

(Biased) Overview of A.I. Topics

David I. Inouye

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](#))

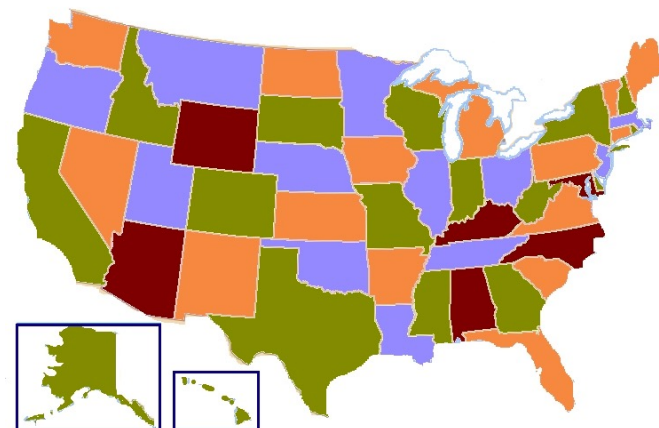
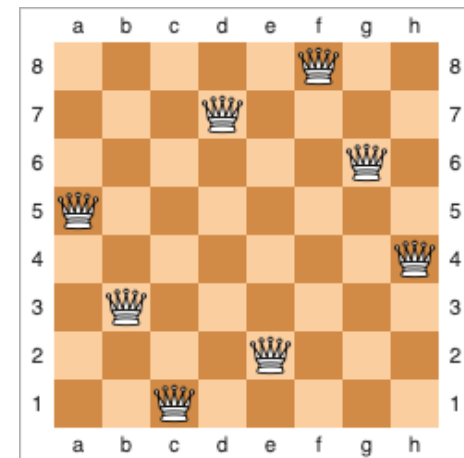
- ▶ Cognitive modeling and systems
- ▶ Constraint Satisfaction/ Optimization
- ▶ Game theory
- ▶ Human + AI
- ▶ Knowledge representation and reasoning
- ▶ Robotics

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

THE PRISONER'S DILEMMA

	B stays silent (cooperates)	B betrays A (defects)
A stays silent (cooperates)	Both serve 1 year	A serves 3 years, B goes free
A betrays B (defects)	A goes free, B serves 3 years	Both serve 2 years

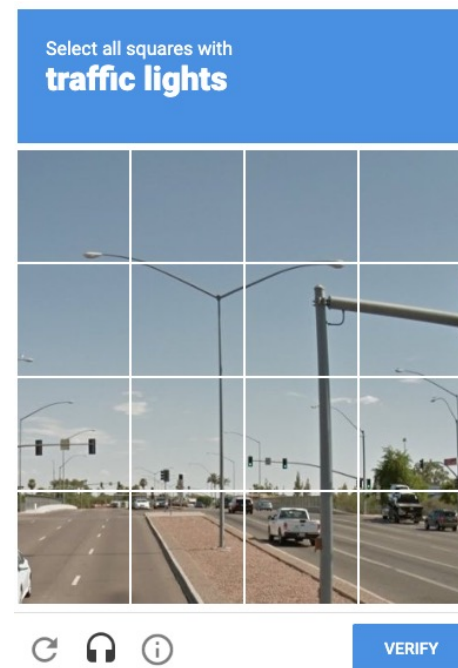
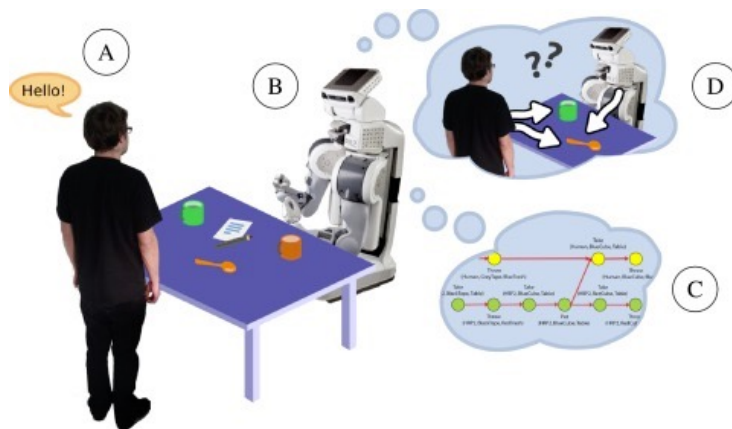
▶ Real-world

- ▶ Google Ads bidding
- ▶ Connections to “Generative Adversarial Networks”

AI: Human + AI

- ▶ Crowdsourcing
 - ▶ “Stop spam, read books”

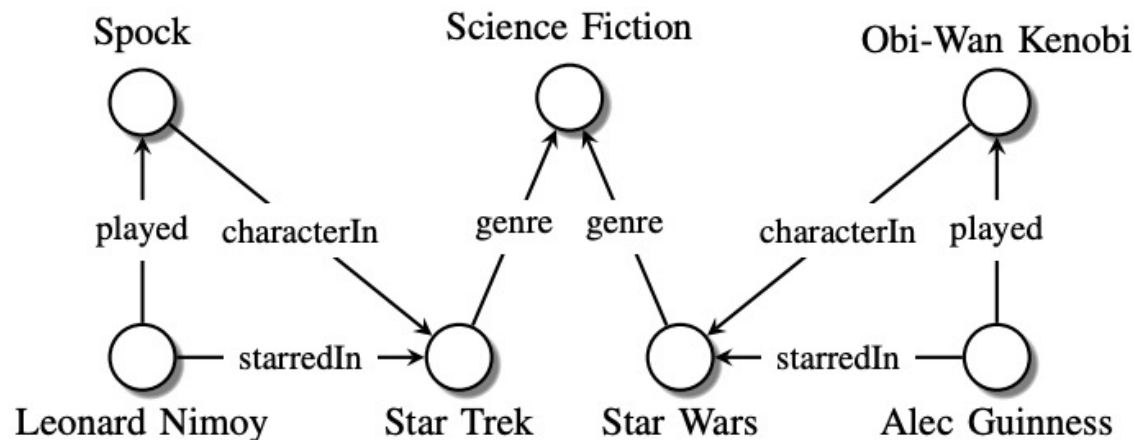
- ▶ Human-robot interactions



<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

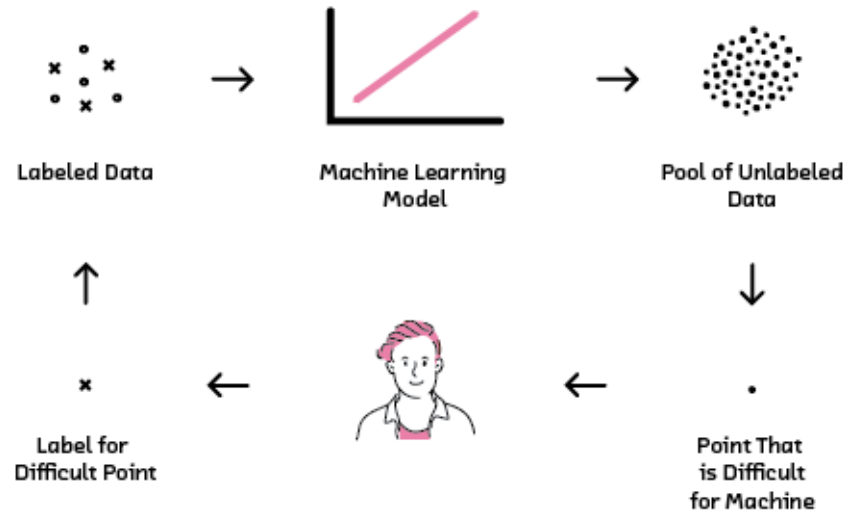
- Did Alec Guinness ever play a Science Fiction character?

2. Machine Learning (based on [NeurIPS Topics](#))

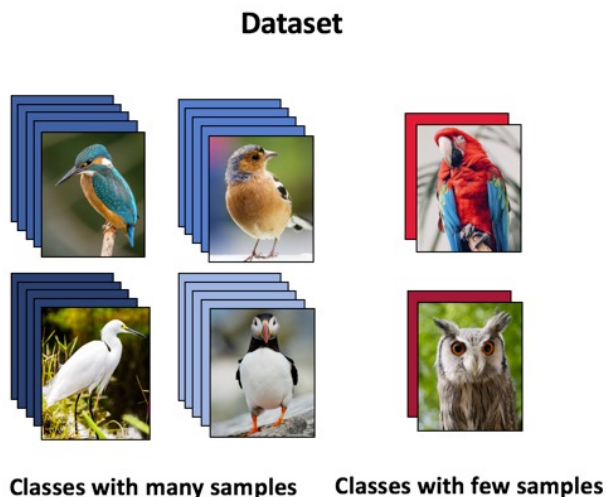
- ▶ Learning with limited labels
- ▶ AutoML / Meta-learning
- ▶ Generative / probabilistic models
- ▶ Reinforcement learning
- ▶ Explainable AI
- ▶ Domain Generalization

ML: Learning with limited labels

▶ Active learning



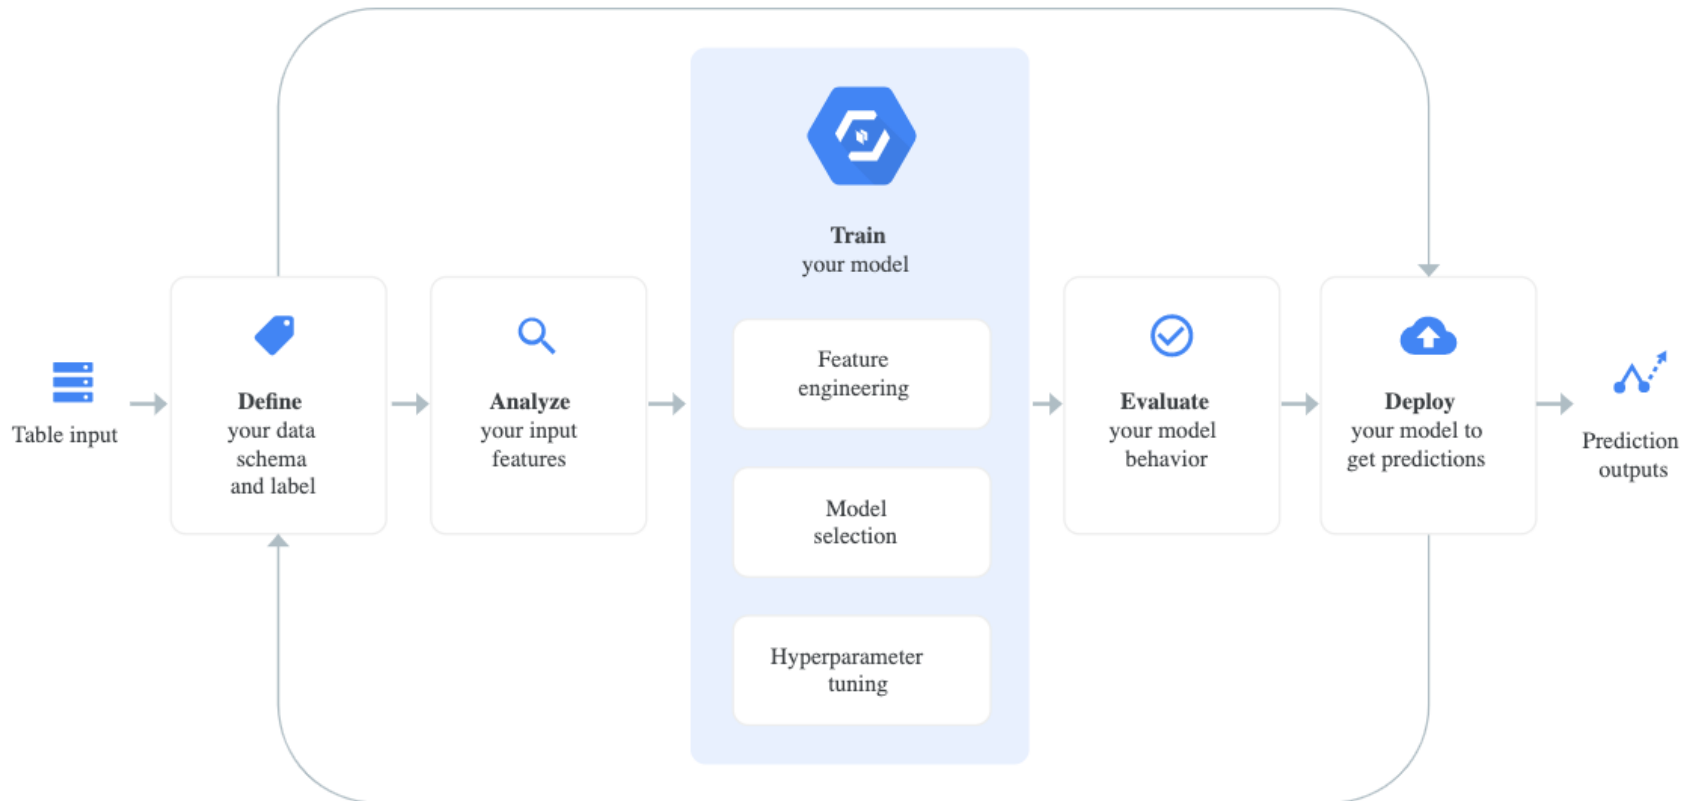
▶ Few-Shot Learning



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

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

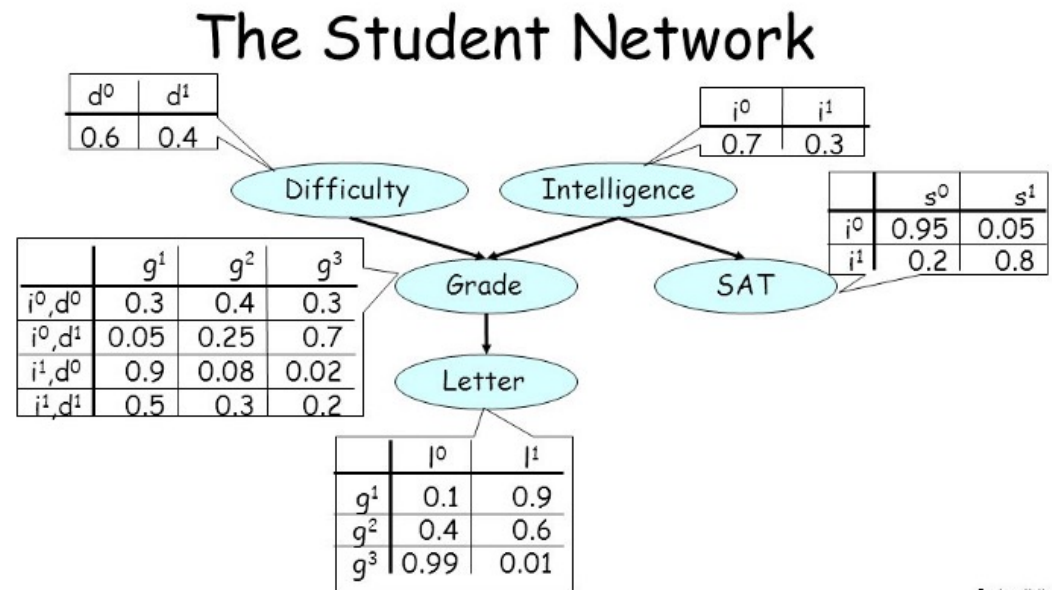
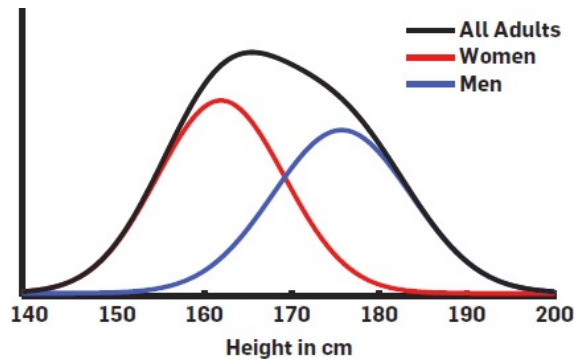
ML: AutoML / Meta-learning



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

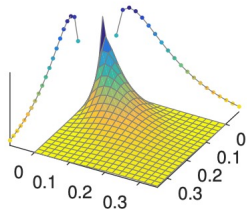
ML: Generative/Probabilistic Models

► Density estimation / Graphical Models

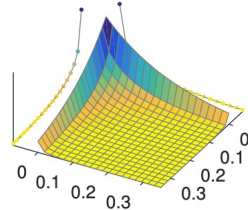


Dephine Koller

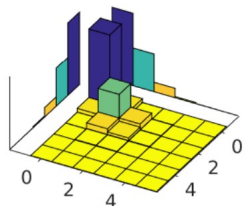
Positive Exp. SQR



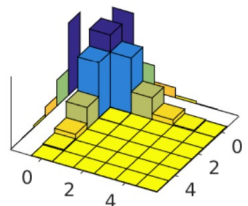
Negative Exp. SQR



Positive Poisson SQR



Negative Poisson SQR



ML: Topic Models

Topics

gene	0.04
dna	0.02
genetic	0.01
...	

life	0.02
evolve	0.01
organism	0.01
...	

brain	0.04
neuron	0.02
nerve	0.01
...	

data	0.02
number	0.02
computer	0.01
...	

Documents

Seeking Life's Bare (Genetic) Necessities

COLD SPRING HARBOR, NEW YORK— How many genes does an organism need to survive? Last week at the genome meeting here,* two genome researchers with radically different approaches presented complementary views of the basic genes needed for life. One research team, using computer analyses to compare known genomes, concluded that today's organisms can be sustained with just 250 genes, and that the earliest life forms required a mere 128 genes. The other researcher mapped genes in a simple parasite and estimated that for this organism, 800 genes are plenty to do the job—but that anything short of 100 wouldn't be enough.

Although the numbers don't match precisely, those predictions

"are not all that far apart," especially in comparison to the 75,000 genes in the human genome, notes Siv Andersson of Uppsala University in Sweden, who arrived at the 800 number. But coming up with a consensus answer may be more than just a genetic numbers game, particularly as more and more genomes are completely mapped and sequenced. "It may be a way of organizing any newly sequenced genome," explains Arcady Mushegian, a computational molecular biologist at the National Center for Biotechnology Information (NCBI) in Bethesda, Maryland. Comparing an

* Genome Mapping and Sequencing, Cold Spring Harbor, New York, May 8 to 12.

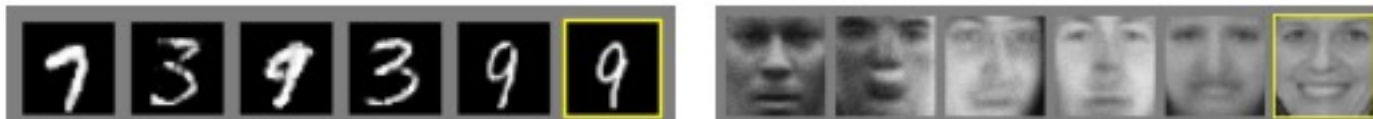
SCIENCE • VOL. 272 • 24 MAY 1996

Topic proportions and assignments

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

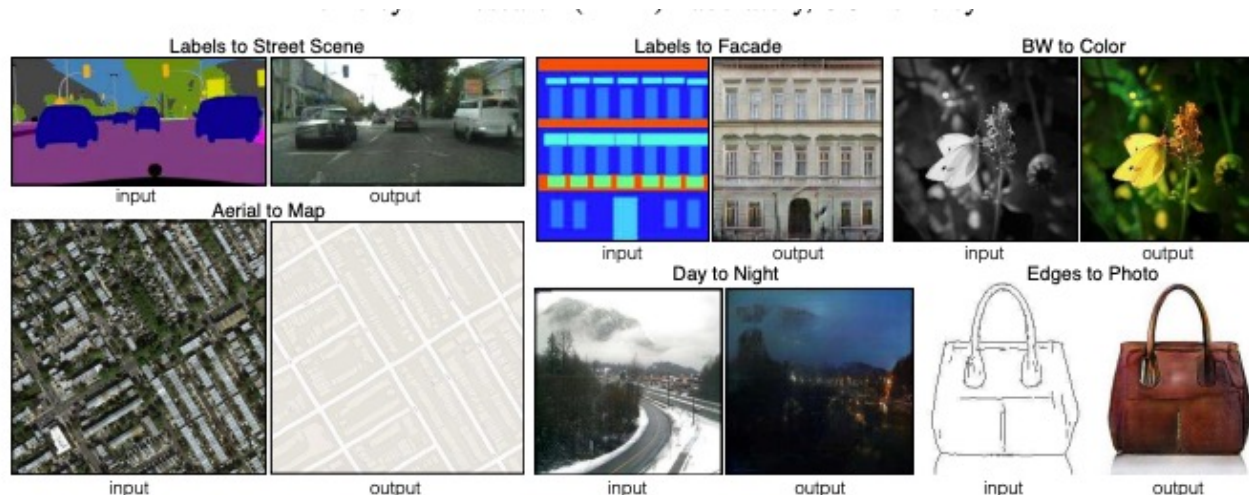
ML: Generative Adversarial Networks (GAN)

► Generative Adversarial Networks (GAN)



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

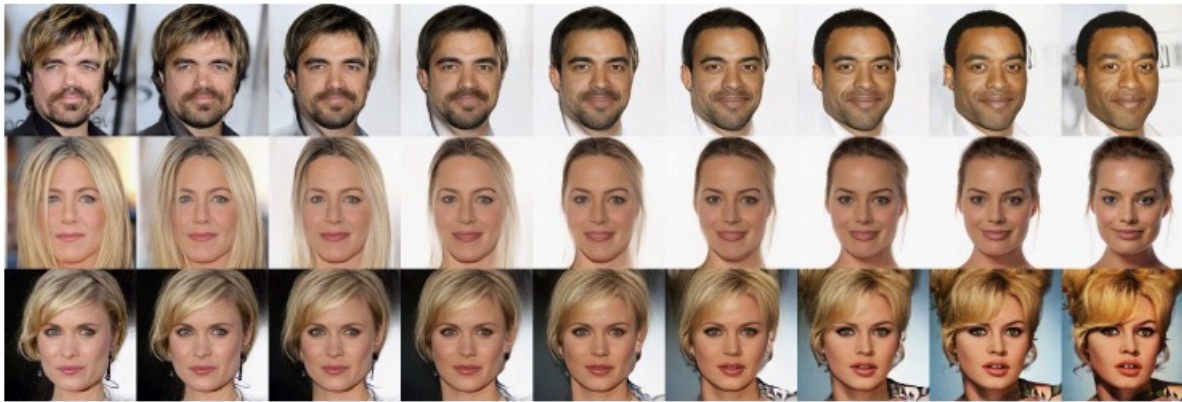
► Image to image translation via GANs



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

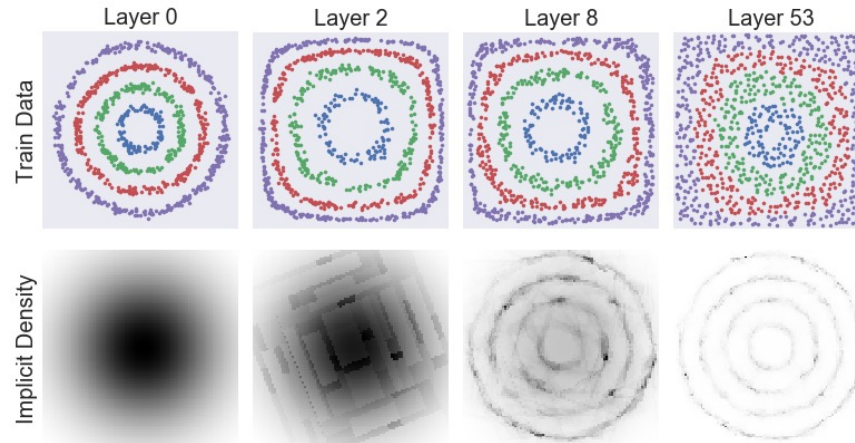
ML: Invertible networks

► Invertible Flows



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

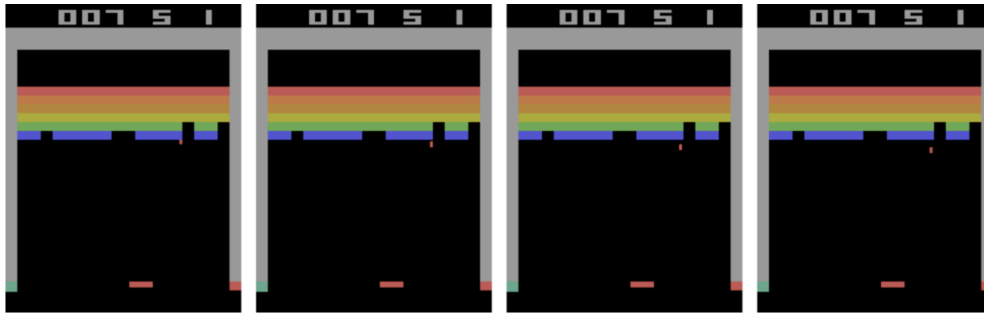
► Deep Density Destructors



<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?”

ML: Domain Generalization

► Distribution shifts in the real-world

	Domain generalization					Subpopulation shift	Domain generalization + subpopulation shift			
Dataset	iWildCam	Camelyon17	RxRx1	OGB-MolPCBA	GlobalWheat	CivilComments	FMoW	PovertyMap	Amazon	Py150
Input (x)	camera trap photo	tissue slide	cell image	molecular graph	wheat image	online comment	satellite image	satellite image	product review	code
Prediction (y)	animal species	tumor	perturbed gene	bioassays	wheat head bbox	toxicity	land use	asset wealth	sentiment	autocomplete
Domain (d)	camera	hospital	batch	scaffold	location, time	demographic	time, region	country, rural-urban	user	git repository
# domains	323	5	51	120,084	47	16	16 x 5	23 x 2	2,586	8,421
# examples	203,029	455,954	125,510	437,929	6,515	448,000	523,846	19,669	539,502	150,000
Train example						What do Black and LGBT people have to do with bicycle licensing?			Overall a solid package that has a good quality of construction for the price.	<pre>import numpy as np ... norm=np.____</pre>
Test example						As a Christian, I will not be patronizing any of those businesses.			I "loved" my French press, it's so perfect and came with all this fun stuff!	<pre>import subprocess as sp p=sp.Popen() stdout=p.____</pre>
Adapted from	Beery et al. 2020	Bandi et al. 2018	Taylor et al. 2019	Hu et al. 2020	David et al. 2021	Borkan et al. 2019	Christie et al. 2018	Yeh et al. 2020	Ni et al. 2019	Raychev et al. 2016

► How can you train models so that they work in new unseen test domains?

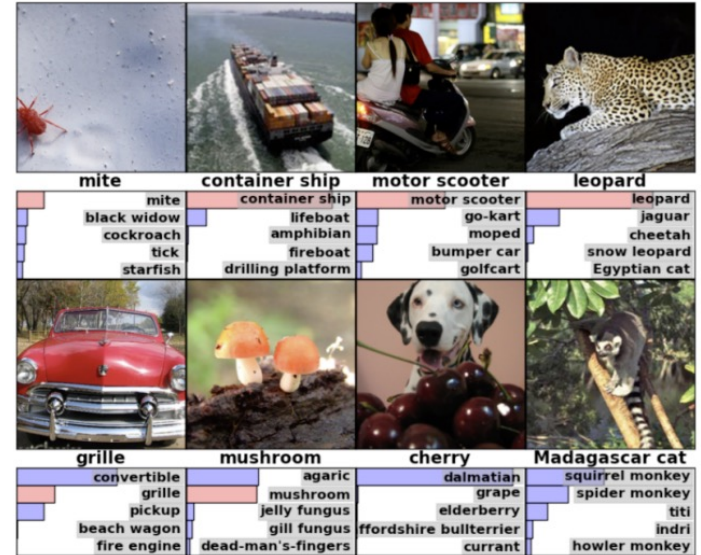
Koh, P. W., Sagawa, S., Xie, S. M., Zhang, M., Balsubramani, A., Hu, W., ... & Liang, P. (2021, July). Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning* (pp. 5637-5664). PMLR. <https://wilds.stanford.edu/datasets/>

3. Computer Vision (Based on CVPR sessions)

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

CV: Classic Tasks

► Recognition



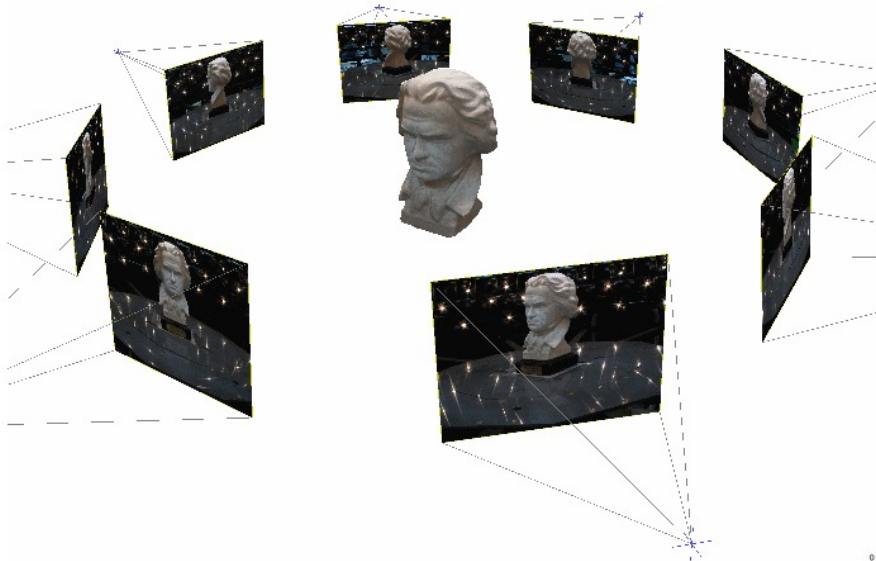
<http://www.cs.toronto.edu/~fritz/absps/imagenet.pdf>

► Segmentation

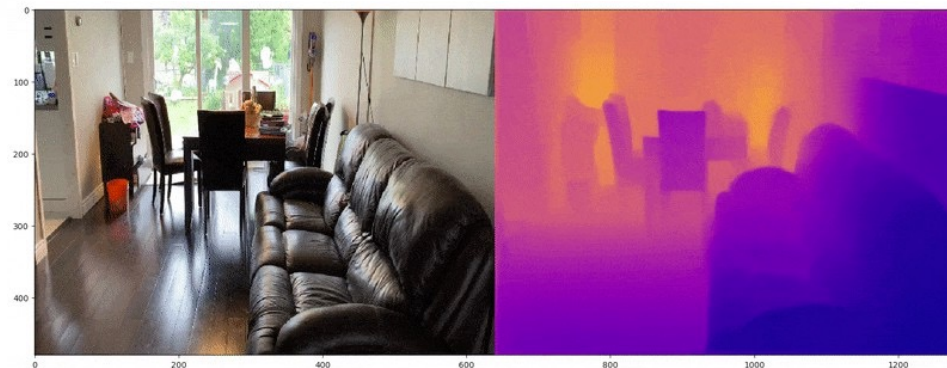


<http://vladlen.info/publications/feature-space-optimization-for-semantic-video-segmentation/>

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)

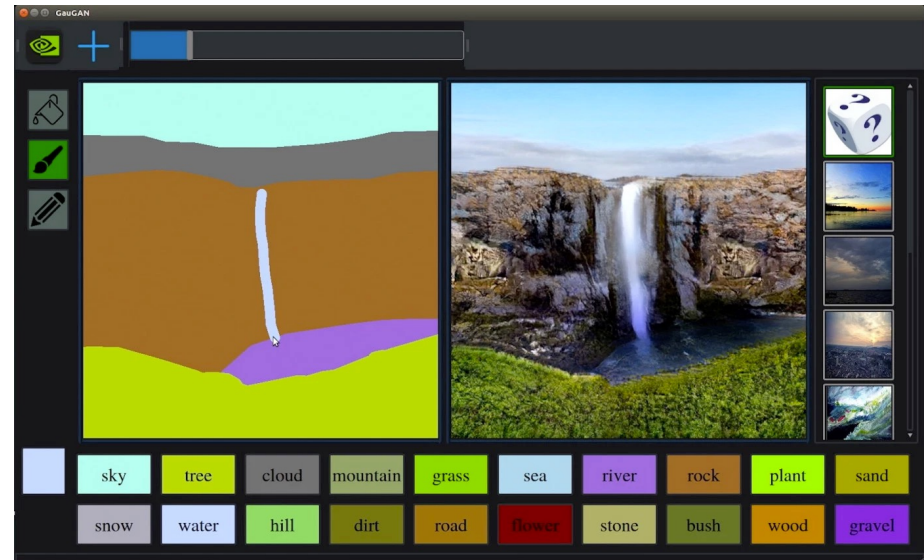
► Style transfer

https://www.cv-foundation.org/openaccess/content_cvpr_2016/paper_s/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

Lexical Ambiguity

The presence of two or more possible meanings within a single word.



"I saw her duck."

Syntactic Ambiguity

The presence of two or more possible meanings within a single sentence or sequence of words.

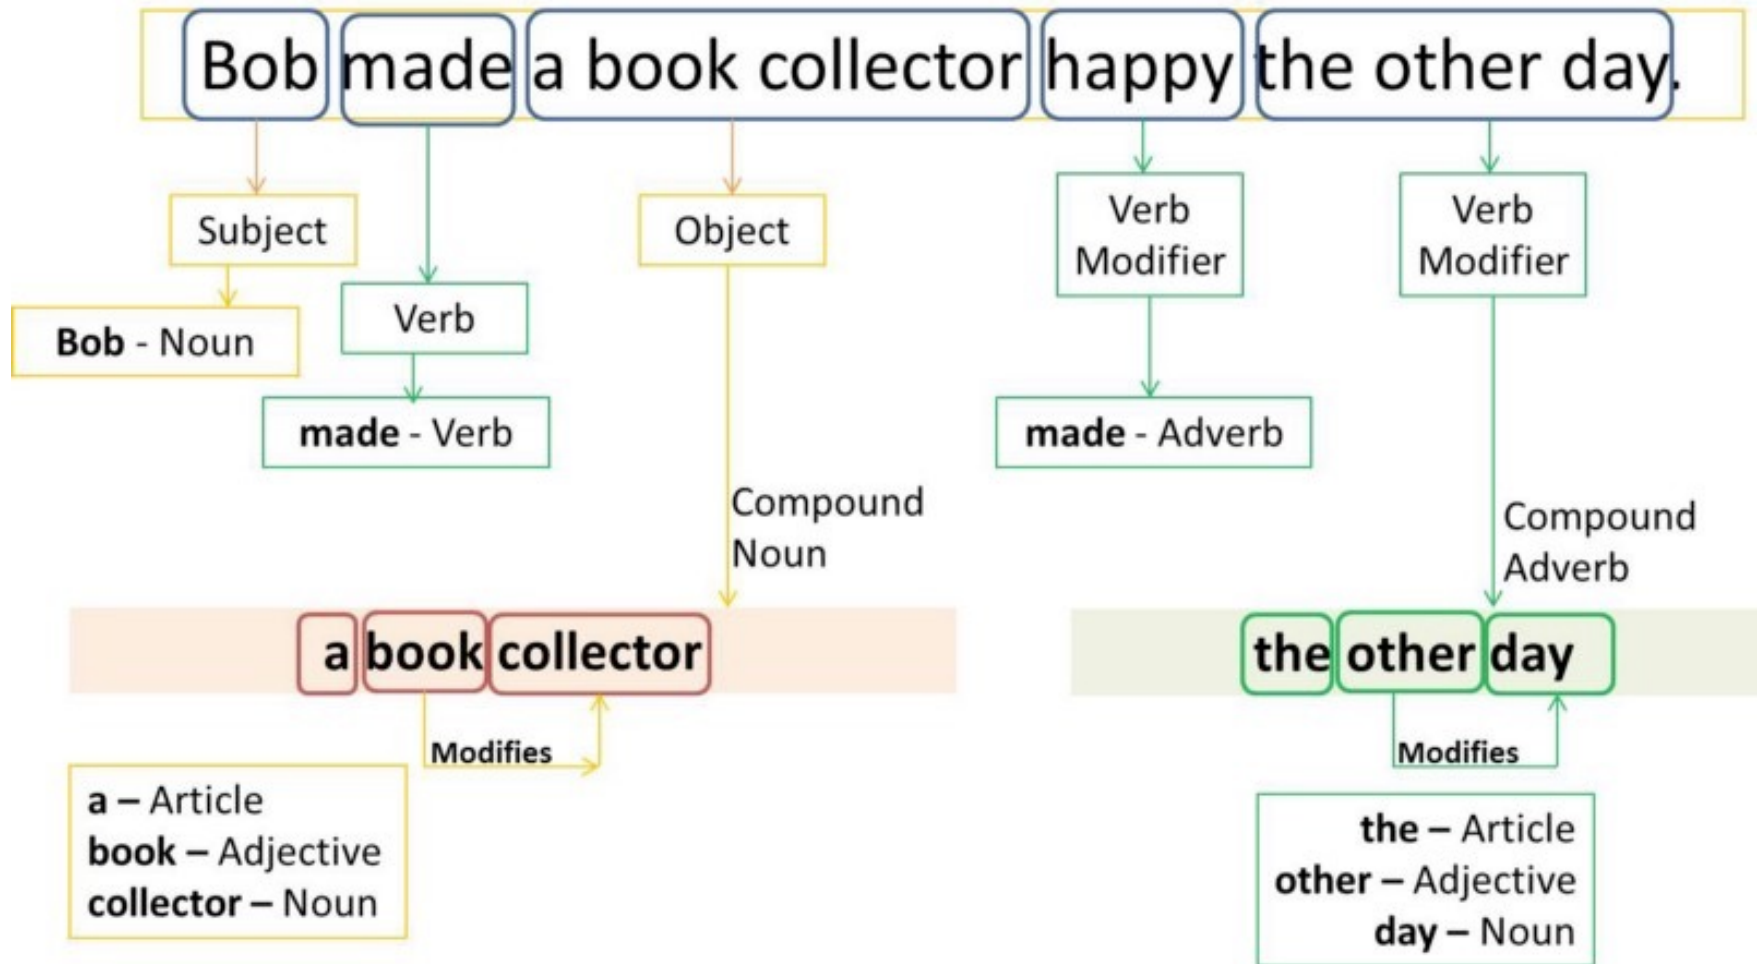


"The chicken is ready to eat."

ThoughtCo.

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

NLP: Tagging and Parsing



NLP: Information Extraction and Text Mining

Text in

Brazil ranks number 5 in the list of countries by population.

The term "**Ibu Negara**" (Lady/Mother of the State) is used for **wife** of the President of **Indonesia**.

Game of Thrones is an adaptation of A Song of Ice and Fire, George R. R. Martin's series of fantasy novels. It ranks **fourth** among the **IMDB Top Rated TV Shows**

Data out

THE COUNTRIES WITH THE LARGEST POPULATION

China	1	1,388,232,693
India	2	1,342,512,706
Unites States	3	326,474,013
Indonesia	4	263,510,146
Brasil	5	174,315,386

THE COUNTRY'S' FIRST LADIES

Brigitte Macron

- Spouse: Emmanuel Macron, President of France (2017 -)

Melania Trump

- Spouse: Donald J. Trump, U.S. President (2017-)

Iriana Widodo

- Spouse: Joko Widodo, **President of Indonesia** (2014 -)

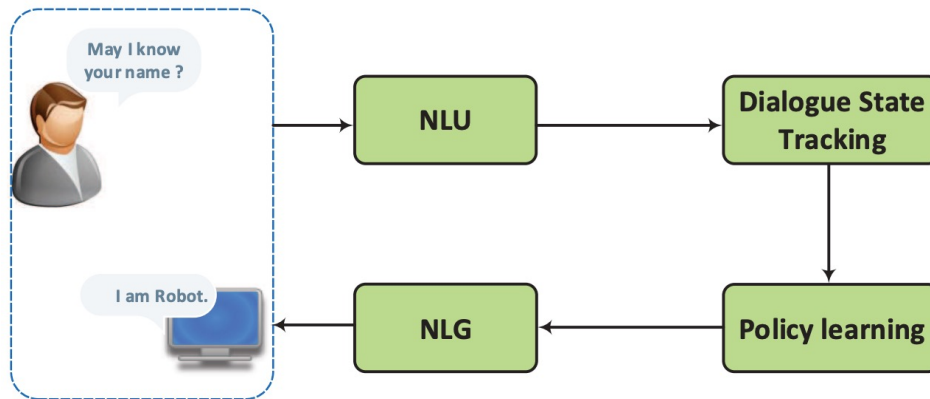
- Also known as: "**Ibu Negara**" (Lady/Mother of the State)

IMDB TOP RATED TV SHOWS

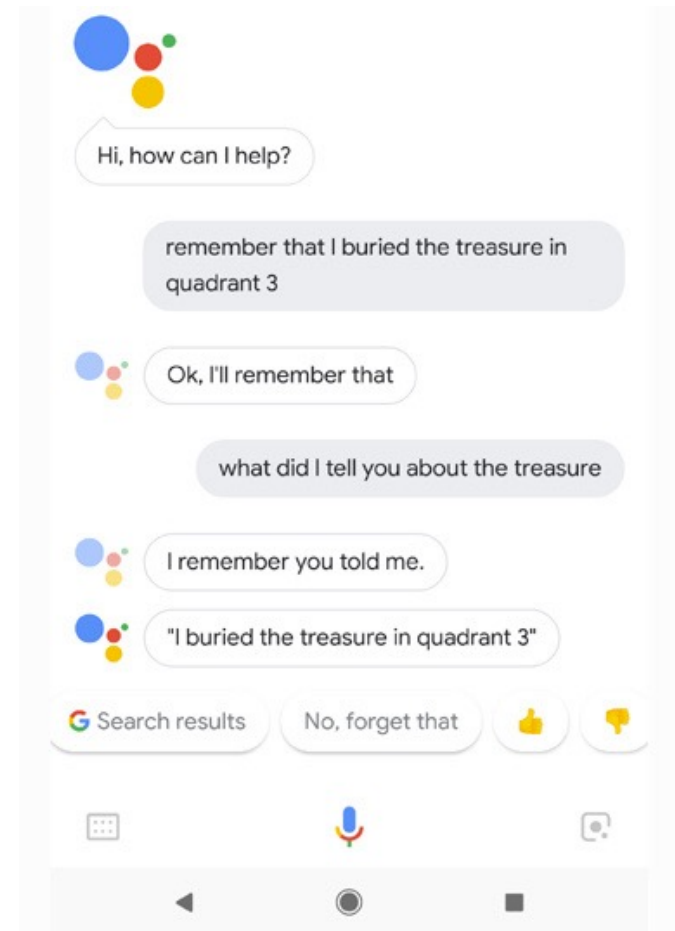
- 1 Planet Earth II (2016) 9.6.
- 2 Band of Brothers (2001) 9.5.
- 3 Planet Earth (2006) 9.5.
- 4 **Game of Thrones** (2011) 9.4.
- 5 Breaking Bad (2008) 9.4.

<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