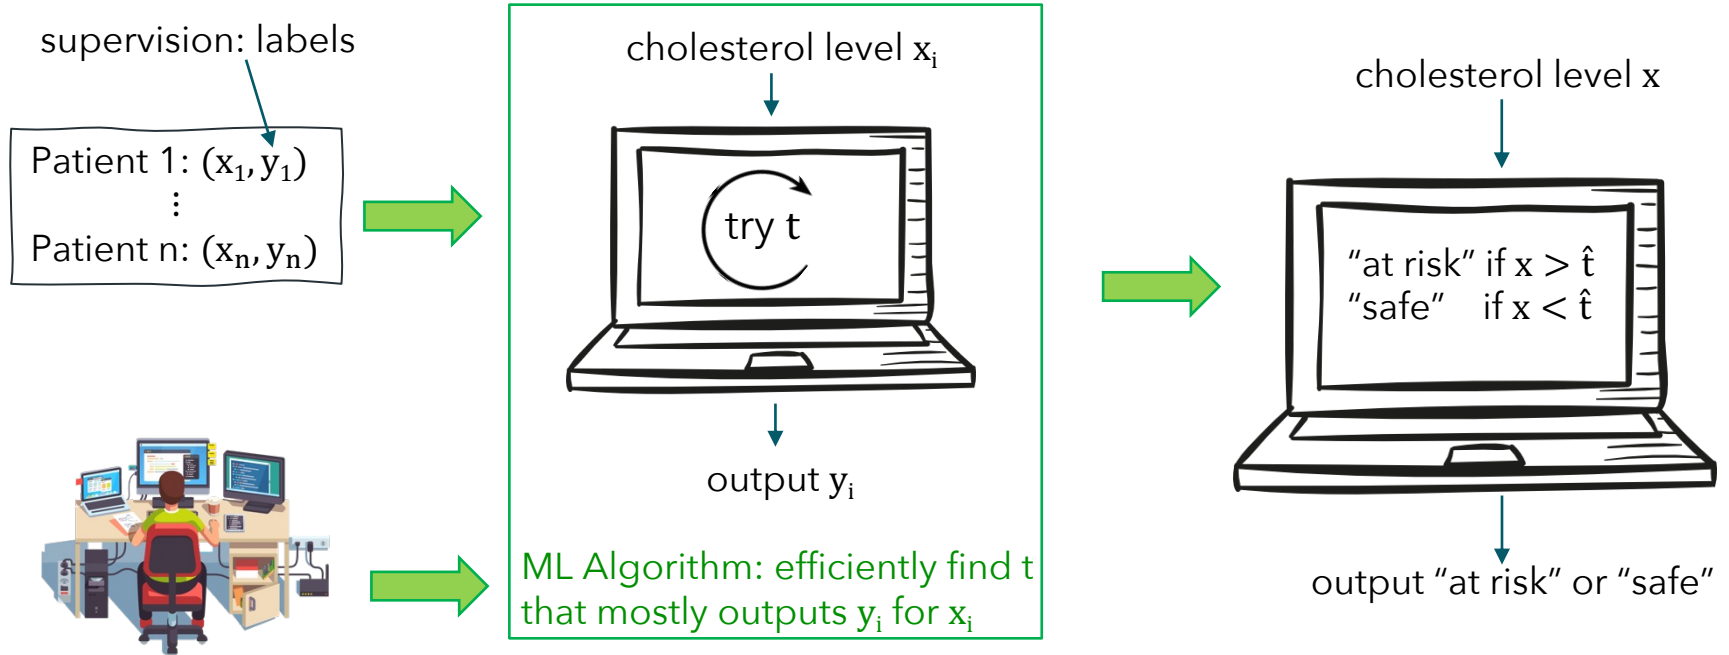
Any other questions?

# Today

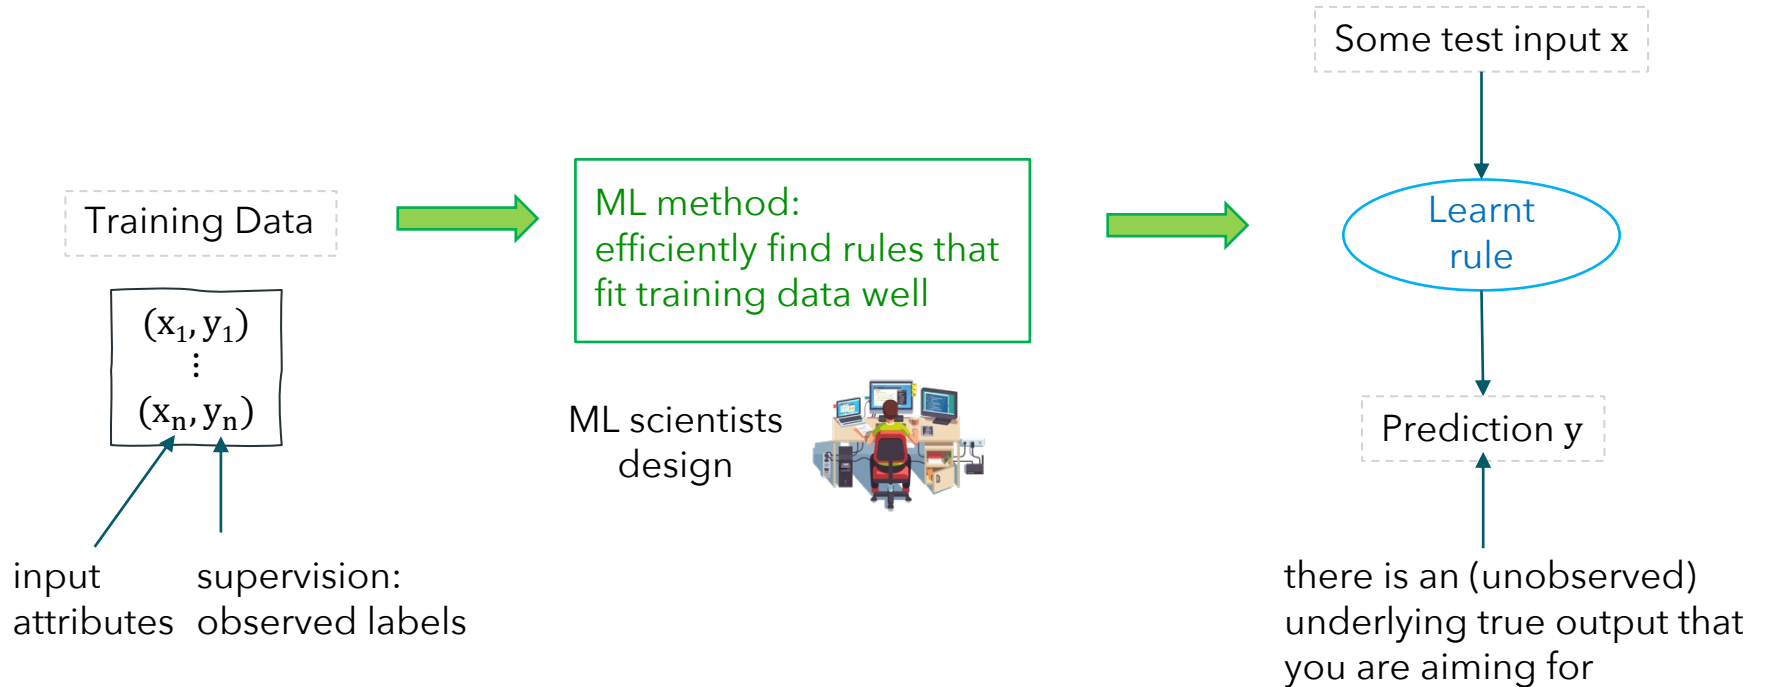
- Who is in this class and what is machine learning
- Learning objectives, prerequisites, other courses
- Logistics
  - Lectures, Tutorials
  - Projects, Moodle
- More examples of machine learning problems
  - Supervised learning
  - Unsupervised learning

# Supervised Learning pipeline

Machine Learning (ML) approach: Machines learn rules using example data



# Simplified diagram of supervised learning



# Different supervised learning tasks

- **Classification:** Predict the class (discrete scalar  $y$ ) of an input  $x$
- **Regression:** Predict a value (continuous scalar  $y$ ) for an input  $x$
- **Structured Prediction:** Predict an output  $y$  beyond scalars  $x$

We will see some examples now...

# Classification of physical measurements



x: blood measurements

Complete blood count		Serum chemistry			
WBC	5400/ $\mu$ L	TP	6.7 g/dL	Na	137 mEq/L
RBC	$443 \times 10^4$ / $\mu$ L	Alb	4.0 g/dL	K	4.4 mEq/L
Hb	12.6 g/dL	T-Bil	0.6 mg/dL	Cl	102 mEq/L
Ht	35.50%	D-Bil	0.1 mg/dL	Ca	9.1 mg/dL
Plt	$11.7 \times 10^4$ / $\mu$ L	BUN	18 mg/dL	CRP	0.55 mg/dL
		Cr	0.90 mg/dL	TC	112 mg/dL
Blood coagulation test		LDH	209 IU/L	HDL-C	16 mg/dL
PT (INR)	1.1	CK	70 IU/L	LDL-C	15 mg/dL
PT	82.80%	AST	21 IU/L	TG	963 mg/dL
APTT	39.6 s	ALT	19 IU/L	SIL-2R	281 U/mL

y: whether patient at risk for heart disease

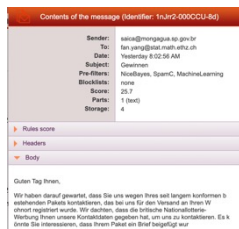

x: acoustic waveform



y: whether listening to music, speech, noise, ...

# Classification of Text

x: E-Mail Messages



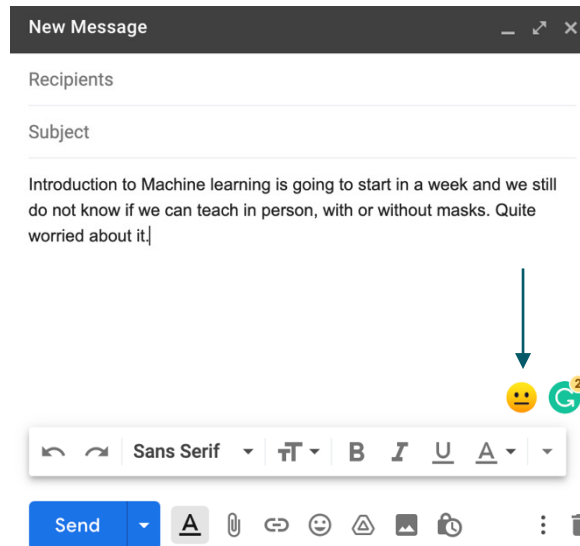
In quarantine: 2 messages

Action	Date	Sender	Subject
<b>Spam</b> 📧 📧 📧	Feb 15, 2022 8:02:57 AM	<a href="mailto:saica@mongagua.sp.gov.br">saica@mongagua.sp.gov.br</a>	Gewinnen

y: label: "spam" or "non-spam"

Example from ETH email spam detector

x: Email draft



y: sentiment

Demo using grammarly used in gmail

# Classification of Images




mite	container ship	motor scooter	leopard
mite	container ship	motor scooter	leopard
black widow	lifeboat	go-kart	jaguar
cockroach	amphibian	moped	cheetah
tick	fireboat	bumper car	snow leopard
starfish	drilling platform	golfcart	Egyptian cat

y: main object in the image



y: pneumonia or healthy

# Classification of Images




ARTIFICIAL INTELLIGENCE, AUTOMATION, AND THE ECONOMY

## Introduction

Recent progress in Artificial Intelligence (AI) has brought renewed attention to questions about automation driven by these advances and their impact on the economy. The current wave of progress and enthusiasm for AI began around 2010, driven by three mutually reinforcing factors: the availability of *big data* from sources including e-commerce, businesses, social media, science, and government,<sup>3</sup> which provided raw material for dramatically *improved machine learning approaches and algorithms*; which in turn relied on the capabilities of *more powerful computers*.<sup>4</sup> During this period, the pace of improvement surprised AI experts. For example, on a popular image recognition challenge that has a 5 percent human error rate according to one error measure,<sup>5</sup> the best AI result improved from a 26 percent error rate in 2011 to 3.5 percent in 2015. This progress may enable a range of workplace tasks that require image understanding to be automated, and will also enable new types of work and jobs. Progress on other AI challenges will drive similar economic changes.

# Regression examples I




x: apartment/house attributes

Size (m <sup>2</sup> )	# rooms	Dist. (km) to pub. transp.	Years since construct.
120	4.5	0.3	30
90	3.5	1.2	87
...	...	...	...

↓

y: selling price (CHF)


x: Molecule structure




↓

y: Functional properties of molecules (e.g., thermostability, binding affinity...)

# Regression examples II



x: User & article (or advertisement)



y: Probability of click / purchase

x: soil properties, weather, ...



y: crop yield


x: vital signs past trajectory




y: short-term forecast

Source: Adobe Stock: sudok1 & zapp2photo

# Structured prediction examples



x: text (in one language)




y: text (in another language)

Demo using google translate

x: amino acid sequence

```
MSEQNNTMTF      SIFS YITESTGTPSN
QIQRIYTKDISFE   ATYTYVIERWDPET
APNAPHVFQKD     SGILNPCYGWPVC
WQPEVKLD...     YVTVNHKHTVNG...
```

AlphaFold 3



● Experimental result  
● Computational prediction

y: protein structure

by GoogleDeepmind

# Structured prediction examples

FORBES > INNOVATION > AI

## AlphaFold Is The Most Important Achievement In AI —Ever

Rob Toews Contributor

*I write about the big picture of artificial intelligence.*

Follow



Oct 3, 2021, 07:34pm EDT

Updated Dec 10, 2021, 09:41am EST


x: amino acid sequence

MSEQNNTMTF  
QIQRIYTKDISFE  
APNAPHVFQKD  
WQPEVKLD...


SIFSYITESTGTPSN  
ATYTYVIERWDPET  
SGILNPCYGWPVC  
YVTVNHKHTVNG...



AlphaFold 3



T1037 / 6vr4  
90.7 GDT  
(RNA polymerase domain)



T1049 / 6y4f  
93.3 GDT  
(adhesin tip)

● Experimental result  
● Computational prediction

y: protein structure

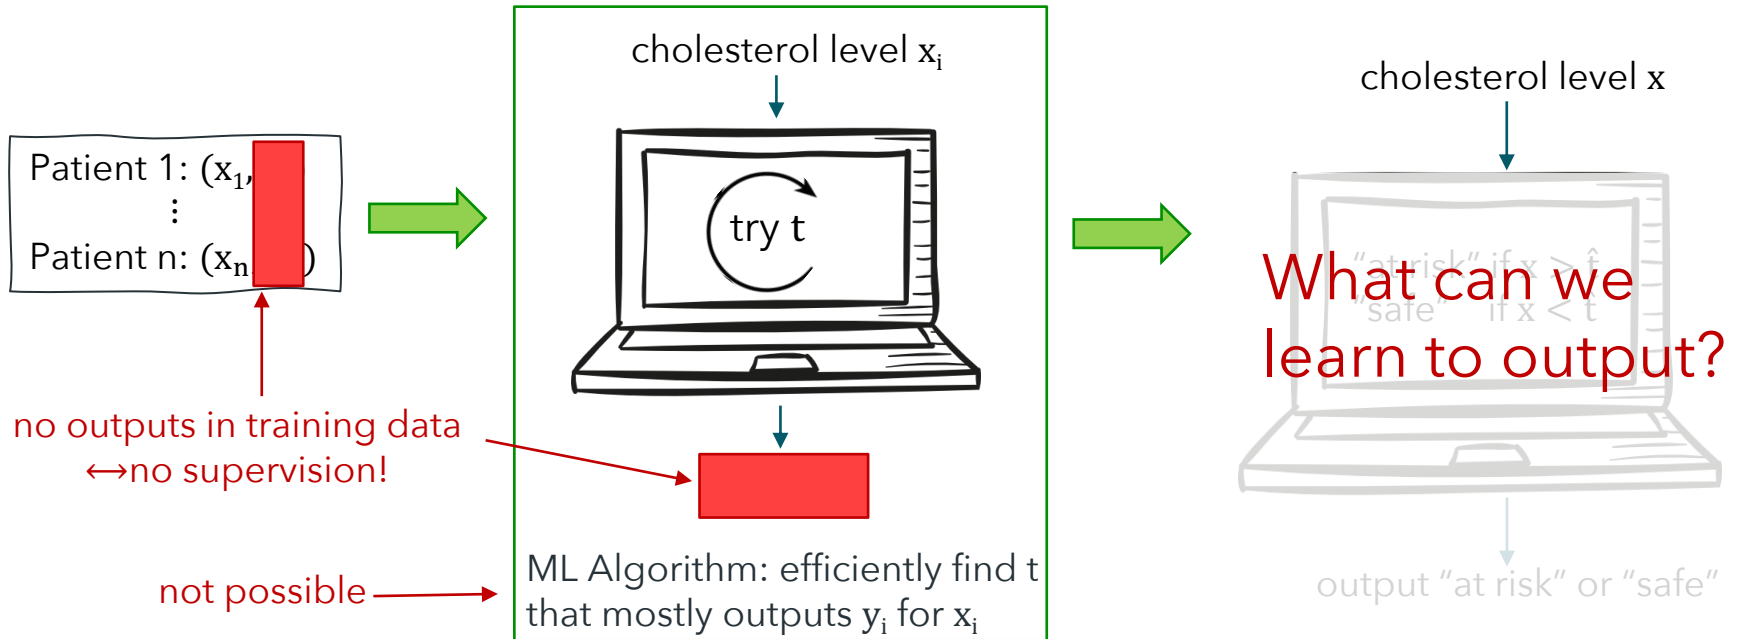
by GoogleDeepmind

# Today

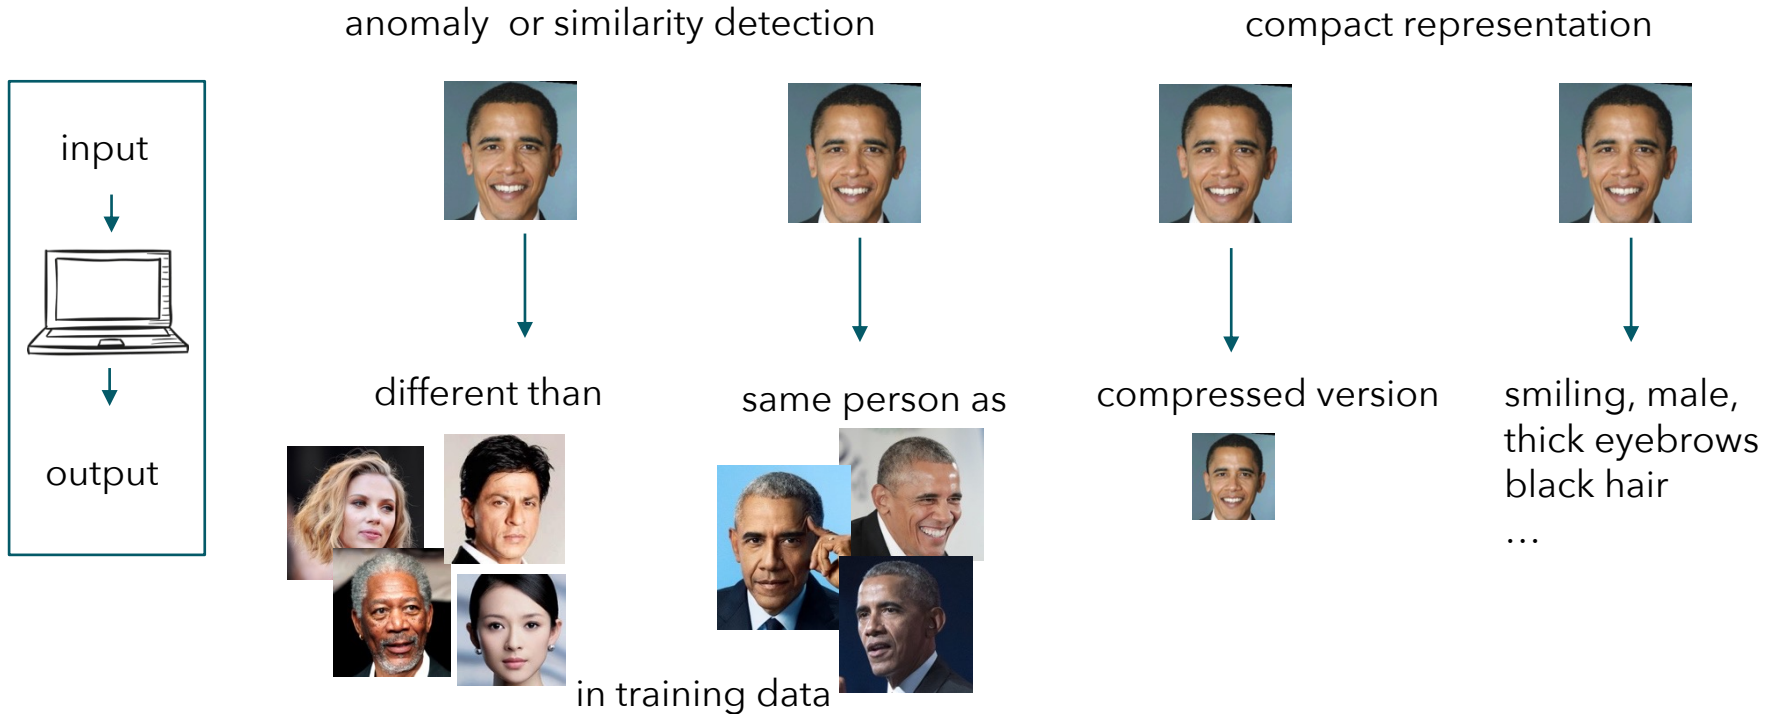
- Who is in this class and what is machine learning
- Learning objectives, prerequisites, other courses
- Logistics
  - Lectures, Tutorials
  - Projects, Moodle
- More examples of machine learning problems
  - Supervised learning
  - [Unsupervised learning](#)

# No labels during training - what can we learn?

Machine Learning (ML) approach: Learn rules using example data




# Unsupervised learning - information retrieval



# Unsupervised learning - completion / generation

You can also try to generate more of the same... no unique best/correct answer



## First sentence

Welcome to the introductory lecture to machine learning at ETH Zurich.



## Completion of paragraph

Welcome to the Introduction to Machine Learning lecture at ETH! In this course, you will embark on an exciting journey through the foundational concepts of machine learning, exploring how algorithms can learn from data and make intelligent decisions. We will cover key topics such as supervised and unsupervised learning, classification, regression, clustering, and neural networks. Throughout the course, you will gain hands-on experience applying machine learning techniques to real-world problems, all while developing a deep understanding of the mathematical principles that drive these technologies. By the end of the lecture series, you'll have the skills to implement basic machine learning models and an appreciation for their impact across various domains, from artificial intelligence to data science.

## Melody



## Harmonization




Demo using ChatGPT

<https://www.youtube.com/watch?v=QiBM7-5hA6o>

# Unsupervised learning - generation

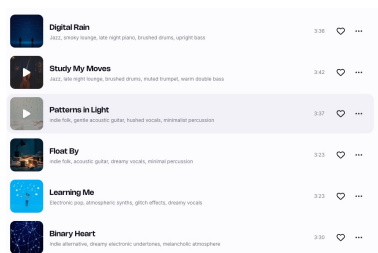
You can also try to generate more of the same... no unique best/correct answer




Songs of a certain music genre



new songs in that style




 Riffusion BETA and many more ...

Pictures of celebrities



new images of that kind



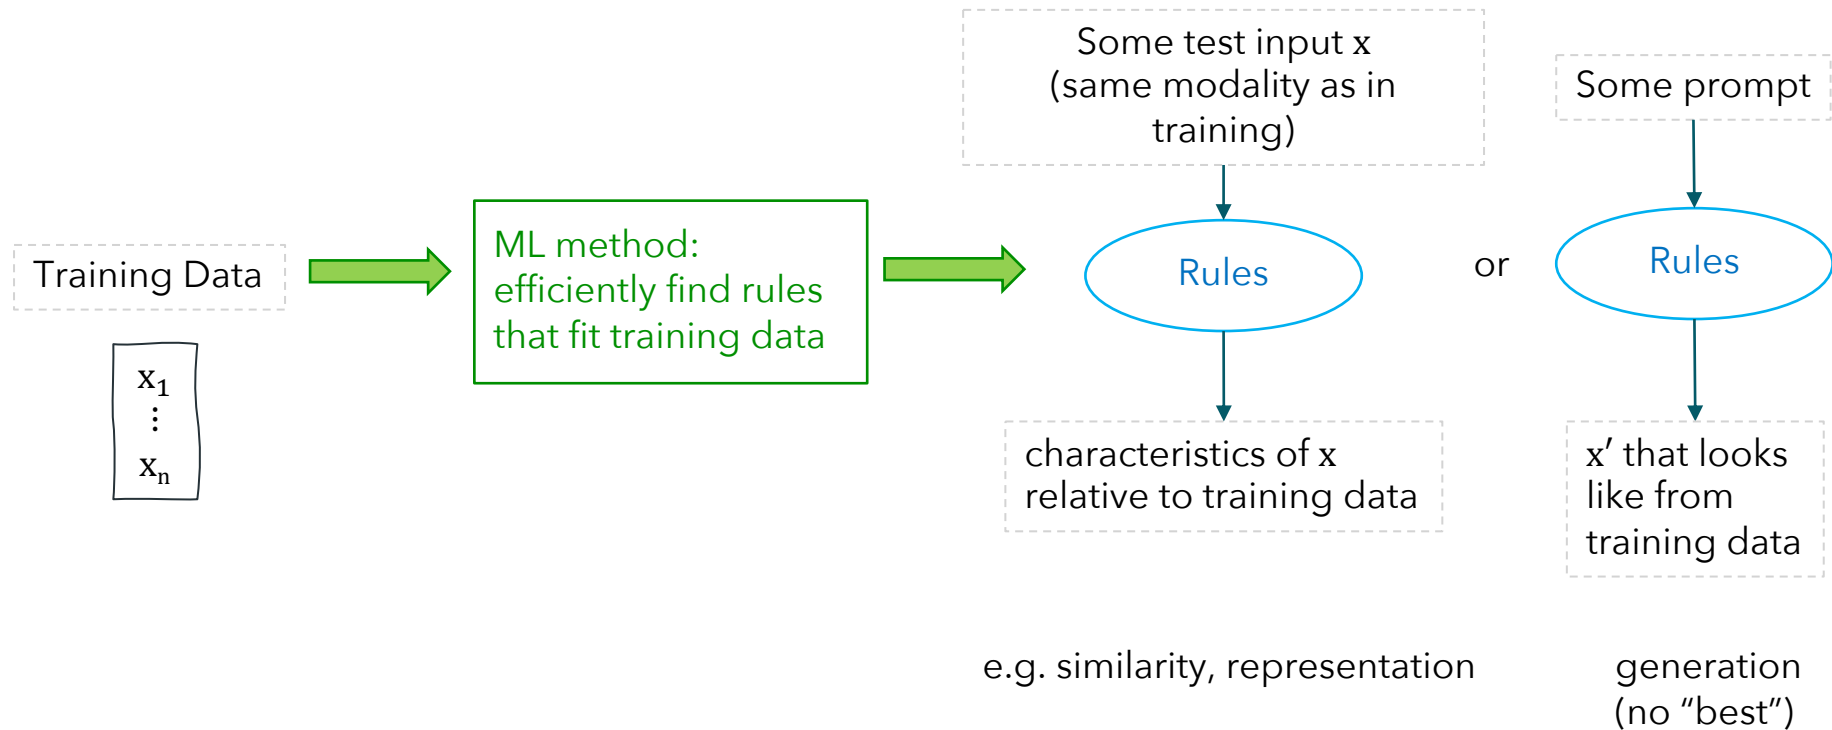
SN-GAN\* '18  
(no prompt)



mage.space '25  
(with specific prompt)

\*Generated using SN-GAN, Miyato et al. '18 "Spectral Normalization for Generative Adversarial Networks"

# Simplified diagram for unsupervised learning



# Unsupervised learning - goals & methods

- Summary of some common goals, using characteristics of the data
  - [Anomaly detection](#) of „unusual“ data points
  - Identification/discovery of (relevant) [unobserved variables \(such as features, classes\)](#)
  - [Compact representation / compression](#) of data sets
  - [Generation](#) of new data
- ML methods to learn characteristics of the data (covered in class):
  - [Clustering](#) e.g. for learning similarity and anomaly of points
  - [Dimensionality reduction](#) e.g. for learning compact representations
  - [Generative modeling](#) e.g. for feature learning and data generation

# A sneak peak into unsupervised methods

(they will be discussed in detail later)

# Clustering

Example: Which images belong together in the training dataset?


Can also be used to answer: Does a new image (yellow) look similar to others in the training dataset?



no: anomaly!

# Dimensionality reduction

Handwritten digits (MNIST)



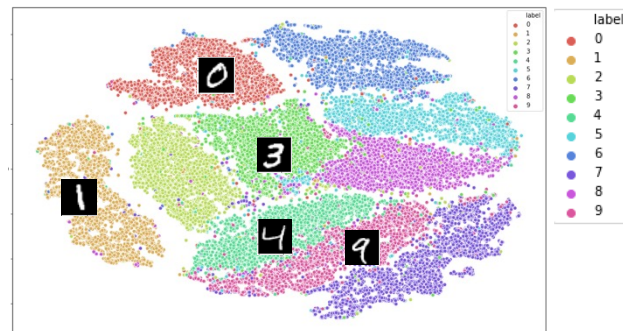
images of size 28 pixels x 28 pixels

→ 784 dimensional!

represented



in 2-d



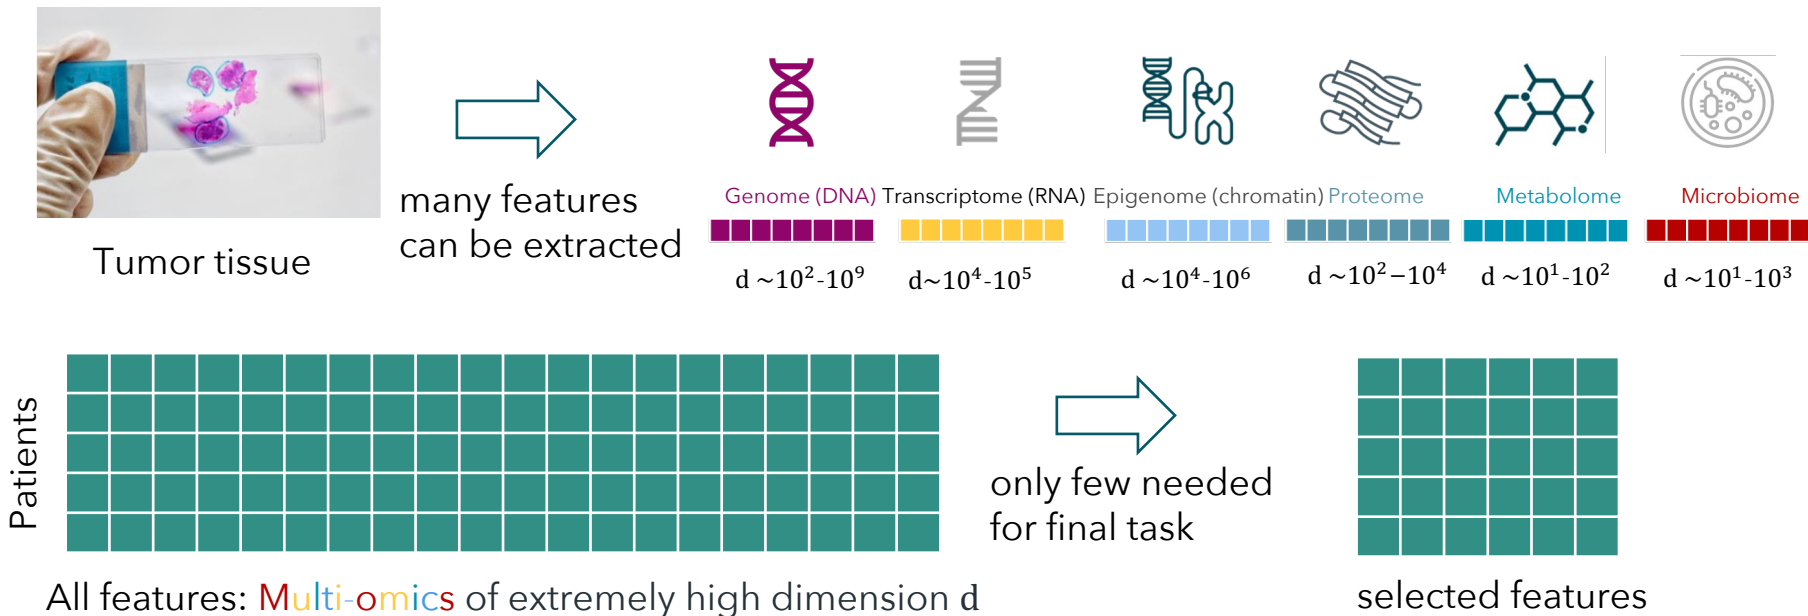
Clustered using projection using PCA + tSNE  
(without using the depicted number as a label!)

(we'll learn what PCA is later)

source: towardsdatascience.com

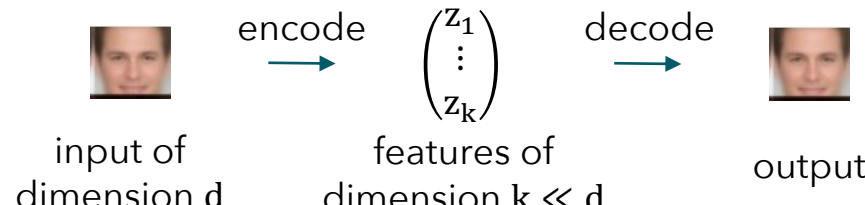
# Dimensionality reduction

Example: Multi-omics for cancer survival prediction have many features of high dimensionality  $d$




# Generative modeling for representation/compression

Example: Facial factors learned from faces of celebrities (CelebA) using variational autoencoders

- ML model can learn rules to   $\begin{matrix} \text{input of} \\ \text{dimension } d \end{matrix} \xrightarrow{\text{encode}} \begin{pmatrix} z_1 \\ \vdots \\ z_k \end{pmatrix} \xrightarrow{\text{decode}} \text{output}$    
 features of dimension  $k \ll d$

- Can visualize these features: (1) encode some image; (2) fix all  $z_1, \dots, z_k$  except  $z_i$ ; (3) vary  $z_i$  and decode

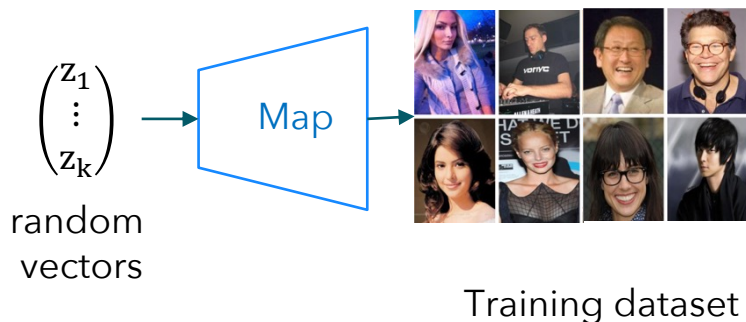


*Note: Latent presentations from VAEs are still used to train e.g. Stable Diffusion*

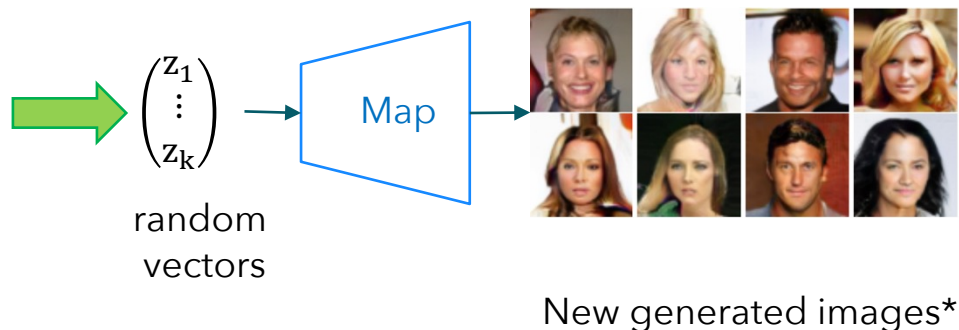
# Generative modeling for generation

Example: GANs that create images like training data

Learn maps (rules) such that different random vectors result in images similar to training set



Use rules to generate new images using more random vectors



*Note: Diffusion models (used in dall-e, midjourney etc.) spiritually similarly also generate from noise, while also allowing prompt usage*

# Generative modeling for generation


Example: LLMs that generate text from text using probabilistic model

**Learn** rules s.t. the observed next word (token) has a high probability (among all possible words)

**Use** rules to generate next words from user-given prompt



Students  
opened  
their



Distribution over possible words  
books (30%), laptops (15%)  
phones (20%) snacks (30%)...

↓ sample from distribution

snacks

Training examples: All parts of all sentences in books, wikipedia, social media ...  
(deployed models also train on human feedback etc.)

L.M. generates new text by repeating this process  
(appending generated words & rerunning pipeline)

More details on language models in the last lecture...

# Other modes of learning (not covered in class)

- Semi-supervised learning
  - Learning from labeled and unlabeled data
- Active learning
  - Acquire most informative labels for learning
- Transfer & meta learning
  - Learn on one domain and test on another
- Online / lifelong / continual learning
  - Learning from data as it arrives over time
- Reinforcement learning
  - Learning by interacting with an unknown environment

# Course syllabus (not in chronological order!)

- Supervised learning
  - Linear regression & classification, optimization
  - model validation, bias-variance trade-off, regularization
  - Kernels, neural networks & deep learning basics
- Unsupervised learning
  - Dimensionality reduction, representation learning
  - Clustering, Generative modeling with neural networks
- The statistical perspective
  - Probabilistic modeling (discriminative vs. generative), estimating the distribution
  - Decision theory, Bayes optimality
  - Formalizing intuitions with mathematical statements and derivations

# What you can do now

- Answer on high level what machine learning is
- Have an overview of problems machine learning can be used for
- Have a rough idea how training data can be used in supervised and unsupervised learning

Tomorrow:

- Learning a simple model from start to finish for a specific application