ML4CC: Lecture 3

Sit with your discussion groups (same as last time)!

Assignments reminder

Keep doing your weekly PMIRO+Q

Your first coding assignment is due before the start of class on Feb 15.

Recap of previous paper

P: Need to be able to predict the energy consumption of commercial buildings based on their features

M: A wide variety of linear and nonlinear regression techniques applied to both “common” and “extended” CBECS features

I: Showing how these techniques can be applied outside the CBECS data (e.g. LL84 and Atlanta) and using feature importance to see which extended features might be worth collecting broadly

R: XGBoost performs best, and there is some benefit to including extended features.

O: May not generalize to other countries, still some bias in XGBoost errors and errors too large for individual buildings

Climate Change in the News

Los Angeles Times

CALIFORNIA

Mudslides, drowned highways, upended homes: Scenes from Southern California's atmospheric river

Clarence Brown assesses damage at his cousin's house caused by a landslide in Studio City that was triggered by heavy rains from the atmospheric river hitting Southern California. (Carlin Stiehl / For The Times)

The record-breaking deluge — which prompted a state of emergency declaration from Gov. Gavin Newsom — triggered mudslides and evacuations, damaged houses, flooded roadways and knocked out power for thousands of people.

In Northern California, three deaths, all from fallen trees, were attributed to the storm, officials said. One was in Santa Cruz County, one in Sutter County and one in Sacramento County.

Paper 2 Discussion

Spotlight Talk at NeurIPS - Tackling Climate Change with Machine Learning workshop 2020

Interpretability in Convolutional Neural Networks for Building Damage Classification in Satellite Imagery

Thomas Y. Chen
The Academy for Mathematics, Science, and Engineering
thomaschen7@acm.org

Abstract

Natural disasters ravage the world's cities, valleys, and shores on a regular basis. Deploying precise and efficient computational mechanisms for assessing infrastructure damage is essential to channel resources and minimize the loss of life. Using a dataset that includes labeled pre- and post- disaster satellite imagery, we take a machine learning-based remote sensing approach and train multiple convolutional neural networks (CNNs) to assess building damage on a per-building basis. We present a novel methodology of interpretable deep learning that seeks to explicitly investigate the most useful modalities of information in the training data to create an accurate classification model. We also investigate which loss functions best optimize these models. Our findings include that ordinal-cross entropy loss is the most optimal criterion for optimization to use and that including the type of disaster that caused the damage in combination with pre- and post-disaster training data most accurately predicts the level of damage caused. Further, we make progress in the qualitative representation of which parts of the images that the model is using to predict damage levels, through gradient-weighted class activation mapping (Grad-CAM). Our research seeks to computationally contribute to aiding in this ongoing and growing humanitarian crisis, heightened by anthropogenic climate change.

Attendance

Select one person from the group to go to this Google Doc and write down the names of all people present in the group (remember to mark who took attendance!)

https://docs.google.com/document/d/1PKhw9E2lJpAnFrFO88DOc2rscZFVlcIv47QNa5h1sGs/edit?usp=sharing (link is in Brightspace under Syllabus content)

Discussion Question 1

What dataset did the author use and what are two positive features of the dataset?

xBD: A Dataset for Assessing Building Damage from Satellite Imagery

Ritwik Gupta[1,2] Richard Hosfelt[1,2] Sandra Sajeev[1,2] Nirav Patel[3,4] Bryce Goodman[3,4]
Jigar Doshi[5] Eric Heim[1,2] Howie Choset[1] Matthew Gaston[1,2]
[1]Carnegie Mellon University [2]Software Engineering Institute [3]Defense Innovation Unit
[4]Department of Defense [5]CrowdAI, Inc.

Figure 2: Pre-disaster imagery (top) and post-disaster imagery (bottom). From left to right: Hurricane Harvey; Joplin tornado; Lower Puna volcanic eruption; Sunda Strait tsunami. Imagery from DigitalGlobe.

3.1. Multiple Levels of Damage

After discussions with disaster response experts from CAL FIRE and the California Air National Guard, it was clear that agencies did not currently have the capacity to classify multiple levels of damage. Many analysis centers simply label buildings as "damaged" or "undamaged" to reduce the amount of expert man-hours needed for assessment, even though it was clear that damage is not a binary status. Discerning between multiple levels of damage is a critical mission need, therefore xBD needed to represent a continuum of damage.

3.2. Image Resolution

Differences between levels of damage are often visually minute. To facilitate the labeling of these types of damage, supporting imagery must be of high fidelity and have enough discerning visual cues. We targeted satellite imagery to be below a 0.8 meter ground sample distance (GSD) mark to fulfill this requirement.

3.3. Diversity of Disasters

One goal of the xView 2 prize challenge is to output models that are widely applicable across a large number of disasters. This will enable multiple disaster response agencies to potentially reduce their workload by using one model with a known deployment cycle. xBD would need to be representative of multiple disaster types and not simply

Discussion Question 2

Explain what the author is saying they are doing here, why they are doing it, and whether or not you think it is a good idea.

“The dataset consists of 1024 by 1024 pixel satellite images....We discard buildings that have a bounding box size of less than 2,000 pixels, as they are too small and blurred to be valuable training data, possibly hindering the model from achieving accurate results”

Discarding small buildings

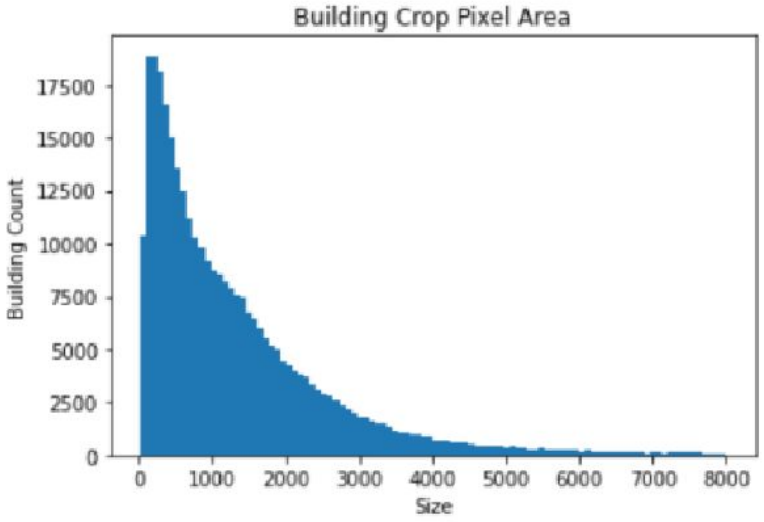

area of 2000 pixels = ~44x44 pixels

or .2% of full 1024x1024 pixel image

While it is true that classifying the damage level of these small buildings would be difficult, they form a large fraction of the dataset, suggesting it is crucial to classify them.

Discussion Question 3

Explain what the author is saying they are doing here, why they are doing it, and whether or not you think it is a good idea.

In order to maintain an equal distribution over JDS classification (damage level) in our training and validation sets so that we can properly assess model accuracy, we provide for an equal number of buildings of the categories "destroyed," "major damage," "minor damage," and "no damage" in each set, while still maintaining a 0.8:0.2 ratio between train and validation. The xBD dataset is deliberately created with a disproportionately large volume of buildings with no damage [10], but training on such a lopsided data distribution would yield artificially high accuracy numbers and not yield valuable results.

The original dataset is heavily imbalanced

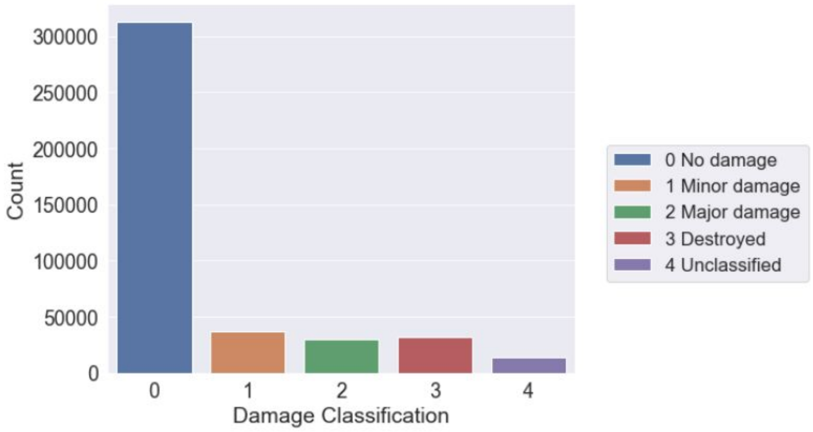

Figure 9: Distribution of damage class labels.

Accuracy can be high on imbalanced data just by chance.

Resampling the data to have balanced numbers in each class solves that problem, but the resampled data no longer represents the original problem.

There are other solutions to this problem...

Discussion Question 4

What are the three loss functions tested here and how do they differ?

Loss functions

Basic cross entropy

Captures classification performance, but treats each category as **totally separate**

model uses the cross-entropy loss function, which is defined as

$$-\sum_{c=1}^{4} y_{o,c} \log(p_{o,c}),$$

*model outputs a vector

where $y_{o,c}$ is a binary indicator (either 0 or 1) of whether c, as a label, correctly classifies observation o, and $p_{o,c}$ is the predicted probability that observation o is of the class c. Cross-entropy loss is defined, in other terms, as the negative sum of the expression $y_{o,c}\log(p_{o,c})$ across all 4 possible classes c: no damage, minor damage, major damage, and destroyed. The network is trained on 12,800

$$H(P^* | P) = -\sum_i \underbrace{P^*(i)}_{\text{TRUE CLASS DISTIRBUTION}} \log \underbrace{P(i)}_{\text{PREDICTED CLASS DISTIRBUTION}}$$

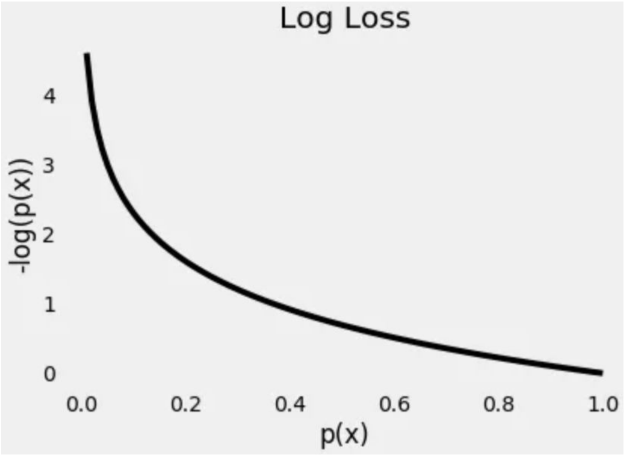

Loss functions

Mean-squared error

We define mean squared error as

$$\frac{1}{b}\sum_{i=1}^{b}(y-\hat{y})^2,$$

*model outputs a scalar

where b is the batch size, y is the ground truth (a class from 0 to 3 representing each damage level), and $\hat{y}$ is the prediction.

Incorporates a more natural relationship between classes, but doesn't account for the fact that true labels are discrete integers.

Loss functions

Ordinal cross entropy loss

level), and $\hat{y}$ is the prediction. Ordinal cross-entropy loss differs from cross-entropy loss in that it takes into account the distance between the ground truth and the predicted class (hence "ordinal"). Since the building damage classification problem involves different and increasing levels of damage from no damage to destruction, this function is useful to distinguish between different categories. To implement ordinal cross-entropy loss as the loss function, we treat it as generic multi-class classification and encode the classes no damage, minor damage, major damage, and destroyed as [0, 0, 0], [1, 0, 0], [1, 1, 0], and [1, 1, 1], respectively [3]. The other aspects of the training process

*model outputs a vector

Allows for categories that have a natural ordering

Discussion Question 5

What are two things the author did to try to understand how these models work (i.e. make the models more “interpretable”)?

Understanding the model by:

-Varying the inputs

Table 1: Comparison of Validation Accuracy on 9 Different Models

Model Accuracy on Validation Set with Chosen Loss (100 epochs)			
	Loss Function		
Model Input	Mean Squared Error	Cross-Entropy Loss	Ordinal Cross-Entropy Loss
Post-Disaster Image Only	45.3%	59.5%	64.2%
Pre-Disaster, Post-Disaster Images	50.2%	68.3%	71.2%
Pre-Disaster, Post-Disaster Images, Disaster Type	49.7%	72.7%	74.6%

Comparing the performance of a model with all available inputs to those with only a subset can help identify the importance of different inputs

...i.e., a feature importance method!

Understanding the model by:

Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization

Ramprasaath R. Selvaraju · Michael Cogswell · Abhishek Das · Ramakrishna Vedantam · Devi Parikh · Dhruv Batra

Uses gradient calculations to identify which units are most important for classifying an image as of a certain class.

Backwardpass

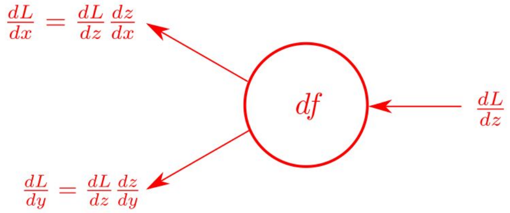

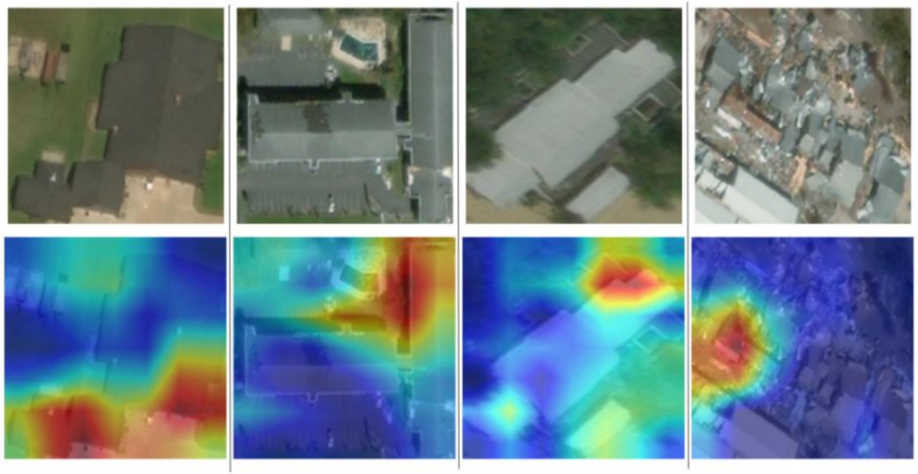

Figure 1: Gradient class activation maps [20] depict which parts of the building crop lead the baseline model to predict a certain classification. On the first row are the original images (crops) and on the second row are the corresponding gradient class activation maps. The images included consist of solely post-disaster images. From left to right: (1) A building with label "no damage," after flooding in the Midwestern United States, (2) A building with label "minor damage," after Hurricane Michael, (3) A building with label "major damage," after Hurricane Harvey, and (4) A building with label "destroyed," after Hurricane Michael.

Discussion Question 6

How specifically were each of the three inputs given to the model?

Inputs

the type of disaster (e.g. volcano, wind, etc.) that caused the building damage. To train a model that takes in both pre-disaster images and their corresponding post-disaster images, we concatenate the RGB channels of the two and use that as input. To train a model that takes in the pre-disaster image, post-disaster image, and disaster type, we do the same, but also concatenate a one-hot encoded representation of the disaster type in one of the later layers of the CNN.

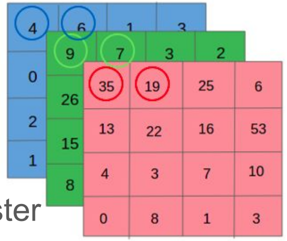

(1) post-disaster image

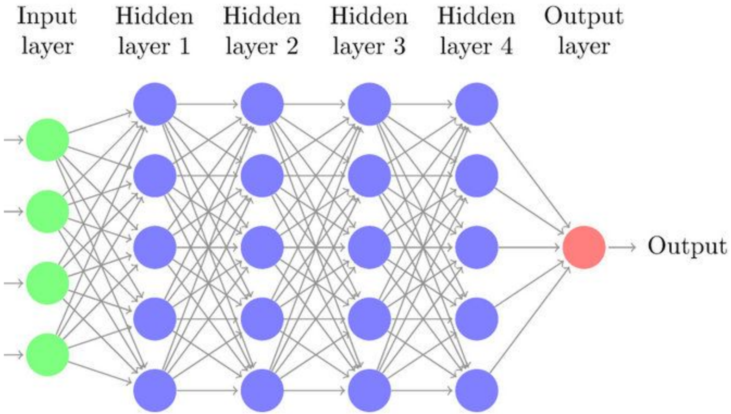

Inputs

the type of disaster (e.g. volcano, wind, etc.) that caused the building damage. To train a model that takes in both pre-disaster images and their corresponding post-disaster images, we concatenate the RGB channels of the two and use that as input. To train a model that takes in the pre-disaster image, post-disaster image, and disaster type, we do the same, but also concatenate a one-hot encoded representation of the disaster type in one of the later layers of the CNN.

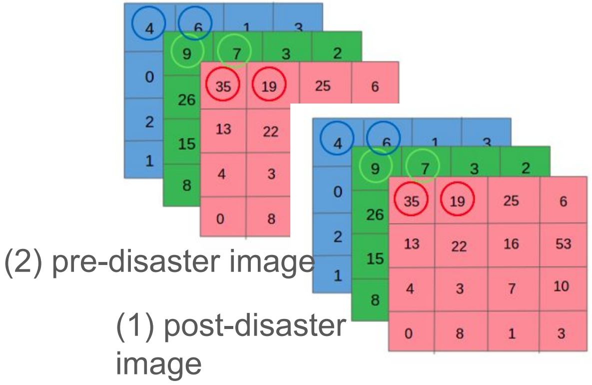

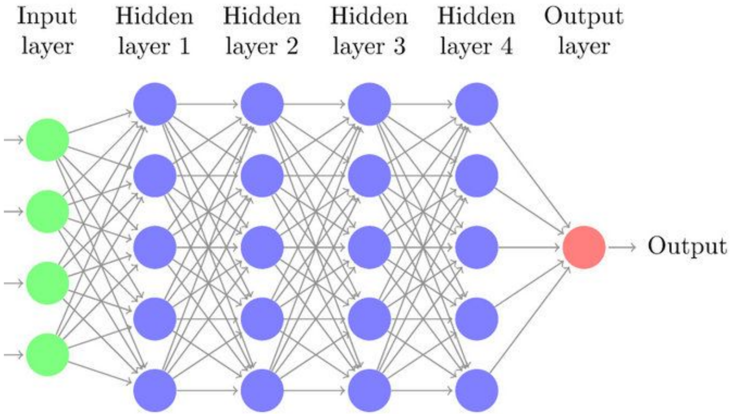

Inputs

the type of disaster (e.g. volcano, wind, etc.) that caused the building damage. To train a model that takes in both pre-disaster images and their corresponding post-disaster images, we concatenate the RGB channels of the two and use that as input. To train a model that takes in the pre-disaster image, post-disaster image, and disaster type, we do the same, but also concatenate a one-hot encoded representation of the disaster type in one of the later layers of the CNN.

(3) disaster type

[0,0,1,...0,0]

(2) pre-disaster image

(1) post-disaster image

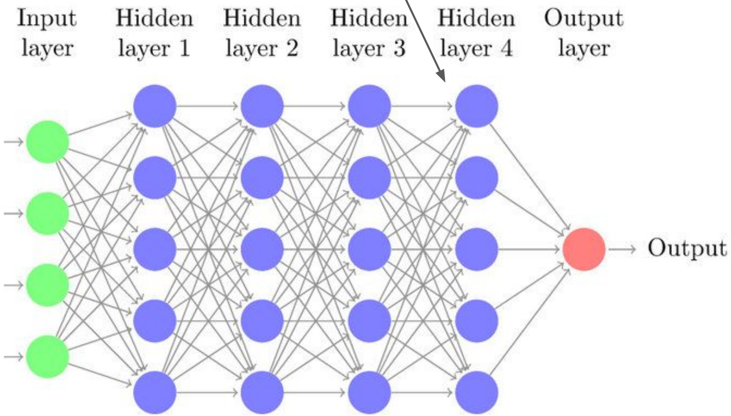

Discussion Question 7

What is notable about the institution that the author is at?

Discussion Question 7

What is notable about the institution that the author is at?

Thomas Y. Chen
The Academy for Mathematics, Science, and Engineering
thomaschen7@acm.org

Academy for Mathematics, Science, and Engineering

Article Talk

From Wikipedia, the free encyclopedia

The **Academy for Mathematics, Science, and Engineering** (**AMSE**) is a four-year magnet public high school program intended to prepare students for STEM careers. Housed on the campus of Morris Hills High School in Rockaway, in the U.S. state of New Jersey, it is a joint endeavor between the Morris County Vocational School District and the Morris Hills Regional District.

6 Acknowledgements

The author thanks Ethan Weber (Massachusetts Institute of Technology) for his mentorship during the ideation, experimental design, and overall research processes. The author also thanks Climate Change AI (CCAI) and the NeurIPS 2020 CCAI Workshop Organizing and Program Committees.

Discussion Question 8

Share what questions you wrote in your PMIRO+Q and decide as a group what you’d like to ask.

Update your PMIRO+Q

Submit a second file to the Brightspace assignment (don’t overwrite the original):

It should:

Update your PMIRO as needed

Answer your own Q

You can be talking with your group during this!

15 min break

Lecture

Climate change content: Impacts of climate change on oceans, Remote Sensing Earth Observation

Machine learning content: Convolutional neural networks, Image segmentation

Essential Climate Variables

Essential variables (EV) are variables known to be critical for observing and monitoring a given facet of the Earth system.

Many fields such as oceanography, climatology, biodiversity studies, and geodiversity have come together to identify these variables.

Having a common set of accurate and sustained measurements with standards for data collection and dissemination ensures the usability of data across multiple platforms and agencies.

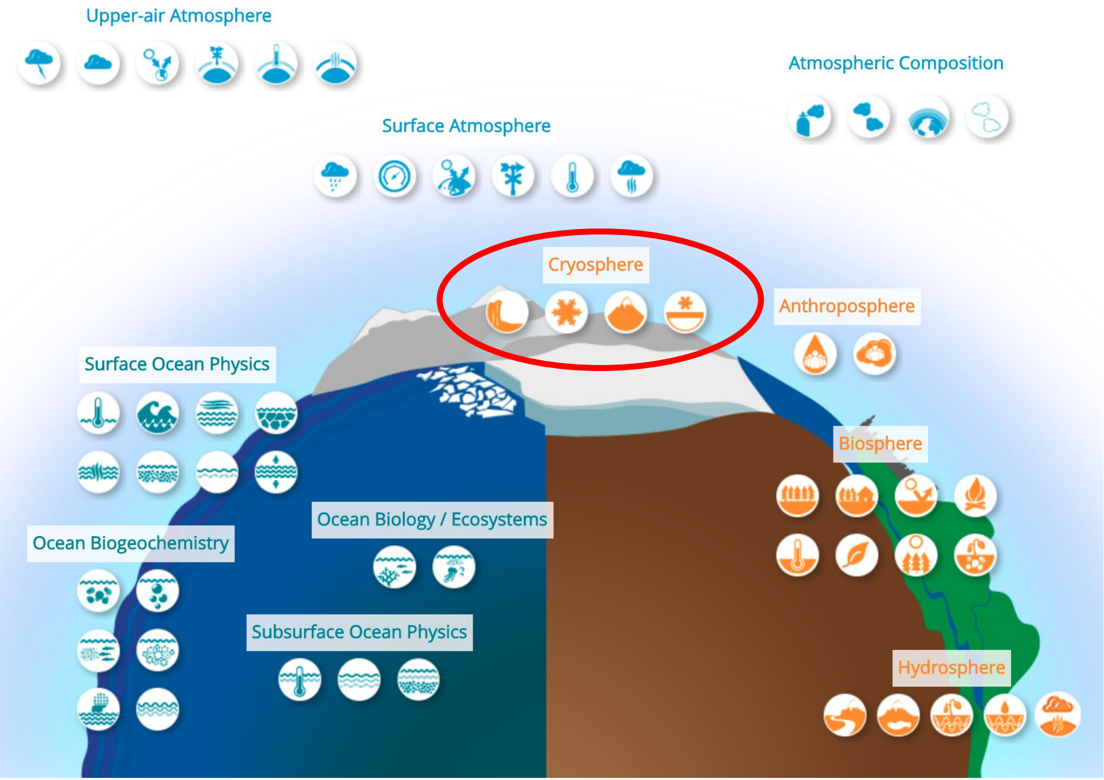

USGS
science for a changing world

FREQUENTLY ASKED QUESTIONS | GEOLOGY

What is a glacier?

A glacier is a large, perennial accumulation of crystalline ice, snow, rock, sediment, and often liquid water that originates on land and moves down slope under the influence of its own weight and gravity. Typically, glaciers exist and may even form in areas where:

1. mean annual temperatures are close to the freezing point
2. winter precipitation produces significant accumulations of snow
3. temperatures throughout the rest of the year do not result in the complete loss of the previous winter's snow accumulation

Over multiple decades this continuing accumulation of snow results in the presence of a large enough mass of snow for the metamorphism from snow to glacier ice process to begin. Glaciers are classified by their size (i.e. ice sheet, ice cap, valley glacier, cirque glacier), location, and thermal regime (i.e., polar vs. temperate). Glaciers are sensitive indicators of changing climate.

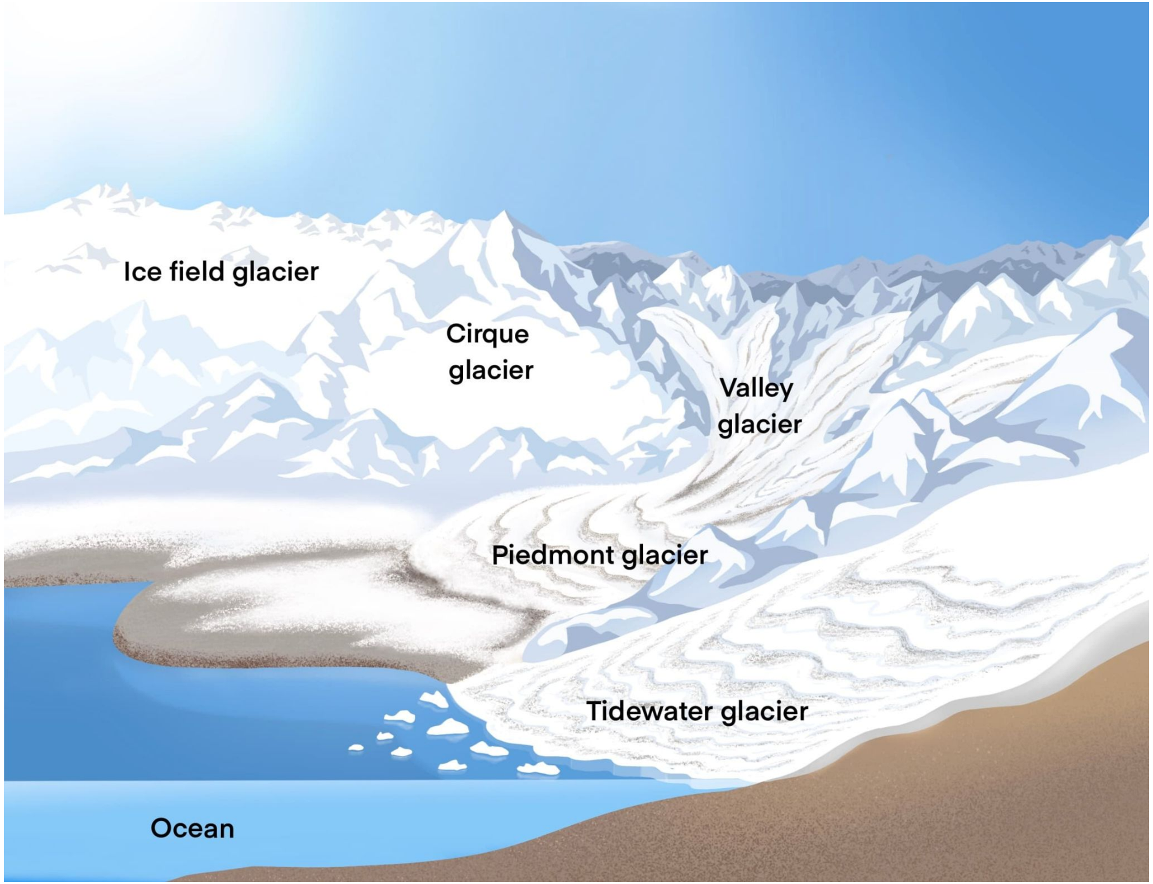

National Snow and Ice Data Center

Cirque glaciers are named for the bowl-like hollows they occupy, which are called cirques. Typically, they are found high on mountainsides and tend to be wide rather than long.

Piedmont glaciers occur when valley glaciers spill into relatively flat plains

How are glaciers impacted by climate change?

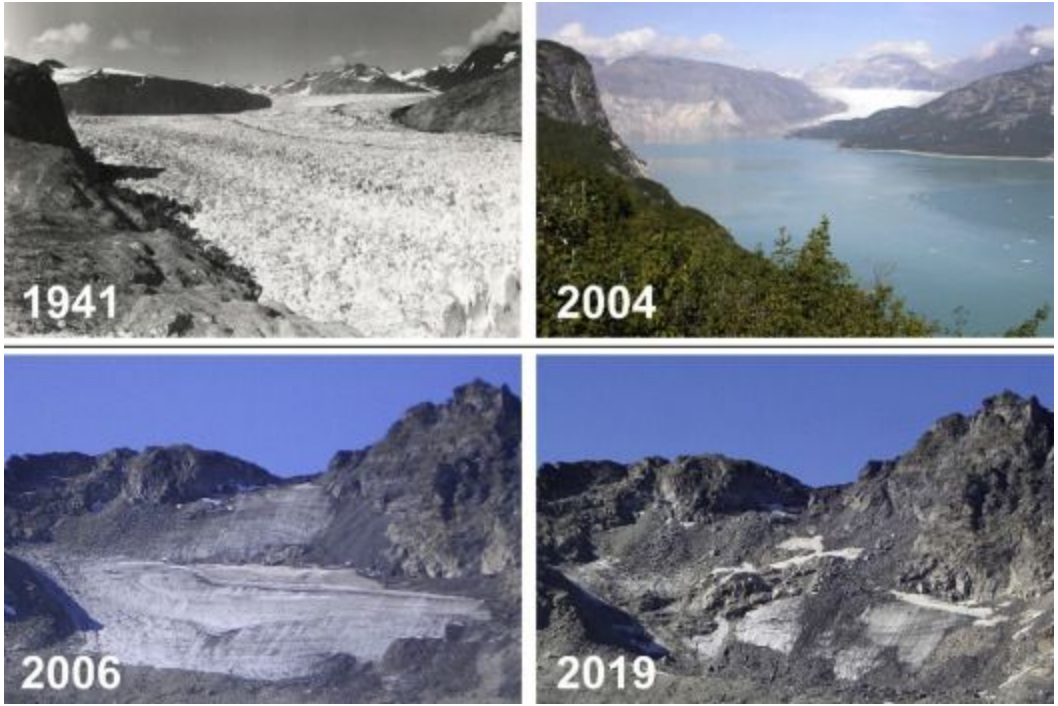

Glaciers retreat and melt in warmer climates

Why does glacier retreat matter?

Glaciers are an important source of freshwater

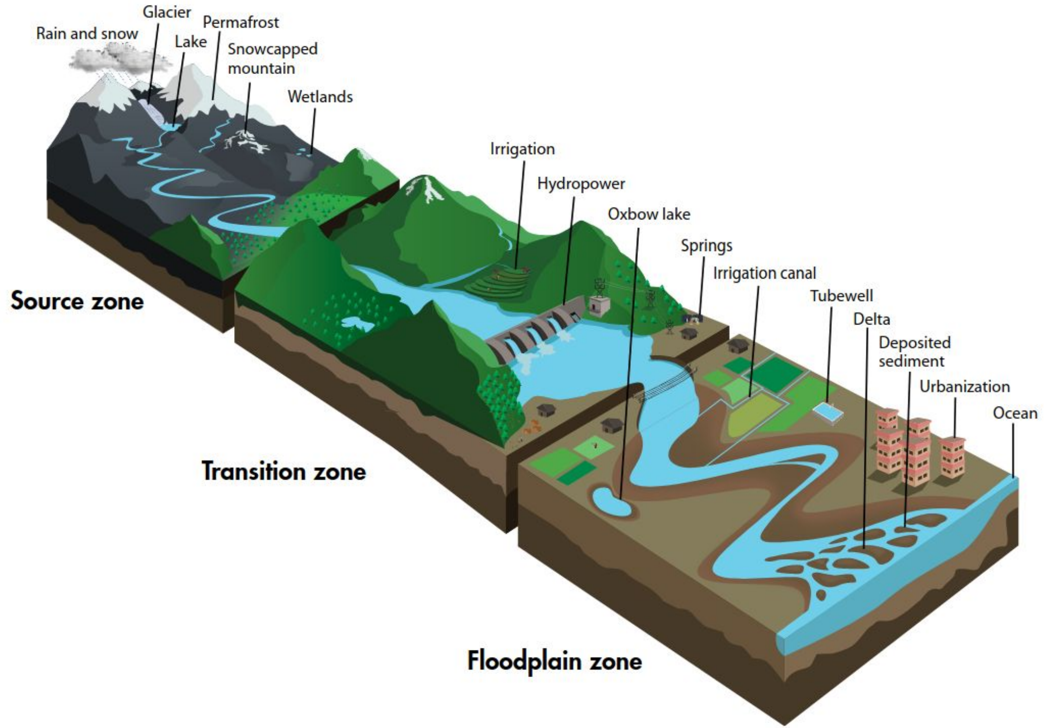

HI-AWARE

Why does glacier retreat matter?

Direct and indirect effects of sea level rise on migration

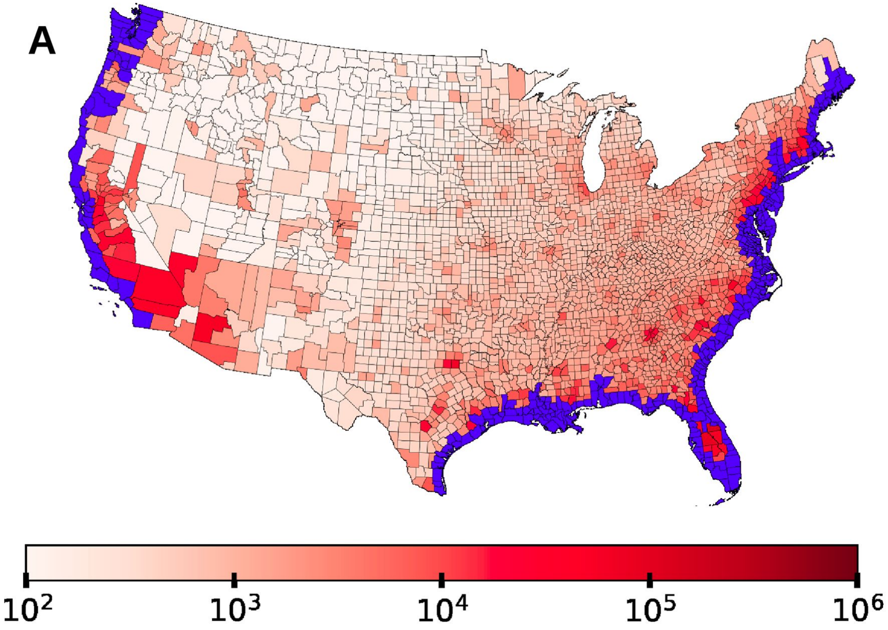

All counties that experience flooding under 1.8m of sea level rise (SLR) by 2100 in blue. Remaining counties are colored based on the number of additional incoming migrants per county that there are in the SLR scenario over the baseline.

OPEN ACCESS PEER-REVIEWED

RESEARCH ARTICLE

Modeling migration patterns in the USA under sea level rise

Caleb Robinson, Bistra Dilkina, Juan Moreno-Cruz

Published: January 22, 2020 • https://doi.org/10.1371/journal.pone.0227436

Feedback loop: Melting glaciers cause more warming

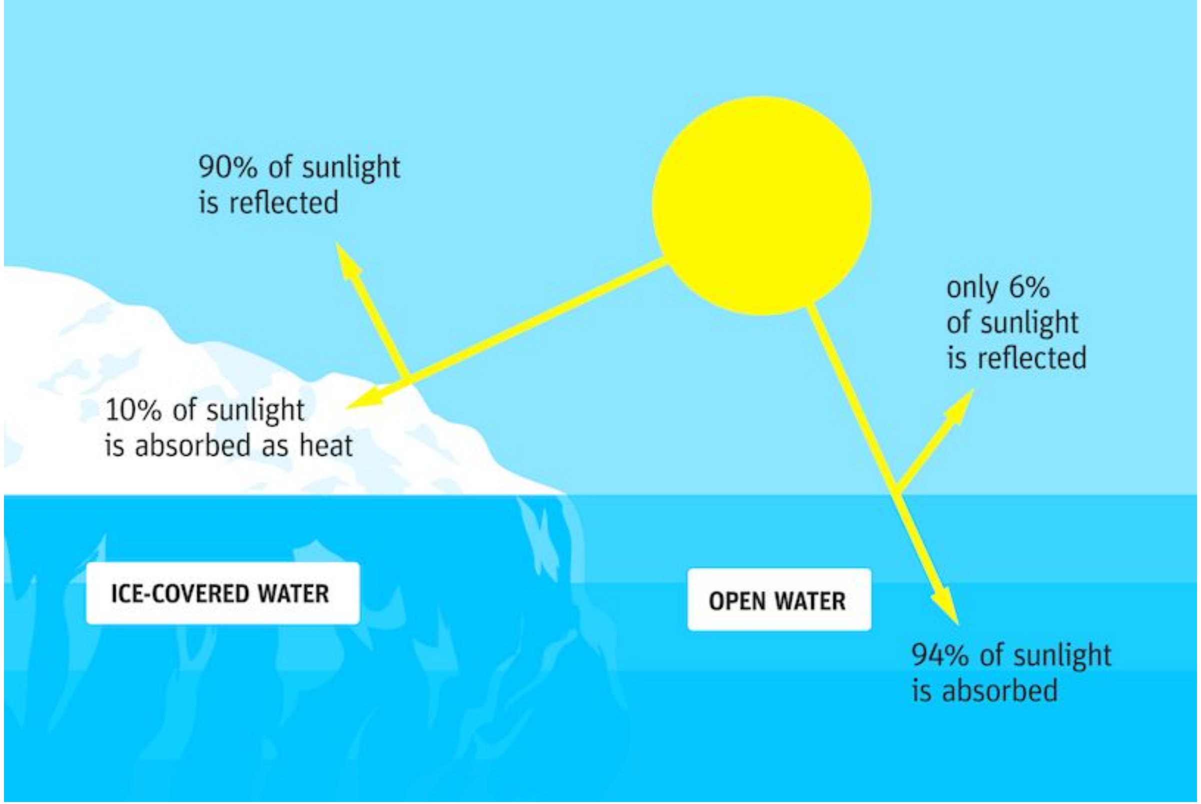

exploratorium.edu

Melting glaciers also impacts ocean currents

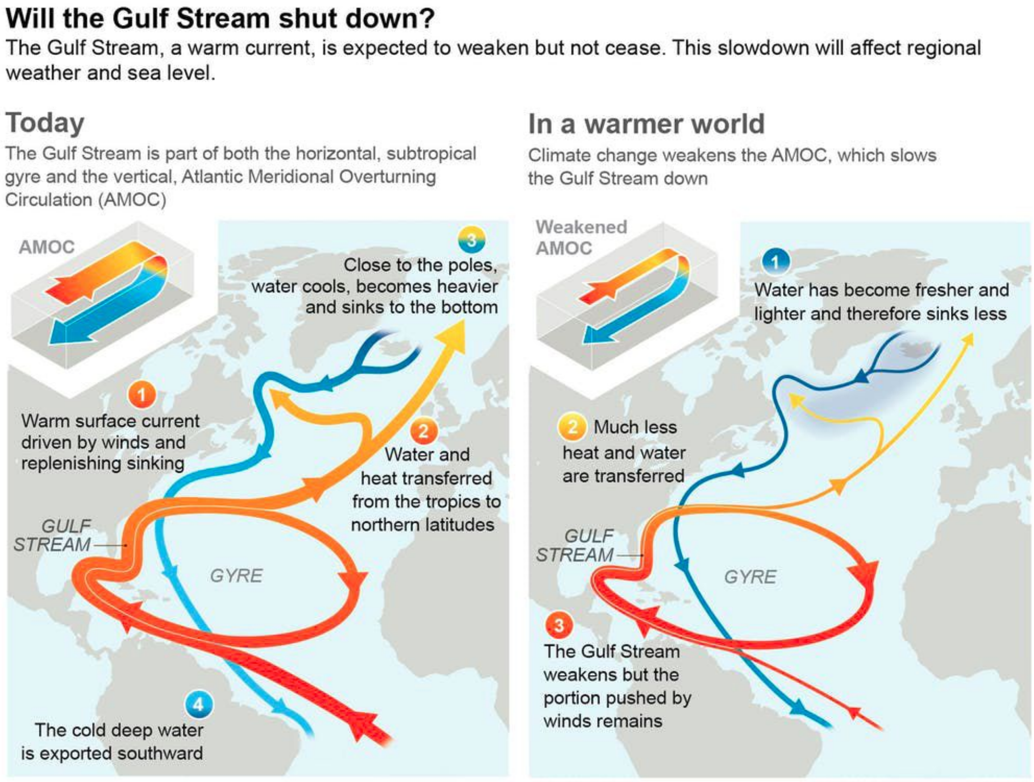

IPCC

Melting glaciers make water fresher (i.e. less salty) and less dense. Therefore it doesn’t sink in the same way as salt water. By weakening this sinking behavior, climate change weakens ocean currents that are driven by such “overturning”

Melting glaciers also impacts ocean currents

If the AMOC shuts down it could cause:

Significant cooling over parts of Europe by as much as 5 or 10 degrees Celsius.

A shift in the position of tropical rain belts

More droughts in some places and more flooding in others

Rising sea levels in North America

A decrease in oxygen and CO2 storage in the ocean, impacting marine life.

SciAm

IPCC

Earth Observation

“the use of remote sensing technologies to monitor land, marine (seas, rivers, lakes) and atmosphere.” -EUSPA

Can be applied to a wide variety of climate-related questions (and non-climate related)

Relies in large part on publicly and privately run satellite systems

Remote Sensing

Features of satellite imagery

Spatial resolution: what size on the ground does each pixel correspond to

Temporal resolution: how frequently does the satellite revisit the same location

Spectral resolution: how many/which spectral bands does the satellite collect

Multi/Hyper Spectral satellites

Collect data across many wavelengths, not just the standard RGB

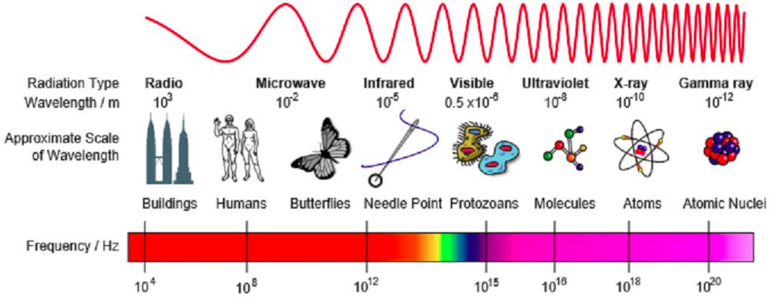

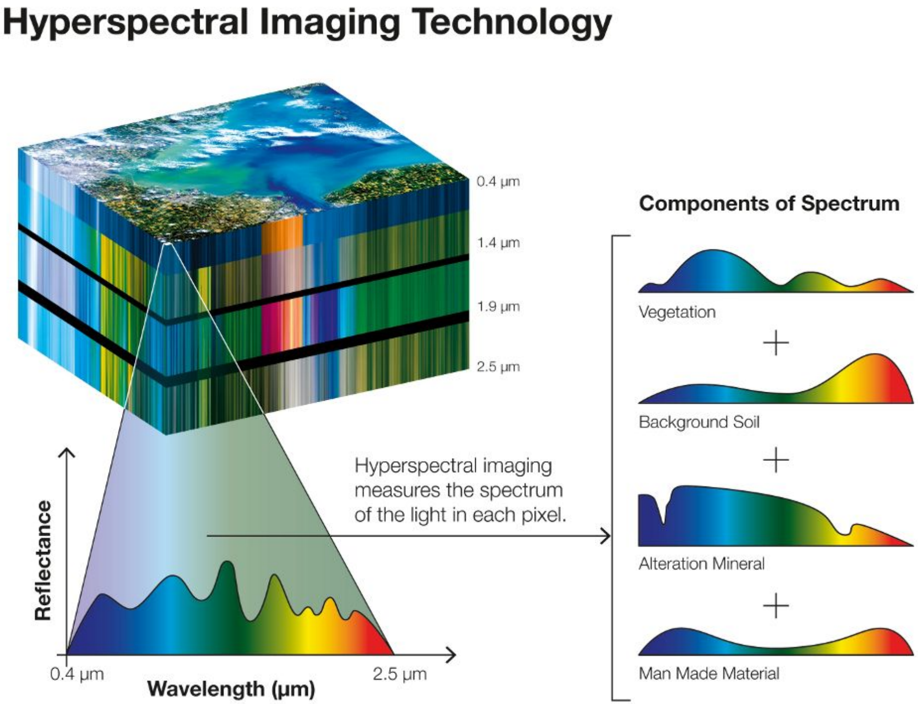

Information from other wavelengths can help identify properties of the materials in the image

Sentinel-2 Satellite

Run by the European Space Agency

10m spatial resolution (for RGB bands)

Revisits every 5 days

13 spectral bands

Sentinel-2 Bands	Central Wavelength (µm)	Resolution (m)
Band 1 - Coastal aerosol	0.443	60
Band 2 - Blue	0.490	10
Band 3 - Green	0.560	10
Band 4 - Red	0.665	10
Band 5 - Vegetation Red Edge	0.705	20
Band 6 - Vegetation Red Edge	0.740	20
Band 7 - Vegetation Red Edge	0.783	20
Band 8 - NIR	0.842	10
Band 8A - Vegetation Red Edge	0.865	20
Band 9 - Water vapour	0.945	60
Band 10 - SWIR - Cirrus	1.375	60
Band 11 - SWIR	1.610	20
Band 12 - SWIR	2.190	20

Sentinel-2 produces about a terabyte of data per day

Images can be fed into artificial neural networks

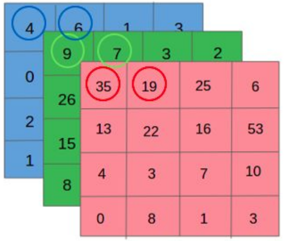

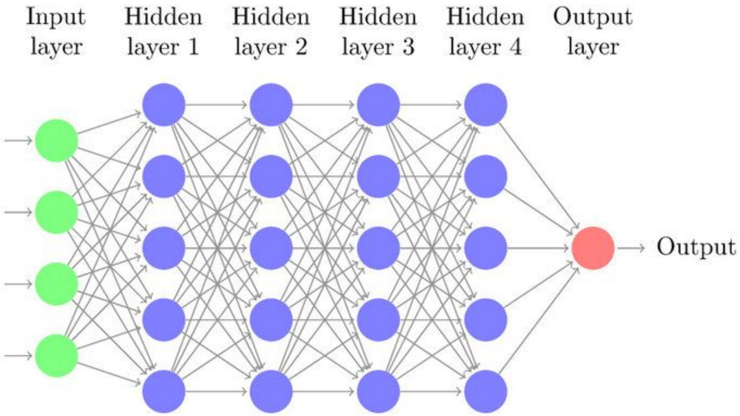

Images can be fed into artificial neural networks

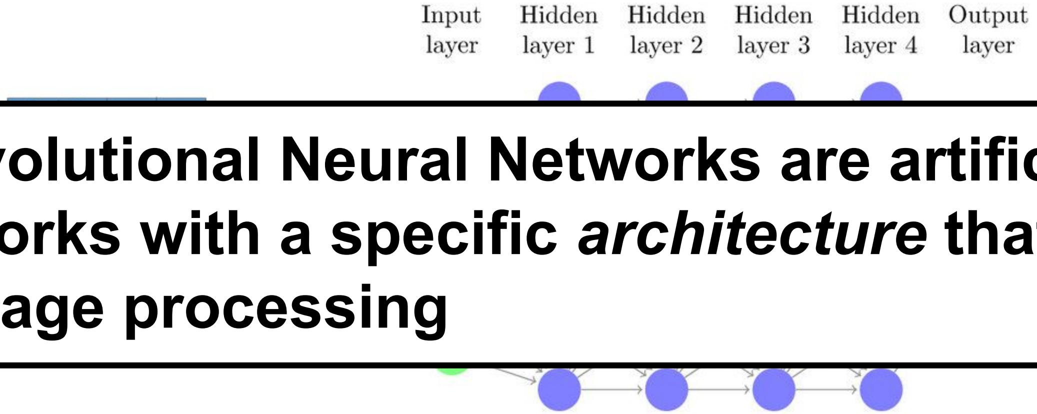

Convolutional Neural Networks are artificial neural networks with a specific *architecture* that is well suited to image processing

Neural network architectures

Defined by:

How many layers and units per layer

What activation functions are used

Any constraints on connections

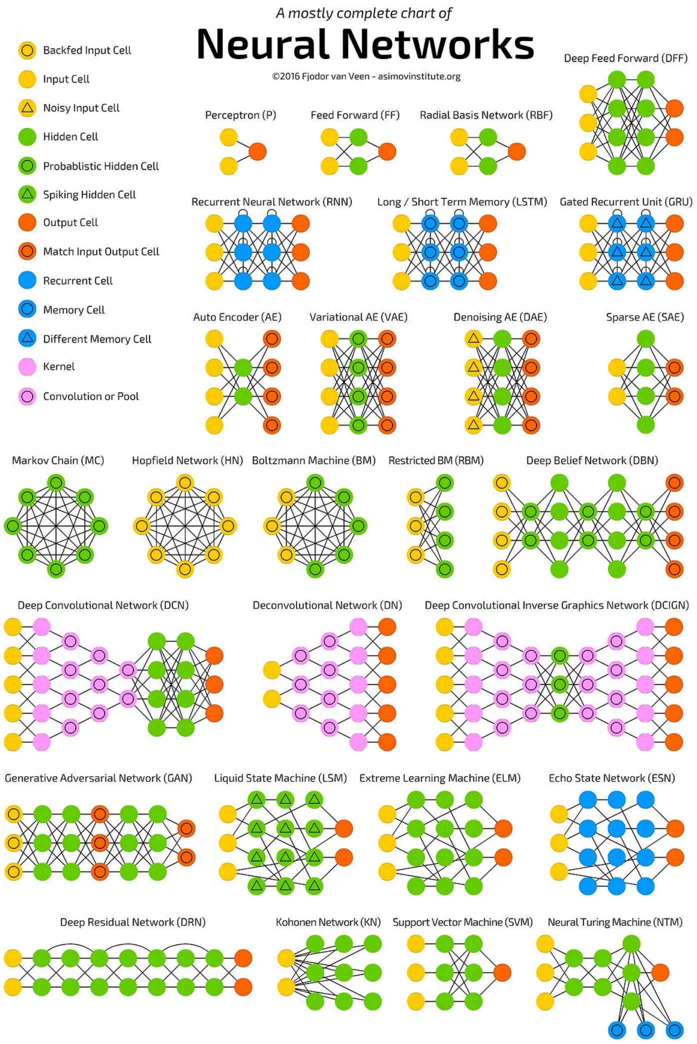

Convolutional neural networks

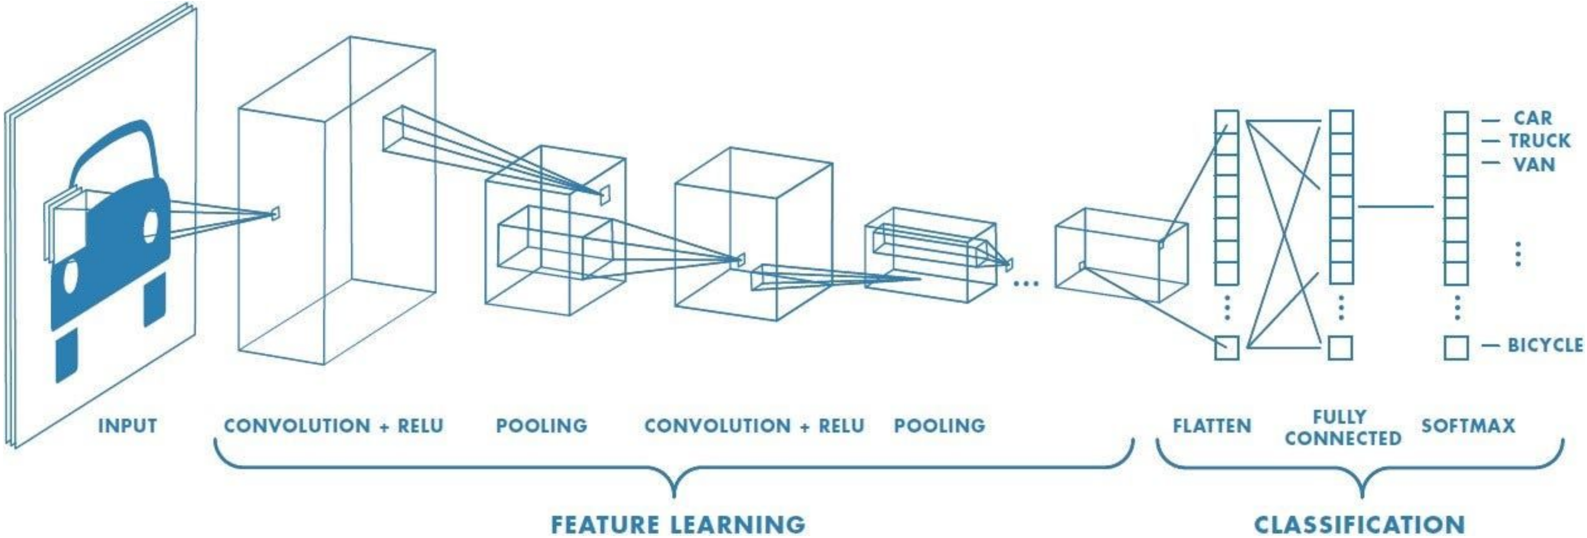

Uses convolution and pooling operations

Convolution

In a convolution the same 2-D grid of weights (called a 'filter') is applied to each location in the image.

Each application of the filter provides the input to a unit at the next layer.

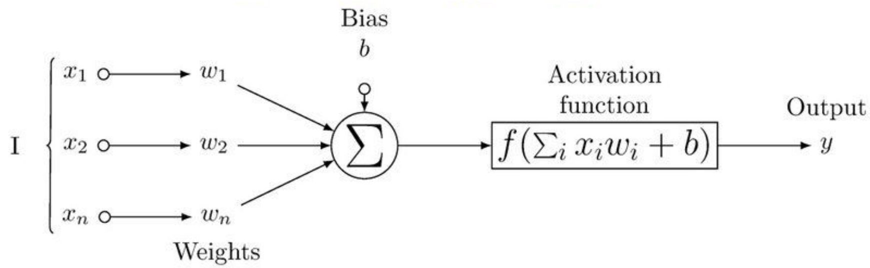

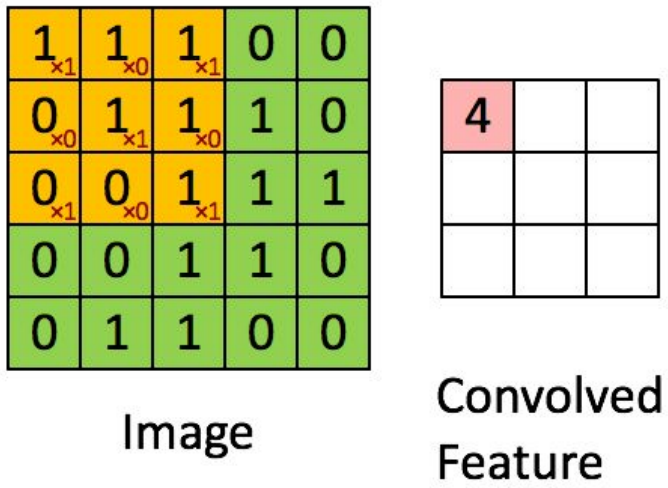

Convolution

Values are highest where the image is most similar to the filter. In this way, convolutions are pattern detectors

$$
\underset{\text{Matrix A}}{\begin{pmatrix} 1 & 2 & -3 & 4 \\ 5 & -6 & 7 & 8 \\ 0 & 1 & 2 & -3 \\ 4 & 5 & 6 & -7 \end{pmatrix}} * \underset{\text{Matrix B}}{\begin{pmatrix} 5 & -6 \\ 0 & 1 \end{pmatrix}} = \underset{\text{Matrix C}}{\begin{pmatrix} -13 & 35 & 31 \\ 62 & -70 & -16 \\ -1 & -1 & 21 \end{pmatrix}}
$$

Pattern/“Filter” — “Map”

Convolution

Values are highest where the image is most similar to the filter. In this way, convolutions are pattern detectors

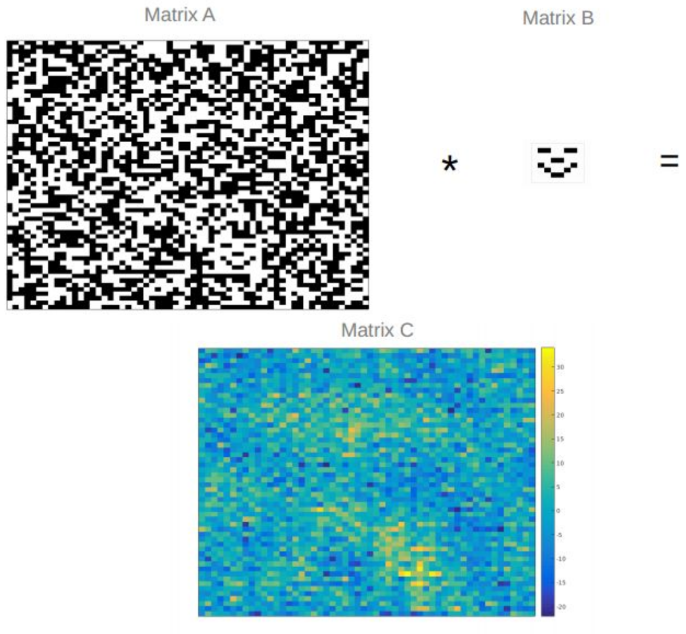

Convolutional neural networks

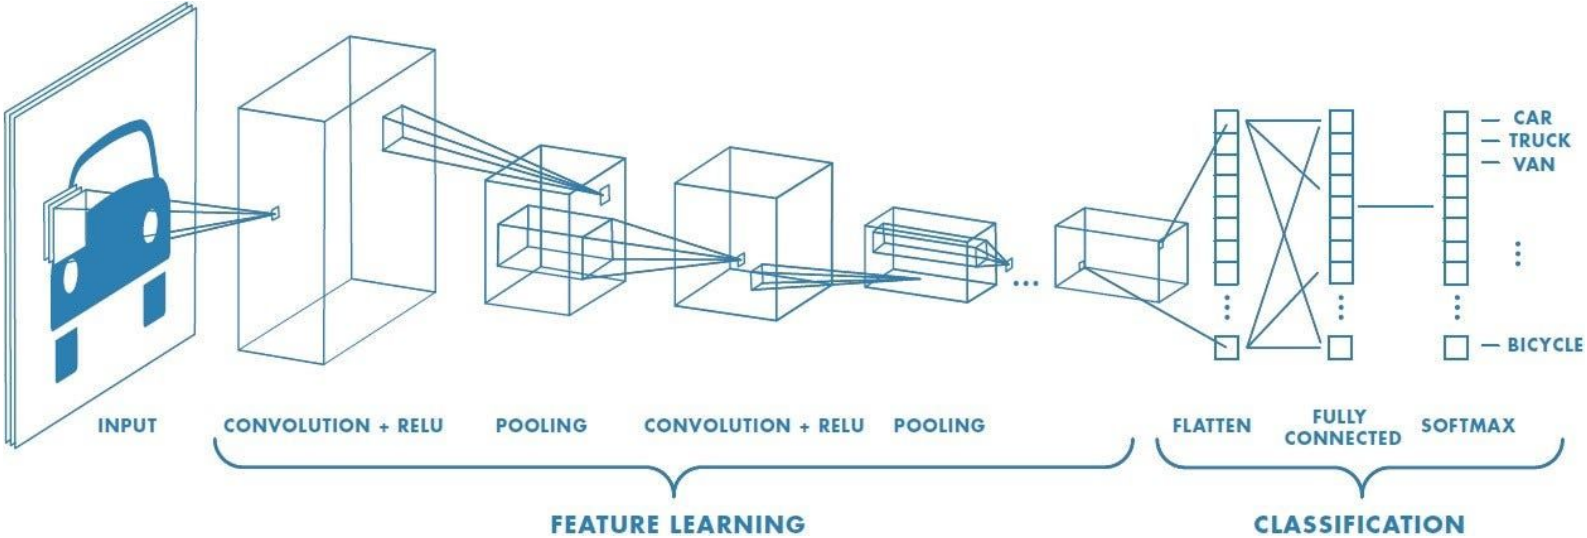

Uses convolution and pooling operations

Pooling

Pooling simply takes the largest value in a specific region of the layer below. This reduces the size of the map/layer.

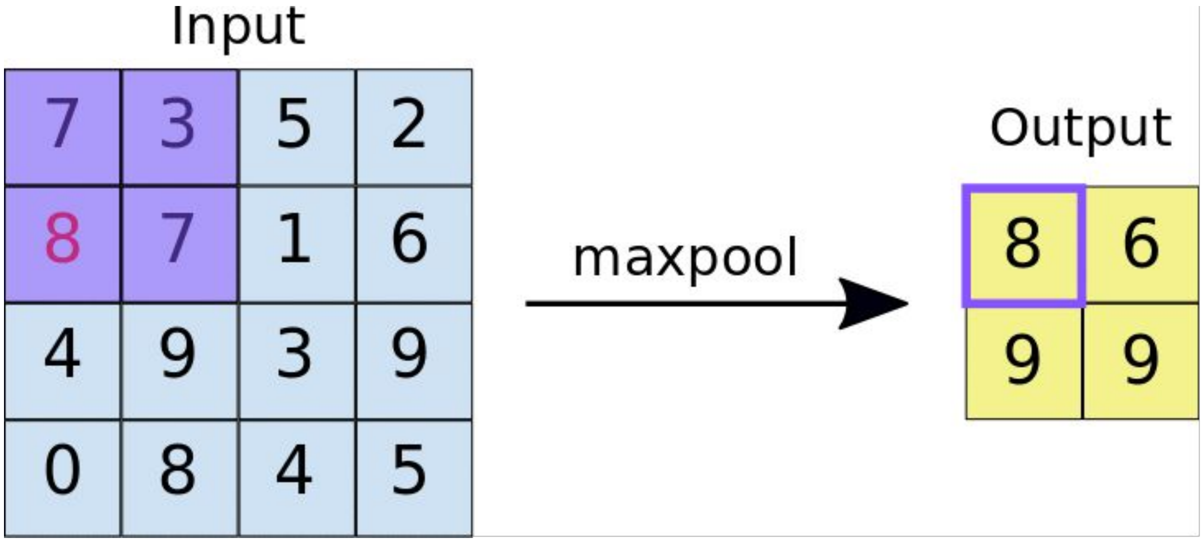

Convolutional neural networks

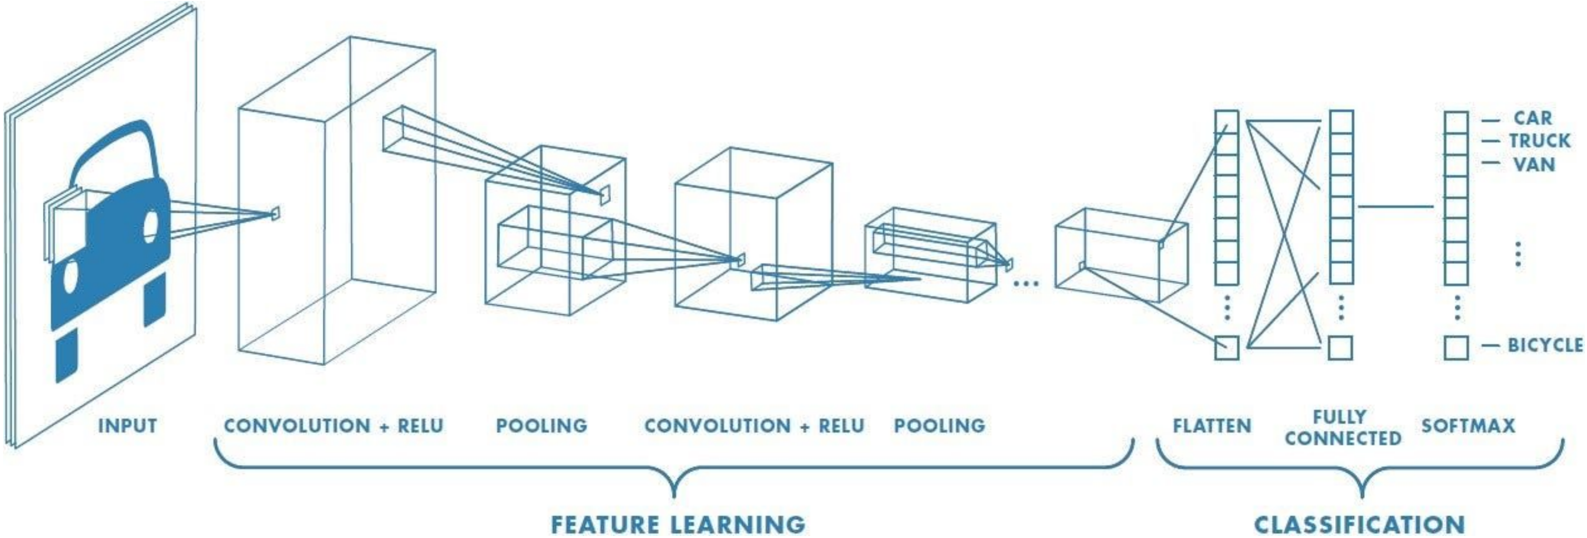

When stacked these operations extract complex spatially invariant image features that can be used for classification

CNNs are inspired by the brain's visual system

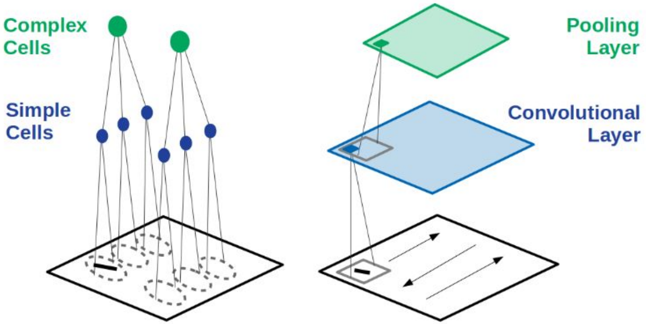

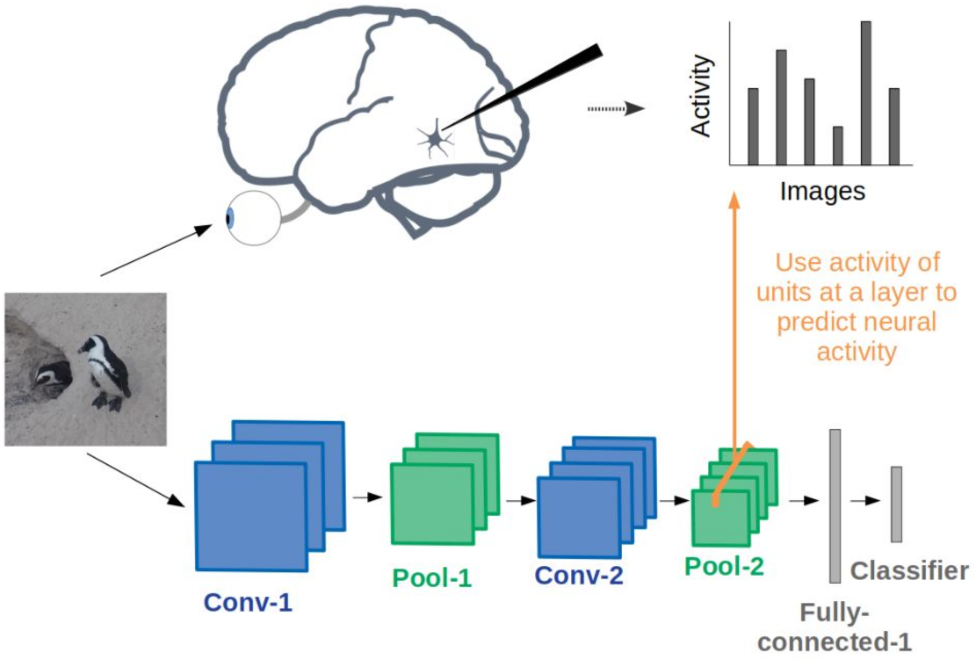

Convolutional neural networks

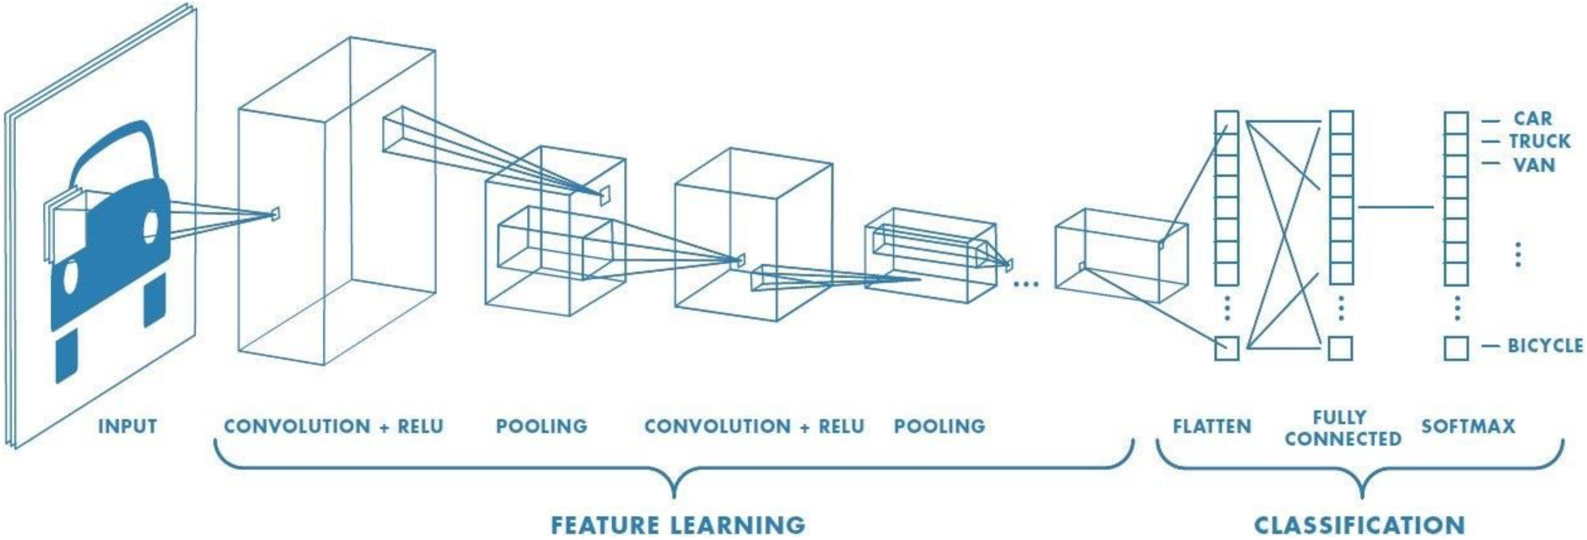

This architecture works for *classification* problems, but what about *segmentation*?

Image Segmentation

Classifying each pixel, based on the image as a whole

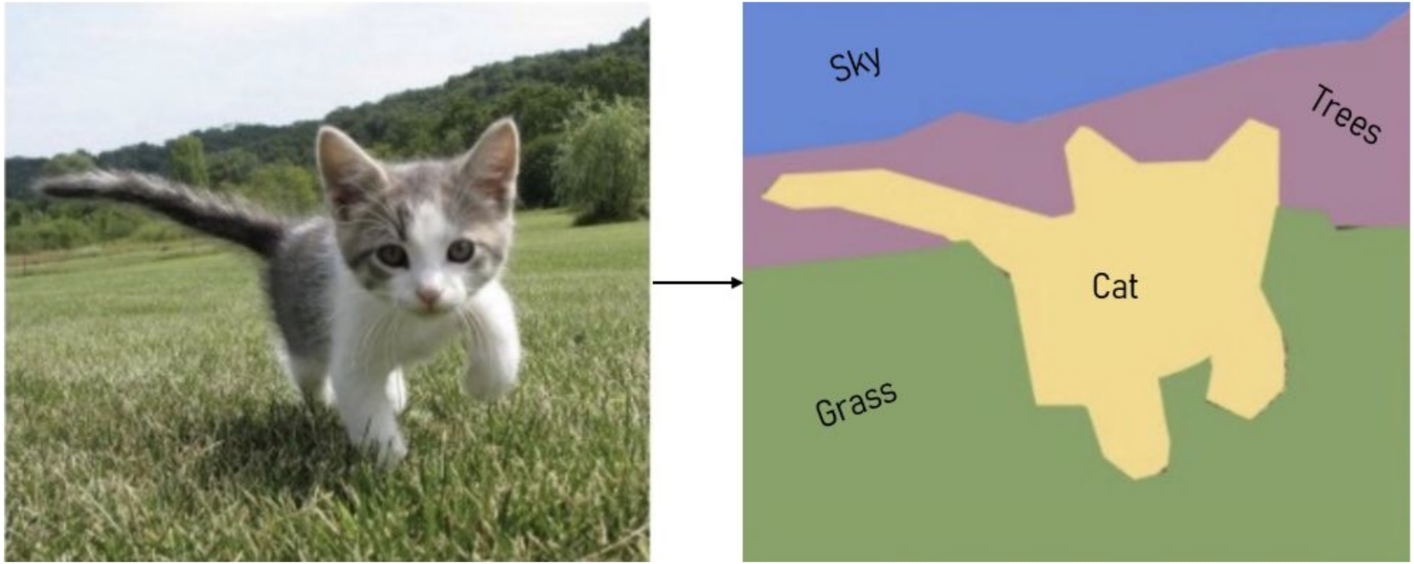

datahacker.rs

The U-Net architecture is used for segmentation

The U-net is like a CNN and a reversed CNN stitched together

It outputs a segmentation mask.

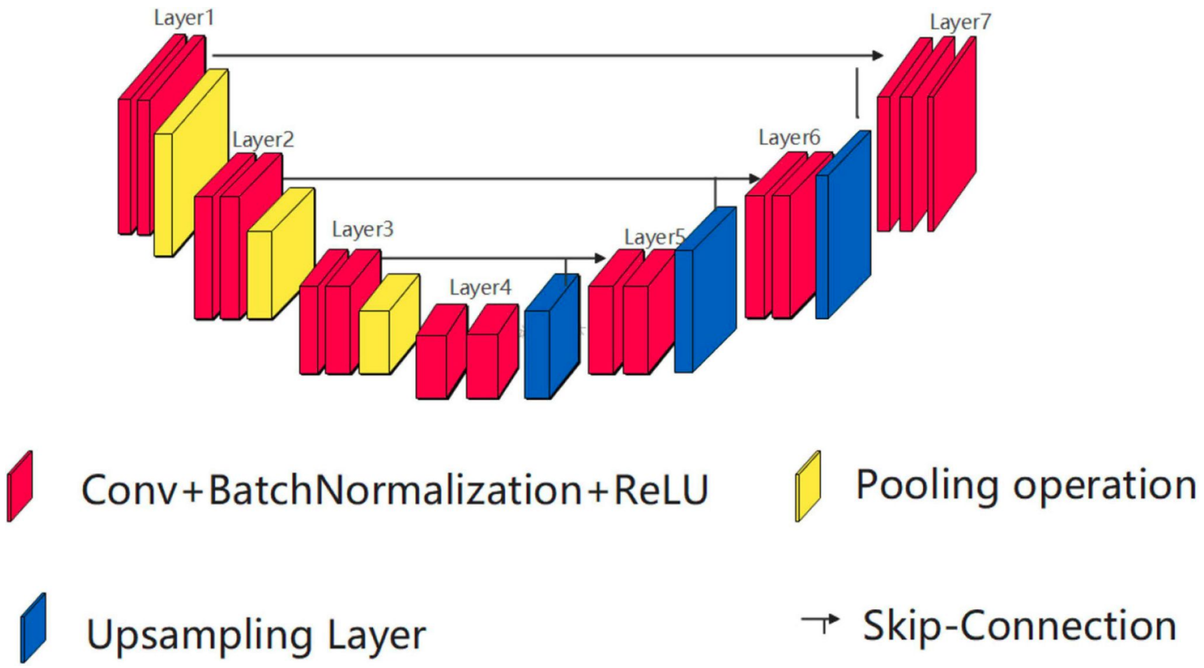

Cai et al

How do we evaluate the quality of a segmentation mask?

Binary classification performance can be assessed according to:

Accuracy: percentage of pixels correctly classified

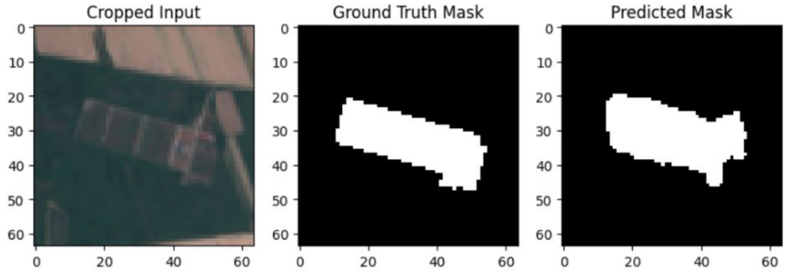

How do we evaluate the quality of a segmentation mask?

Binary classification performance can be assessed according to:

Accuracy: percentage of pixels correctly classified

Precision and Recall: gives more insight into types of error

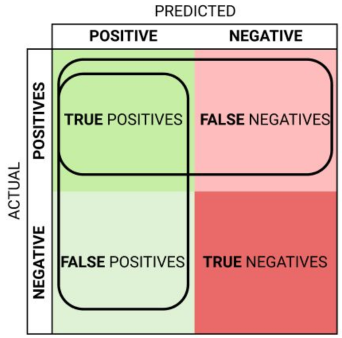

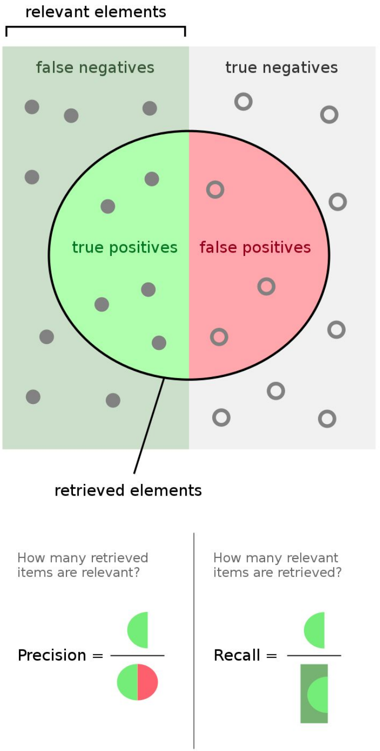

How do we evaluate the quality of a segmentation mask?

Binary classification performance can be assessed according to:

Accuracy: percentage of pixels correctly classified

Precision and Recall: gives more insight into types of error

F1 Score: summary of P & R

$$2\frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$$

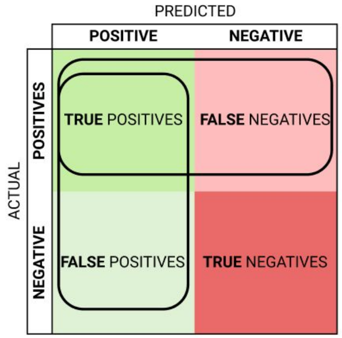

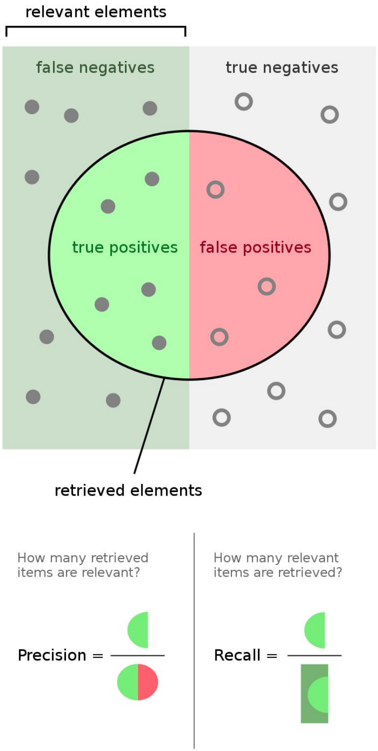

How do we evaluate the quality of a segmentation mask?

- More specific metric:

Intersection Over Union (IoU)

- 1 is best

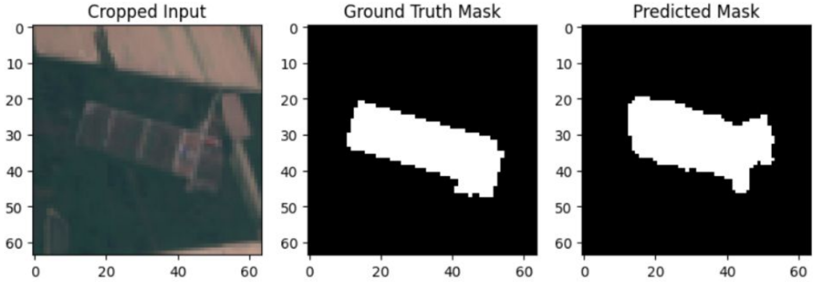

$$\text{IoU} = \frac{\text{Area of Overlap}}{\text{Area of Union}}$$

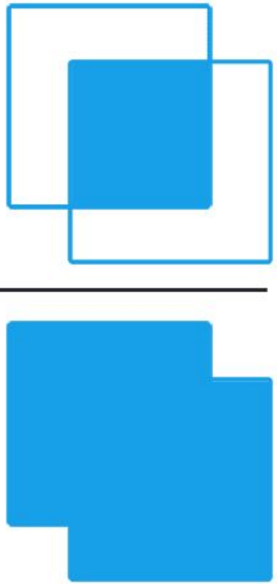

For your reading:

“ResNet” is a particular CNN architecture. It comes in various sizes.

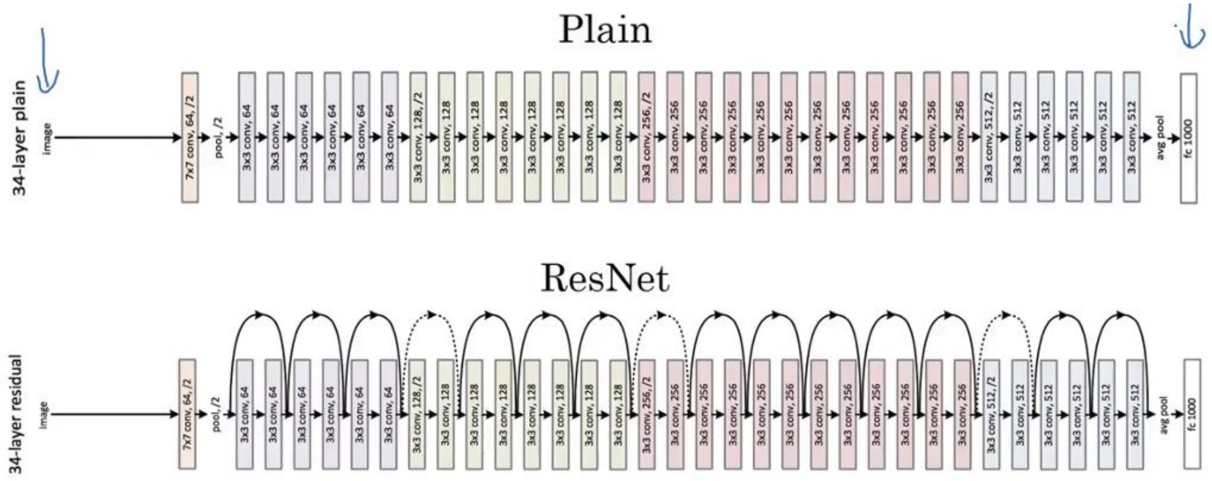

For your reading:

NASADEM is a surface elevation map collected via remote sensing

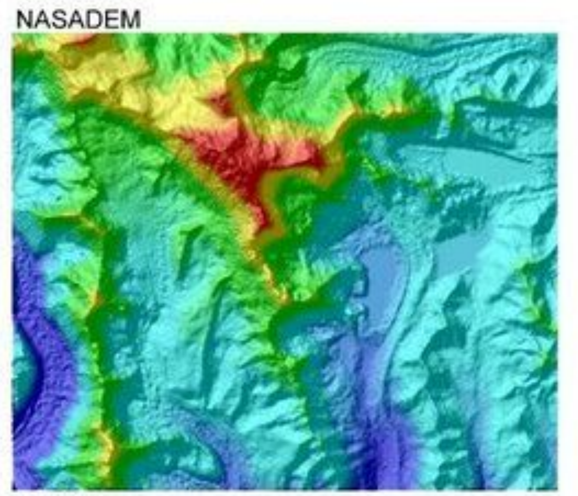

Chen et al, 2021