



Elemental Learning Theory

Alexander G. Ororbia II
Introduction to Machine Learning
CSCI-335
2/13/2026

ML in a Nutshell: Core Ideas

Representation / Modeling

Data format, organization

Model structure, “architecture”

Evaluation

“Goodness” of model on data sample

Guides optimization / model fitting

Optimization

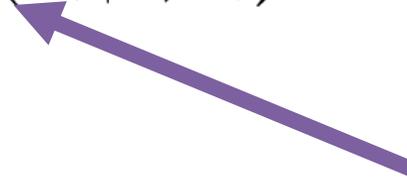
“Model fitting” → evolution/adjustment of parameters, error correction

Learning Algorithms

- Definition (well-posed learning problem):
 - A computer program is said to learn from *experience* E
 - with respect to some *class of tasks* T and *performance measure* P ,
 - if its performance at task T , as measured by P , improves with experience E

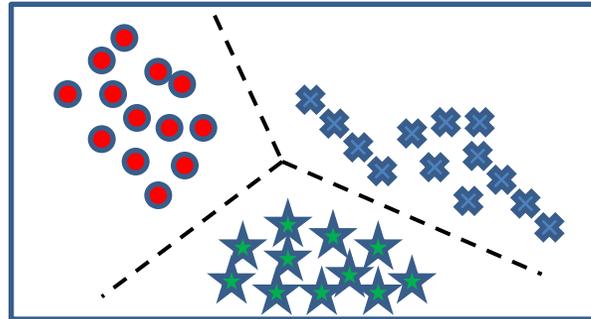
Types of Learning

- **Supervised (inductive) learning** $\{\mathbf{x}_n \in \mathbb{R}^D, \mathbf{y}_n \in \mathbb{R}^C\}_{n=1}^N$
 - Training data includes desired outputs, $(\mathbf{x}_i, \mathbf{y}_i = f(\mathbf{x}_i))$
 - Prediction / Classification (discrete labels), Regression (real values)
- **Unsupervised learning** $\{\mathbf{x}_n \in \mathbb{R}^D\}_{n=1}^N$
 - Training data does not include desired outputs
 - Clustering / probability distribution estimation
 - Finding association (in features)
 - Dimension reduction
- **Semi-supervised learning** $\{\mathbf{x}_n \in \mathbb{R}^D, \mathbf{y}_n \in \mathbb{R}^C\}_{n=1}^N \cup \{\mathbf{x}_m \in \mathbb{R}^D\}_{m=1}^M, M \gg N$
 - Training data includes a few desired outputs
- **Reinforcement learning** $\{\mathbf{x}_t \in \mathbb{R}^D, r(\mathbf{x}_{t+1}, a_t) \in \mathbb{R}\}_{t=1}^{T=\infty}$
 - Rewards from sequence of actions
 - Decision making (robot, chess machine)

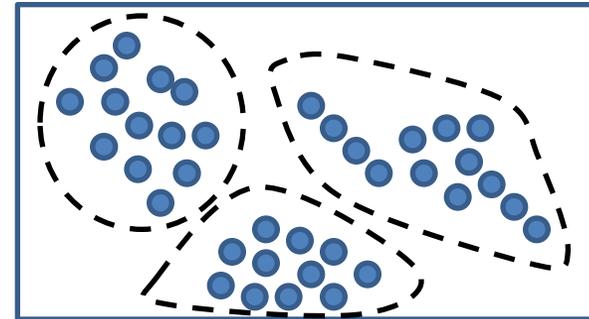


Reward function

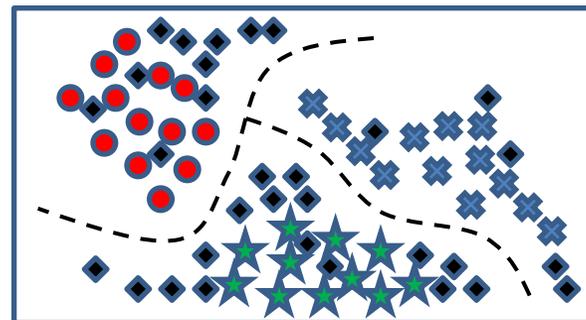
Visualizing the Types of Learning



Supervised
learning



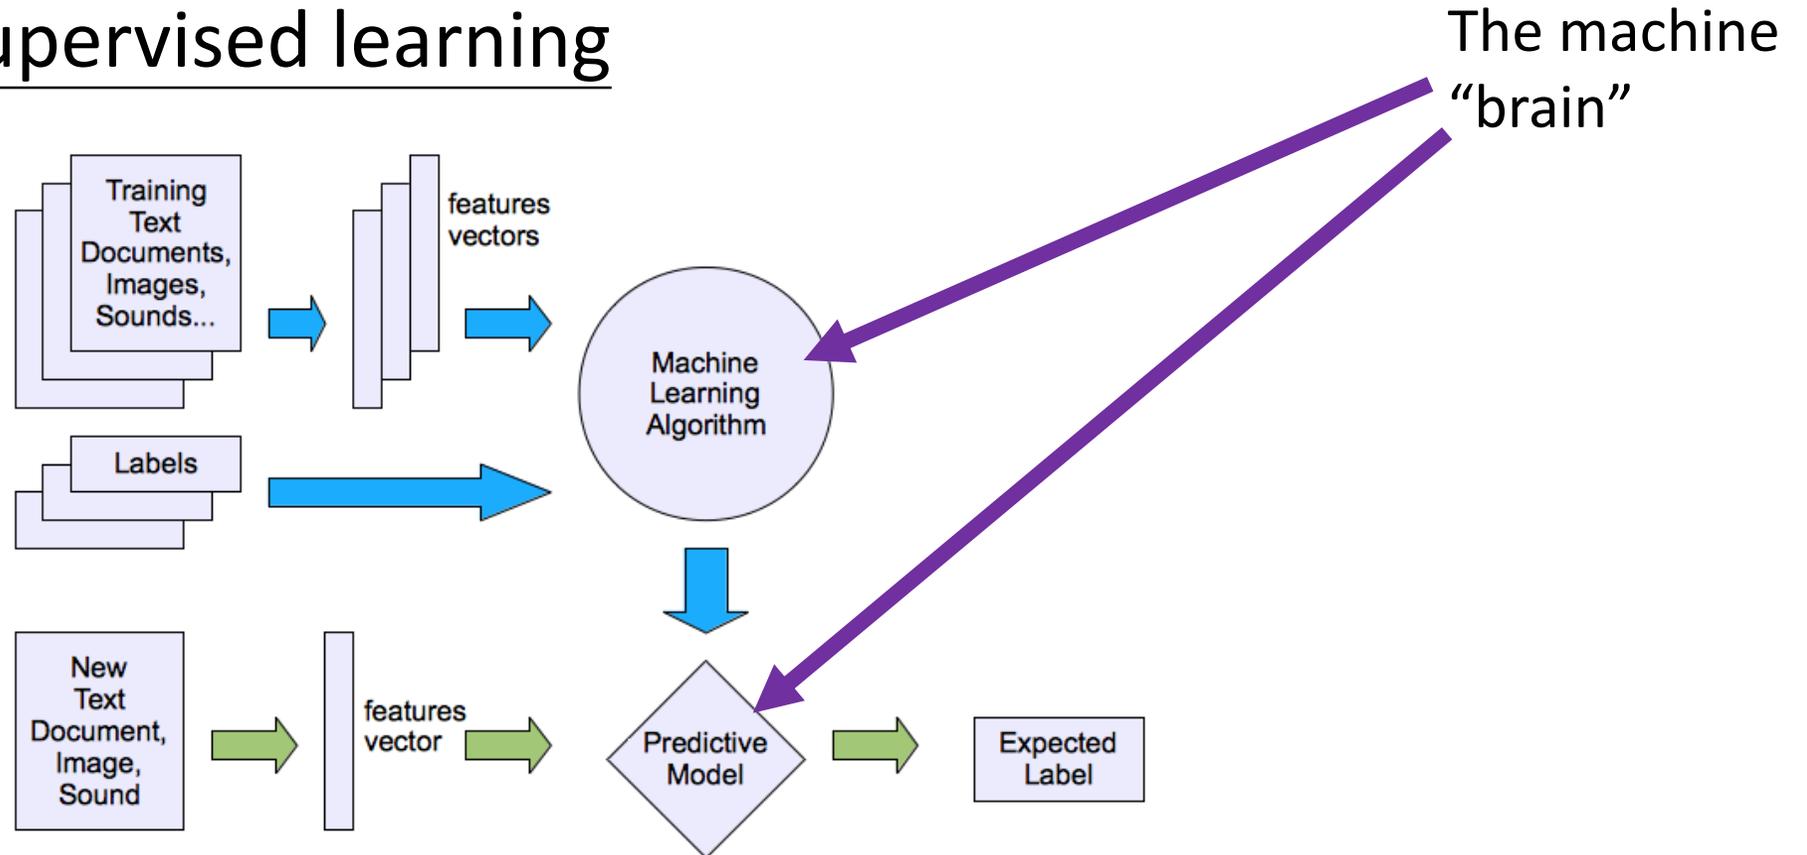
Unsupervised
learning



Semi-supervised
learning

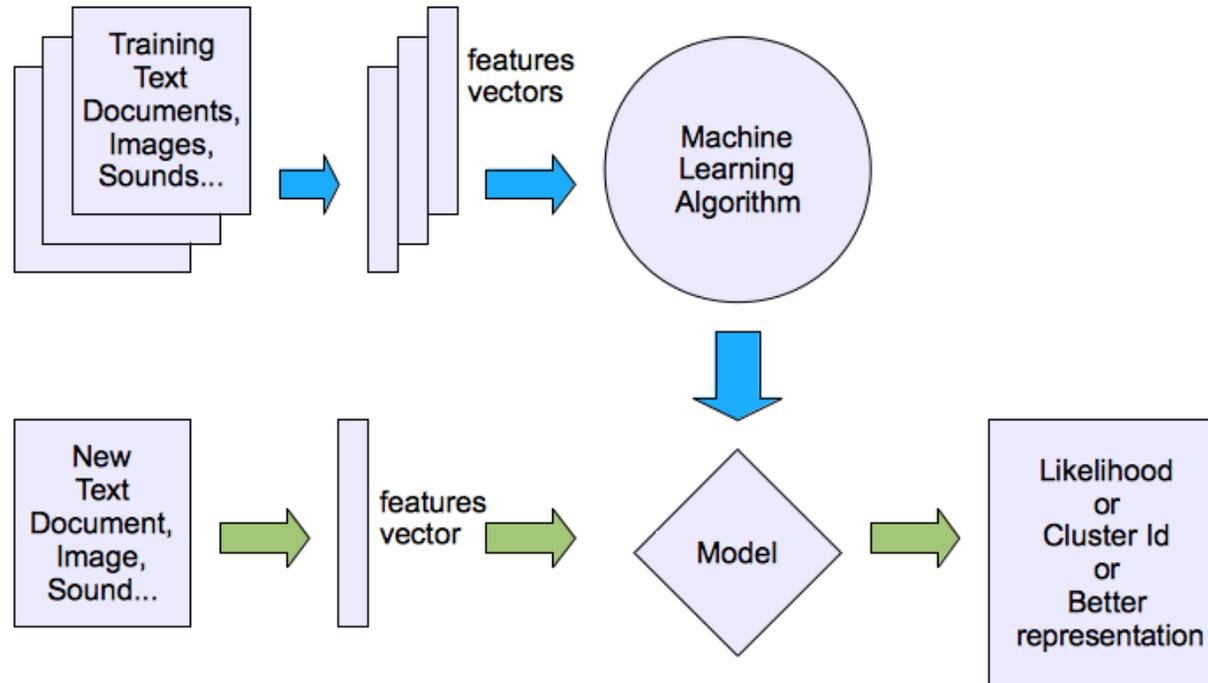
A Typical Supervised Learning Pipeline

- Supervised learning



A Typical Unsupervised Learning Pipeline

- Unsupervised learning



Statistical Learning in Practice

- Understanding domain, prior knowledge, and goals
- Data integration, selection, cleaning, pre-processing, etc.
- Learning model(s)
- Interpreting results
- Consolidating and deploying discovered knowledge
- ***Loop***

Statistical Learning in Practice

The science of the problem



- Understanding domain, prior knowledge, and goals
- Data integration, selection, cleaning, pre-processing, etc.
- Learning model(s)
- Interpreting results
- Consolidating and deploying discovered knowledge
- **Loop**

CSCI 635
(You are here!)



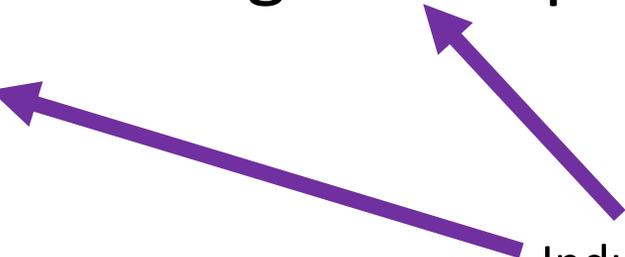
What big data/infrastructure courses help with



Defining the problem helps with this!



Industry dictates this
(machine learning engineering)



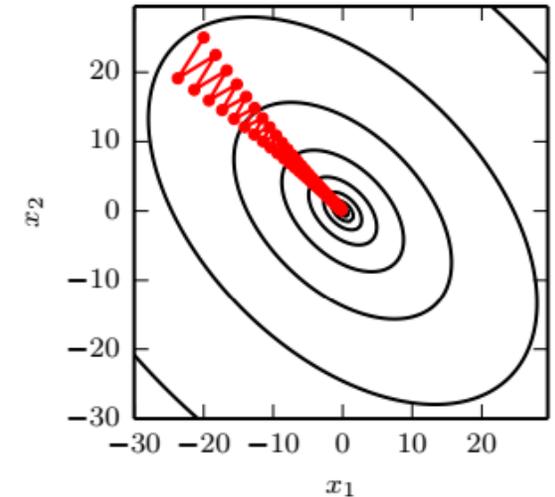
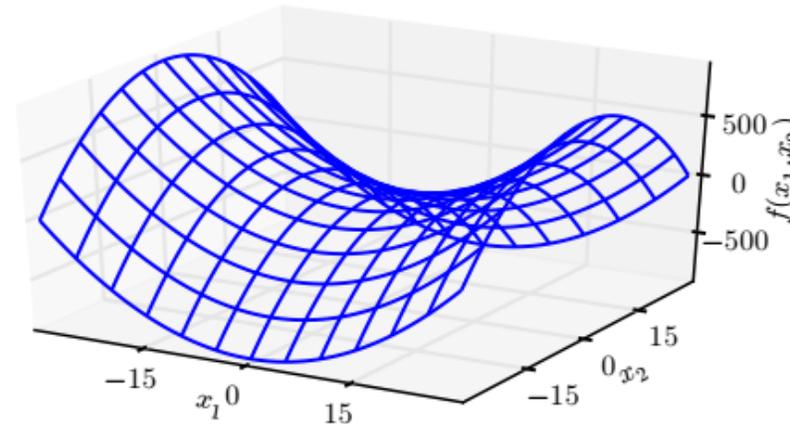
Training Experience

- Direct experience: Given sample input and output pairs for a useful target function
 - Checker boards labeled with the correct move, e.g. extracted from record of expert play
- Indirect experience: Given feedback which is *not* direct I/O pairs for a useful target function
 - Potentially arbitrary sequences of game moves and their final game results (i.e., a score or cumulative function output)
- Credit/Blame Assignment Problem: How to assign credit or blame to individual moves given only indirect feedback?
 - The **problem of credit assignment** = appears everywhere in statistical learning

So, Statistical Machine Learning is...

- **Three Pillars:**
 - Representation
 - Evaluation
 - Optimization

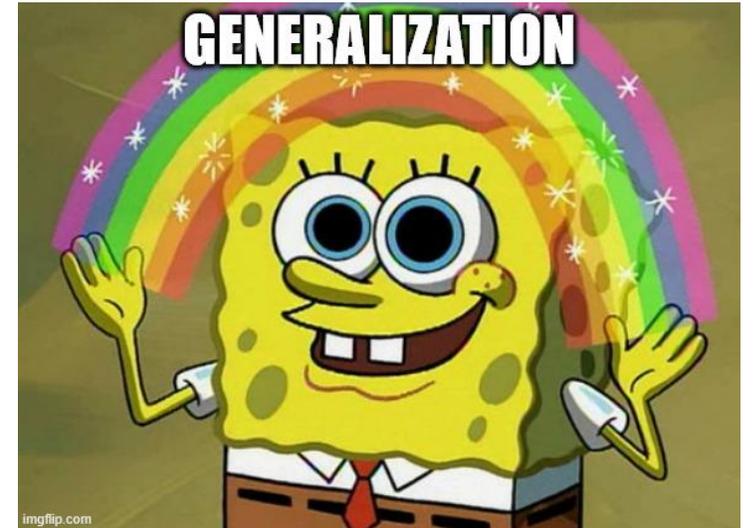
It is not only just fitting/optimization...



**Key
Point:**

...it is also about generalization!

So,...what is
generalization?



Training vs. Test Distribution

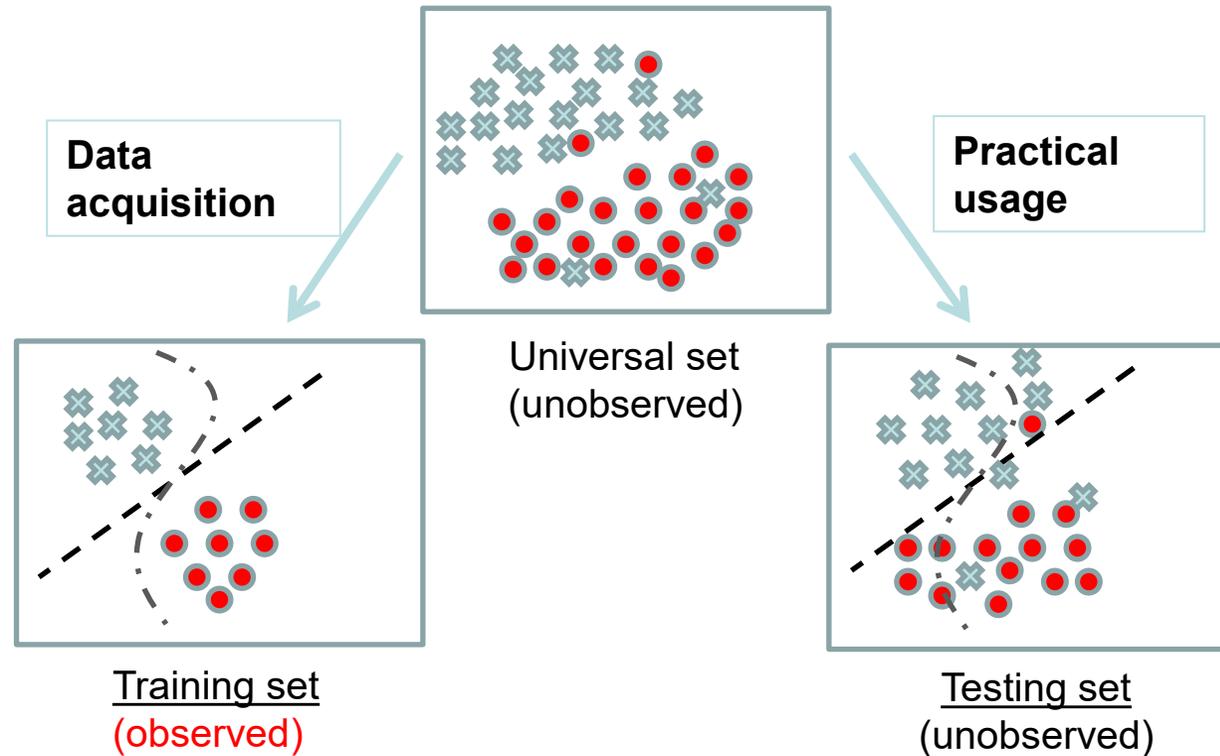
- *Generally*, assume that the training and test examples are independently drawn from the same overall distribution of data
 - **IID: Independently and identically distributed**
- If test distribution is different, requires *transfer learning*

At the core of the assumption:

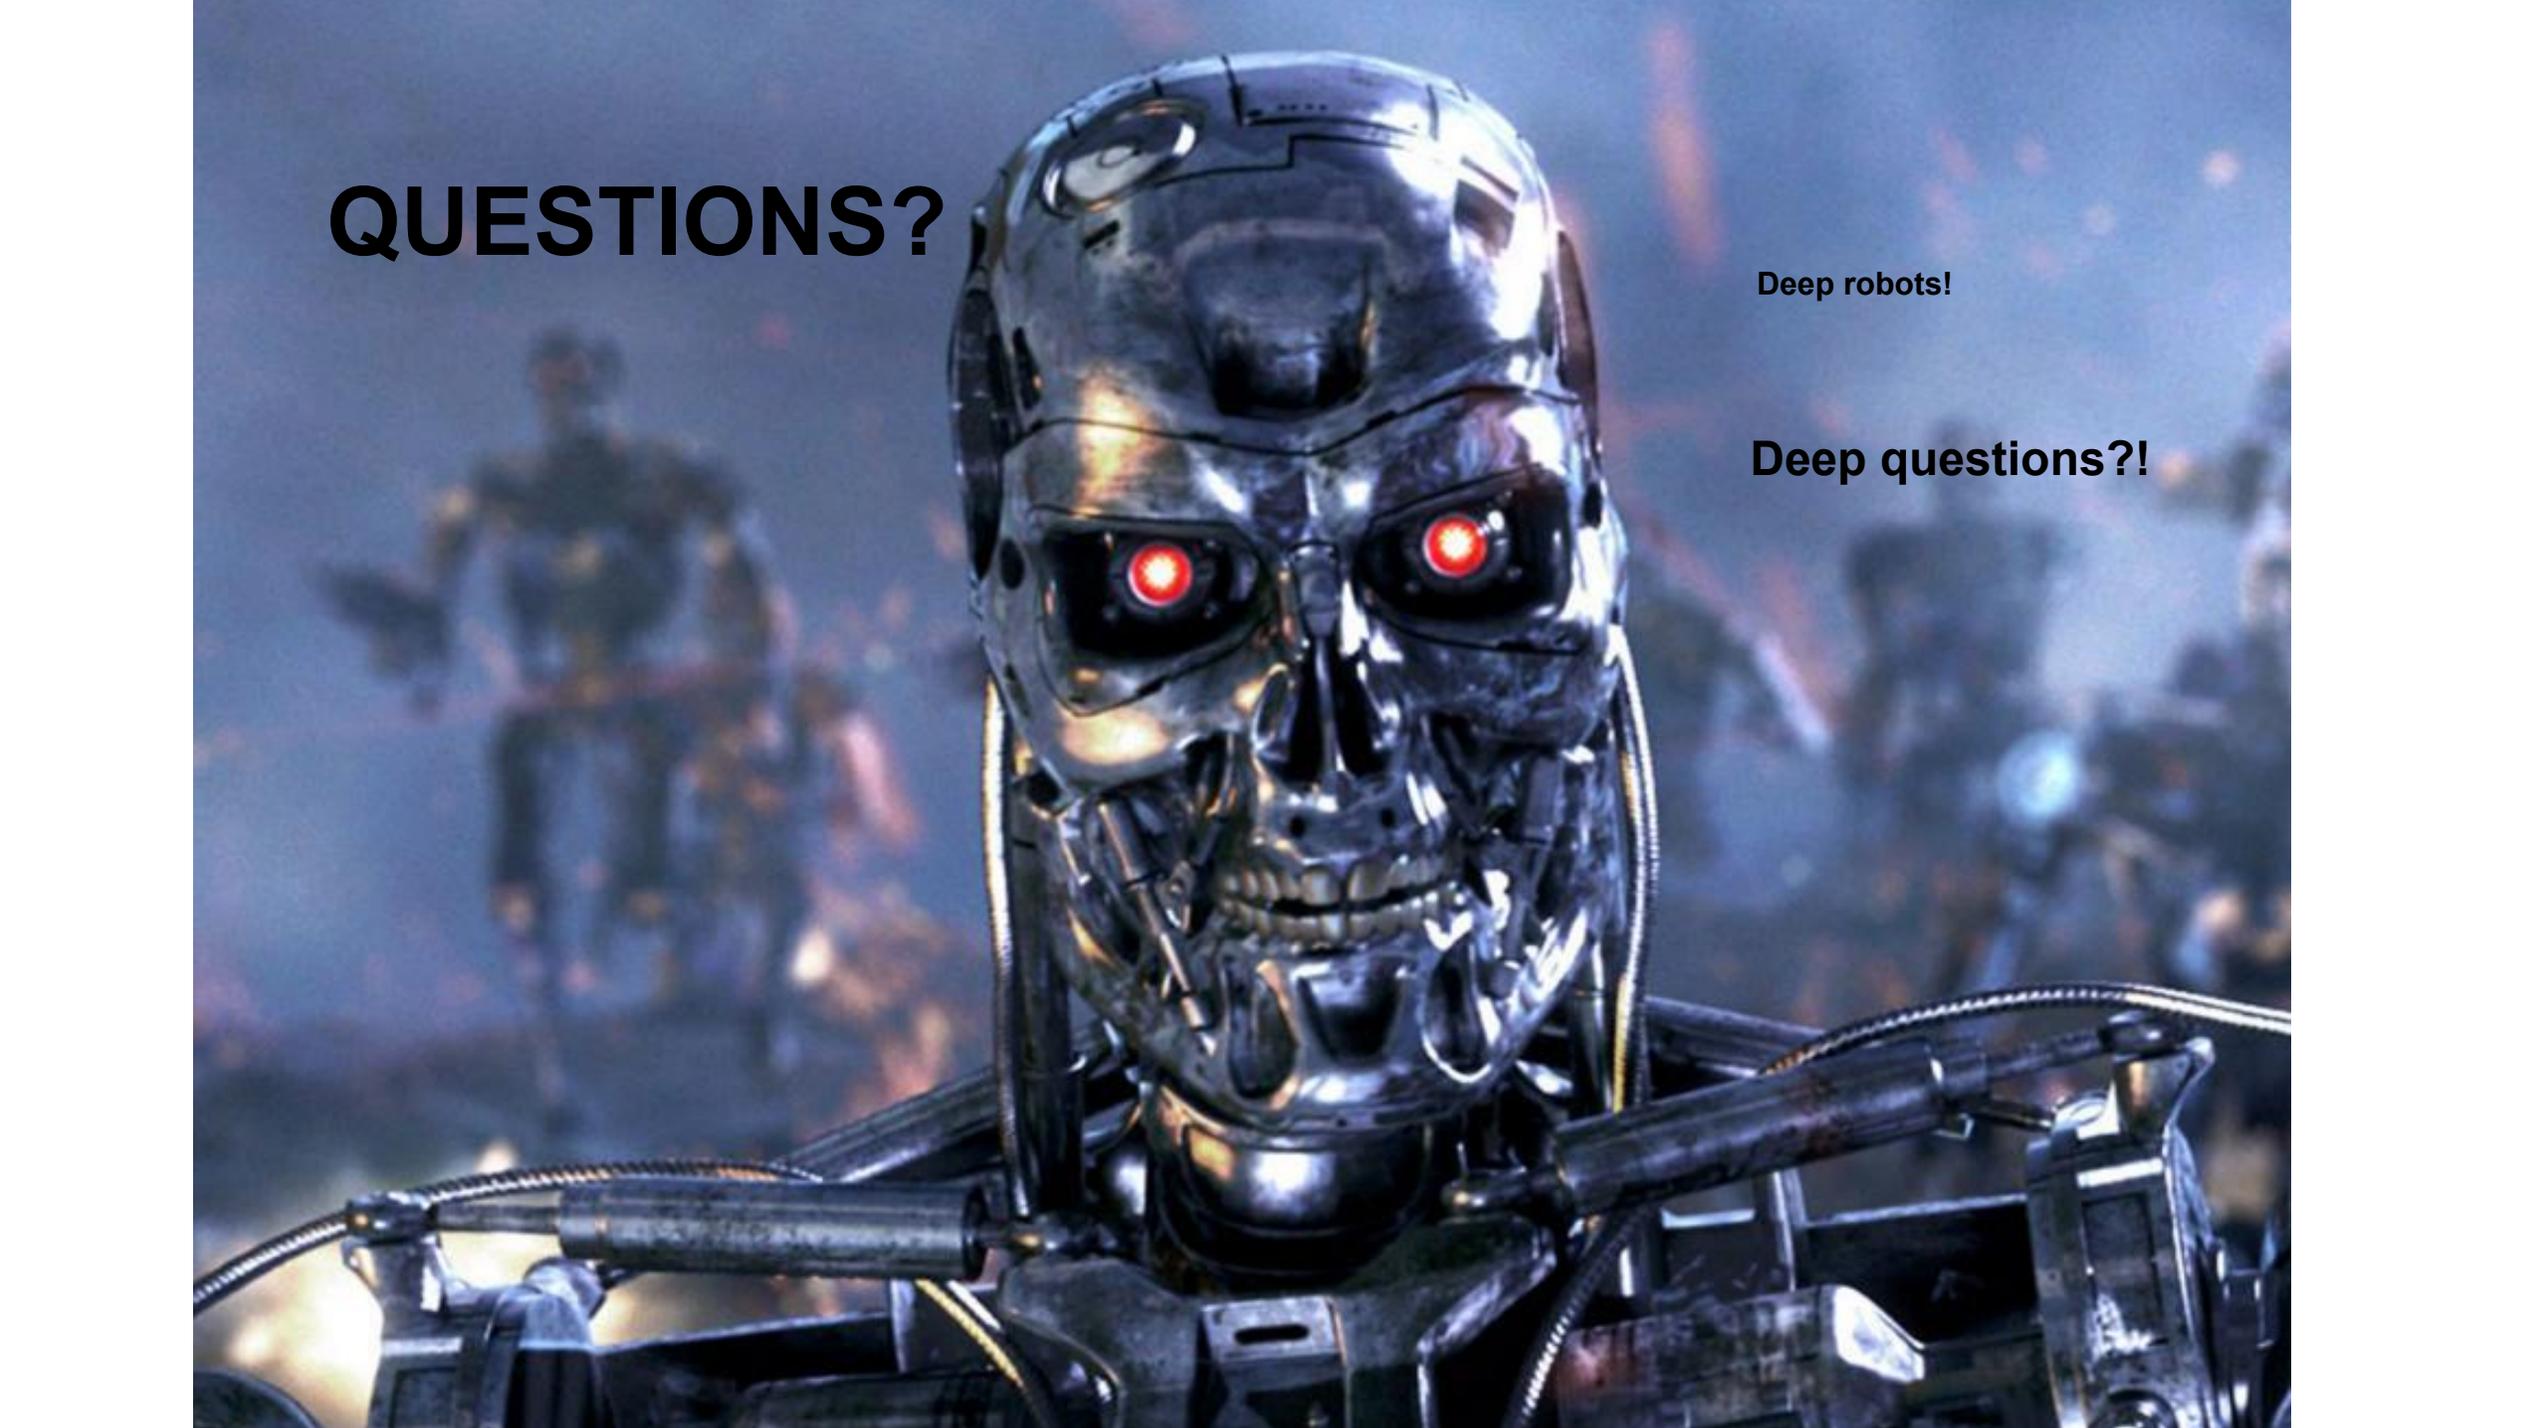
A collection of random variates *must* fall under same probability distribution *and* are mutually independent

Identical = no overall trends/fluctuations in (same) distribution of collected objects

Independent = collected objects are all independent events; value of one item gives no knowledge about values of others (& vice versa)



- **Training** = process of making the system able to learn
- **Testing** = process of seeing how well the system learned
 - Simulates “real world” usage
 - Training set and testing set should come from same distribution
 - Need to make some **assumptions** or introduce a bias
- **Deployment** = actually using the learned system in practice



QUESTIONS?

Deep robots!

Deep questions?!

