Machine Learning based Smoke Detection

May 29, 2019

Kinshuk Govil

Overview

- Satellite vs. Terrestrial imagery
- Machine Learning setup
- Sample images
- Results
- Future

Satellite vs. Terrestrial imagery

- Satellites imagery
  - Pros: Can scan large area in single image
  - Con: Coarse grain resolution (single pixel is > 500 meters wide or more)
  - Geostationary satellites don’t have time gaps, but resolution is worse
  - Early detection requires looking for single pixel with abnormal value
- Terrestrial imagery from vantage points
  - Visible light optical cameras installed on fire towers on mountain tops
  - Pros: Good resolution (single pixel is ~10 meters wide at 10 miles away)
  - Pros: ~1 minute time gap
  - Cons: Terrain/topography and visibility limit range to ~10 miles => tower every ~200 sq miles
  - Early detection requires recognizing shape, color, or motion with >1000 pixels of smoke plume
  - Commercial smoke detection software uses hand coded algorithms (been around for decade)
  - We have achieved better accuracy using modern machine learning based image object recognition technology

Machine Learning Setup

- Supervised training with two classes: smoke and not-smoke
- Training Google's Inception v3 model architecture with wildland images
  - Originally designed to detect 1000 objects such as dog, cat, cup, car, barn, castle, etc..
  - Inception v3 expects 299x299 pixel images (images are resized if needed)
  - Shrinking large images would lose smoke, so segment into overlapping 299x299 squares
- Training data for smoke
  - Match Calfire's historical fire data with camera locations to search archived images
  - Volunteers mark smoke boundary rectangle
  - Generate 10 segments (flipped and recentered) per smoke image
    - 2 (Flip + original) x 5 (center, top left, top right, bottom left, bottom right)
    - ~6,000 manually labeled images => ~60,000 smoke segments
- Training data for not-smoke
  - Segments of first smoke image for each fire where segments don't overlap smoke rectangle
  - False positive segments from earlier trained models

Sample test set true positive result

- Two squares > .5
- Clouds have low scores

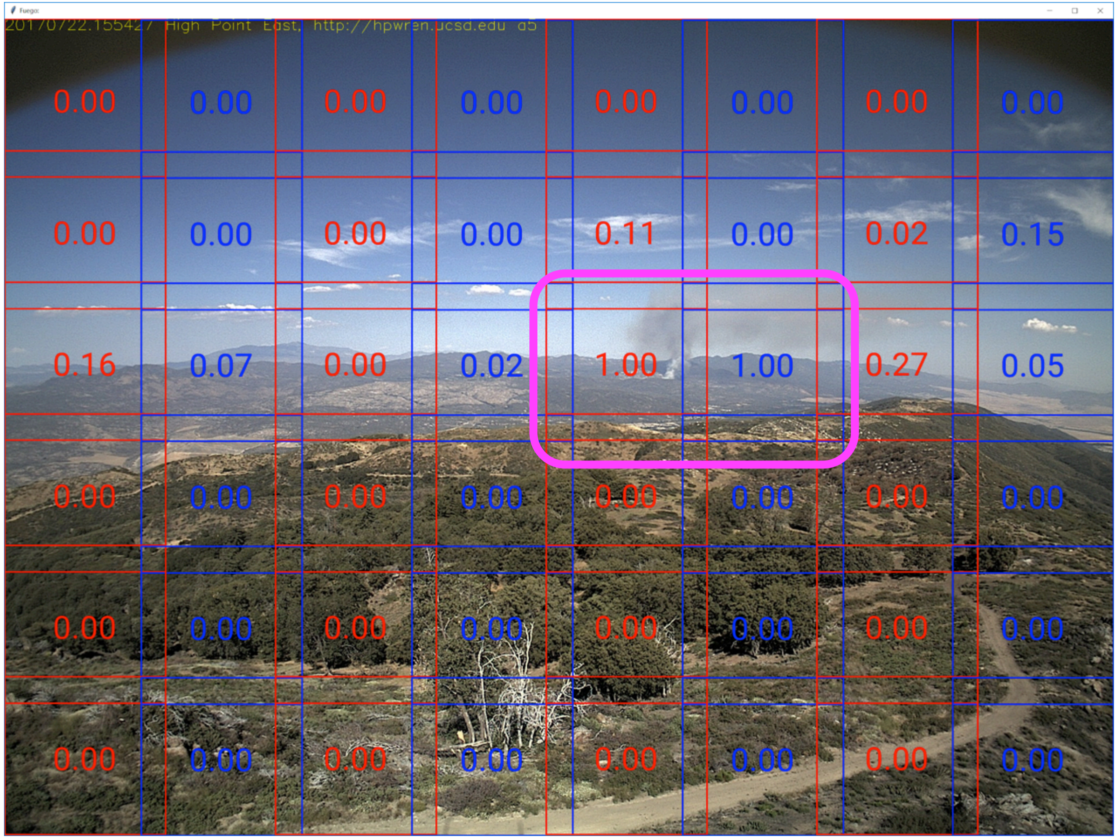

Sample test set true negative result

- Highest score: 0.14
- Fog correctly ignored

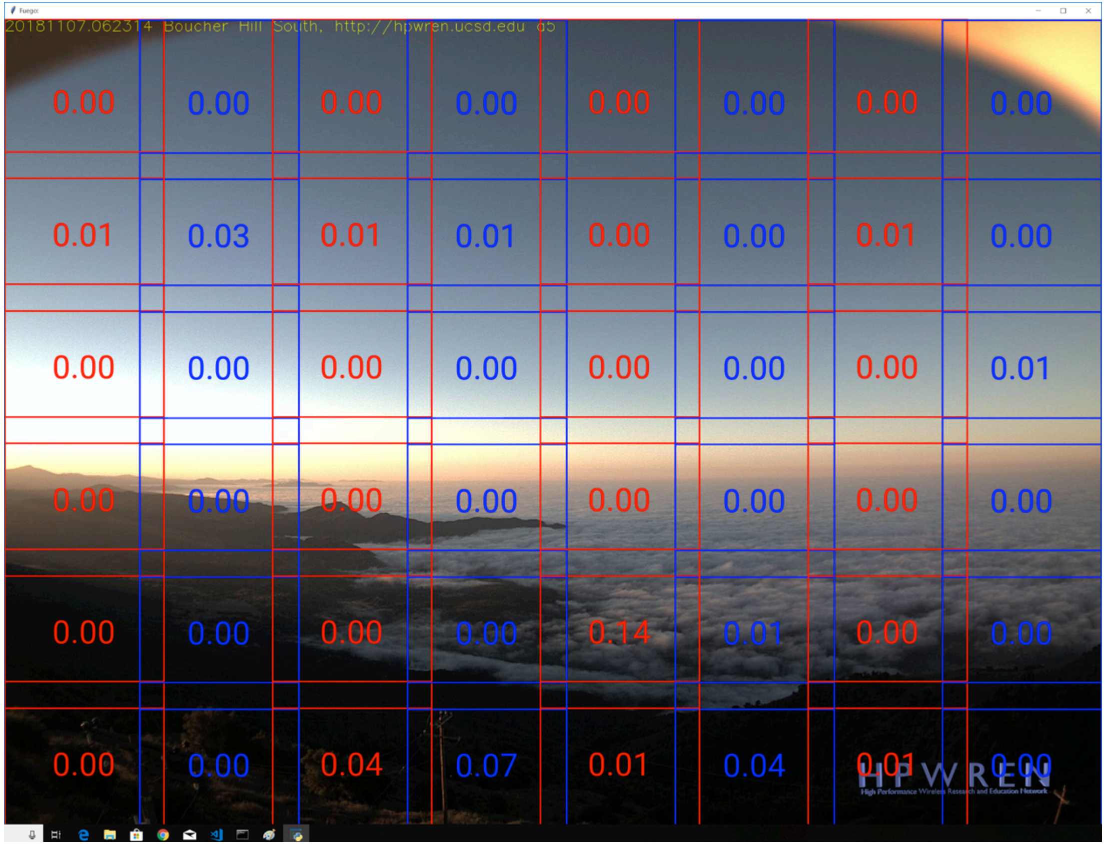

Sample test set false negative result

- Too similar to fog?

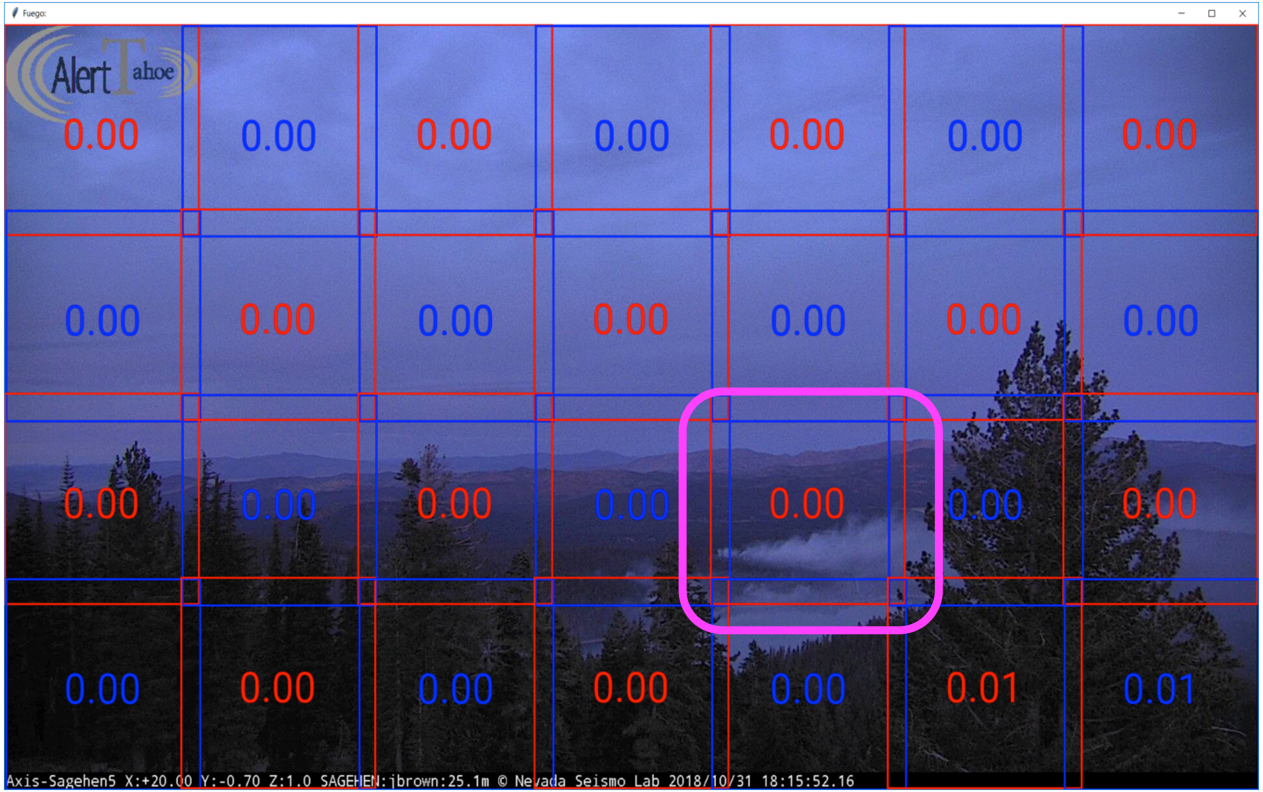

Sample test set false positive result

- Highest score: 0.65
- Such segments are sent for retraining
- Increasing threshold based on last few days of data filters out 60% of these

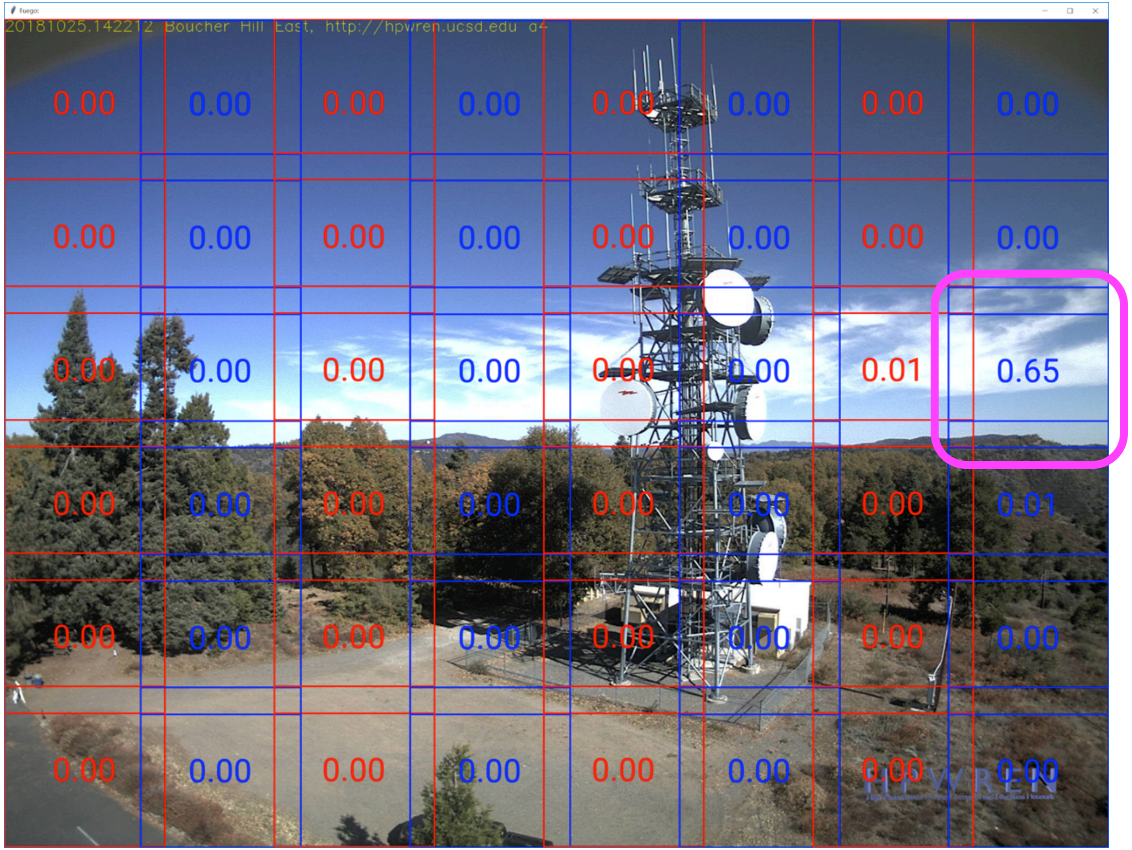

Results

- ML accuracy on test set
  - 250 full sized images **not** used in training
    - 100 smoke (ideally should be much bigger)
    - 150 non smoke
  - Requirements
    - Non smoke: Every segment must be classified as not smoke
    - Smoke: At least one segment must be classified as smoke
  - Types of models:
    - Top most layer, fine-tune, full training
  - F-1 score: 0.85
- Able to detect 2018 Holy fire from image 2 minutes before 9-1-1
- False positive rate: once per camera field of view per 2 days

Future

- Improve accuracy of this model
  - Continue to retraining from false positives
  - Get more smoke images from volunteers
- Experiment combining with new models for terrestrial images
  - Subtract images to capture motion of smoke
  - Very Near IR up to 1um (standard silicon without IR cut filter): 1 pixel ~ 10m
  - IR 0.7-1.7um (InGaAs or CQD), 7-14um (micro bolometer): 1 pixel ~ 100 m
    - Image object recognition unsuitable for detecting few pixels
    - Different ML approach may help
    - Satellite image detection techniques may help
- Combine with satellite imagery to leverage best of both
  - Higher confidence alerts if something detected in both systems
  - For one sided detections: Sending both sets of images will help people decide