Deep Learning and GeoSpatial Data

@o_courtin

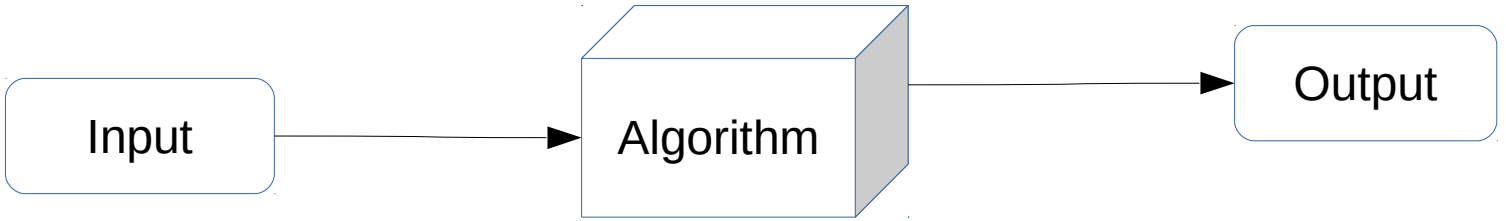
Input
Algorithm
Output

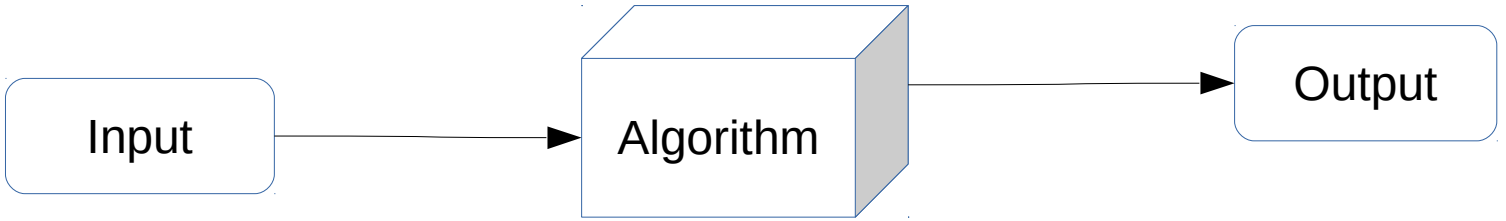
Input
Algorithm
Output

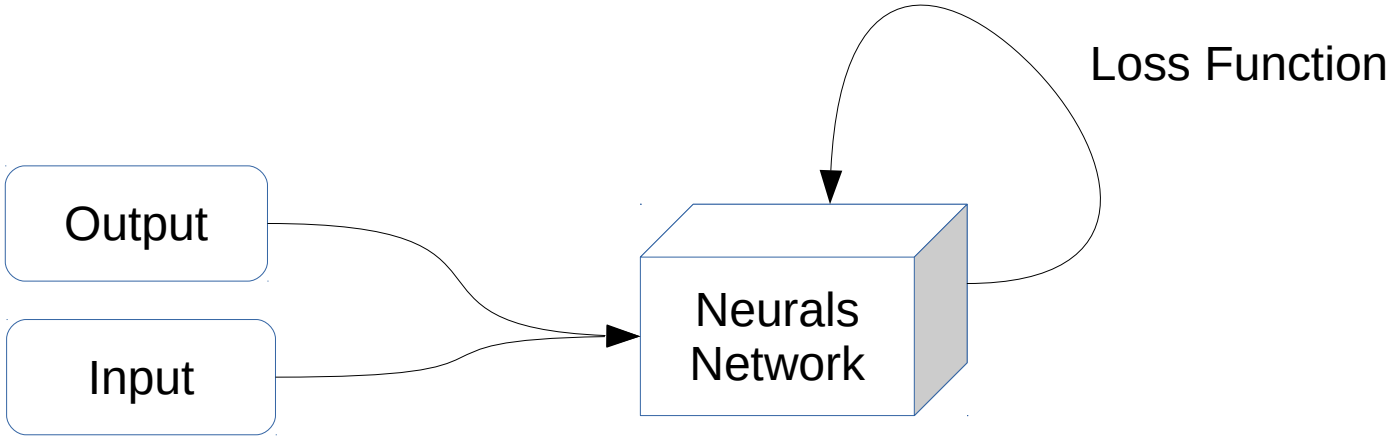
Loss Function
Output
Input
Neurals
Network

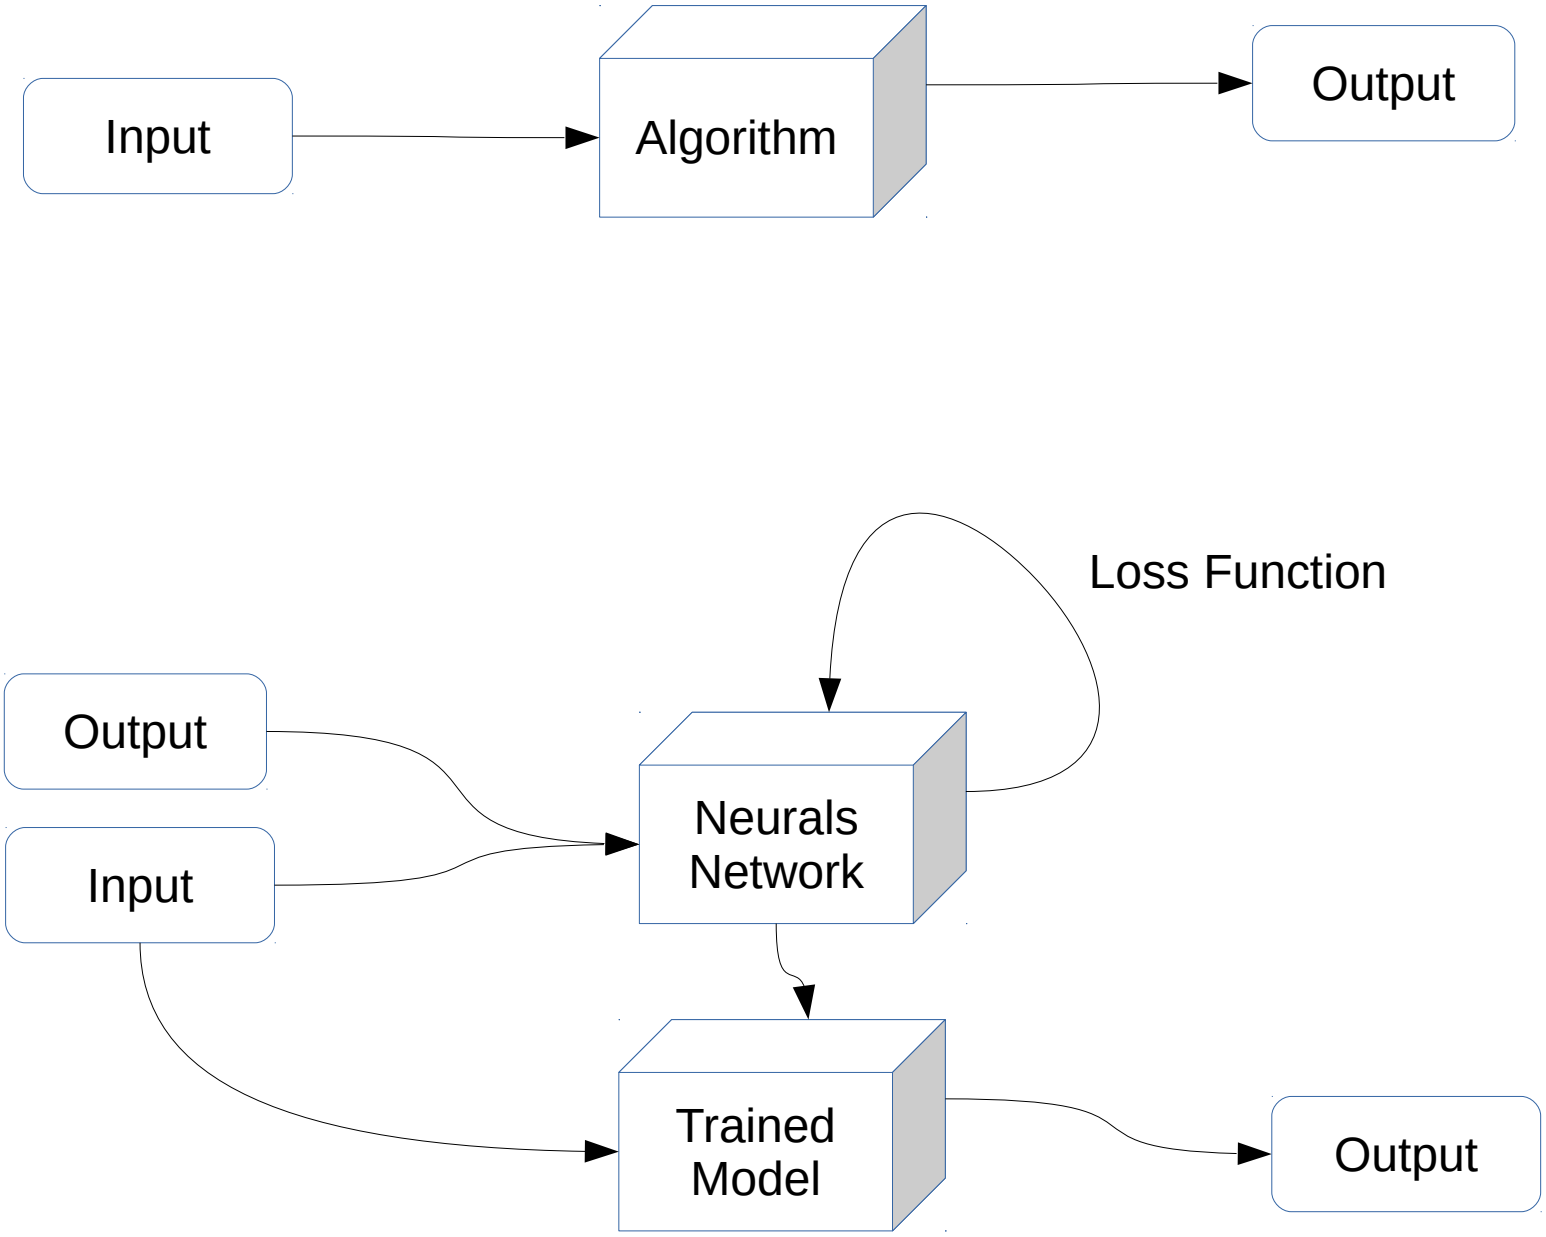
Input
Algorithm
Output
Loss Function
Output
Input
Neurals
Network
Trained
Model
Output

Computer Vision at scale with RoboSat.pink

Goals

Detect inconsistencies between two DataSets

Train on a small area, predict on a larger one.

RoboSat.pink

DataSet Quality Analysis

Change Detection highlighter

Features extraction

@RoboSatPink Computer Vision ecosystem for GeoSpatial Imagery

RoboSat != RoboSat.pink

https://github.com/mapbox/robosat

https://github.com/datapink/robosat.pink

https://github.com/mapbox/robosat/issues/184

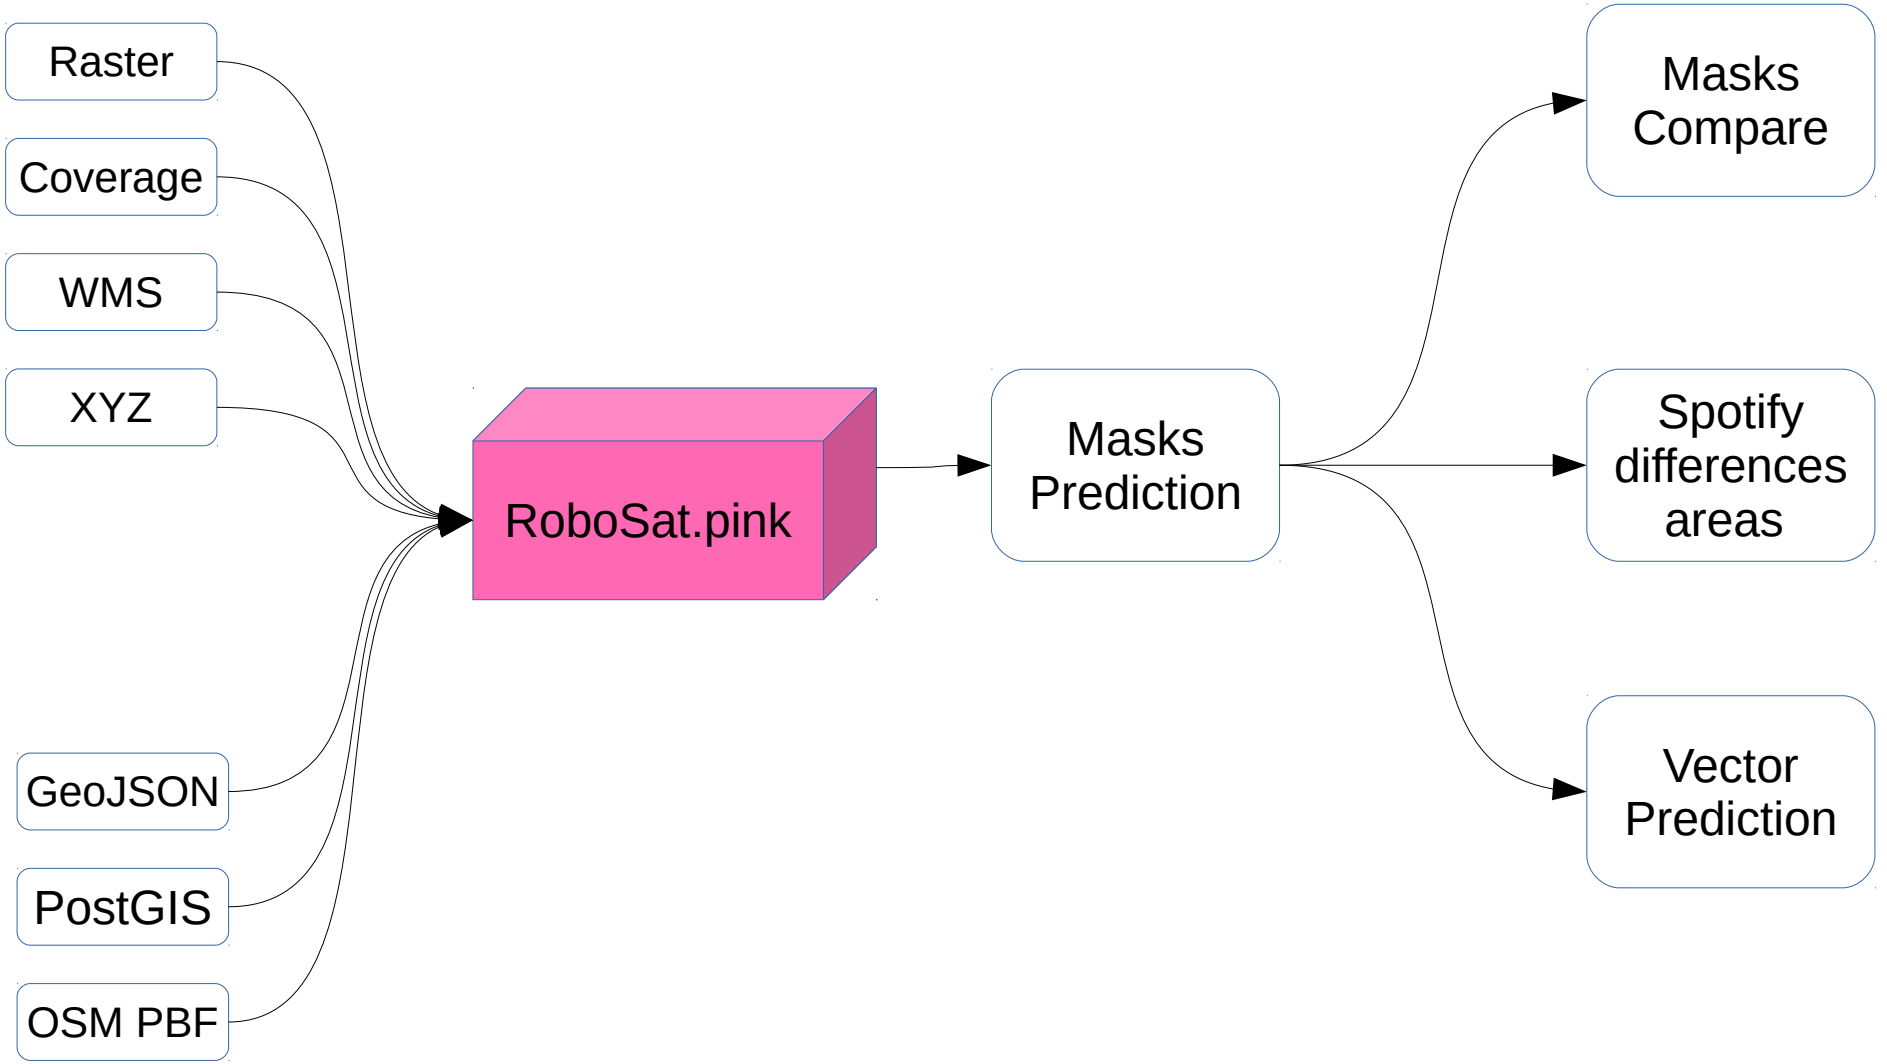
Raster
Coverage
WMS
XYZ
GeoJSON
PostGIS
OSM PBF
RoboSat.pink
Masks Prediction
Masks Compare
Spotify differences areas
Vector Prediction

Compare Predicts against alternate datasets

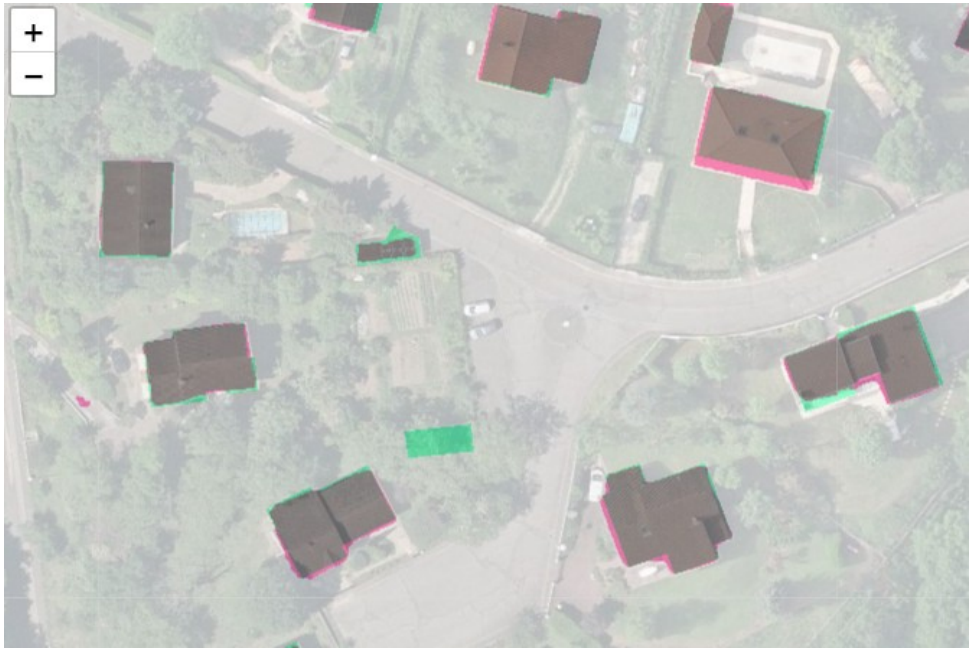

Pink : Predicted by trained model

Green : Alternate dataset

Grey : Both agree

Pink squares : Significant differences

Command Line Interface

Tools:

- `rsp cover` Generate a tiles covering, in csv format: X,Y,Z
- `rsp download` Downloads tiles from a remote server (XYZ, WMS, or TMS)
- `rsp extract` Extracts GeoJSON features from OpenStreetMap .pbf
- `rsp rasterize` Rasterize vector features (GeoJSON or PostGIS), to raster tiles
- `rsp subset` Filter images in a slippy map dir using a csv tiles cover
- `rsp tile` Tile raster coverage
- `rsp train` Trains a model on a dataset
- `rsp export` Export a model to ONNX or Torch JIT
- `rsp predict` Predict masks, from given inputs and an already trained model
- `rsp compare` Compute composite images and/or metrics to compare several XYZ dirs
- `rsp vectorize` Extract simplified GeoJSON features from segmentation masks
- `rsp info` Print RoboSat.pink version informations

RoboSat.pink 101

https://github.com/datapink/robosat.pink/blob/master/docs/101.md

Easy to deploy

```
pip3 install RoboSat.pink
```

So all you need is :

- Imagery
- GPU
- Labels

So all you need is :

- Imagery → *any file format readable by GDAL*
- GPU → *NVIDIA > 8Go RAM*
- Labels → *that’s often the key point*

From OpenData to OpenDataSet

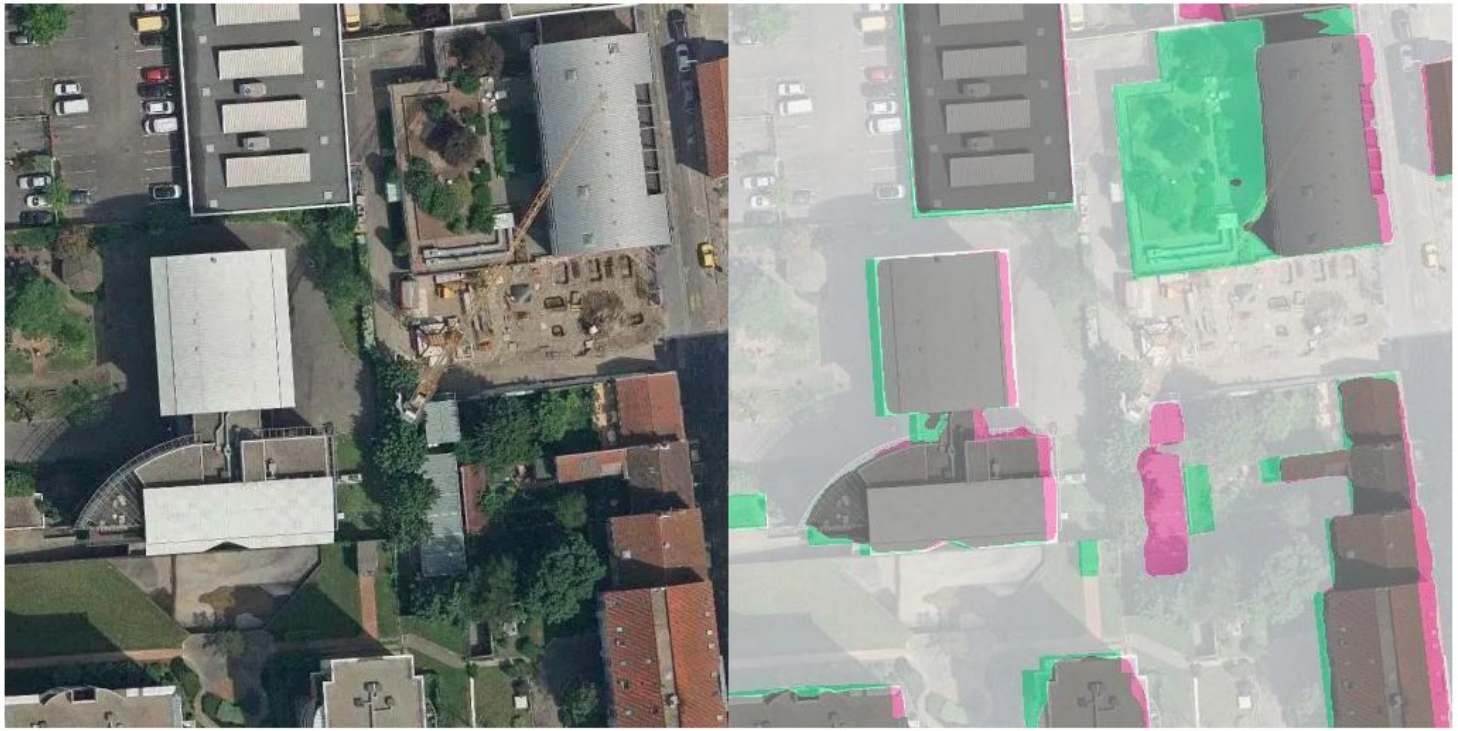

https://github.com/datapink/robosat.pink/blob/master/docs/from_opendata_to_opendataset.md

Unet Like with ResNet50 Encoder

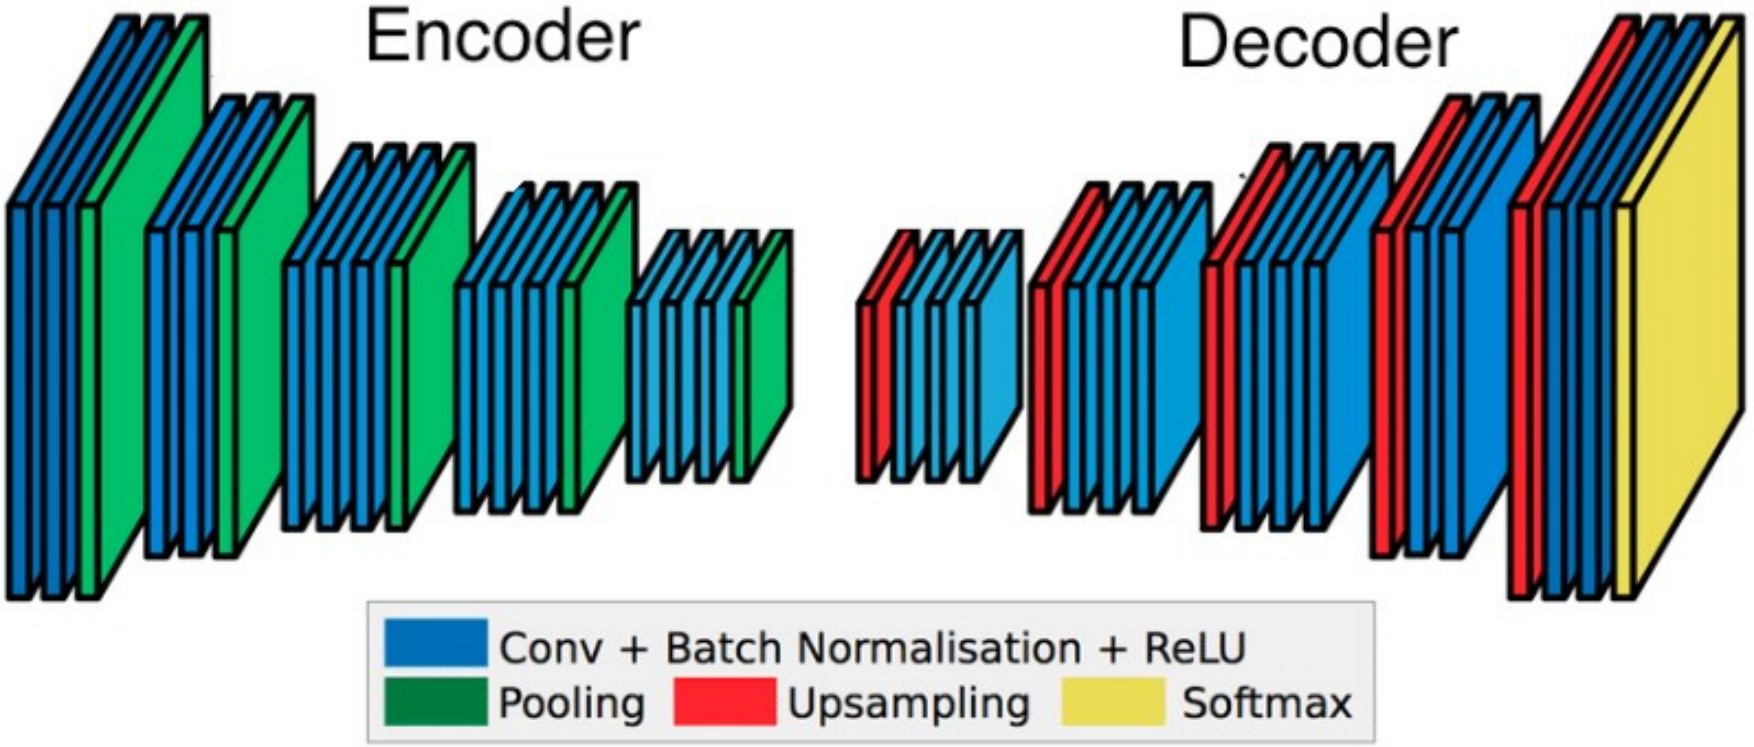

Encoder–Decoder architecture

Image Credit: https://divamgupta.com

Surface based semantic Loss

Image Label Cross Entropy mIoU Lovasz

http://www.cs.toronto.edu/~wenjie/papers/iccv17/mattyus_etal_iccv17.pdf
http://www.cs.umanitoba.ca/~ywang/papers/isvc16.pdf
https://arxiv.org/abs/1705.08790

Data Augmentations

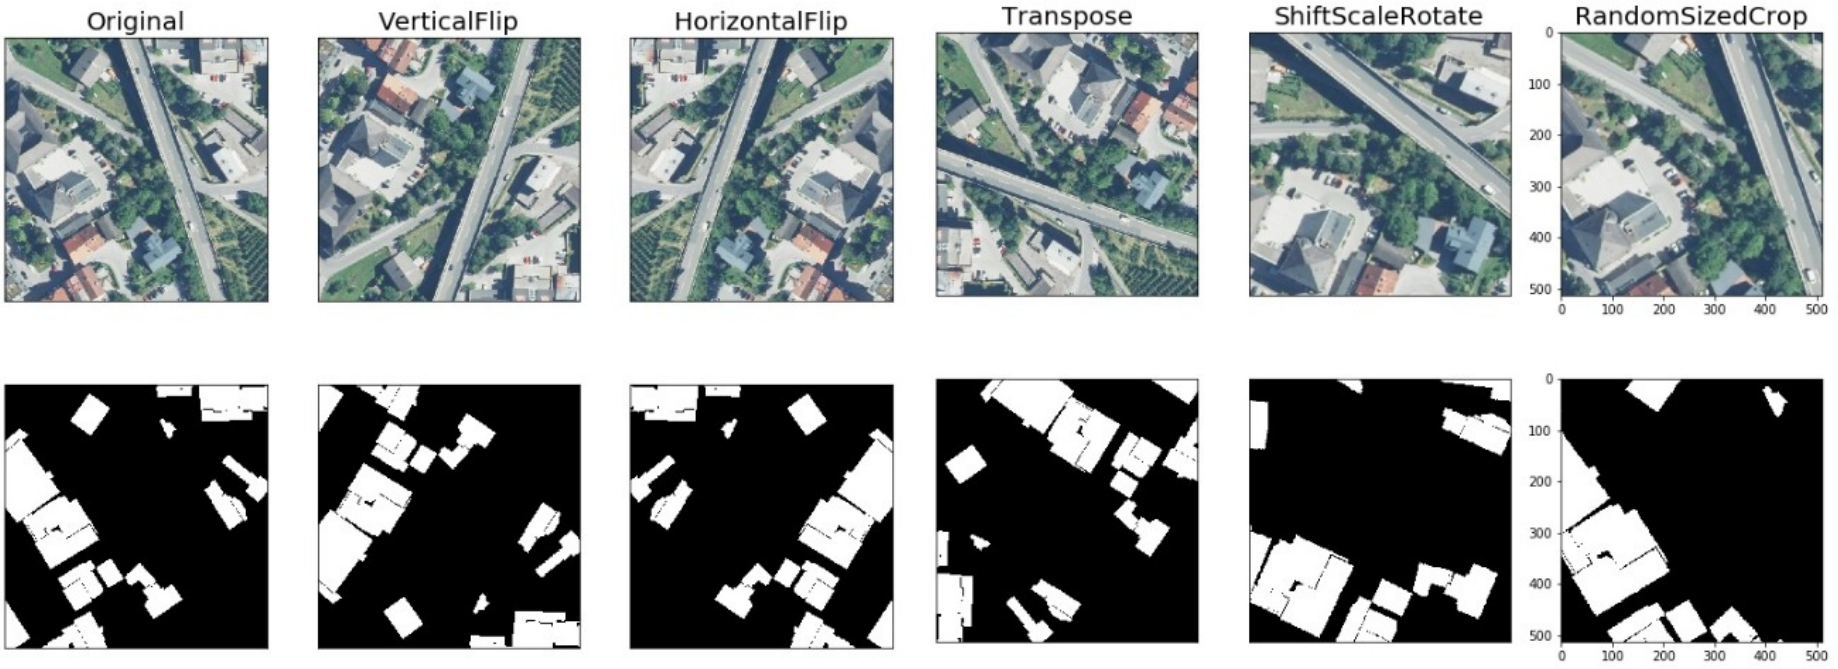

https://arxiv.org/abs/1809.06839
https://github.com/albu/albumentations

More than an application, an easy to extent framework

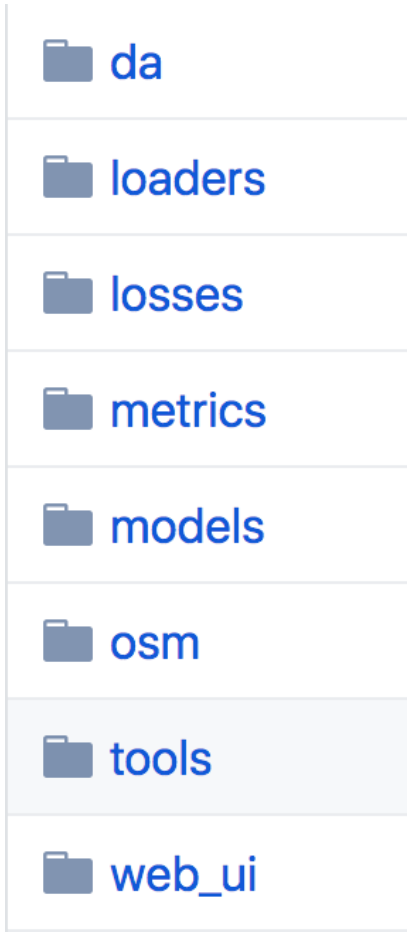

https://github.com/datapink/robosat.pink/blob/master/docs/extensibility_by_design.md

Stacks

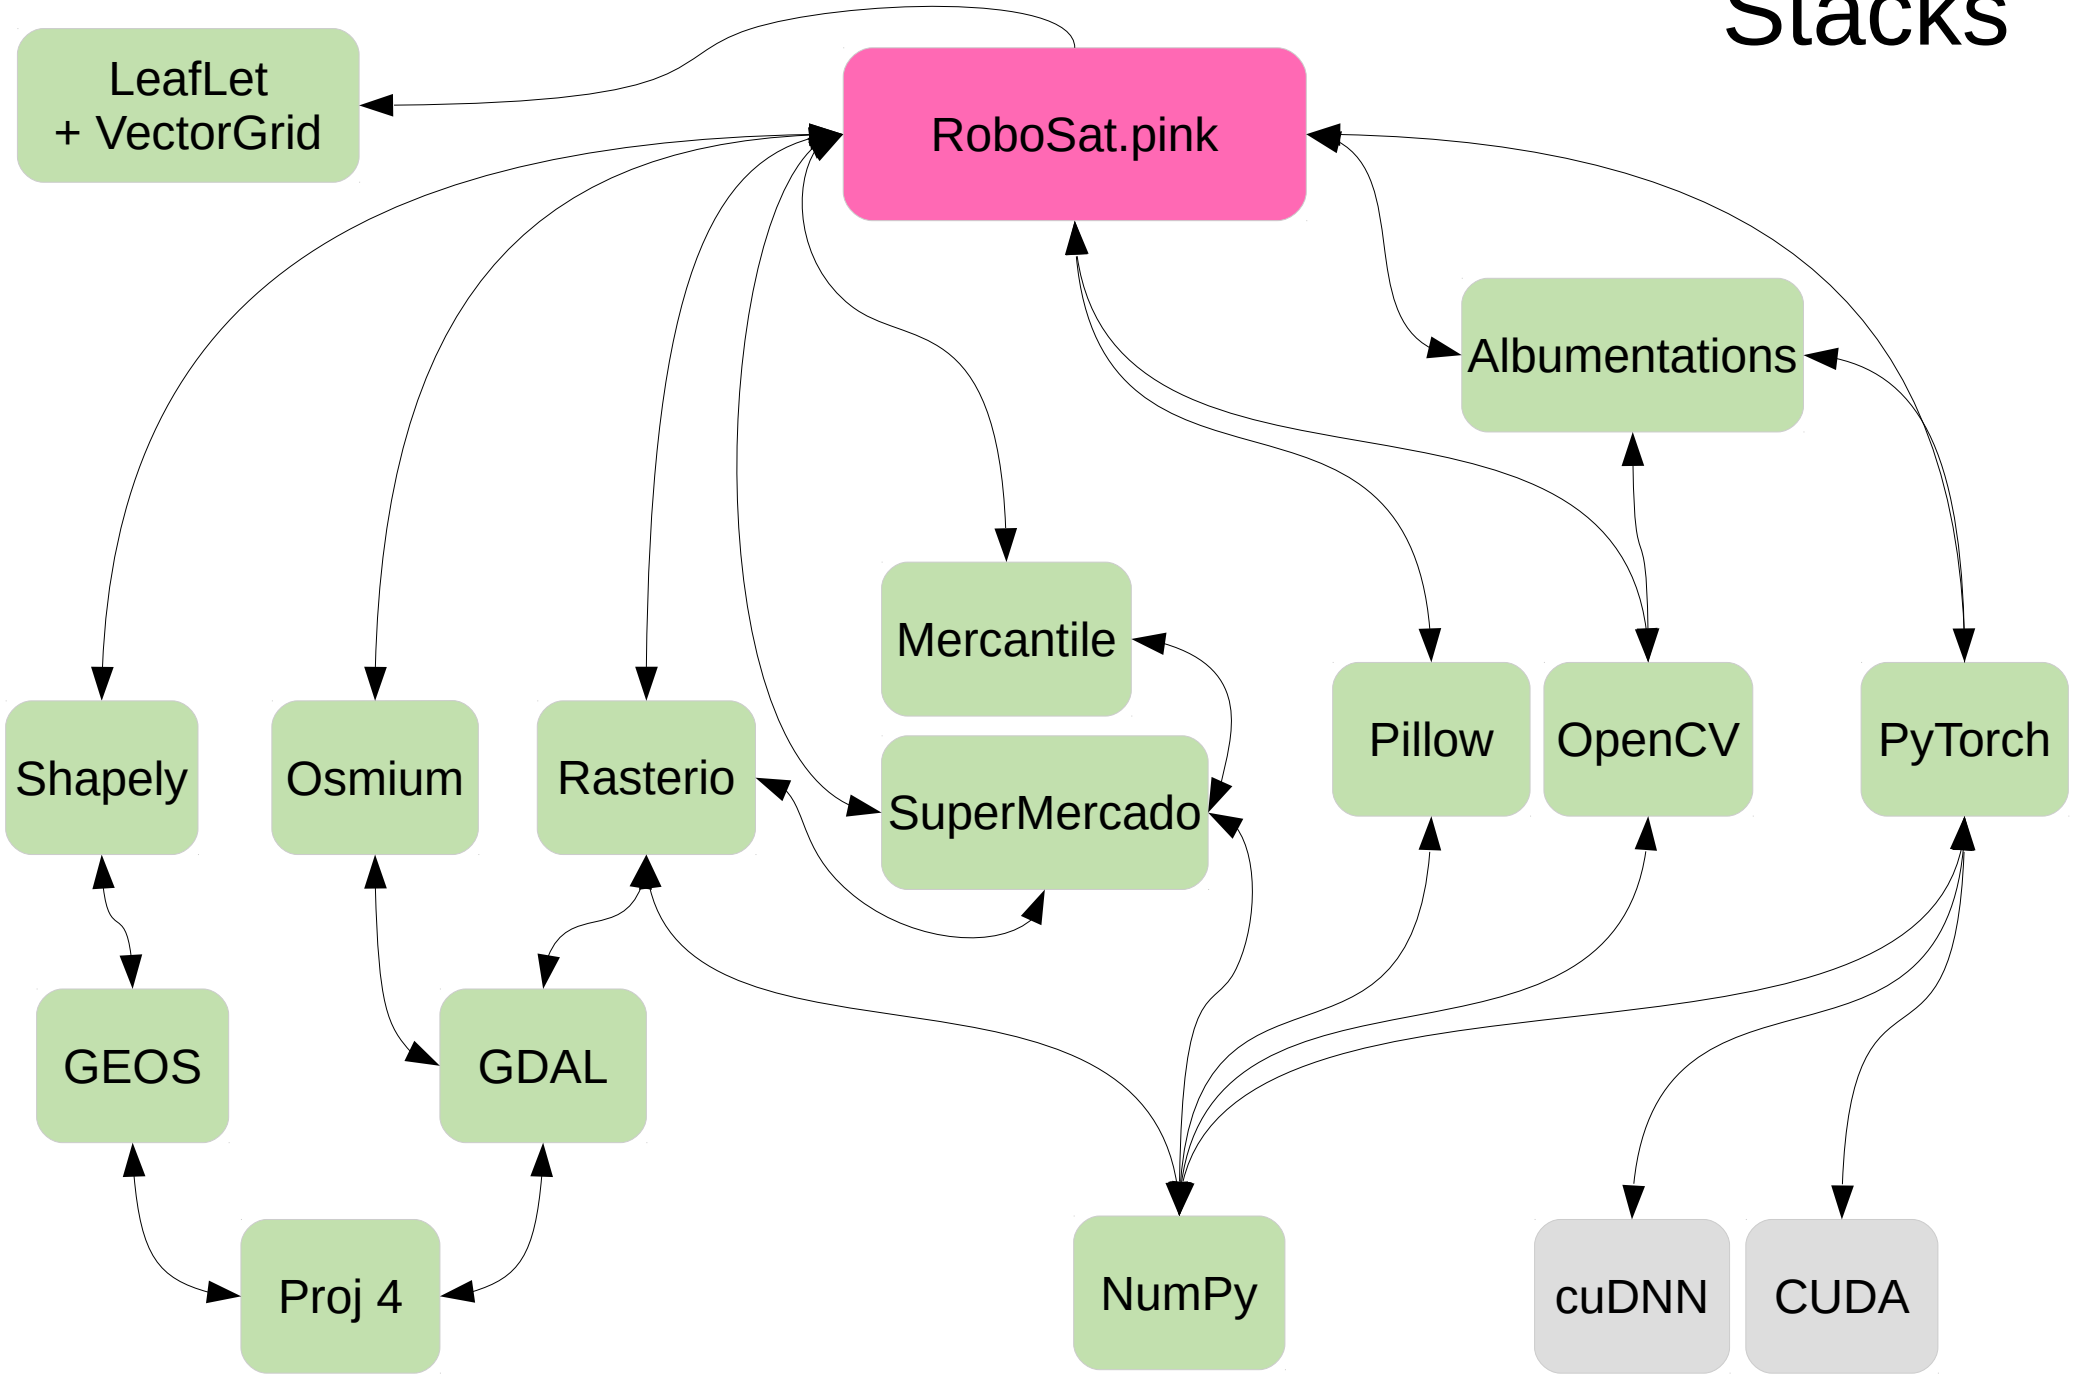
LeafLet
+ VectorGrid
RoboSat.pink
Albumentations
Mercantile
Shapely
Osmium
Rasterio
SuperMercado
Pillow
OpenCV
PyTorch
GEOS
GDAL
Proj 4
NumPy
cuDNN
CUDA

Open Source

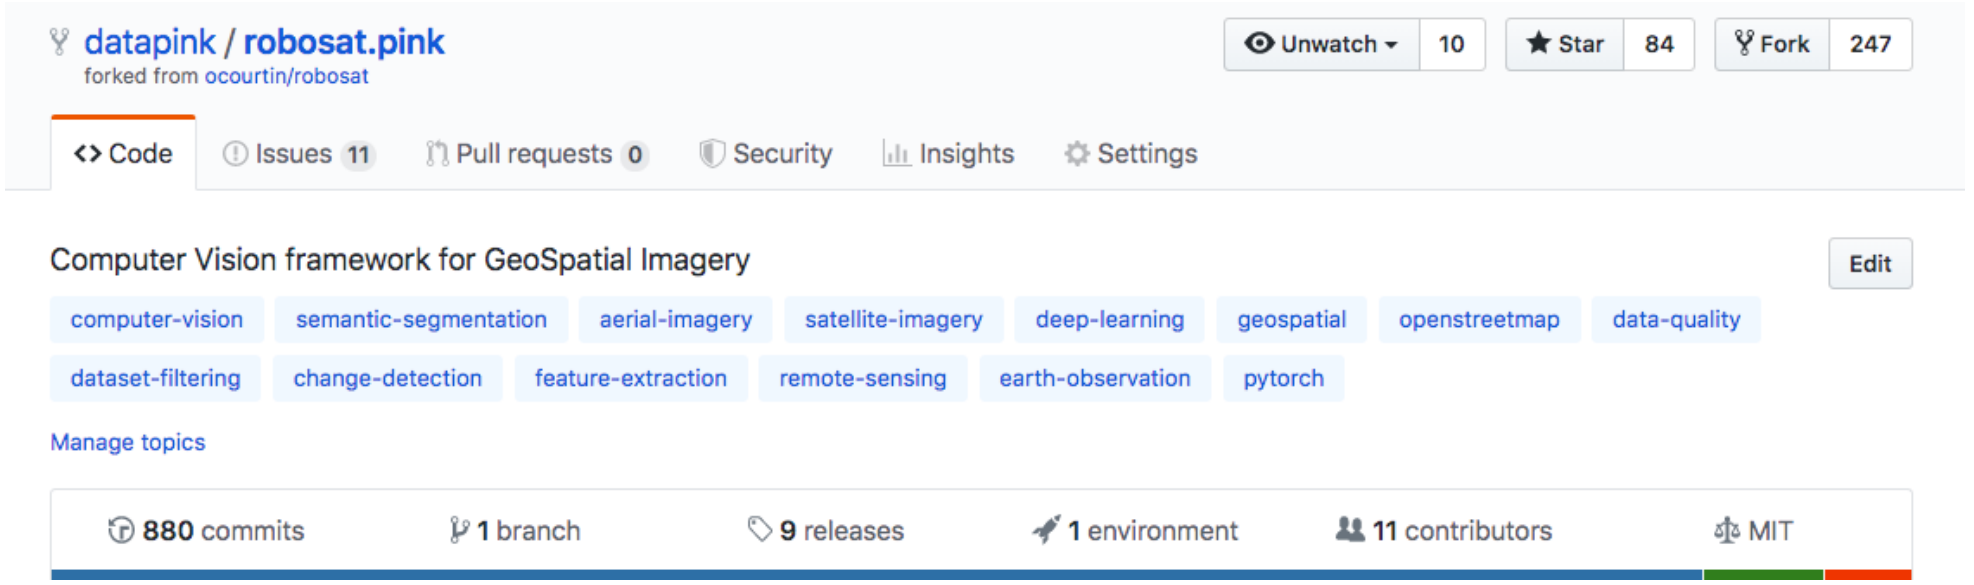
datapink / robosat.pink
forked from ocourtin/robosat
Unwatch 10
Star 84
Fork 247
Code
Issues 11
Pull requests 0
Security
Insights
Settings
Computer Vision framework for GeoSpatial Imagery
Edit
computer-vision
semantic-segmentation
aerial-imagery
satellite-imagery
deep-learning
geospatial
openstreetmap
data-quality
dataset-filtering
change-detection
feature-extraction
remote-sensing
earth-observation
pytorch
Manage topics
880 commits
1 branch
9 releases
1 environment
11 contributors
MIT

Request For Funding

- Increase (again) prediction accuracy :
    - on low resolution imagery
    - even with few labels
    - feature extraction when they are (really) close
    - with multibands and Data Fusion

- Add support for :
    - MultiClass
    - PointCloud data support
    - Time Series Imagery

- Improve (again) performances

Few performances Metrics

rsp train	~5 Mp/s, -per epoch-
rsp download	~1 to 5 Mp/s
rsp tile	~5 Mp/s
rsp predict	~10 Mp/s
rsp compare	~50 Mp/s
rsp rasterize	~50 Mp/s
rsp vectorize	~50 Mp/s

8 cores CPU, single GPU (either RTX or V100), SSD

How to scale it, or improve it again ?

rsp train	add more GPU, reduce dataset redundancy, improve model, loss or optimizer
rsp tile	add more CPU use raster compression
rsp predict	export model to ONNX or JIT, then use an high performance inference solution.

Why performances matters ?

- Playful and Human Learning
- Time and money saver
- No Planet B

Computer Vision Take Away

- Industrial OSS state of art GeoSpatial Imagery framework available

- Plain OpenData can be use to train accurate model

- Performances already OK to use it for region / small country area, even on cheap GPU server

- Scale if you provide better hardware…

From text to map, a state of art !

["The Port of Paulsboro is located on the Delaware River and Mantua Creek in and around Paulsboro, in Gloucester County, New Jersey, US, approximately 78 miles (126 km) from the Atlantic Ocean. Traditionally one of the nation's busiest for marine transfer operations, notably for crude oil and petroleum products, such as jet fuel and asphalt, it is a port of entry with several facilities within a foreign trade zone.\nA part of the port is being redeveloped as an adaptable deep water omniport able to handle a variety of bulk and break bulk cargo, as well as shipping containers. It is targeted to become a manufacturing/assembly center for wind turbines for the development of wind power in New Jersey and other offshore wind power projects along the East Coast of the United States. The Paulsboro Marine Terminal, as it is known, is owned by the South Jersey Port Corporation and operated by Holt Logistics. The first ship is expected to arrive at the new facility in early 2017 carrying steel for NLMK. The first ship to call at the port, the Doric Warior, carrying steel for NLMK, arrived March 3, 2017, marking the opening of the new facility."]

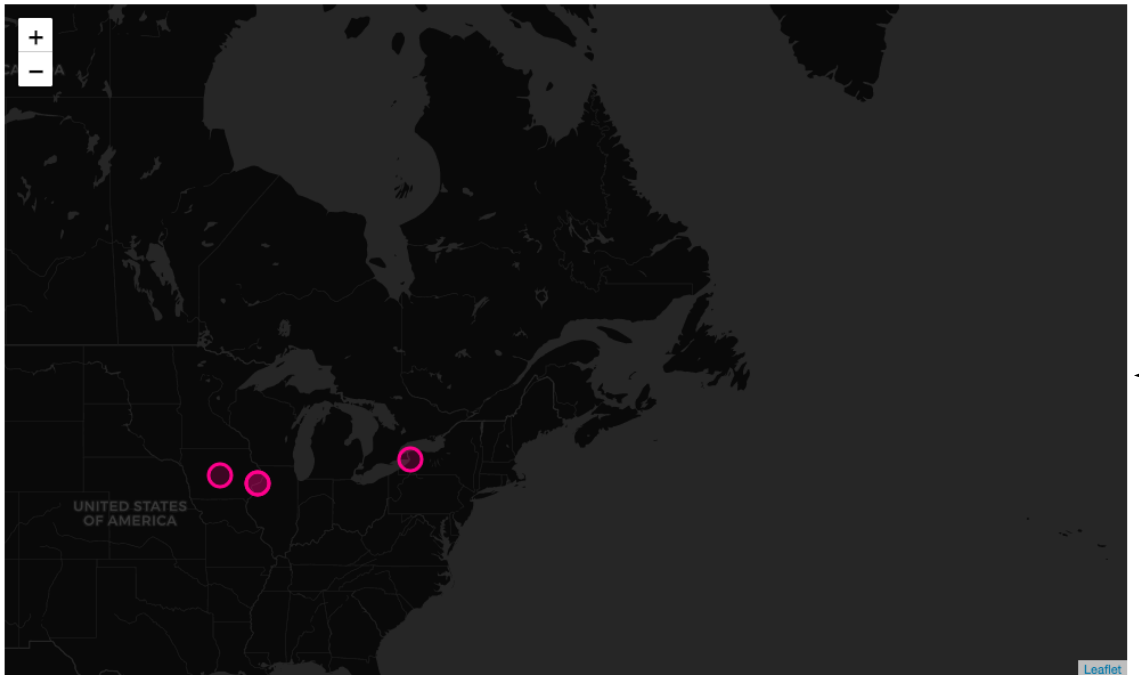

Text 2 Map

Which part of the text refer to a spatial stuff ?

Which part of the text refer to a spatial stuff ?

Then, is this toponym, findable in a Gazetter ?

Which part of the text refer to a spatial stuff ?

Named Entity Recognition

Then, is this toponym, findable in a Gazetter ?

Which part of the text refer to a spatial stuff ?

Named Entity Recognition

Then, is this toponym, findable in a Gazetter ?

GeoNames query

Which part of the text refer to a spatial stuff ?

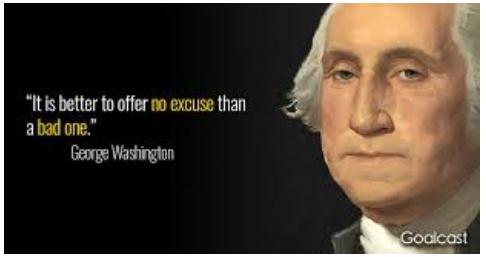

Named Entity Recognition

Then, is this toponym, findable in a Gazetter ?

GeoNames query

Which part of the text refer to a spatial stuff ?

Named Entity Recognition

Then, is this toponym, findable in a Gazetter ?

GeoNames query

Why NLP hard ?

Let’s eat grandma !

Let’s eat, grandma !

Irony

« The meaning of a word is its use in the language. »

Ludwig Wittgenstein

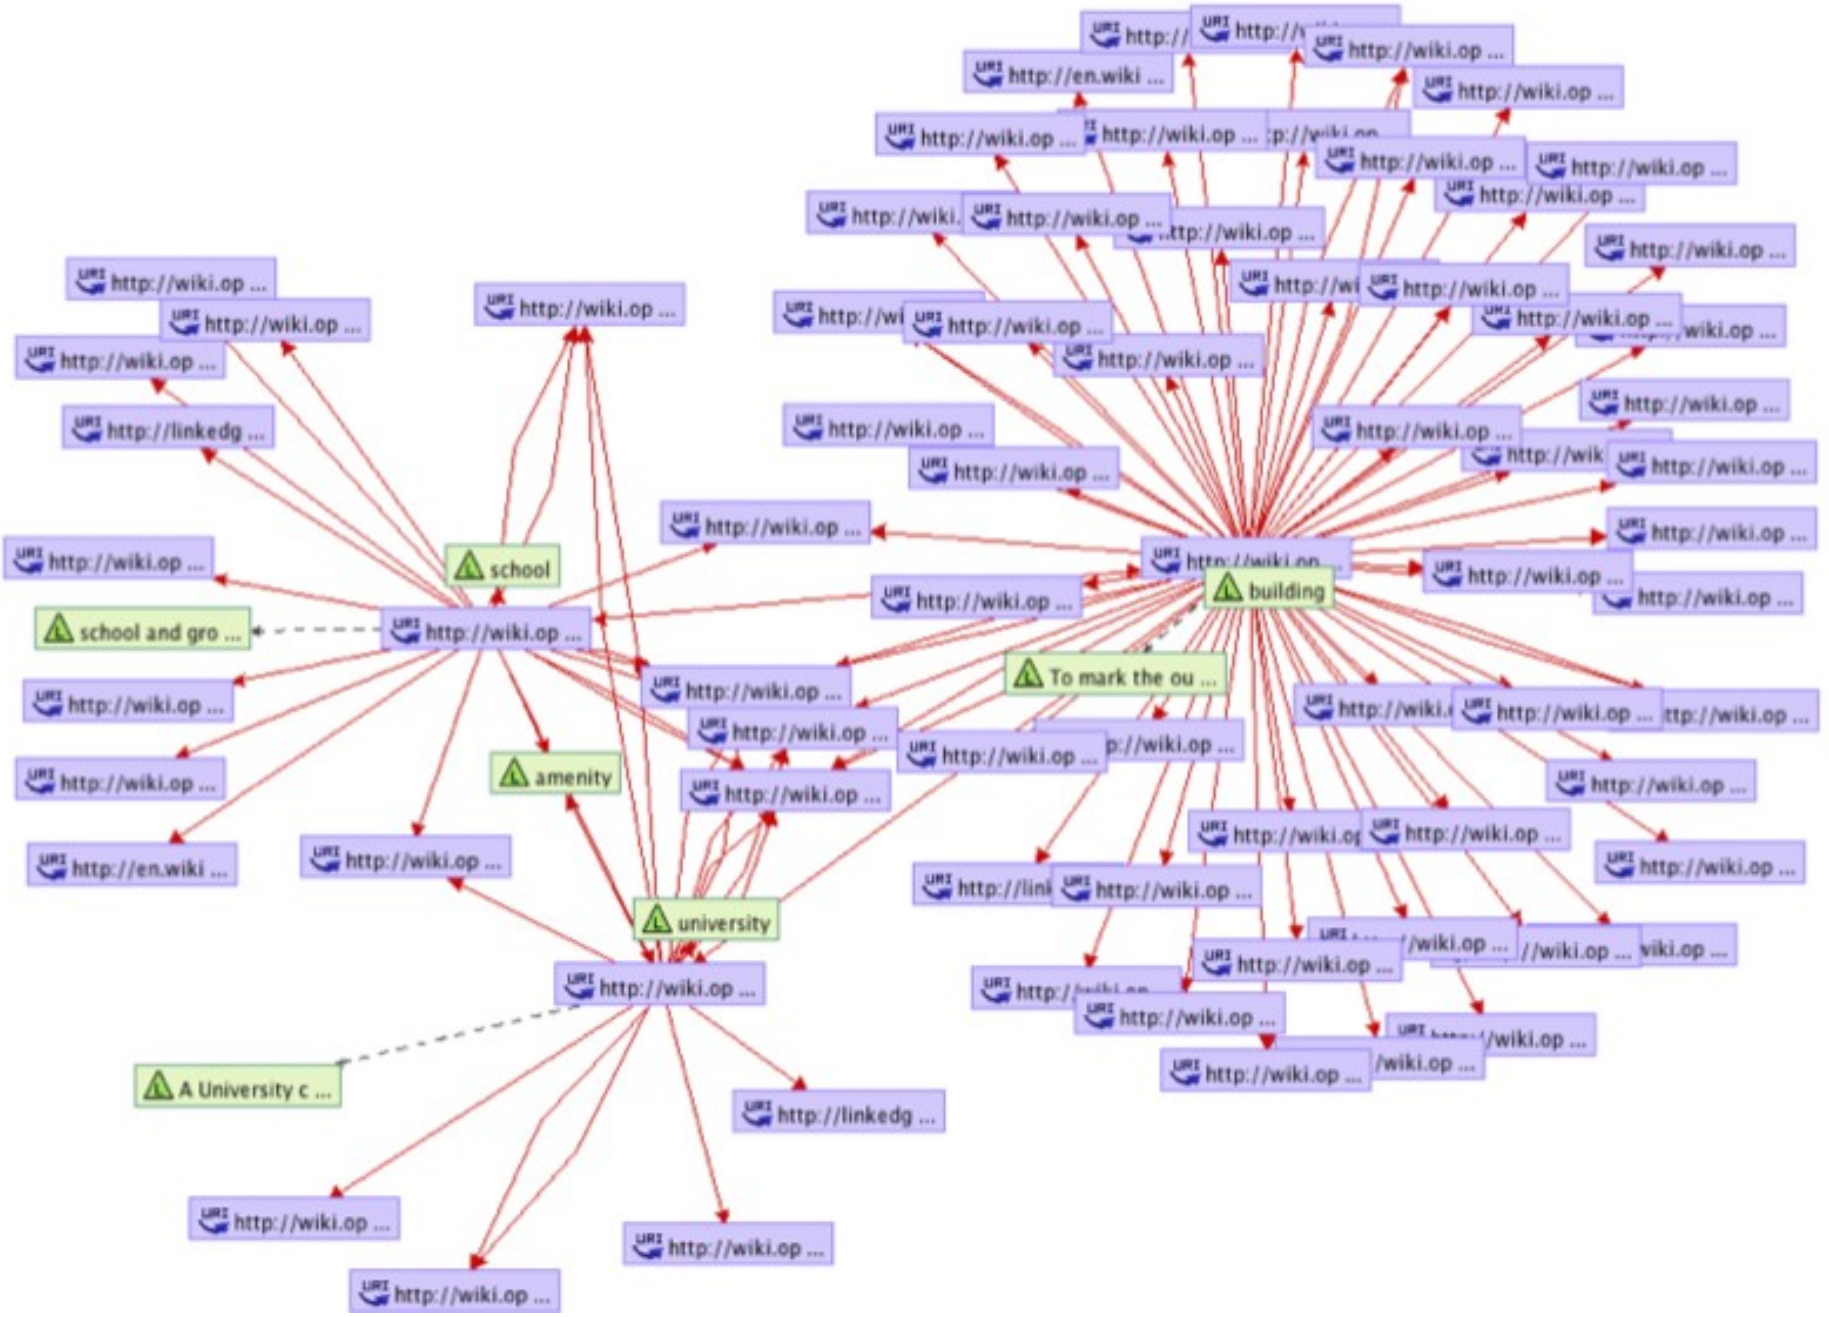

Source: https://wiki.openstreetmap.org/wiki/OSM_Semantic_Network

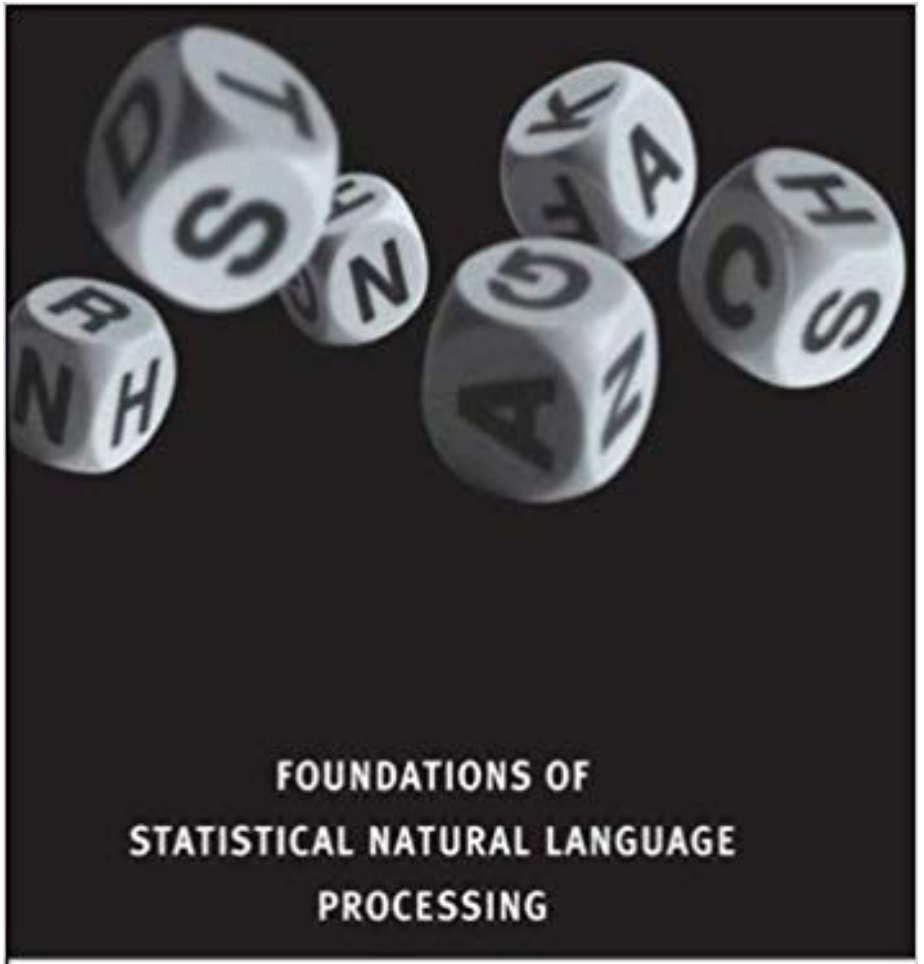

FOUNDATIONS OF STATISTICAL NATURAL LANGUAGE PROCESSING

CHRISTOPHER D. MANNING AND
HINRICH SCHÜTZE

NLP Open Source libs

	SPACY	NLTK	CORENLP
Programming language	Python	Python	Java / Python
Neural network models	✓	✗	✓
Integrated word vectors	✓	✗	✗
Multi-language support	✓	✓	✓
Tokenization	✓	✓	✓
Part-of-speech tagging	✓	✓	✓
Sentence segmentation	✓	✓	✓
Dependency parsing	✓	✗	✓
Entity recognition	✓	✓	✓
Entity linking	✗	✗	✗
Coreference resolution	✗	✗	✓

Source: https://spacy.io/usage/facts-figures

	SPACY	NLTK	CORENLP
Programming language	Python	Python	Java / Python
Neural network models	✓	✗	✓
Integrated word vectors	✓	✗	✗
Multi-language support	✓	✓	✓
Tokenization	✓	✓	✓
Part-of-speech tagging	✓	✓	✓
Sentence segmentation	✓	✓	✓
Dependency parsing	✓	✗	✓
Entity recognition	✓	✓	✓
Entity linking	✗	✗	✗
Coreference resolution	✗	✗	✓

Source: https://spacy.io/usage/facts-figures

SacreMoses https://github.com/alvations/sacremoses

PyTorch Text https://github.com/pytorch/text

DeepPavlov https://github.com/deepmipt/DeepPavlov

```
# NER classification with an EN Large model

ner_en_lg = nlp_en_lg(abstract)
sp.displacy.render(ner_en_lg, jupyter=True, style='ent')
```

['The Davenport Community School District ORG Is a public school district in Scott County GPE , Iowa GPE . The school district covers 109 square miles QUANTITY (280 km2 QUANTITY) that includes the city of Davenport GPE , where it is based, and the western Scott County GPE communities of Blue Grass ORG , Buffalo GPE and Walcott PERSON . Founded in 1858 DATE it established one CARDINAL of the first ORDINAL publicly funded high schools in the United States GPE , the third ORDINAL teachers' ORG training school and hired the first ORDINAL female superintendent in the country. It serves nearly 16,000 CARDINAL students in 31 CARDINAL school buildings.']


MORDECAI
mordecai
full text geoparsing
github.com/openeventdata/mordecai

Perspectives

NER :

- Use more rich structured text, as dbpedia to train models
- Anything helpful for multilang handling
- Use latest papers: https://nlpprogress.com/english/named_entity_recognition.html

Gazetteer

- OCR existing maps for completion ?
- Other kind of data alowing massive completion ?

Toponym matching

- Batch to use text context
- Use others dimensions than only population,
  and obviously geographical one...

NLP TakeAway

NLP is harder than any others patterns recognition domains

NER softwares and models already availables

Academic NLP R&D is an hot and vivid topic

Firsts Text2Map tools arises

But still lot to do !

Extract insights from GeoSpatial data with Deep Learning

Development, Expertise, Support and Training, on :

- Computer Vision

- Natural Language Processing

- Time Series Analysis

@data_pink

www.datapink.com