SAND2022-6511C

Cloud Segmentation & Motion Tracking in Sky Images

A Case Study using Machine Learning for PV Applications

Ben Pierce, Sandia National Labs

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Introduction & Objective

- Forecasting in the near future useful for PV to balance inputs to the grid
- Short term power prediction impacted by cloud behavior.
  - Clouds in front of Sun = change in irradiance received

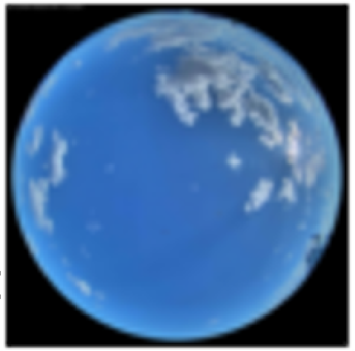

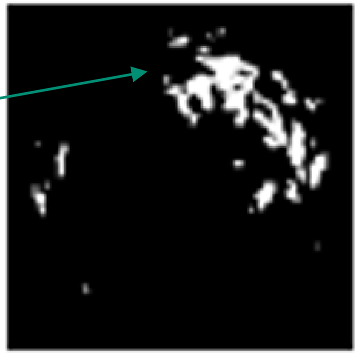

Task I: Identify clouds from image

- Easy for humans
- Hard to do automatically due to various sky conditions

Task II: Track & predict cloud movement through time

- Will there be clouds in front of the Sun in the next 10 minutes?

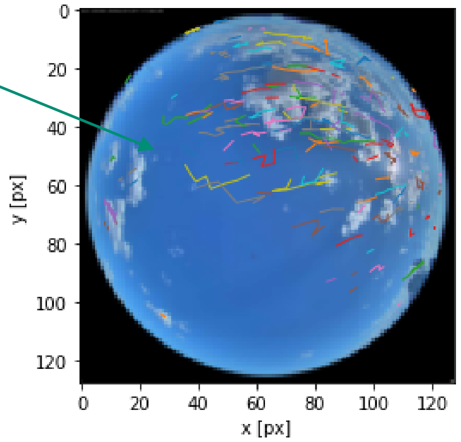

Part I: Cloud Segmentation

Cloud Segmentation

- Need some method to isolate clouds within an image
  - "Segmentation"
- Multiple classical (i.e. non-ML methods)
  - Differ in various ways
  - Sensitive to changes in image type
    - Camera distortion
    - Luminosity
    - Obstructions
- Filtering approach
  - Find some image transform to isolate clouds
- **How do we find a general, robust model that can adapt to the data?**

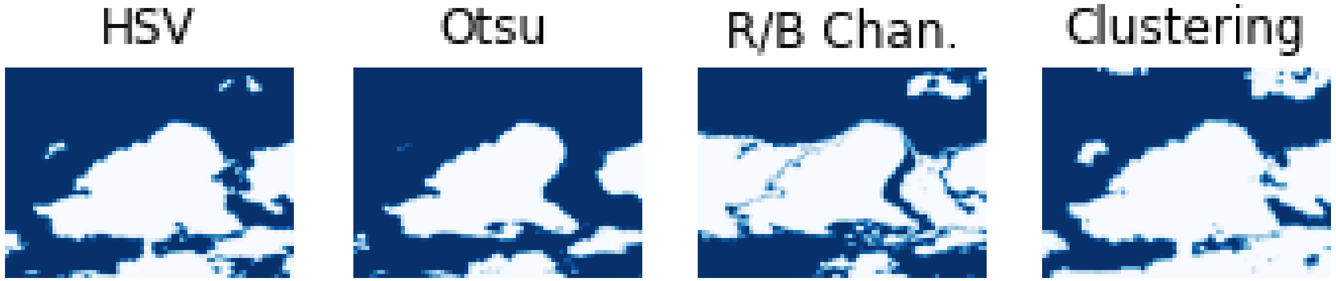

Method	Authors	Source	Date
Grey Level Threshold	Otsu	[2]	1979
RBG Ratio Threshold	Long *et al*	[3]	2006
Hue-Saturation-Value Threshold	Pierce *et al*	[1]	2022
Color channel K-means	Shirazi *et al*	[4]	2021

TABLE I
UNSUPERVISED PREPROCESSING METHODS

Training UNet

- CNNs are **supervised**
  - Require labeled **training data**
- Problem: training data is **expensive** to get
  - Manual (human) labeling
- Idea: use (all) classical methods as training data
  - Learn from strengths/weakness from all methods
- Challenge: **Can UNet model learn from flawed, auto-generated training data?**
  - Also: does this process generalize well to different sources of data?

Input

RBG (128,128,3)

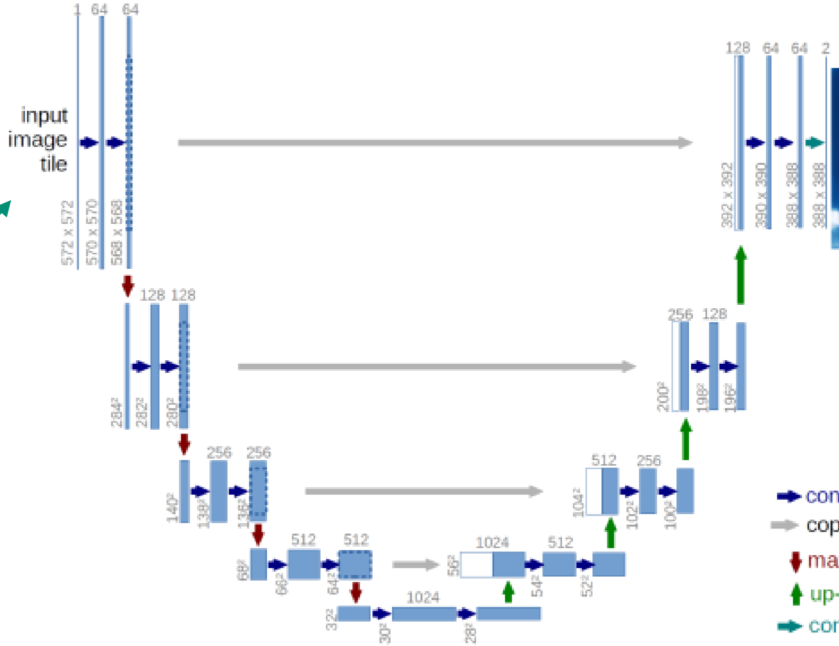

Output

(128,128,1)
probability map

Adjust UNet by taking gradient with respect to *L*

HSV

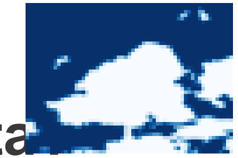

Otsu

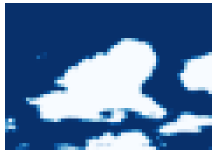

R/B Chan.

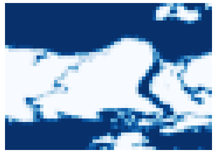

Clustering

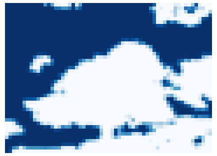

(128,128,1) Binary {0,1}

Results

- Three subsets
  - In sample truth
  - In sample masks
  - Out of sample masks
- Masks predicted via classical algo.
- Differences
  - Cameras
  - Cloud types
  - Locations

Dataset 1

Dataset 2

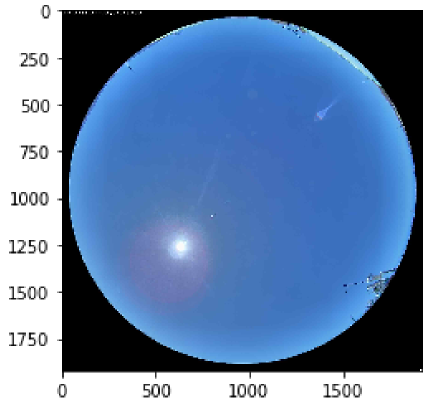

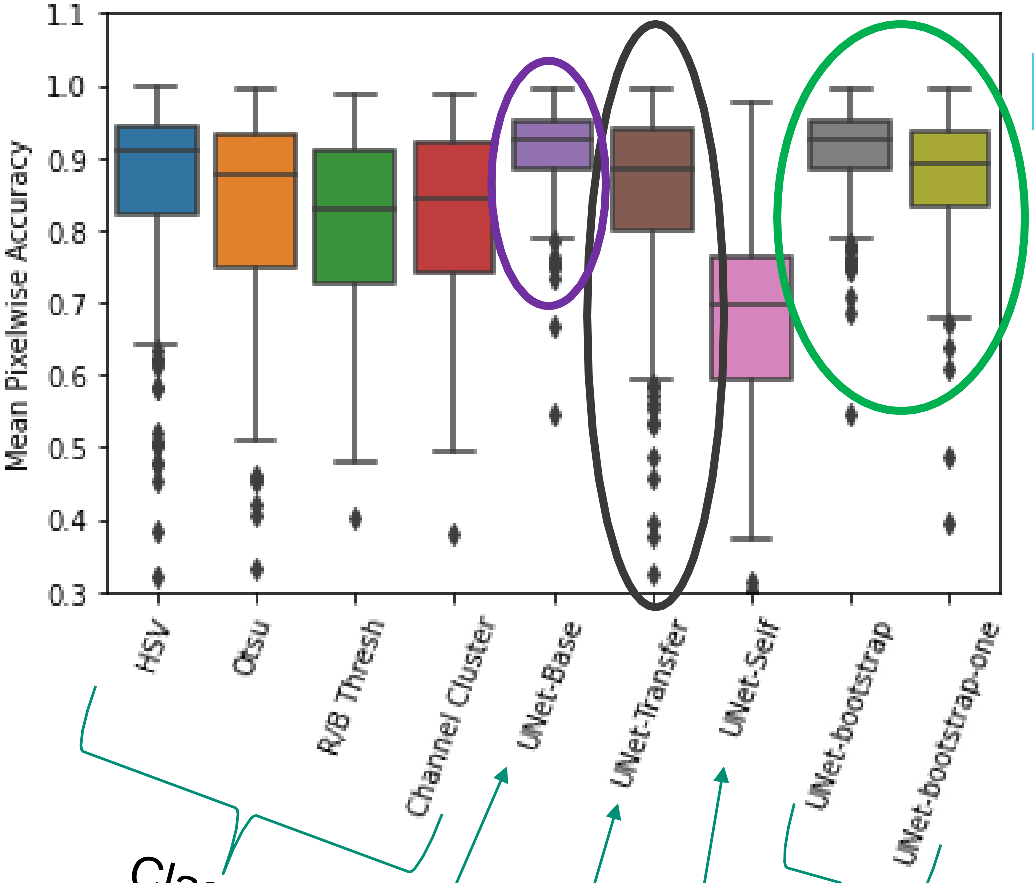

Classical

In-sample truth

Out of sample masks

In-sample masks

Out of sample masks + in sample truth

- When trained, model **performs better** then previous state of the art
- Model **can generalize** out of sample
  - Location & camera independent
- Masks created by classical algo. can be used to **boost initial performance** w/o ground truth training data

Part II: Motion tracking & prediction

Motion tracking & prediction

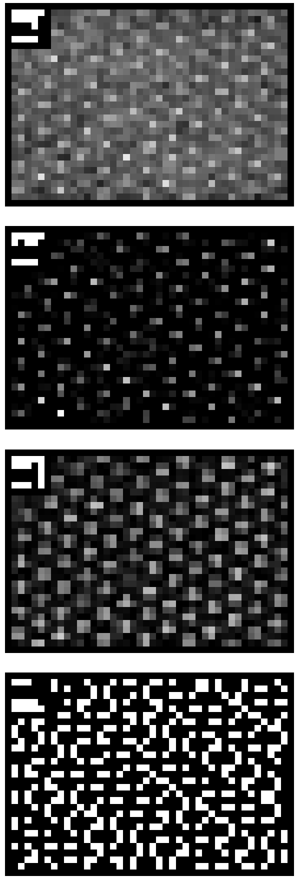

- Crocker and Weeks, 1996 present a method for tracking colloidal particles
  1. Isolate particles from binary mask
  2. Find brightness maxima
  3. Use maxima to locate centroids
  4. Link though time

- Implemented in open-source Python trackpy package
  - Easy to apply to new data
  - Provides exactly the statistics we need

Locate objects

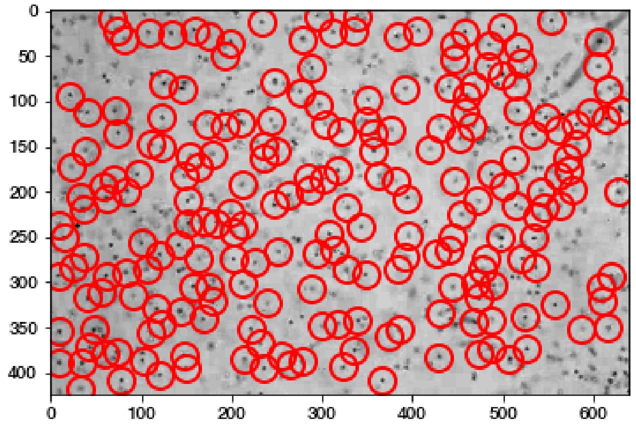

Object level drift

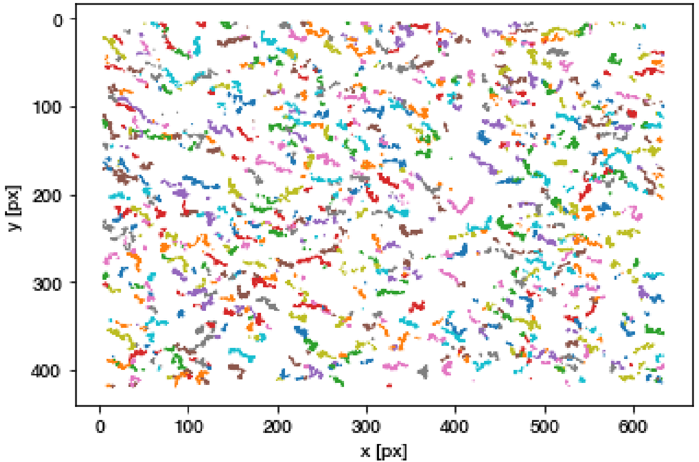

Overall drift

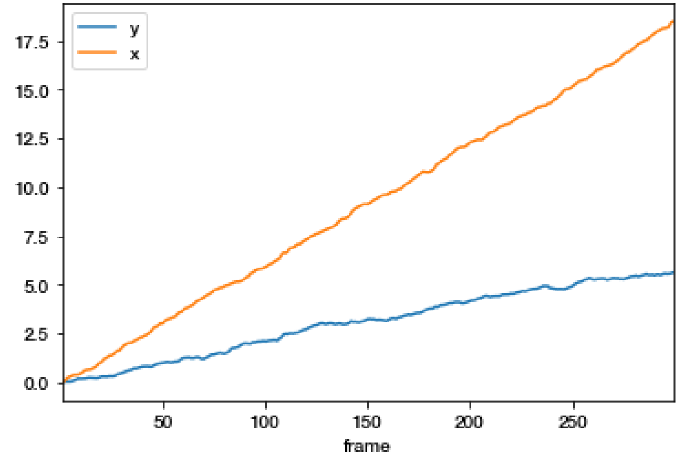

Using motion linkages to predict future

- Trackpy links particles through time
- This provides us with their *xy* position over *n* minutes
- How do we predict the future?
  - Numerical gradients via Taylor expansion/central differences
- Once gradient is calculated, take mean with respect to *x* and *y* directions.
  - This is projected cloud drift

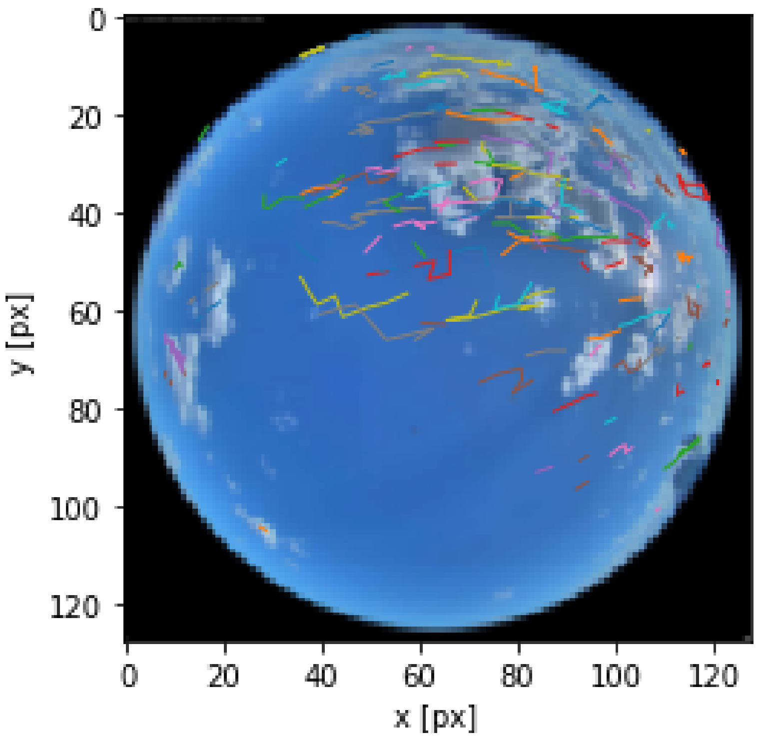

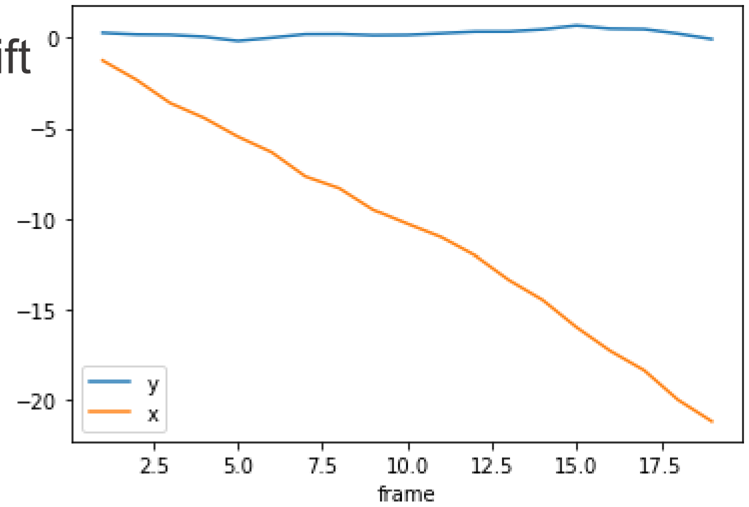

Conclusions

- Deep Learning achieves great accuracy on cloud segmentation, but can be hindered by lack of a training dataset
  - This can be addressed with imperfect, autogenerated labels
- Transfer learning, even with imperfect labels
  - Can be used to ease training a deep learning model
- Simple colloidal tracking algorithms can be repurposed for motion tracking and prediction of clouds
- These methods useful to calculate projected input to grid as short-term power forecasts

Questions?

Extra slides for questions

Machine Learning in PV

- Machine Learning (ML) is a catch-all term for a variety of statistical methods, including (but not limited to)
  - Regression, clustering, feature extraction, dimensionality reduction, etc.
  - **Neural Networks/Deep Learning**
- Many uses in PV
  - Power forecasting
  - Material performance
  - Fault detection
- Strength: **Using trends in high-dimensional data to solve problems with no known closed-form equations**
  - No existing physics-based equations
  - Data in hard-to-quantify formats (e.g. images)
  - Large quantities of data

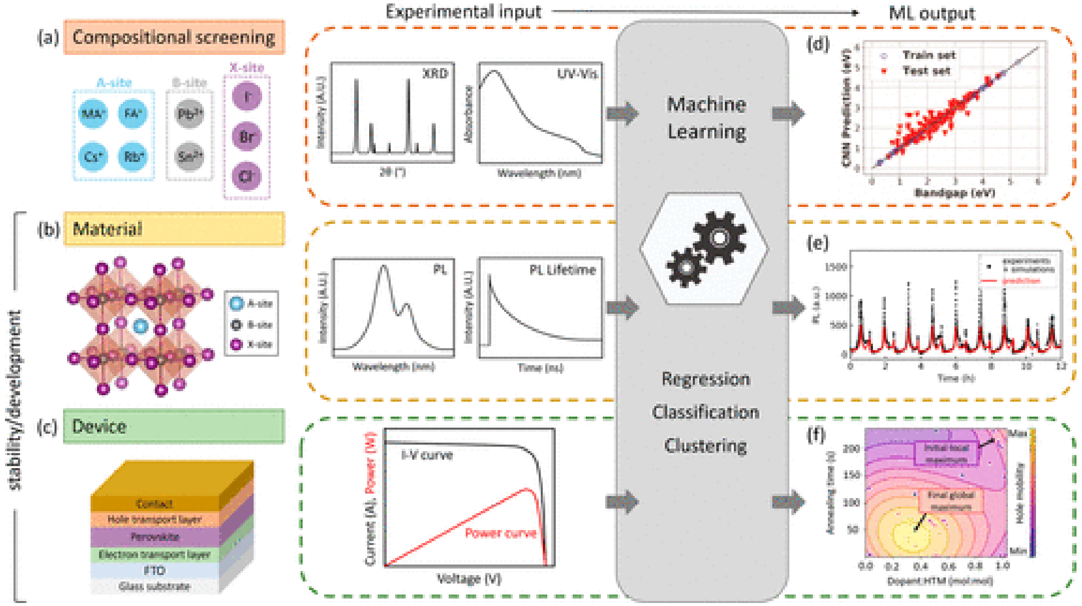

Srivastava et al, *Machine Learning Roadmap for Perovskite Photovoltaics*

Kahn et al, *Improved solar photovoltaic energy generation forecast using deep learning-based ensemble stacking approach*

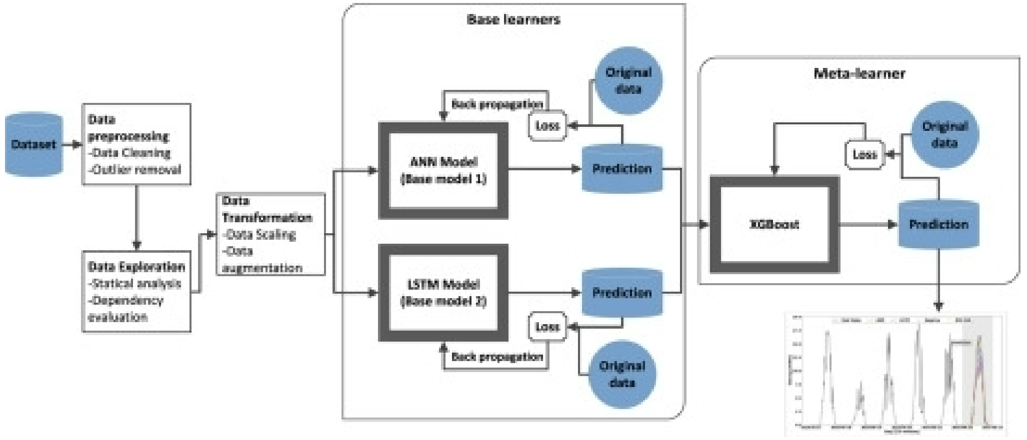

Convolutional Neural Networks

- Idea: use data to learn general filters
- Filter is called the “kernel”
  - Which has elements called “weights”
- Learning task
  - **Find weights for useful filters**

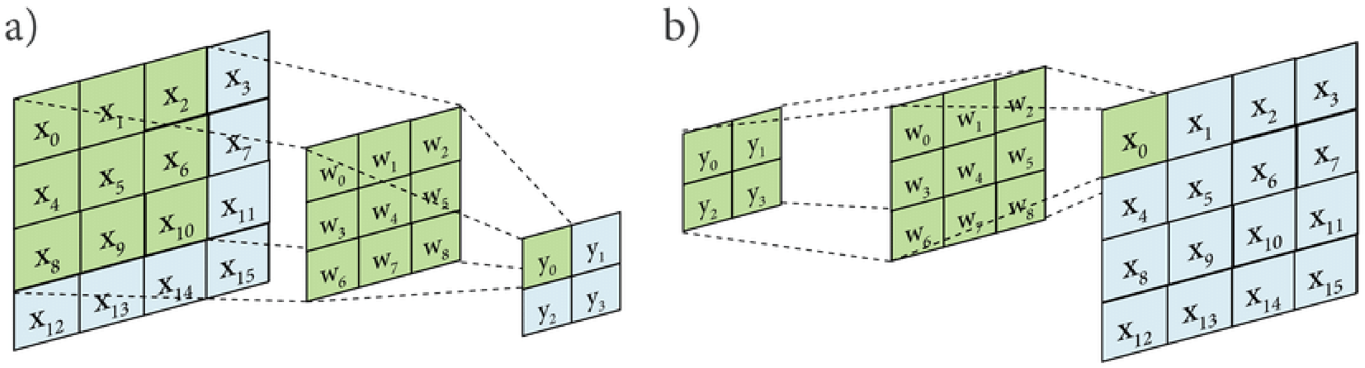

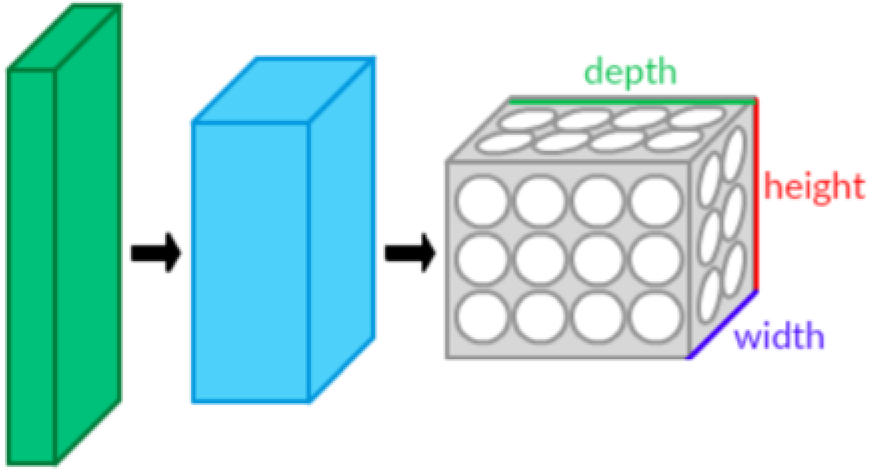

Input image

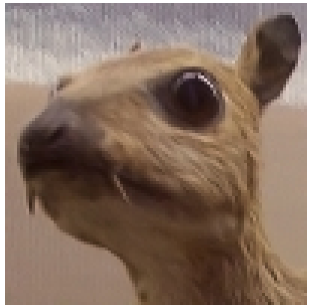

Convolution Kernel

$$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$

Feature map

Fully Convolutional NNs: UNet

- Idea: use only conv layers to alter input to produce binary mask as output
- UNet Model
  - Downsample, then upsample
  - Skip connections between mirror layers
- Start with RGB image-> Binary mask of cloud/not cloud encoding

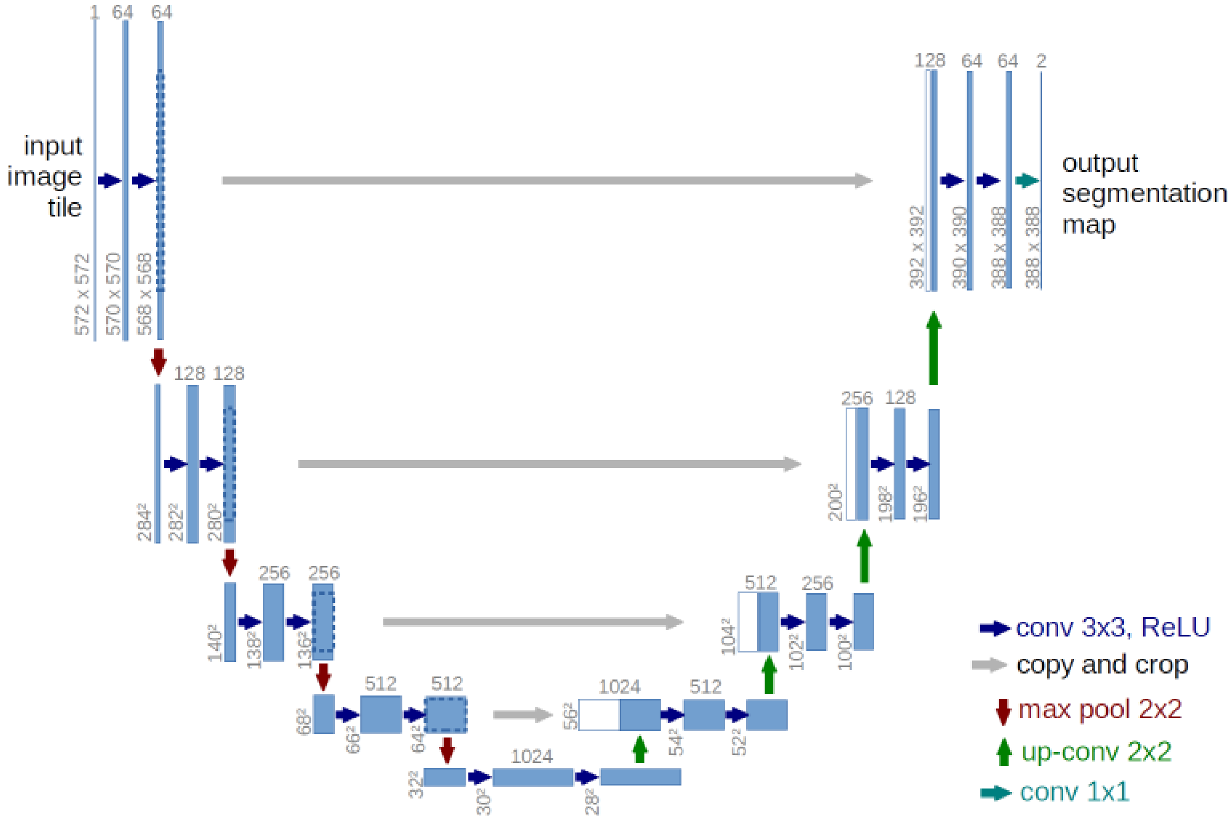

Original HSV mask CAE prediction Difference

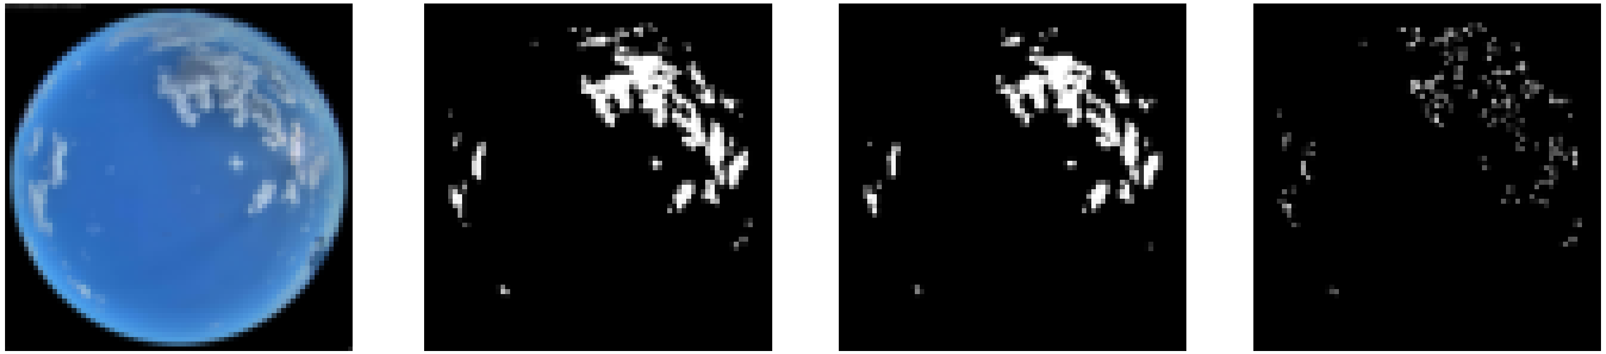

Part II: Motion tracking & prediction

- Problem: how do we track clouds through time to predict their movements?
  - Can we do this in real-time on low-power hardware?
- Inspirations from **biology**
  - Tracking animals & cells in the lab
- Deep learning based approaches show promise, but lack
  - Explain-ability
  - Consistent performance
  - Speed
  - Brevity
- This may be a situation where ML is **not** needed

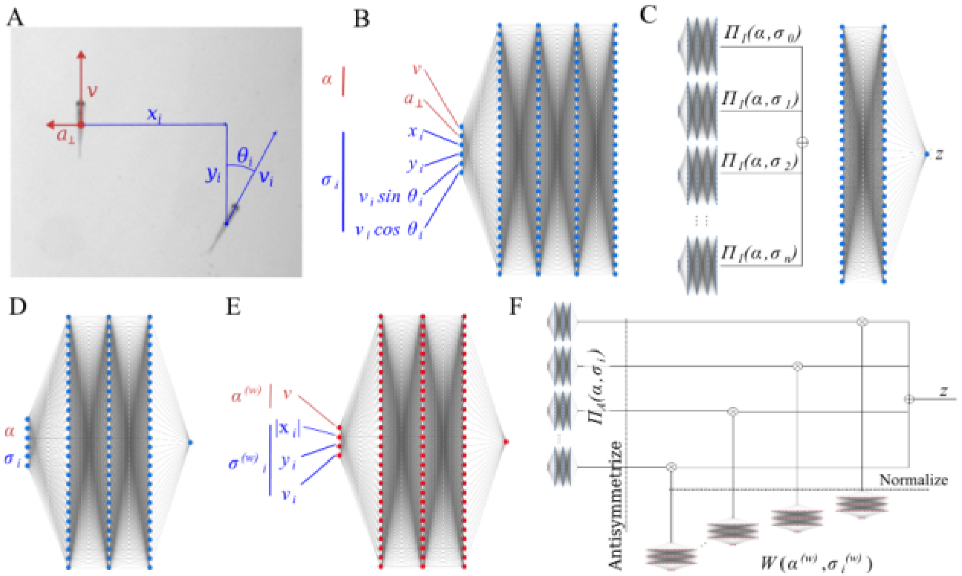

Fig 1. Deep-learning a model of collective behaviour. (A) Variables used to predict future turns. Asocial variables, those only involving the focal, in red. Social variables, those involving both the focal and a neighbour, in blue. (B) Pair-interaction subnetwork, receiving asocial variables α and social variables σ_i from a single neighbour i, and outputting a vector of 128 components. All pair-interaction networks share the same weights. (C) Interaction network, showing how the outputs of the pair-interaction subnetworks, one for each neighbour, are summed and then fed to an interaction subnetwork. The output, z is the logit of the focal fish turning right after 1 s. (D) Pair-interaction subnetwork of the attention network. (E) Aggregation subnetwork of the attention network. Same structure as D, but the input is a restricted symmetric subset of the variables and the output is passed through an exponential function to make it positive. (F) Attention network, showing how the inputs of the pair-interaction and aggregation subnetworks are integrated to produce a single logit z for the focal fish turning right after 1 s.

https://doi.org/10.1371/journal.pcbi.1007354.g001

Heras et al, *Deep attention networks reveal the rules of collective motion in zebrafish*

Effects of cloud cover on PV

- Cloud cover causes radiation to become more diffuse
  - Ex: where's your shadow?
- Difficult to predict quantitatively without specialized sensors
- Easy to observe– just look up
  - Sky camera
- **How do we make quantitative predictions from image data?**