CS839 Special Topics in Deep Learning
Course Overview

Sharon Yixuan Li
University of Wisconsin-Madison

September 3, 2020

Part I: Logistics

Instructor

- Prof. Sharon Li

  Email: sharonli@cs.wisc.edu

  Office: 5393 Computer Sciences

  Virtual office hours: TBD

Use piazza for questions:
piazza.com/wisc/fall2020/cs839/home

For emails, please include [**CS839**] in the subject title!

Teaching Assistant

- Yiyou Sun

  Email: sunyiyou@cs.wisc.edu

  Virtual office hours: Tuesday 3-4pm (BB Collab)

  Piazza: piazza.com/wisc/fall2020/cs839/home

Course Enrollment

Course capacity: ~**40** students due to

- Limited computing resources & first offering

Waiting list has >60 students

- Enroll on a **first come first serve** if registered students drop the course.
- This class will be offered again next fall!

This course will allow you to:

- **Advance** your knowledge in deep learning
  - In-depth read papers on cutting-edge topics of AI and deep learning
- **Project**
  - Explore new research directions and applications of deep learning
  - Ability to start original research in a collaborative team
- **Practice**
  - Write code in Python / Jupyter
  - Solve real problems

Course Schedule

- Time: **Tuesday** and **Thursday** 4:00-5:15pm CT
- Location: **BlackBoard Collaborate** for Fall 2020
- Schedule is available on the course website:

  http://pages.cs.wisc.edu/~sharonli/courses/cs839_fall2020
- Slides online on website

Prerequisites

- This course assumes that you **already** have a basic understanding of deep learning.
- Prerequisites
  - CS760: *Machine Learning*
  - *(preferred)* CS761: *Mathematical Foundations of Machine Learning*
- Familiarity with linear algebra, statistics, optimization is expected.

Textbooks

- *Deep Learning*. I. Goodfellow, Y. Bengio, and A. Courville.

  https://www.deeplearningbook.org/front_matter.pdf

- Dive into Deep Learning. Aston Zhang and Zachary C. Lipton and Mu Li and Alexander J. Smola.

- *Pattern Recognition and Machine Learning*. C. Bishop. Springer, 2011.

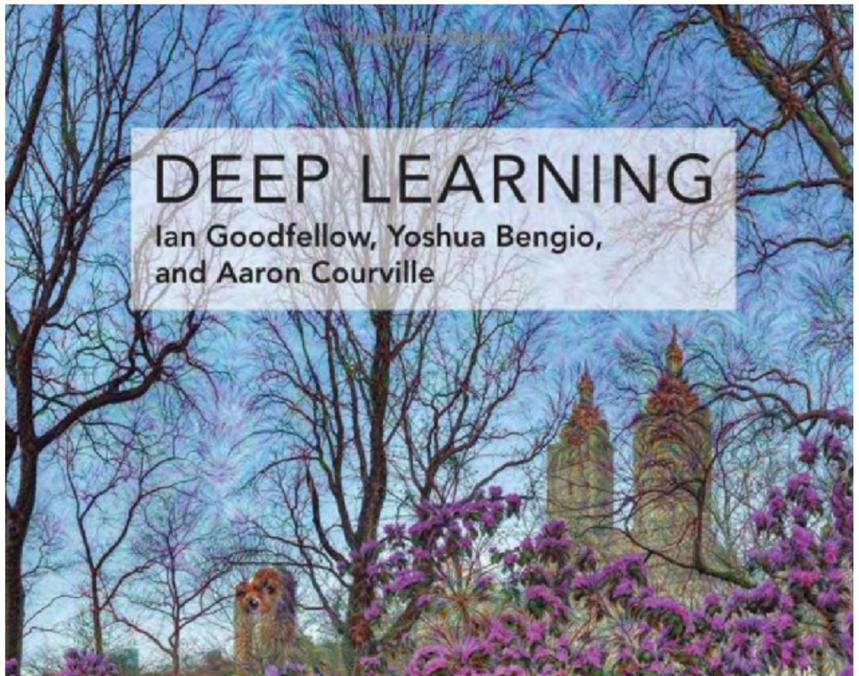

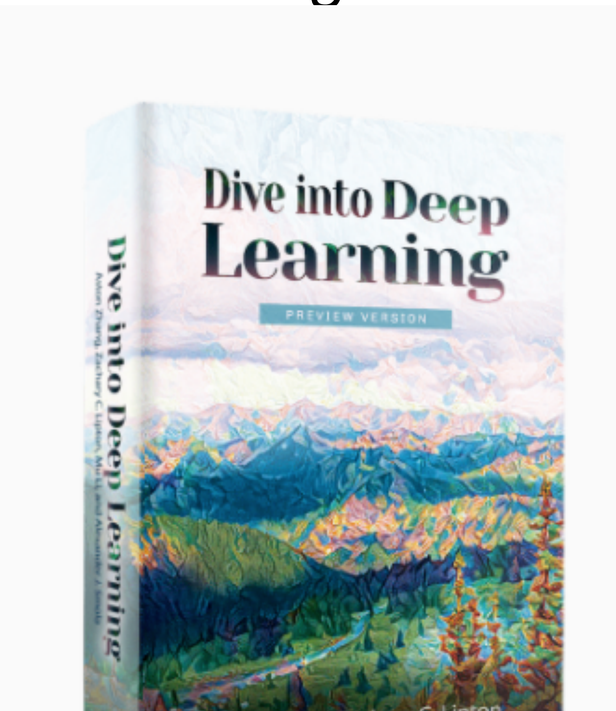

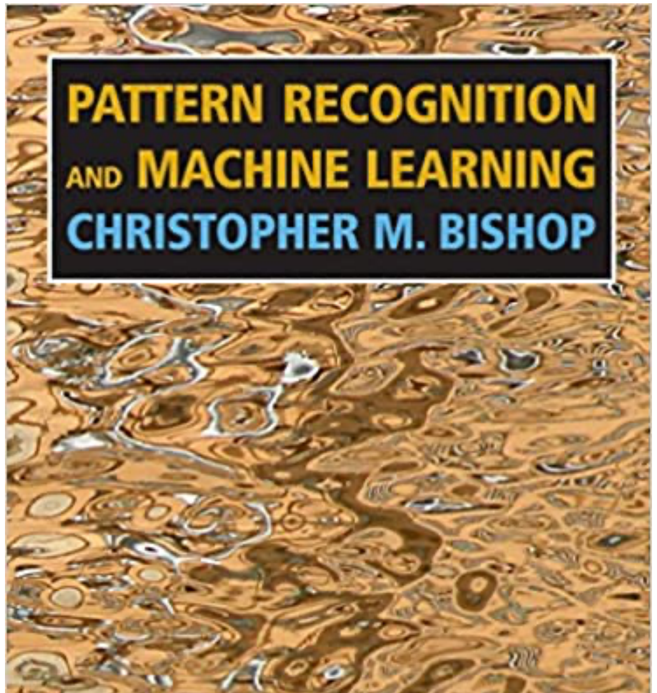

Course readings

- Most readings will be recent papers, articles and book chapters
- Available on course website (will be updated from time to time)

Week	Date	Topic	Reading materials
1	Thursday, September 3	Course overview and introduction [Slides]	
2	Tuesday, September 8	Neural Architecture Design I (lecture) [Slides]	Goodfellow-Bengio-Courville Chapter 6
2	Thursday, September 10	Neural Architecture Design II (paper presentation)	Deep Residual Learning for Image Recognition Densely Connected Convolutional Networks (to present) Mask R-CNN
3	Tuesday, September 15	Neural Architecture Design III (paper presentation)	Neural Architecture Search with Reinforcement Learning Efficient Neural Architecture Search via Parameter Sharing
3	Thursday, September 17	Trustworthy Deep Learning I (lecture)	Goodfellow-Bengio-Courville Chapter 7.5, 7.13
4	Tuesday, September 22	Trustworthy Deep Learning II (paper presentation)	Out-of-distribution Reliability Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images On Calibration of Modern Neural Networks Enhancing The Reliability of Out-of-distribution Image Detection in Neural Networks A Simple Unified Framework for Detecting Out-of-Distribution Samples and Adversarial Attacks Robust Out-of-distribution Detection via Informative Outlier Mining (to present)

Grading scheme

- In-class quizzes: **10**% (you can skip up to 2 of them)
- Paper presentation: **20**%
- Project proposal: **10**% (2 pages, due end of September)
- Final project presentation: **15**%
- Final project report (written): **45**%
- **No final exam**

Paper presentation (20%)

- Sign up **today:** 2-3 students each presentation

https://docs.google.com/spreadsheets/d/18hCfFDD3ahPJfed_nkk4nWtzzFnc_ynRqr10000RgX4

Course schedule: http://pages.cs.wisc.edu/~sharonli/courses/cs839_fall2020/schedule.html

* Each presentation group can have a maximum 3 students. As a team, every student in the same group will receive the same score for the presentation.
* It's up to the presentation team to arrange the presentation & scribes in a way that maximize the outcome.

						(only fill this if C & E columns are full)	
Date	**Topic**	**Presenter 1 Name**	**Presenter 1 Email**	**Scribe person name**	**Scribe person email**	**Presenter 2 Name (optional)**	**Presenter 2 Email (optional)**
September 10	Neural Architecure Design (+10% bonus in final grade)						
September 15	Neural Architecure Design (+5% bonus in final grade)						
September 22	Trustworthy Deep Learning						
September 24	Trustworthy Deep Learning						
September 29	Trustworthy Deep Learning						
October 13	Interpretable Deep Learning						
October 15	Interpretable Deep Learning						
October 22	Deep Learning Generalization and Theory						
October 27	Deep Learning Generalization and Theory						
November 3	Learning with less supervision						
November 5	Learning with less supervision						
November 10	Learning with less supervision						
November 17	Lifelong learning						
November 19	Lifelong learning						
December 1	Deep generative modeling						
December 3	Deep generative modeling						

Paper presentation (20%)

- Sign up **today:** 2-3 students each presentation
    - 1-2 persons will present and lead the discussion
    - Interactive discussion (everyone should do the reading ahead of class)
    - One person will take notes and synthesize the discussion
    - Compile three quiz questions for in-class testing (send to TA, who will upload to Canvas)
- First presentation (September 10) gets extra **10%** in final grade.

  Densely Connected Convolutional Networks by Gao et al. 2017
- A great guide by Prof. Kayvon Fatahalian on giving clear talks:

  https://graphics.stanford.edu/~kayvonf/misc/cleartalktips.pdf
- Deadlines:
    - **Day before presentation**: email TA the slides + quiz questions by 6pm
    - **Day following the presentation**: email TA the notes by 6pm (10% per-day late penalty)

During class

- Start with quiz questions on Canvas (10-15mins)
    - You may skip up to 2 quizzes throughout the semester
- Presenter(s):
    - Time the presentation to last 1 hour, including QA
- All:
    - Ask questions during the presentation

Presentation rubric

- **Technical**:
    - Depth of content
    - Accuracy of content
    - Paper criticism
    - Discussion lead
- **Soft presentation skills**
    - Time management
    - Responsiveness to audience
    - Organization
    - Presentation aids (slides etc)

Project (70%)

- **Original work** in deep learning
  - Existing tools applied to novel problem
  - Novel algorithms/theory/tools
  - Choose research topic covered by this course.
- **Academic research process**
  - Research in a team (**2-4** students)
  - End result is a paper/report (ICML template) + academic presentation
  - Ask instructor & TA for advice if you are stuck - we are here to help

Project (70%)

- 9/17 **Register team** (names, working title)
- 9/29 **Project proposal** (2 pages, excluding references)
- 10/1 or 10/6 **Talk to instructor** to discuss (5-min talk with 10min discussion)
- 12/8-12/17 **Final presentation & report**

Start **early** (last minute projects often fail)**!**

Integrity

Any instance of sharing or plagiarism, copying, cheating, or other disallowed behavior will constitute a breach of ethics. Students are responsible for reporting any violation of these rules by other students, and failure to constitutes an ethical violation that carries with it similar penalties.

GPU access

instgpu-01.cs.wisc.edu
instgpu-02.cs.wisc.edu
instgpu-03.cs.wisc.edu
instgpu-04.cs.wisc.edu

- Every student enrolled will be granted access to instructional GPU servers.
- 4 servers (8 GPUs each) for ALL.
- Job submitted through SLURM to ensure fair resource usage.
- Recommend using 1 GPU at a time.
- Ask TA on Piazza for GPU related questions.
- Account will be deleted after the end of semester.

Part II: Topic Overview

Topics covered in this course

1. Neural architecture design
2. Trustworthy deep learning
3. Interpretable deep learning
4. Deep learning generalization and theory
5. Learning with less supervision
6. Lifelong learning
7. Deep generative modeling

Each topic will be covered by Lecture + Paper presentations (Overview & deep dive)

1. Evolution of neural net architectures

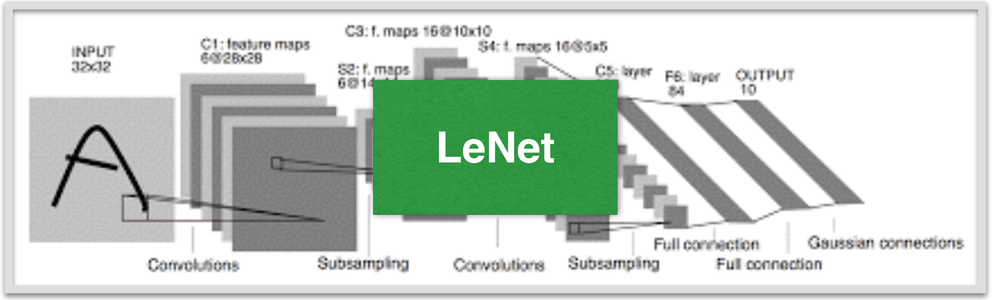

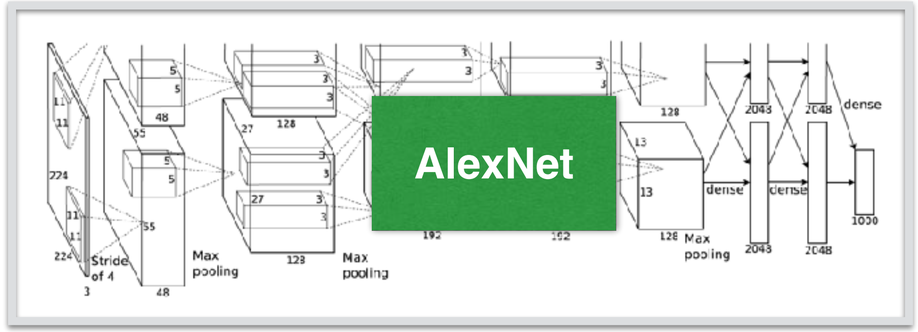

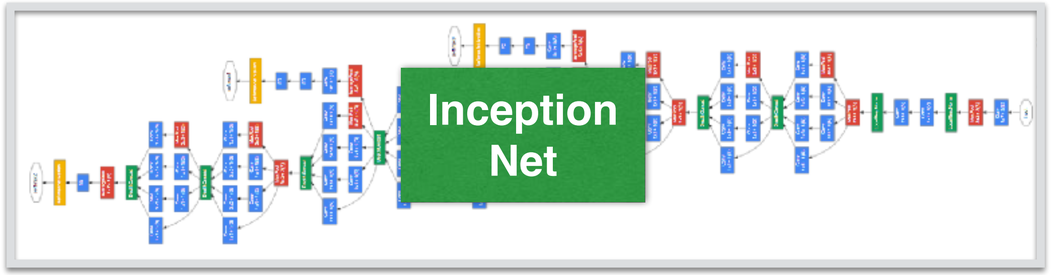

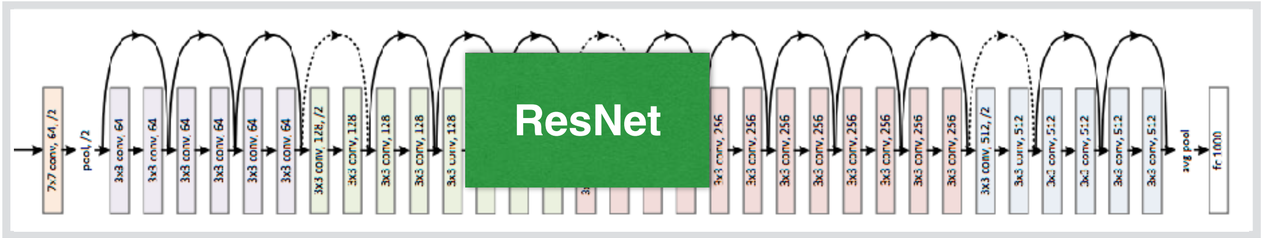

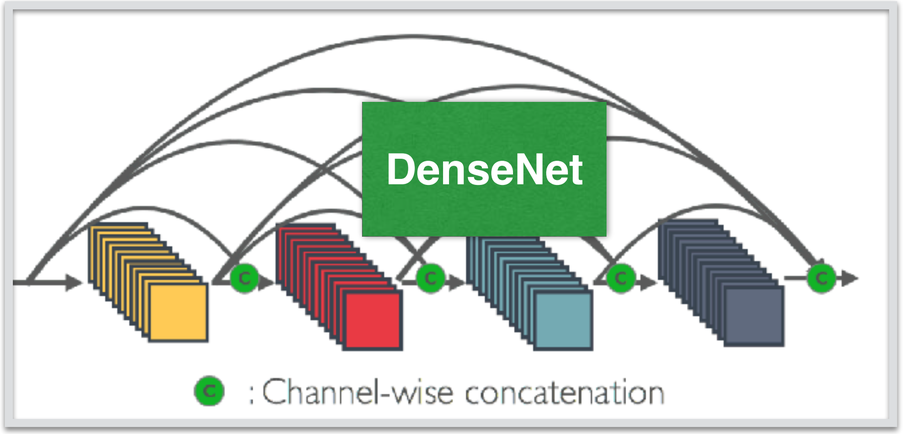

1. Evolution of neural net architectures

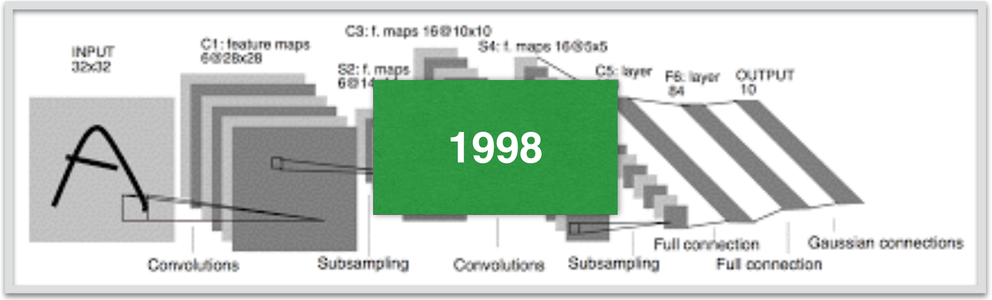

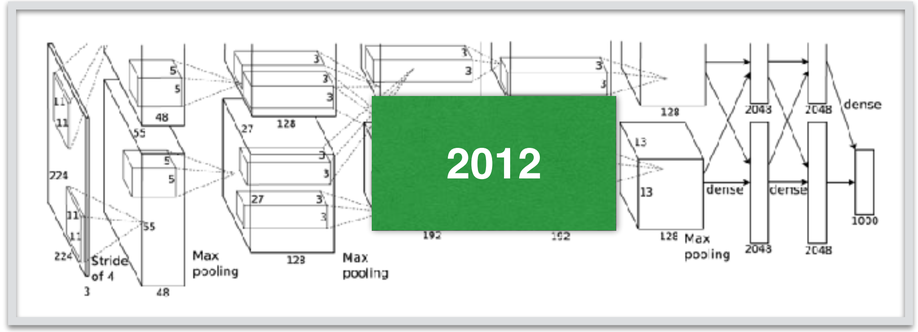

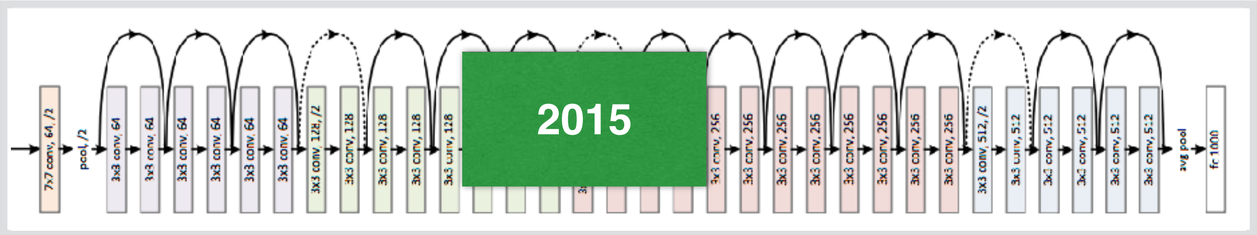

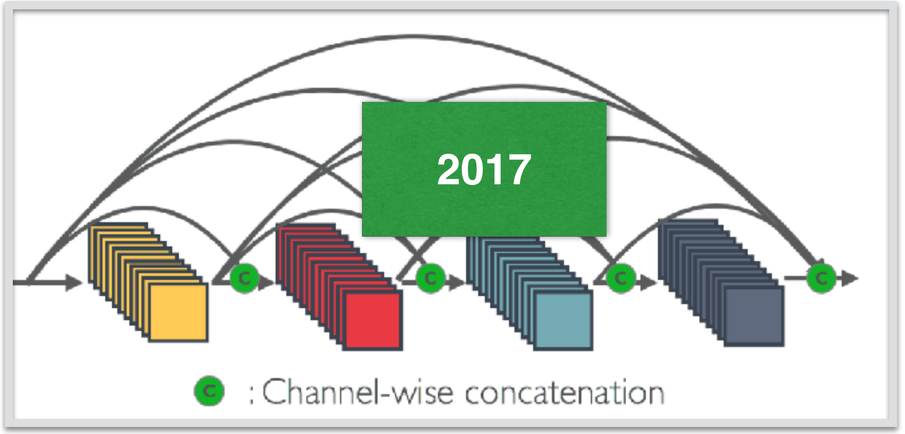

1. Evolution of neural net architectures

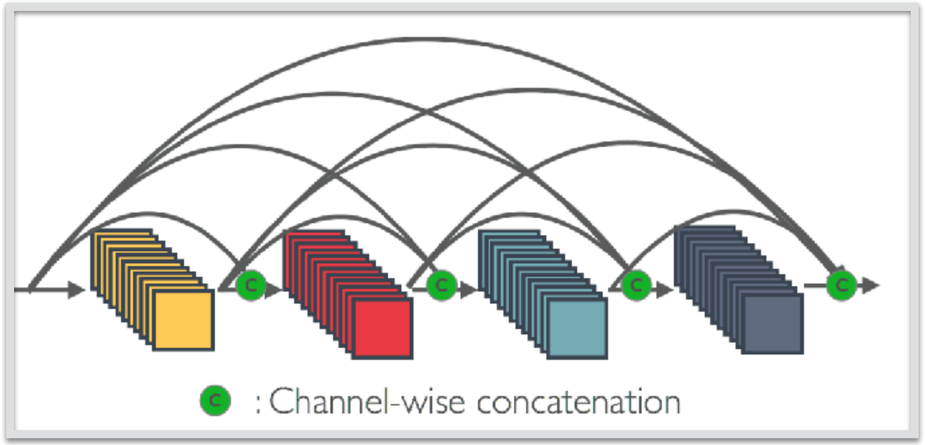

DenseNet

AutoML

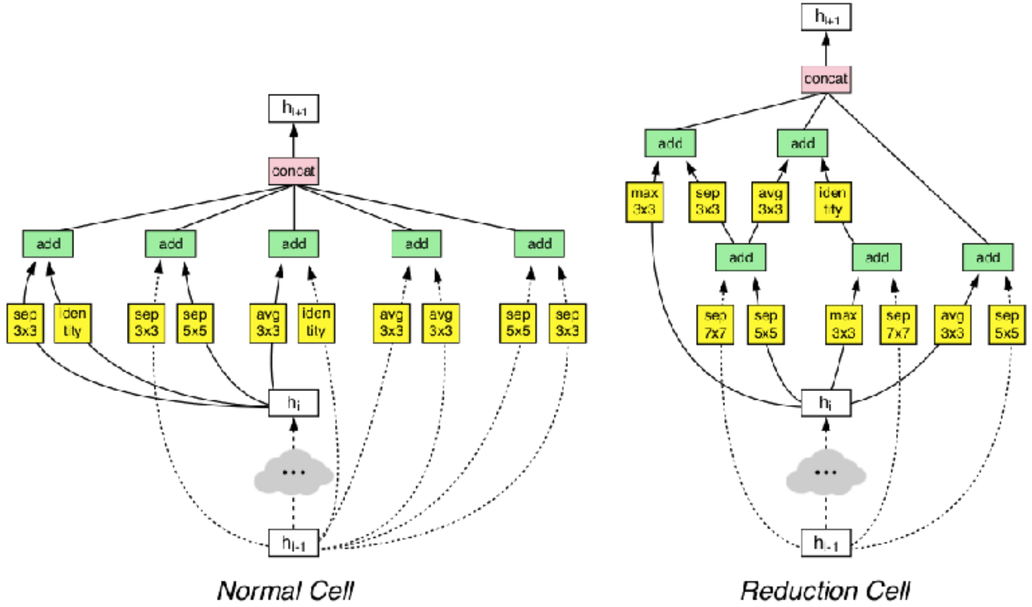

NasNet

[Zoph et al. 2017]

2. Trustworthy Deep Learning

Out-of-distribution reliability

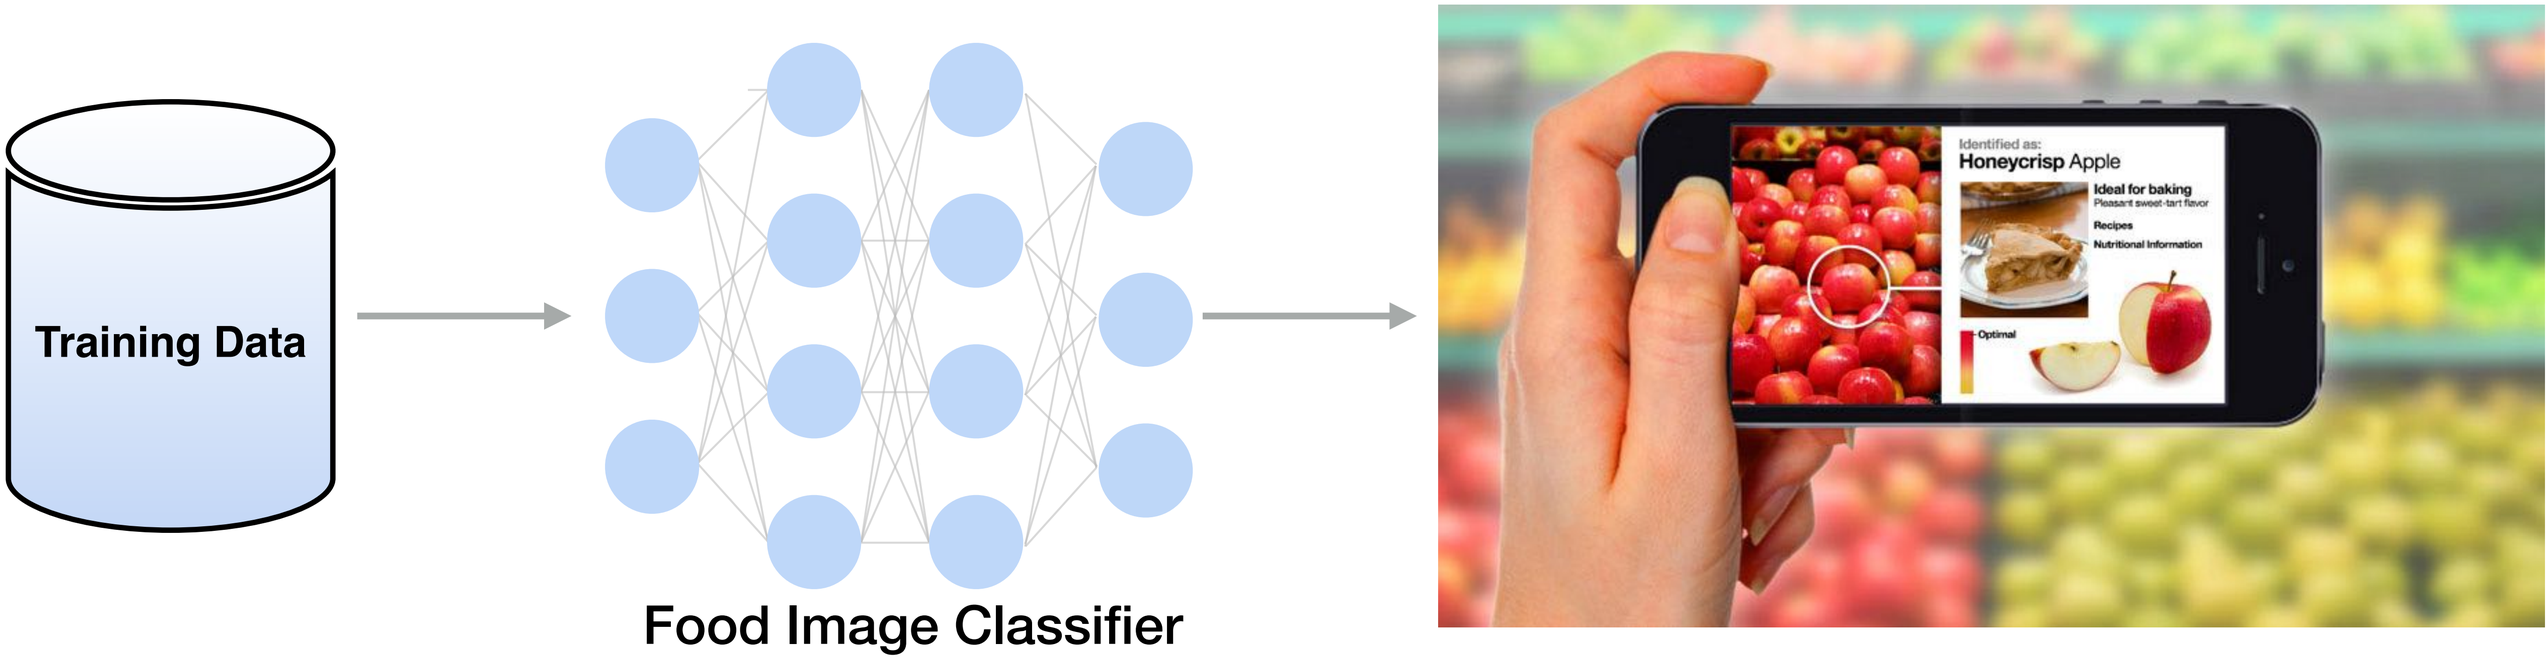

Closed-world**:** Training and testing distributions **match**

Open-world**:** Training and testing distributions **differ**

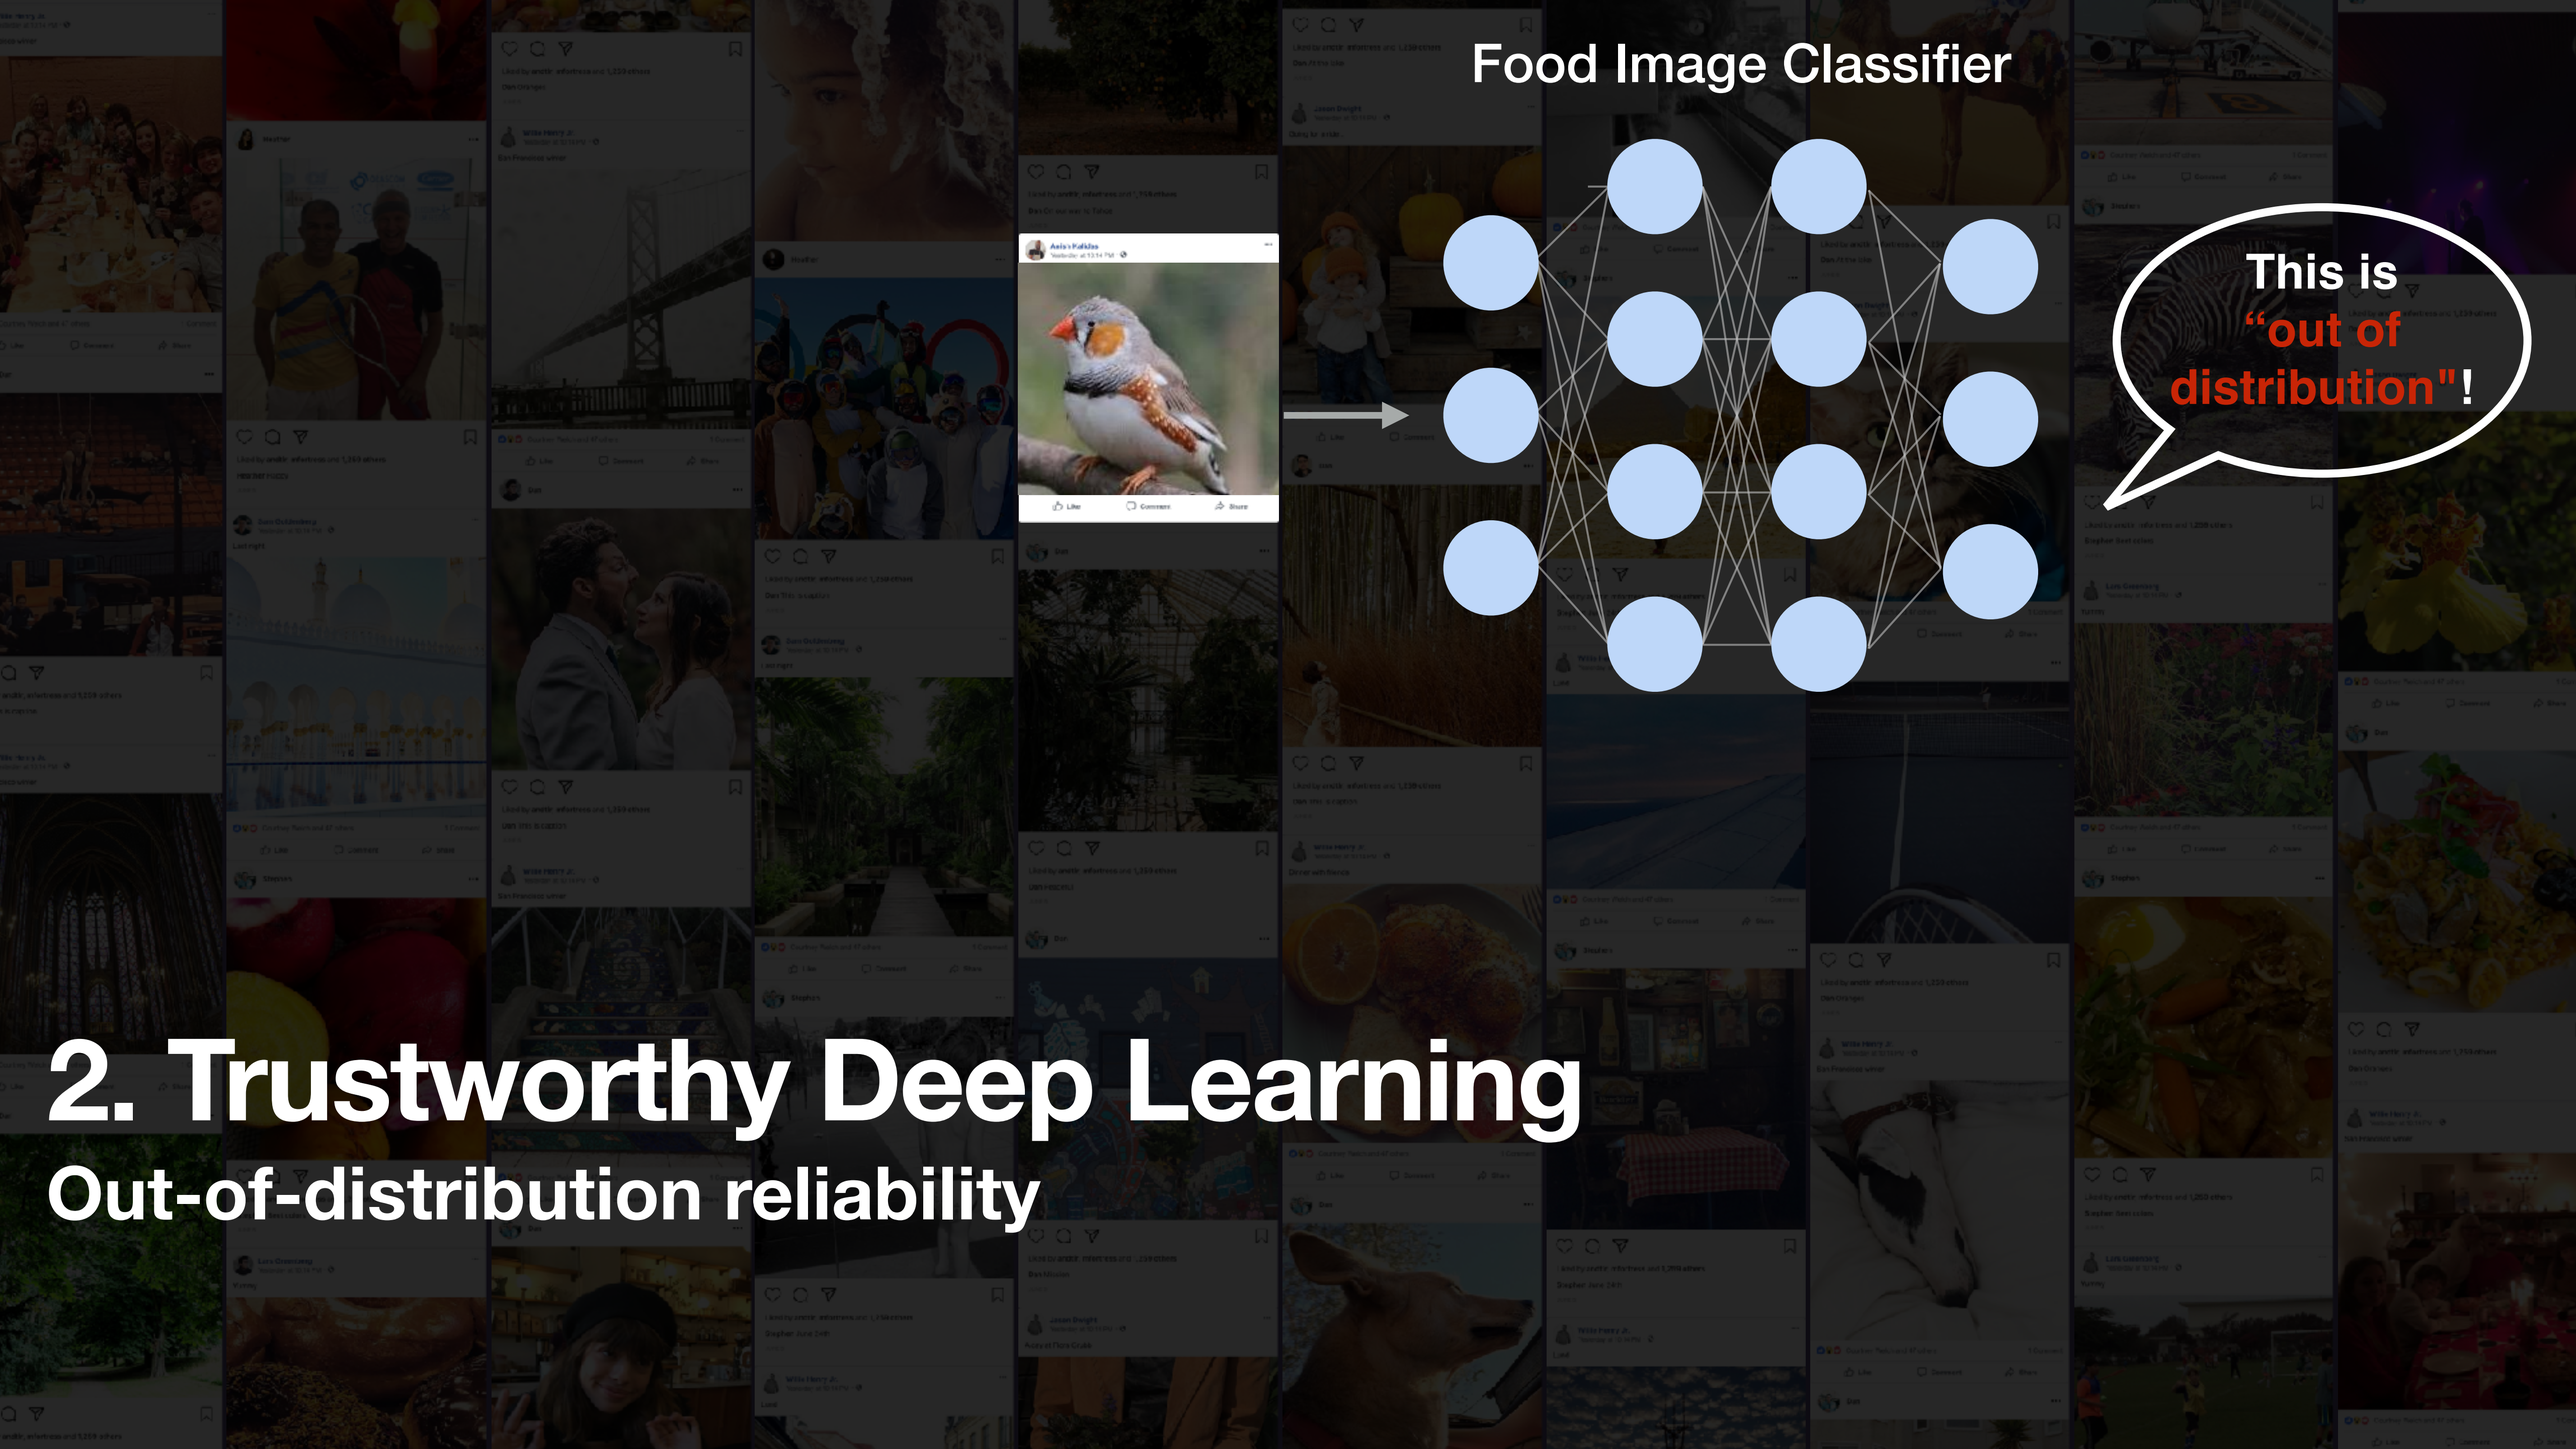

2. Trustworthy Deep Learning

Out-of-distribution reliability

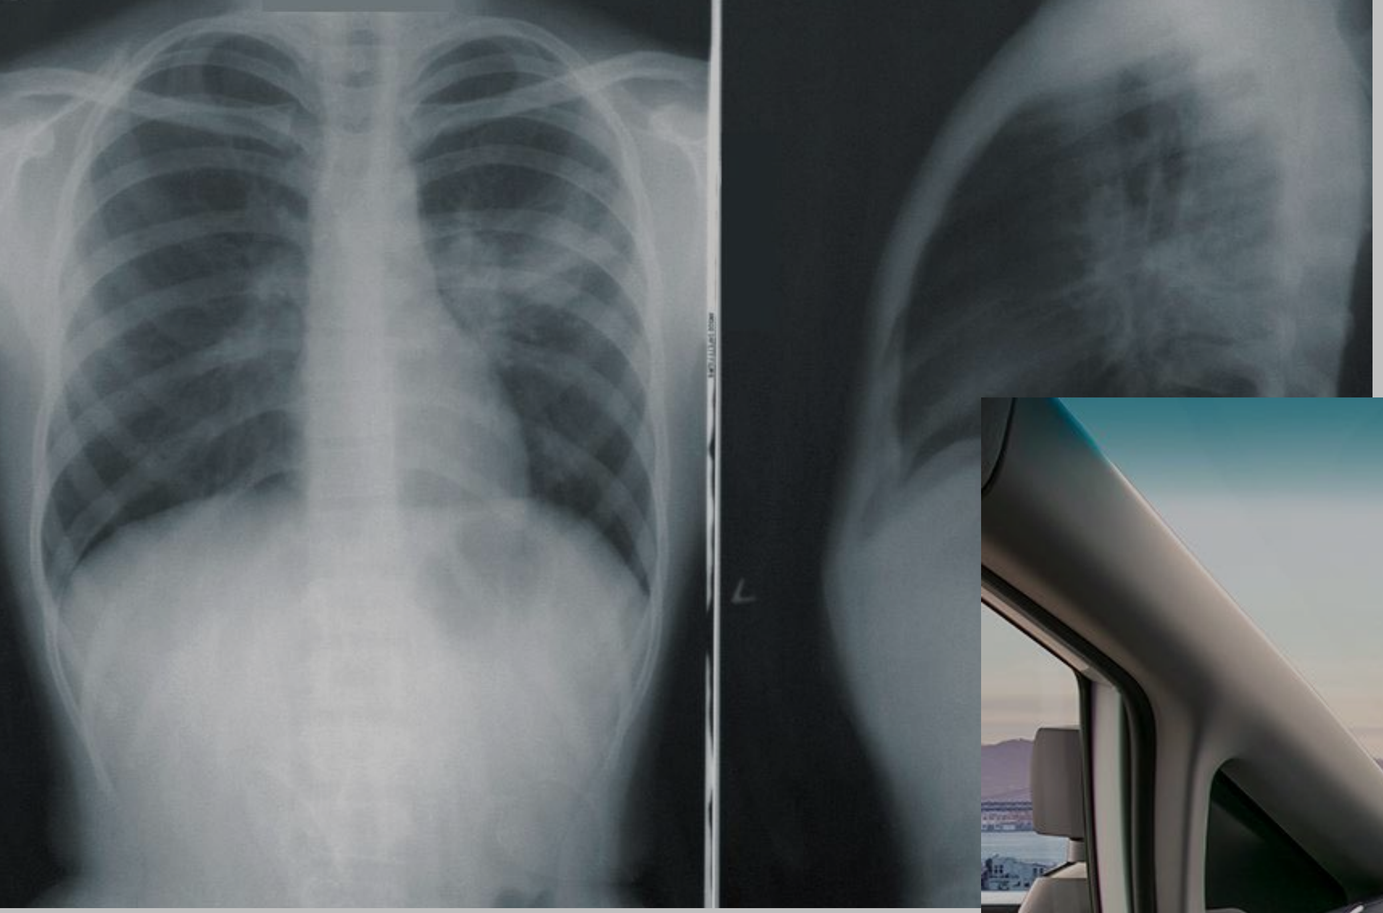

Photos from: CDC/GM

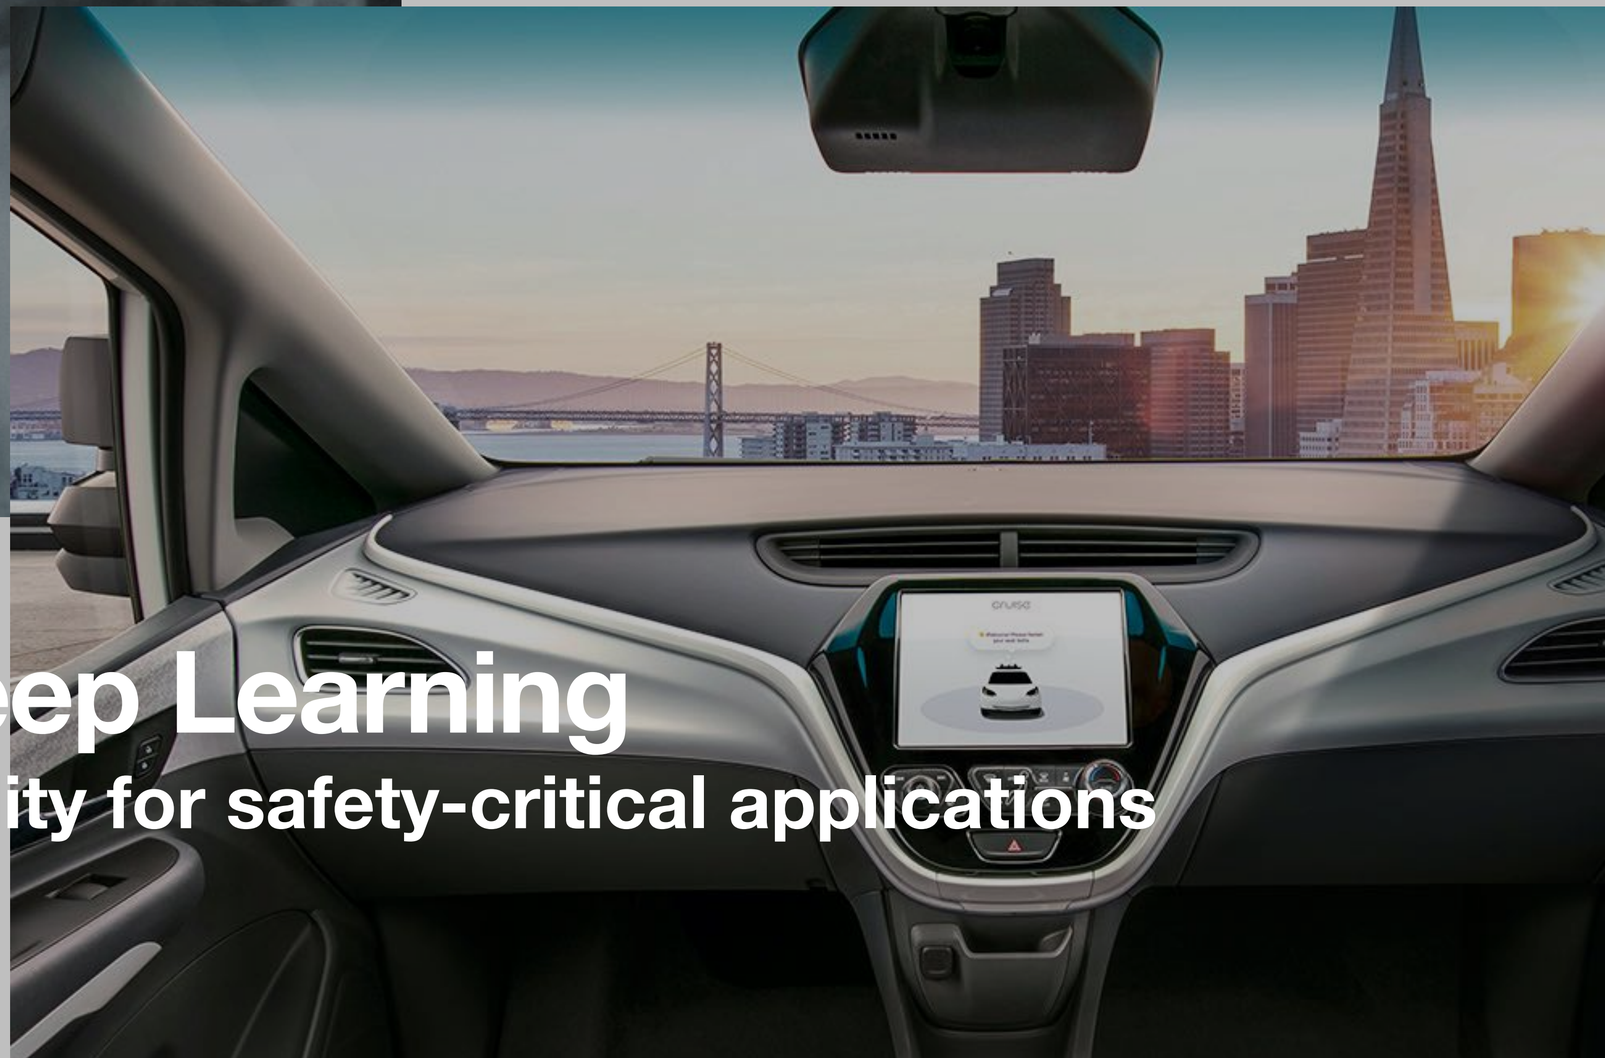

2. Trustworthy Deep Learning

Out-of-distribution reliability for safety-critical applications

2. Trustworthy Deep Learning

Adversarial Robustness

$\boldsymbol{x}$

"panda"
57.7% confidence

[Goodfellow et al. 2015]

2. Trustworthy Deep Learning

Fairness / Group Robustness

	Common training examples		Test examples
Waterbirds	y: waterbird a: water background 	y: landbird a: land background 	y: waterbird a: land background 
CelebA	y: blond hair a: female 	y: dark hair a: male 	y: blond hair a: male 
MultiNLI	y: contradiction a: has negation (P) The economy could be still better. (H) The economy has never been better.	y: entailment a: no negation (P) Read for Slate's take on Jackson's findings. (H) Slate had an opinion on Jackson's findings.	y: entailment a: has negation (P) There was silence for a moment. (H) There was a short period of time where no one spoke.

		Average Accuracy	
		ERM	DRO
Waterbirds	Train	100.0	100.0
	Test	97.3	97.4
CelebA	Train	100.0	100.0
	Test	94.8	94.7
MultiNLI	Train	99.9	99.3
	Test	82.5	82.0

		Worst-Group Accuracy	
		ERM	DRO
Waterbirds	Train	100.0	100.0
	Test	60.0	76.9
CelebA	Train	99.9	100.0
	Test	41.1	41.1
MultiNLI	Train	99.9	99.0
	Test	65.7	66.4

[Sagawa et al. 2020]

3. Interpretable Deep Learning

3. Interpretable Deep Learning

The big picture

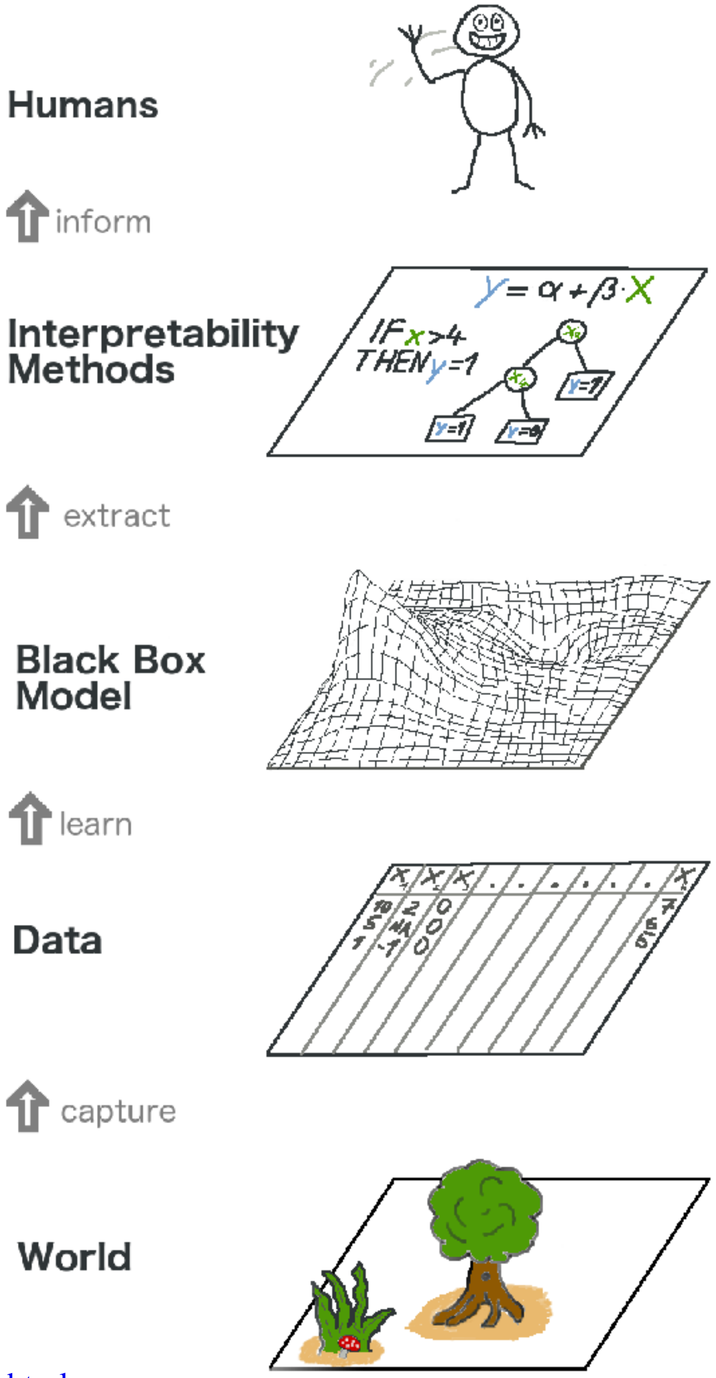

https://christophm.github.io/interpretable-ml-book/agnostic.html

3. Interpretable Deep Learning

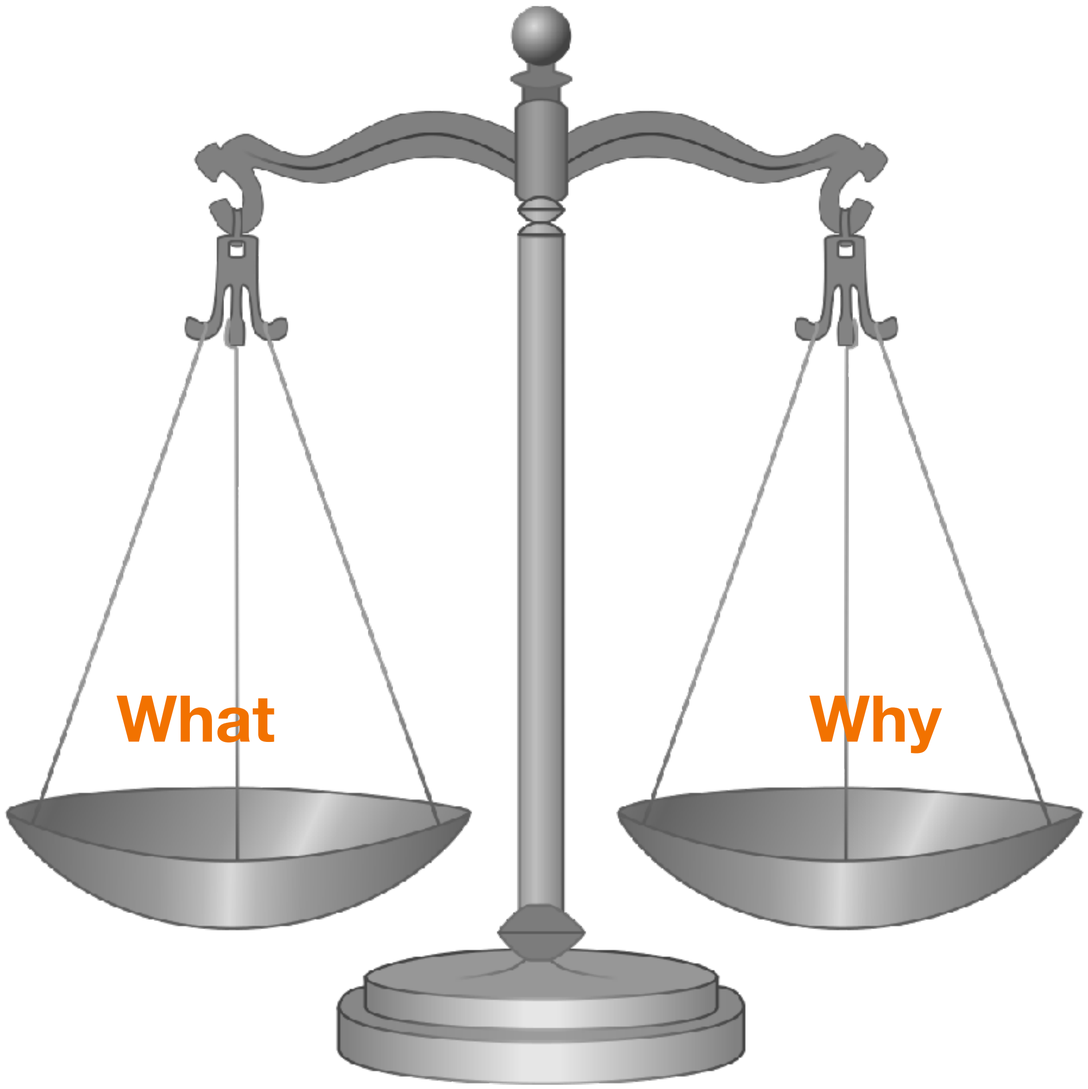

3. Interpretable Deep Learning

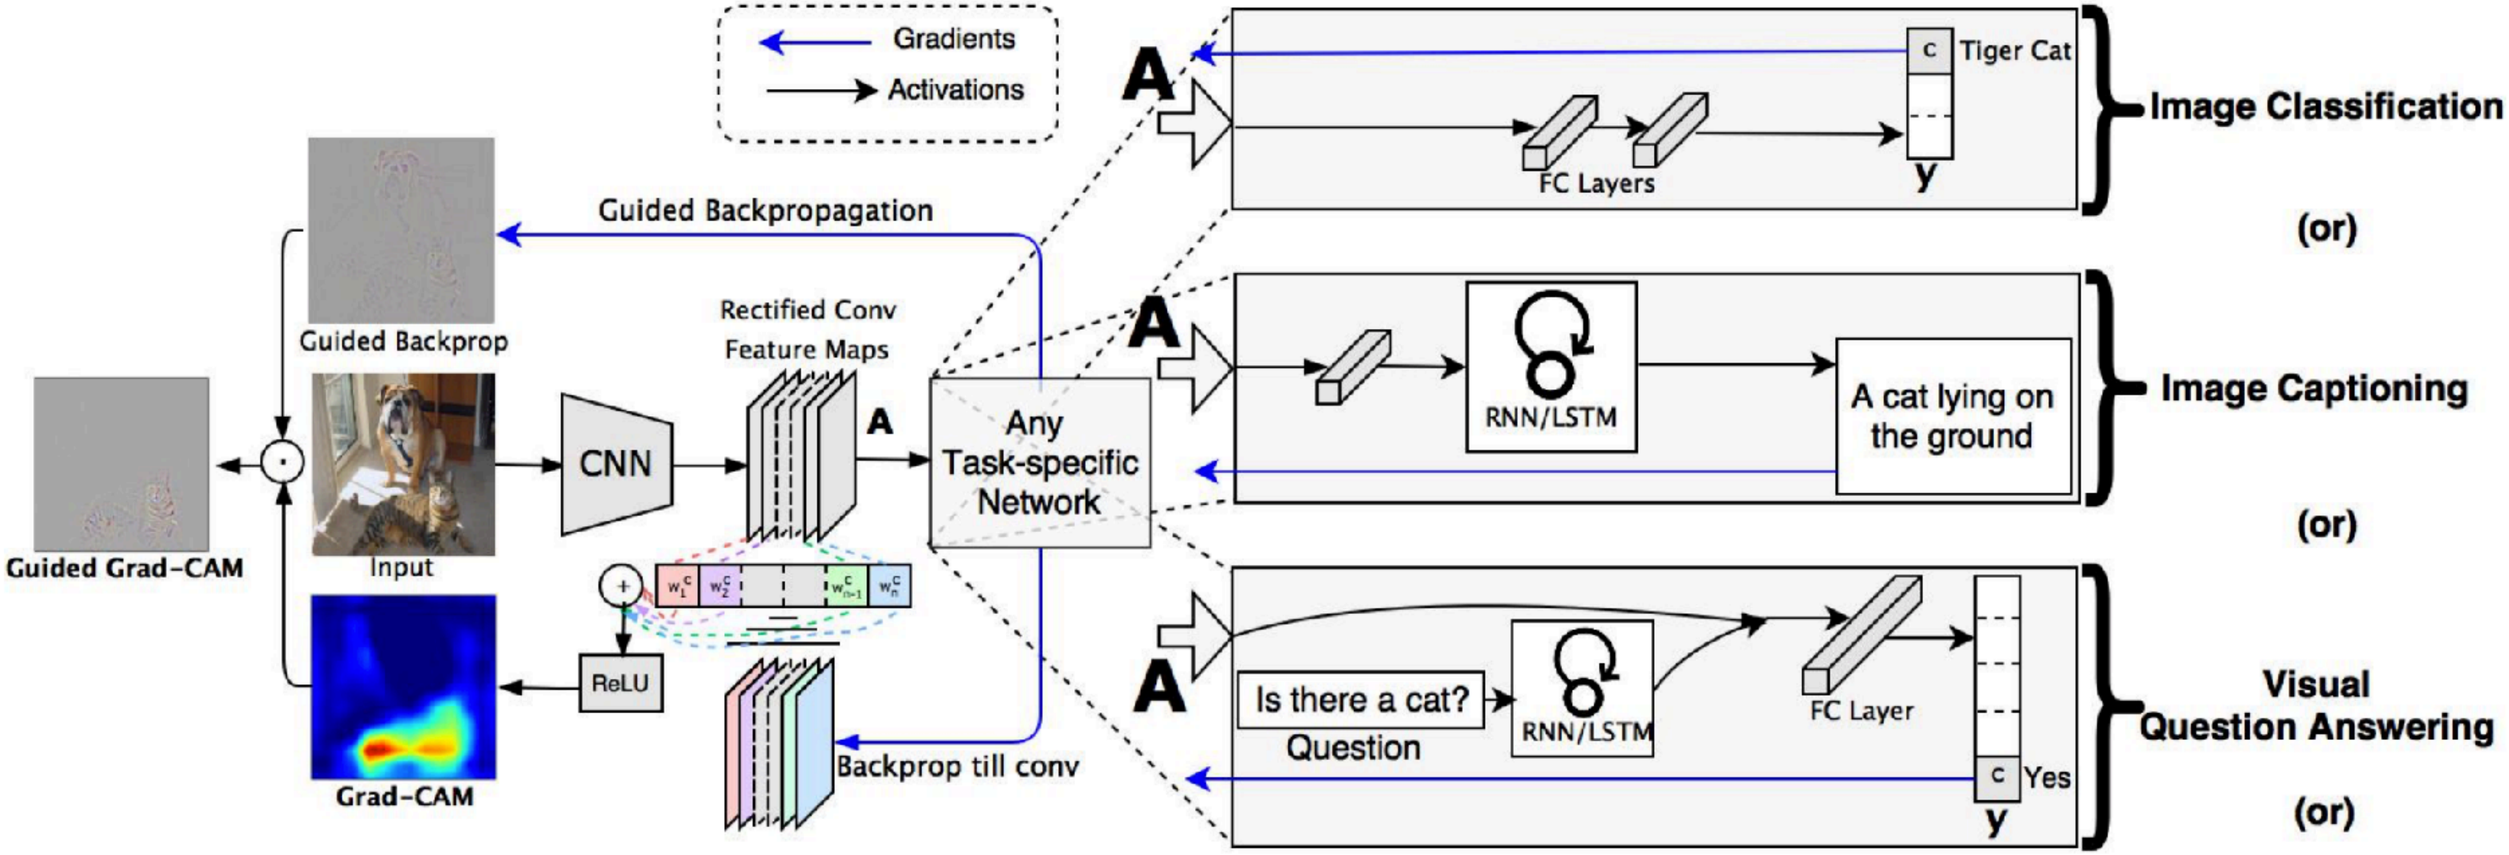

[Selvaraju et al. 2016]

3. Interpretable Deep Learning

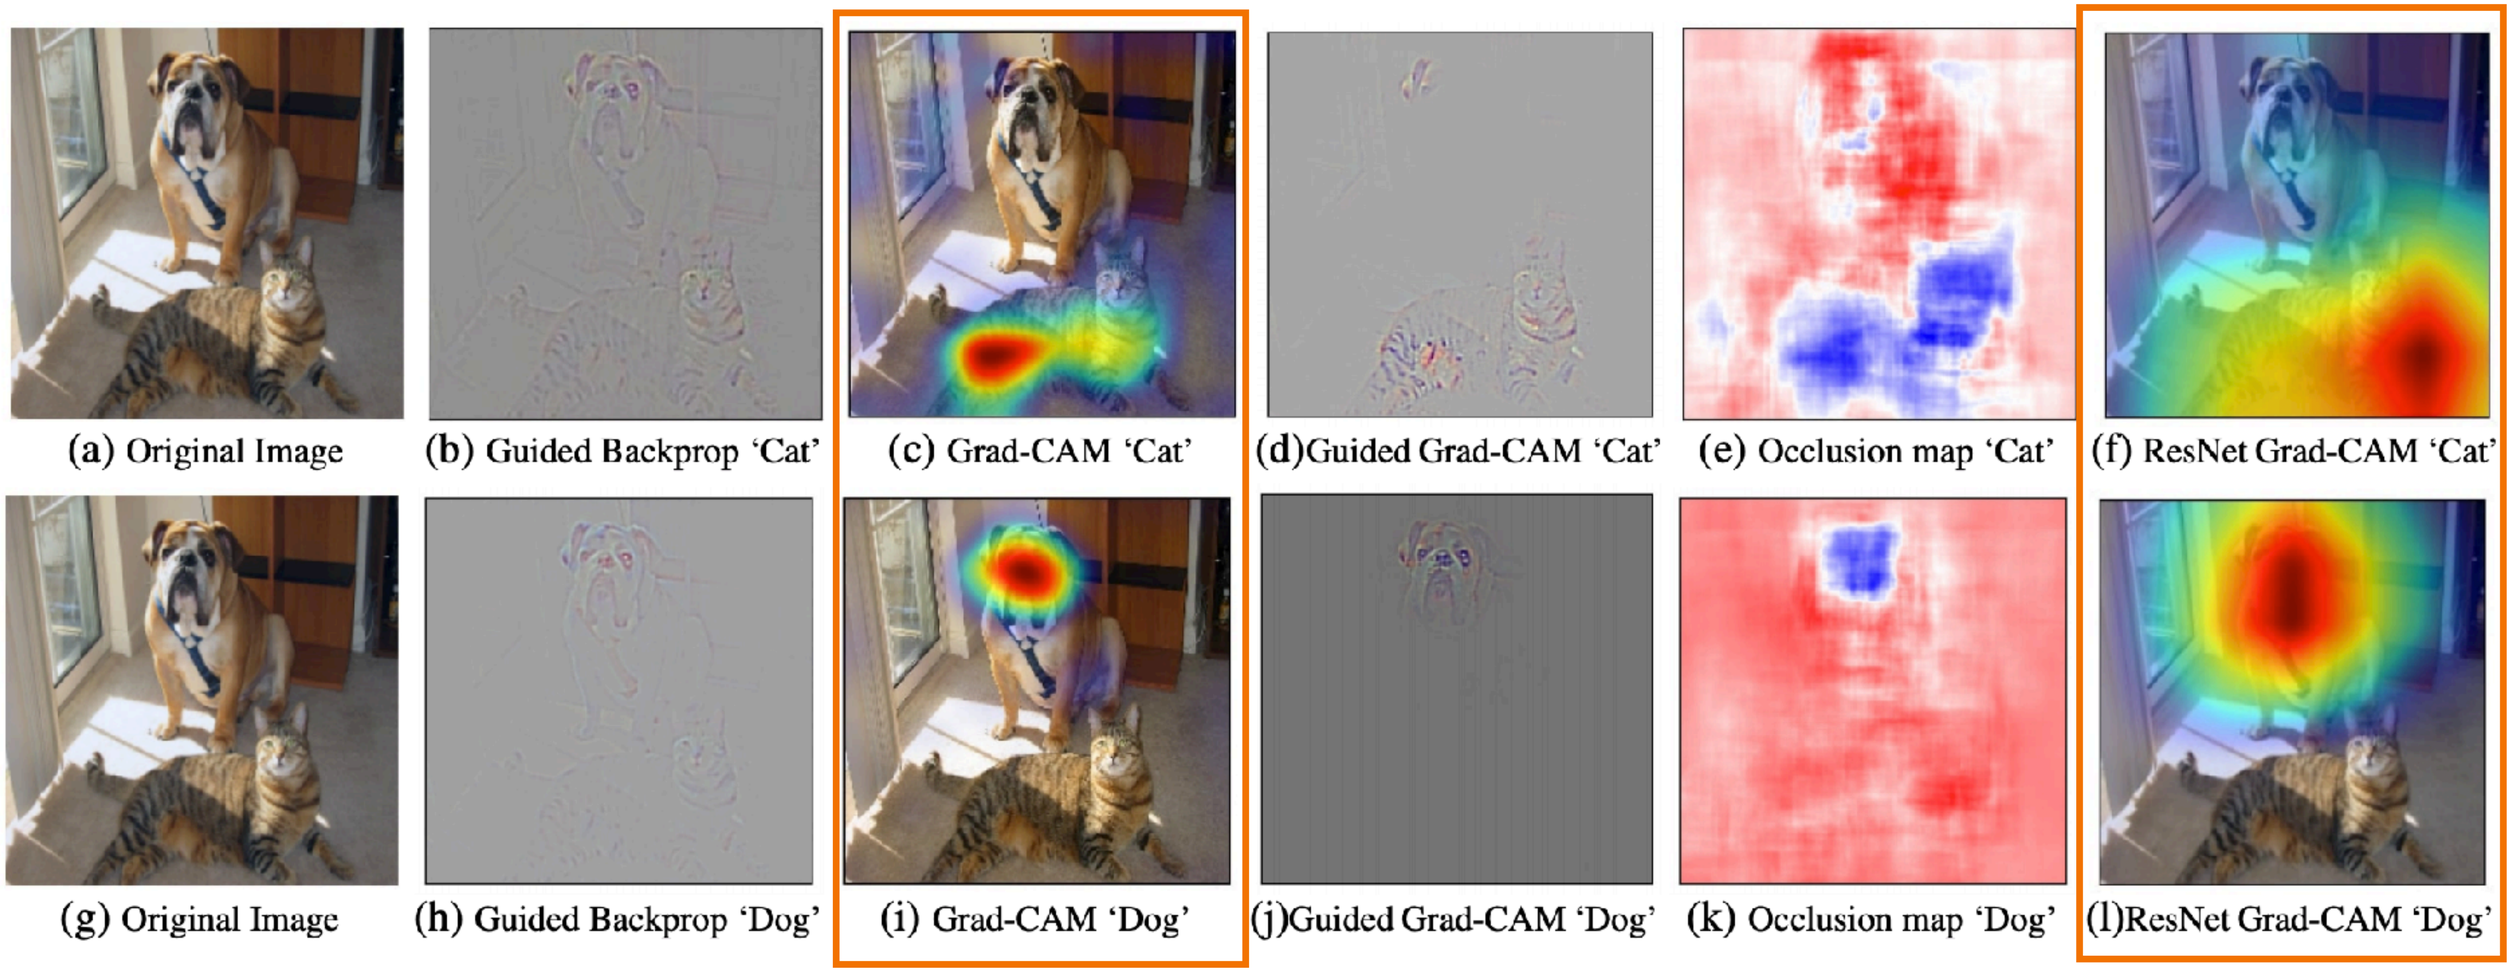

[Selvaraju et al. 2016]

4. Deep Learning Generalization and Theory

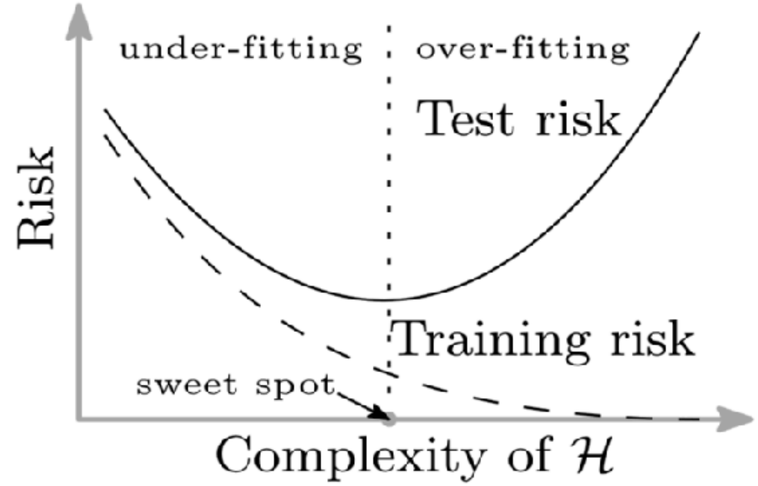

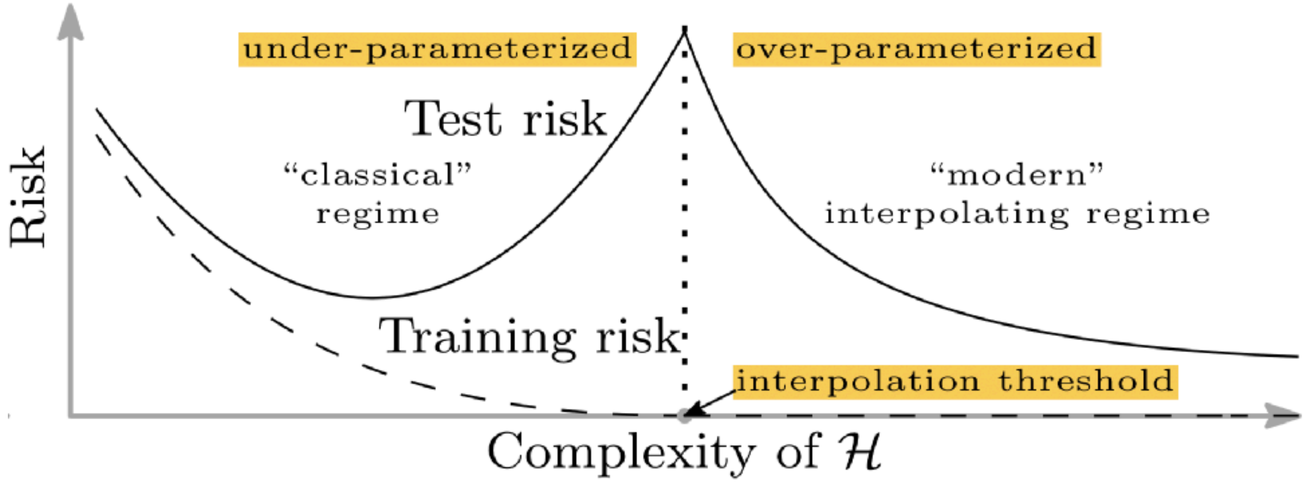

[Belkin et al. 2018]

5. Learning with less supervision

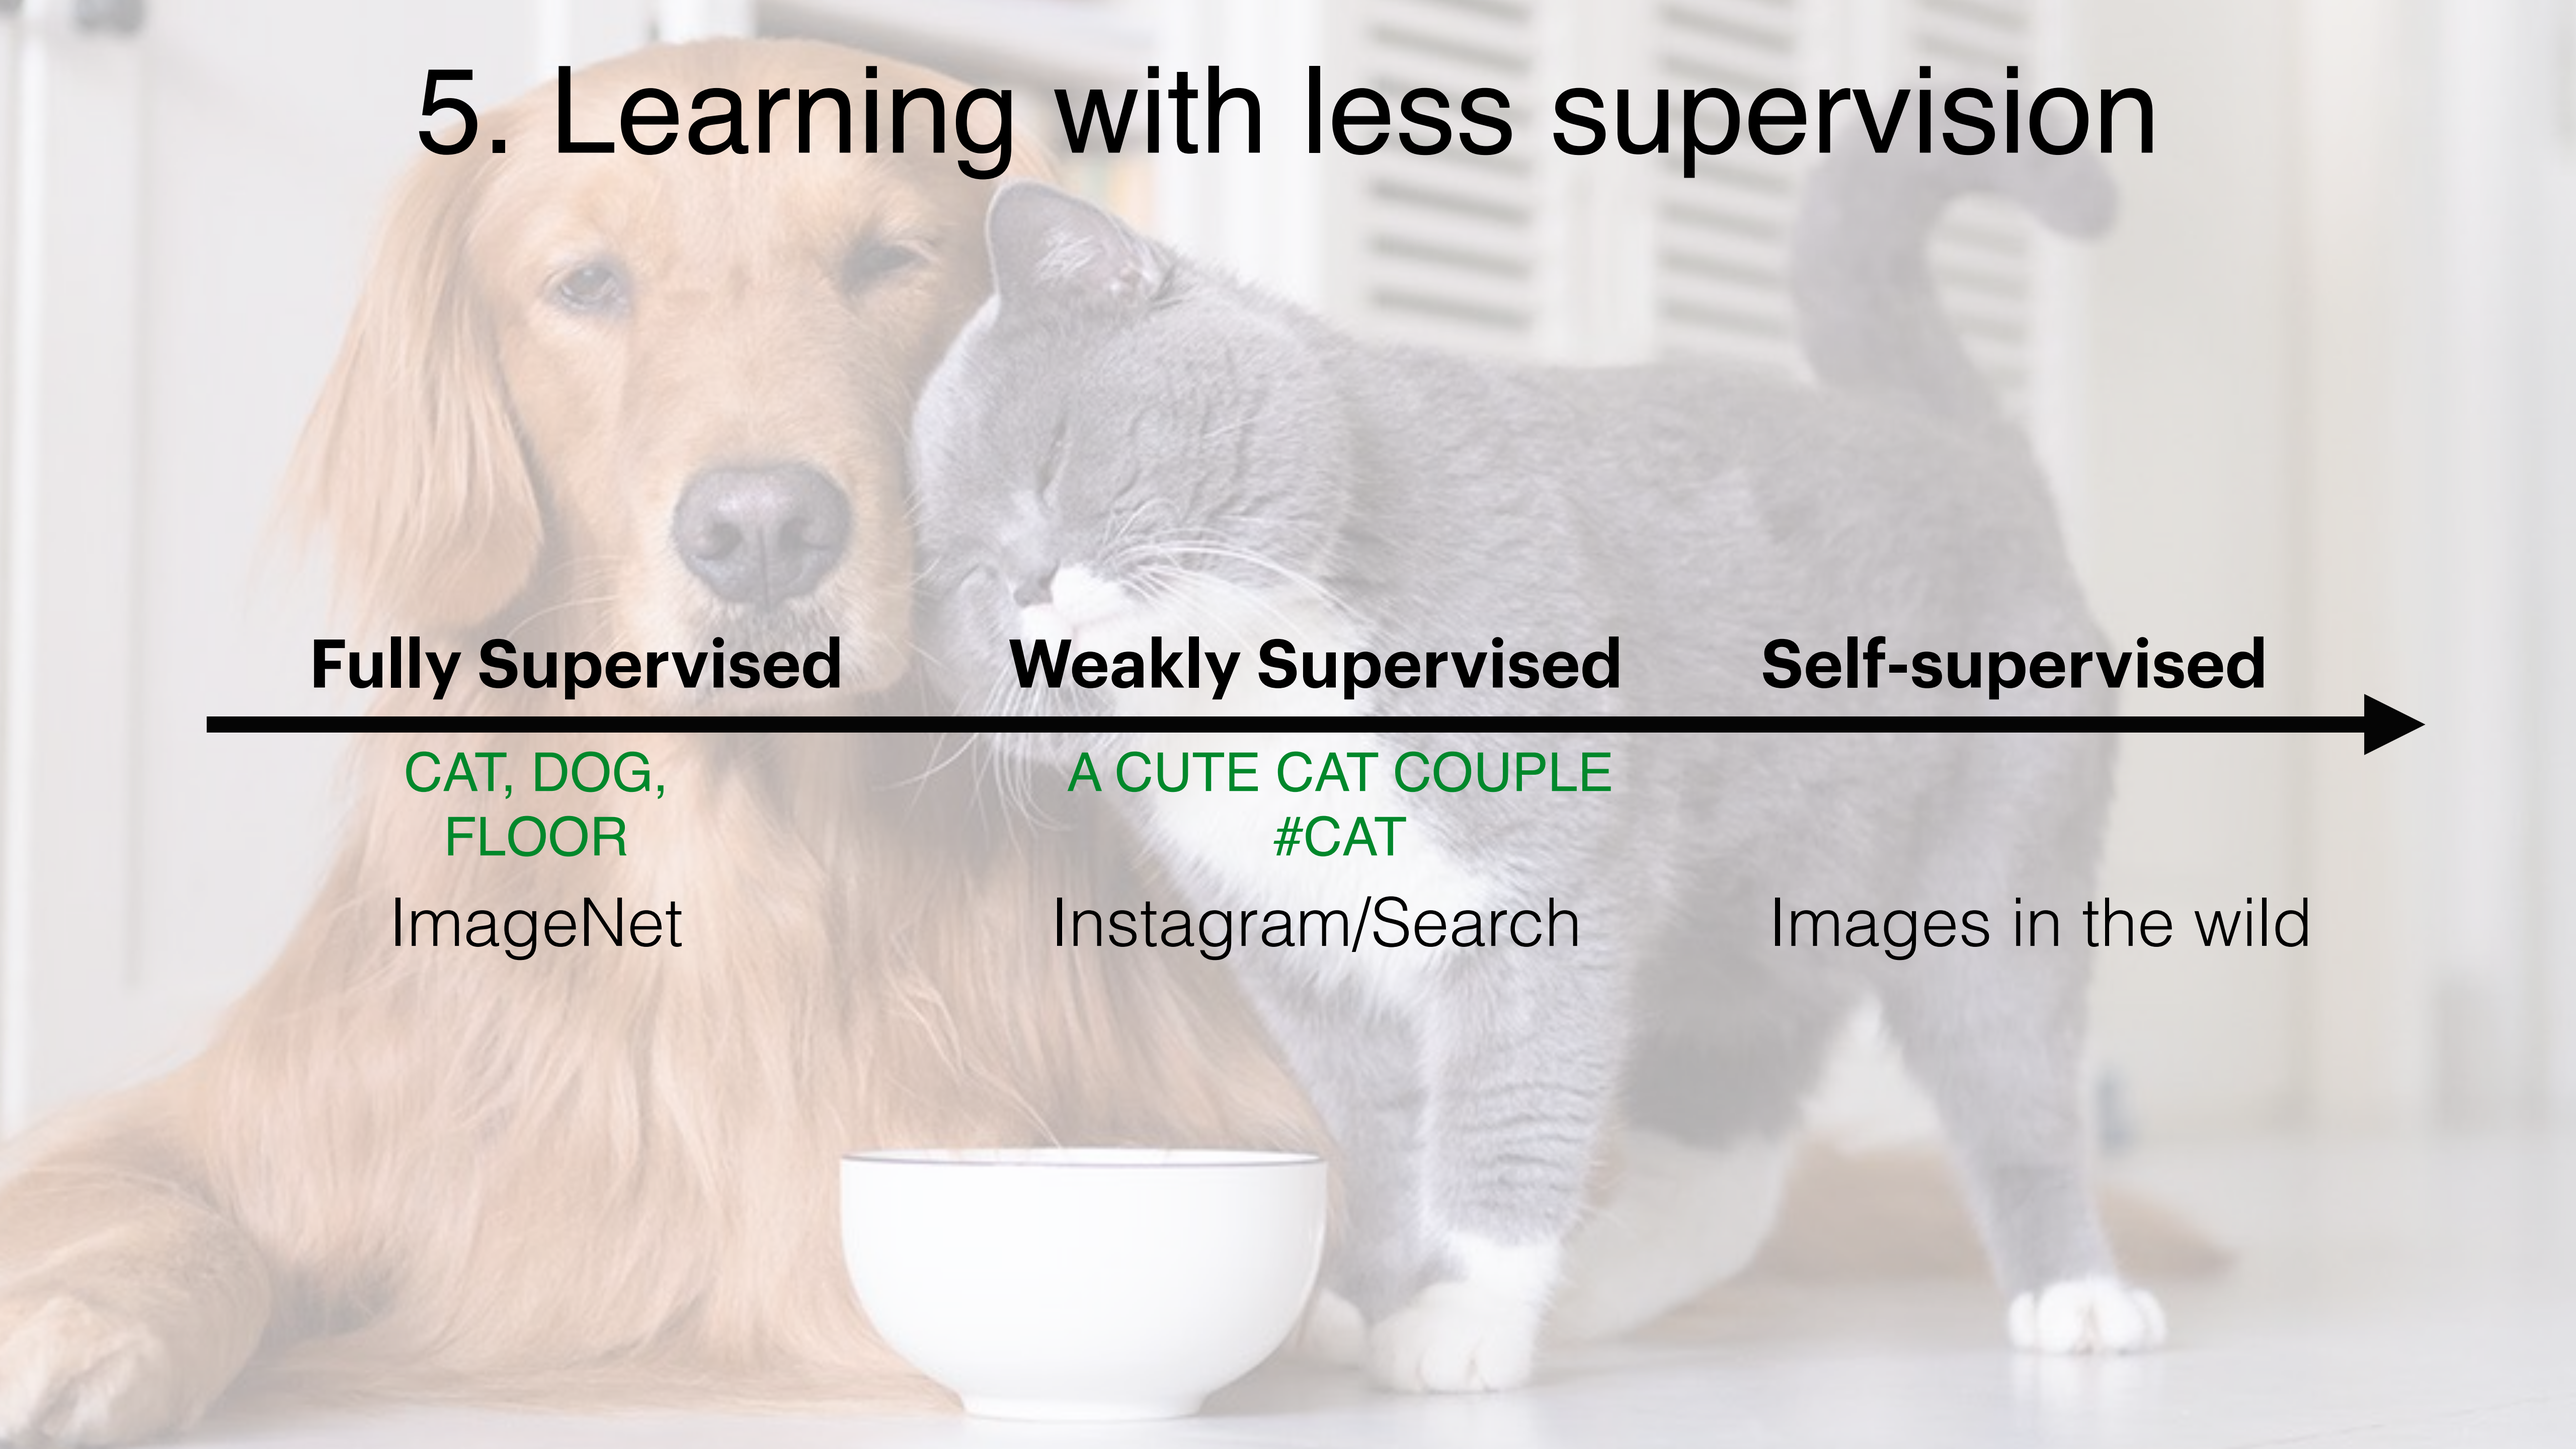

Fully Supervised	**Weakly Supervised**	**Self-supervised**
CAT, DOG, FLOOR	A CUTE CAT COUPLE #CAT	
ImageNet	Instagram/Search	Images in the wild

6. Lifelong Learning

Machines that improve with experience and become “smarter” over time.

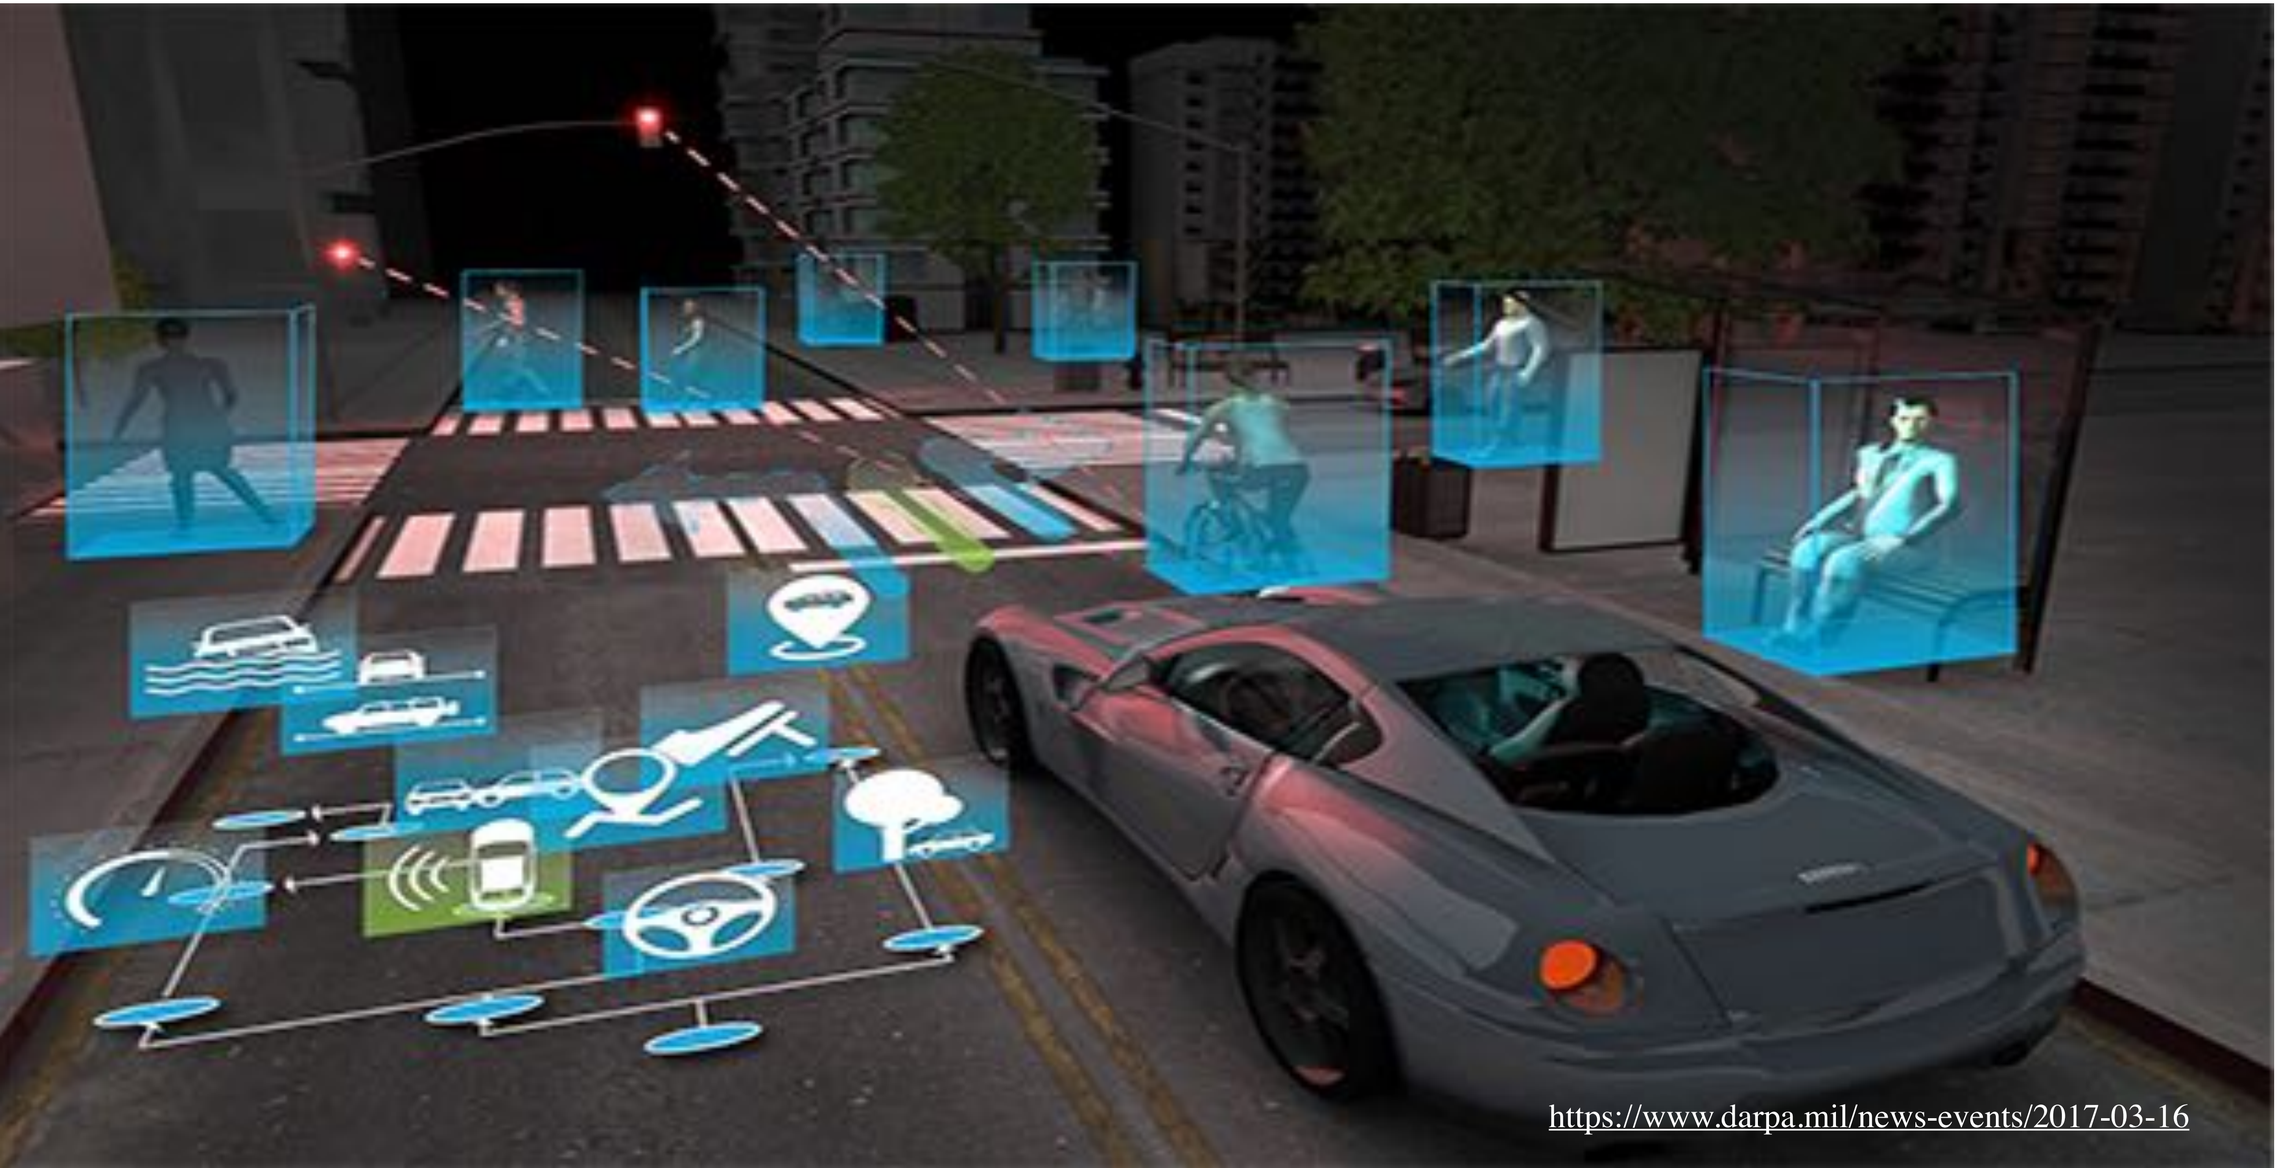

https://www.darpa.mil/news-events/2017-03-16

7. Deep Generative Modeling

4.5 years of face generation

2014

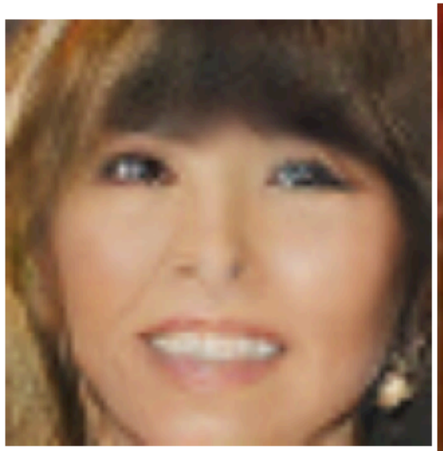

2015

2016

2017

2018

7. Deep Generative Modeling

Synthesize the images

http://www.whichfaceisreal.com/methods.html

7. Deep Generative Modeling

Style transfers

https://github.com/StacyYang/MXNet-Gluon-Style-Transfer

Part III: Get to know EVERYONE

Remember to sign up the paper presentation TODAY!

Thanks!