Extreme Weather Event Prediction

Enhancing weather forecasts

Assignments

Keep working on your projects!

Poll on project work day

Homework #4 - Dengue fever cases

Due Monday the 17th by midnight

Climate change in the news

Climate change in the news

PRESS RELEASE

H&M MOVE PARTNERS WITH LANZATECH TO LAUNCH CAPSULE COLLECTION USING CAPTURED CARBON EMISSIONS

This is the stuff of science fiction: LanzaTech diverts carbon emissions heading for the atmosphere, traps them, and turns them into thread. In a leap towards innovating sportswear, H&M Move partners with the breakthrough material science company for a drop arriving at hm.com/move on April 6.

Using three simple steps, LanzaTech captures carbon emissions from steel mills, traps them in bioreactors and converts them into the same building blocks that conventional polyester is made of. This revolutionary solution helps reduce pollution and limits the use of virgin fossil resources needed to make new products.

H&M Move's upcoming three-piece drop for women includes a jumpsuit, a top and a pair of tights, partly made of LanzaTech CarbonSmart™ polyester with a DryMove™ finish. The sleek, all-black garments draw on the contemporary athluxury movement and feature contrasting seams and creative cut-outs.

Recap

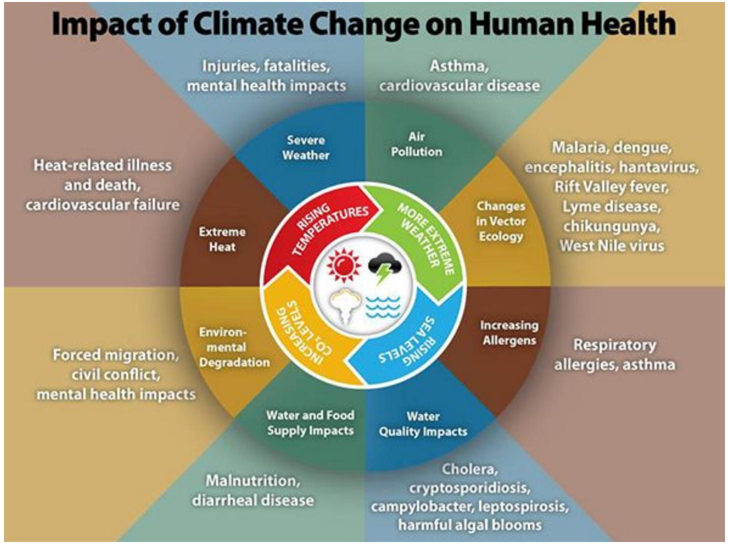
Impact of Climate Change on Human Health
Injuries, fatalities, mental health impacts
Asthma, cardiovascular disease
Severe Weather
Air Pollution
Heat-related illness and death, cardiovascular failure
Malaria, dengue, encephalitis, hantavirus, Rift Valley fever, Lyme disease, chikungunya, West Nile virus
Extreme Heat
Changes in Vector Ecology
RISING TEMPERATURES
MORE EXTREME WEATHER
INCREASING CO2 LEVELS
RISING SEA LEVELS
Environmental Degradation
Increasing Allergens
Forced migration, civil conflict, mental health impacts
Respiratory allergies, asthma
Water and Food Supply Impacts
Water Quality Impacts
Malnutrition, diarrheal disease
Cholera, cryptosporidiosis, campylobacter, leptospirosis, harmful algal blooms


NASA
www.nasa.gov
TIPS AND TRICKS
NASA GLOBE OBSERVER
MOSQUITOES
Learn to IDENTIFY mosquito habitats
Tip #1: Wear long sleeves, pants, socks, and shoes. Apply an effective insect repellent to exposed skin.
GLOBE
mosquito habitat mapper
Observations: 1
New Mosquito Habitat Observation
Review / Send My Mosquito Observations
See My Data
See Today's Mosquito Measurements
GLOBE Observer
observer.globe.gov

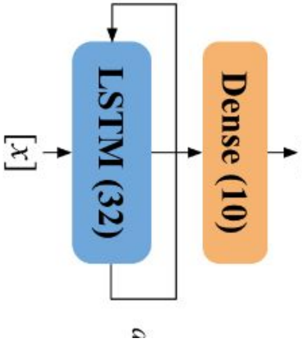
[x]
LSTM (32)
Dense (10)

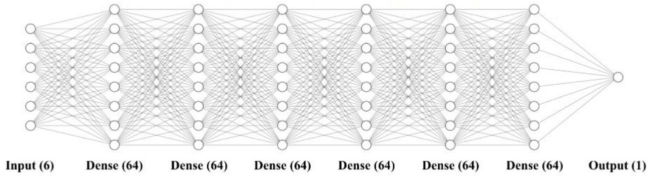
Input (6)
Dense (64)
Dense (64)
Dense (64)
Dense (64)
Dense (64)
Dense (64)
Output (1)

Recap

Prof. Laure Zanna

Geophysical Research Letters®

Research Letter | Free Access

Data-Driven Equation Discovery of Ocean Mesoscale Closures

Laure Zanna, Thomas Bolton

First published: 06 August 2020 | https://doi.org/10.1029/2020GL088376 | Citations: 55

https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2020GL088376

Goal: find an equation that relates small scale physics to large scale, in order to correct coarse-grain simulations

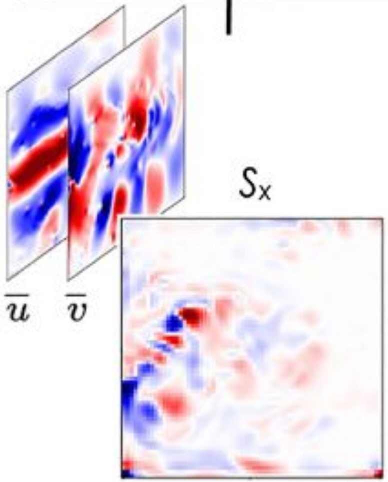

Diagnose **S** from the high-resolution model velocities

Plan: Use “relevance vector machine” regression

- “Equation discovery” - provide a list of terms and it will find weights for each term
- Requires some ‘hand-engineering’
- Produces more interpretable result
- some memory constraints

“filtered velocities using up to second order for both spatial derivatives and polynomial products”

Basis functions:	
Divergence:	$\sigma = \overline{\nabla} \cdot \overline{\mathbf{u}}$
Vorticity:	$\zeta = \overline{\nabla} \times \overline{\mathbf{u}}$
Shearing Deformation:	$D = \frac{\partial \overline{u}}{\partial y} + \frac{\partial \overline{v}}{\partial x}$
Stretching Deformation:	$\tilde{D} = \frac{\partial \overline{u}}{\partial x} - \frac{\partial \overline{v}}{\partial y}$
Library of functions as divergence of flux:	$\phi_i(\overline{u}, \overline{v})$

RVM regression

- Gaussian prior distributions for each regression weight
- Allows for a measure of uncertainty
- Terms are iteratively pruned from the equation if width of the distribution is too large
- This uncertainty threshold is the only parameter that requires setting

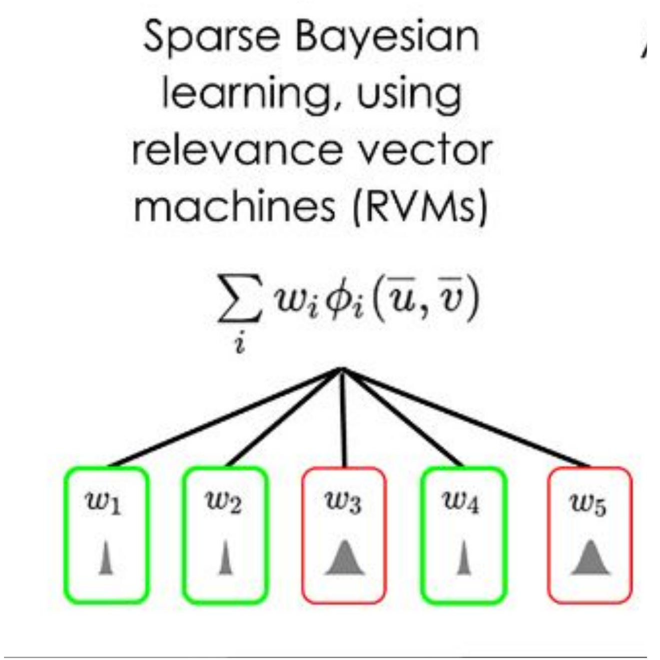

Overall process

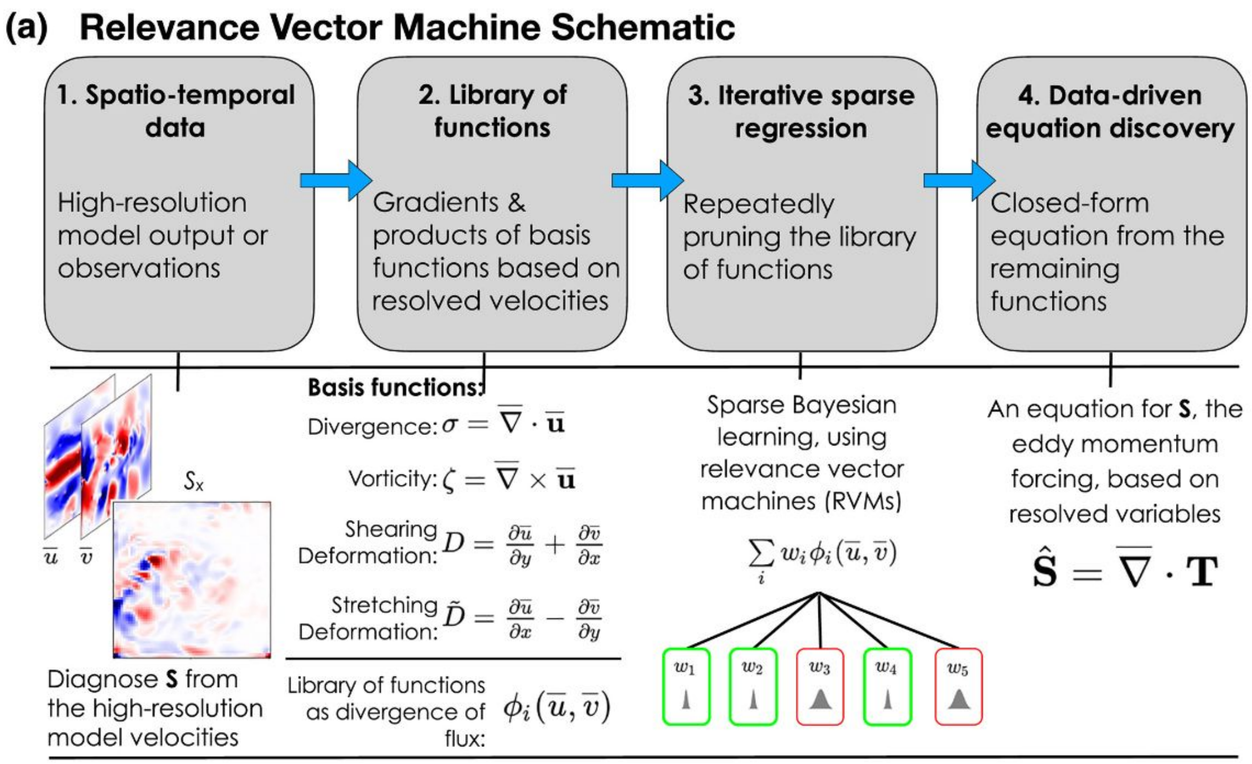

Results

RVM versus convolutional neural network approach has some tradeoffs

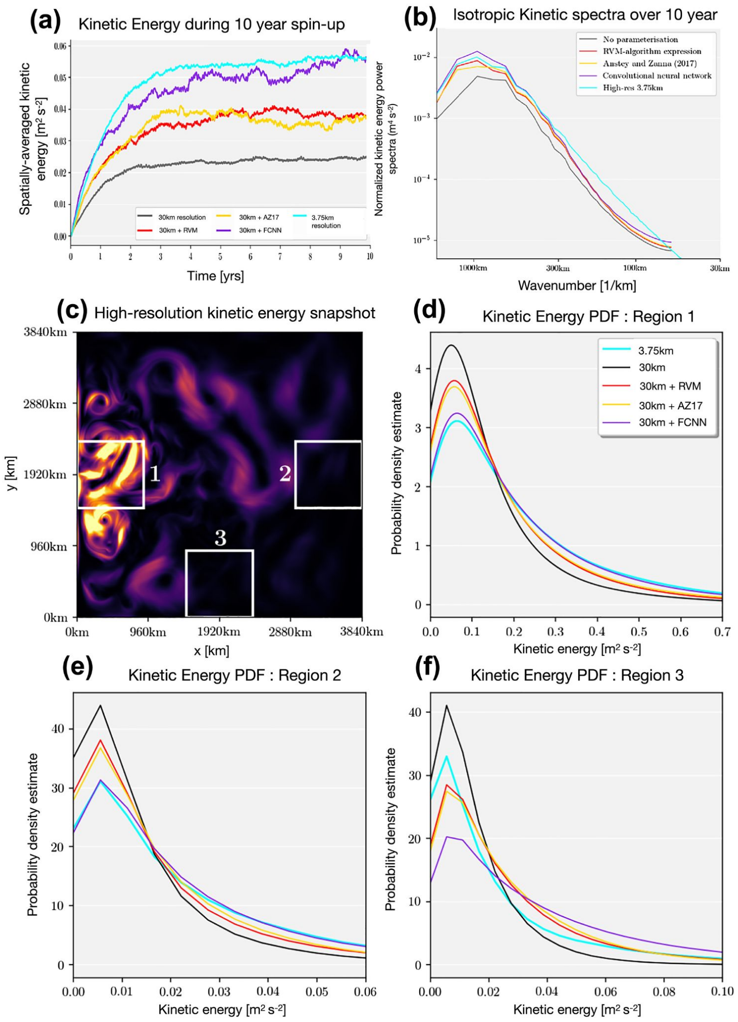

Extreme weather events will become more common

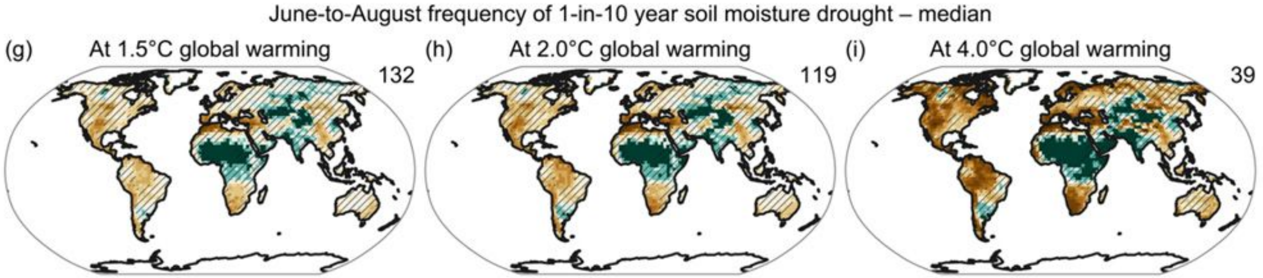

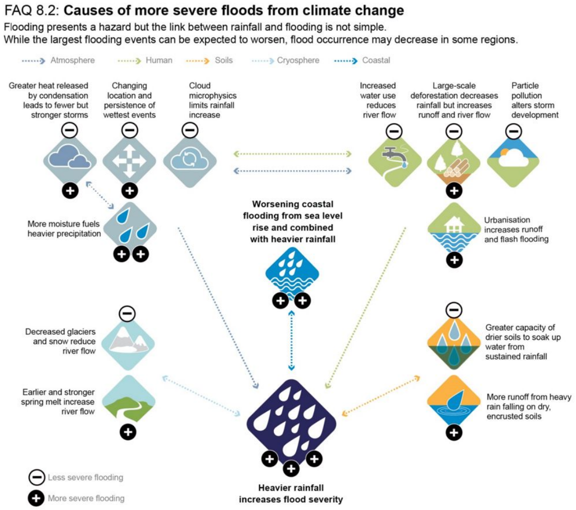

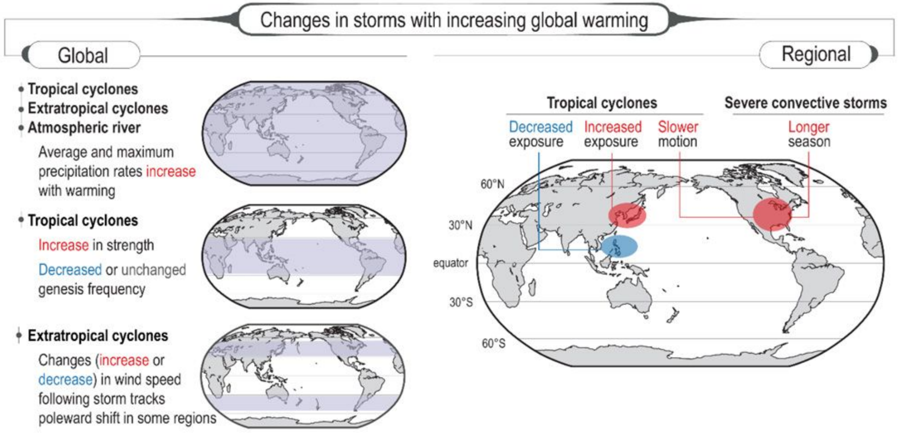

Importance of prediction

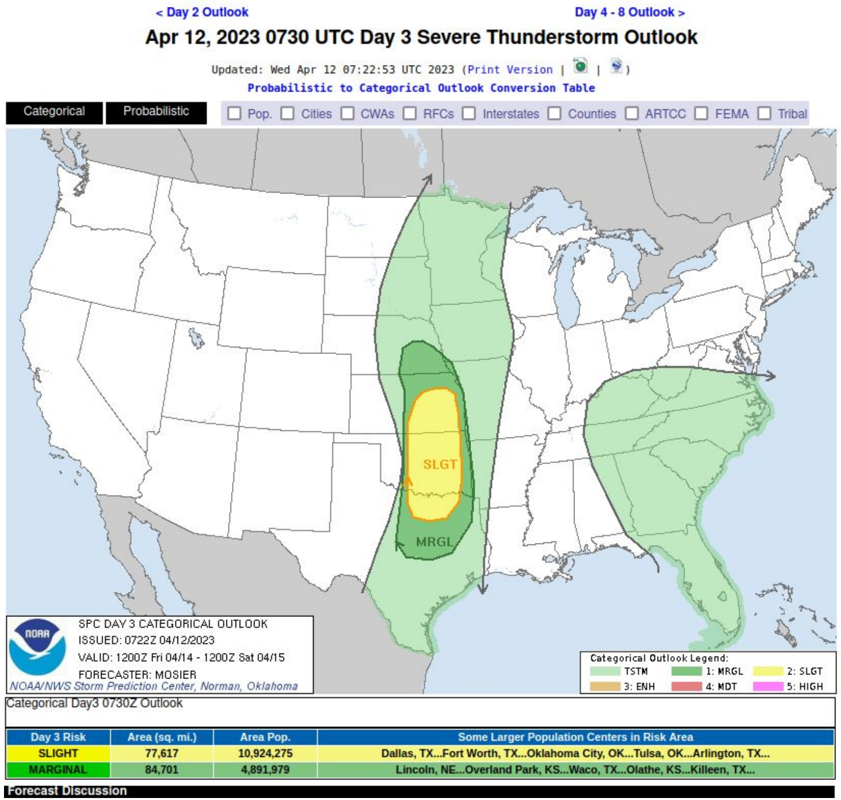

Day 3 Risk	Area (sq. mi.)	Area Pop.	Some Larger Population Centers in Risk Area
SLIGHT	77,617	10,924,275	Dallas, TX...Fort Worth, TX...Oklahoma City, OK...Tulsa, OK...Arlington, TX...
MARGINAL	84,701	4,891,979	Lincoln, NE...Overland Park, KS...Waco, TX...Olathe, KS...Killeen, TX...

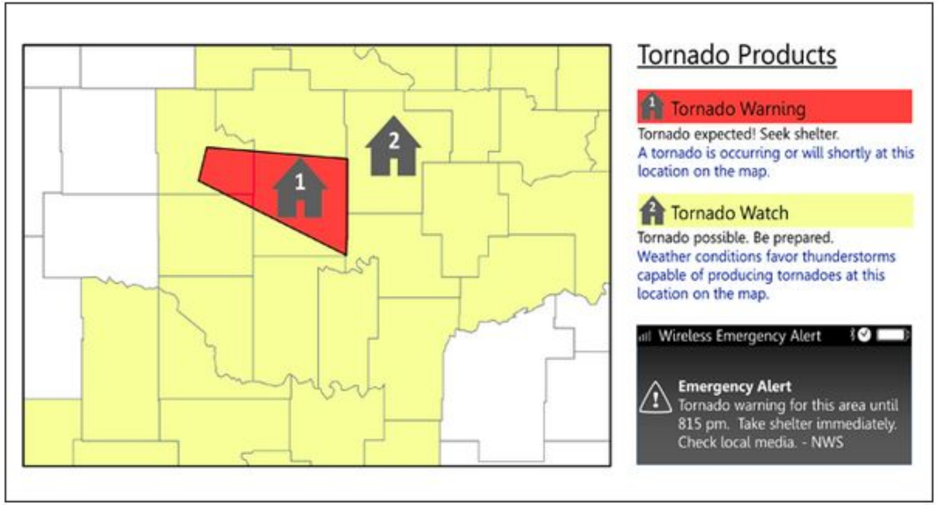

Prediction and alert systems give people time to prepare. The earlier the better, but accuracy matters too.

Challenge of weather prediction

How predictions are done

Numerical simulations predict future weather using physics and dynamics to mathematically describe the atmosphere's behavior

Models cover points over a large area, from the Earth's surface to the top of the atmosphere.

They start with current weather data and run forward

Data is gathered from weather balloons launched around the globe twice each day, in addition to measurements from satellites, aircraft, ships, temperature profilers and surface weather stations.

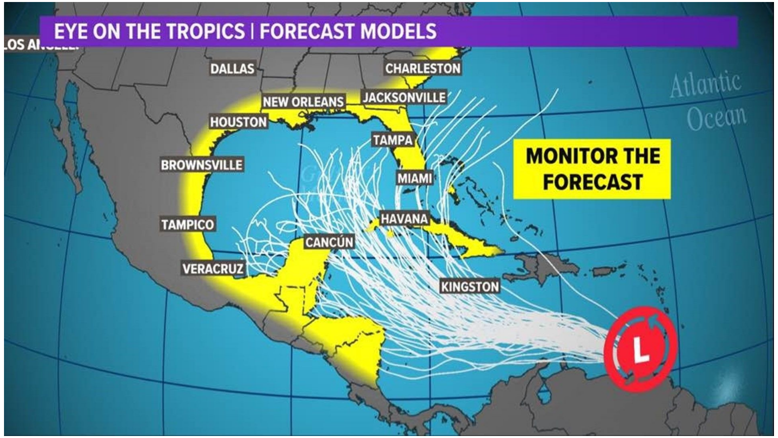

Ensembles based on different initial conditions or model

www.nssl.noaa.gov

Paper Deep Dive

Deep Learning Models for Predicting Wildfires from Historical Remote-Sensing Data

Fantine Huot[*, 1, 2] fantine@stanford.edu
R. Lily Hu[1] rlhu@google.com
Matthias Ihme[1, 2]
Qing Wang[1]
John Burge[1]
Tianjian Lu[1]
Jason Hickey[1]
Yi-Fan Chen[1]
John Anderson[1]

[1]Google Research
[2]Stanford University

https://arxiv.org/pdf/2010.07445.pdf

Motivation

In 2019, 775,000 residences across the United States were flagged as at an "extreme" risk of destructive wildfire, amounting to an estimated reconstruction cost value of $220 billion dollars.

Wildland fires have a significant impact on the global climate, representing 8 billion tons or 10% of global CO_2 emissions per year.

Furthermore, the health impact due to wildfire aerosols is estimated at 300,000 premature deaths globally per year.

There is a real need for novel wildfire warning and prediction technologies that enable better fire management, mitigation, and evacuation decisions.

Goal

Historically, wildfire likelihood has been based on fire behavior modeling across *simulations* by varying feature parameters (e.g. weather, topography) that contribute to the probability of a fire occurring.

Goal: Predict wildfire burning (over the next few days) in a specific location using machine learning

Brainstorm

What kind of data would you want to have to be able to approach this problem?

What kind of methods would you apply?

How would you measure success?

What difficulties might you face?

Data

Aggregated nearly a decade of satellite observations, combining historical wildfire, terrain, vegetation, and weather data to train a deep learning model.

Need data that has

-historical record for training

-high spatial resolution and regular updates for future prediction

Data

Shuttle Radar Topography Mission (SRTM), sampled at 30 m resolution

collection of drought indices derived from the GRIDMET dataset, sampled at 4 km resolution every 5 days since 1979

Collection of daily surface fields of temperature, precipitation, winds, and humidity at 4 km resolution since 1979

daily fire mask composites at 1 km resolution since 2000

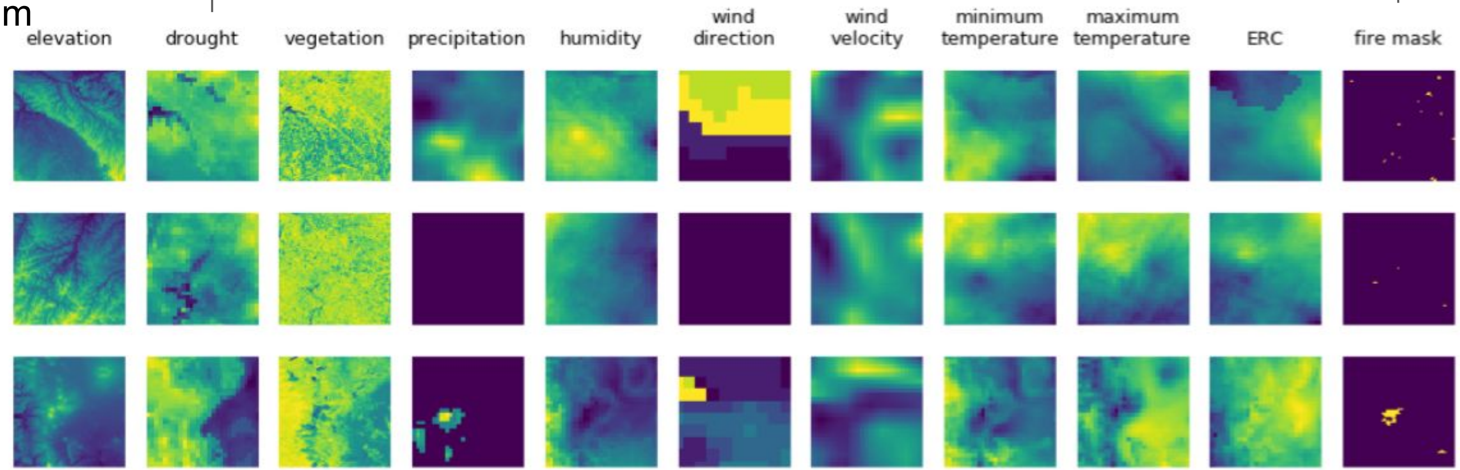

a collection of vegetation indices sampled at 500 m resolution every 8 days since 2012

Data problem

Different spatial resolutions, temporal resolutions, and temporal extents.

Data problem

Different spatial resolutions, temporal resolutions, and temporal extents.

The data is resampled to 1 km resolution, the resolution of the labels. From this data we sample 128 × 128 tiles at 1km resolution. The region of study is restricted to the contiguous United States from 2012 to 2020.

Temporal resolution ??

Framing the problem

Framing the problem

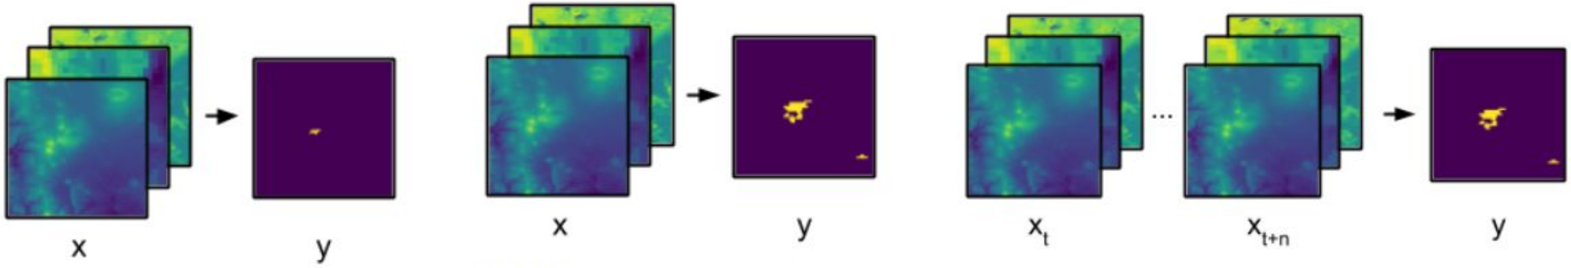

(a) Image segmentation

(b) Image segmentation with aggregated labels.

(c) Sequence segmentation

Figure 2: The fire likelihood estimation problem is framed as three machine learning tasks: (a) an image segmentation problem, where the input features are taken one day prior to a fire. (b) an image segmentation problem with aggregated fire masks, where the input features are taken one day prior to a fire, and the fire masks are aggregated over a week to capture the total burn area, and (c) a sequence segmentation problem. The input features are arranged in week-long sequences to capture the time-component of the data.

“segmentation”?

Framing the problem

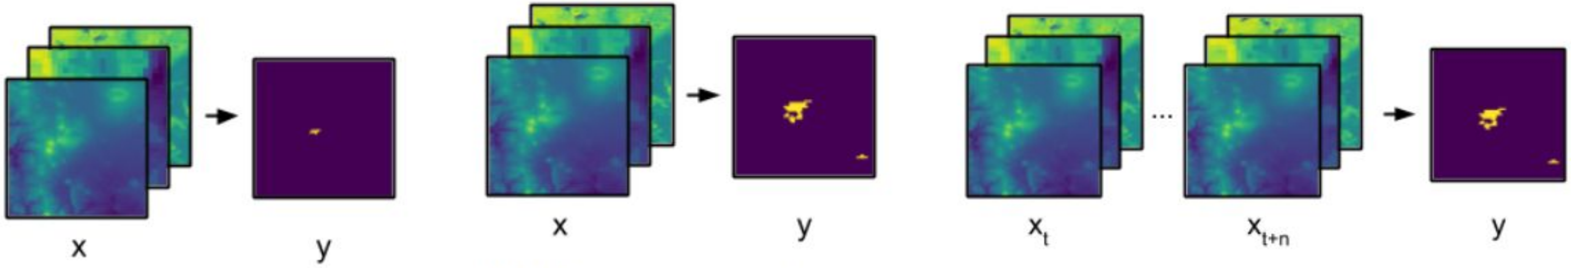

(a) Image segmentation

(b) Image segmentation with aggregated labels.

(c) Sequence segmentation

The dataset for experiment 1 has about 110,000 samples in the training set, 10,000 samples in the validation set, and 10,000 samples in the testing set.

The data sets for experiments 2 and 3 have about 35,000 samples in the training set, 5,000 in the validation set, and 5,000 in the testing set.

Model(s)

Model(s)

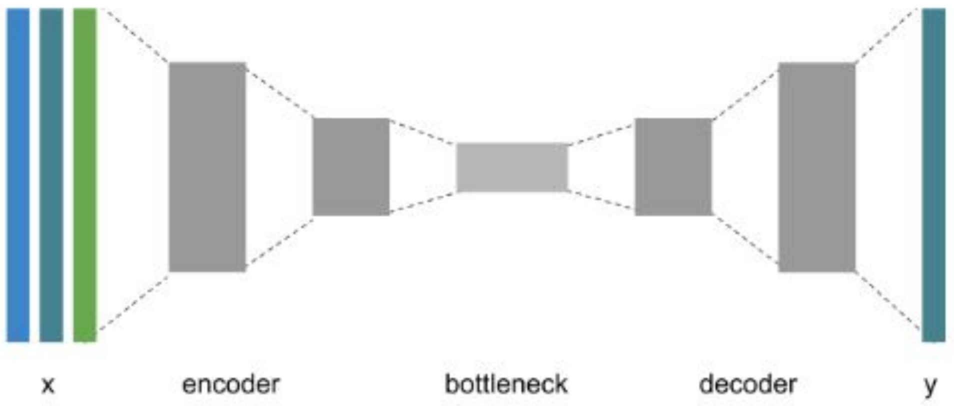

Convolutional “auto”encoder

What an autoencoder really is

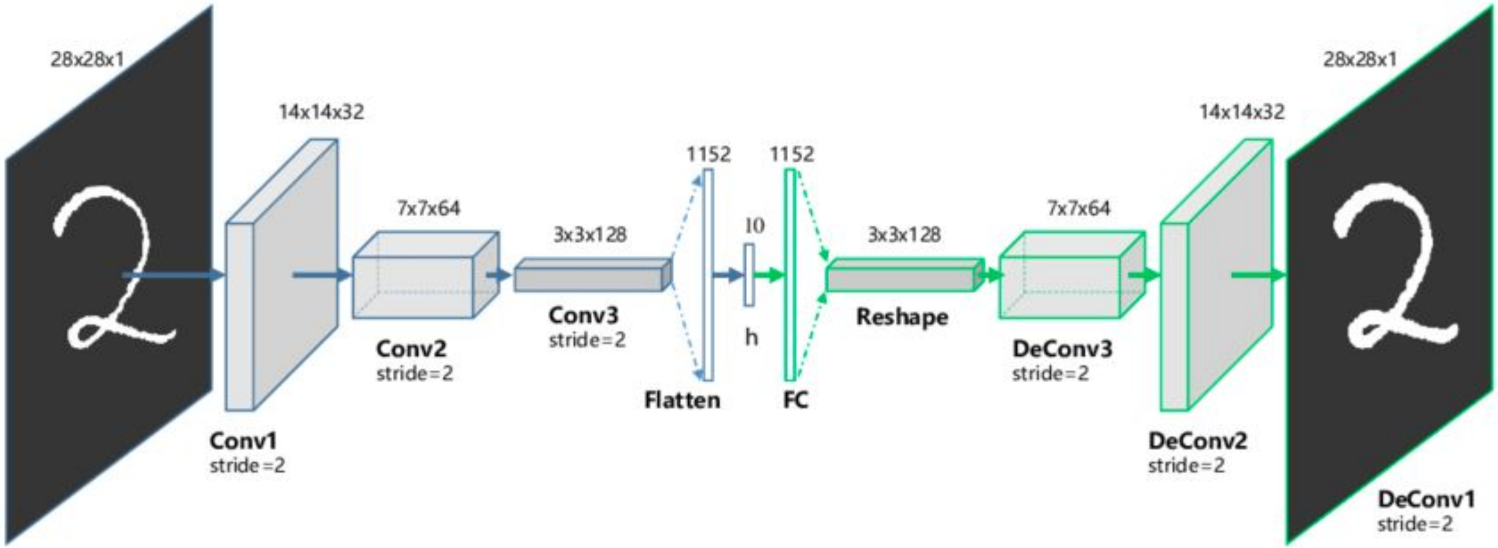

In a true autoencoder, the model is trained to reproduce its input

Model(s)

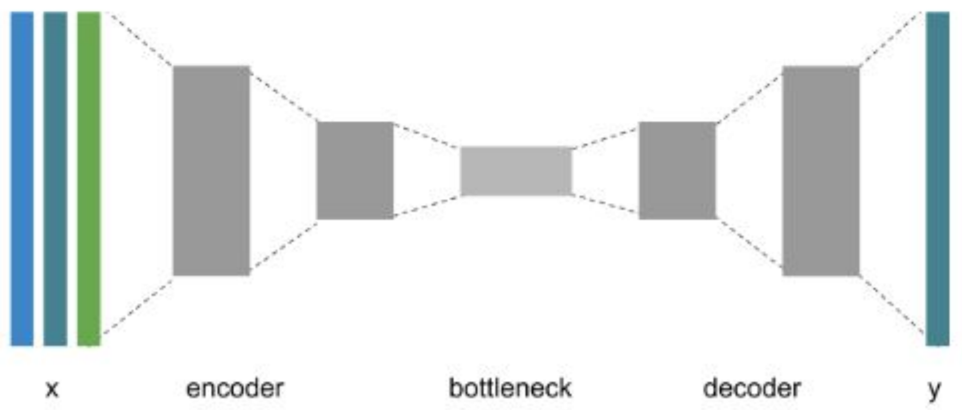

Convolutional ~~“auto”encoder~~ encoder/decoder

Model(s)

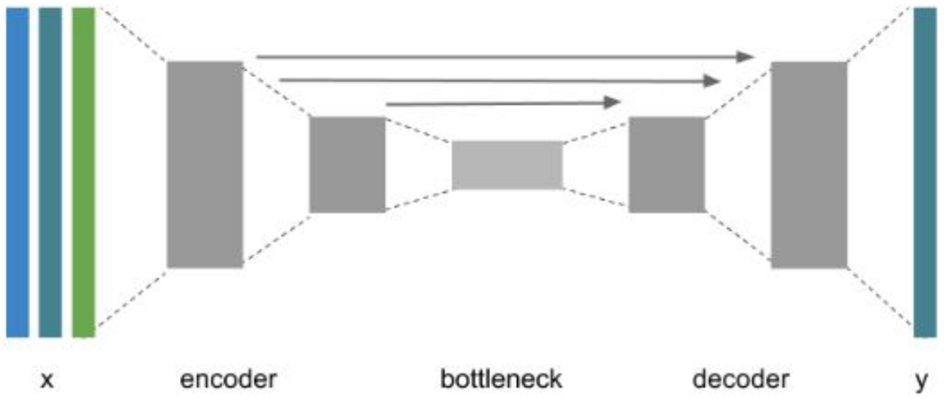

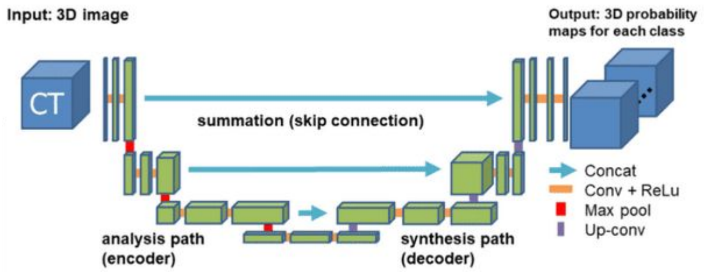

Convolutional U-net

Model(s)

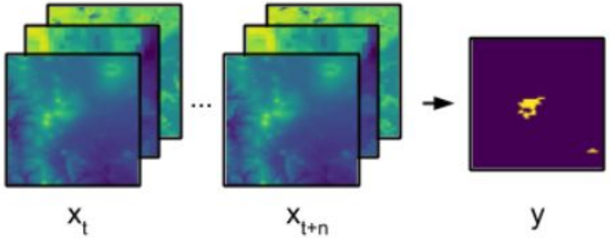

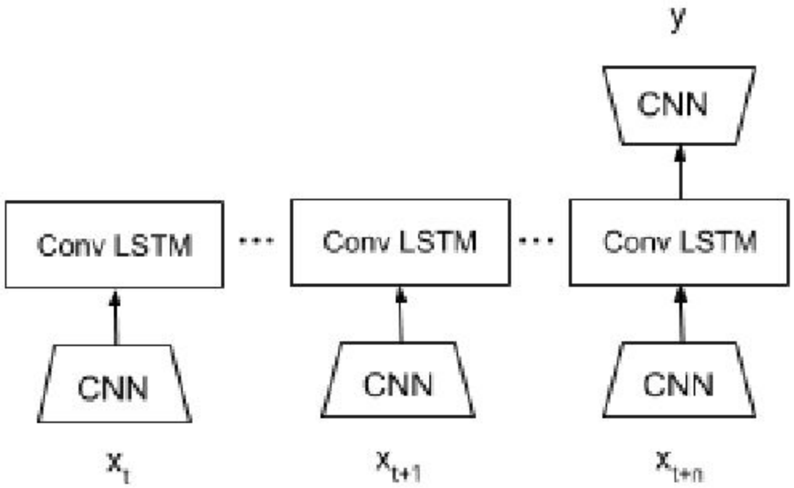

Convolutional LSTM

Model(s)

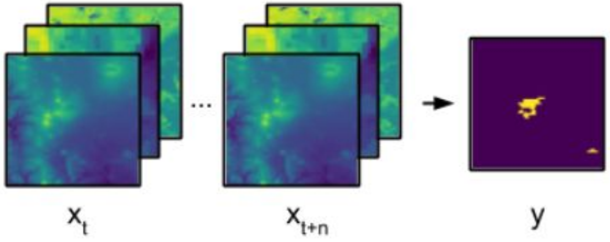

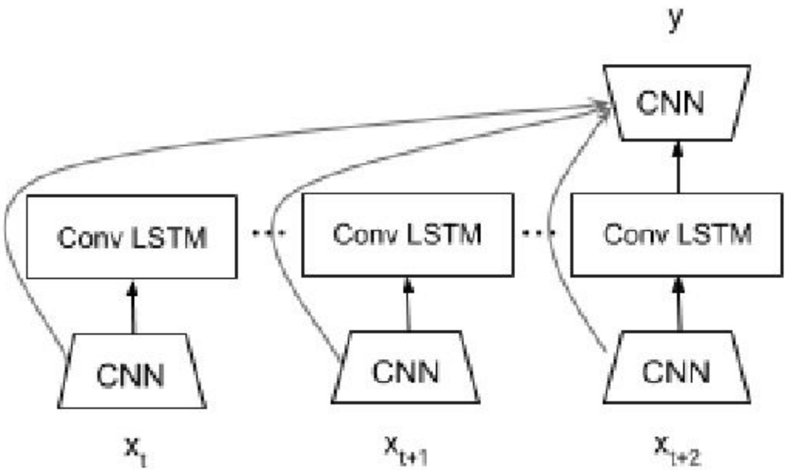

Convolutional U-net LSTM

Training problem!

The vast majority of pixels do not have fire.

Training problem!

The vast majority of pixels do not have fire.

Use “weighted cross-entropy loss” to overcome the class imbalance.

$$CE = -(y \log \hat{y} + (1-y) log(1-\hat{y}))$$

$$WCE = -(\beta y \log(\hat{y}) + (1-y) \log(1-\hat{y}))$$

true class label

model prediction

Training

- Hyperparameters explored via grid search:
  - Model depth
  - Number of filters
  - Batch size
  - Learning rate
  - Beta, weight of positive class

Evaluated on validation set according to AUC

Evaluation

Area under the ROC

Precision

Recall

Intersection over Union

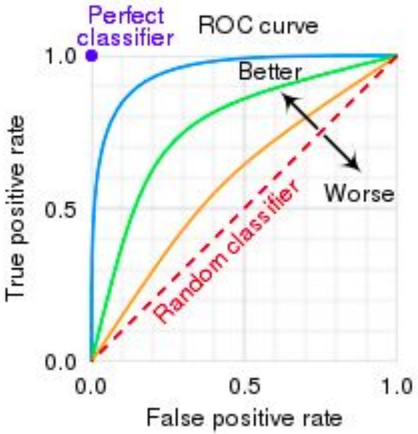

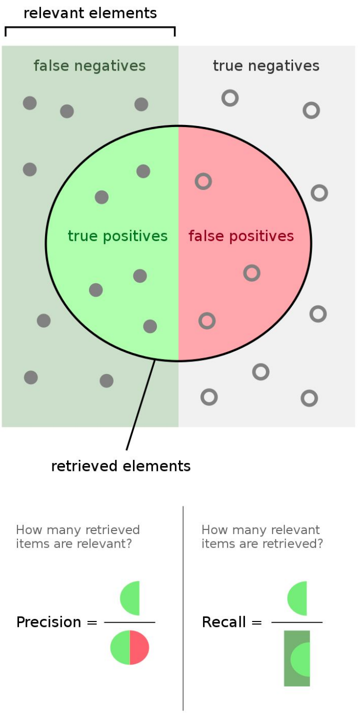

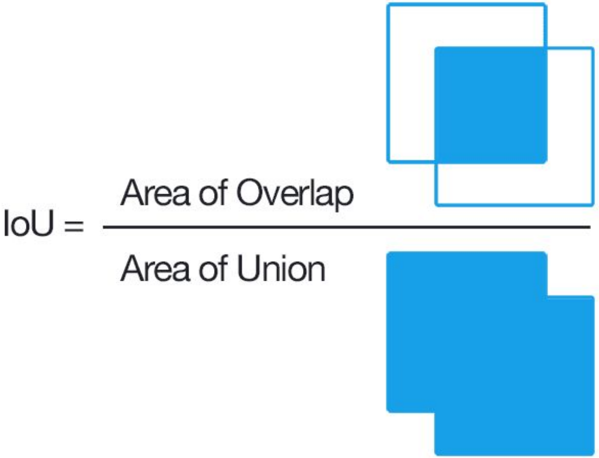

Results

Experiment	Model
Daily Segmentation	Autoencoder U-Net
Aggregated Segmentation	Autoencoder U-Net
Aggregated Segmentation with Sequential input	Autoencoder LSTM U-Net LSTM

Results

Experiment	Model	AUC	Precision	Recall	IoU
Daily Segmentation	Autoencoder	**0.83**	0.53	0.12	0.52
	U-Net	0.72	0.29	0.08	0.52
Aggregated Segmentation	Autoencoder	0.55	0.30	0.05	0.50
	U-Net	0.54	0.29	0.05	0.50
Aggregated Segmentation with Sequential input	Autoencoder LSTM	0.58	0.49	0.05	0.51
	U-Net LSTM	0.60	0.48	0.03	0.50

Results

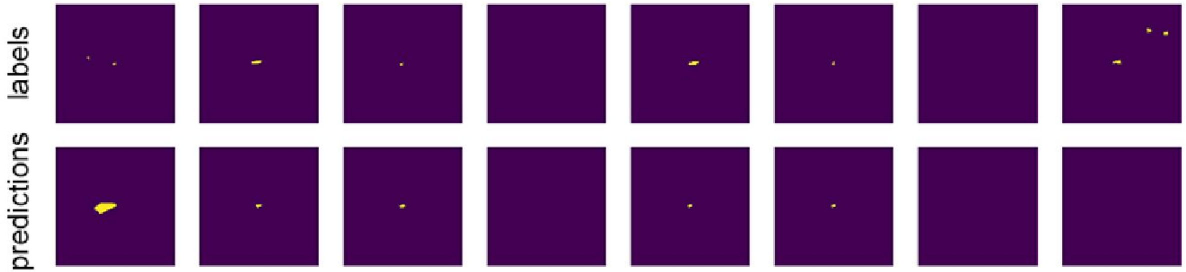

(a) Image segmentation results from the autoencoder model on the test data with daily active fire labels.

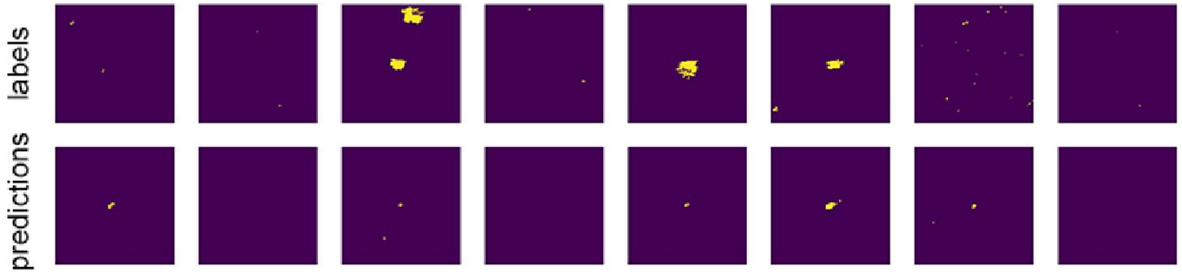

(b) Image segmentation results from the autoencoder LSTM model on the test data with aggregated fire labels.

“The low values are indicative of misclassified pixels at the segmentation boundary between fire and non-fire, or very small fires not being detected, *and thus not representative of the overall performance of the models to predict the presence of fire.”*

“Even without exact segmentation predictions, identifying these regions ahead of time would allow forestry management to allocate resources to specific target regions.”

Results

Experiment	Model	AUC	Precision	Recall	IoU
Daily Segmentation	Autoencoder	**0.83**	0.53	0.12	0.52
	U-Net	0.72	0.29	0.08	0.52
Aggregated Segmentation	Autoencoder	0.55	0.30	0.05	0.50
	U-Net	0.54	0.29	0.05	0.50
Aggregated Segmentation with Sequential input	Autoencoder LSTM	0.58	0.49	0.05	0.51
	U-Net LSTM	0.60	0.48	0.03	0.50

“We observed more severe overfitting with the aggregated segmentation map, likely due to the a smaller dataset size.”

“These experiments would likely improve from data augmentation techniques, for example, by adding noise to the input features. Another idea would be to increase the time span of the data sets, but this would involve combining multiple data sources for some of the input variables due to limited temporal coverage – many satellites started collecting data only within the past few years. Another future step is to expand these experiments beyond the United States to a global scale.”

Framing the problem

“Identifying regions that have high **likelihood** for wildfires is a key component of land and forestry management and disaster preparedness.”

“In particular, assessing the fire **likelihood** – the **probability** of wildfire burning in a specific location – would provide valuable insight for forestry and land management, disaster preparedness, and early-warning system”

“A use case for models trained on this compiled data is to predict the **likelihood** of fires given recent remote sensing data.”

“Positive events are only labeled when there are actual fires, but the dataset also contains tiles with high fire likelihood that did not have actual fires, such as days before fires or neighboring areas that did not catch fire because firefighters intervened”

5 Conclusions

We demonstrate the potential of deep learning approaches for estimating the fire likelihood from remote-sensing data. We create a data set by aggregating nearly a decade of remote-sensing data, combining features including weather, drought, topography, and vegetation, with historical fire records. Our trained models can successfully distinguish between fire and non-fire conditions. Going forward, this data-driven approach could be valuable for wildfire risk estimation, and could be incorporated into wildfire warning and prediction technologies to enable better fire management, mitigation, and evacuation decisions. Beyond the wildfire likelihood problem, the described workflow and methodology could be expanded to other problems such as estimating the likelihood of regions to droughts, hurricanes, and other phenomena from historical remote-sensing data.

Further resources

https://www.quantamagazine.org/machine-scientists-distill-the-laws-of-physics-from-raw-data-20220510/

Why rain is hard to predict (DeepMind)
https://www.youtube.com/watch?v=snCo0Z0dt-k

Summary

Sparse Bayesian learning, using relevance vector machines (RVMs)

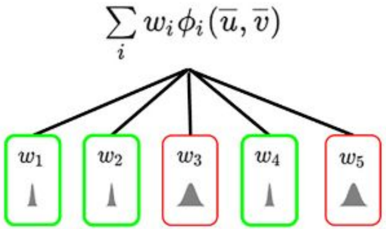

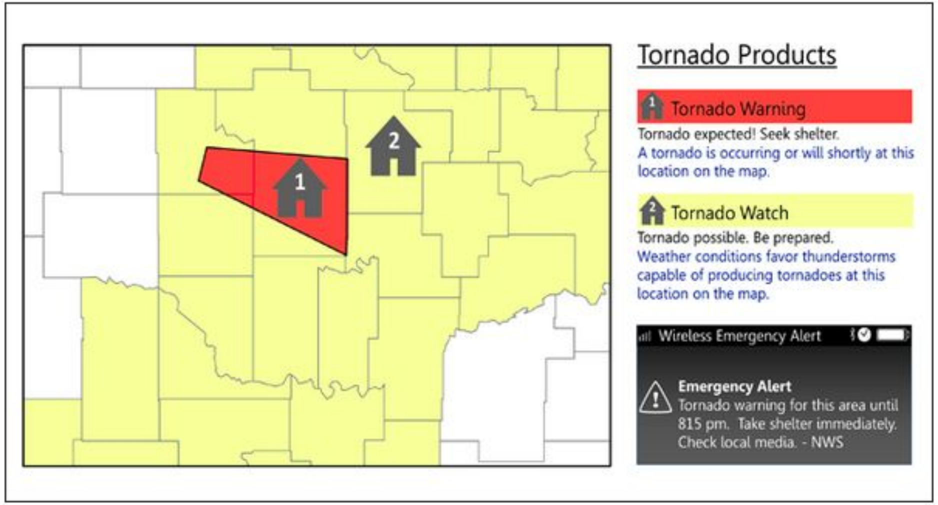

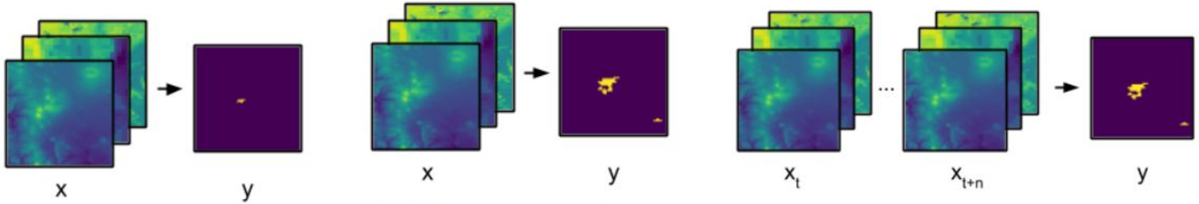

(a) Image segmentation (b) Image segmentation with aggregated labels. (c) Sequence segmentation