living planet symposium | BONN 23–27 May 2022

TAKING THE PULSE OF OUR PLANET FROM SPACE

Exploring specificities of ML algorithms for Fire Risk Prediction

Stella Girtsou, Alexis Apostolakis, Dr. Giorgos Giannopoulos, Dr. Charalampos Kontoes

25.05.2022

Problem formulation

- Predict the risk of fire occurrence in an area for a day *k*, exploiting information for the area exclusively gathered up until day *k-1*
- Essentially handled as binary *{fire, no-fire},* due to label availability (historical fires)
  - Ideally, a reliable confidence (probability of risk) level should be output
- Each area corresponds to a 500m cell of a grid
  - Grid covers the whole Greek territory
- Detailed historical data from 2010-2020
  - > 800M instances

Problem formulation - Features

EO: NDVI, EVI, LST

Meteorological features: Temperature (max, min, mean), Dew Temperature (max, min, mean), Wind speed (max, dominant), Wind direction (wind_direction, dominant_direction), Cumulative Precipitation

Geomorphological/natural features: DEM (DEM, aspect, slope, curvature), Land use/Land cover

Fire history, Spatially smoothed fire history, Month of the year, Week Day

Motivation

- Essential tool for **daily operational organization** of fire services
- Current service of Civil Protection daily maps
- Need for higher spatial resolution

ΧΑΡΤΗΣ ΠΡΟΒΛΕΨΗΣ ΚΙΝΔΥΝΟΥ ΠΥΡΚΑΓΙΑΣ ΠΟΥ ΙΣΧΥΕΙ ΓΙΑ
Κυριακή 01/08/2021

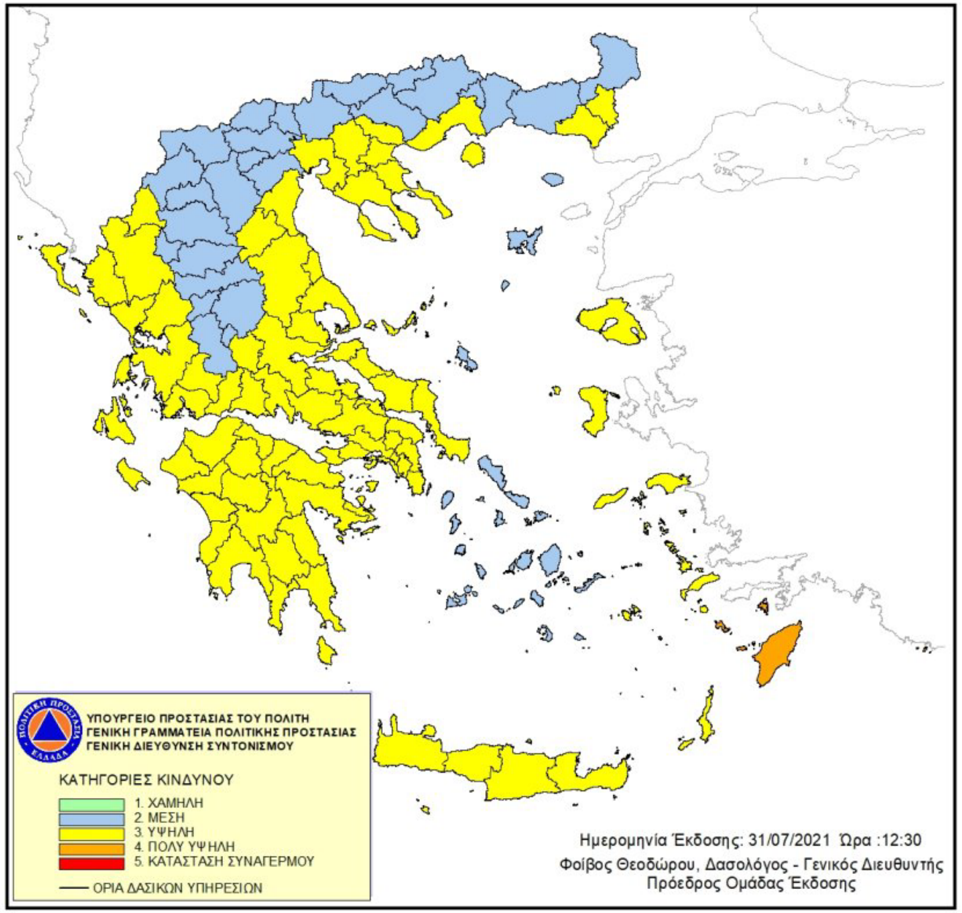

Motivation – deliver meaningful results

- Predict most of the fires
- Try not to predict the majority of the territory (country) as fire

Translation: a good balance between sensitivity/specificity*

- Ideal: {>95%, >90%}
- Realistic:
  - {>90%, >70%}
  - {>80%, >80%}
  - Depends on the exact application setting/needs

- *Percentage of actual fires (resp. no-fires) we correctly predict

BEYOND Centre of EO Research and Satellite Remote Sensing

esa

Domain specificities

Extreme data imbalance → Ratio of ~1:100K between fire/no-fire cells

Large data scale → Challenging to properly perform model selection

830M instances for 11 years

Year	August No Fire	August Fire	Sum Jun-Sep No Fire	Sum Jun-Sep Fire
2010	11687055	347	45995051	607
2011	11685953	1468	45993489	2202
2012	11685532	1816	45992810	2806
2013	11686833	599	45994470	1233
2014	11687130	304	45994809	899
2015	11687290	144	45994915	793
2016	11687188	246	45993758	1950
2017	11686508	926	45994210	1498
2018	11687345	87	45995092	598
2019	11562808	386	45100739	631
2020	11560400	221	44926467	749

Ratio ~1/17.000

Domain specificities

Absence of fire → Areas that should have a fire occurrence but did not by *chance*- lack of impossible to capture features (i.e. a person's decision to start a fire, a cigarette thrown by a driver, a lightning)

Spatiotemporal correlations → Adjacent cells are expected to be nearly-identical
Previous years' incidents might affect the short-term behaviour of an area

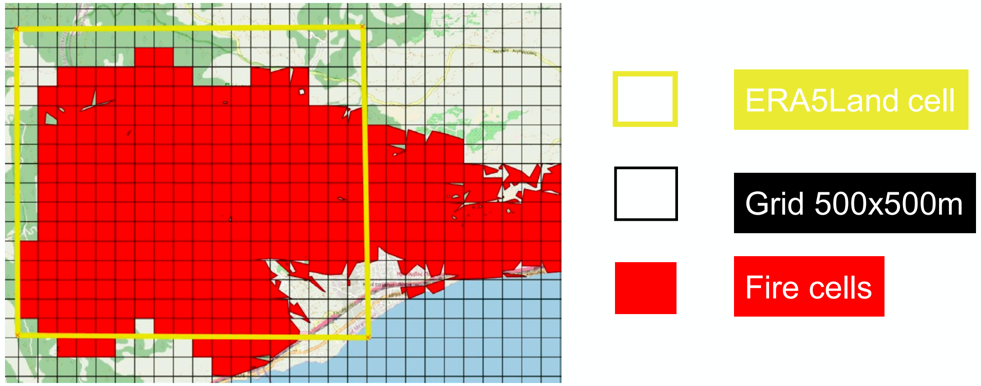

Domain specificities

Concept drifts

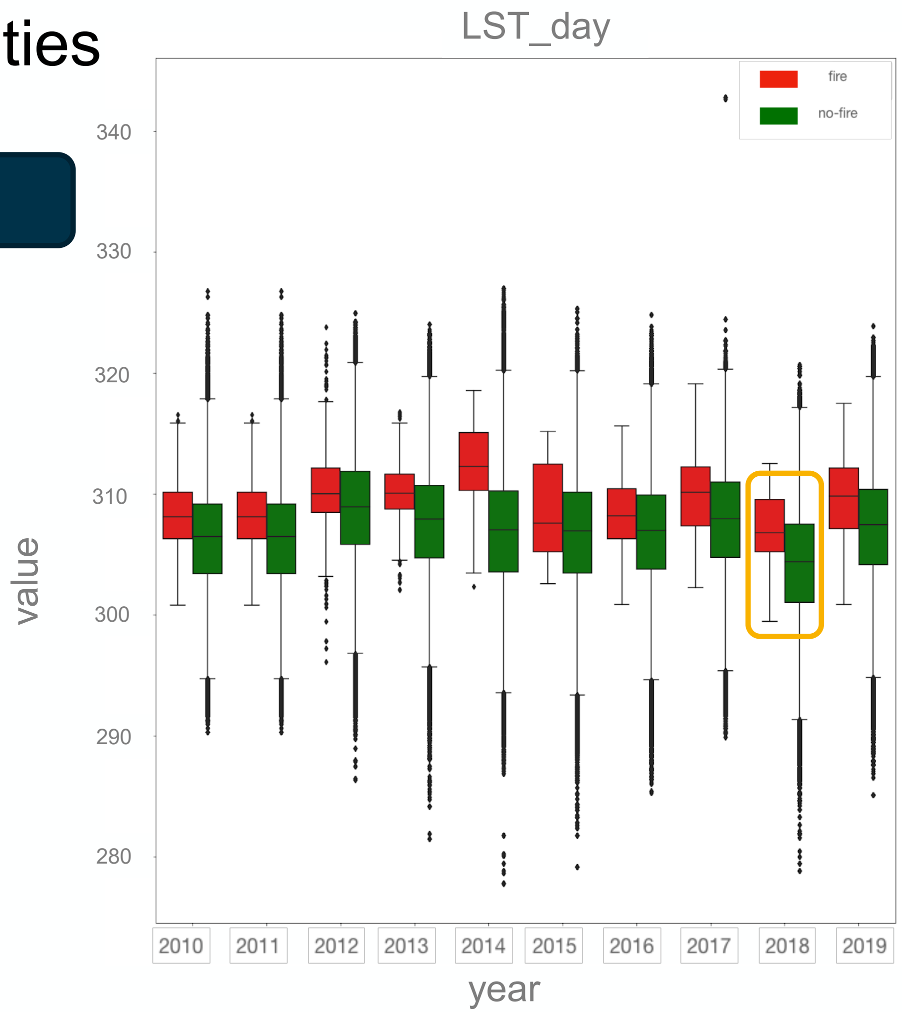

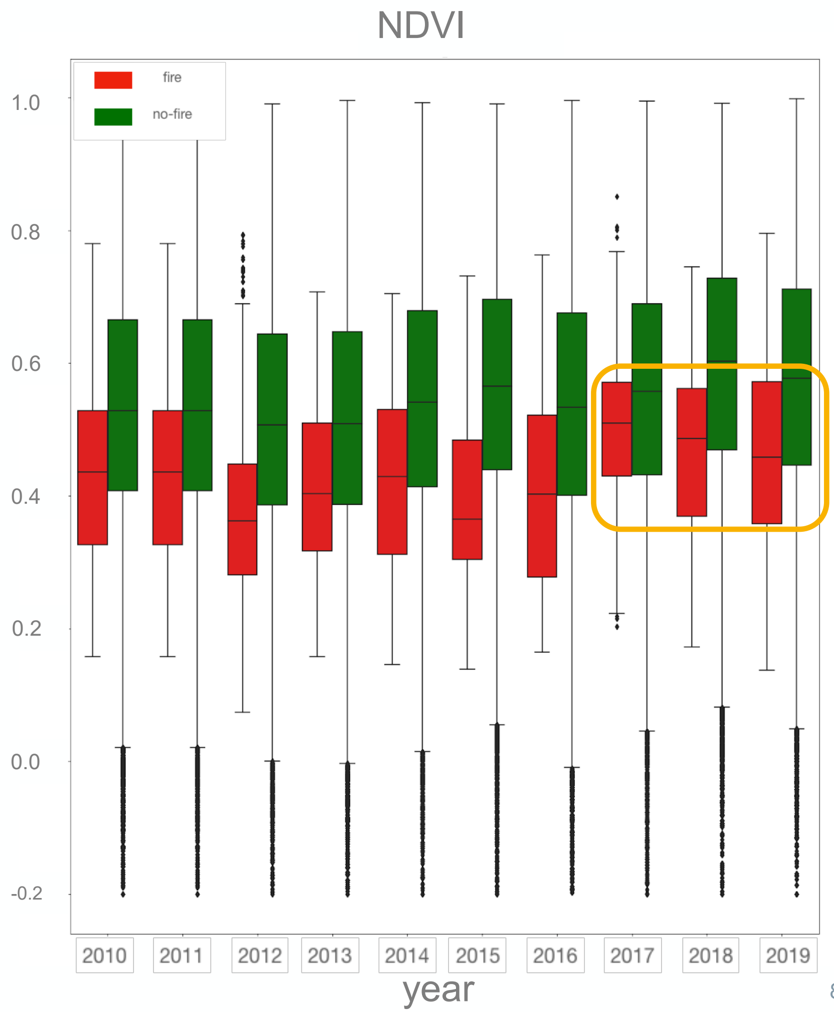

Current approach

Establishment of a complete ML workflow

- Two alternative schemes for cross validation
  - Default:
    - Consider all the fire (minority) instances of the training set
    - Geographically sample the no-fire (majority) instances to create a balanced set
    - Perform k-fold cross-validation and select models on the average best validation scores
  - Alternative:
    - Make the training set balanced, but keep the validation sets highly imbalanced (1/10)
    - Adjust so that each training set precedes the respective validation set on a yearly level
    - Perform model selection on highly imbalanced folds closer to the real distribution
- Proper dataset splitting for model selection and evaluation
  - Ensure that events from the same day/fire event are not distributed in different folds

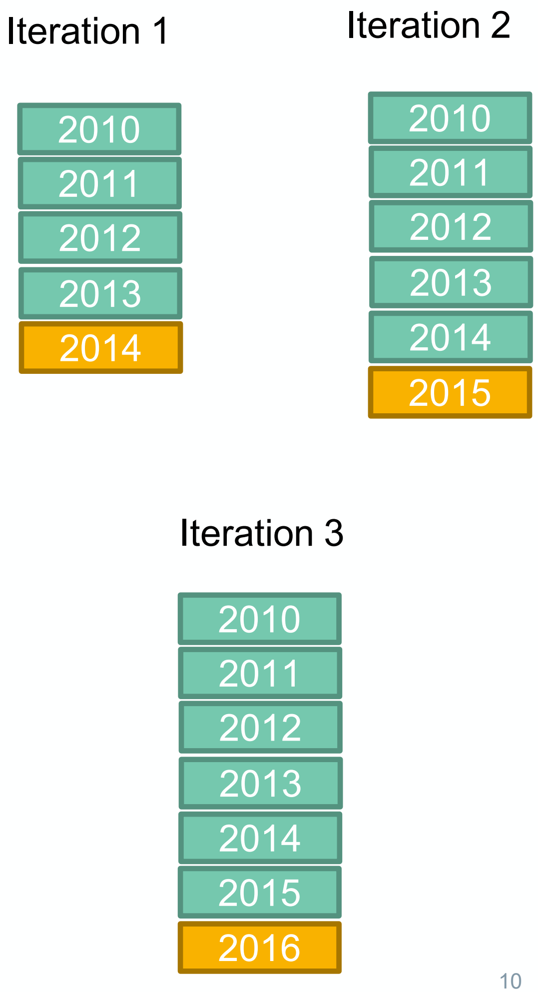

Establishment of a complete ML workflow

- Adjusted evaluation measures for model selection
  - Evaluated on the validation set
  - Variable weighting between sensitivity and specificity
- Exploration of a large hyperparameter space for each adopted soa algorithm
  - RF, XGBoost, ExtraTress, shallow NNs

Results

- Targets (sensitivity/specificity):
  - {>90%, >70%}
  - {>80%, >80%}
- Achieved:
  - {90%, 66%}, {93%, 62%}
  - {82%, 71%}, {79%, 76%}

- Agility on balancing the trade-off between sensitivity/specificity
  - Via combinations of cross-validation schemes and model selection evaluation measures
- A proper problem formulation and baseline methodology

Exploring specificities of ML algorithms for Fire Risk Prediction

Success stories

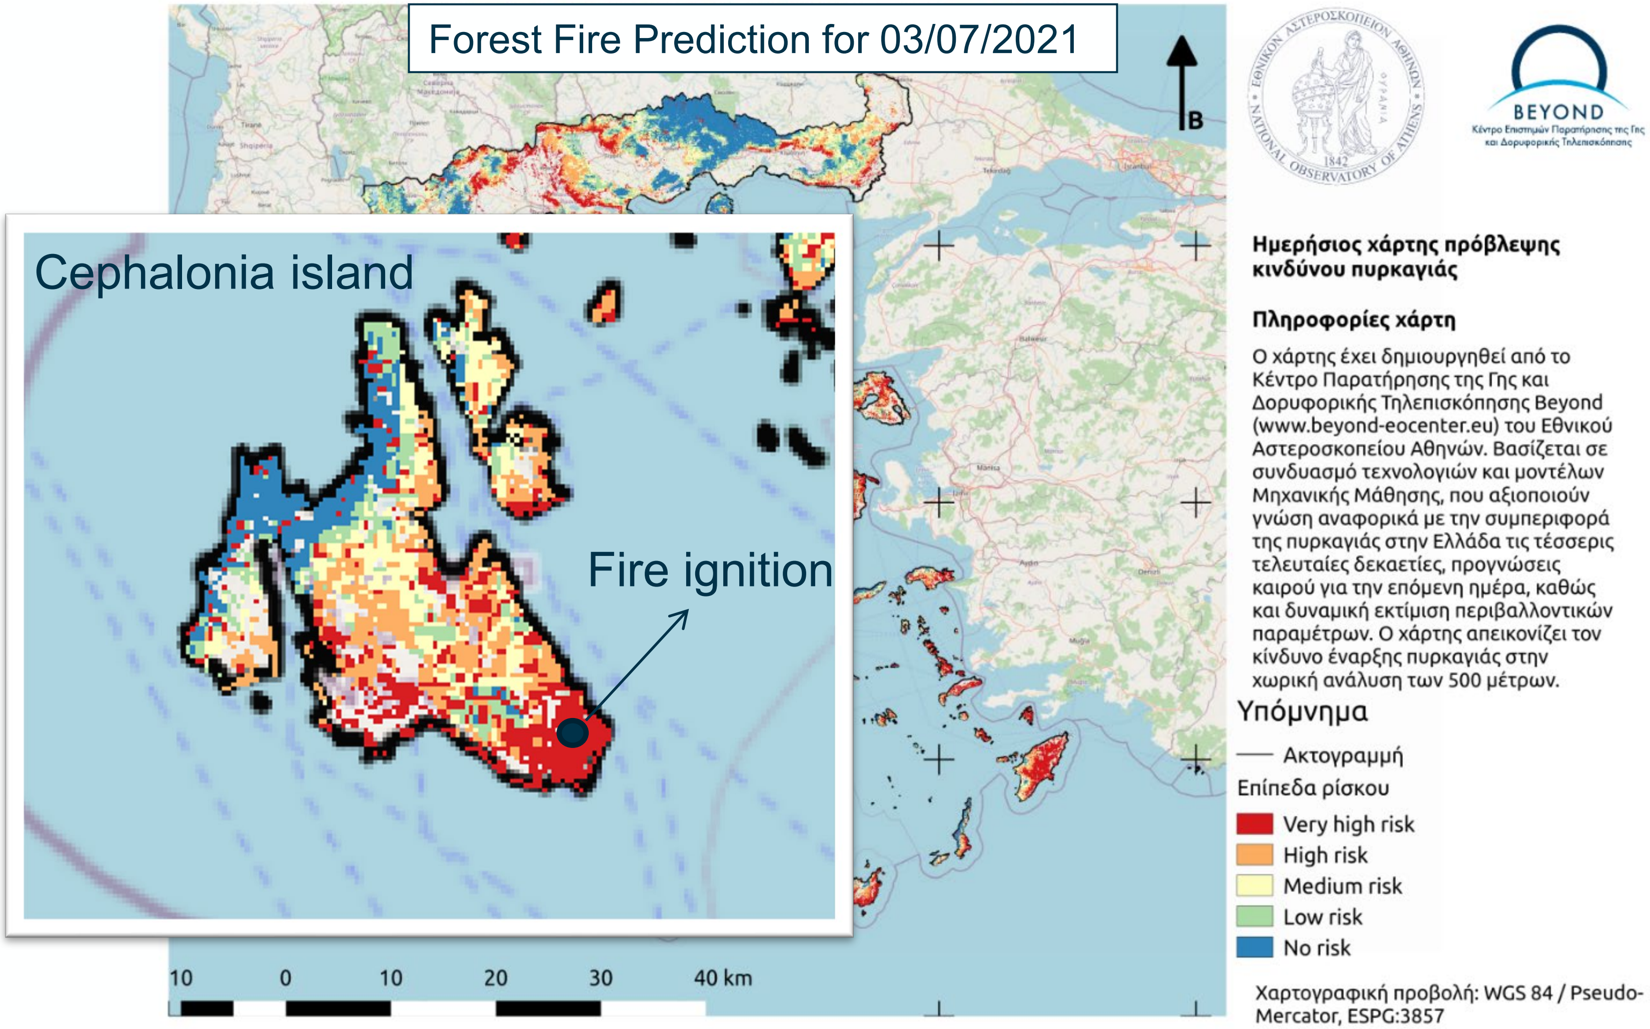

Success stories

Burned area: ~50,000 ha
Active for 8 days

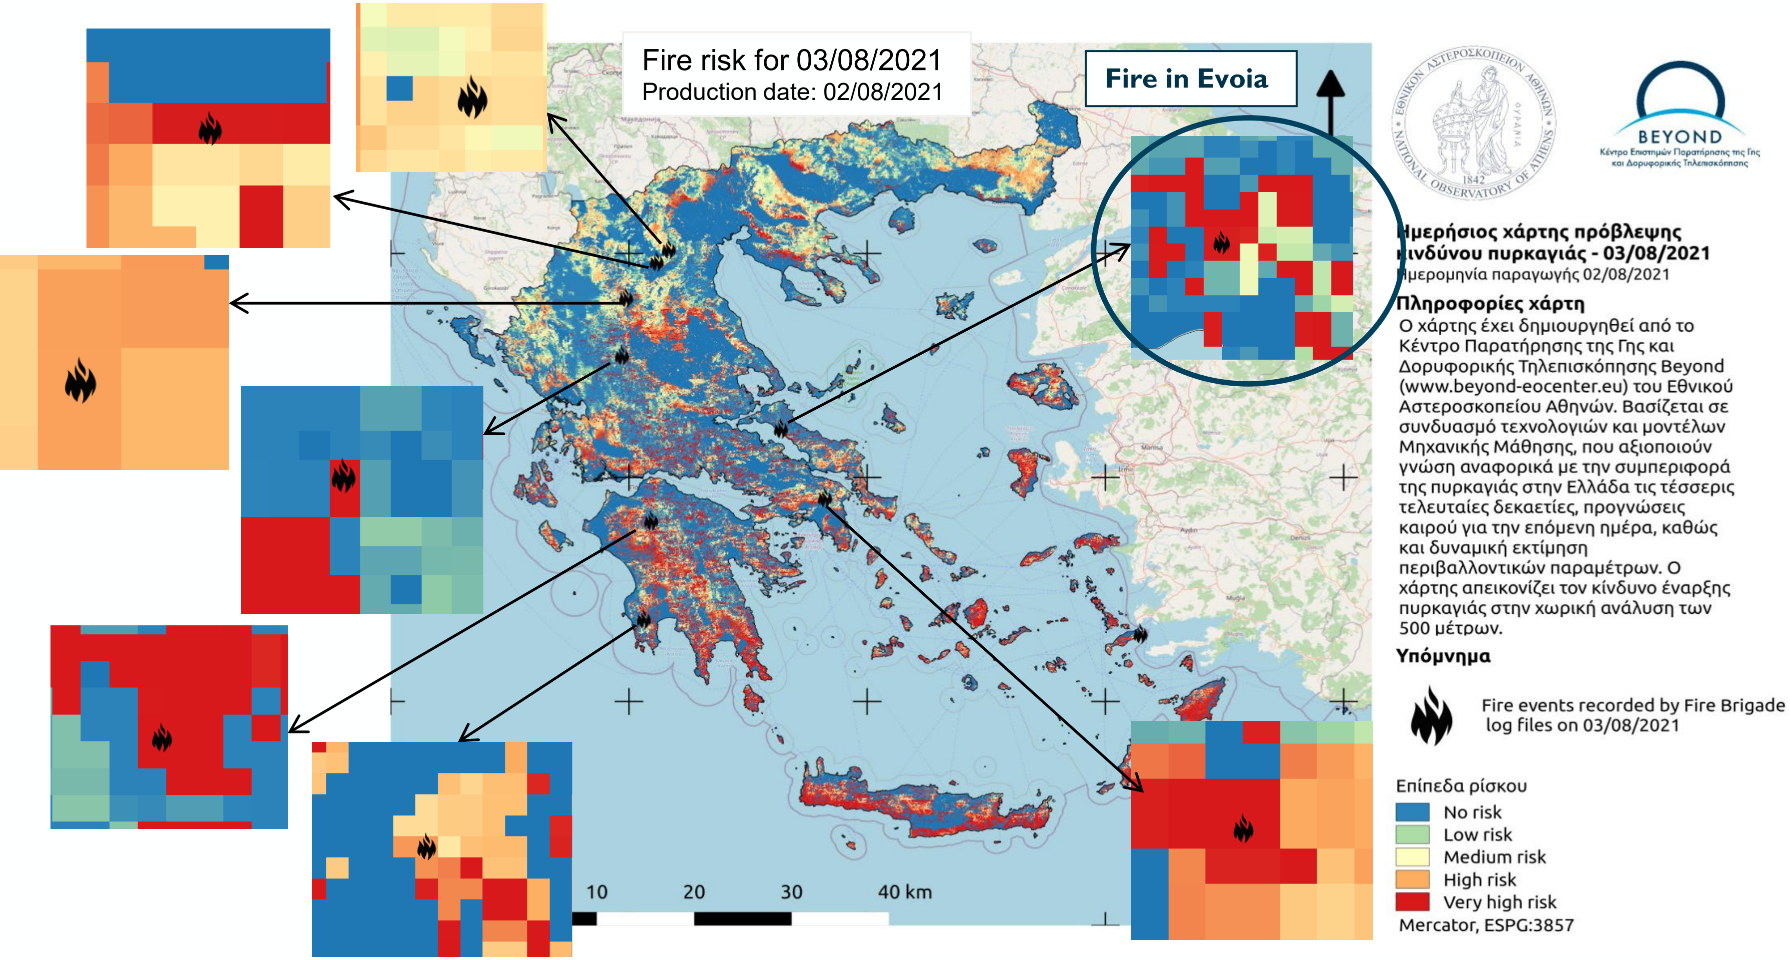

Success stories

Burned area: ~7,000 ha
Active for 3 days

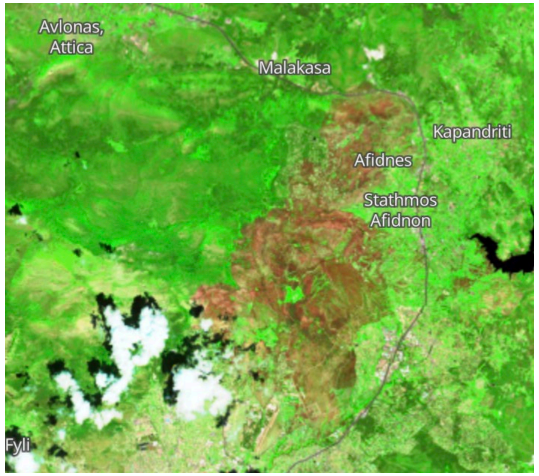

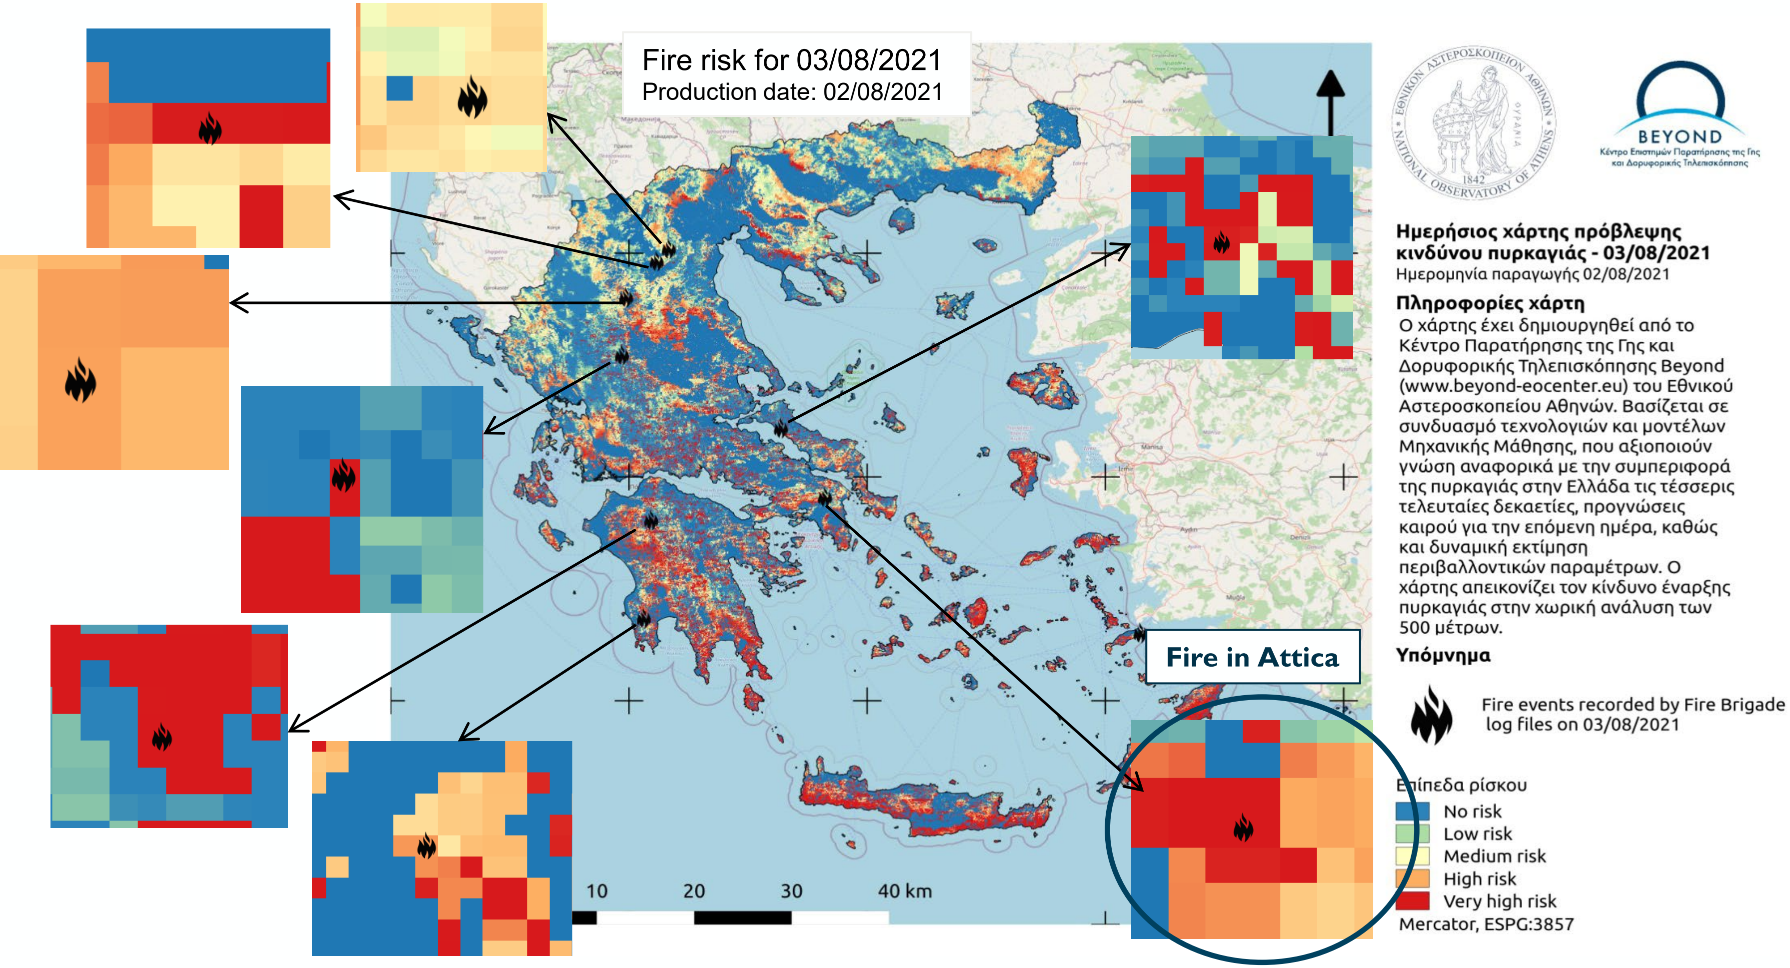

Ongoing work

Directions

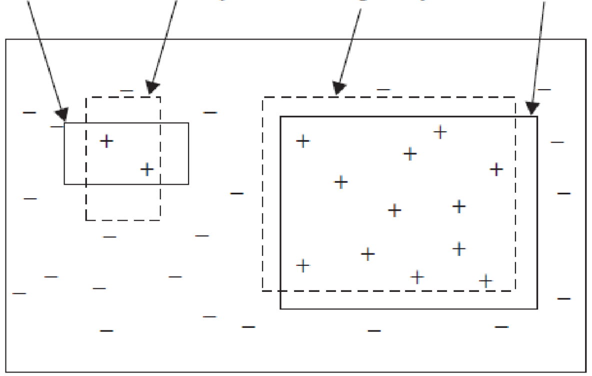

Imbalanced Learning_ Foundations, Algorithms, and Applications (2013, Wiley-IEEE Press)

- Handle absence of fire phenomenon
  - No-fire instances that are very close to fire instances
    - Problematic for learning proper boundaries
    - Reduces specificity by default
- Better handle imbalance
  - Existing schemes are only half-measures
  - Training/validation/test on different distributions
- Examine rare cases and small disjuncts
  - Indications that fire instances form discrete clusters within the hyperspace
- Handle data sizes
  - Try to limit undersampling as much as possible to exploit the whole dataset

Approach: Siamese NNs

- Architectures that aim at learning a similarity function
  - Comprise of parallel NN architectures that receive different inputs but learn the same parameters
- SNNs provide the framework for handling several of the aforementioned issues
  - Particularly triplet loss based SNNs
  - Input as triplets of {anchor, positive, negative}

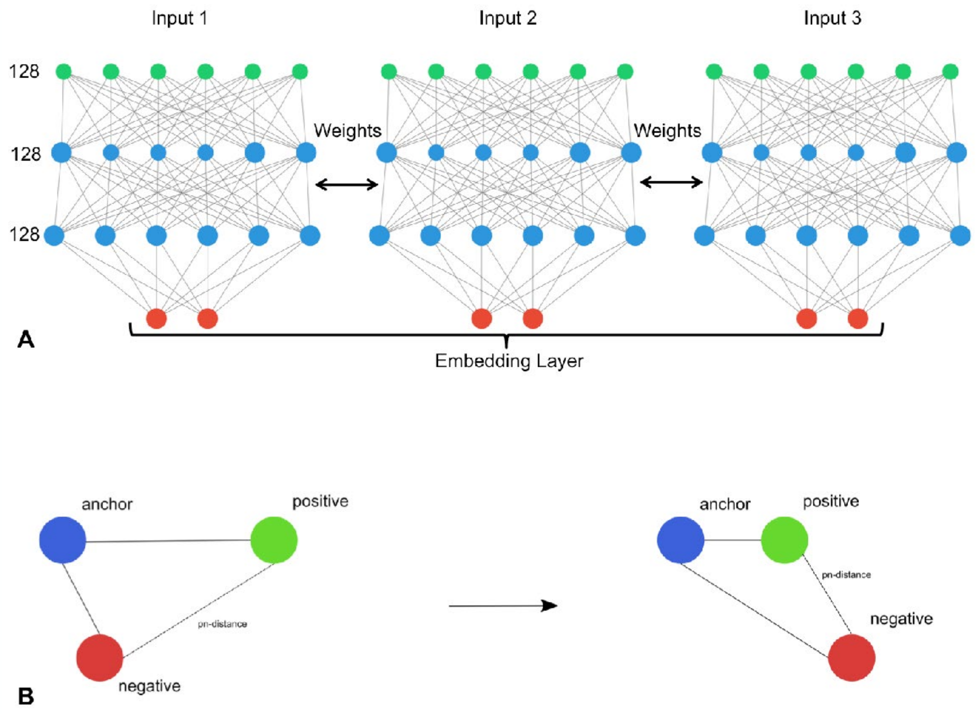

Structure-preserving visualisation of high dimensional single-cell datasets

Approach: Siamese NNs

- **Absence of fire** and **extreme imbalance** can be handled to some extent by properly constructing {anchor, positive, negative} triplets
  - Properly adjust the triplet generation function
  - Hard negatives can be ignored or transformed into positives
  - Semi-hard should probably be emphasized
- Variations of undersampling techniques can be combined
  - E.g. Tomek links
  - Removing majority instances
  - Transforming majority into minority instances

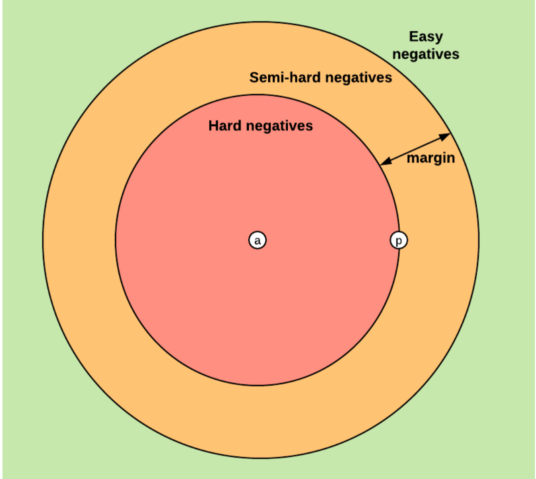

https://medium.com/@enoshshr/triplet-loss-and-siamese-neural-networks-5d363fdeba9b

Illustration of a Tomek link

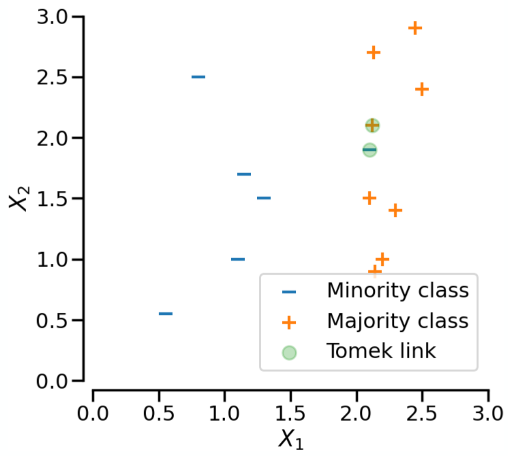

https://imbalanced-learn.org/stable/under_sampling.html

Approach: Siamese NNs

- Initial findings
  - Vanilla Siamese* reached similar effectiveness scores with tuned baseline ML models
    - Without any triplet tuning or over/undersampling
    - With moderate network tuning

* https://bering-ivis.readthedocs.io/en/latest/metric_learning.html

Thank you!

Questions?