(Biased) Overview of A.I. Topics

David I. Inouye

Wednesday, August 26, 2020

High-Level Categorization of AI Topics

1. Artificial Intelligence (other than topics below)
2. Machine Learning (& Deep Learning)
3. Computer Vision
4. Natural Language Processing

1. Artificial Intelligence (Based on AAAI topic list)

- Cognitive modeling and systems
- Constraint Satisfaction/ Optimization
- Game theory
- Human + AI
- Knowledge representation and reasoning
- Robotics

AI: Cognitive Modeling

- Models of human/animal cognition
- Based on psychological theory and experiments
- 2 Goals
  - AI -> Cognitive Science: Understand/test underlying cognitive mechanisms by computational modeling
  - Cognitive Science -> AI: Improve computational methods via insights from cognitive science

AI: Constraint Satisfaction / Heuristic Optimization

- Eight queens puzzle

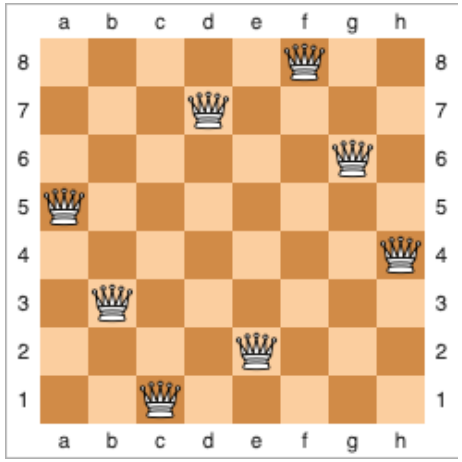

- Map coloring problem

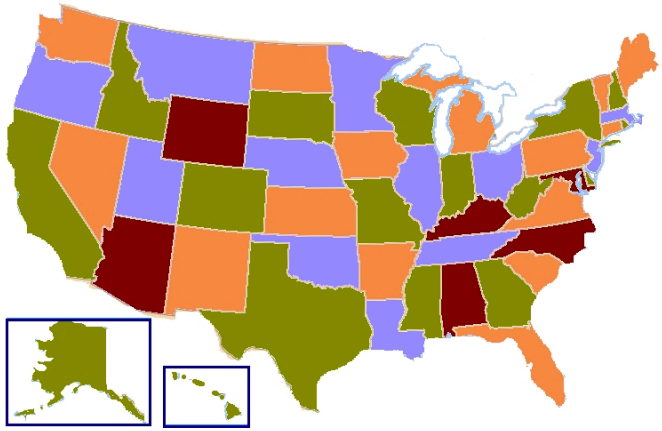

- Real-world
  - Resource allocation
  - Scheduling

AI: Game Theory

- Prisoner's dilemma

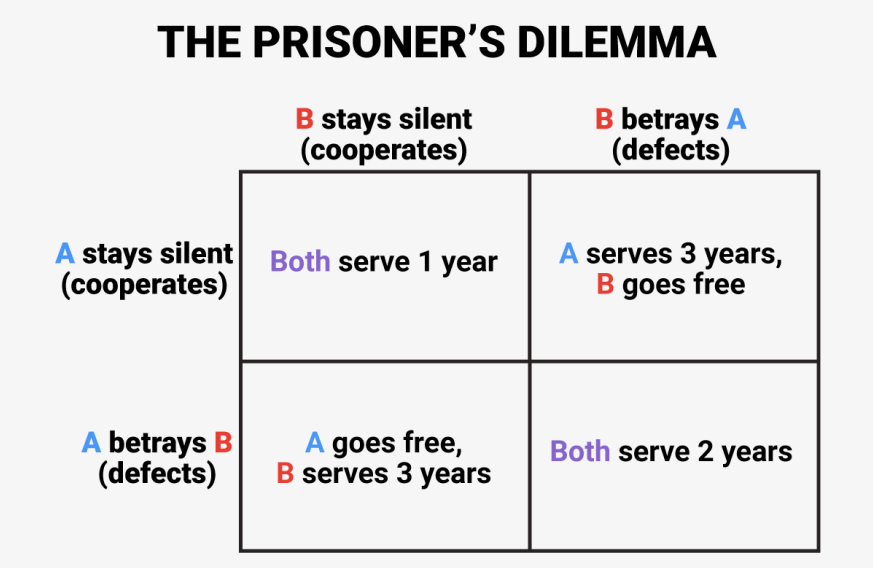

- Real-world
  - Google Ads bidding
  - Connections to "Generative Adversarial Networks"

AI: Human + AI

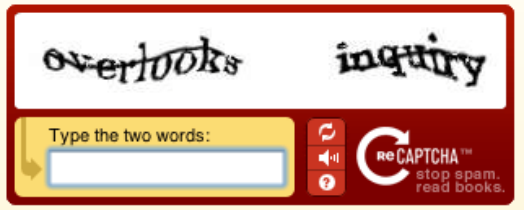

- Crowdsourcing
  - "Stop spam, read books"

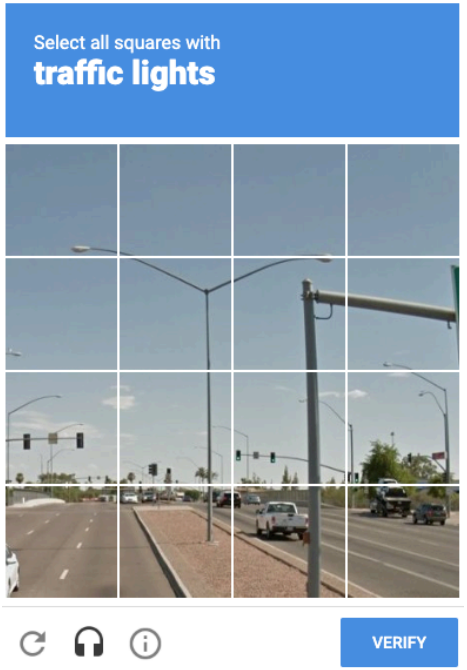

- Human-robot interactions

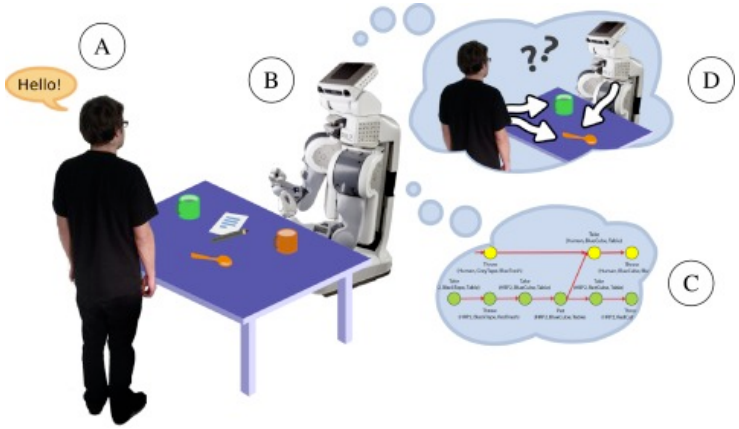

https://www.sciencedirect.com/science/article/pii/S0004370216300790

AI: Knowledge representation and reasoning

- Knowledge graphs

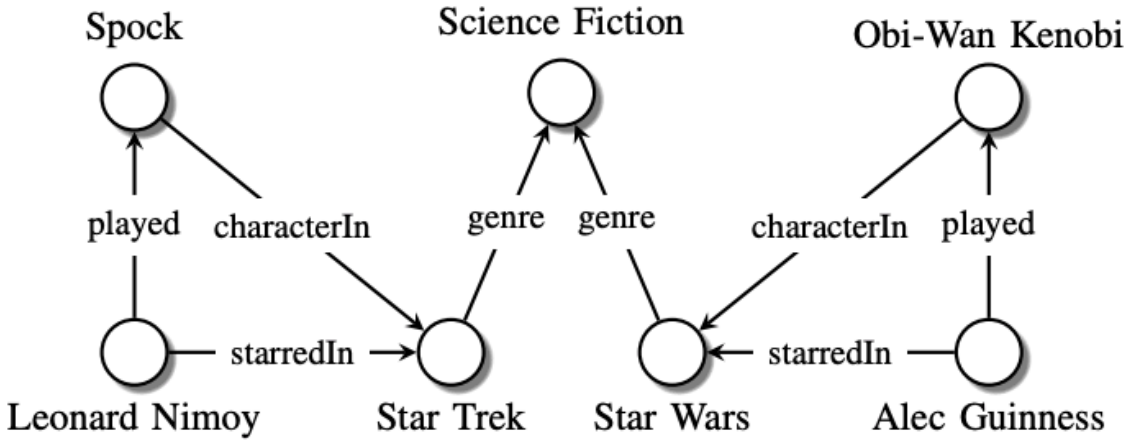

https://arxiv.org/pdf/1503.00759.pdf

- Inferences in knowledge graphs
  - Did Alec Guinness ever play a Science Fiction character?

2. Machine Learning (based on NeurIPS Topics)

- Learning with limited labels
- Generative / probabilistic models
- Reinforcement learning
- Explainable AI

ML: Learning with limited labels

- Active learning

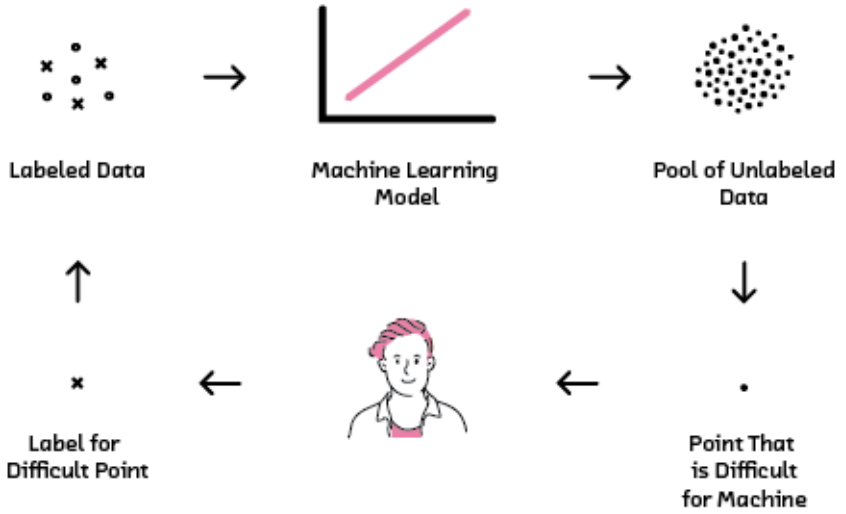

https://blog.cloudera.com/a-guide-to-learning-with-limited-labeled-data/

- Few-Shot Learning

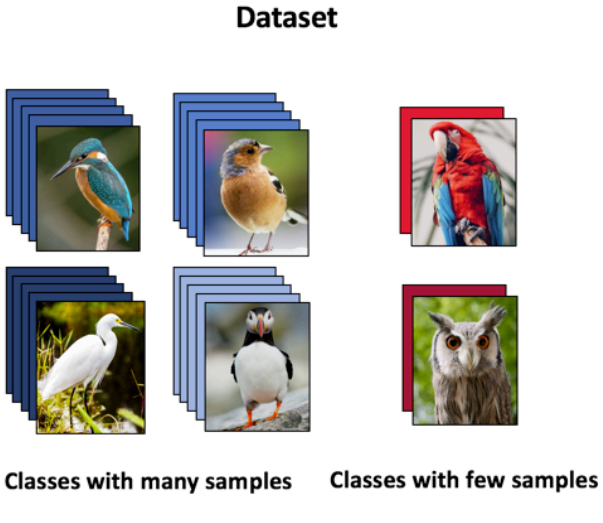

https://medium.com/sap-machine-learning-research/deep-few-shot-learning-a1caa289f18

ML: AutoML / Meta-learning

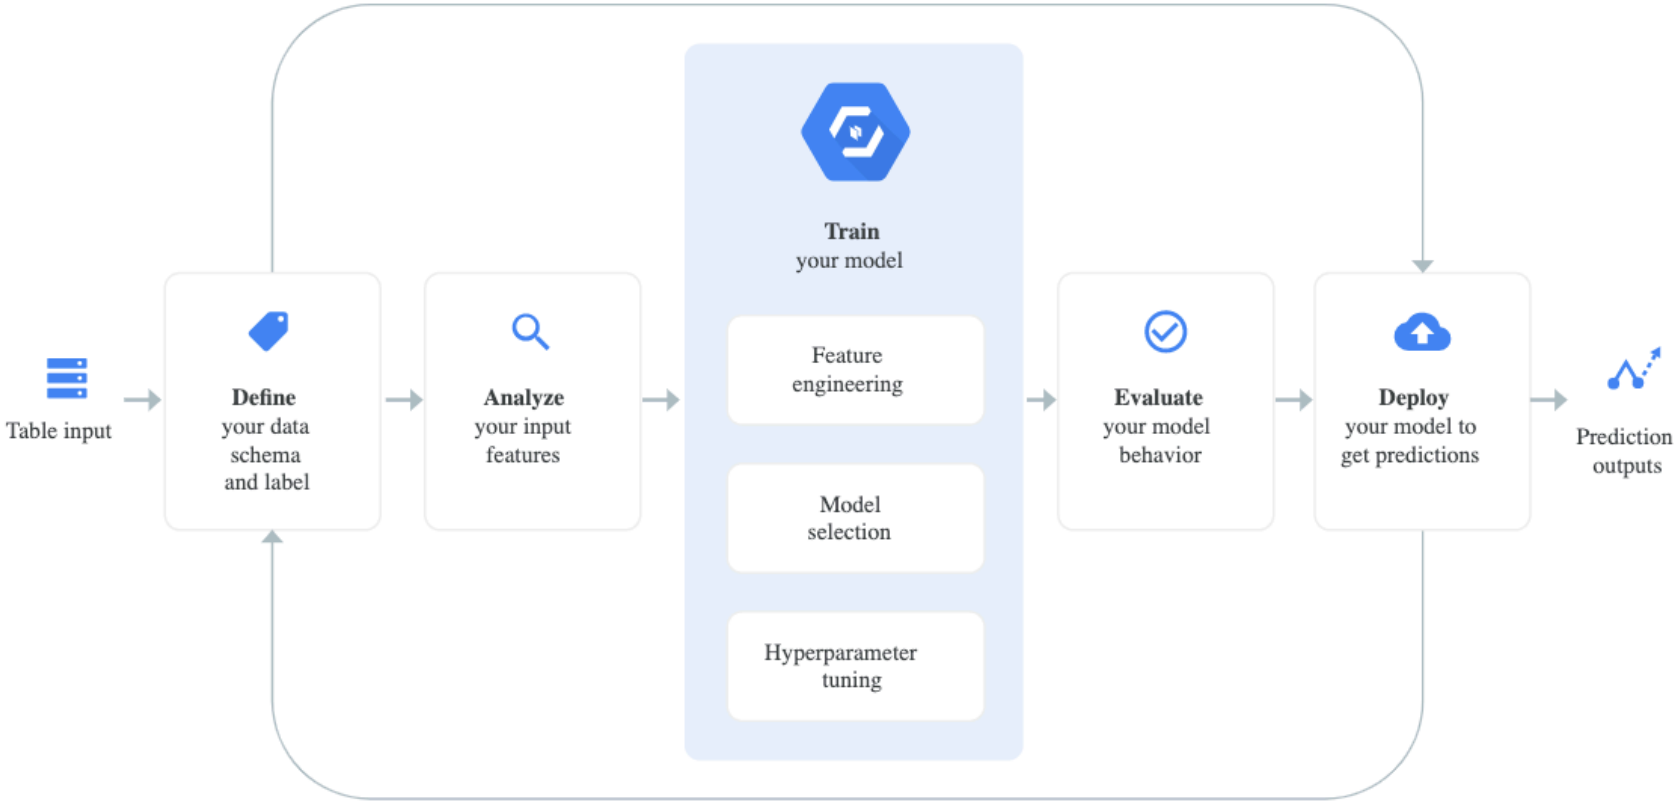

https://cloud.google.com/automl-tables/?hl=vi

ML: Generative/Probabilistic Models

- Density estimation / Graphical Models

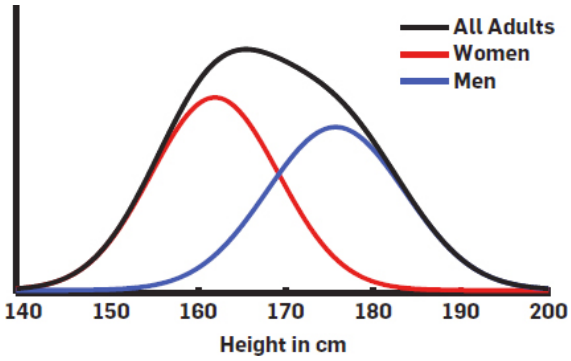

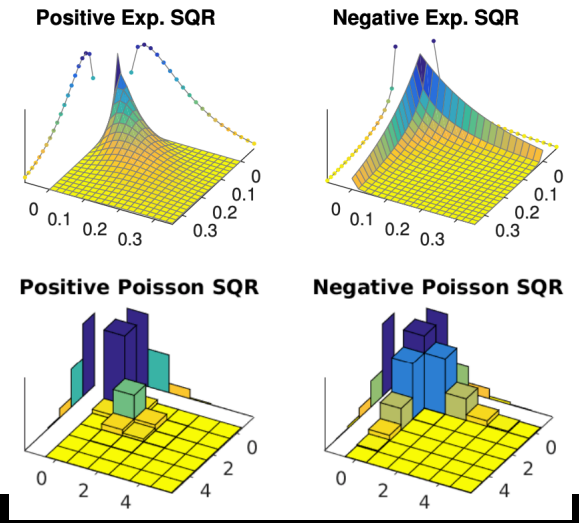

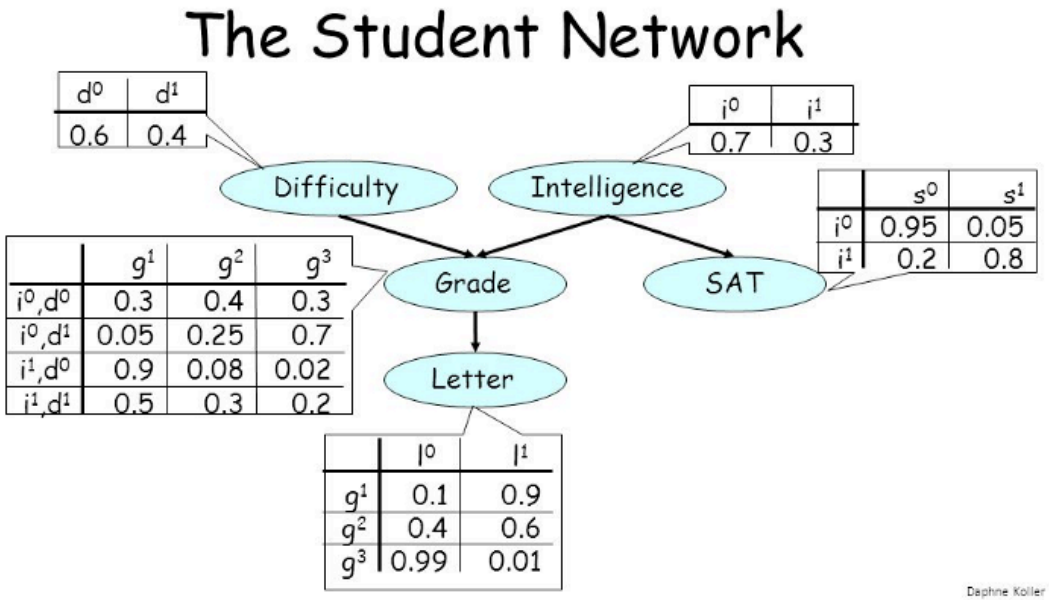

d^0	d^1
0.6	0.4

i^0	i^1
0.7	0.3

	s^0	s^1
i^0	0.95	0.05
i^1	0.2	0.8

	g^1	g^2	g^3
i^0,d^0	0.3	0.4	0.3
i^0,d^1	0.05	0.25	0.7
i^1,d^0	0.9	0.08	0.02
i^1,d^1	0.5	0.3	0.2

	l^0	l^1
g^1	0.1	0.9
g^2	0.4	0.6
g^3	0.99	0.01

ML: Topic Models

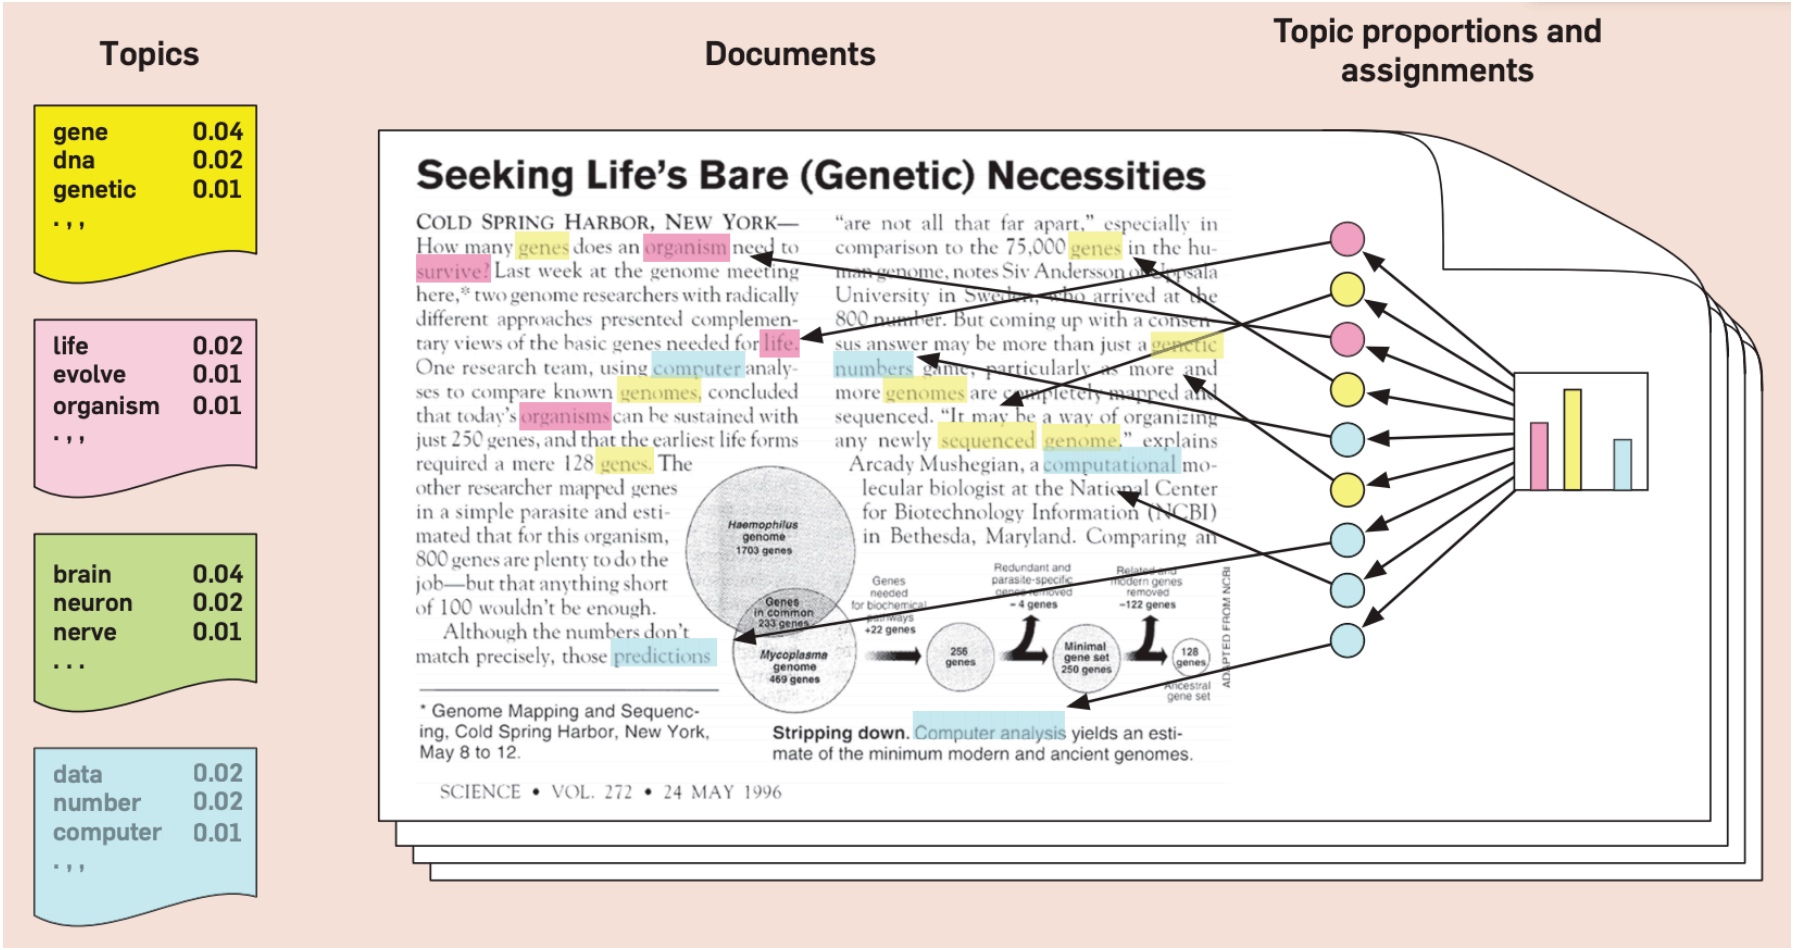

http://www.cs.columbia.edu/~blei/papers/Blei2012.pdf

ML: Generative Adversarial Networks (GAN)

- Generative Adversarial Networks (GAN)

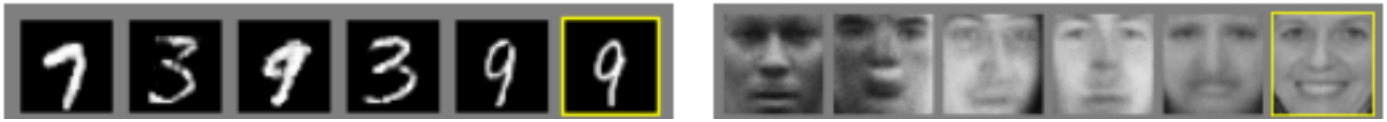

http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf

- Image to image translation via GANs

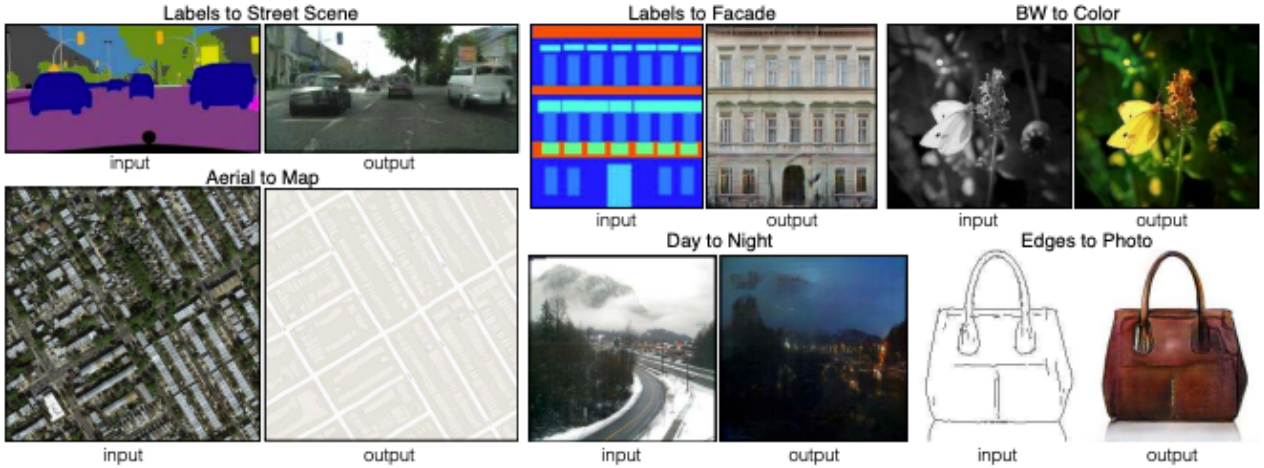

http://openaccess.thecvf.com/content_cvpr_2017/papers/Isola_Image-To-Image_Translation_With_CVPR_2017_paper.pdf

ML: Invertible networks

- Invertible Flows

https://papers.nips.cc/paper/8224-glow-generative-flow-with-invertible-1x1-convolutions.pdf

- Deep Density Destructors

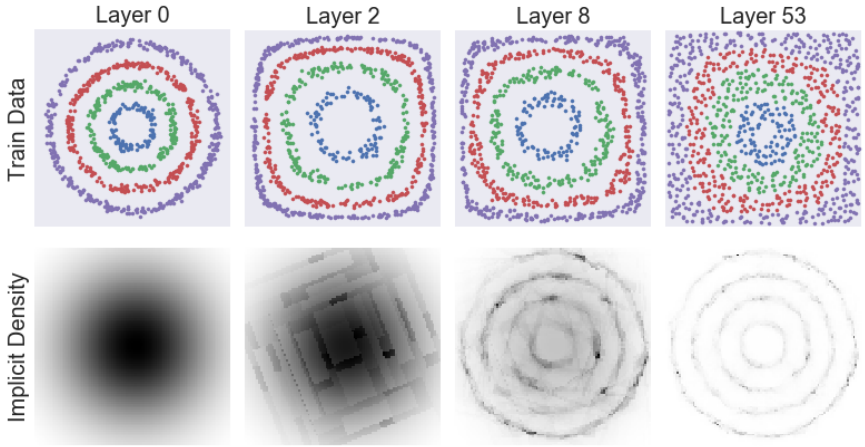

https://www.davidinouye.com/publication/inouye-2018-deep/inouye-2018-deep.pdf

ML: Reinforcement Learning

- Game playing

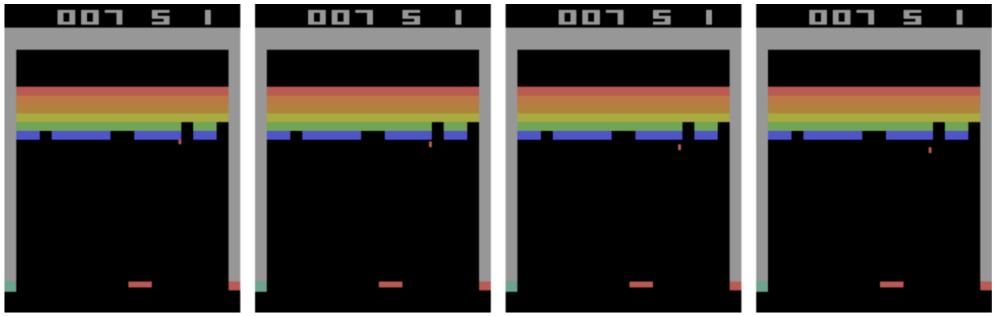

- Bandit algorithms (simpler form of RL)
  - Which Google search result should I show?

ML: Explainable AI

Why model explanations? Accuracy is insufficient for many applications

Loan approval: “Could the model make a catastrophic mistake?”

Self-driving cars: “Does the model obey common sense intuitions?”

Bail decisions: “Is the model biased because of historical discrimination?”

Healthcare: “Does the model agree with doctor’s knowledge?”

Military strategy: “How will the model perform in adversarial settings?”

3. Computer Vision (Based on CVPR sessions)

- Classic tasks
- 3D Multiview / Depth estimation
- Synthesis

CV: Classic Tasks

- Recognition

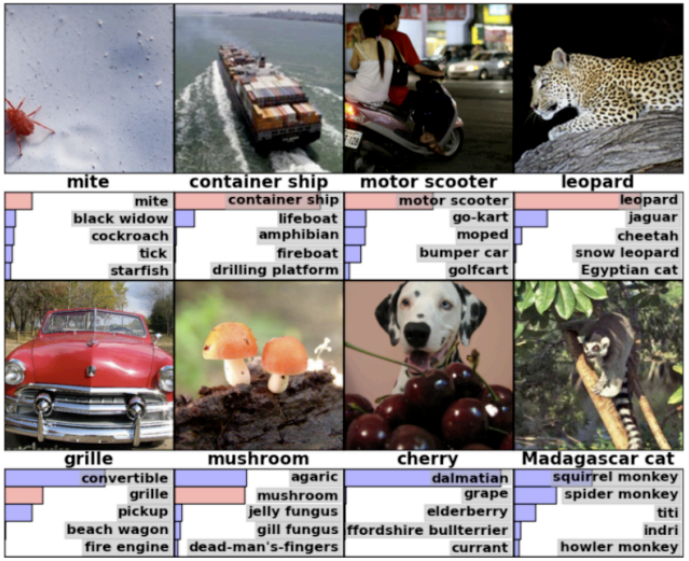

http://www.cs.toronto.edu/~fritz/absps/imagenet.pdf

- Segmentation

http://vladlen.info/publications/feature-space-optimization-for-semantic-video-segmentation/

CV: 3D Multiview / Depth estimation

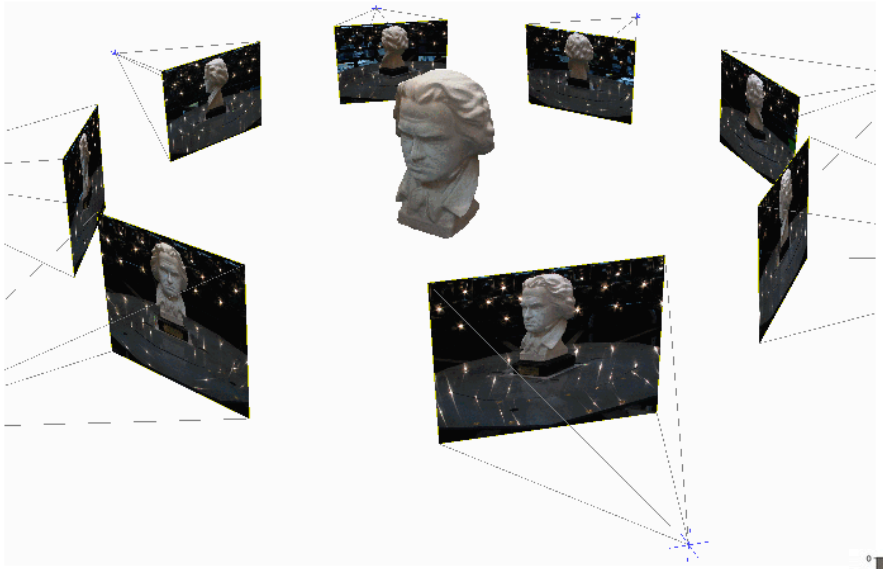

https://vision.in.tum.de/research/image-based_3d_reconstruction/multiviewreconstruction

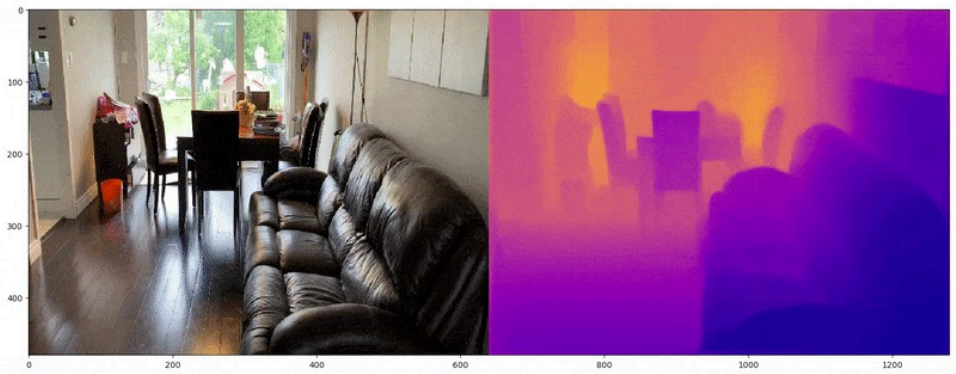

https://towardsdatascience.com/depth-estimation-on-camera-images-using-densenets-ac454caa893

CV: Image / Video Generation (Synthesis)

- Style transfer

https://www.cv-foundation.org/openaccess/content_cvpr_2016/papers/Gatys_Image_Style_Transfer_CVPR_2016_paper.pdf

- Sketch to draw

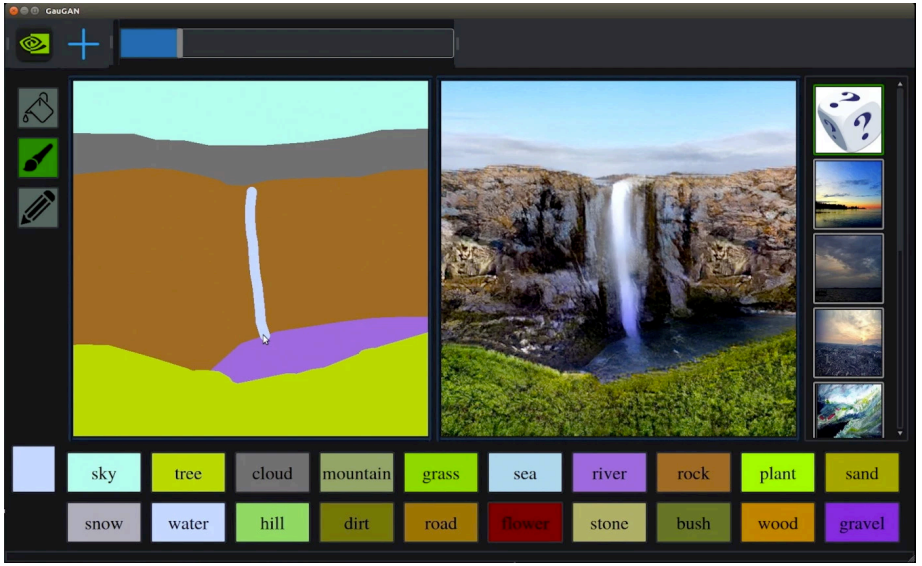

https://techcrunch.com/2019/03/18/nvidia-ai-turns-sketches-into-photorealistic-landscapes-in-seconds/

4. Natural Language Processing (based on ACL 2019 Call for Papers (CFP))

- Tagging and Parsing
- Information Extraction and Text Mining
- Dialogue Systems / Question Answering
- Applications
  - Summarization
  - Sentiment Analysis
  - Machine Translation

NLP: Ambiguity is huge challenge in NLP

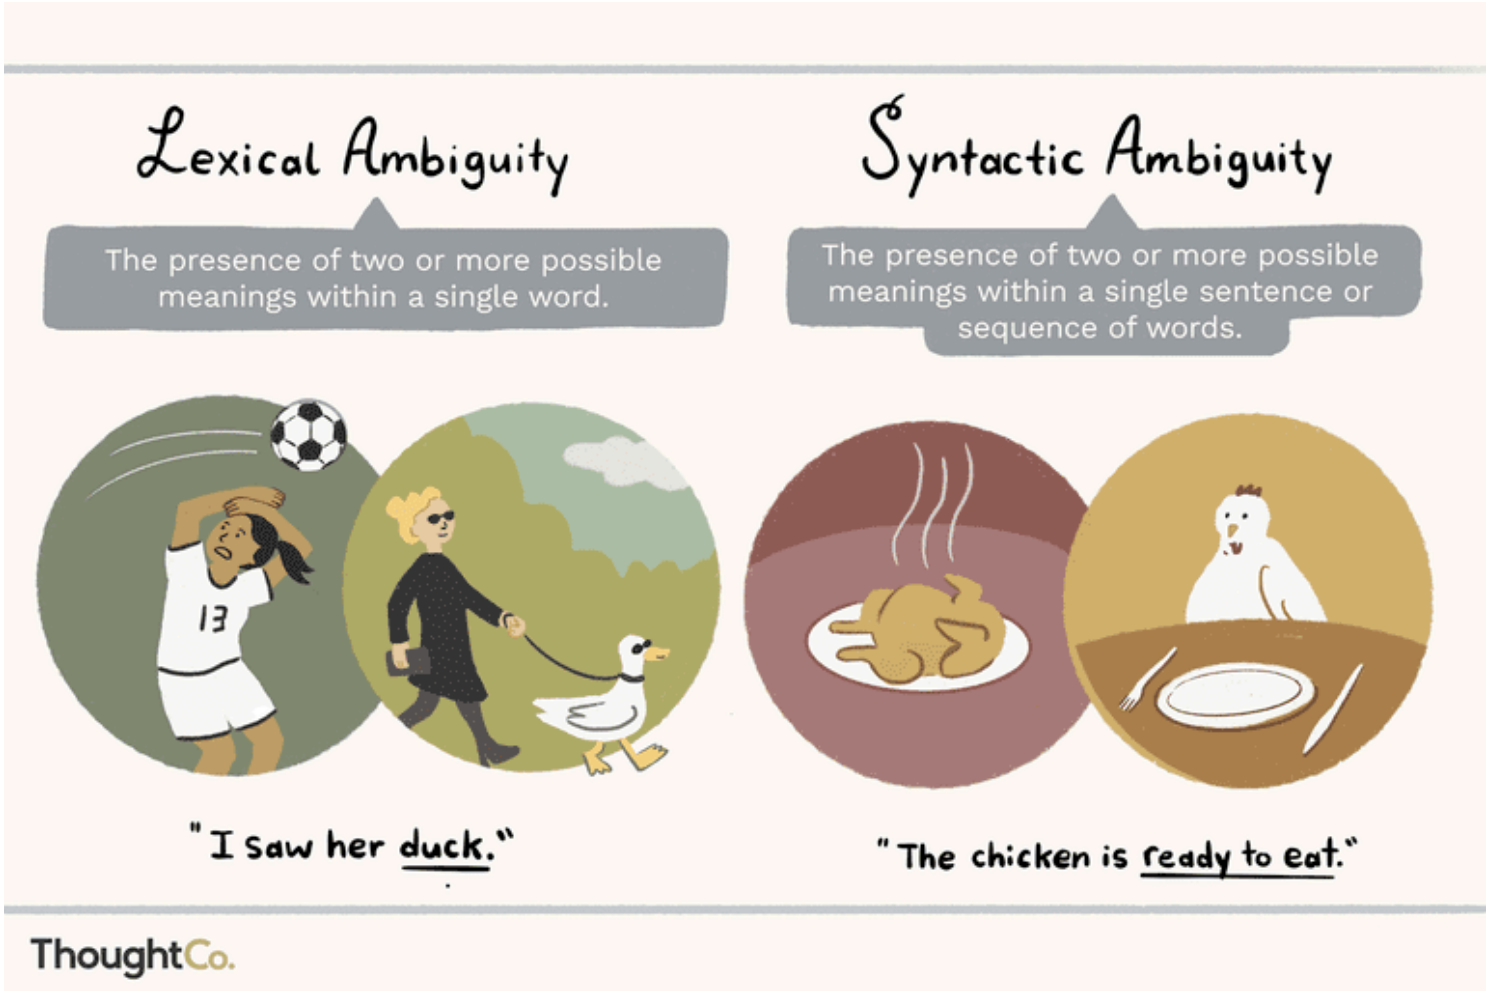

https://www.thoughtco.com/ambiguity-language-1692388

NLP: Tagging and Parsing

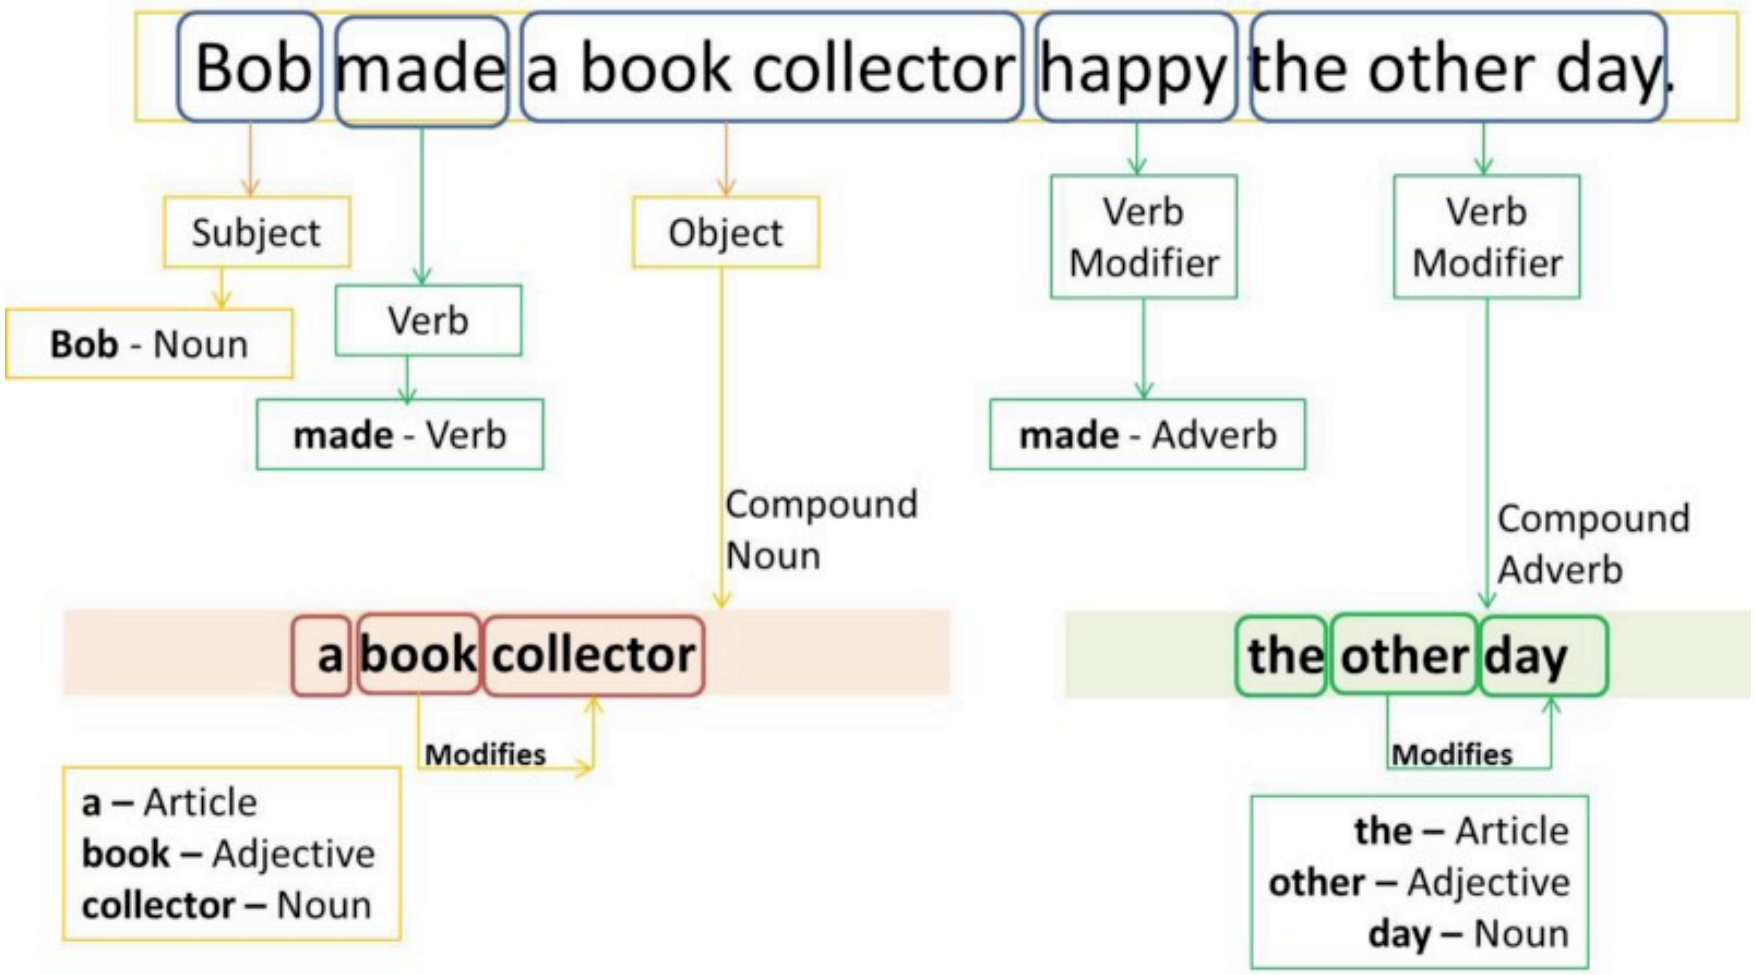

NLP: Information Extraction and Text Mining

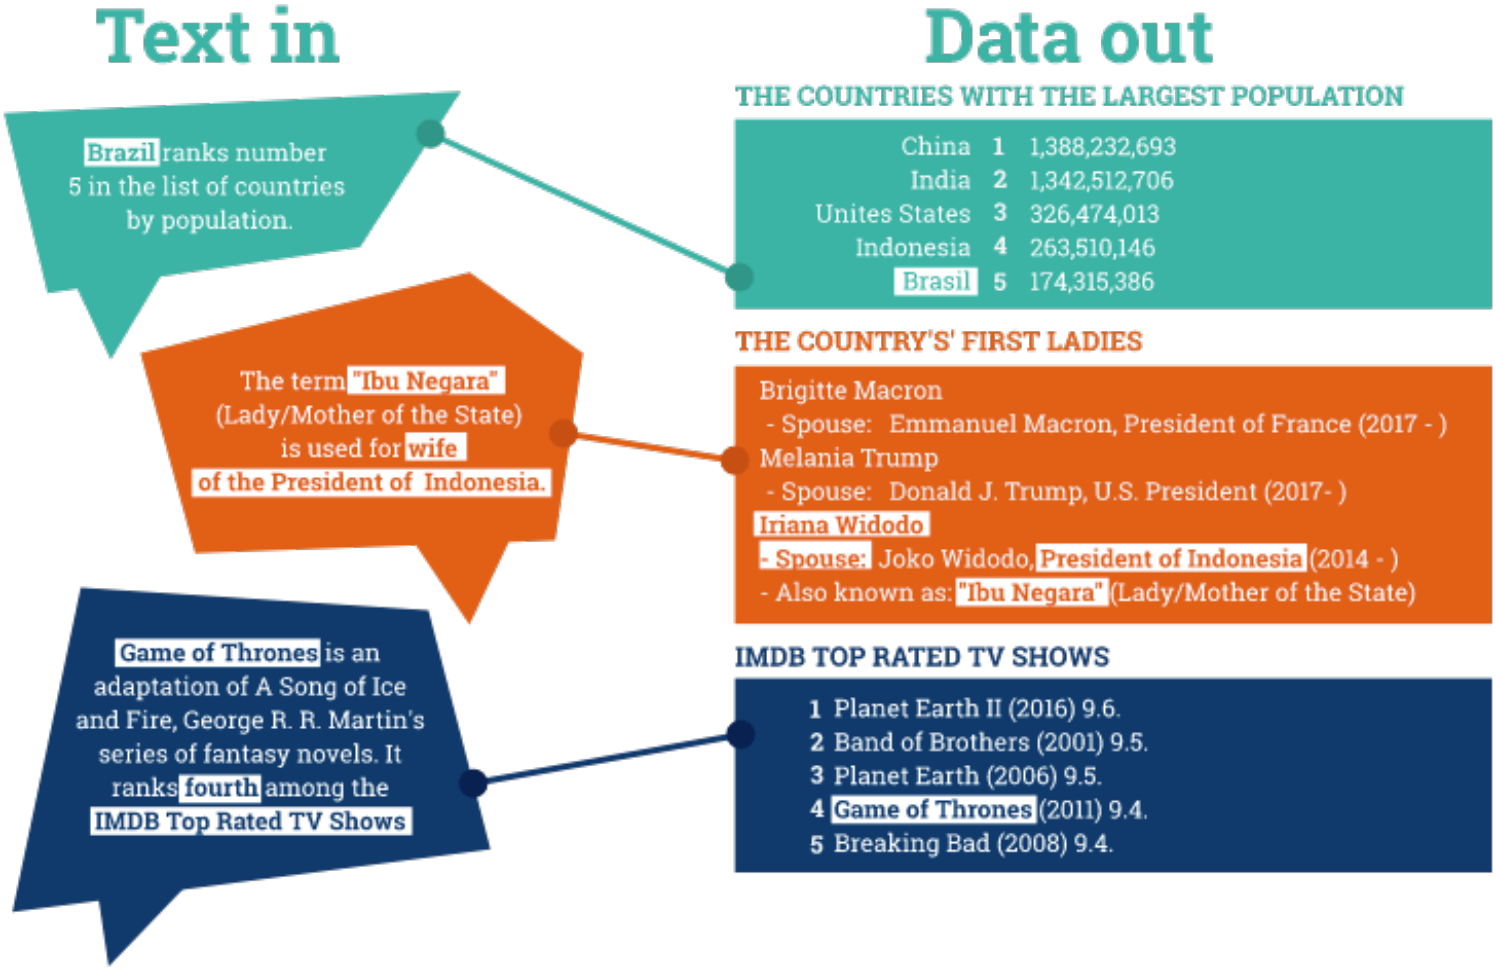

https://www.ontotext.com/knowledgehub/fundamentals/information-extraction/

NLP: Dialogue Systems / Question Answering

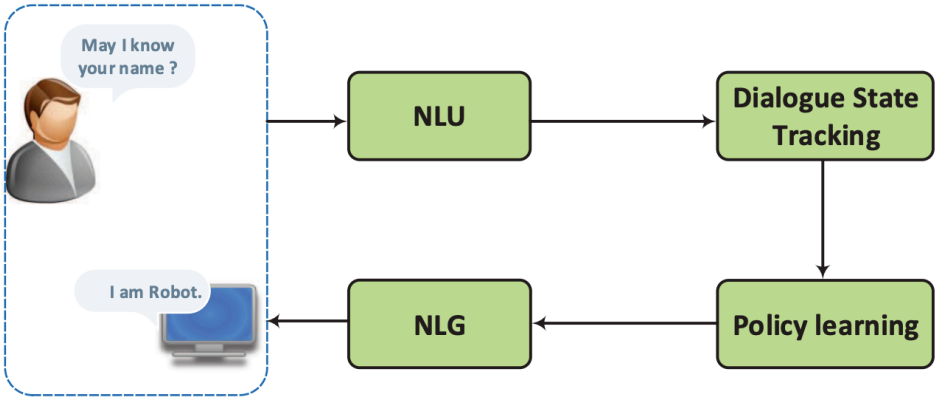

https://arxiv.org/pdf/1711.01731.pdf

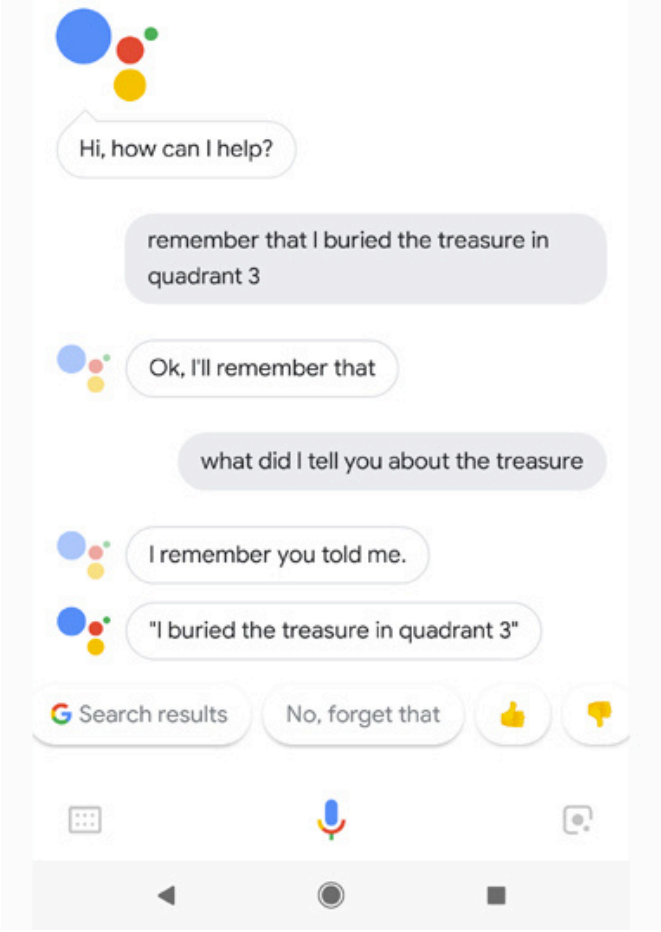

High-Level Categorization of AI Topics

1. Artificial Intelligence (other than topics below)
2. Machine Learning (& Deep Learning)
3. Computer Vision
4. Natural Language Processing