



Conclusion

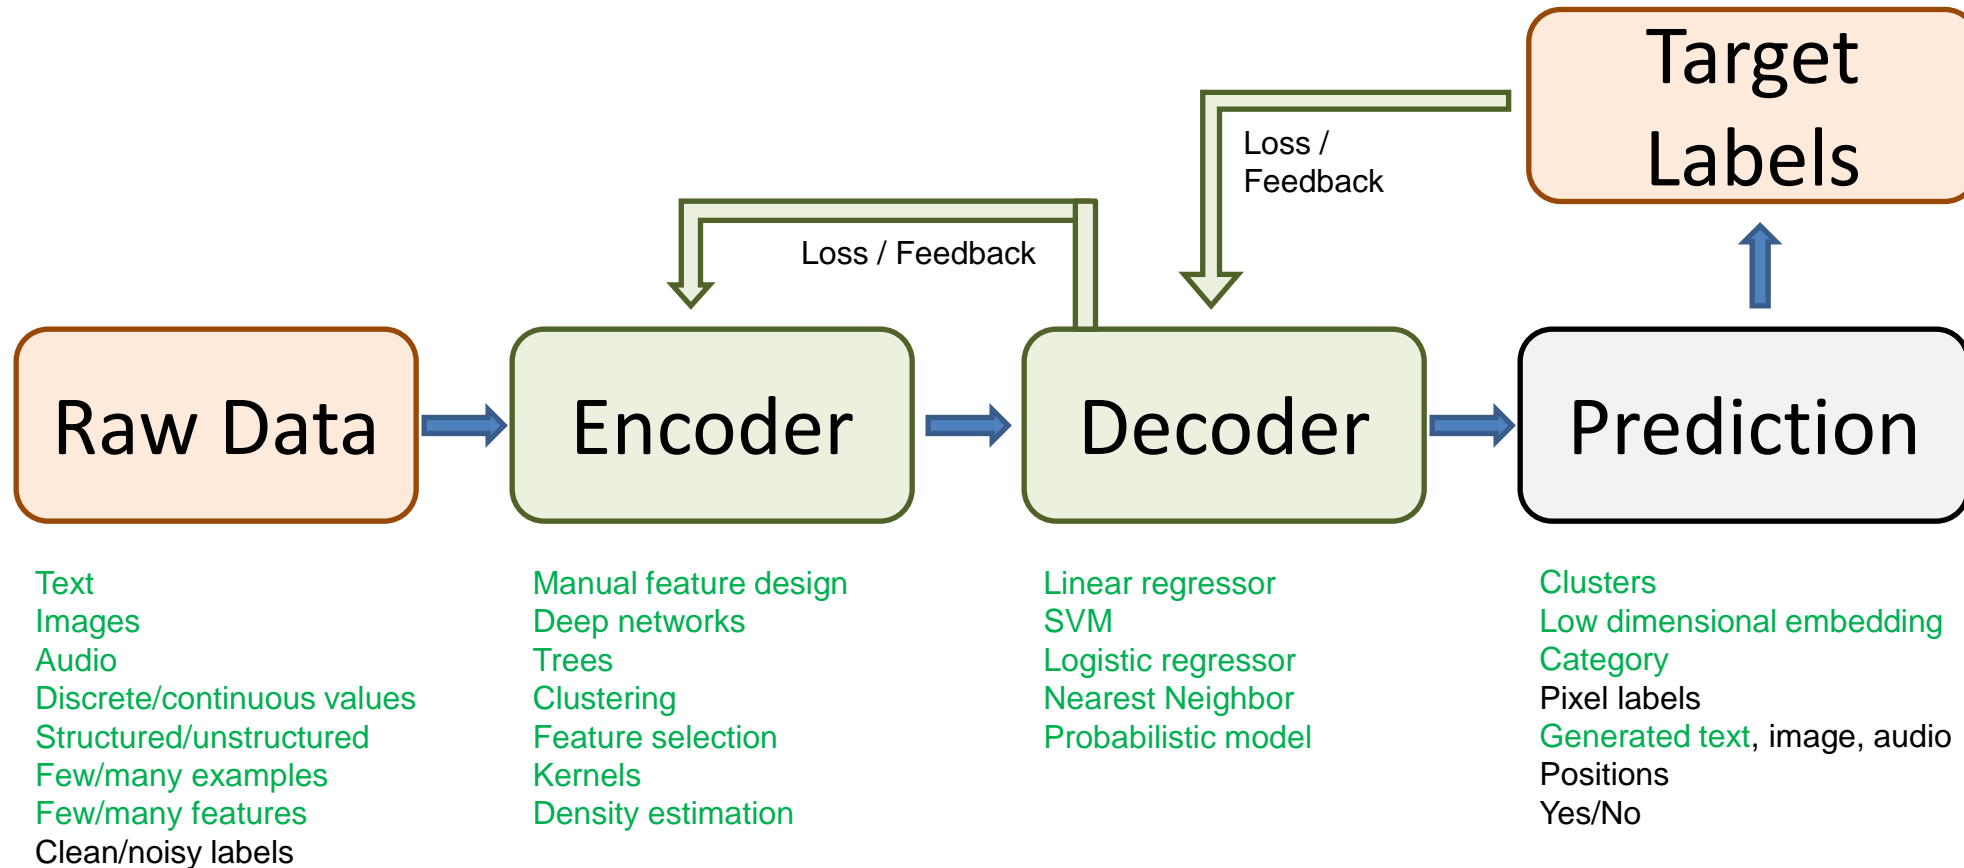
Applied Machine Learning
CS 441

Derek Hoiem

This class: Conclusion

- Birds-eye view of machine learning
- Where to learn more
- Trends and future of machine learning
- Feedback and closing remarks

We've learned a lot



We've covered what you need to get started for most ML applications

- Solid foundation in core ML concepts
 - Classification, clustering, dimensionality reduction, objectives, losses, regularization, generalization, experimental setups
- Intuition, some math, and application of classic methods
 - Logistic regression, Linear Regression, SVM, KNN, Neural Nets, Naïve Bayes, Boosting, Trees, Ensembles, PCA, K-means, Kernel Density Estimation, EM algorithm, PCA
- Exposure to state of the art recent methods
 - Deep convolutional networks, transformers, CLIP, GPT, UMAP
- Data representation and application domains
 - Images, text, audio, general data, 1D time series
- Practical considerations for deployment of ML applications
 - Societal impact, bias and mitigation, process for creating ML applications

Predictors: Three basic kinds

1. Linear (Regression, Logistic Regression, SVM)

- General purpose for classification/regression
- Benefits from feature learning

2. Instance-based (KNN)

- Zero training time, flexible
- Benefits from feature learning

3. Probabilistic: $P(\text{features} \mid \text{class})$

- General, but more complicated and often less accurate than linear with feature learning

Feature learning: Two basic kinds

- Tree
 - Learn how to partition input space to group examples together in a discriminative way
 - Followed by instance/linear/probabilistic prediction
 - Easily combine unnormalized continuous and discrete features
 - Especially powerful with ensembles, as in Random Forests or Boosted Trees
- Neural network
 - Jointly train feature representation with linear predictor
 - Major advantages for structured data with convolutional or transformer architectures, or use for retrieval/similarity
 - Representations can be learned on one problem and tuned or applied for another problem in the same domain

Go-to approaches for **classification**

- K-NN
 - Super easy and sometimes surprisingly effective, good sanity check
- Linear SVM or linear logistic regression
 - Easy to optimize, often works well
- Random forest or boosted trees
 - Highly effective with minimal fuss, very flexible
- Linear probe or fine-tune deep network
 - Best for common structured data like images, text, audio

Go-to approaches for **regression**

- Linear regression
 - Interpretable and often works well
 - Sometimes important to transform features or targets so they are better captured by a linear model
- Random forest
 - Highly effective with minimal fuss
- K-NN
 - Especially helpful when predicting multiple correlated values, e.g. which patch can be used to fill a hole in the image

In supervised learning, always remember

- **Clean experiments**
 - training set to learn model parameters
 - validation set to select method and hyperparameters
 - test set for final performance evaluation
- **Bias/variance trade-off**
 - Avoidable error is due to challenge in fitting parameters (variance) and inability to perfectly fit the data (bias)
 - Model designs and hyperparameters often trade off between these, e.g. increasing model complexity can increase variance but reduce bias
 - Ensembles work around this trade-off, and modern deep networks often act like ensembles
- People and algorithms use mental shortcuts that lead to a kind of **bias that can be harmful to society** – transparency and fairness require conscientious sourcing, development, and evaluation

Data organization

- PCA is used to compress a vector into fewer values in a way that can be decompressed with minimal mean squared difference
- Clustering reveals common modes of data
 - K-means is an essential algorithm
- Search for similar data items is an important application and is the computational foundation for clustering
 - FAISS is a very useful library for efficient search
 - Approximate search, e.g. trees and LSH, are needed when speed is a priority
 - Representation learning (e.g. with deep networks or latent semantic analysis) is needed to make similarity meaningful

It's all about the data

- Model architectures and computational techniques get all the attention, but there are often many reasonable choices that perform similarly
- To get best performance, data requires more thought and effort than algorithms
- Need creative ways to get freely supervised data, as well as careful curation of evaluation sets

Machine learning is easier and more effective than ever

- Great libraries like sklearn and pytorch
- Great models like CLIP, GPT, BERT
- Leverage existing models whenever possible to minimize development cost

Two minute break

Then, how to learn more and trends/future of ML

Where to learn more: courses

Many related 400-level courses, e.g.

- CS 412: Data Mining
- CS 442: Trustworthy ML
- CS 443: Reinforcement Learning (SP)
- CS 444: Deep Learning for Vision (SP)
- CS 445: Computational Photography
- CS 446: Machine Learning
- CS 447: NLP (SP)
- CS 448: Audio Computing Lab (SP)
- CS 449: Robot Perception

Also many 500-level courses, often requiring background in ML and some domain, e.g.

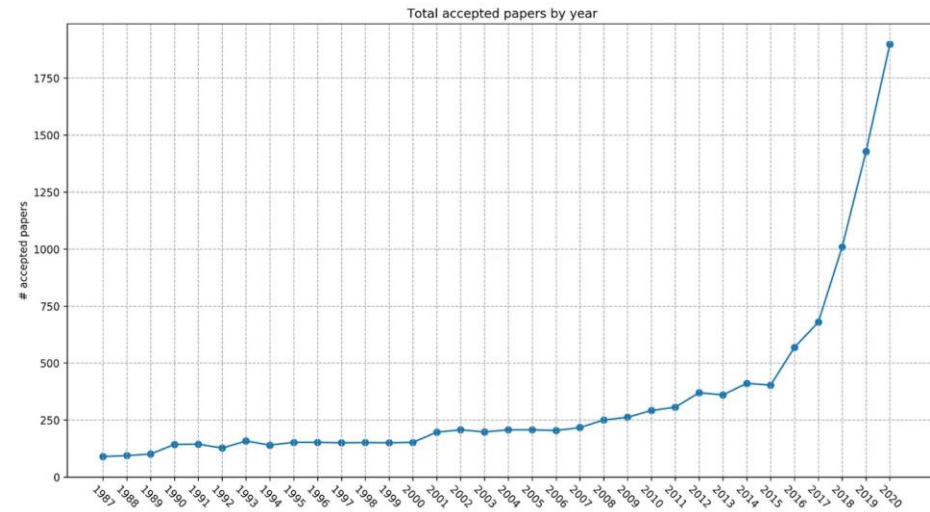

- CS 543 (Computer Vision)
- CS 545 (ML for Signals)
- CS 598 (Adv Robotics)
- CS 562 (Topics in Security and Privacy)

Where to learn more: try things

- Many datasets, github repos, and libraries available
- Easy to check out state-of-art in many problems or to try out different related algorithms

ML is changing

Google Trends: AI over past 5 years



Papers published in NeurIPS

2672 in 2022

10x growth in 12 years

ML is changing

- Machine learning has been used for decades, e.g. reading addresses for the post office, but it's now everywhere
- Major increases in conference attendance, press coverage, investment by companies, number of authors per paper
- Much better infrastructure for ML in terms of libraries and models, as well as tutorials and other online resources
- Short-term value in AI is more clear, causing many companies to focus their investments

AI/ML will change the world in unpredictable ways

- We have seen transformations of agriculture and manufacturing – AI will transform services
 - In a three sector breakdown, the first two sectors of raw materials (including agriculture) and production have already been transformed by technology
 - In consequence, the majority of the workforce has move into services in advanced countries, as well as people working less on average
 - AI may dramatically transform services, but it's not clear where the workforce can move
- Previous products were predictable; AI, not so
 - Objects and computer software are under direct control of their users and operate in predictable ways
 - ML-based products depend on data in non-transparent ways

How to prepare yourself for a fast-changing tech world?

Master the fundamentals

- Tools and algorithms change, but the ideas behind them are relatively constant
- Fundamentals are math, ability to frame problems computationally, ability to draw insights and predict from data, communication, and empathy
- These abilities enable you to learn the latest, design solutions, and apply effectively

Current state of AI/ML

- Fundamental ideas of classification and regression are well established and haven't changed in many years
- Recent advances enable integrating lots of data to create powerful predictors/generators
- AI capability is still far inferior to humans
 - Requires clearly defined tasks
 - Cannot perform complex tasks
 - Requires many examples
 - Difficult to adapt to new tasks

Future directions for machine learning

- Using ML to perform complex tasks
 - Perform software tasks, such as image editing, drafting slides based on an outline, or summarizing and reporting information from diverse sources
 - Learn from watching how people perform tasks and through natural training interfaces
- Integrating ML into physical devices
 - IoT, sensors
 - Robots
- Multimodal and broadly knowledgeable AI
 - Models increasingly defined in terms of input/output modalities rather than narrow tasks
 - Use vision, language, situation-specific, and broader knowledge to summarize and recommend
- Complementary learning systems
 - Integrate predictive/generative models and retrieval of individual memories and associations
 - Integrate with planning, task decomposition, and prioritization

One example: Visual Programming

- Use LLM to generate code that calls pretrained vision models and functions

Instruction: Replace the ground with white snow and the bear with a white polar bear



```
OBJ0=Seg(image=IMAGE)
OBJ1=Select(image=IMAGE,object=OBJ0,query='ground')
IMAGE0=Replace(image=IMAGE,object=OBJ1,prompt='white snow')
OBJ2=Seg(image=IMAGE0)
OBJ3=Select(image=IMAGE0,object=OBJ2,query='bear')
IMAGE1=Replace(image=IMAGE0,object=OBJ3,prompt='white polar bear')
RESULT=IMAGE1
```



Prediction:



Compositional Visual Question Answering

IMAGE:



Question: Are there both ties and glasses in the picture?

Program:

```

BOX0=Loc(image=IMAGE, object='ties')
ANSWER0=Count(box=BOX0)
BOX1=Loc(image=IMAGE, object='glasses')
ANSWER1=Count(box=BOX1)
ANSWER2=Eval("'yes' if {ANSWER0} > 0 and {ANSWER1} > 0 else 'no'")
RESULT=ANSWER2
Prediction: no

```

Natural Language Visual Reasoning

LEFT:



RIGHT:



Statement: The left and right image contains a total of six people and two boats.

Program:

```

ANSWER0=Vqa(image=LEFT, question='How many people are in the image?')
ANSWER1=Vqa(image=RIGHT, question='How many people are in the image?')
ANSWER2=Vqa(image=LEFT, question='How many boats are in the image?')
ANSWER3=Vqa(image=RIGHT, question='How many boats are in the image?')
ANSWER4=Eval('{ANSWER0} + {ANSWER1} == 6 and {ANSWER2} + {ANSWER3} == 2')
RESULT=ANSWER4
Prediction: False

```

Factual Knowledge Object Tagging

IMAGE:



Prediction: IMAGE0



Instruction: Tag the 7 main characters on the TV show Big Bang Theory

Program:

```

OBJ0=FaceDet(image=IMAGE)
LIST0=List(query='main characters on the TV show Big Bang Theory', max=7)
OBJ1=Classify(image=IMAGE, object=OBJ0, categories=LIST0)
IMAGE0=Tag(image=IMAGE, object=OBJ1)
RESULT=IMAGE0

```

Natural Language Image Editing

IMAGE:



Prediction: IMAGE1



Instruction: Hide Daniel Craig with 8) and Sean Connery with ;)

Program:

```

OBJ0=FaceDet(image=IMAGE)
OBJ1=Select(image=IMAGE, object=OBJ0, query='Daniel Craig', category=None)
IMAGE0=Emoji(image=IMAGE, object=OBJ1, emoji='smiling_face_with_sunglasses')
OBJ2=Select(image=IMAGE, object=OBJ0, query='Sean Connery', category=None)
IMAGE1=Emoji(image=IMAGE0, object=OBJ2, emoji='winking_face')
RESULT=IMAGE1

```

IMAGE:



Prediction: IMAGE0



Instruction: Replace desert with lush green grass

Program:

```

OBJ0=Seg(image=IMAGE)
OBJ1=Select(image=IMAGE, object=OBJ0, query='desert', category=None)
IMAGE0=Replace(image=IMAGE, object=OBJ1, prompt='lush green grass')
RESULT=IMAGE0

```

IMAGE:



Prediction: IMAGE0



Instruction: Create a color pop of Barack Obama (person)

Program:

```

OBJ0=Seg(image=IMAGE)
OBJ1=Select(image=IMAGE, object=OBJ0, query='Barack Obama', category='person')
IMAGE0=ColorPop(image=IMAGE, object=OBJ1)
RESULT=IMAGE0

```

Finishing up the semester

- Wed May 3: Final project due at 11:59pm
- May 9: Final exam covers entire semester
 - Covers all material for semester
 - You may start between 9:30am on May 9 to 9:30am on May 10
- See CampusWire pinned posts for details
- Only one more office hour after May 3: Derek's at 11am on May 8
 - We will still monitor and respond to CampusWire until the exam period
 - We will not answer questions during the exam period, unless it pertains to access or technical difficulty with PrairieLearn

ICES Feedback

- We put much time and energy into this course
- Please take some time to provide ratings and feedback
 - go.illinois.edu/ices-online
- Good things to know about ICES in general
 - Considered mainly as a reflection of the instructor and considered for promotions and course improvements
 - Reflect aspects of the course that are controllable by instructor, e.g. content, structure, and behavior. Can comment on facilities, time, and your external workload, but it's best not to factor them into overall ratings
 - Unconscious bias that negatively affects women, people of color, and other nationalities is well documented; try to consciously counter-act it

Thank you to the course staff for all of their hard work to create materials, grade, and answer questions.

Thank you for your hard work and engagement!