Deep Learning: Teaming up Machines and Humans to Peek into the Black Box

Andrés Muñoz-Jaramillo
www.solardynamo.org

Southwest Research Institute
High Altitude Observatory
National Solar Observatory

What is deep learning?

What techies think it is

What advocates think it is

What it really is

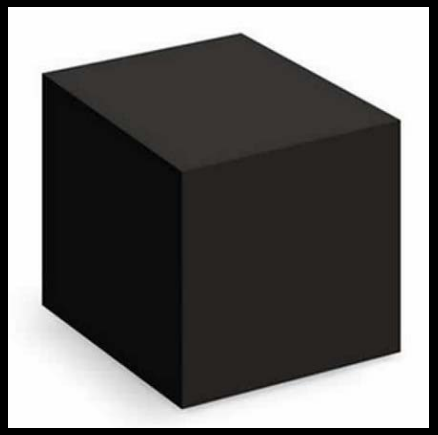

What skeptics think it is

What the public thinks it is

What is deep learning?

What techies think it is

What advocates think it is

What it really is

Getting past preconceptions is a big challenge

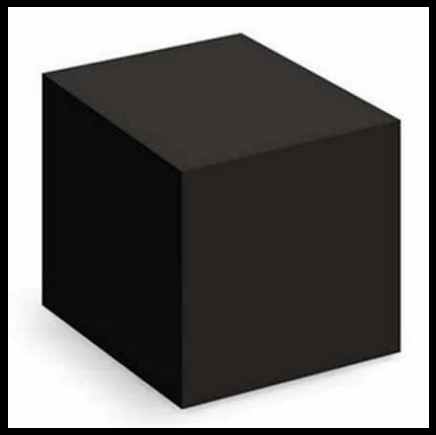

What skeptics think it is

What the public thinks it is

A tool that helps us find new things in our data

What is deep learning?*

A class of machine learning algorithms that:

- Use a cascade of multiple layers of nonlinear processing units for feature extraction and transformation.
- Learn multiple levels of representations that correspond to different levels of abstraction (i.e. the levels form a hierarchy of concepts).

* Taken from Wikipedia

What is deep learning?*

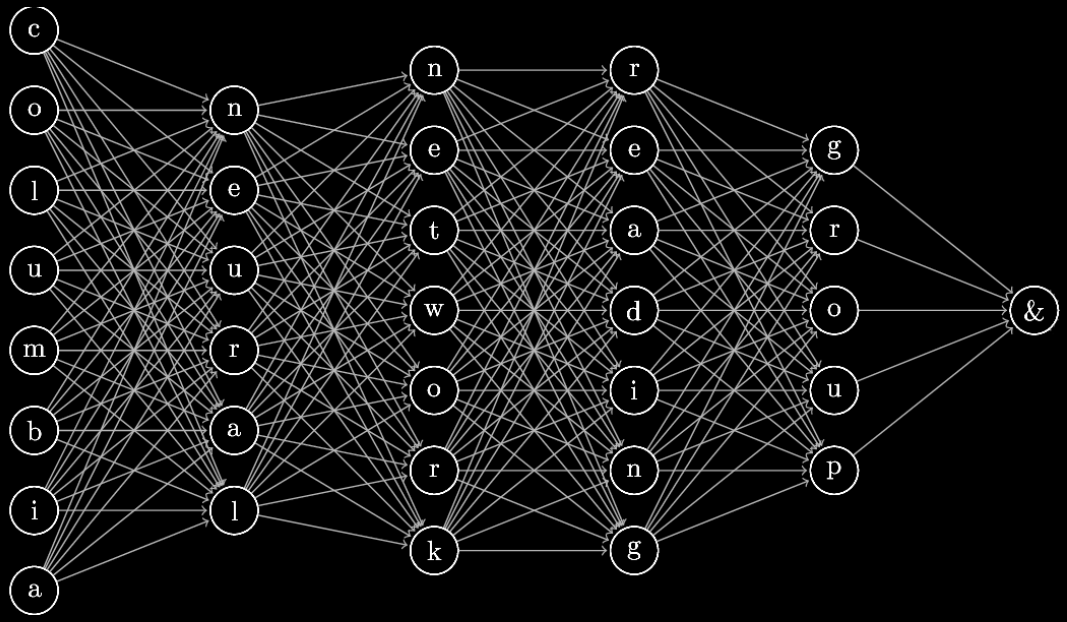

- Learn multiple levels of representations that correspond to different levels of abstraction (i.e. the levels form a hierarchy of concepts).

* Taken from Wikipedia

What is deep learning?*

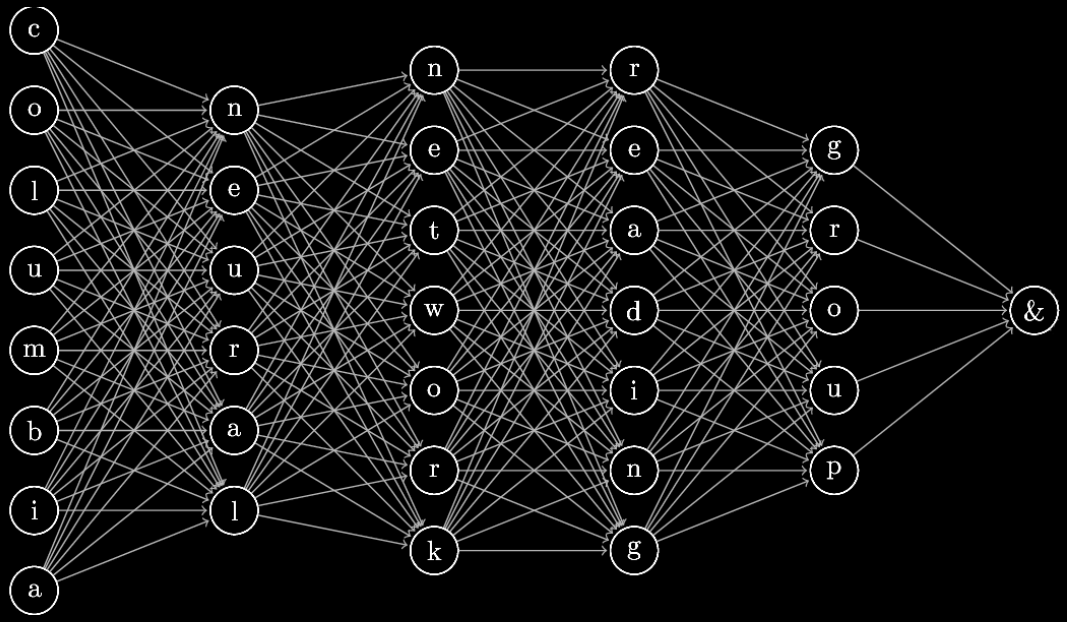

Neural networks are not black boxes.

* Taken from Wikipedia

Why deep learning?

Deep learning has Important limitations too

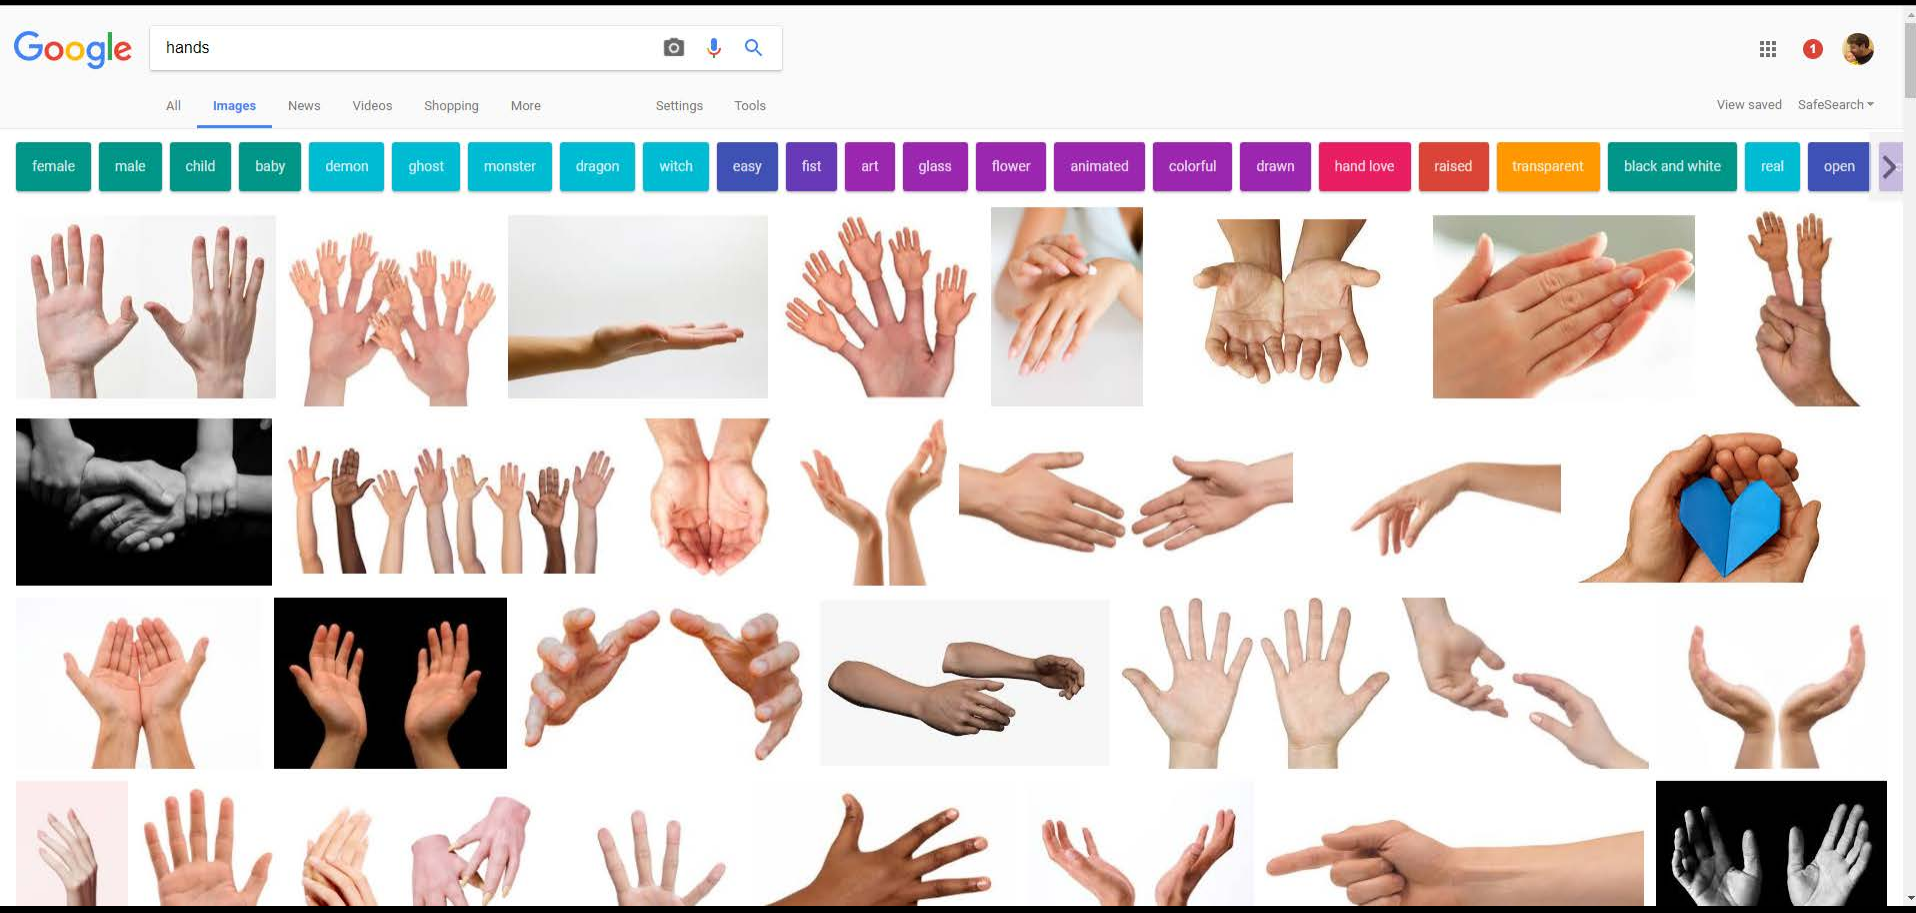

Deep learning has important limitations too

TL;DR \ MICROSOFT \ WEB \

Twitter taught Microsoft's AI chatbot to be a racist asshole in less than a day

68

by James Vincent | @jjvincent | Mar 24, 2016, 6:43am EDT

SHARE TWEET LINKEDIN

It took less than 24 hours for Twitter to corrupt an innocent AI chatbot. Yesterday, Microsoft unveiled Tay — a Twitter bot that the company described as an experiment in "conversational understanding." The more you chat with Tay, said Microsoft, the smarter it gets, learning to engage people through "casual and playful conversation."

NOW TRENDING

The Pixel 2 XL drama is undermining Google's entire Pixel project

Nintendo announces Animal Crossing: Pocket Camp for smartphones

Deep learning has limitations too

- Deep learning algorithms are completely naïve and single-minded in the way they learn.
- Training data selection is absolutely critical for their success.

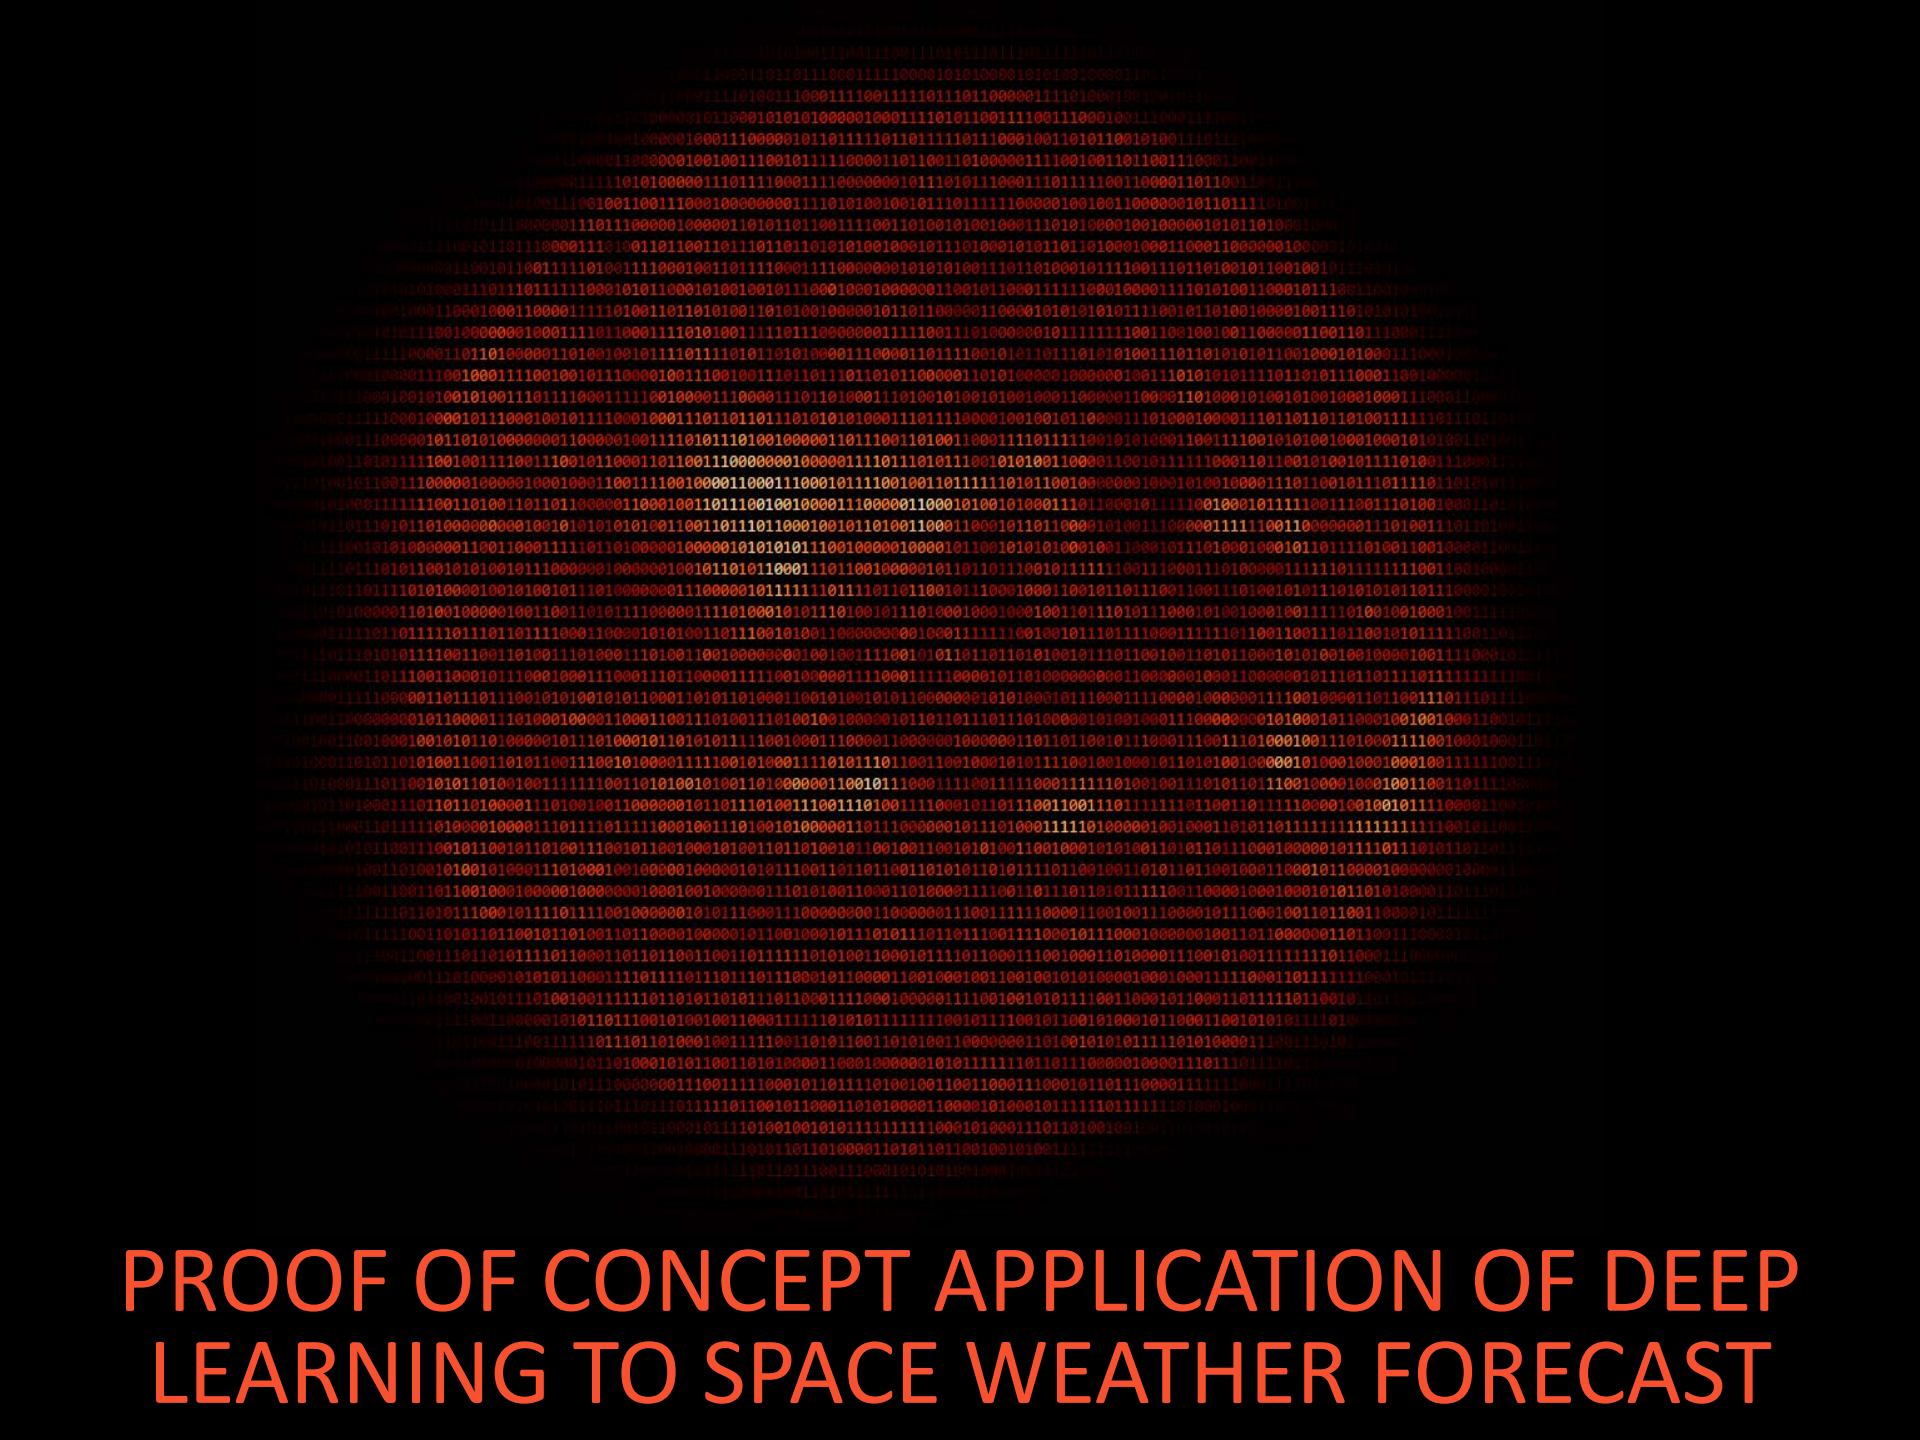

PROOF OF CONCEPT APPLICATION OF DEEP LEARNING TO SPACE WEATHER FORECAST

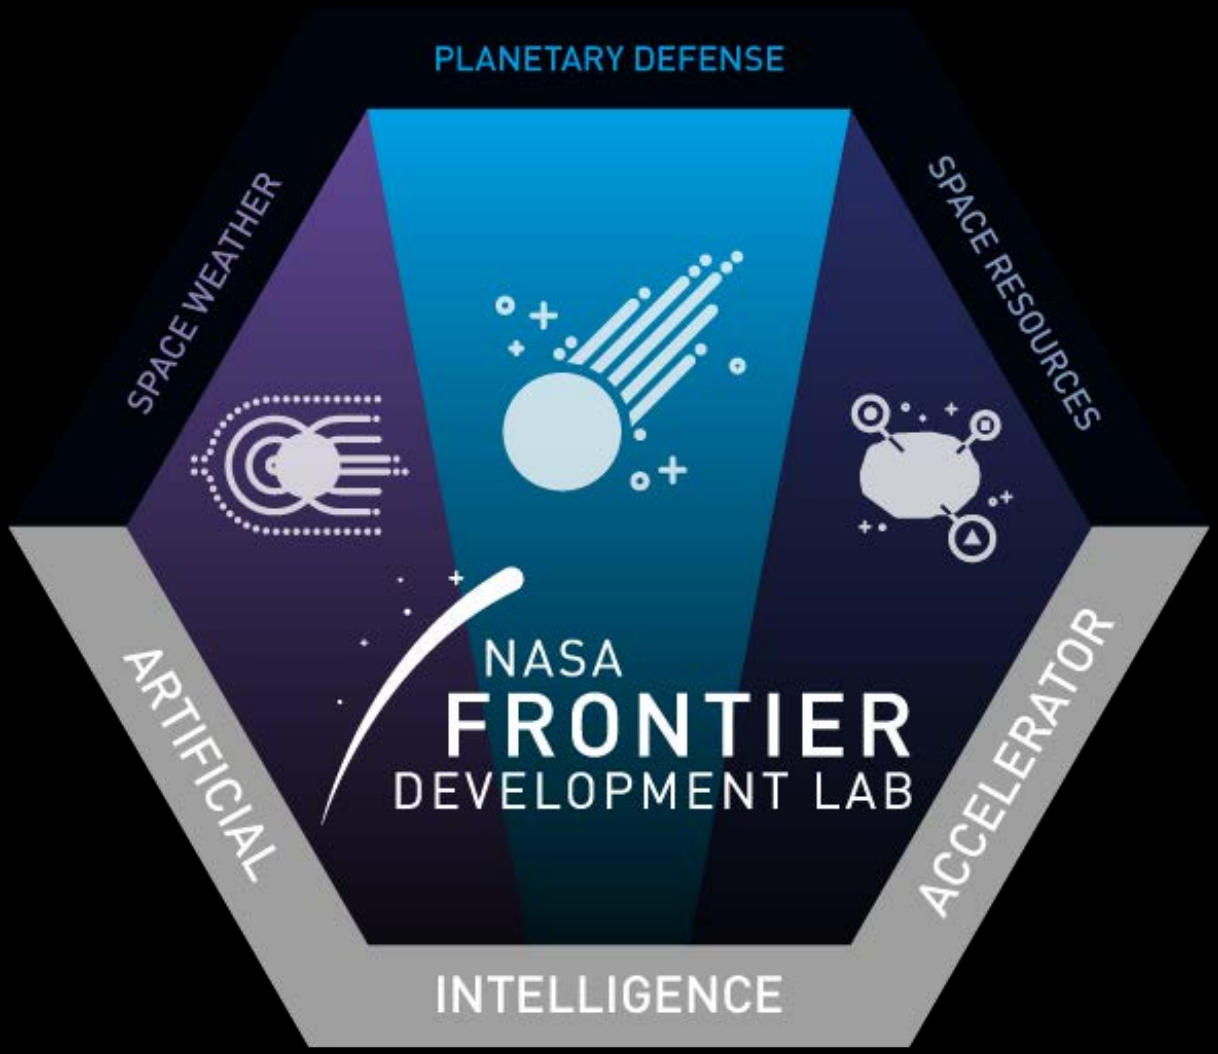

AN APPLIED RESEARCH ACCELERATOR DESIGNED TO ENHANCE NASA'S CAPABILITIES BY COMBINING THE EXPERTISE OF NASA, ACADEMIA, AND THE PRIVATE RESEARCH COMMUNITY.


LONG PERIOD COMETS
FRONTIER DEVELOPMENT LAB 2017

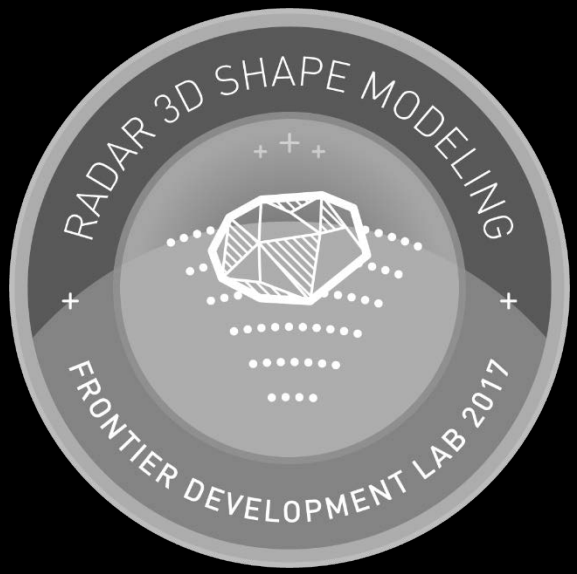
RADAR 3D SHAPE MODELING
FRONTIER DEVELOPMENT LAB 2017

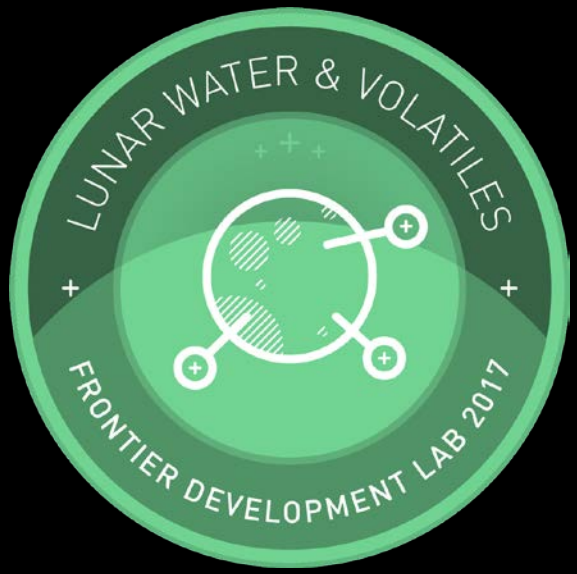
LUNAR WATER & VOLATILES
FRONTIER DEVELOPMENT LAB 2017

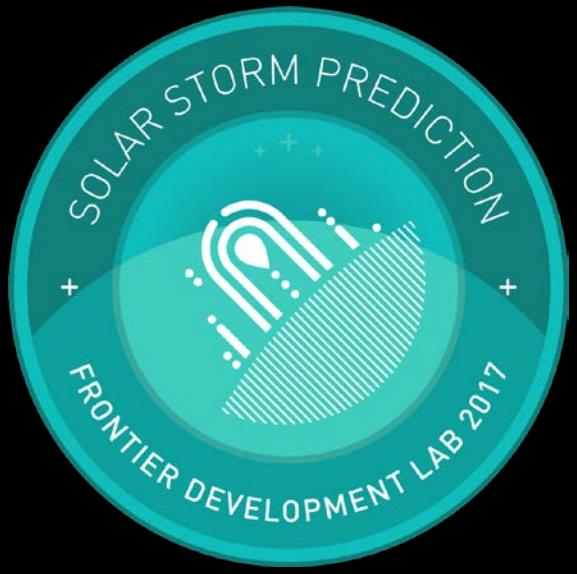
SOLAR STORM PREDICTION
FRONTIER DEVELOPMENT LAB 2017


SOLAR-TERRESTRIAL INTERACTIONS
FRONTIER DEVELOPMENT LAB 2017

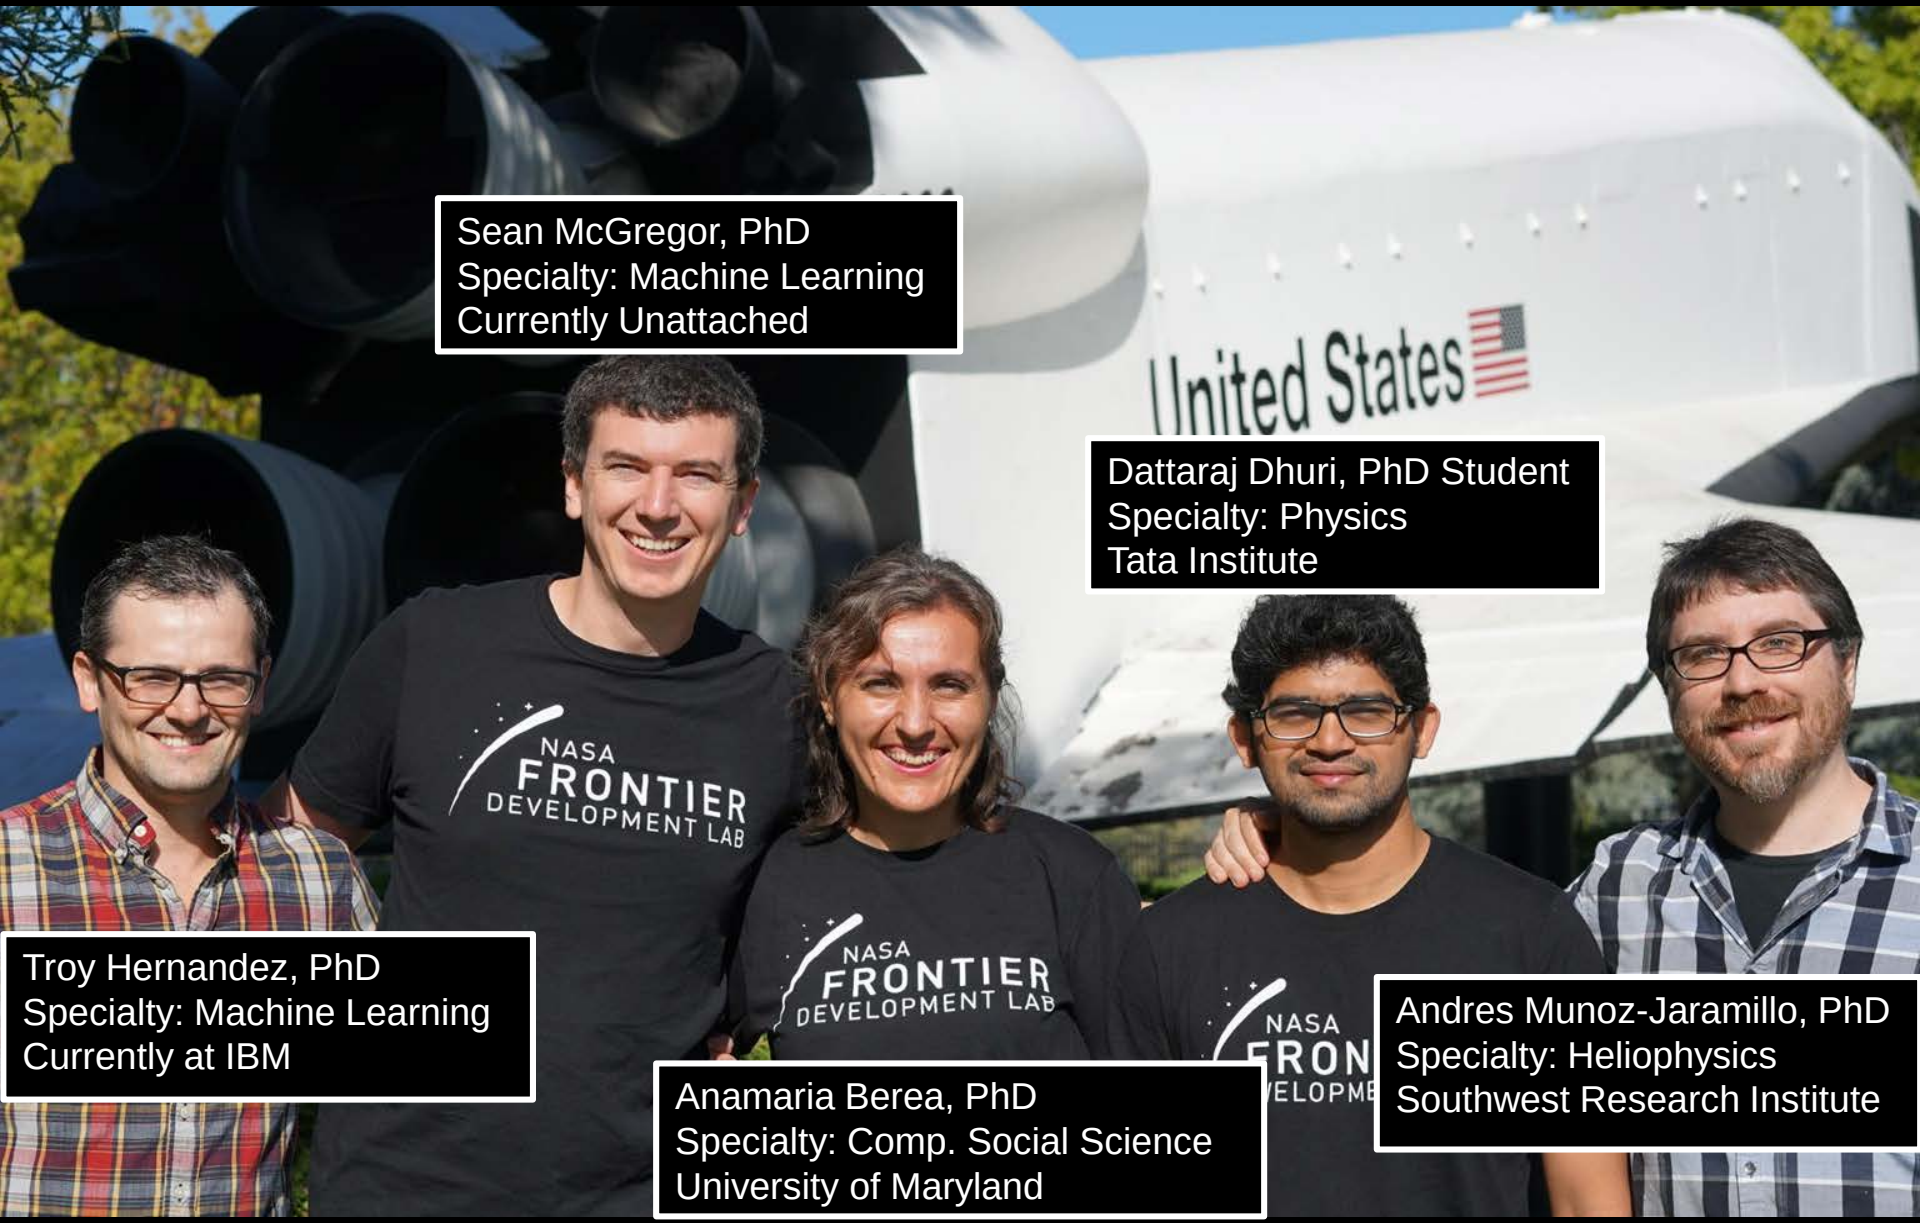
Sean McGregor, PhD
Specialty: Machine Learning
Currently Unattached
Dattaraj Dhuri, PhD Student
Specialty: Physics
Tata Institute
Troy Hernandez, PhD
Specialty: Machine Learning
Currently at IBM
Anamaria Berea, PhD
Specialty: Comp. Social Science
University of Maryland
Andres Munoz-Jaramillo, PhD
Specialty: Heliophysics
Southwest Research Institute
United States
NASA
FRONTIER
DEVELOPMENT LAB

Interdisciplinary research is fun!

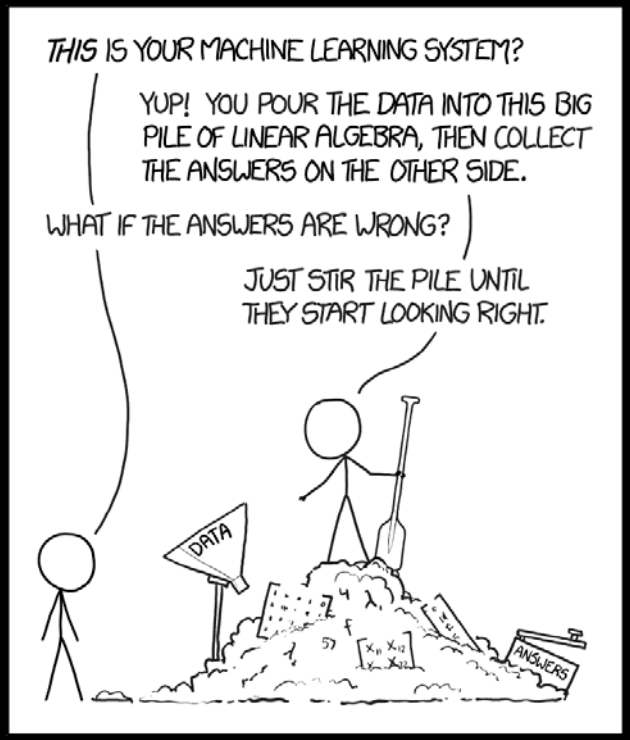

Heliophysicist's view of ML

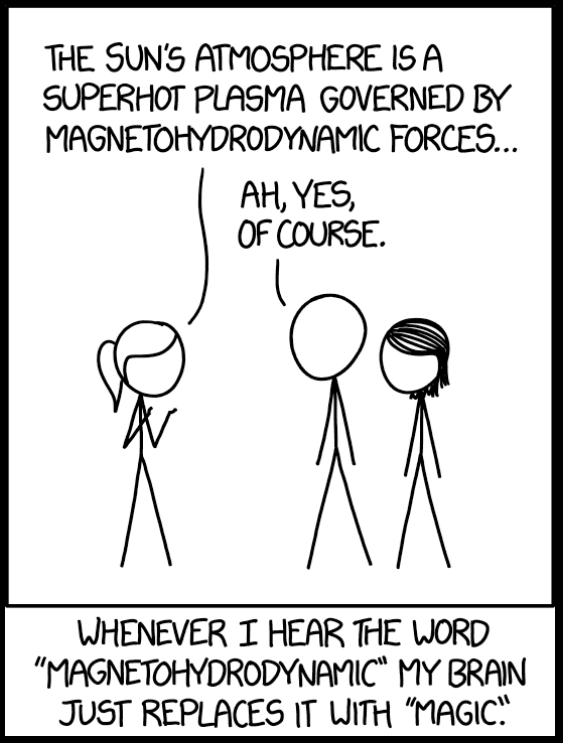

Data scientist's view of HP

Which aspect of space weather to tackle?

FLARES

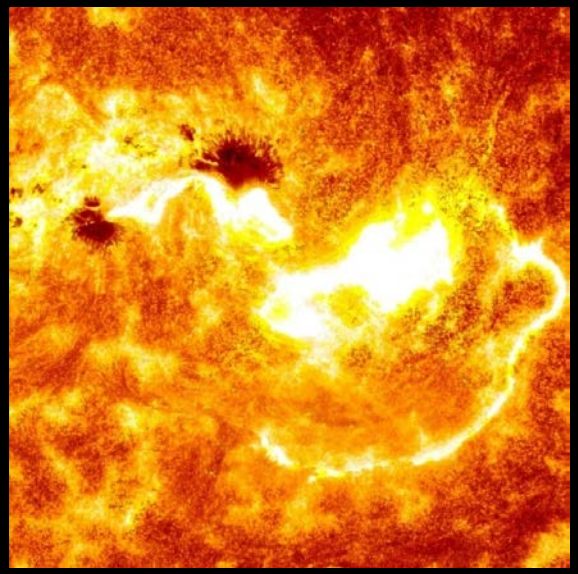

Speed of Light

No warning

ENERGETIC PARTICLES

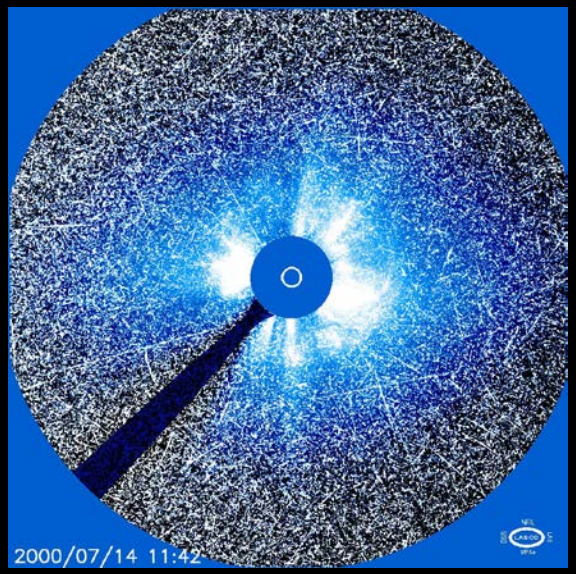

Relativistic speeds

20 minute warning

MASS EJECTIONS

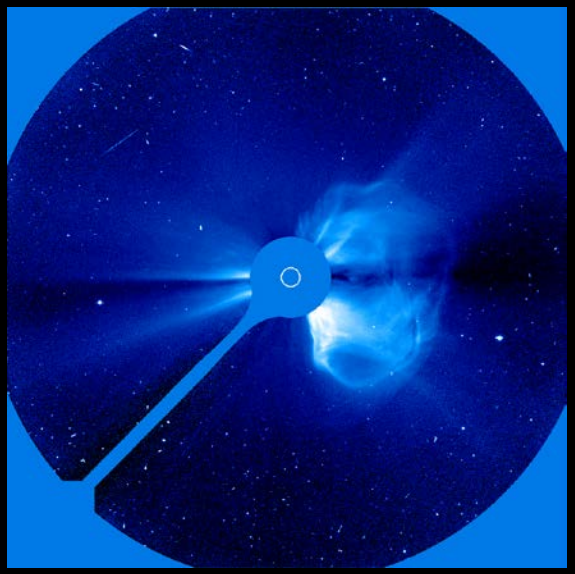

Non-relativistic speeds

20 hour warning

Deep learning and image data

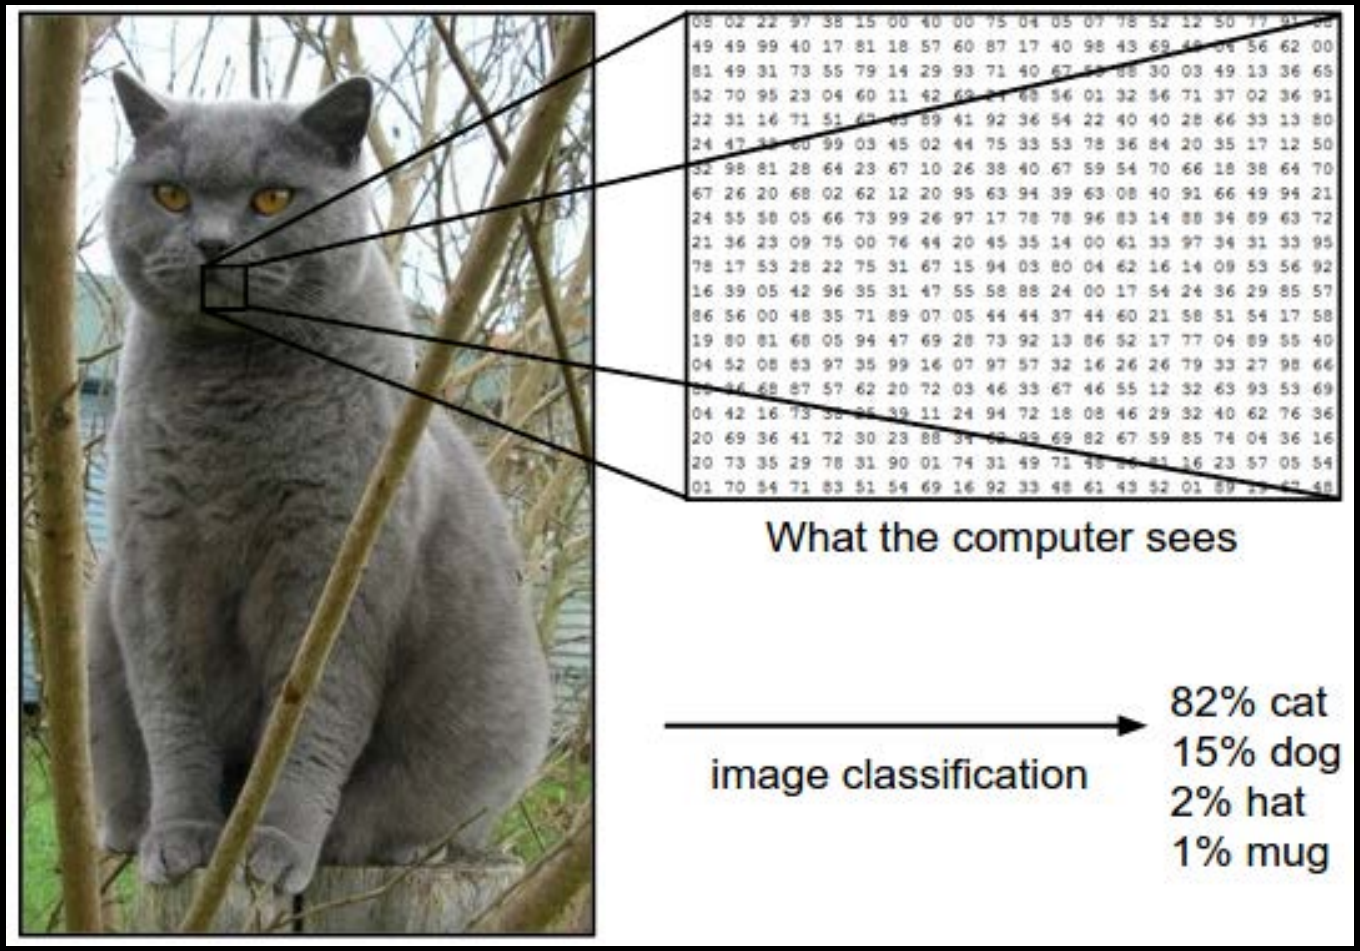

cs231n.github.io/classification

Deep learning and image data

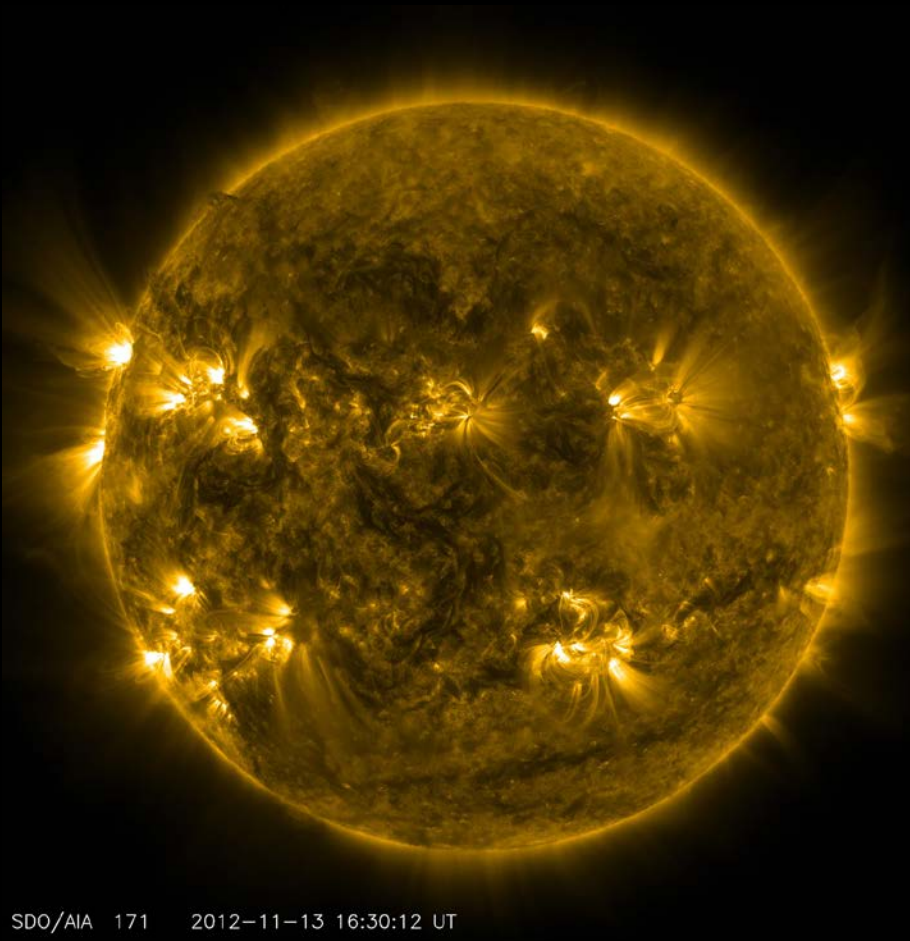

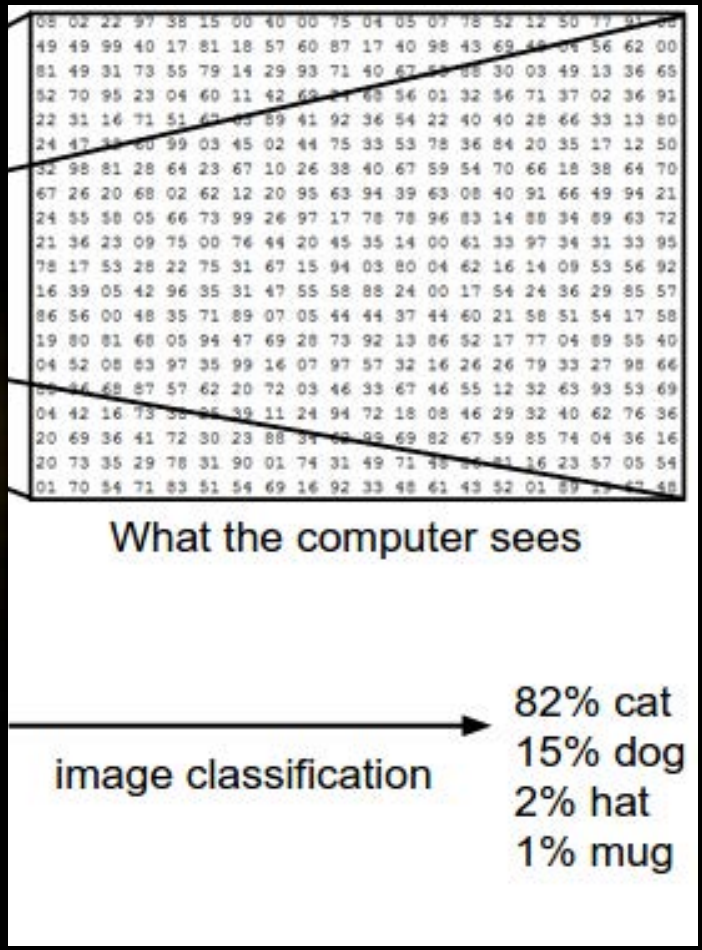

cs231n.github.io/classification

Deep learning and image data

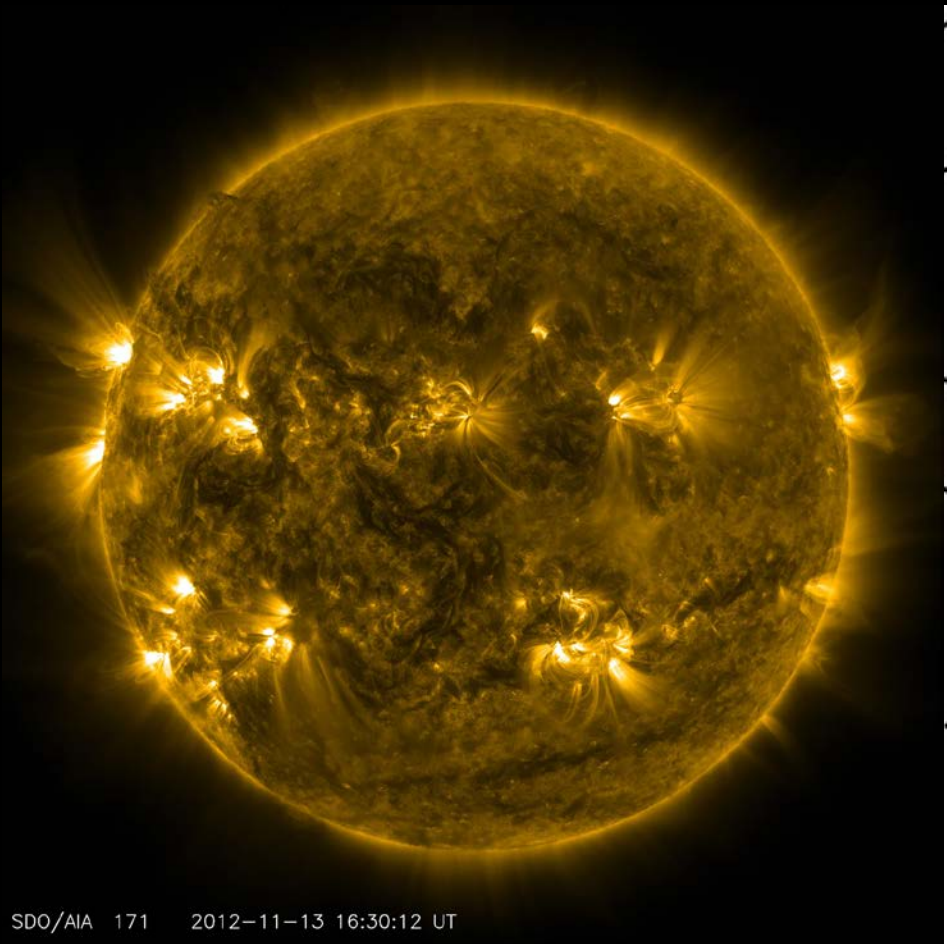

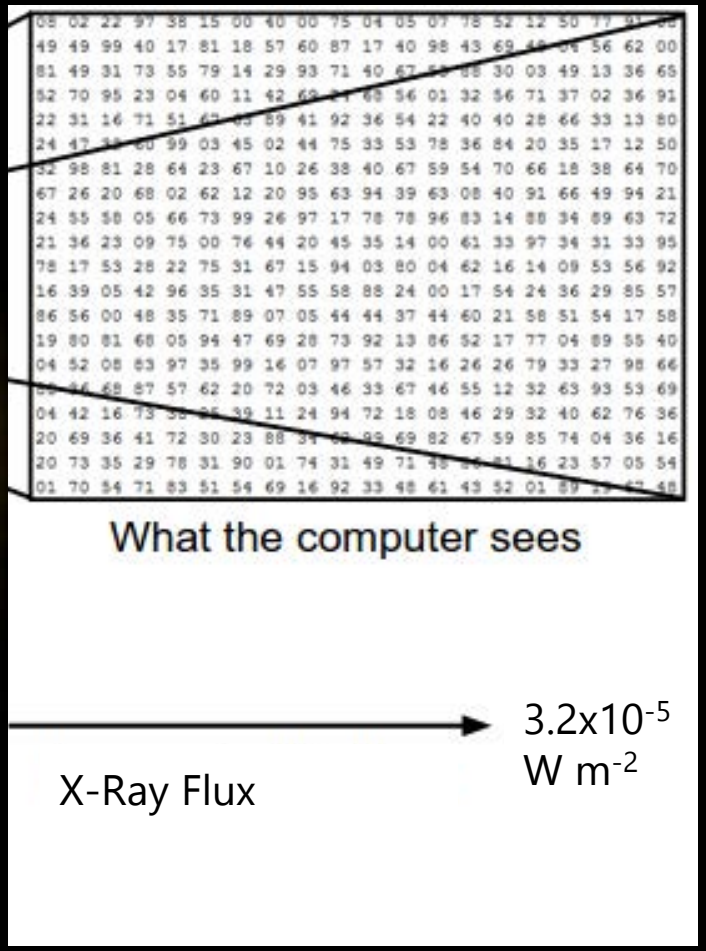

cs231n.github.io/classification

Convolutional Neural Networks

Neural networks with layers made of tunable convolution filters

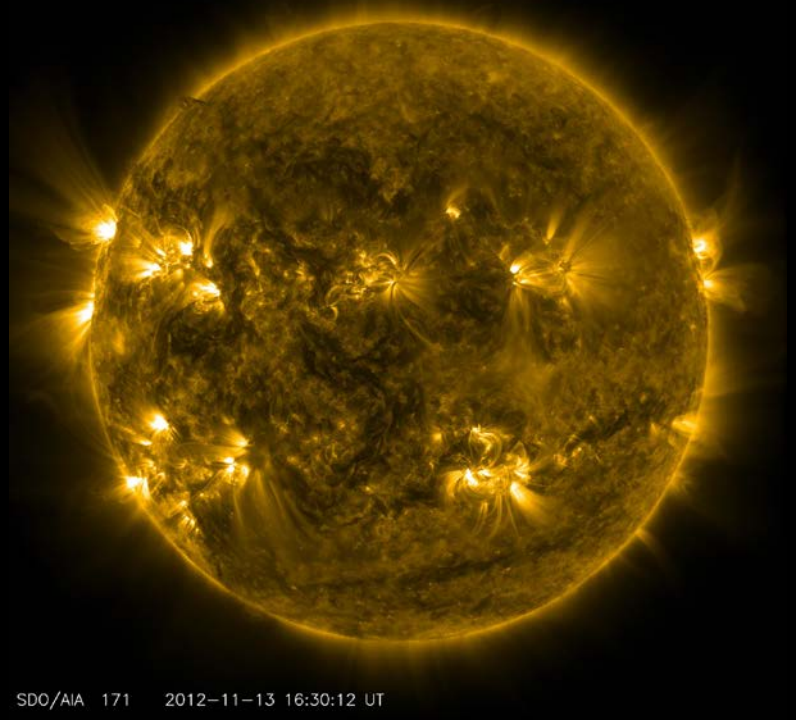

Convolutional Neural Networks

Neural networks with layers made of tunable convolution filters

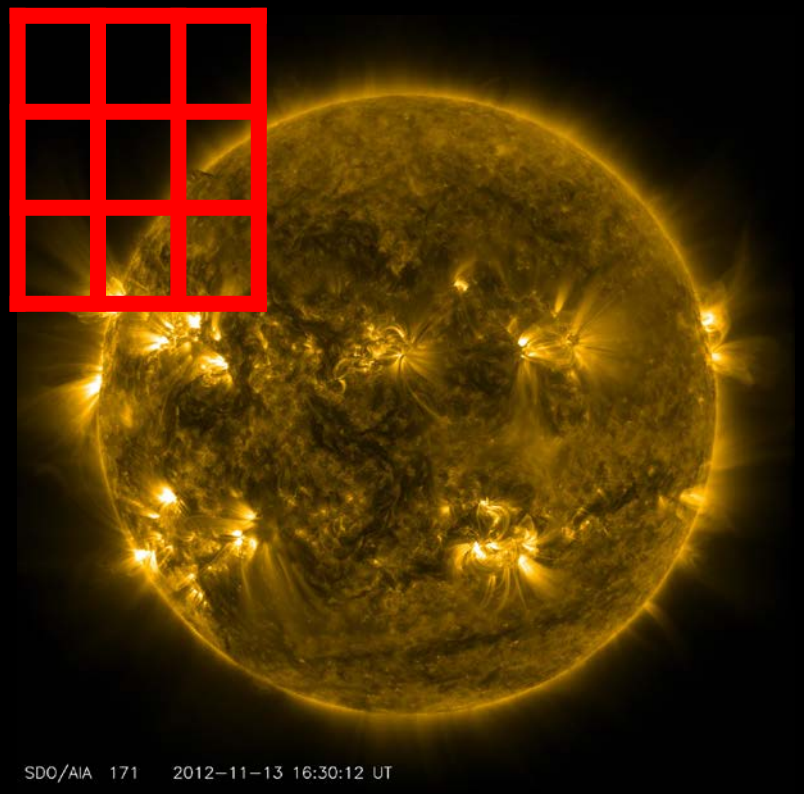

Convolutional Neural Networks

Neural networks with layers made of tunable convolution filters

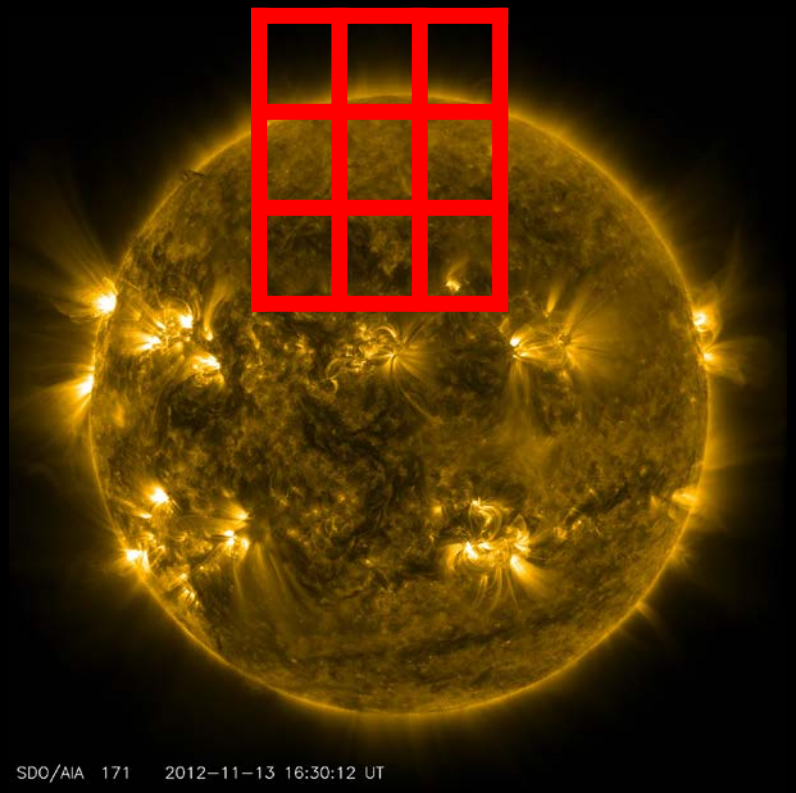

Convolutional Neural Networks

Neural networks with layers made of tunable convolution filters

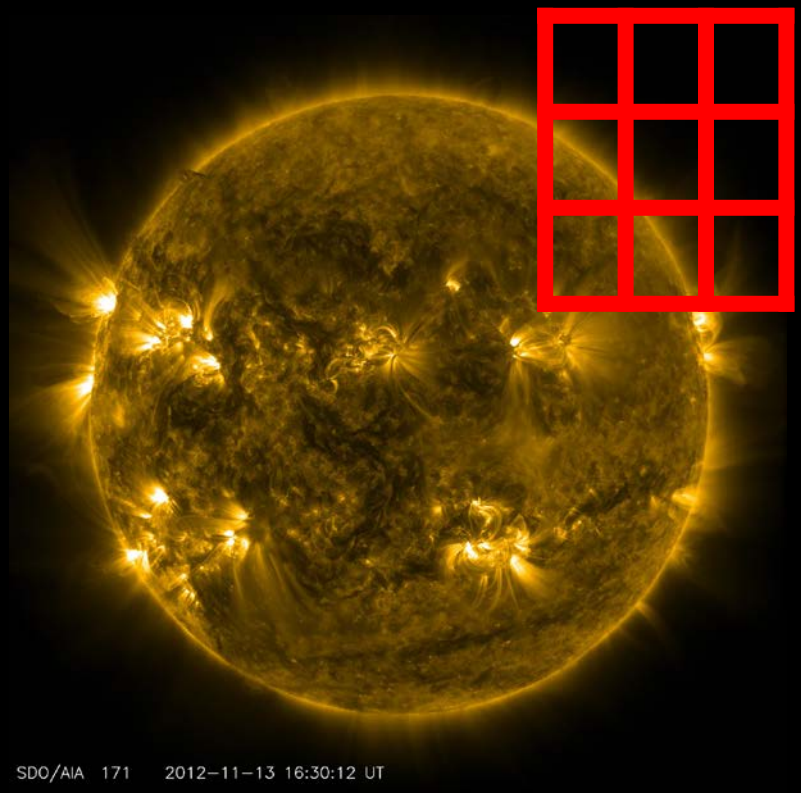

Convolutional Neural Networks

Neural networks with layers made of tunable convolution filters

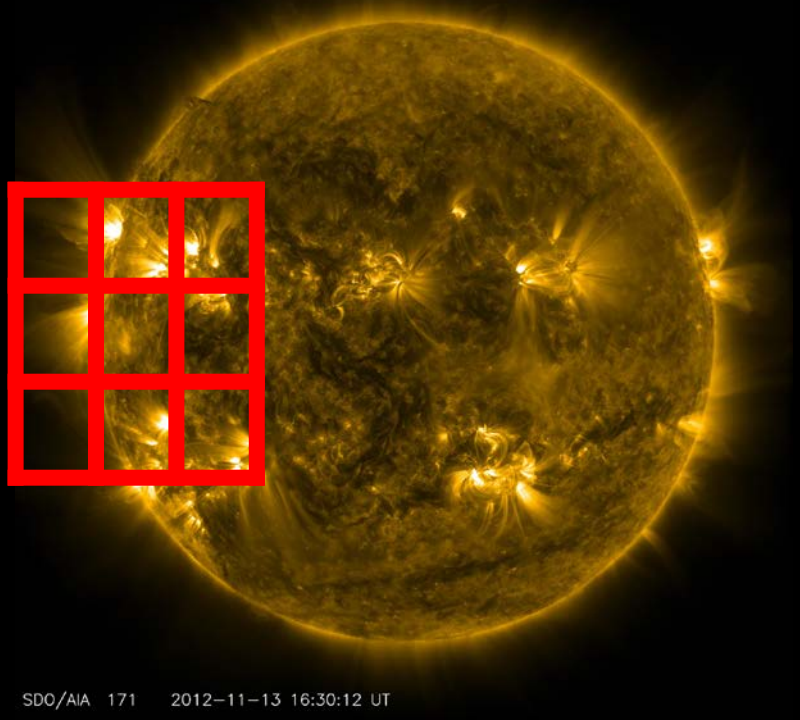

Convolutional Neural Networks

Neural networks with layers made of tunable convolution filters

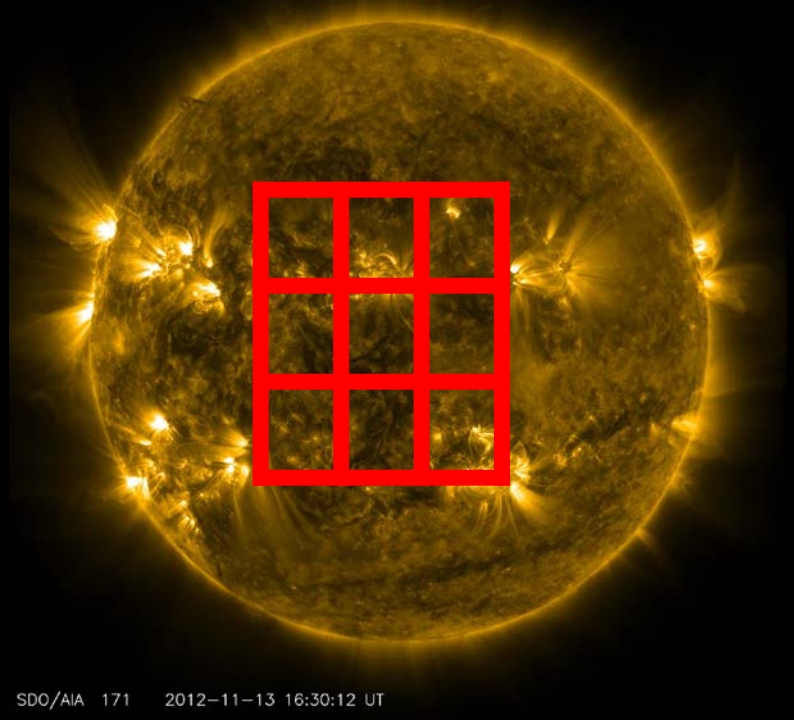

Convolutional Neural Networks

Neural networks with layers made of tunable convolution filters

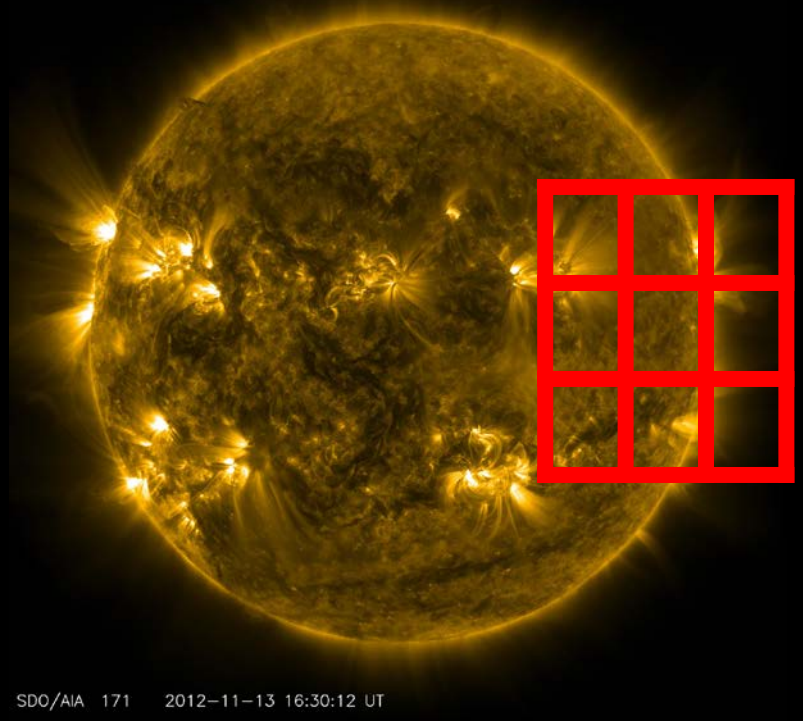

Convolutional Neural Networks

Neural networks with layers made of tunable convolution filters

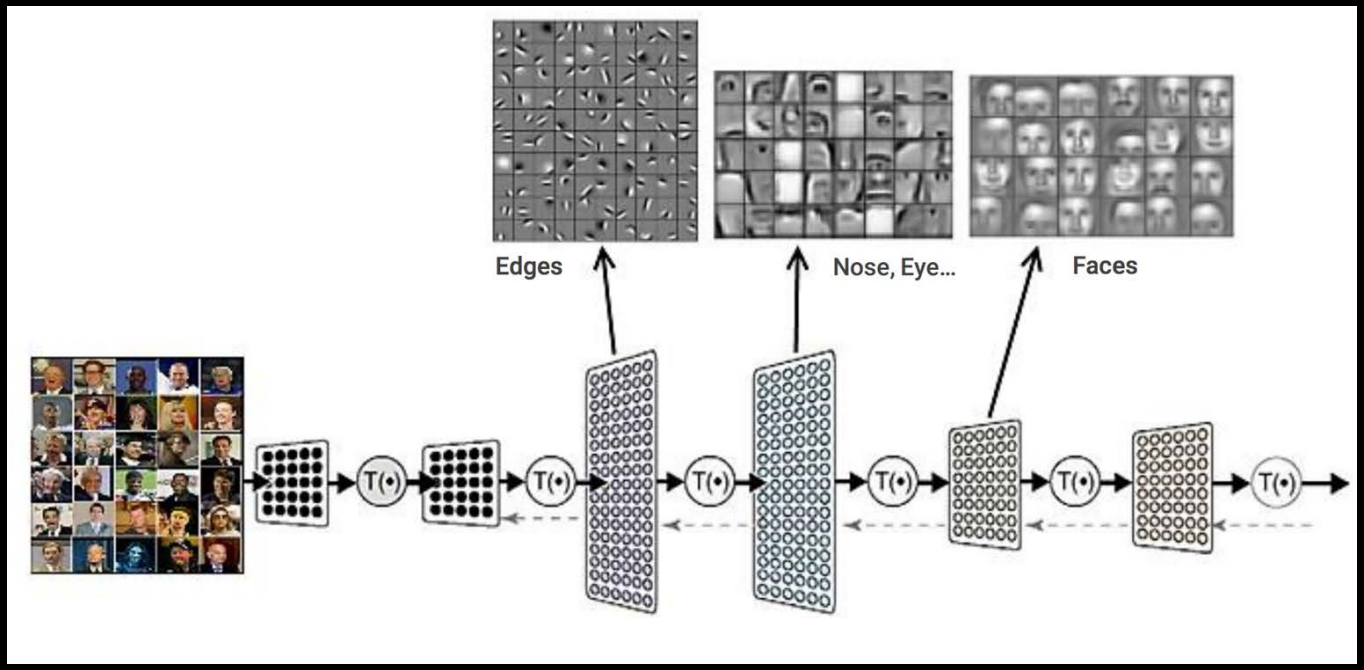

Several convolutional layers allow the neural network to recognize features of increased complexity

Convolutional Neural Networks

Neural networks with layers made of tunable convolution filters

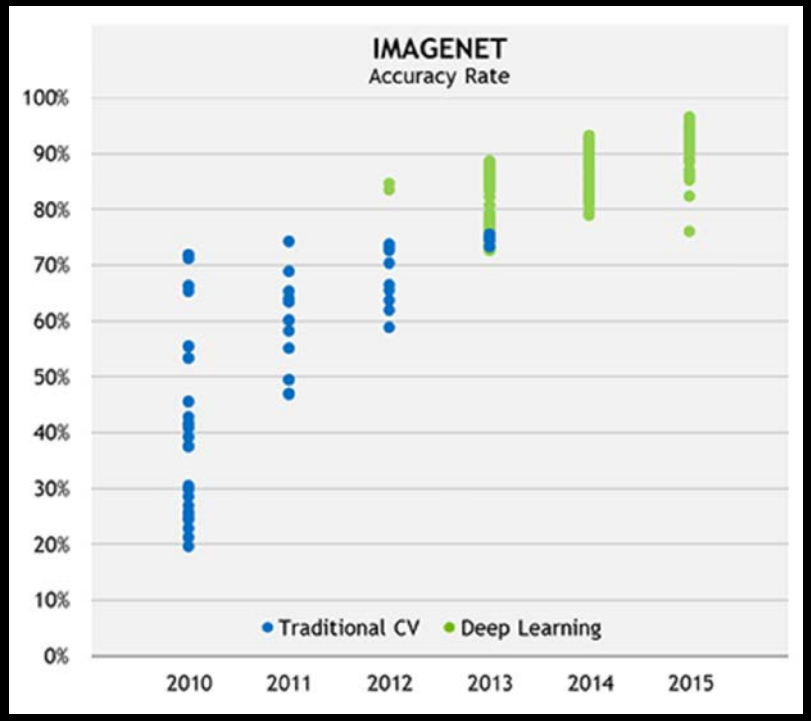

CNNs have revolutionized the way we do image classification.

FlareNet

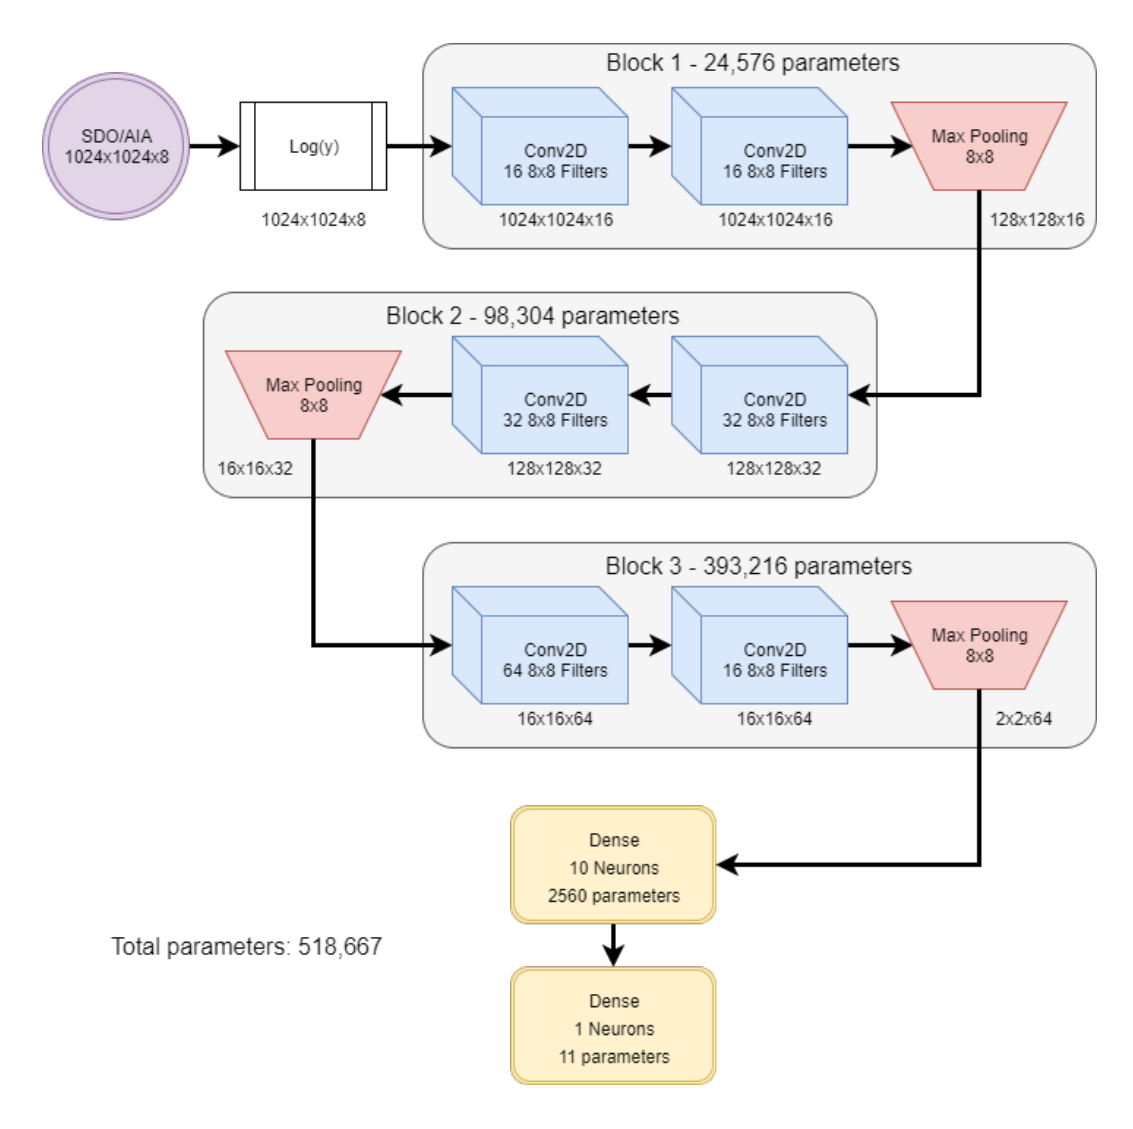
SDO/AIA
1024x1024x8
Log(y)
1024x1024x8
Block 1 - 24,576 parameters
Conv2D
16 8x8 Filters
1024x1024x16
Conv2D
16 8x8 Filters
1024x1024x16
Max Pooling
8x8
128x128x16
Block 2 - 98,304 parameters
Conv2D
32 8x8 Filters
128x128x32
Conv2D
32 8x8 Filters
128x128x32
Max Pooling
8x8
16x16x32
Block 3 - 393,216 parameters
Conv2D
64 8x8 Filters
16x16x64
Conv2D
16 8x8 Filters
16x16x64
Max Pooling
8x8
2x2x64
Dense
10 Neurons
2560 parameters
Dense
1 Neurons
11 parameters
Total parameters: 518,667

Input and output data

Input data: SDO/AIA

Hotter and Higher

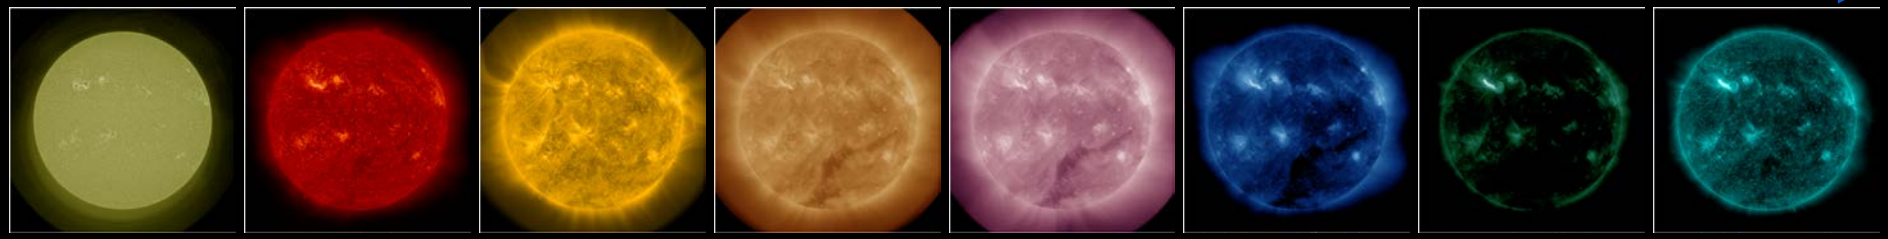

Photosphere | Chromosphere | Corona | Hot Corona

Output data: Goes X-Ray flux

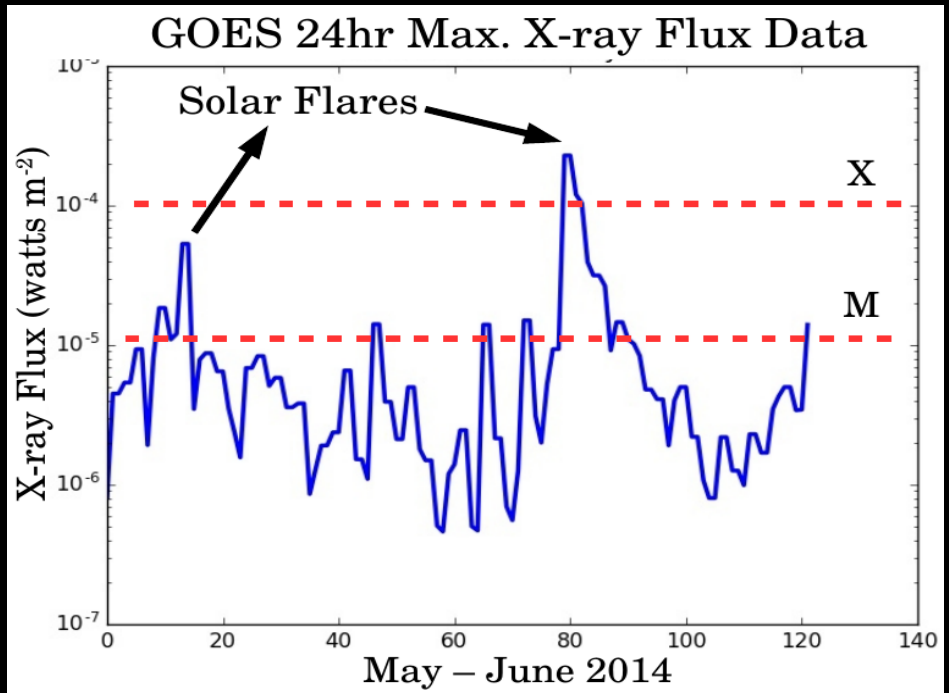

Training, Validation and Test

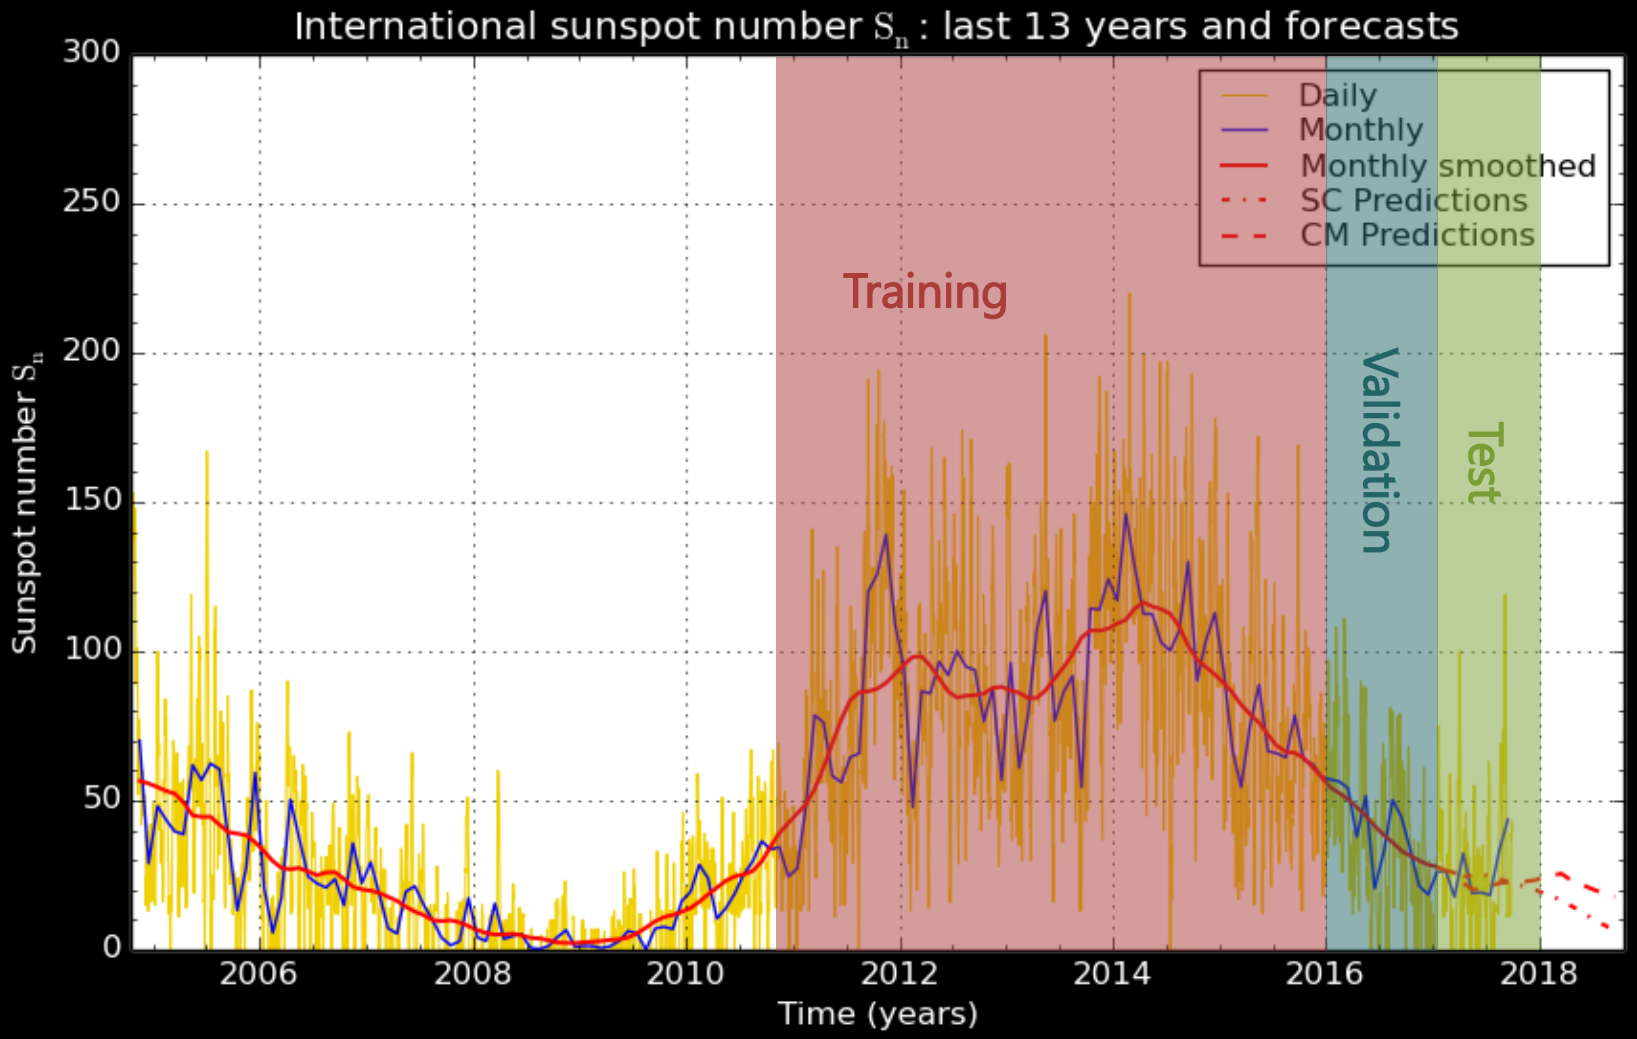

Results: Expectations vs. Reality

Expectations

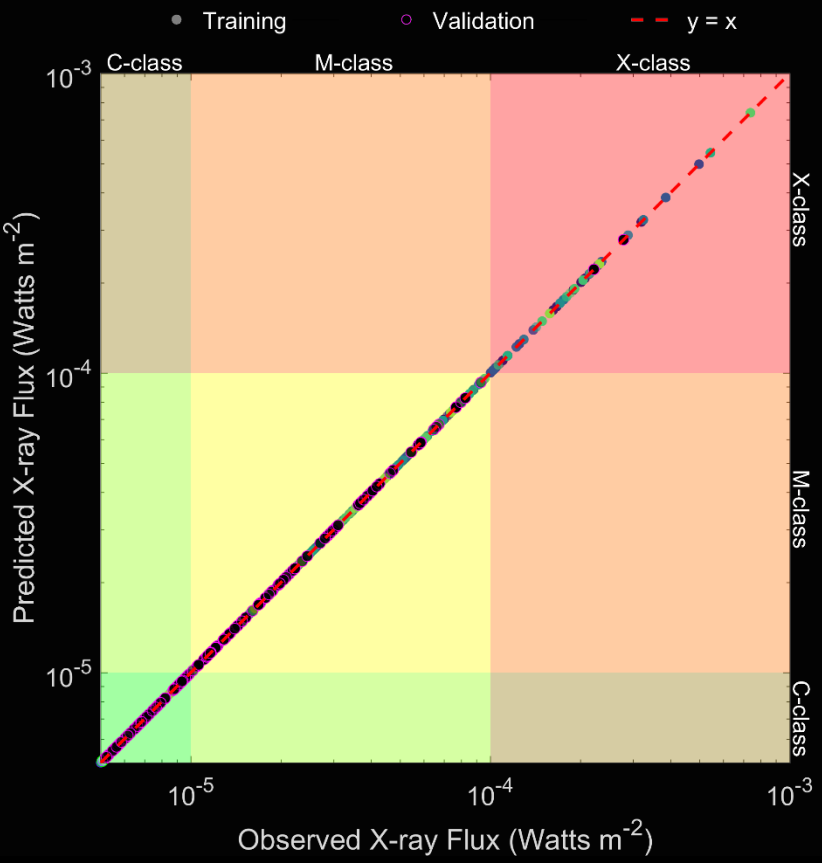

Reality

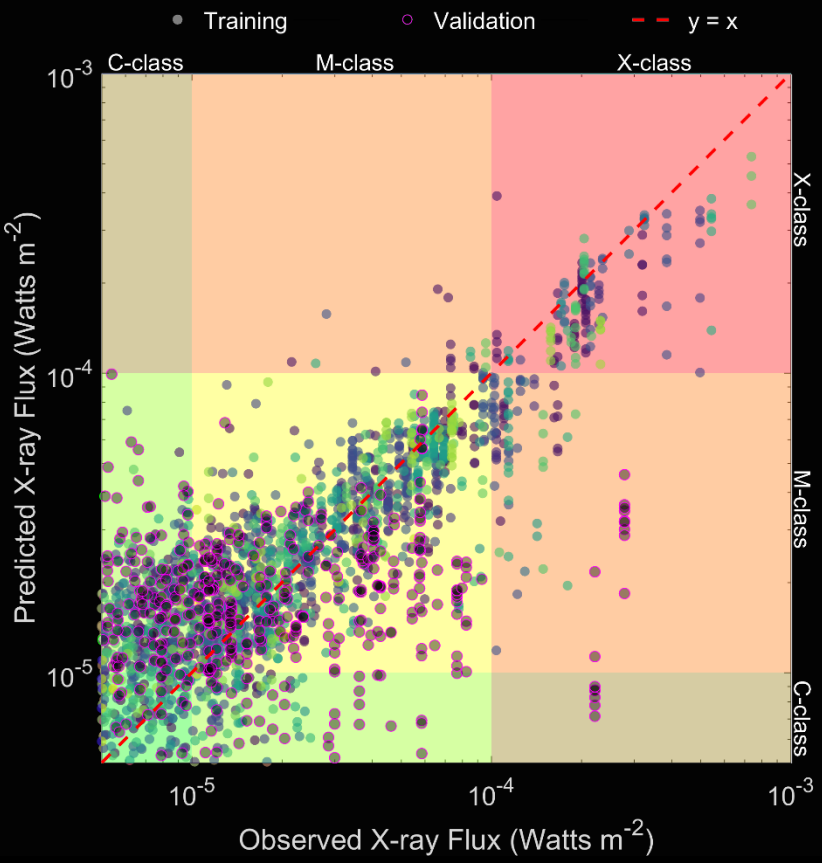

We are achieving some generalization, but finding challenging to make a clean connection between images and flares due to class imbalance

Results: Expectations vs. Reality

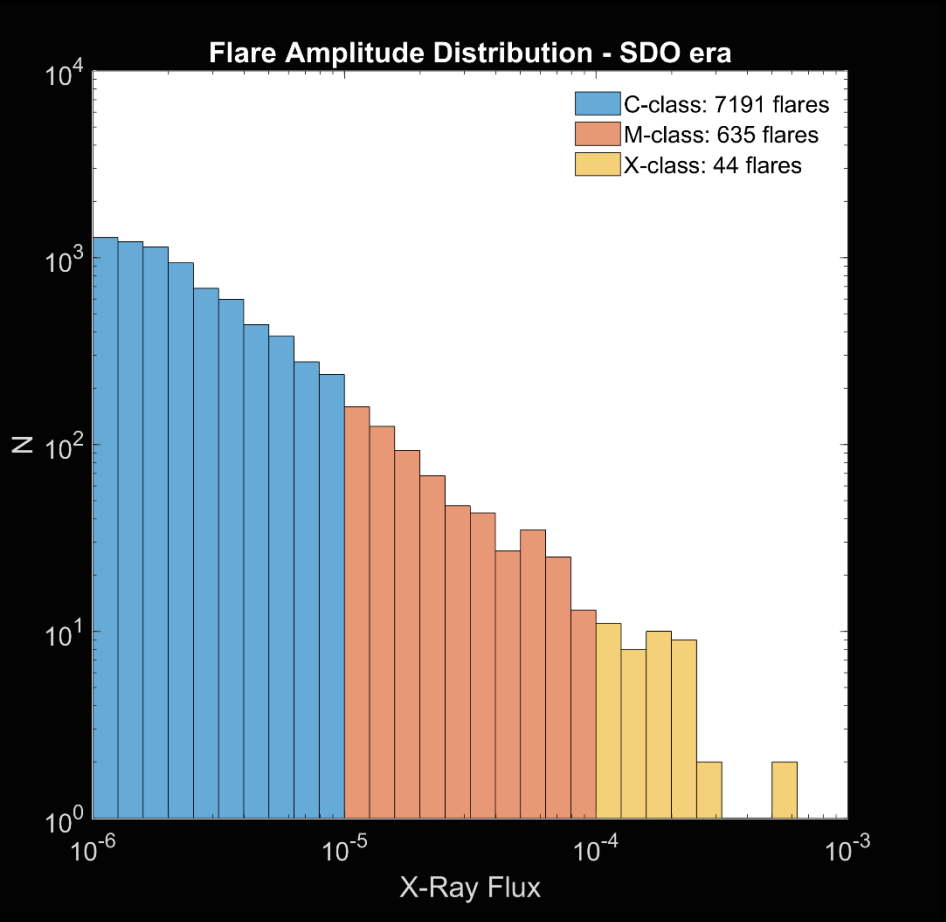

We are achieving some generalization, but finding challenging to make a clean connection between images and flares due to class imbalance

Results: Expectations vs. Reality

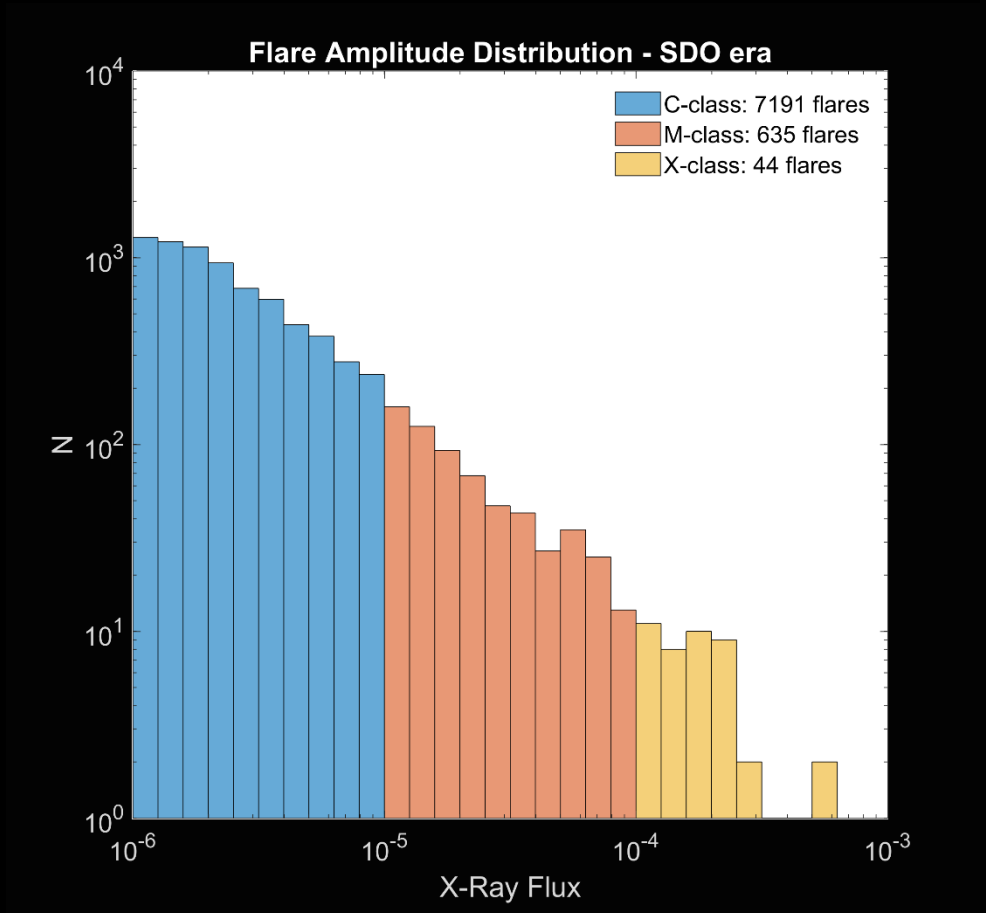

Data shouldn‘t be only big,
it should also be long!

Grad-CAM: Why did you say that?
Visual Explanations from Deep Networks via Gradient-based Localization

Ramprasaath R. Selvaraju Abhishek Das Ramakrishna Vedantam Michael Cogswell
Devi Parikh Dhruv Batra
Virginia Tech
{ram21, abhshkdz, vrama91, cogswell, parikh, dbatra}@vt.edu

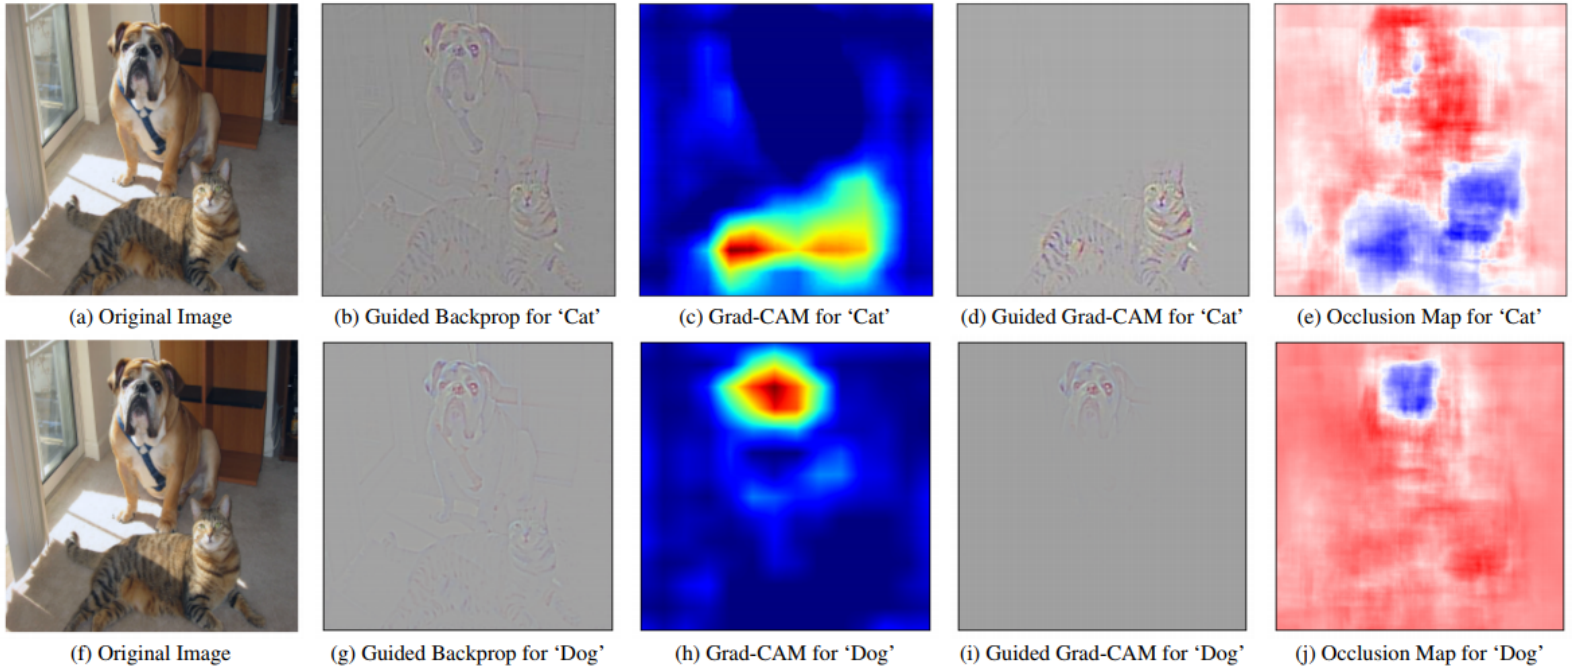

(a) Original Image (b) Guided Backprop for 'Cat' (c) Grad-CAM for 'Cat' (d) Guided Grad-CAM for 'Cat' (e) Occlusion Map for 'Cat'

(f) Original Image (g) Guided Backprop for 'Dog' (h) Grad-CAM for 'Dog' (i) Guided Grad-CAM for 'Dog' (j) Occlusion Map for 'Dog'

Figure 1. (a) Original image with a cat and a dog. (b-e) Support for the cat category according to various visualizations. (b) Guided Backpropagation [38]: provides high-resolution visualization of contributing features, (c) Grad-CAM (Ours): localizes class-discriminative regions, (d) Combining (b) and (c) gives Guided Grad-CAM, which gives high-resolution visualizations that are class-discriminative. Interestingly, the localizations achieved by our Grad-CAM technique (c) are very similar to results from occlusion sensitivity (e), while being much cheaper to compute. Note that in (e), blue corresponds to evidence *for* the class while in (c) blue indicates regions with low score for the class. Figure best viewed in color.

MINING NEURAL NETWORKS

Mining Neural Networks: Saliency

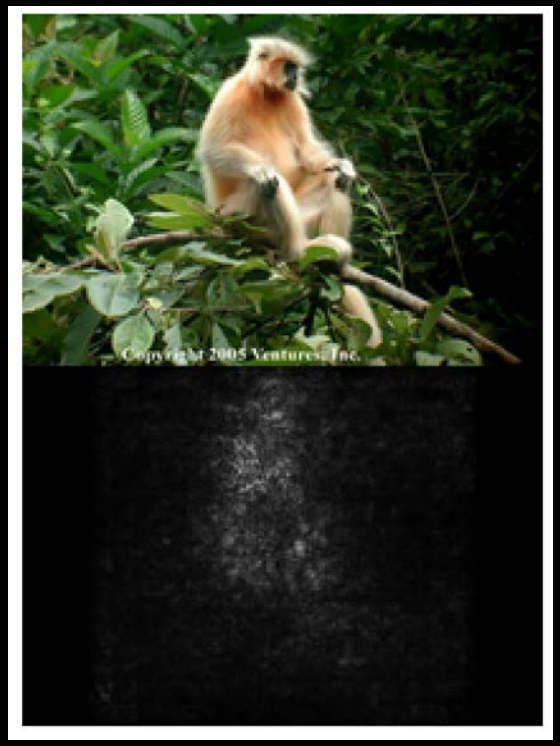

What does a convolutional neural network pay attention to?

Simonyan, K., Vedaldi, A., & Zisserman, A. (2013). Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps.

Mining Neural Networks: Saliency

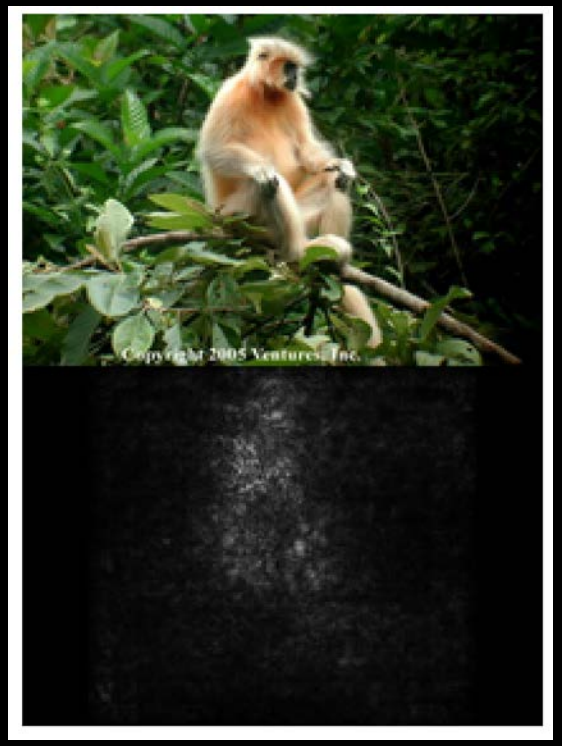

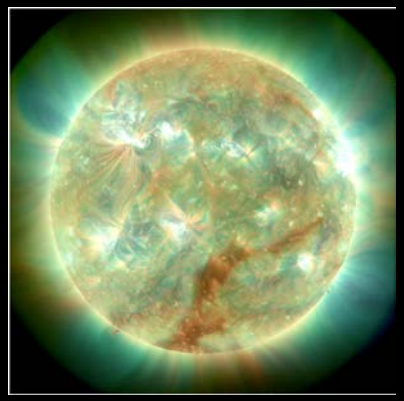

?

Simonyan, K., Vedaldi, A., & Zisserman, A. (2013). Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps.

Mining Neural Networks: Saliency

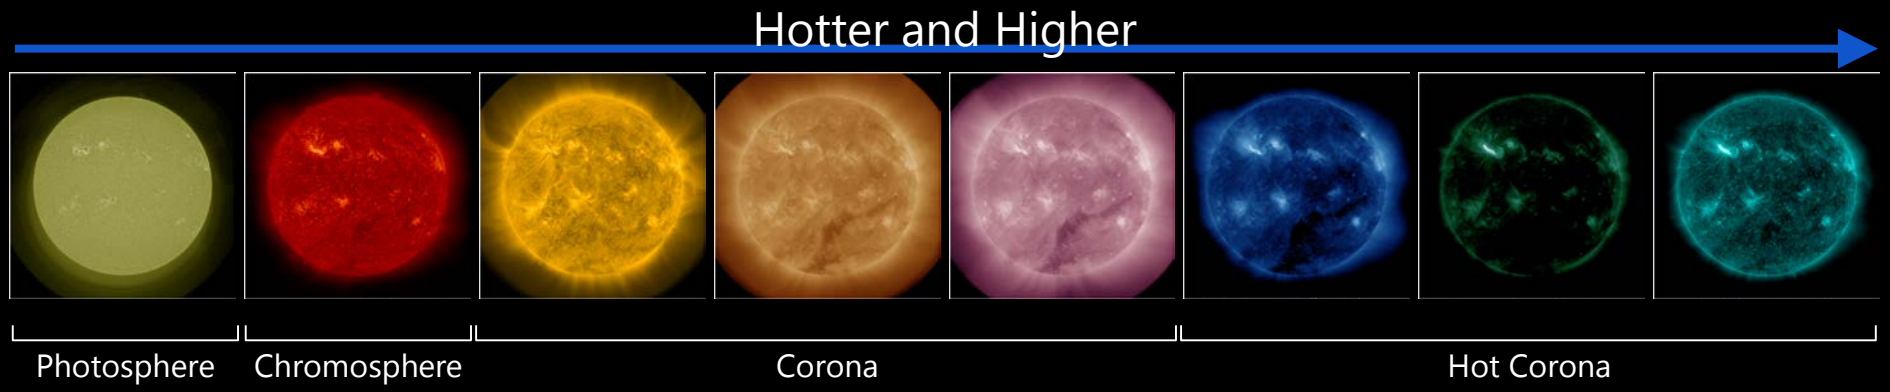

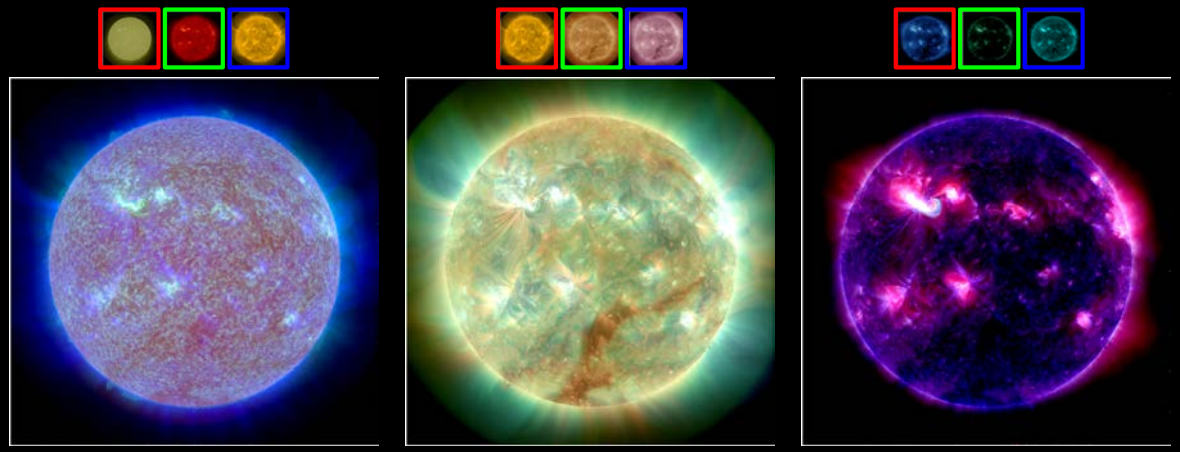

Mining Neural Networks: Saliency

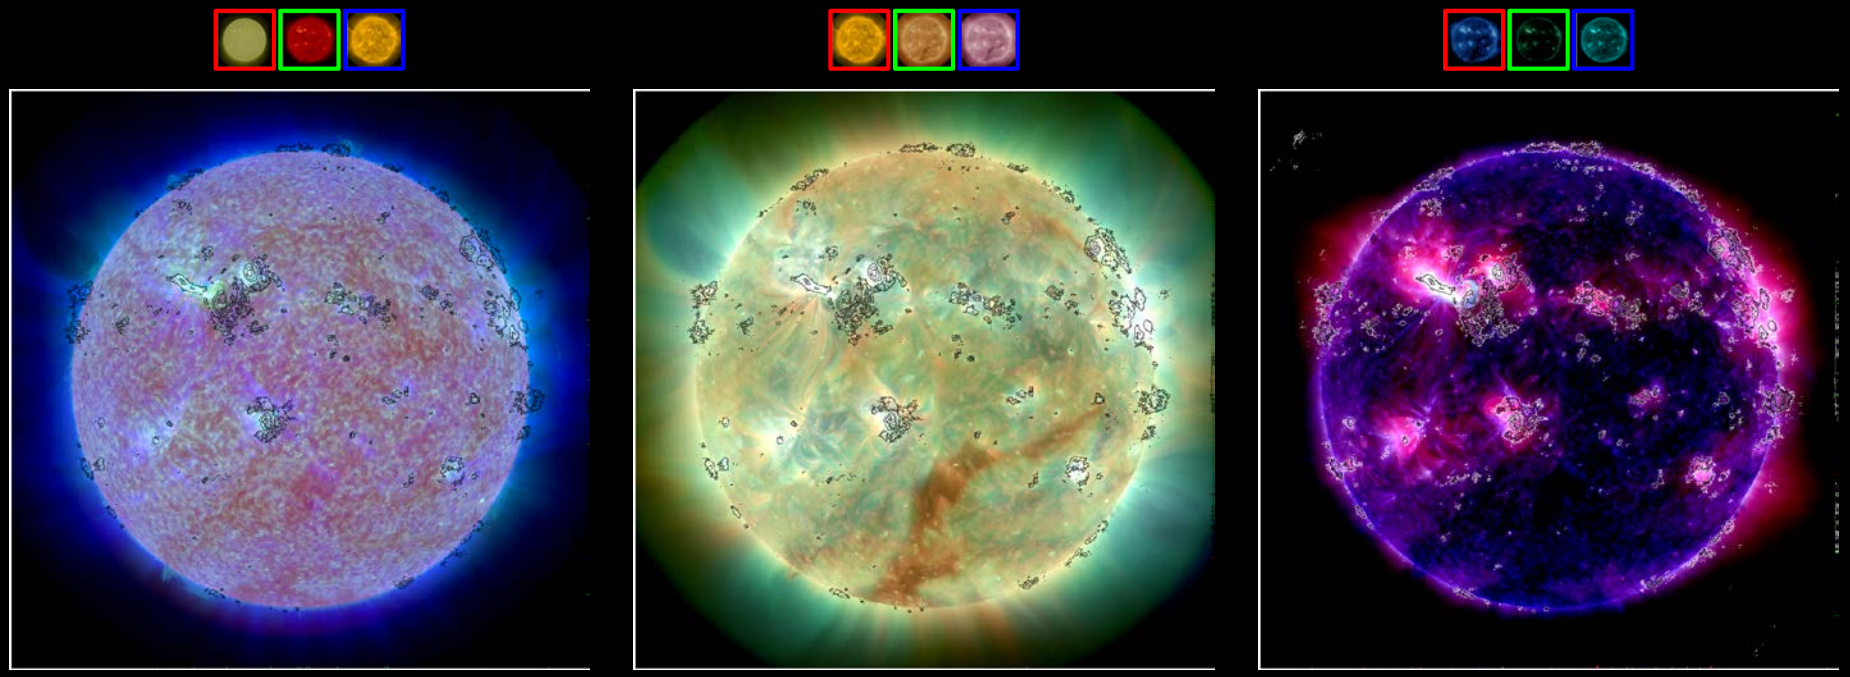

Mining Neural Networks: Saliency

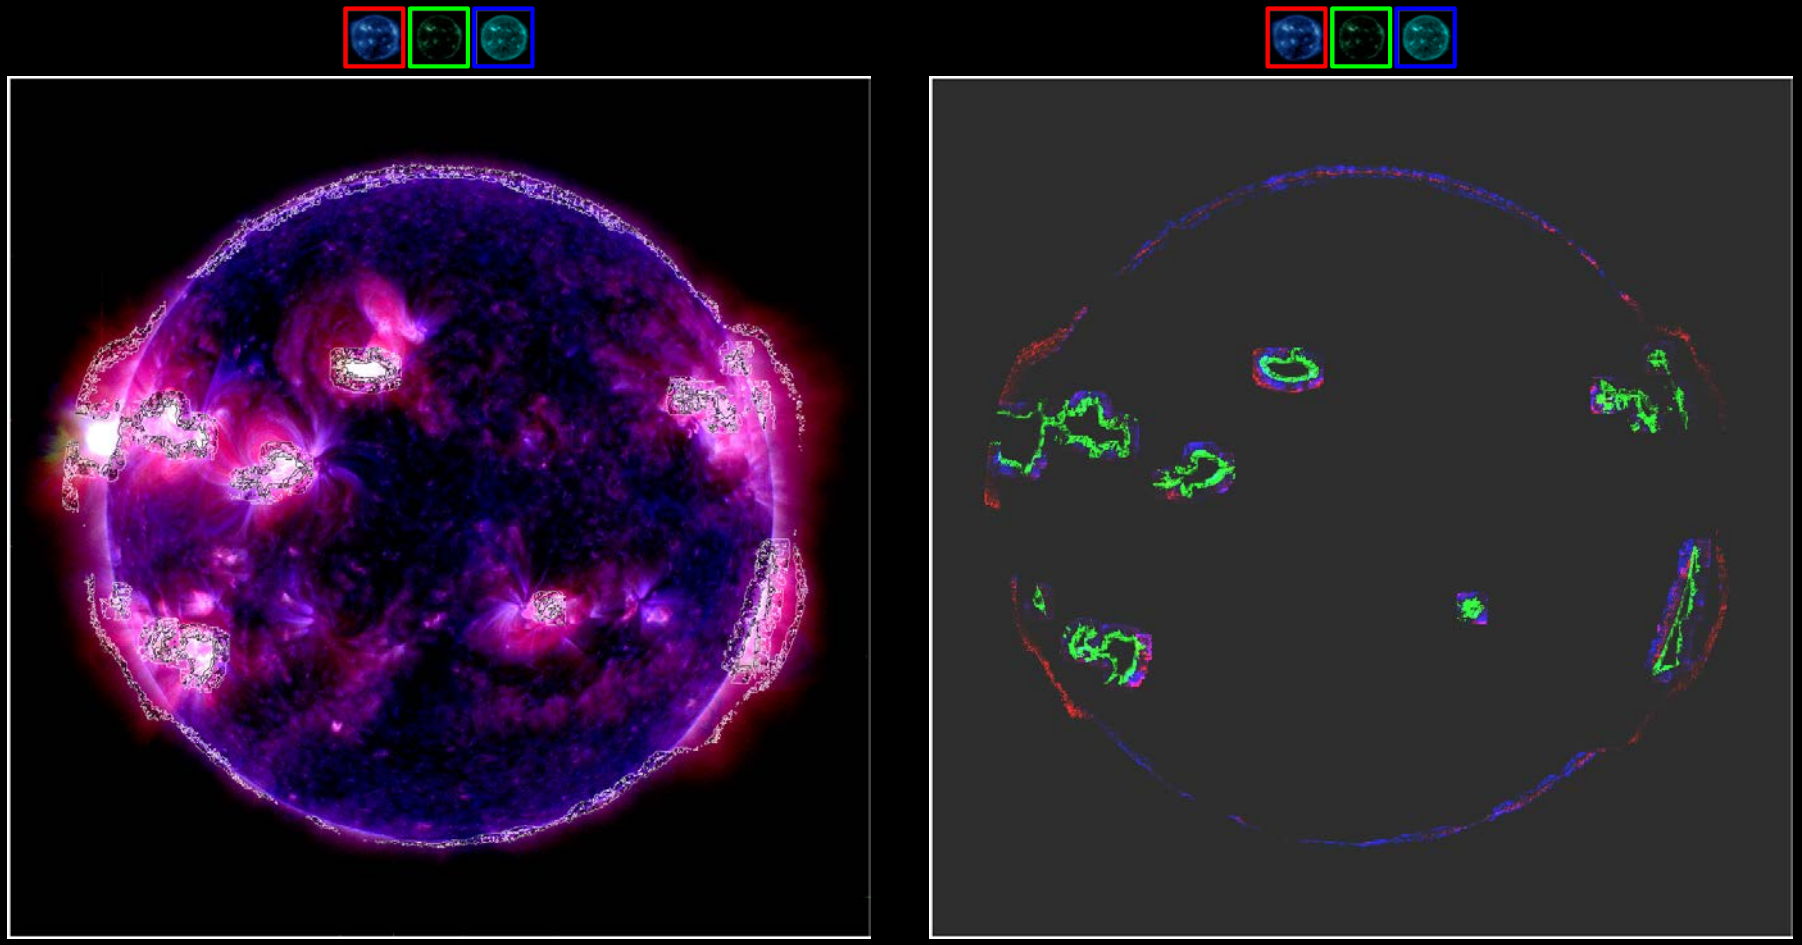

FlareNet is paying attention to the relative location of structures in different channels

Mining Neural Networks: Filter Activation

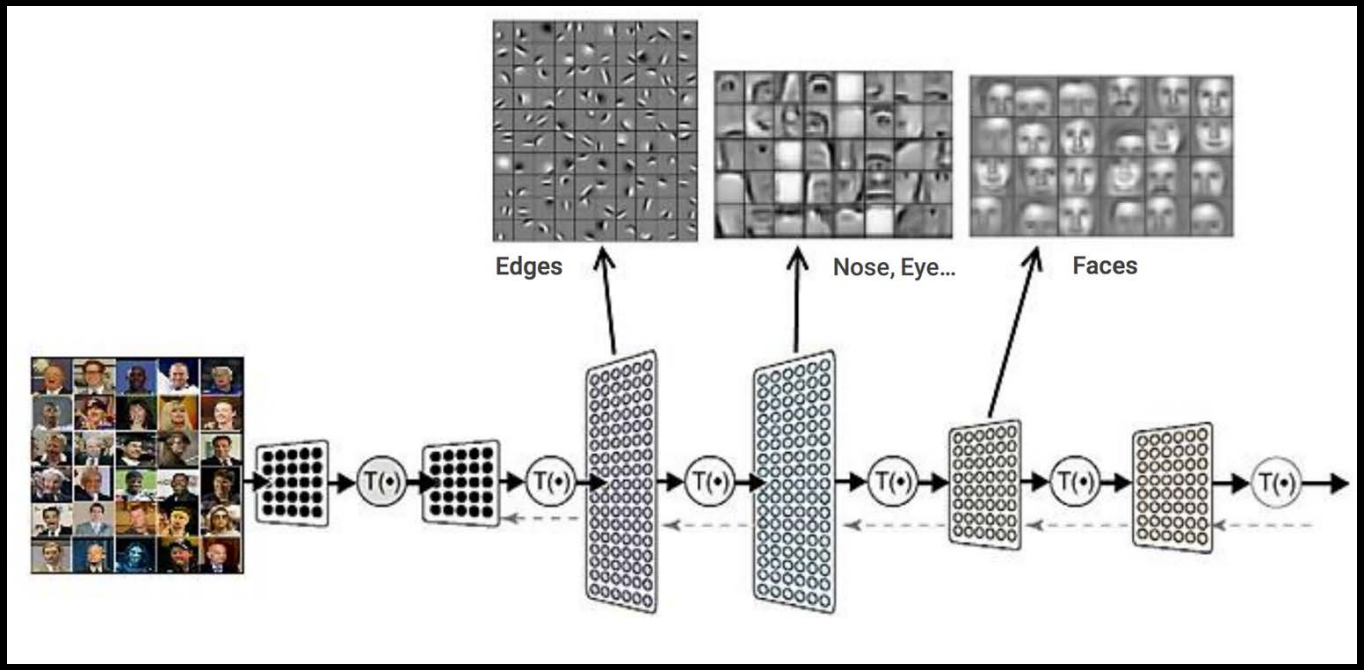

Several convolutional layers allow the neural network to recognize features of increased complexity

Mining Neural Networks: Filter Activation

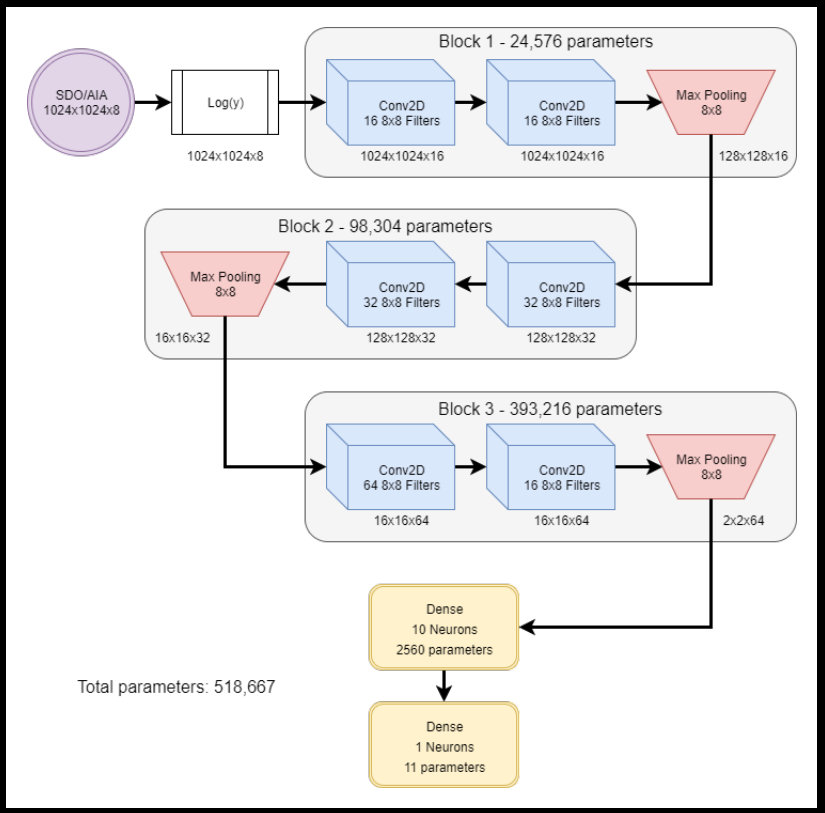

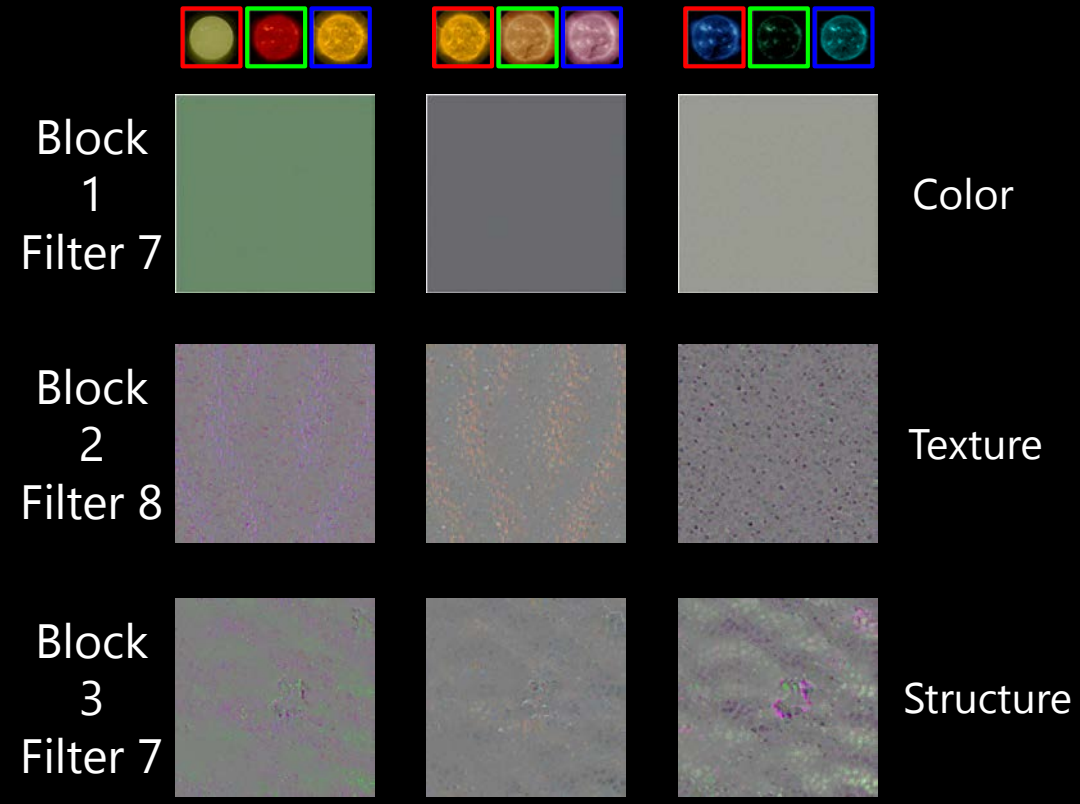

Mining Neural Networks: Filter Activation

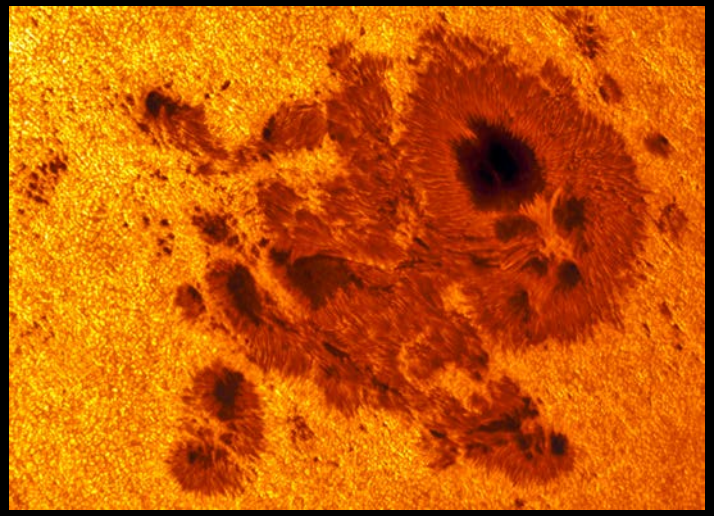

FlareNet learned the importance of active regions

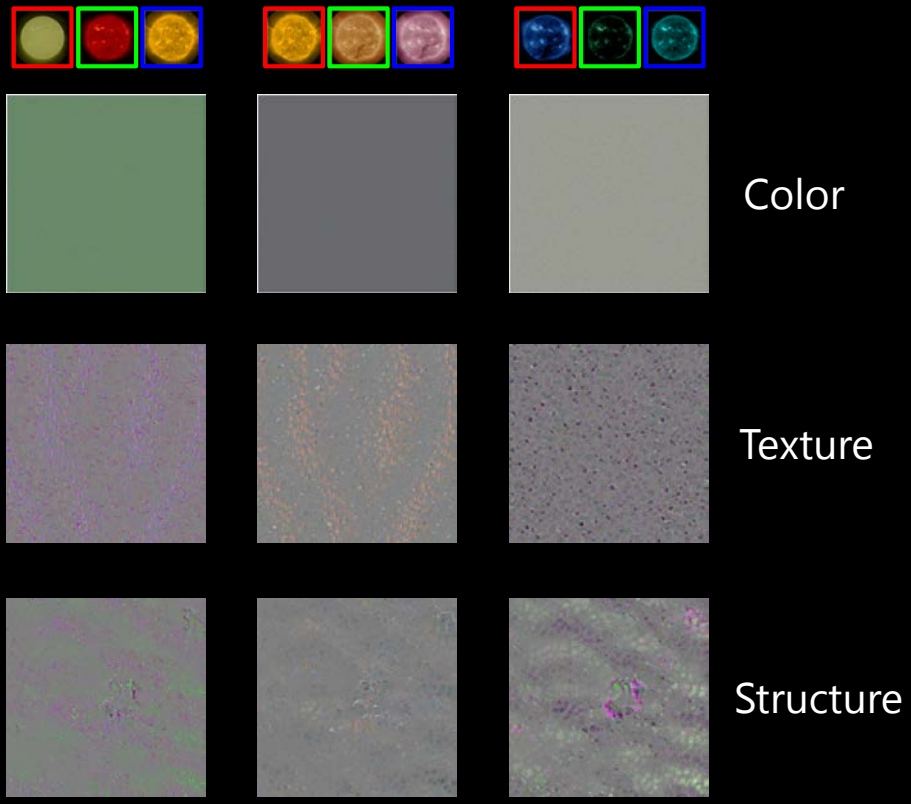

What is deep learning?

What techies think it is

What advocates think it is

What it really is

What skeptics think it is

What the public thinks it is

What is deep learning?

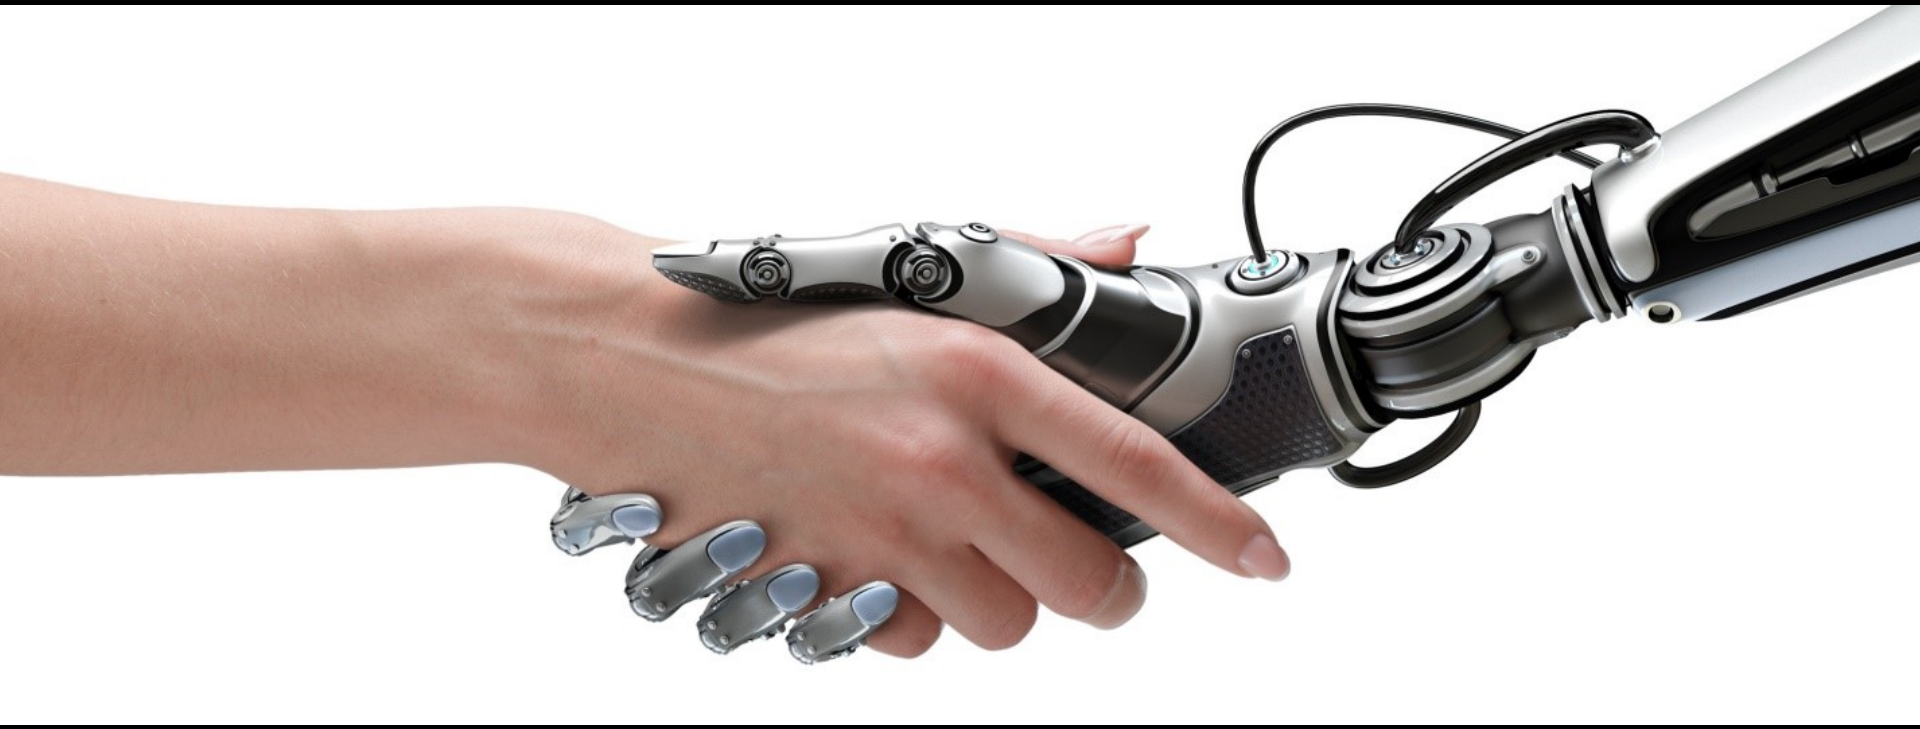

A revolutionary technique that will transform the way we interact with data (big or small)