Crowd driven AI: Bilderkennung zur Rettung des Regenwaldes

Gerhard Svolba

Data Scientist, SAS Austria

My Data Science Webinar Series:
Youtube|Medium|Linkedin
https://twitter.com/gsvolba
https://github.com/gerhard1050
https://www.linkedin.com/in/gerhardsvolba/

Links

- https://www.sas.com/de_at/data-for-good/rainforest.html
- https://developer.sas.com/home.html
- https://developer.sas.com/guides/dlpy.html
- https://github.com/sassoftware/python-dlpy
- **IIASA, SAS. (2022) Crowd-driven deep learning tracks Amazon deforestation. *In prep,***
  - ***Ian McCallum, International Institute for Applied Systems Analysis, Laxenburg, Austria,***
  - ***Jon Walker, SAS Campus Drive, Cary, NC 27513, United States***

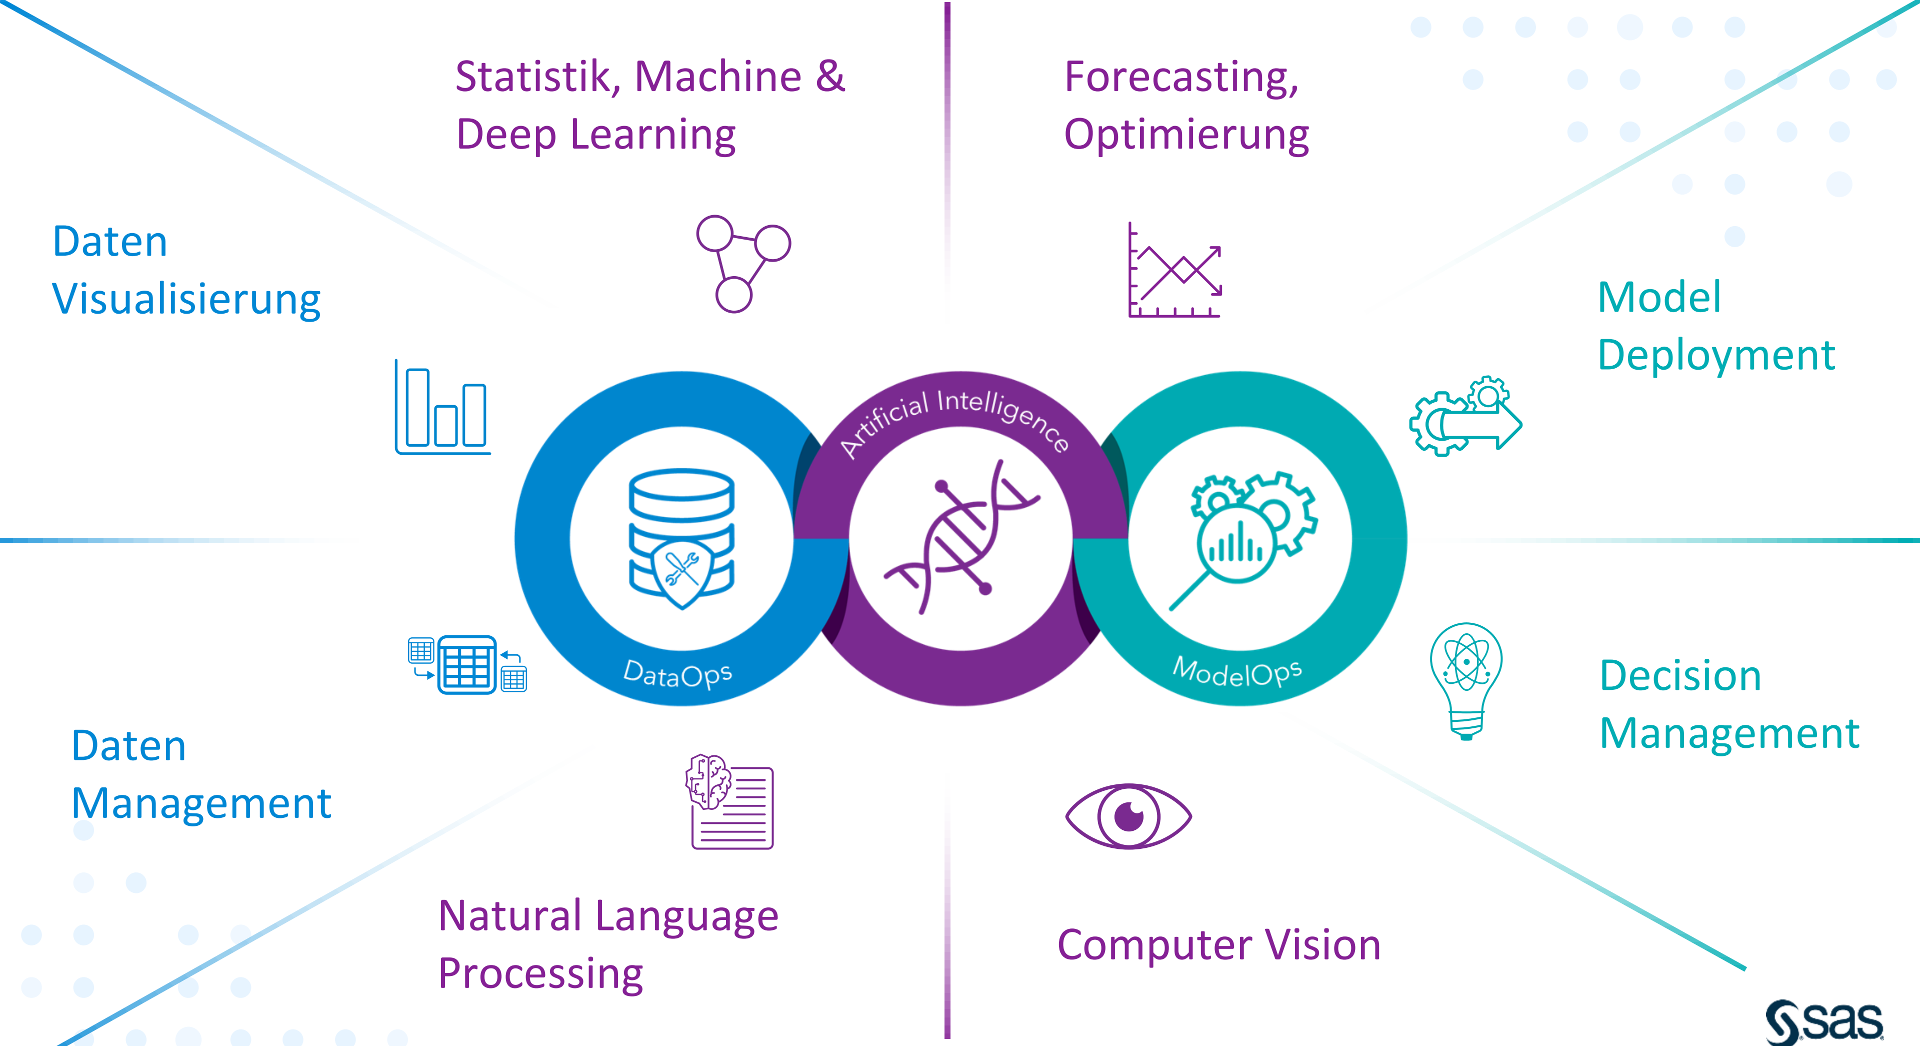
Statistik, Machine & Deep Learning
Forecasting, Optimierung
Daten Visualisierung
Model Deployment
Artificial Intelligence
DataOps
ModelOps
Decision Management
Daten Management
Natural Language Processing
Computer Vision

Künstliche Intelligenz

ist die Disziplin, Systeme zu trainieren, um Abläufe der menschlichen Tätigkeit durch LERNEN und AUTOMATISIERUNG zu emulieren.

Lernen aus Erfahrungen/ Daten

Anpassen an neue Fakten

Automatisierung des Prozesses

Analyse der Abholzung des Amazonas Regenwalds auf Basis der automatischen Klassifikation von Satellitenbildern

(Kooperation mit der IIASA, International Institute for Applied System Analysis)

https://www.sas.com/de_at/data-for-good/rainforest.html

Background and Project Goal

- Amazon Rainforest 5,5 Mio km^2 (larger that EU)
- Deforestation:
  - 29000km^2/year in 2003
  - 6000km^2/year in 2014,
  - now raising again
- Project Goal:
  Train a computer vision model to automatically detect areas where deforestation is taking place

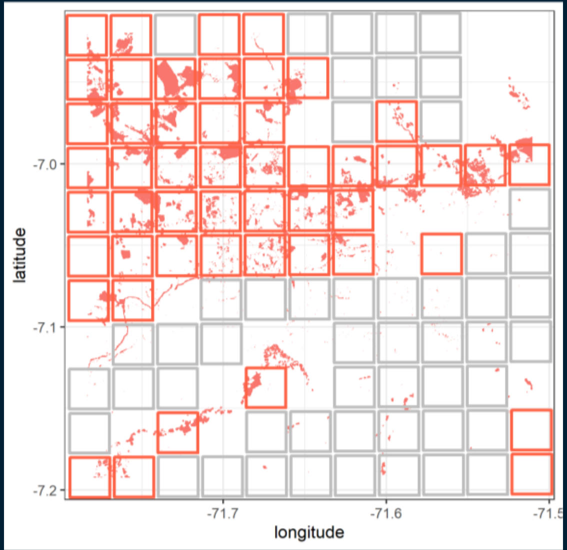

Classification Map

Based on previous input, this map shows smaller areas of the Amazon rainforest where we are focusing efforts to identify more recent deforestation. By using image data from this ecologically diverse territory, we are giving our model a wide variety of examples so that it can one day learn to detect human impact anywhere in the Amazon. Here's our current progress to date for phase two.

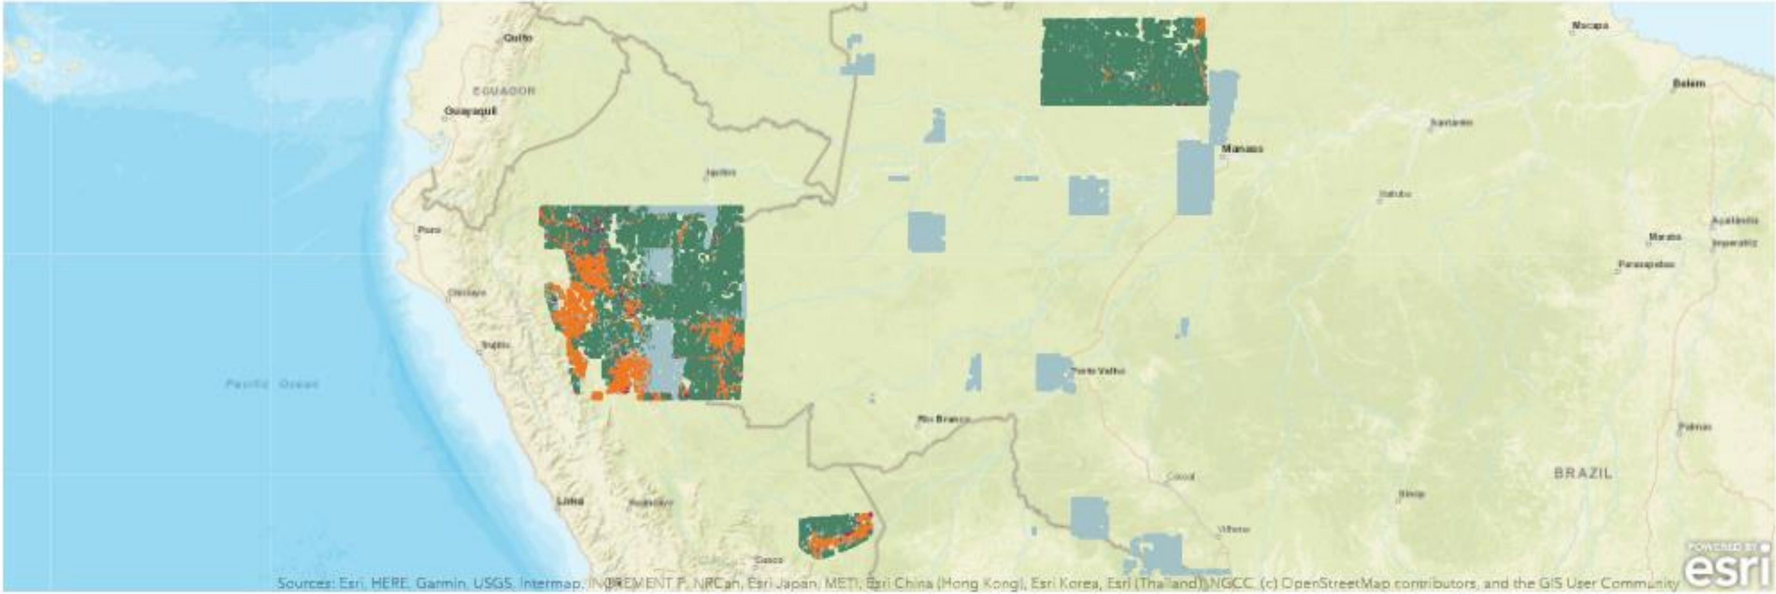

Human Impact | No Human Impact | Recent Human Impact | Not Yet Classified

Überblick über den Modellierungsablauf

Satellite Images

Perform Expert Rating (Subset)

Train a first Model

Model 1 “Expert”

Validate Model Accuracy

Provide Data for Crowd Sourcing

Images

Crowd Sourcing (GatherIQ)

Crowd Sourcing

- Only few rainforest deforestation pictures exist in ImageNet, CIFAR-10, CIFAR-100. And: not in 10m resolution as from Sentinel-2 satellite
- Studies show: crowd performs equally well as experts

Menschlicher Engriff oder natürliche Entwaldung?

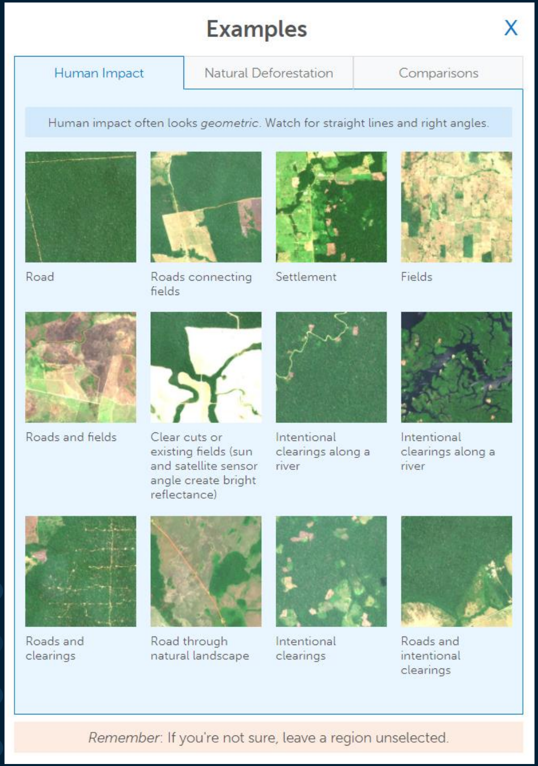

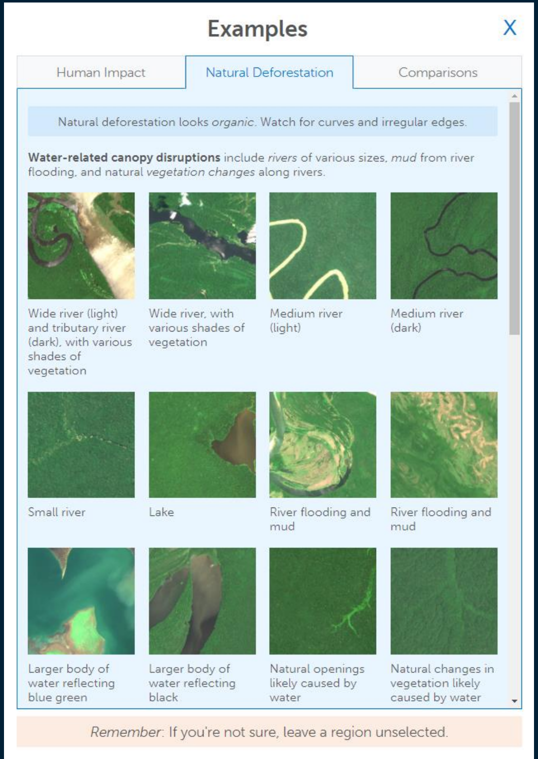

Beispiele für die Anwendung der Klassifikations-App

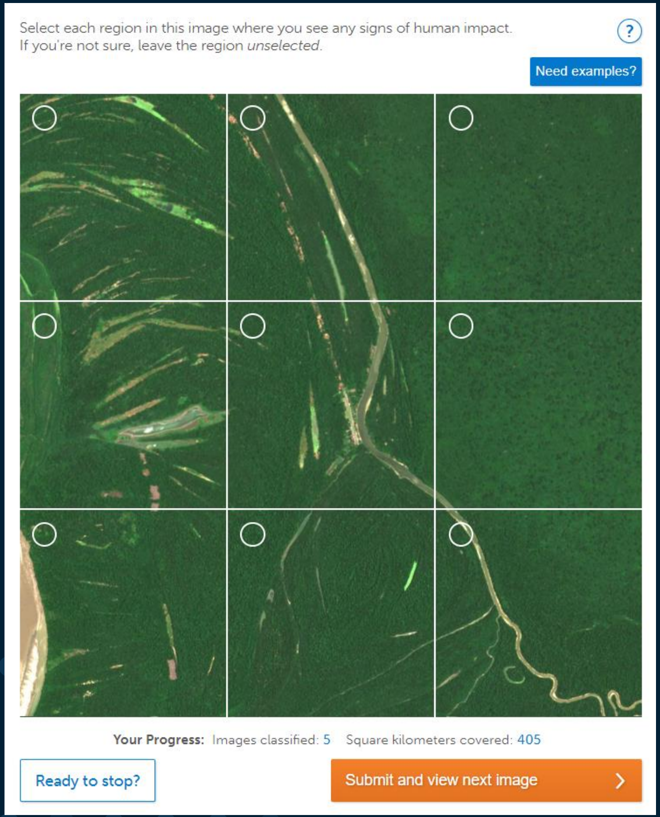

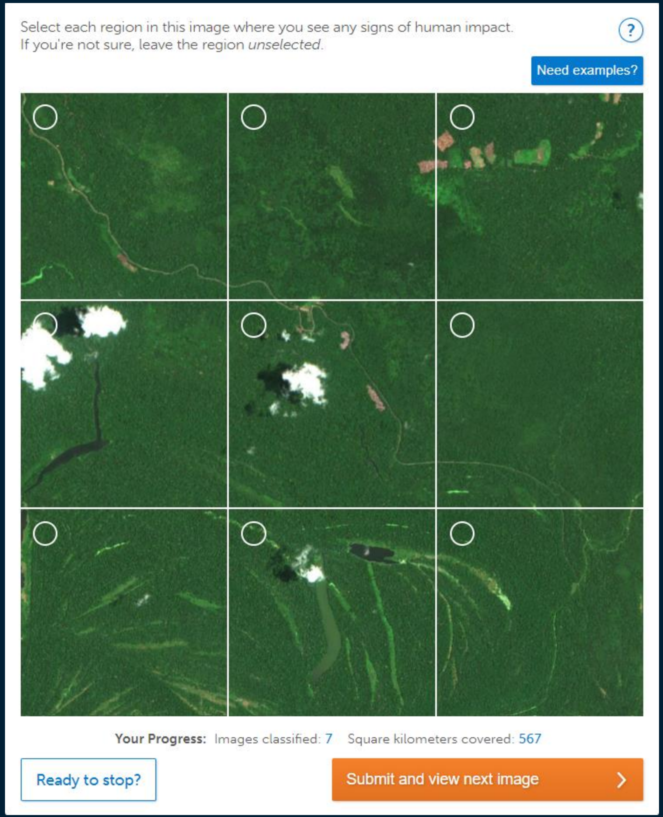

Crowd Sourcing

- Only few rainforest deforestation pictures exist in Imag
  CIFAR-100. And: not in 10m resolution as from Sentinel-2 satellite
- Studies show: perform equally well as experts
- Accuracy increasing when consensus or majority voting is used
- Deforestation project: 6 months, 5500 active participants from 96 countries, 389.988 km2
- Used cloud free samples from sentinel-2 satellite images
- Consensus among the crowd: >80% for majority of pictures
- Agreement with expert review on a sample of 200 pictures: 88%

Crowd Scouring Process

- Each Image shown to multiple users
- 300 pictures removed as no consensus could be obtained

$$\hat{p}_{UB} = \hat{p} + 1.96\sqrt{\frac{\hat{p}(1 - \hat{p})}{n_{total}}}$$

$$\hat{p}_{LB} = \hat{p} - 1.96\sqrt{\frac{\hat{p}(1 - \hat{p})}{n_{total}}}$$

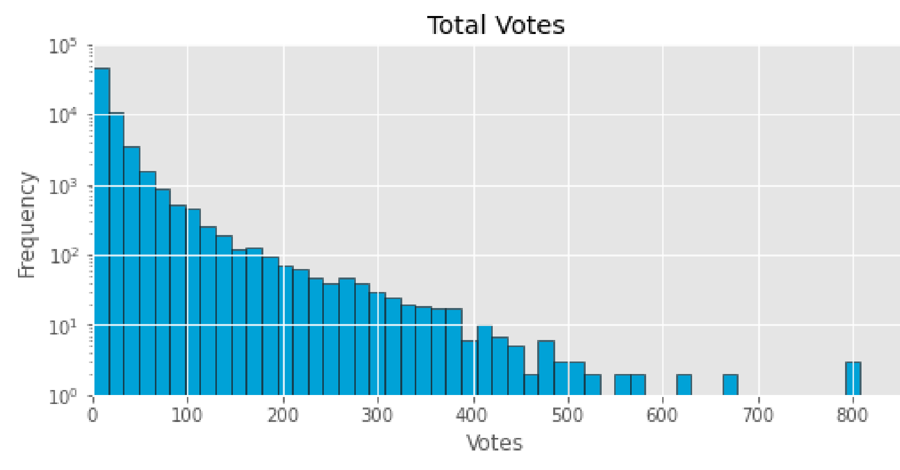

Results from Crowd Sourcing

Figure 1. Results of the crowdsourcing campaign over the Brazilian Amazon between June and November 2020. Map of the 390,000 km^2 (43,100 images) classified by the crowd as having either evidence of deforestation or no deforestation. Individual pixels represent a 3 x 3 km image.

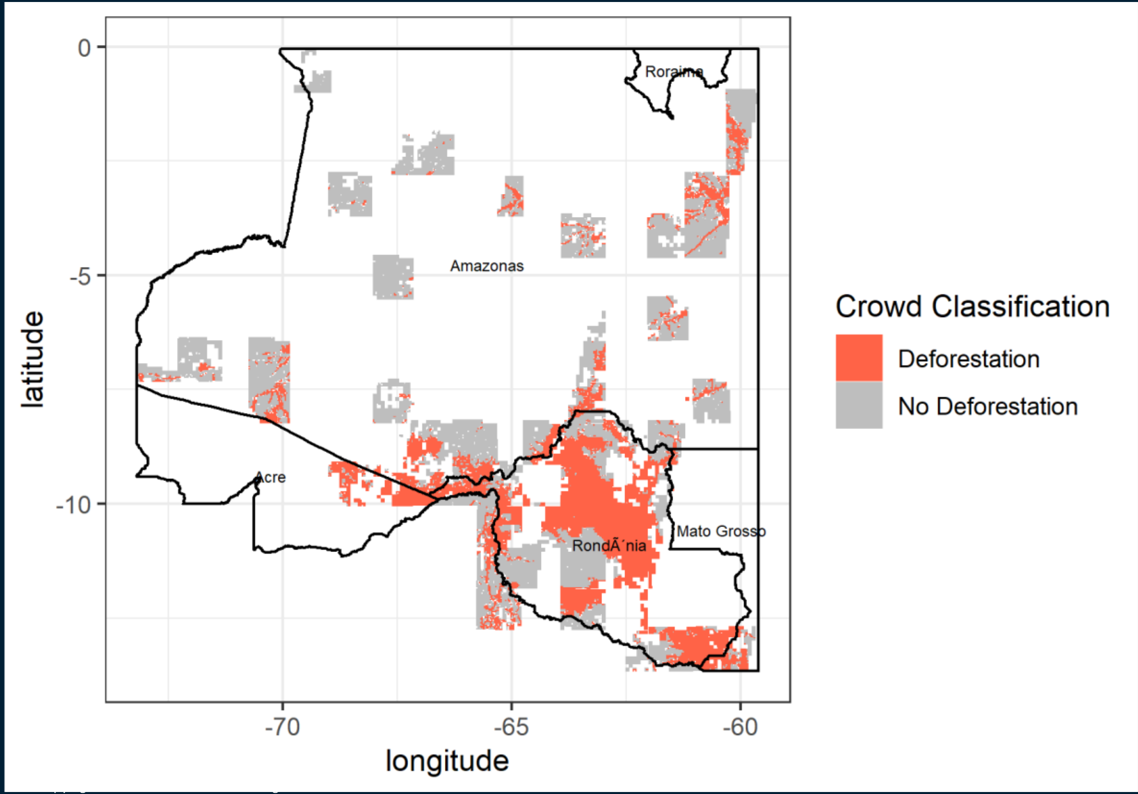

Comparing with existing products

- Comparison with existing products (PRODES and Maryland) difficult as they have differences in spatial details.
- After searching for >1km2 signs to compare with crowd results: 92% accuracy with PRODES and 89% with Maryland.

Überblick über den Modellierungsablauf

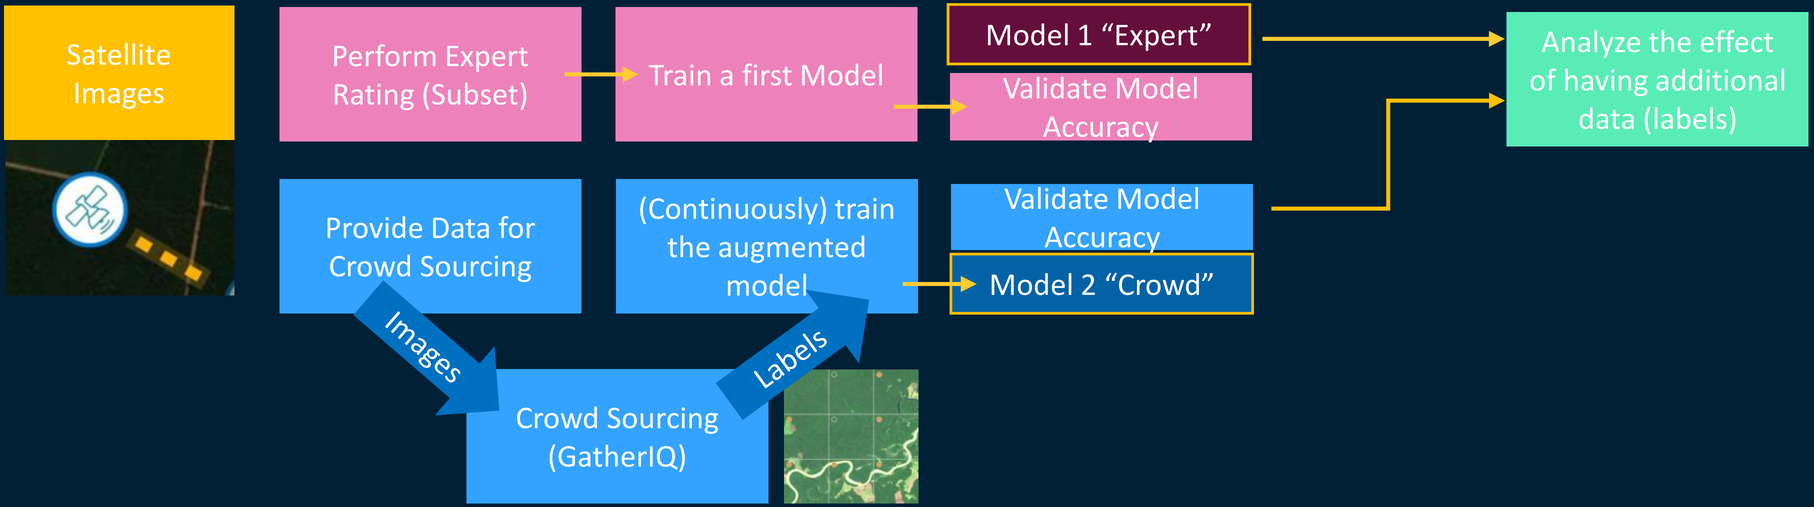

https://github.com/sassoftware/python-dlpy

README.md

DLPy - SAS Viya Deep Learning API for Python

SAS Deep Learning Icon

An efficient way to apply deep learning methods to image, text, and audio dat

SAS VIYA 3.4 | PIP INSTALL SAS DLPY | PYTHON 3+

Overview

DLPy is a high-level Python library for the SAS Deep learning features available in SAS Viya. D to provide an efficient way to apply deep learning methods to image, text, and audio data. D created following the Keras APIs with a touch of PyTorch flavor.

What's Recently Added

- Text, audio, and time series support in addition to image
- New APIs for:
  - RNN based tasks: text classification, text generation, and sequence labeling
  - Object detection
  - Image segmentation
  - Time series processing and modeling

master | python-dlpy / examples / | Go to file | Add fi

ChipRobie updated audioTraining example runs end-to-end with new supplied data 7a425c7 on Apr 29, 2021

..	
auto_encoder	Updated autoencoder example. (#319)
functional_model	Adding bulding_model.rst
image_captioning	Hide formatting tags in image captioning example. (#316)
image_classification	Update image classification with EfficientNet example. (#314)
image_denoising	New autoencoder denoising example. (#309)
image_embedding	Notebook Edits
keras_model_conversion	Merge pull request #293 from ChipRobie/DEEPLRN-245
keypoints_detection	Edit keypoints example
learning_rate_policy	Update learning rate example notebook. (#313)
misc	- adding mist folder to examples as well as an example of video proce...
multitask_learning	adding another example notebook
object_detection	Fast RCNN Soccer: update input dataset and model weights
onnx	Revised and clarified import ONNX model example. (#317)

Load Data

```
cas = swat.CAS(hostname, port)
cas.sessionprop.setsessopt(timeout=60 * 60 * 12) # 12 hour timeout

for a in ('deeplearn', 'sampling'):
    cas.loadactionset(a)

train_tbl = ImageTable.load_files(cas, TRAIN_IMAGE_FOLDER)
valid_tbl = ImageTable.load_files(cas, VALID_IMAGE_FOLDER)
test_tbl = ImageTable.load_files(cas, TEST_IMAGE_FOLDER)

train_tbl.show()
```

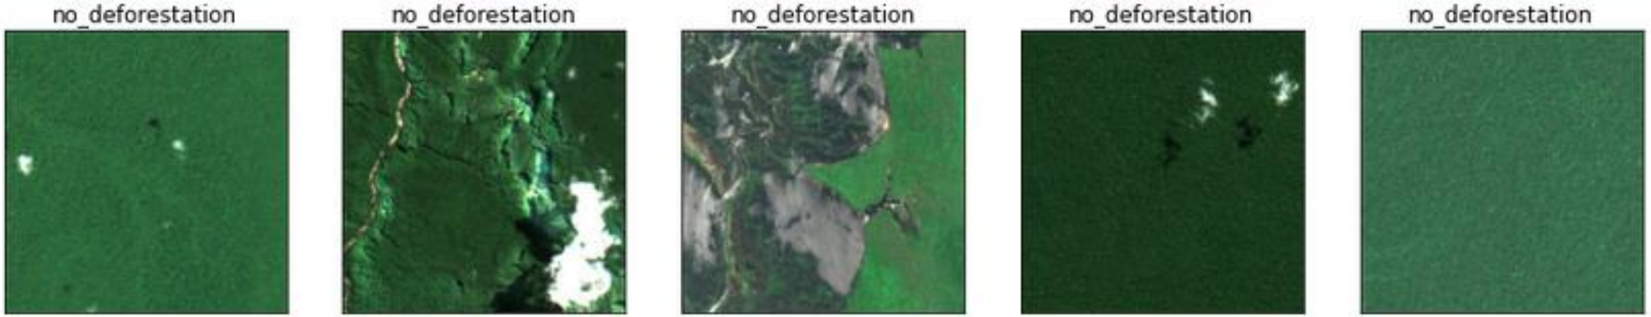

Model Training (Python Syntax, SAS DLPY Library)

```
from datetime import datetime

import dlpy
from dlpy.applications import MobileNetV2, Sequential, ResNet18_SAS, ResNet34_SAS, ResNet50_SAS, DenseNet
from dlpy.layers import Conv2d, Dense, InputLayer, OutputLayer, Pooling
from dlpy.model import AdamSolver, VanillaSolver


def train_model(model):
    start_d = datetime.now()
    print(start_d)

    # Train on the sample
    train_history = model.fit(
        data=train_filter_tbl,
        valid_table=valid_filter_tbl,
        optimizer=Optimizer(algorithm=AdamSolver(learning_rate=1e-4, clip_grad_max=50, clip_grad_min=-50),
                            mini_batch_size=32,
                            reg_l2=1e-5,
                            stagnation=5,
                            max_epochs=50),
        gpu=GPU,
        save_best_weights=True,
        train_from_scratch=True)

    end_d = datetime.now()
    print('Finished in %s' % (end_d - start_d))
    return train_history
```

Deep Learning Methods

- 43000 images in the crowdsourced library
- 60/20/20 split for train, validation, test
- Models trained: VGG1629, ResNet18, ResNet3430, MobileNet31.
- ResNet18 model slightly outperformed all others

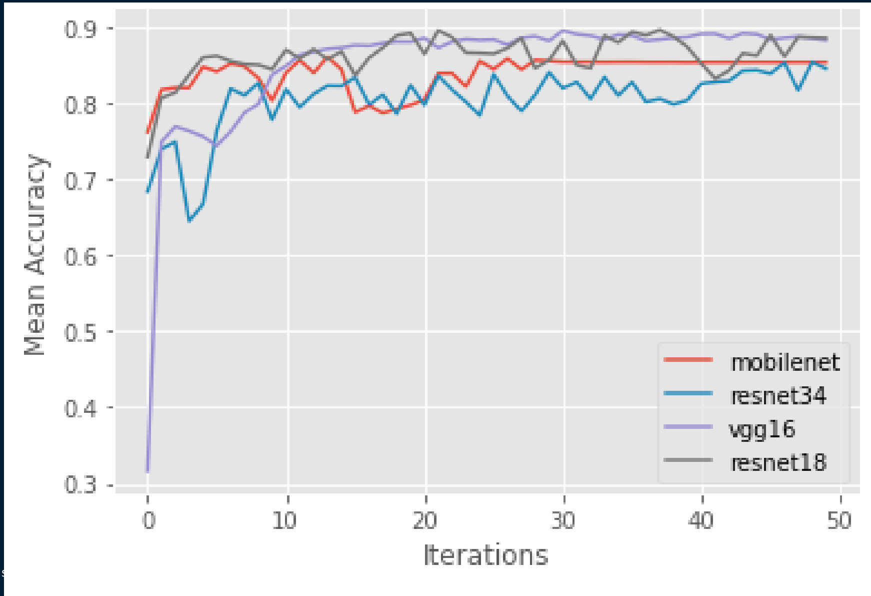

Data Selection and Preprocessing

- Crowd sourcing is based on the visible spectrum (RGB). Sentinel-2 contains 13 optical bands of which we used three to create RGB images (Blue (~493nm), Green (560nm), and Red (~665nm)) with a 10 m resolution.
- Data pre-processing:
  - 5% cloud filter has been used
  - Picture with excessive coloud coverage or with majority of missing data have been removed
- Computer vision model is not restricted to visible wavelengths.
- Models that use the near-infrared spectrum or the actual computed indices, such as NDVI, may be able to more accurately distinguish between natural deforestation, water features, and otherwise disturbed or developed land.

Example results and their interpretation

Test dataset | occlusion sensitivity | confidence of the crowd

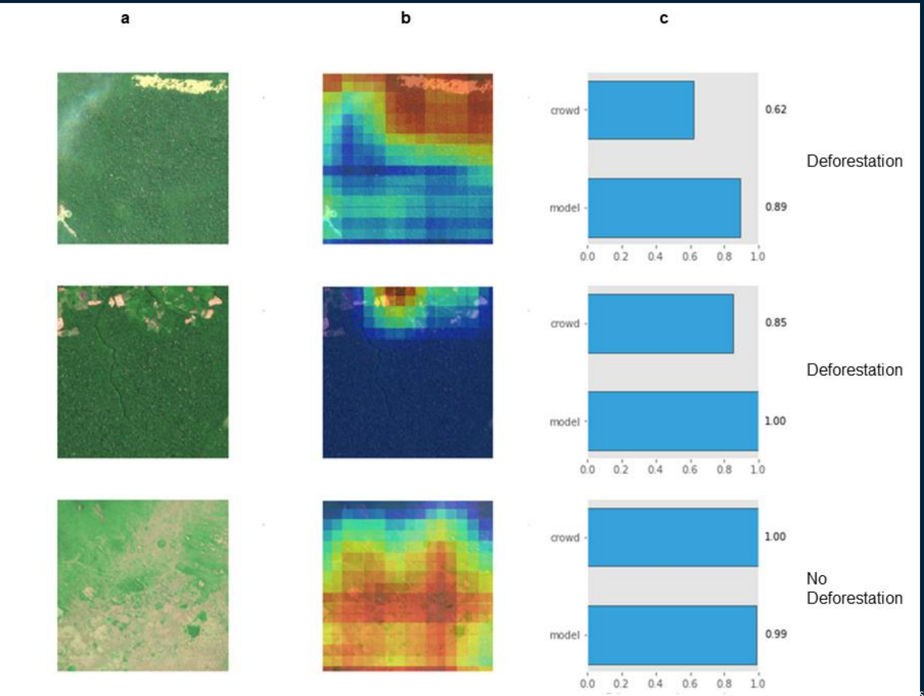

signs of deforestation in the image that trigger the activation layers, giving the model high confidence that human impact has occurred even though the crowd was less confident

signs of human impact at the top of the image trigger the activation layer accordingly

both the crowd and the model identify the canopy disturbance as non-human impact with high confidence

Validation of the AI Model

- 94,8 agreement with the crowd on 8774 images
- 100 (unseen) sentinel-2 satellite images: comparing results between the AI model and the GLAD-S2 deforestation dataset

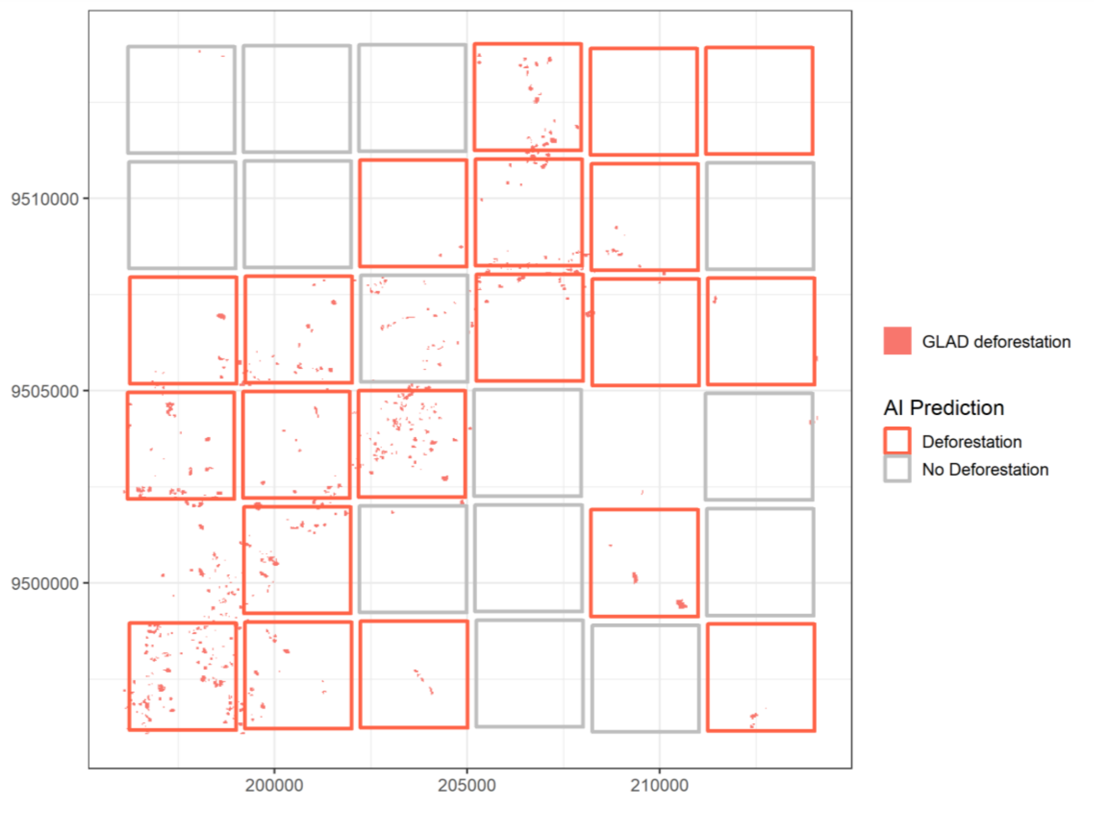

Überblick über den Modellierungsablauf

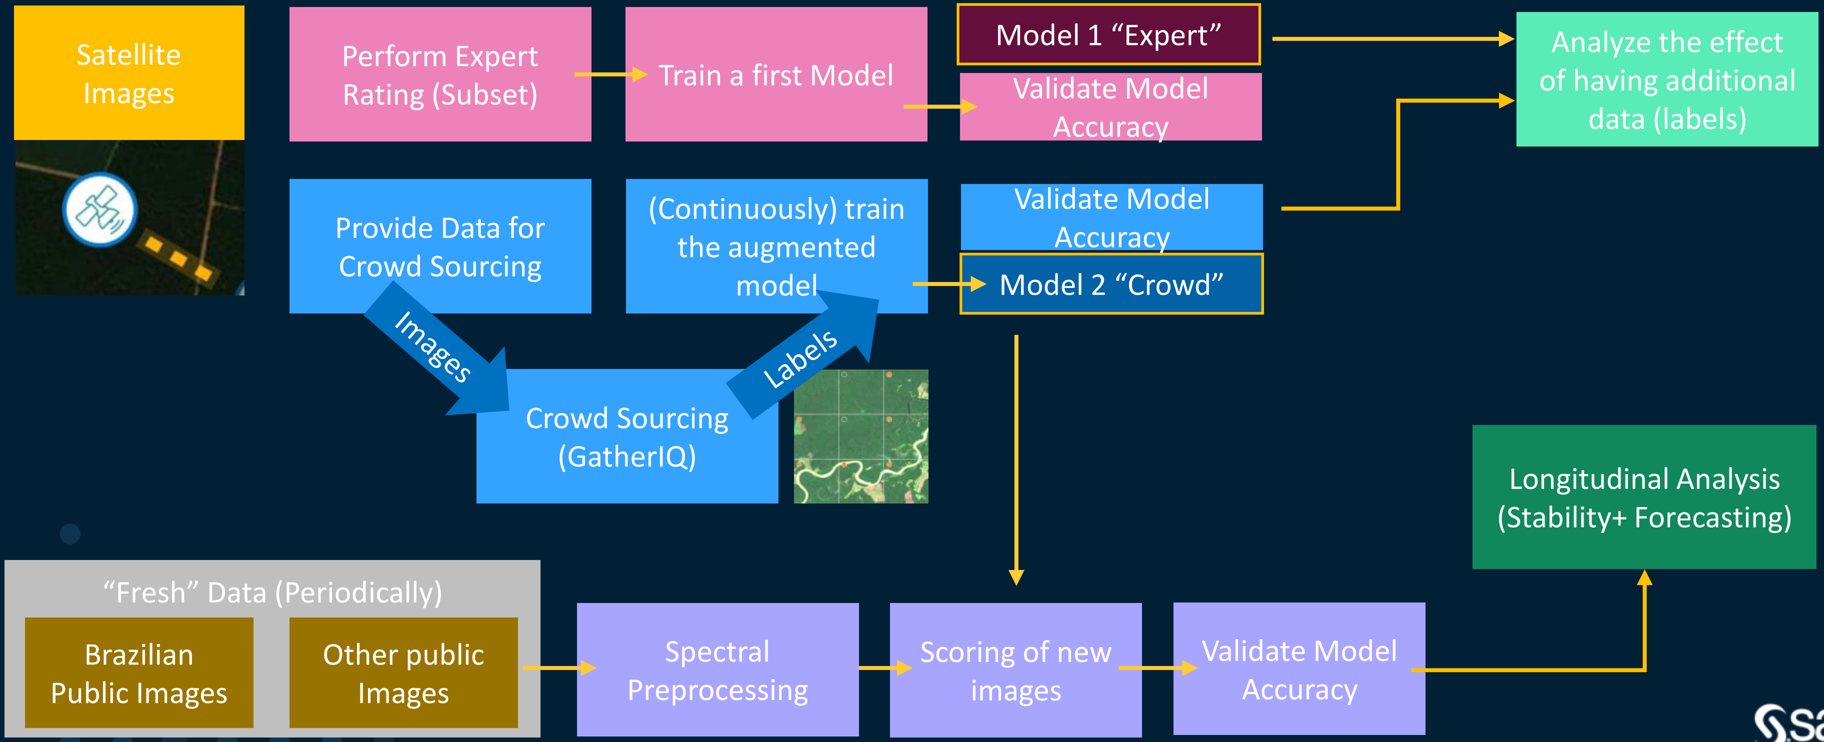

Motivation and advantage of crowd sourcing

- established as an alert system complementary to the existing monitoring systems in place
- address the limited uptake of existing monitoring systems into actionable change or policy developments, through inclusion of the global civil community

Daily Progress Report

While computer vision models can be trained to quickly identify areas of the rainforest that have been significantly damaged, signs of deforestation can be challenging for a computer to see at first. It takes human eyes to properly classify images in order to build models that can detect the subtle differences between satellite imagery. For every image that you select as an area of deforestation in our crowdsourcing project, SAS and our partners get closer to building a model that can alert governments and conservation organizations.

From Earth Day, April 22, 2020, through February 2021, citizen scientists like yourself classified nearly 90,000 satellite images of the Amazon rainforest. The results from the first phase can be reviewed here.

Our work is not finished. During the second phase, which started March 2021, we will continue combining SAS AI technology, your human input, and our partner's expertise to identify where changes are occurring over time. This could one day help predict where deforestation is likely to happen next.

Note: The results below are based on crowdsourced consensus, meaning multiple agreements from citizen scientists like yourself are required to be reflected as one result. The results are provided in the charts below, which are updated twice daily.

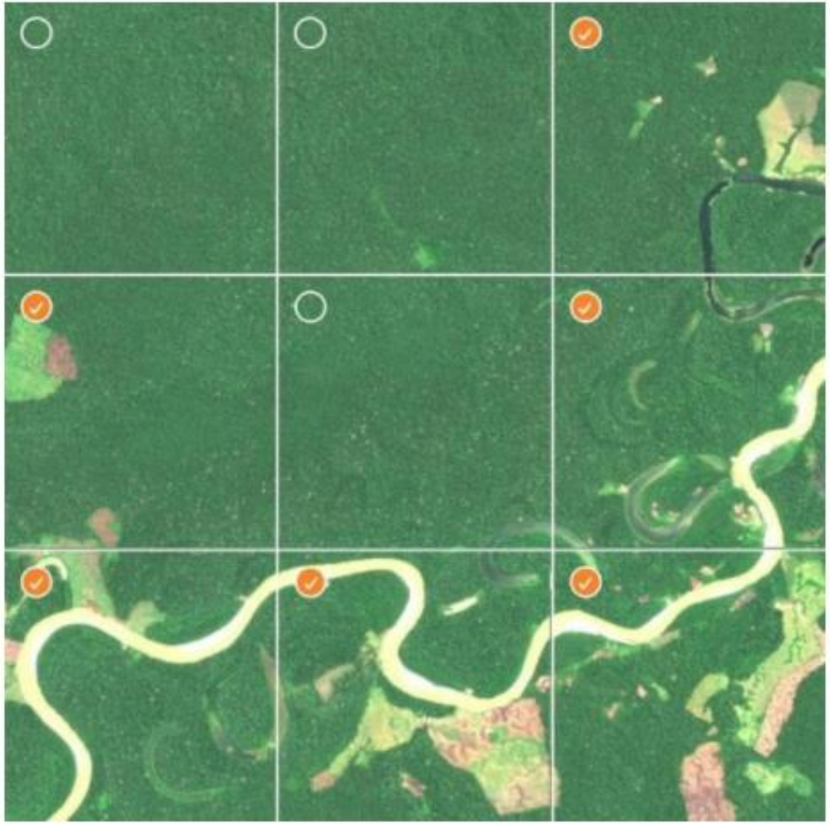

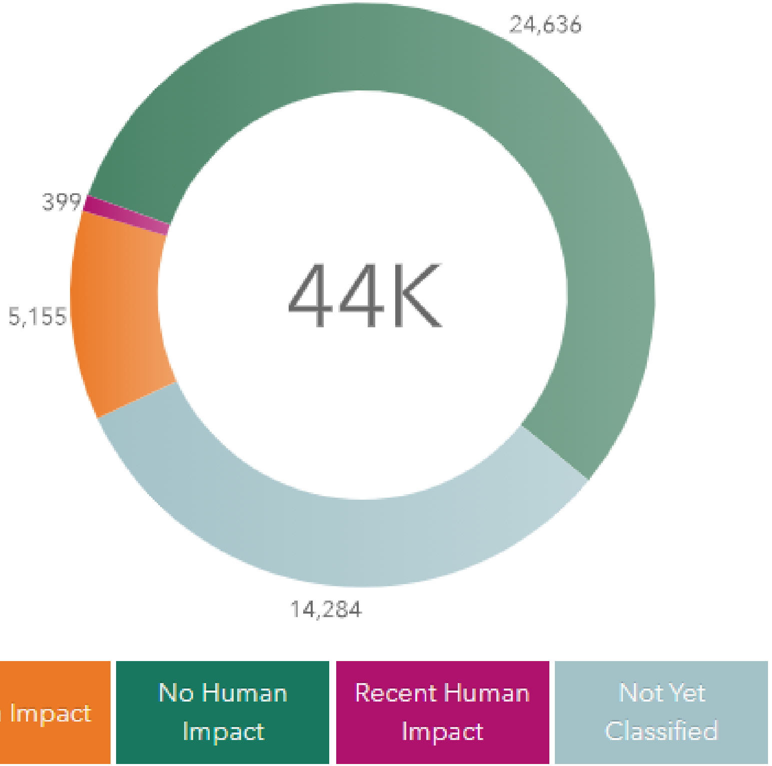

Current Assessments

Through crowdsourcing, citizen scientists from 110 countries helped identify human impact in 90,000 satellite images of the Amazon rainforest during phase one.

Today, in phase two, our goal is to classify 45,000 more images to identify recent human impact in these areas. Here's our progress towards reaching our goal for phase two and how much more we need to do.

Crowd vs. Model: Comparing the classifications

- established as an alert system complementary to the existing monitoring systems in place
- address the limited uptake of existing monitoring systems into actionable change or policy developments, through inclusion of the global civil community
- crowd appears to perform better in terms of classifying rivers and roads (water can appear in different colors, reflection)
- Additional filtering on image data could be applied
- Image size for crowd sourcing: 3 x 3 km for visual quality and identification of large areas. However less precise about the exact location

Analysis over time

- System detects actual status of deforestation, irrespective when it occurred. Known areas can be specially flagged and removed in the system

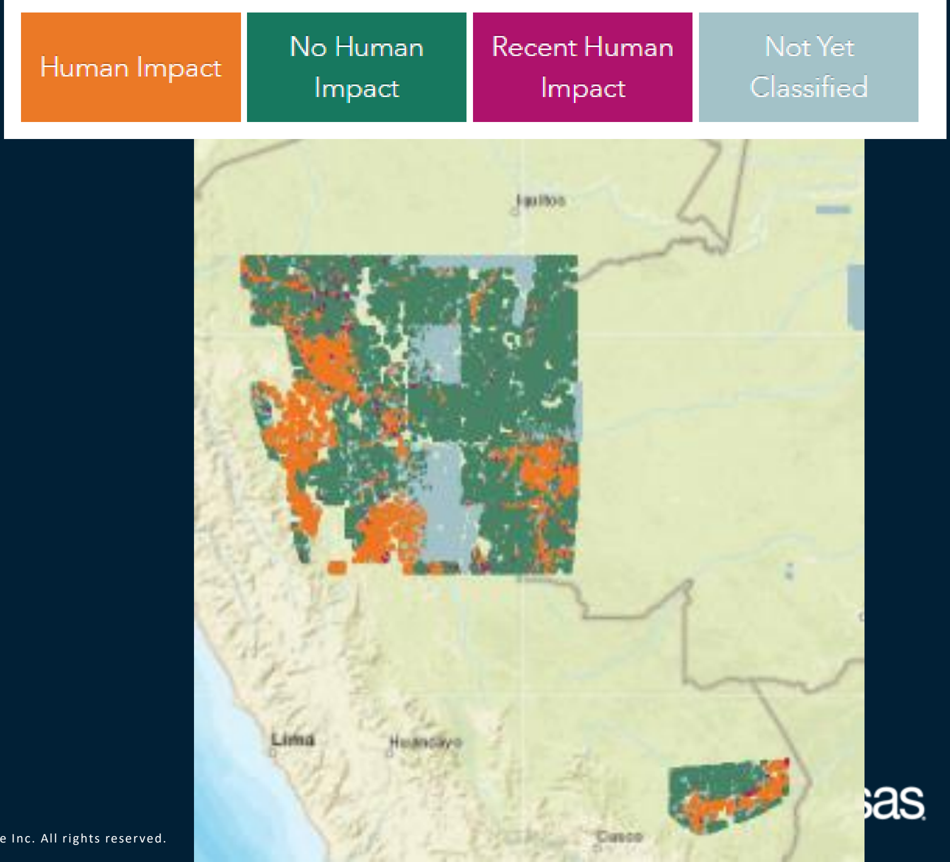

Possible future developments

- Select how images are selected for labeling by the crowd and how many times each image should be labelled to optimize the efforts of the crowd (e.g. apply a Bayesian approach that would remove images once a minimum level of confidence in the answers was reached)
- Users were not required to log-in to eliminate hurdles. However challenging to identify unique labelers and Inter-Rater Reliability cannot be used
- Approach applicable to areas like migration, food security, natural hazards and pollution

Der analytische Lebenszyklus

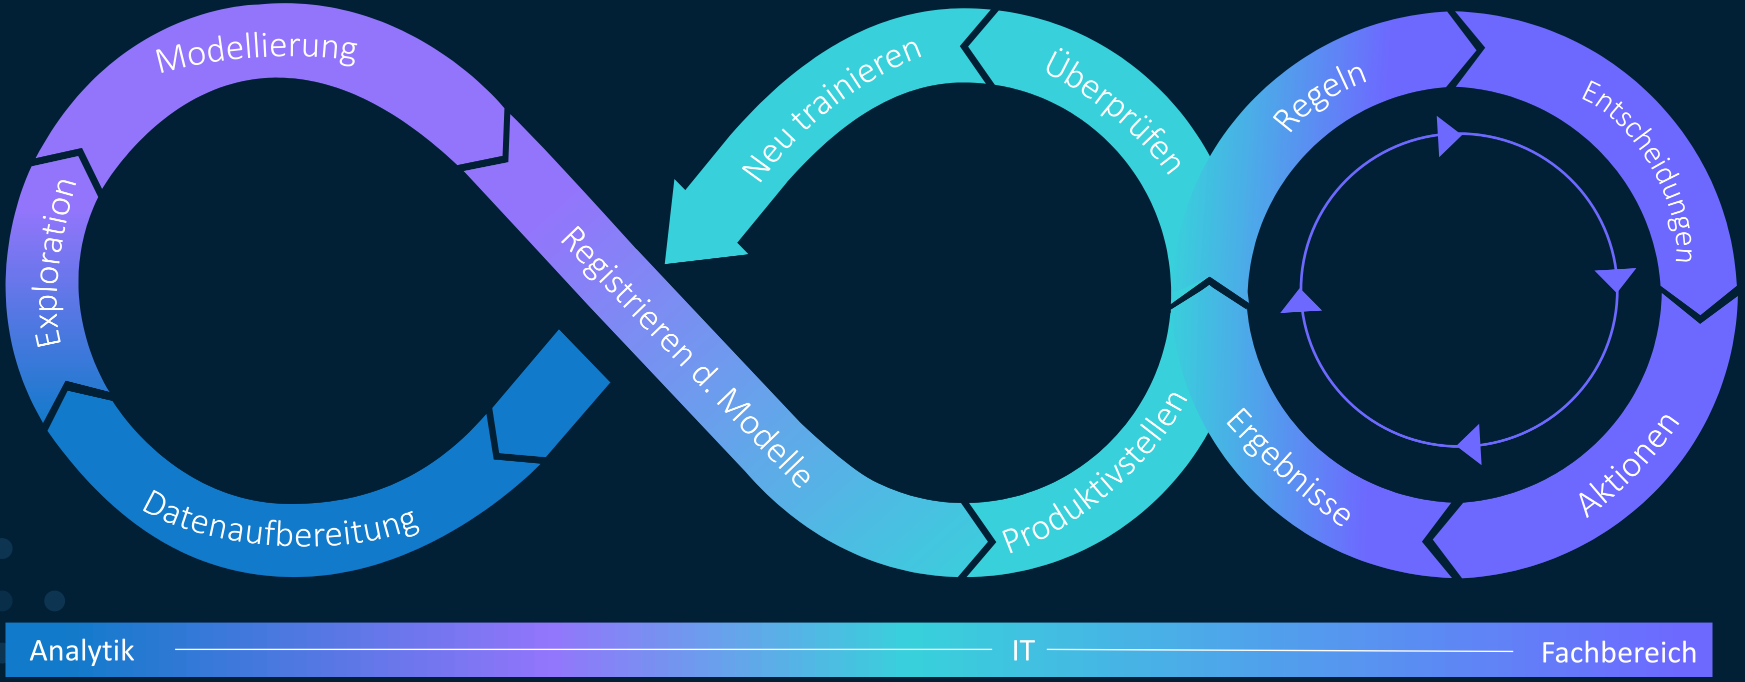

Links

- https://www.sas.com/de_at/data-for-good/rainforest.html
- https://developer.sas.com/home.html
- https://developer.sas.com/guides/dlpy.html
- https://github.com/sassoftware/python-dlpy
- **IIASA, SAS. (2022) Crowd-driven deep learning tracks Amazon deforestation. *In prep,***
  - ***Ian McCallum, International Institute for Applied Systems Analysis, Laxenburg, Austria,***
  - ***Jon Walker, SAS Campus Drive, Cary, NC 27513, United States***