Jet Propulsion Laboratory
California Institute of Technology

COSMIC: Content-based Onboard Summarization to Monitor Image Change

Detecting fresh impacts in HiRISE images with fully-convolutional networks

---

Asher Trockman
Mentor: Dr. Lukas Mandrake

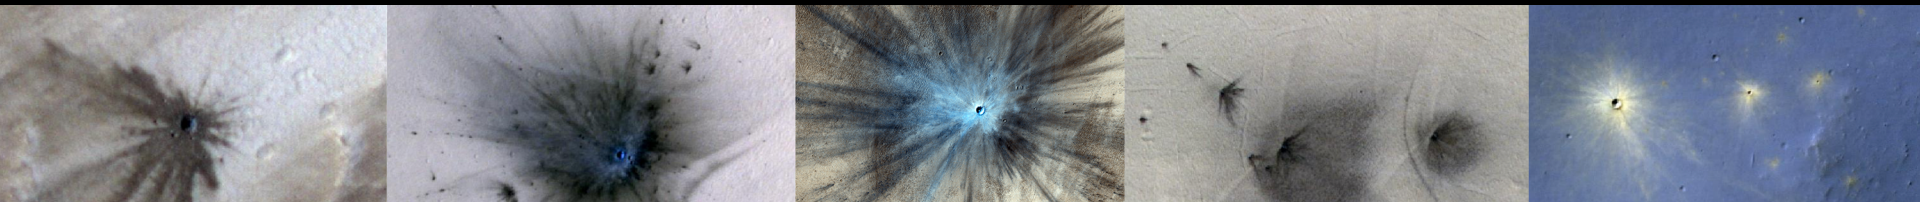

Fresh impacts

The current martian cratering rate, Daubar et al. (2013)

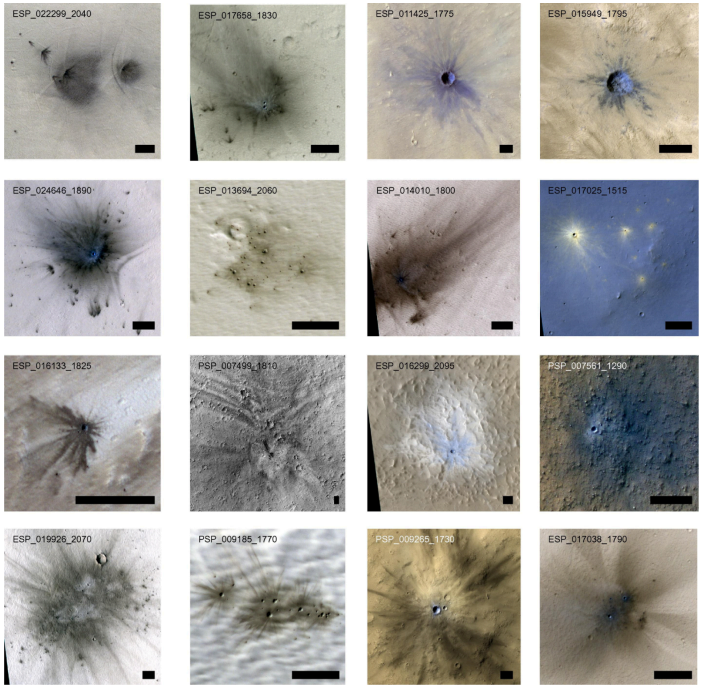

from Daubar et al. (2013)

- Found by looking at before/after CTX images (Daubar et al.)

- Crater diameters range from ~2.1—33.8 meters
- Craters > 3.9m occur at a rate of 1.65×10^{-6} craters/km^2/year
  - (or about 240 per year)

- Often appear in clusters
- **Blast zones are 10-100x the diameter of the crater itself**

- Relatively rare event, little training data
  - We have **308 images** of confirmed impacts

Labeling fresh impacts

Enlisting the help of citizen scientists to generate labels en masse with *Zooniverse*

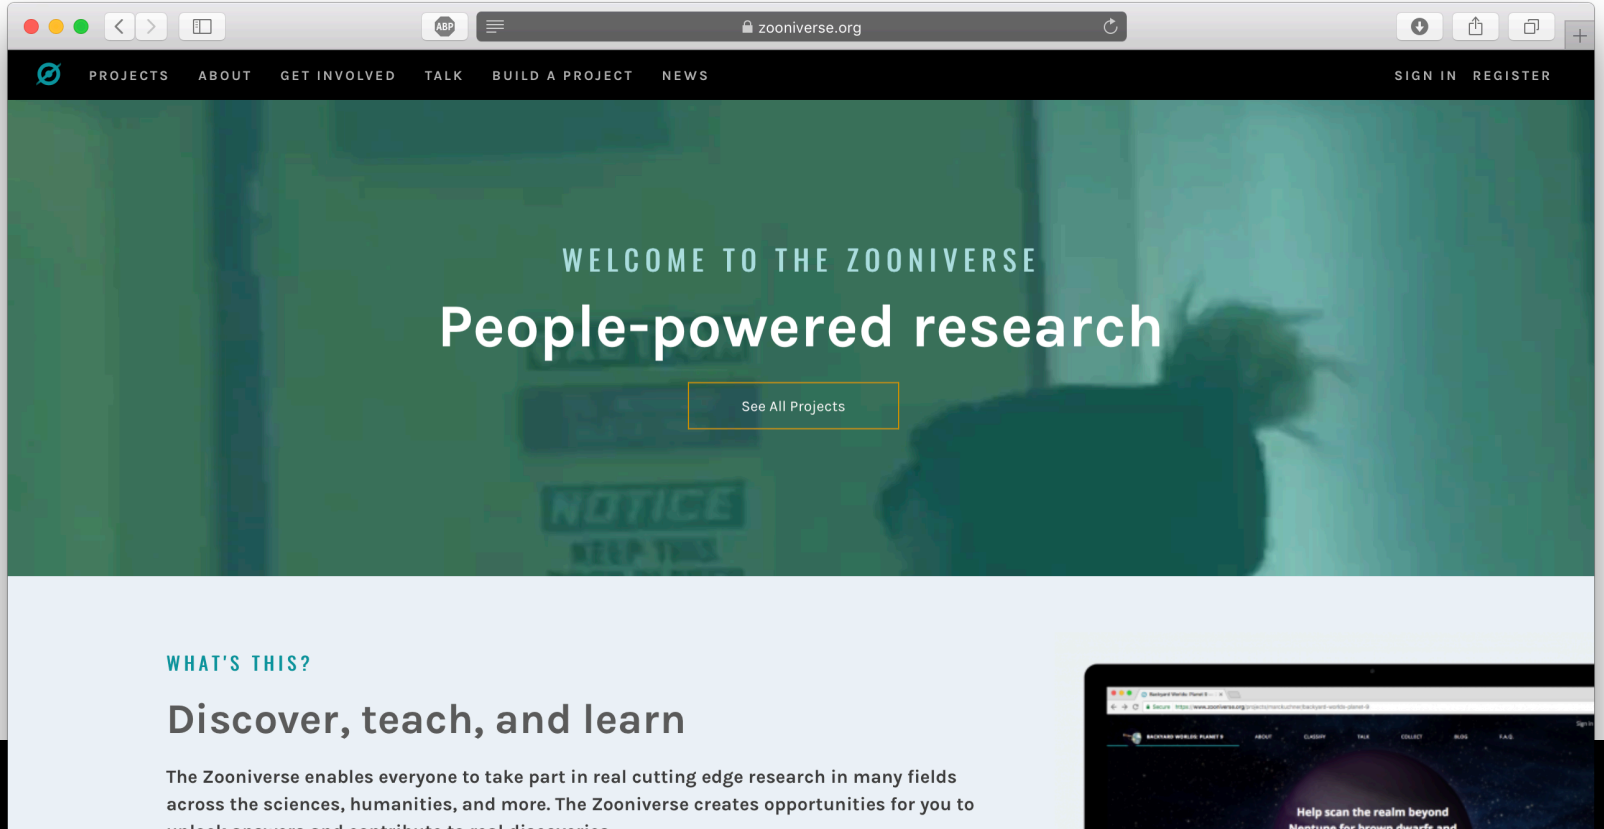

Labeling fresh impacts

Enlisting the help of citizen scientists to generate labels en masse with *Zooniverse*

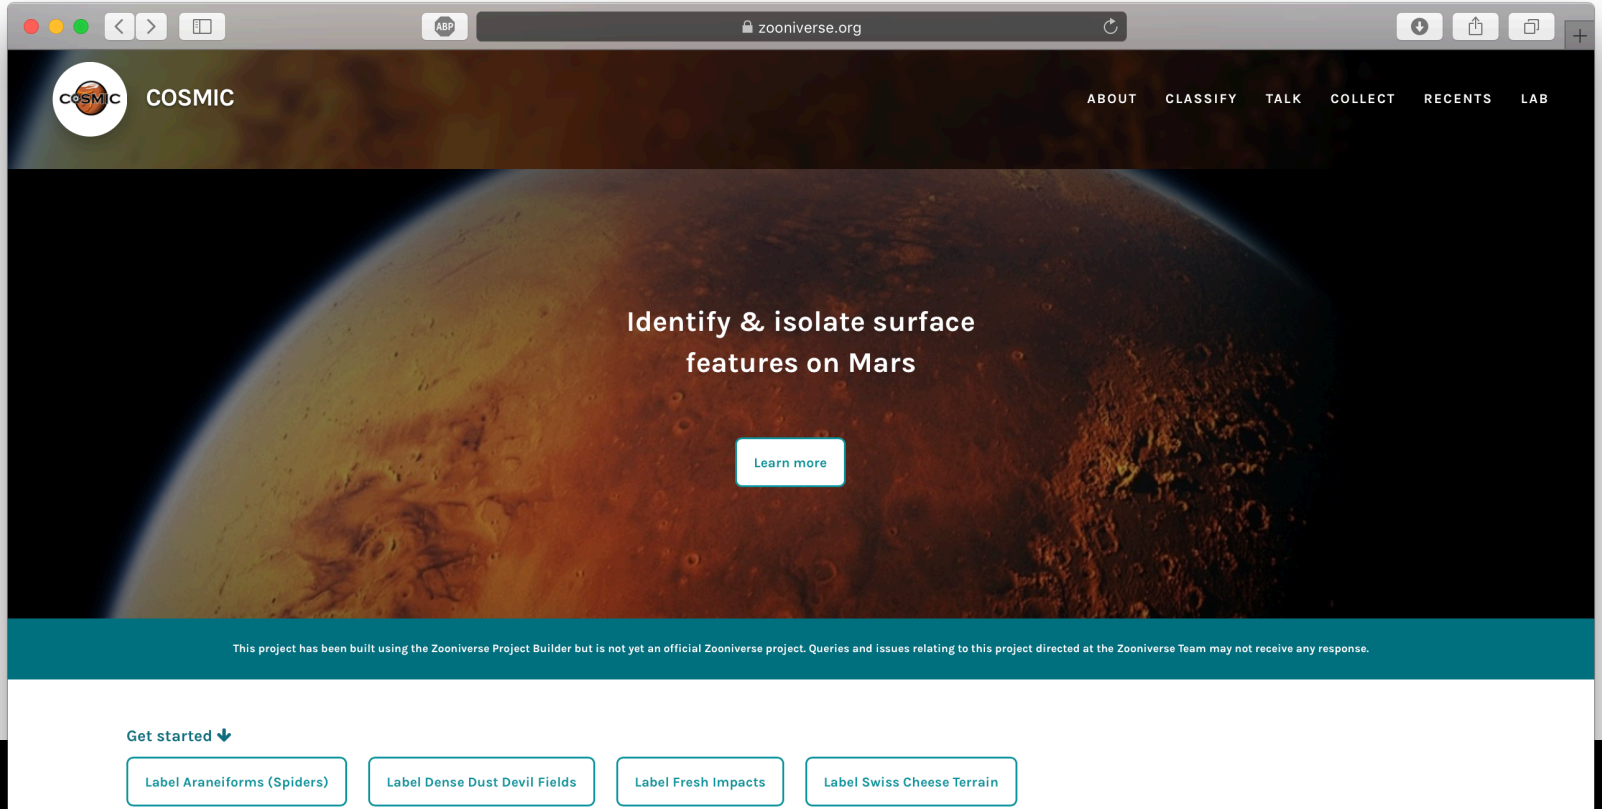

Labeling fresh impacts

Enlisting the help of citizen scientists to generate labels en masse with *Zooniverse*

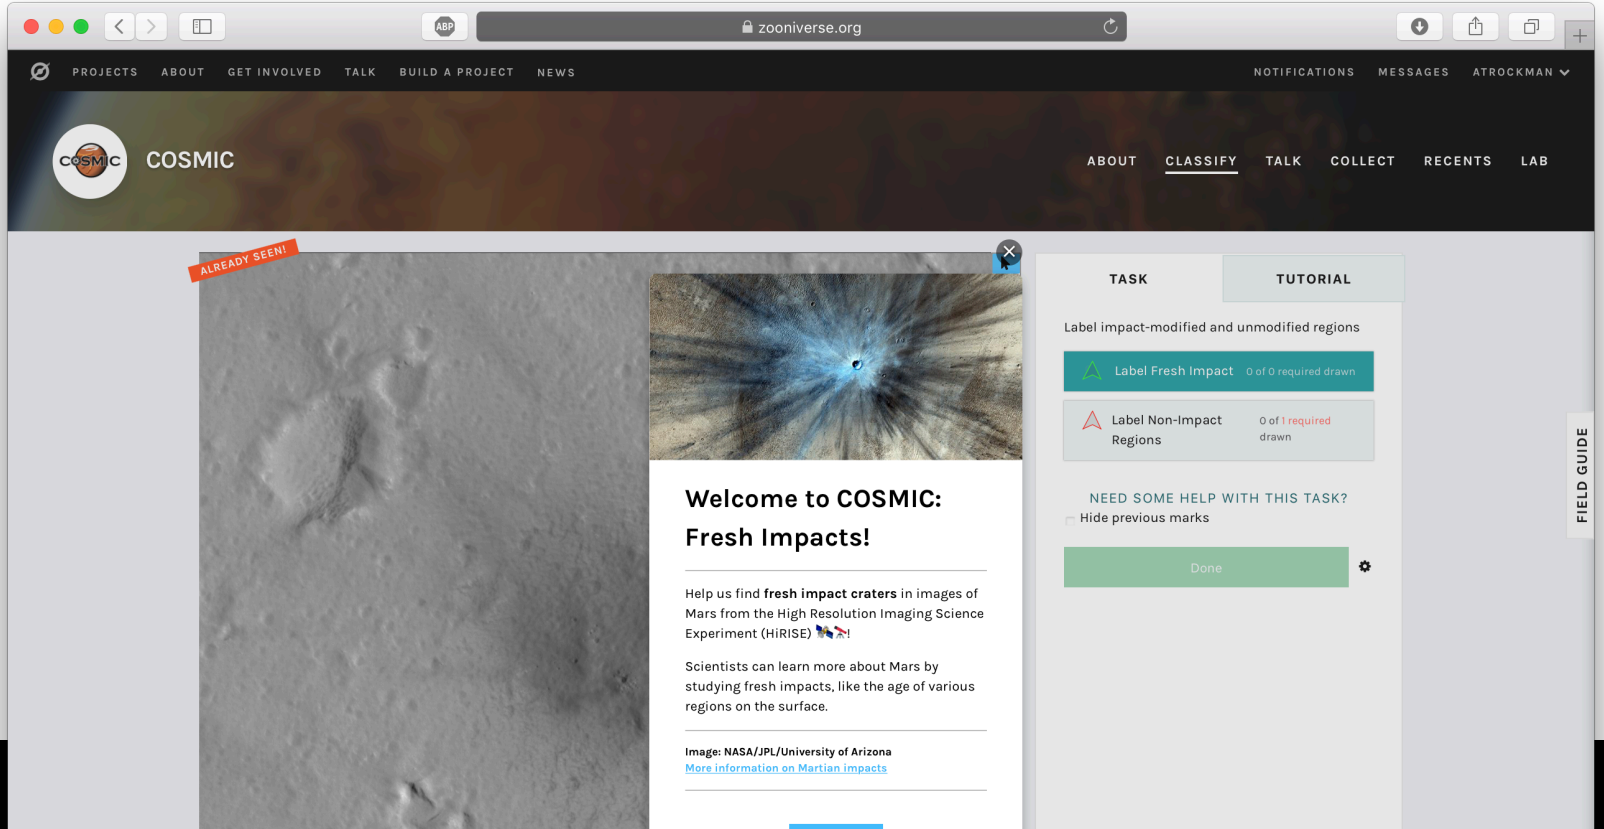

Labeling fresh impacts

Enlisting the help of citizen scientists to generate labels en masse with *Zooniverse*

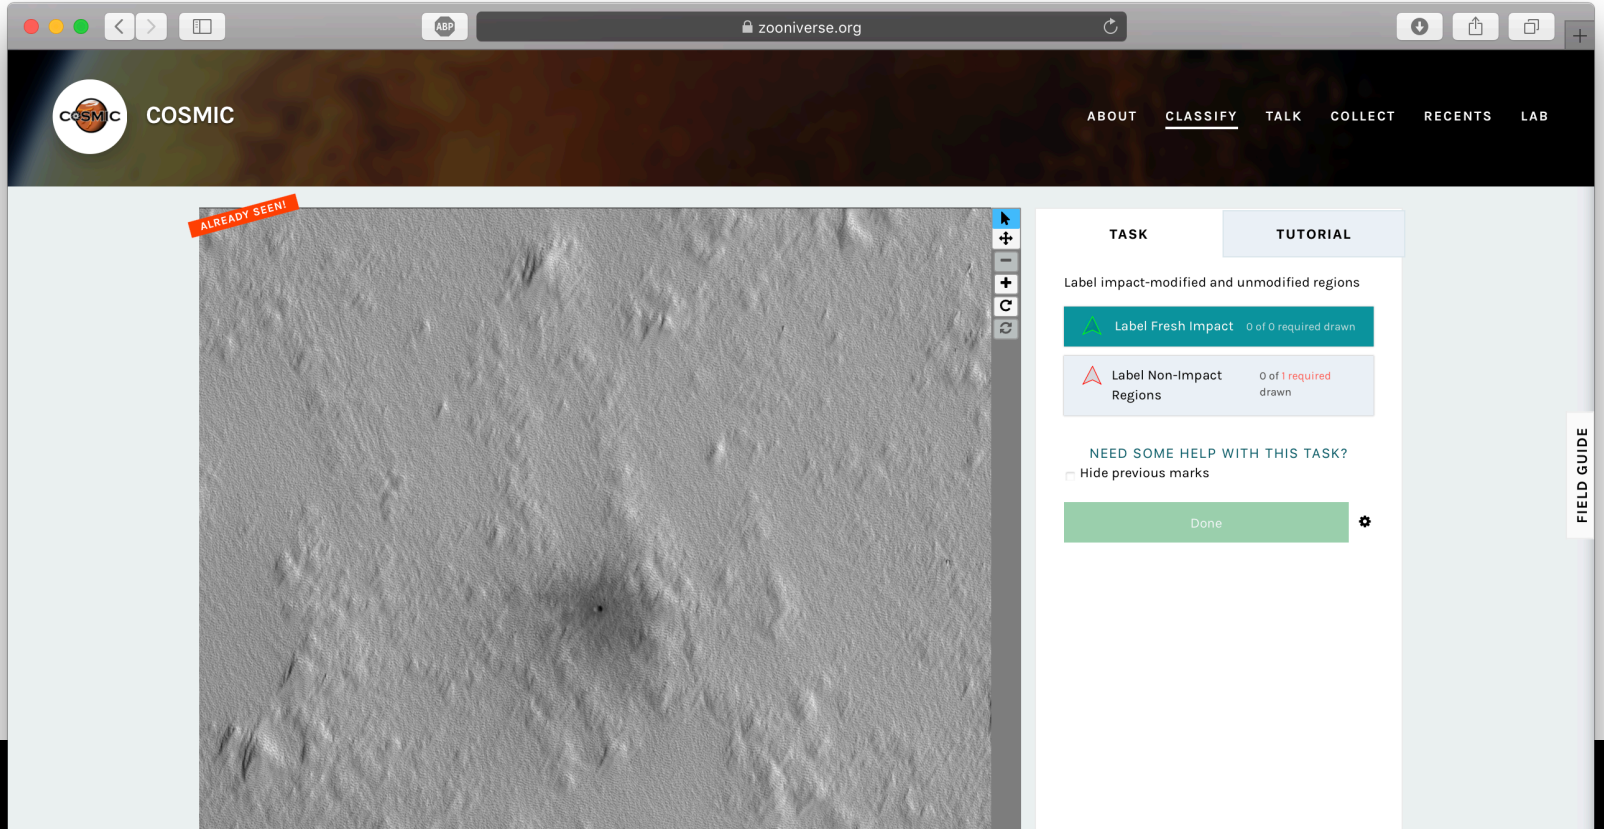

Labeling fresh impacts

Enlisting the help of citizen scientists to generate labels en masse with *Zooniverse*

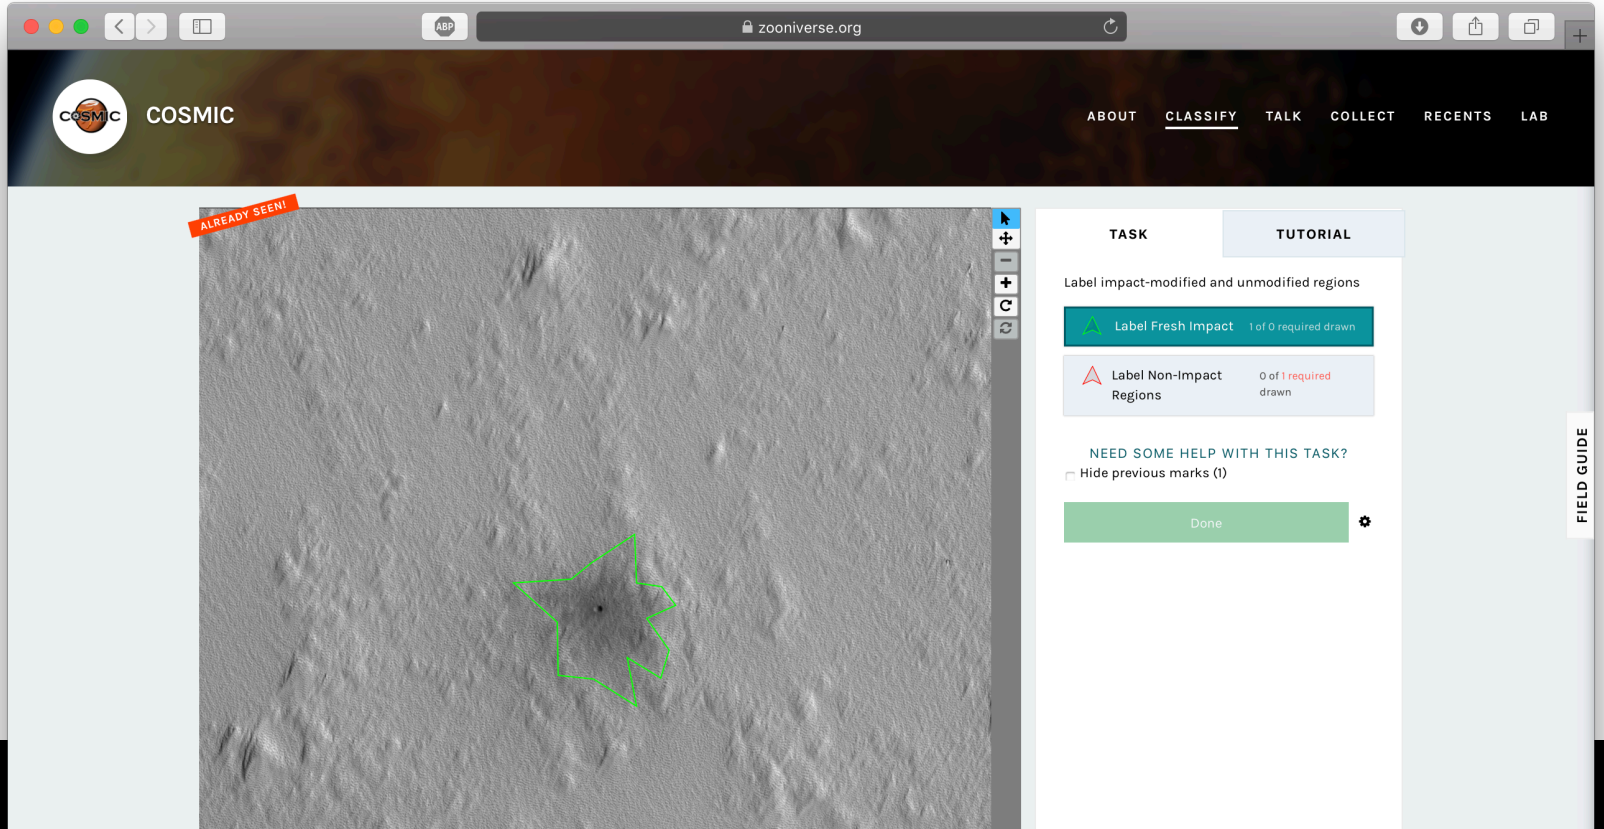

Labeling fresh impacts

Enlisting the help of citizen scientists to generate labels en masse with *Zooniverse*

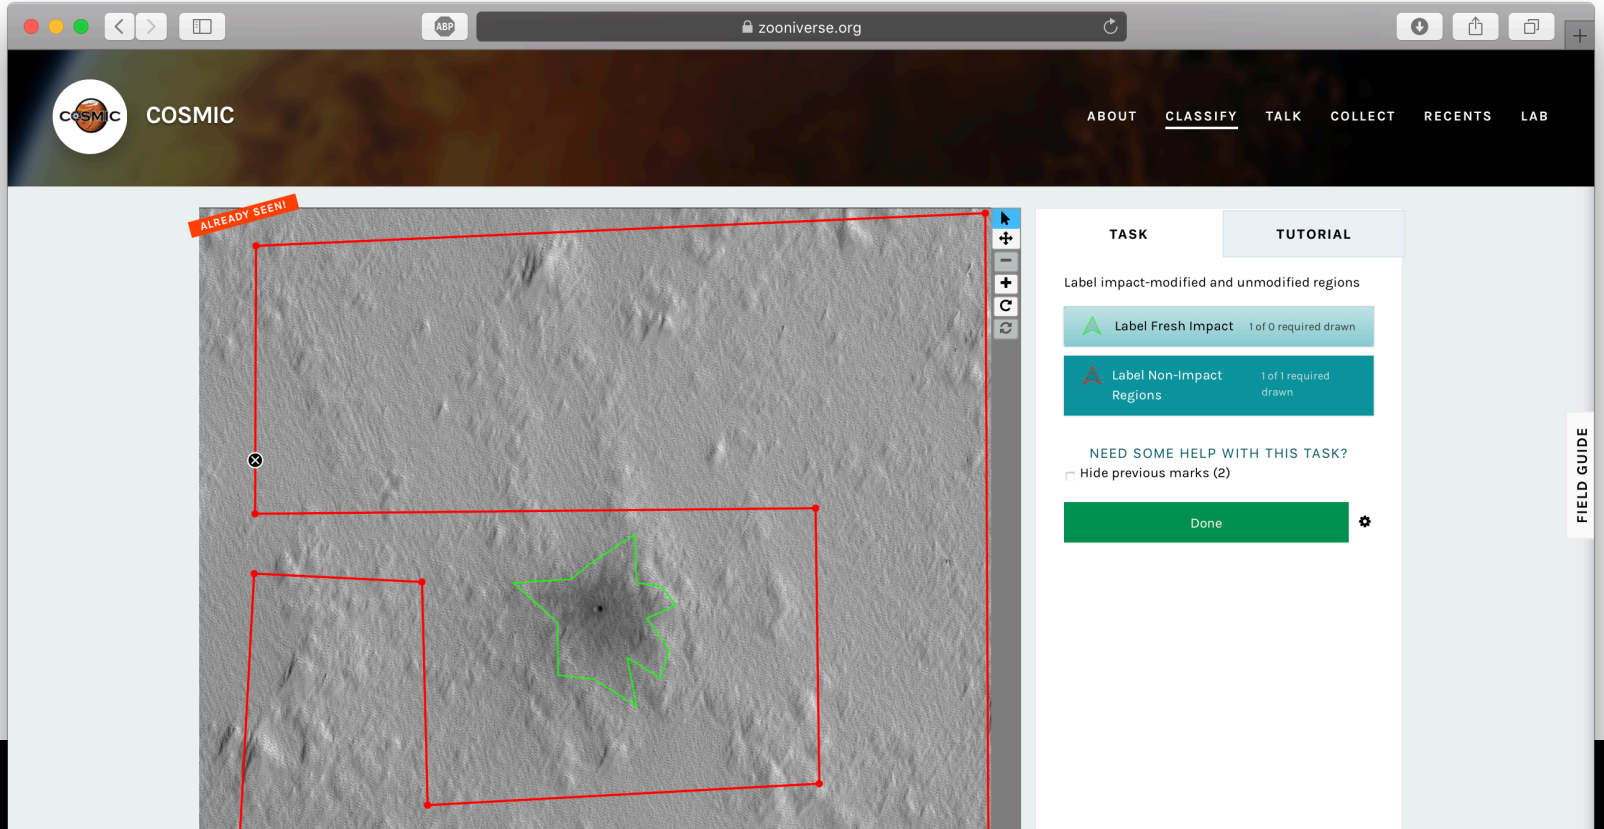

Multi-label problem

Merging labels by 50% inter-rater agreement (voting)

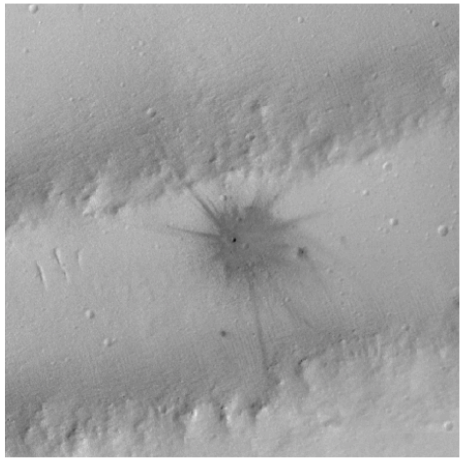

Fresh Impact Example

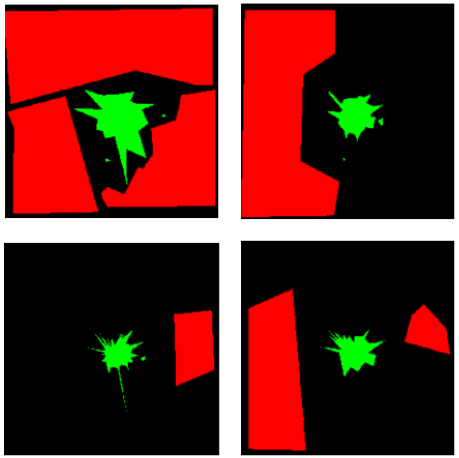

Labels from 4 Users

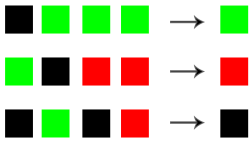

75% fresh impact agreement

50% non-impact agreement

no agreement

Merged Label
(Voting)

Multi-label problem

Extracting labeler confidence: real-valued labels

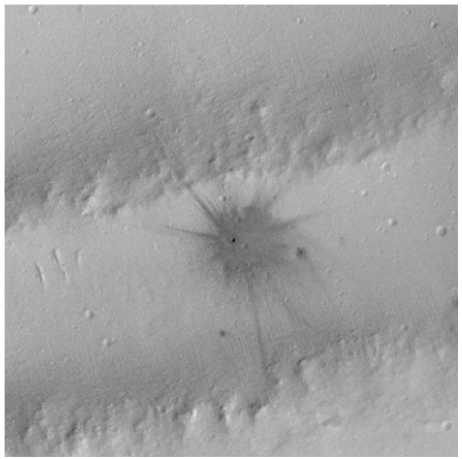

Fresh Impact Example

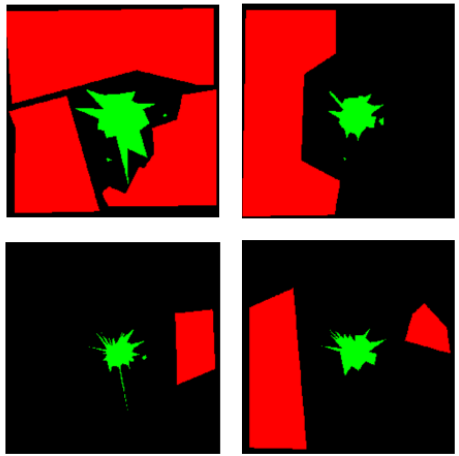

Labels from 4 Users

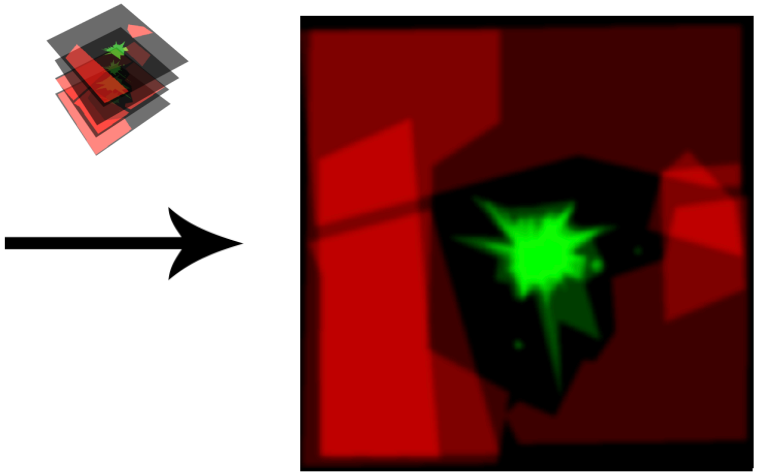

Merged Label
(Frequency)

Semantic Segmentation

Labeling image regions at the pixel-level. Example: MSRC dataset.

Example segmented images:

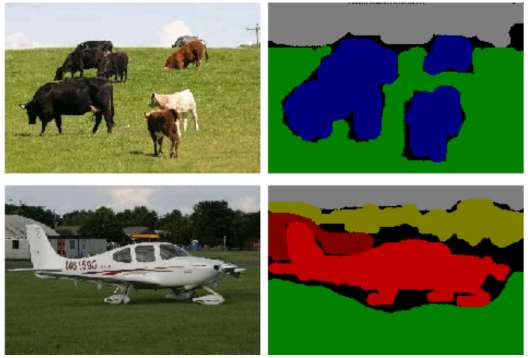

The goal of this work:

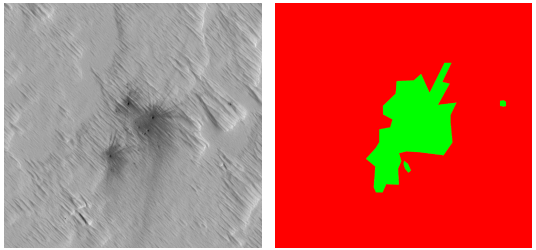

Some approaches:

- TextureCam (Semantic Texton Forests, Shotton et al.)
- Region-based CNNs (R-CNNs, Girshick et al.)
- **Fully-convolutional networks** (FCNs, Long et al.)

Fully Convolutional Networks (*Deep Filters*)

CNNs that work on input of *any size*

Most CNNs: Fixed size input

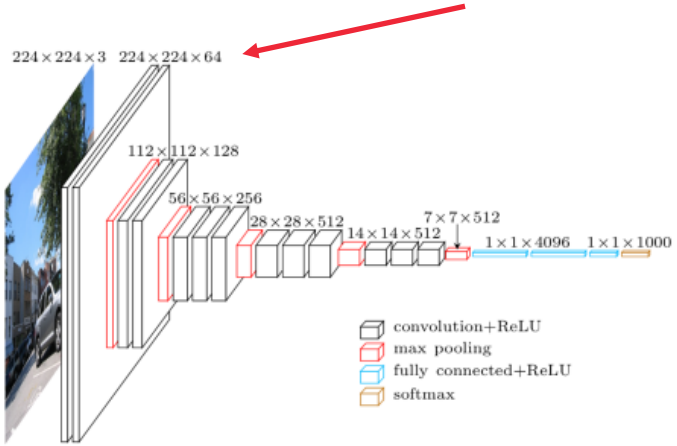

Convolution: Any size input

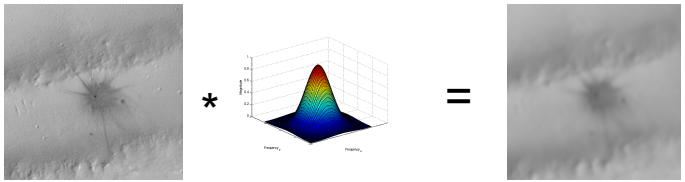

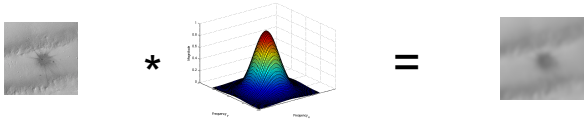

Problem: Fully connected (FC) layers → ***Solution:*** Convert FC to conv layers.

Fully Convolutional Networks *for Semantic Segmentation*

Long et al., 2015

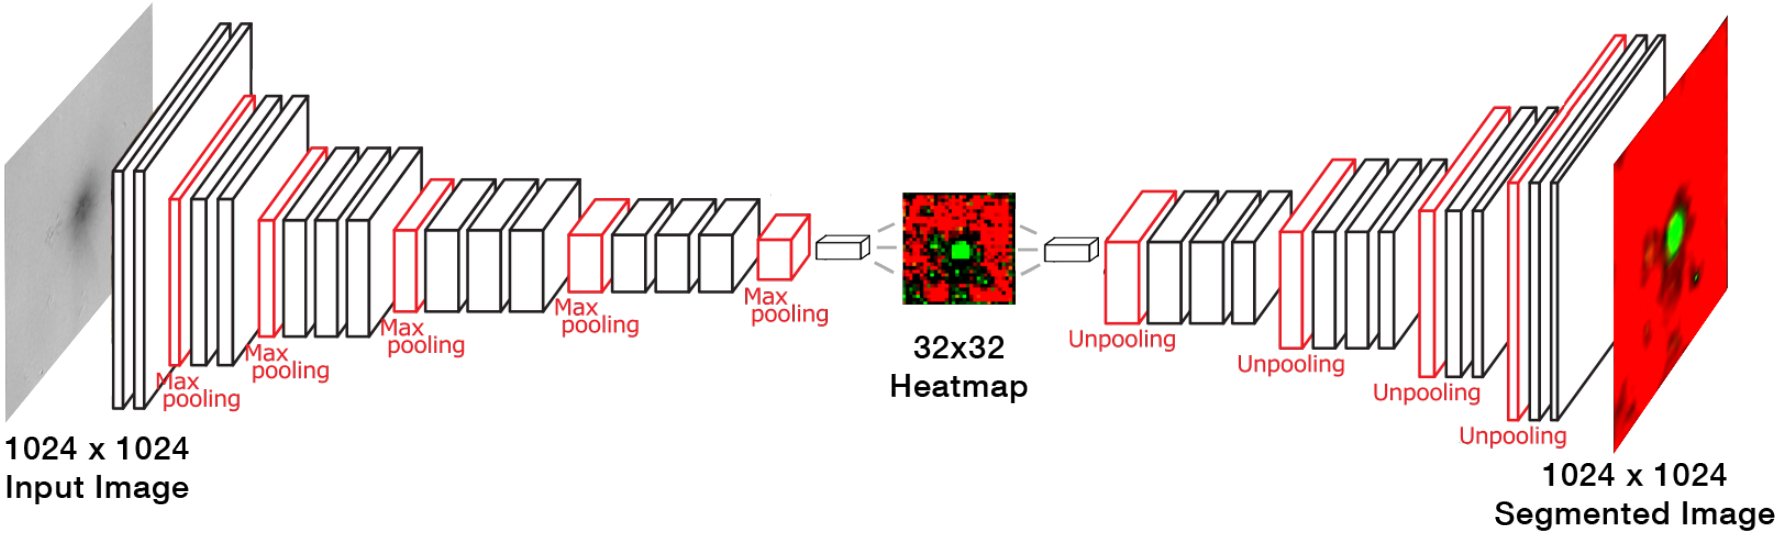

Fully Convolutional Networks *for Semantic Segmentation*

Long et al., 2015

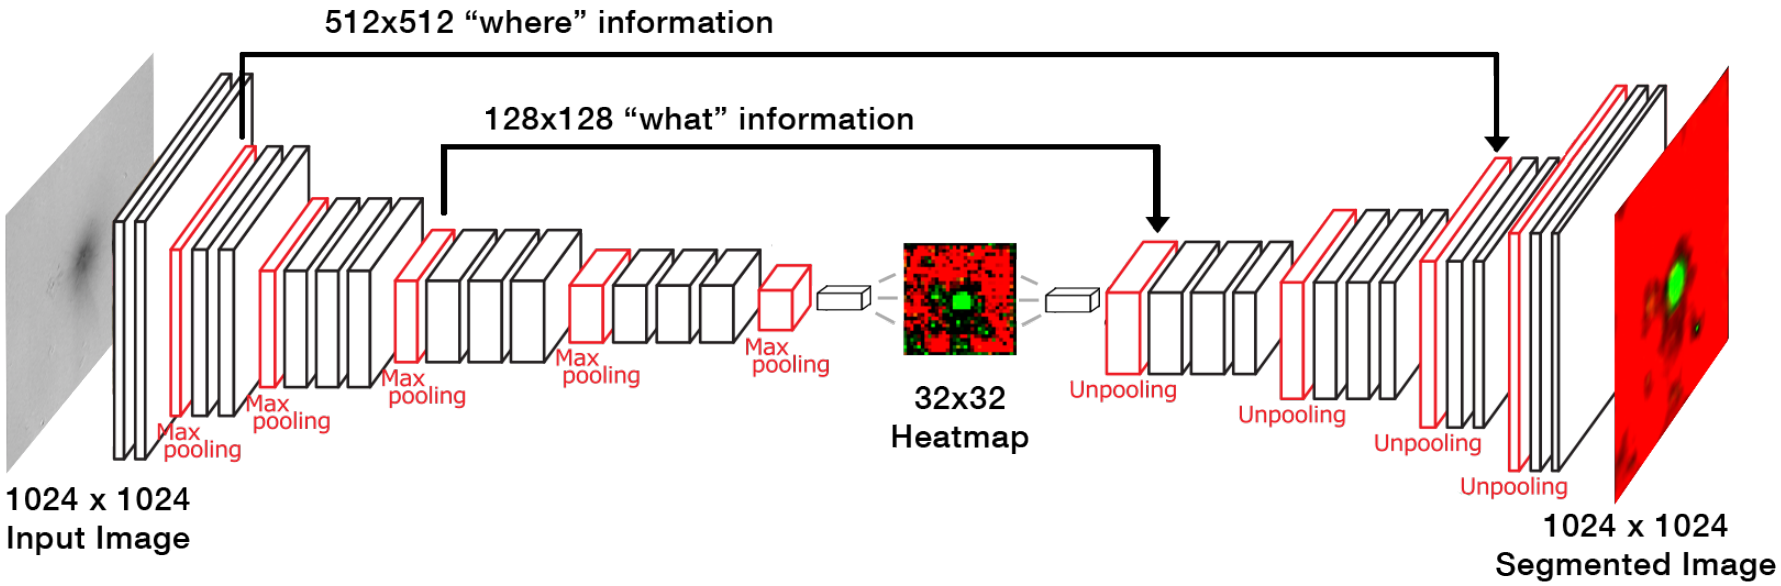

Fully Convolutional Networks *for Semantic Segmentation*

Long et al., 2015

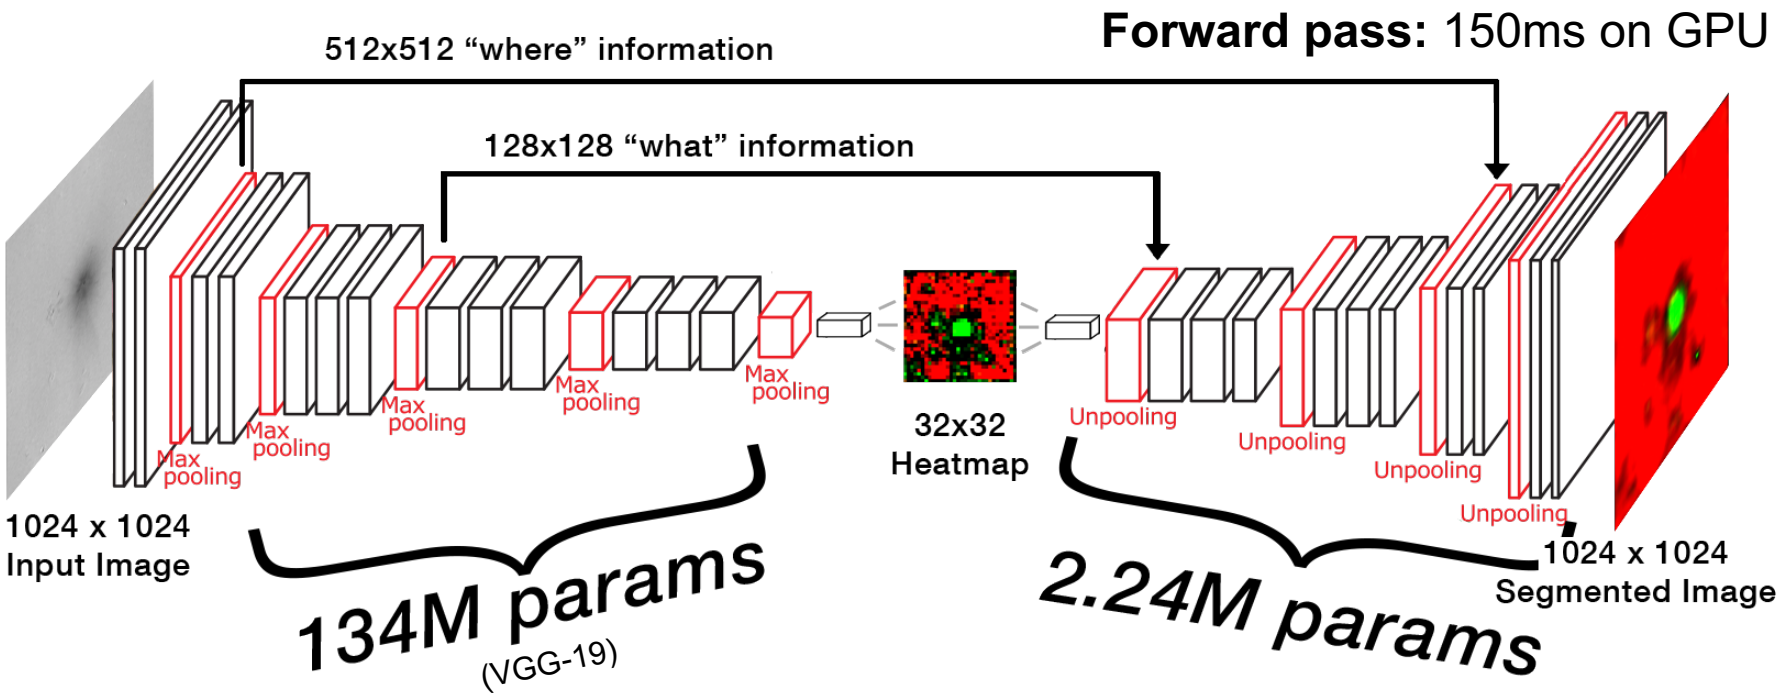

Training with real-valued labels

Instead of pixel classification, do regression and choose confidence thresholds

$$Loss \approx \left\| \text{Network Output} - \text{Merged Labels} \right\|_2$$

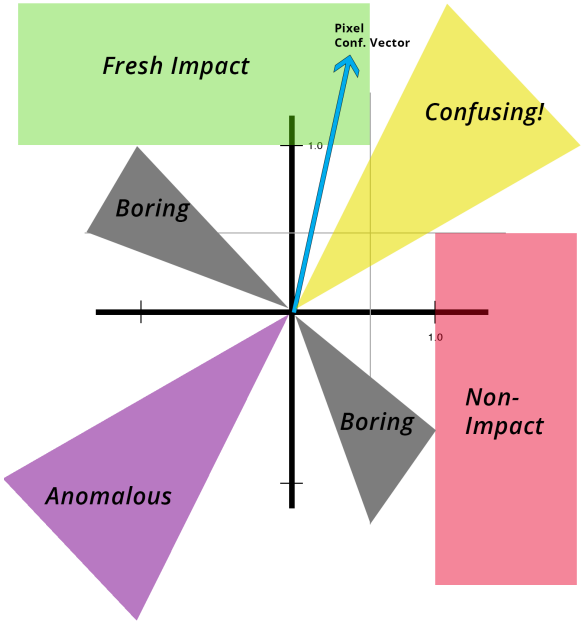

Per-pixel loss is the mean-squared error of the network's pixel and the label's pixel

Each pixel is a 2-vector

Training with real-valued labels

Instead of pixel classification, do regression and choose **confidence thresholds**

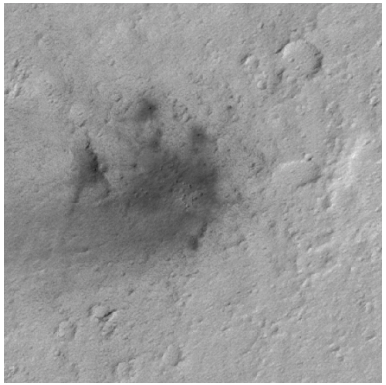

Fresh Impacts

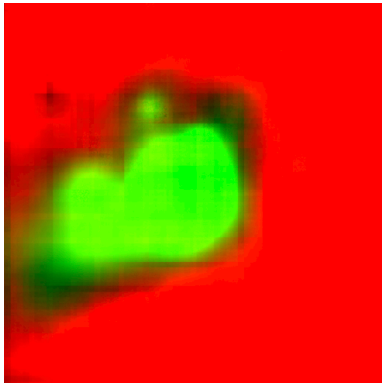

Network Heatmap

Binary Segmentation

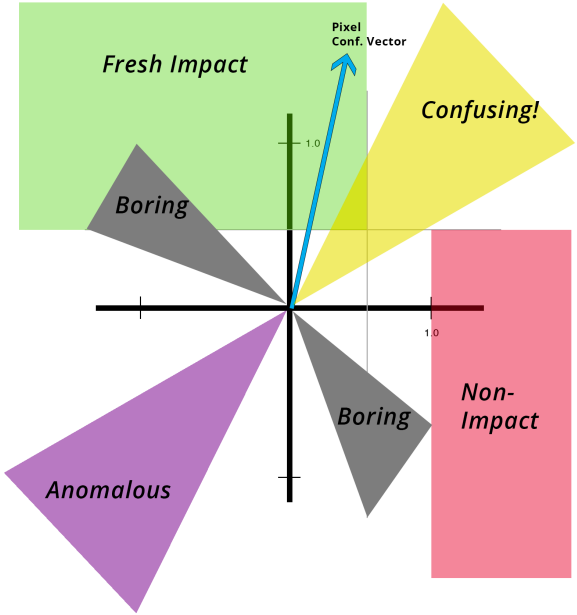

Each pixel is a 2-vector

Training with real-valued labels

Instead of pixel classification, do regression and choose **confidence thresholds**

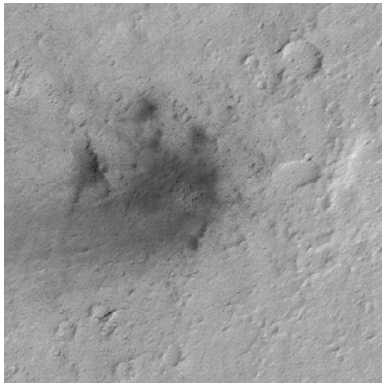

Fresh Impacts

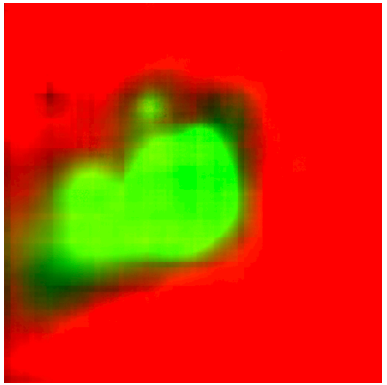

Network Heatmap

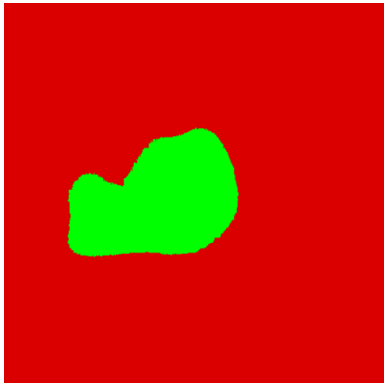

Binary Segmentation

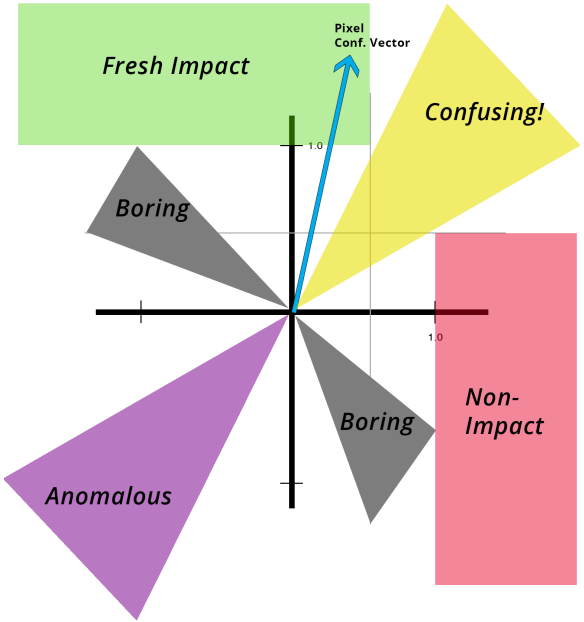

Each pixel is a 2-vector

Training with real-valued labels

Instead of pixel classification, do regression and choose **confidence thresholds**

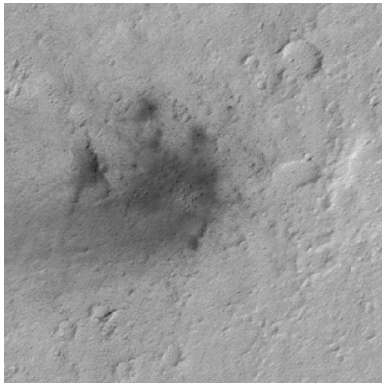

Fresh Impacts

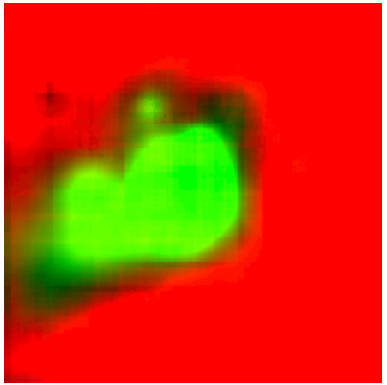

Network Heatmap

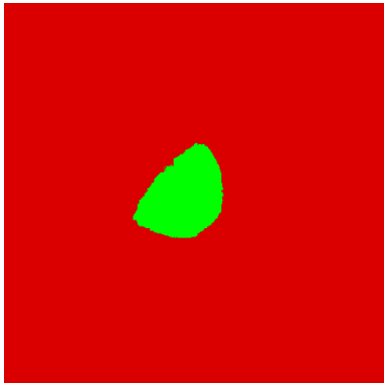

Binary Segmentation

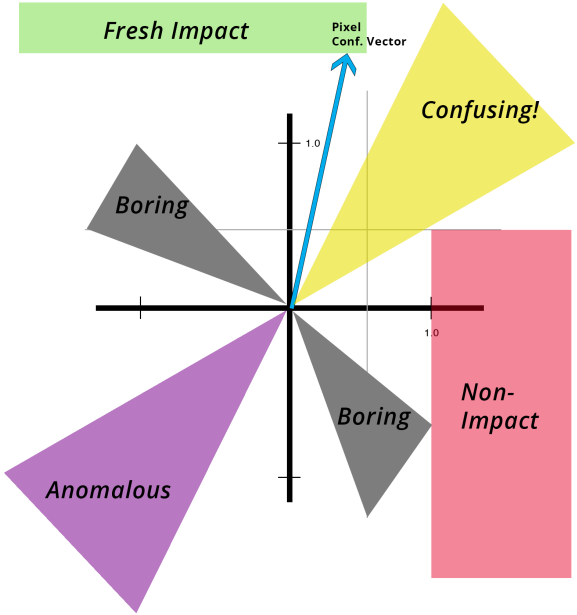

Each pixel is a 2-vector

Training with real-valued labels

Instead of pixel classification, do regression and choose **confidence thresholds**

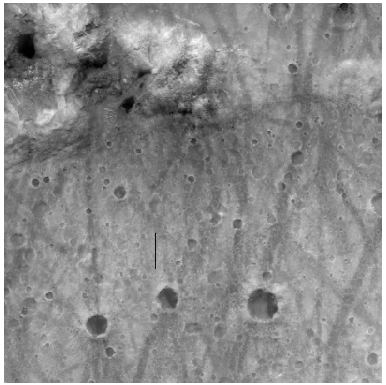

Old Impacts

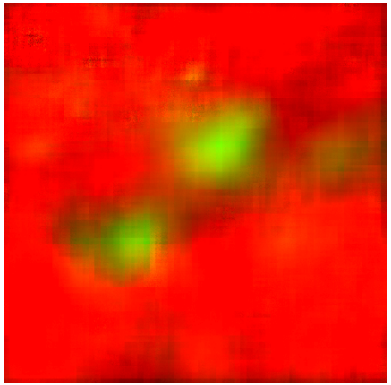

Network Heatmap

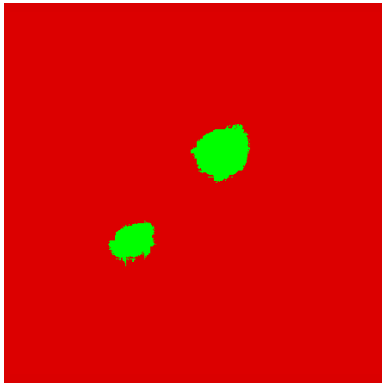

Binary Segmentation

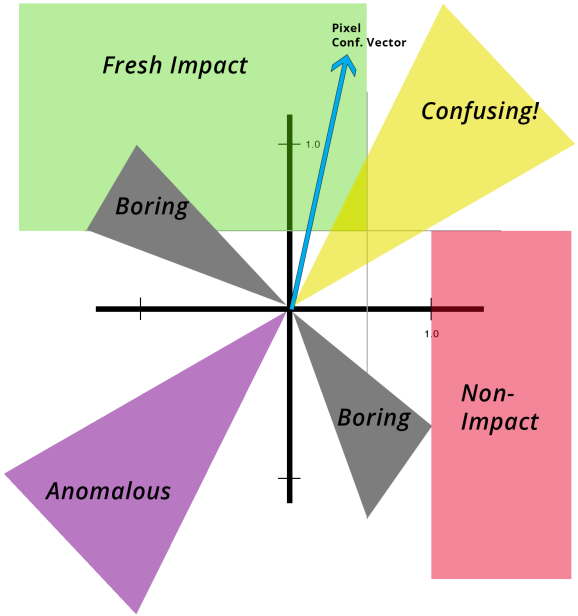

Each pixel is a 2-vector

Training with real-valued labels

Instead of pixel classification, do regression and choose **confidence thresholds**

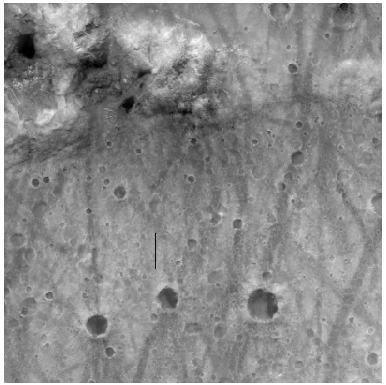

Old Impacts

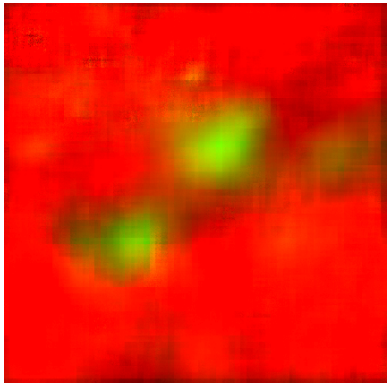

Network Heatmap

Binary Segmentation

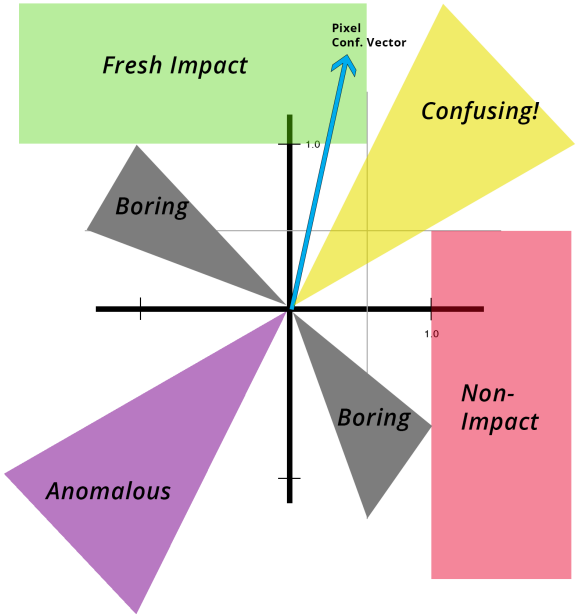

Each pixel is a 2-vector

Training with real-valued labels

Instead of pixel classification, do regression and choose **confidence thresholds**

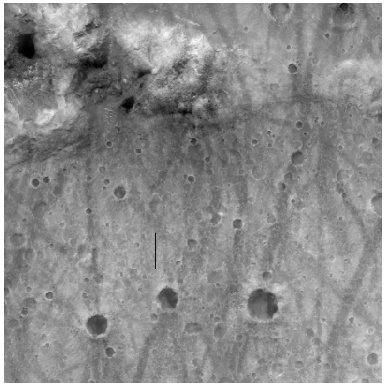

Old Impacts

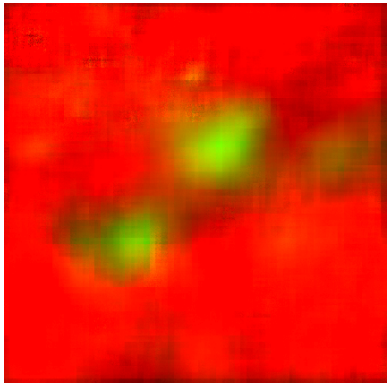

Network Heatmap

Binary Segmentation

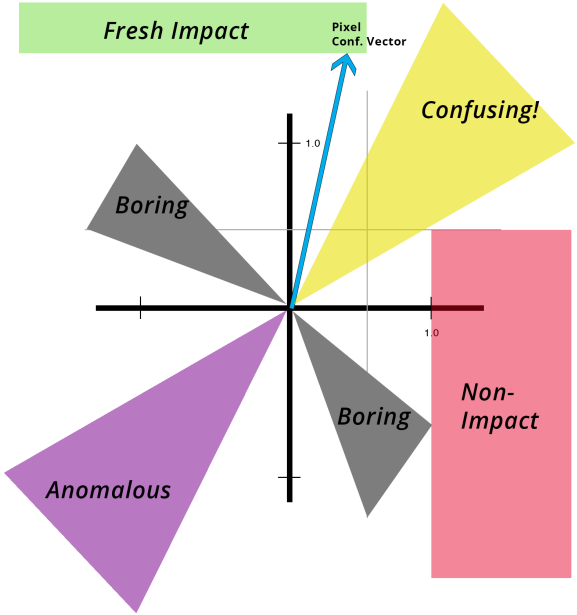

Each pixel is a 2-vector

Early difficulty: troublesome priors

Does the fully-convolutional network really know what a fresh impact looks like?

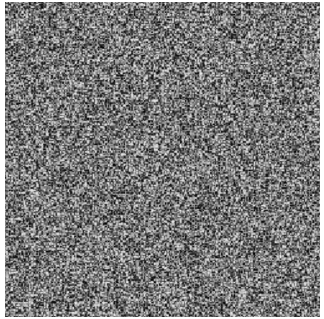

Uniform Noise

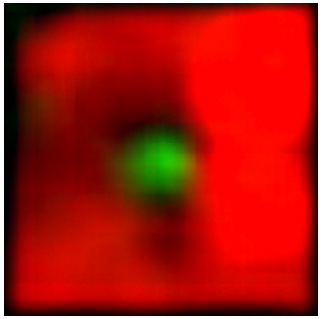

Network Heatmap

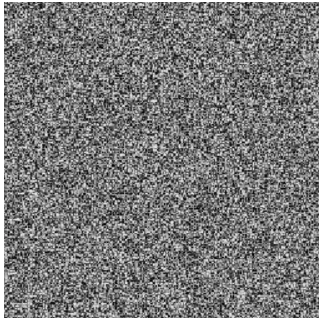

Uniform Noise

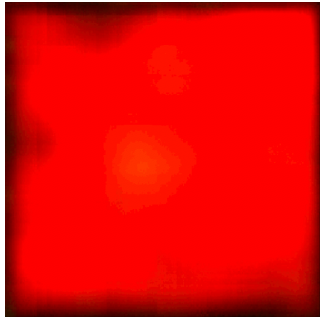

Network Heatmap

Every image the network saw had a fresh impact somewhere.

Solution: Add images *without* fresh impacts to the dataset.

Early difficulty: troublesome priors

Does the fully-convolutional network really know what a fresh impact looks like?

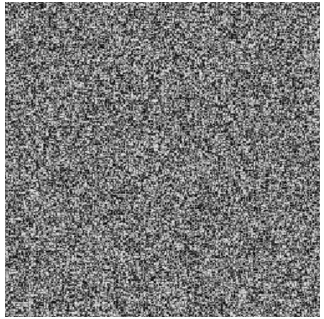

Uniform Noise

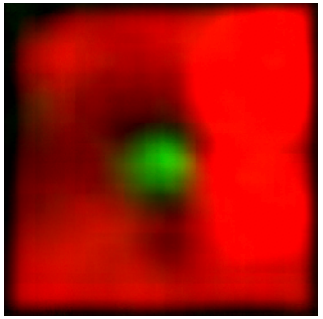

Network Heatmap

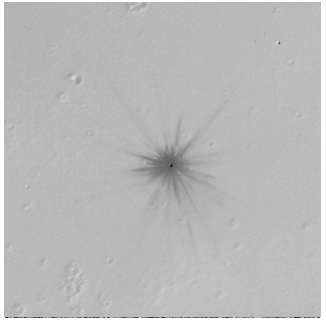

Uniform Noise

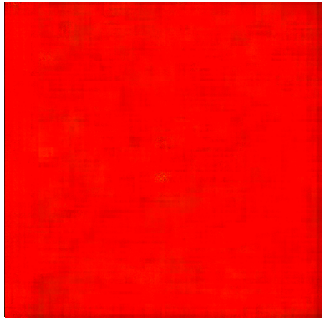

Network Heatmap

Every image the network saw had a fresh impact somewhere.

New problem: Unbalanced data biased towards non-impacts.

Real solution: Curriculum Learning (Bengio et al., 2009)

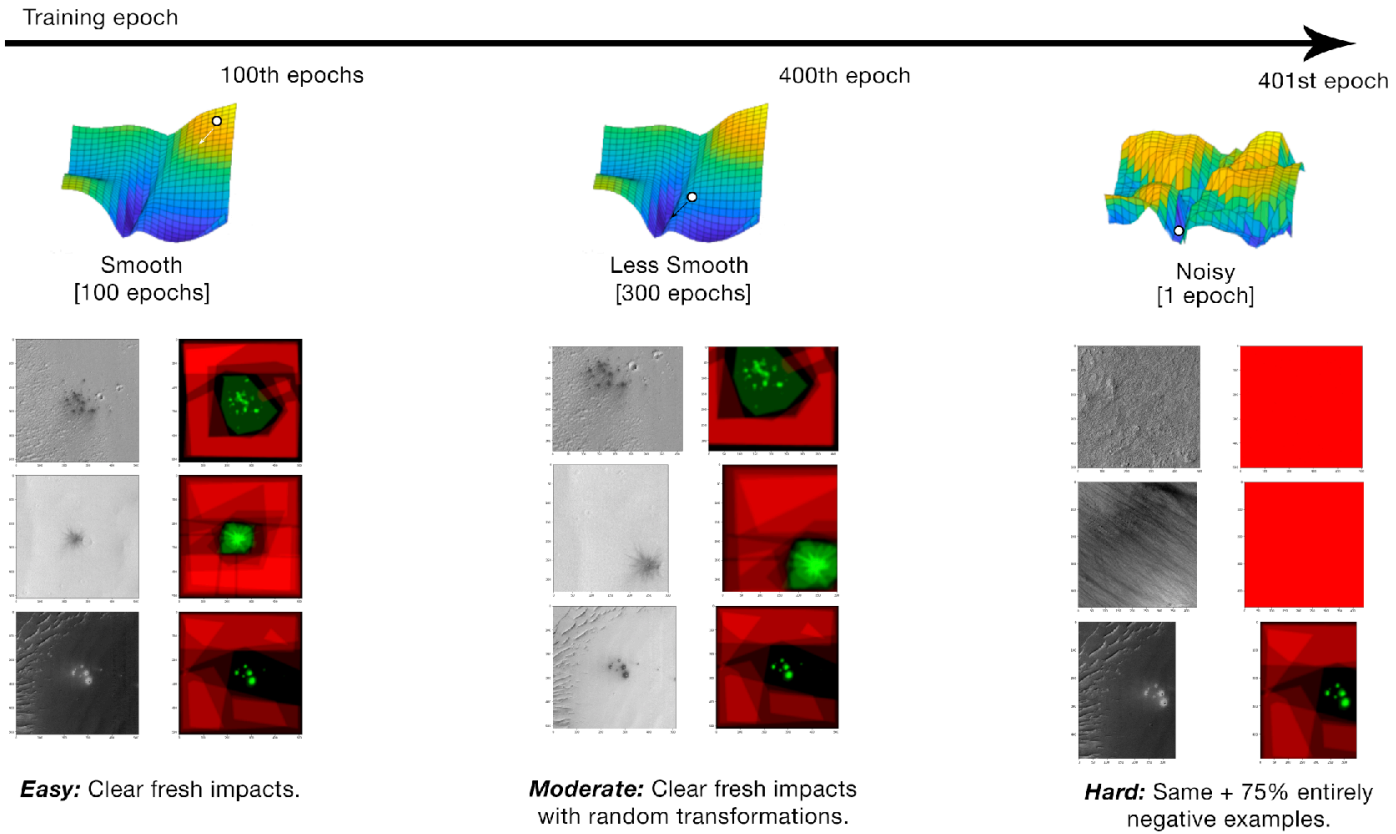

Evaluation in terms of precision and recall

Model **trained** with *real-valued* merged labels, **tested** on *discrete* merged labels

Training set: 280 fresh impact images + 1020 impact-free images
Testing set: 15 fresh impact images + 132 impact-free images

After 100 “easy” epochs → after 300 “moderate” epochs → after **1** “hard” epoch

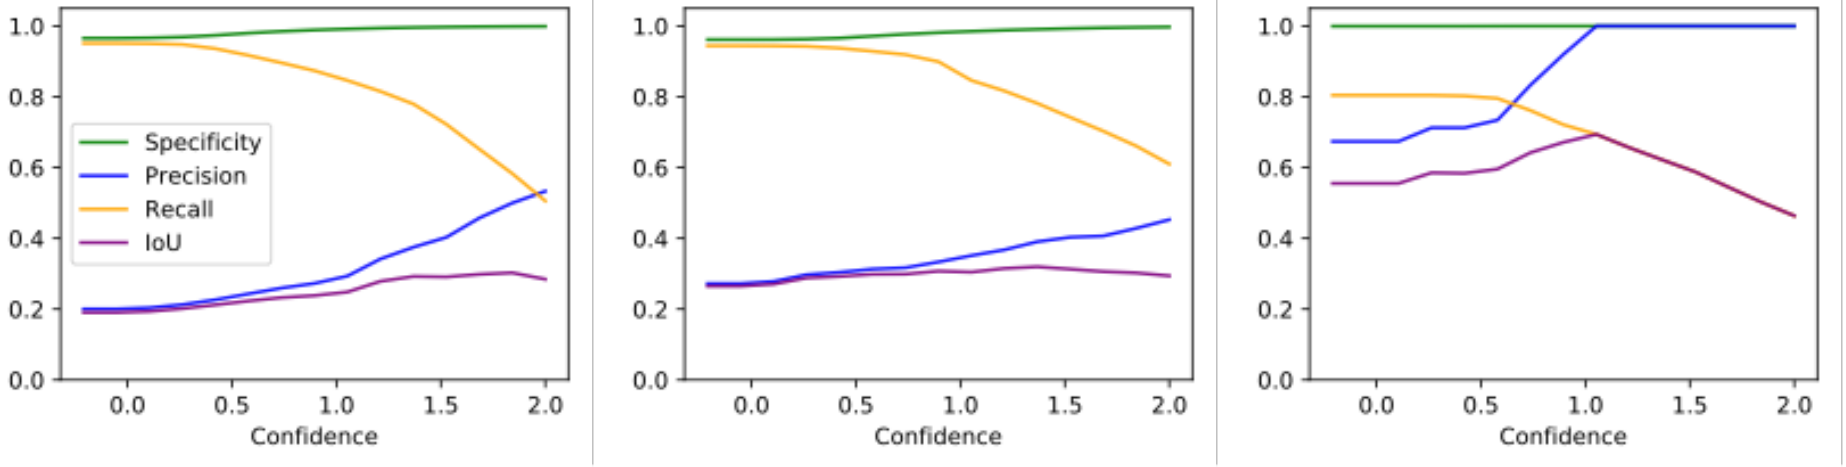

Robustness checks

1. Sensitivity to noise

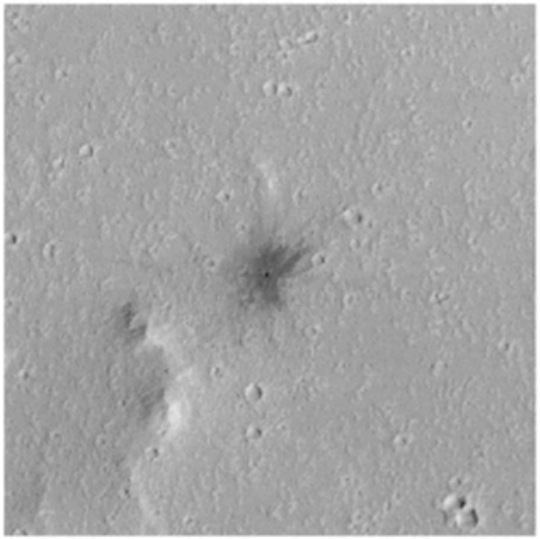

Input Image
Increasingly Strong Noise

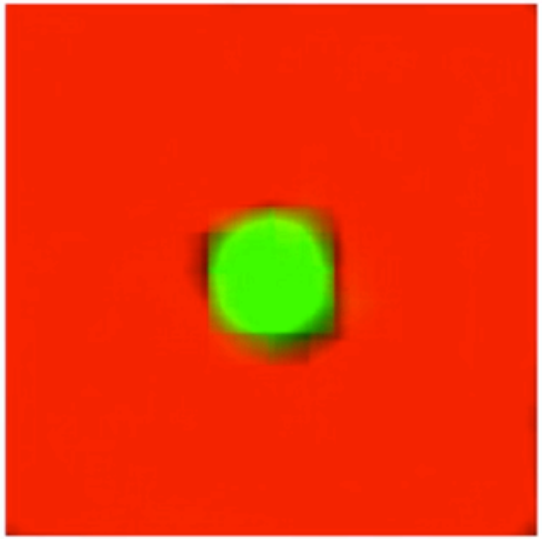

Network Heatmap

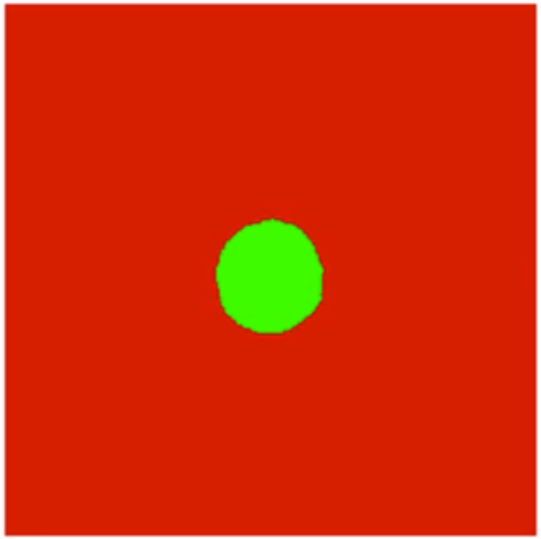

Binary Segmentation

Robustness checks

2. Sensitivity to translation

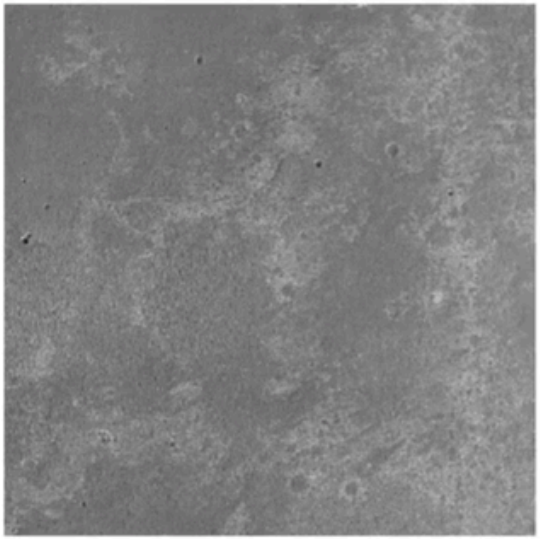

Input Image Stream

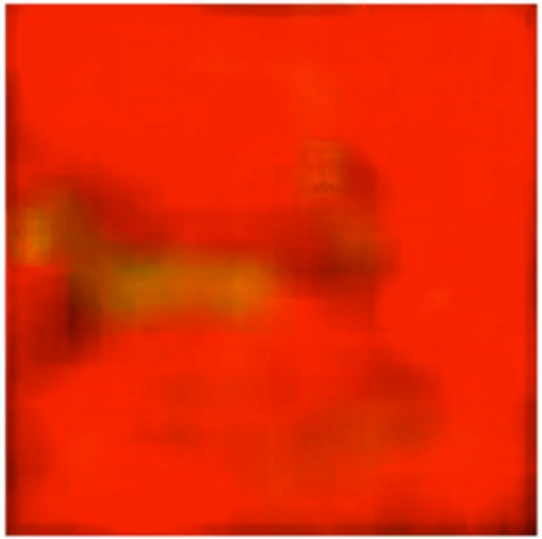

Network Heatmap

Binary Segmentation

Robustness checks

2. Sensitivity to similar landforms (araneiforms)

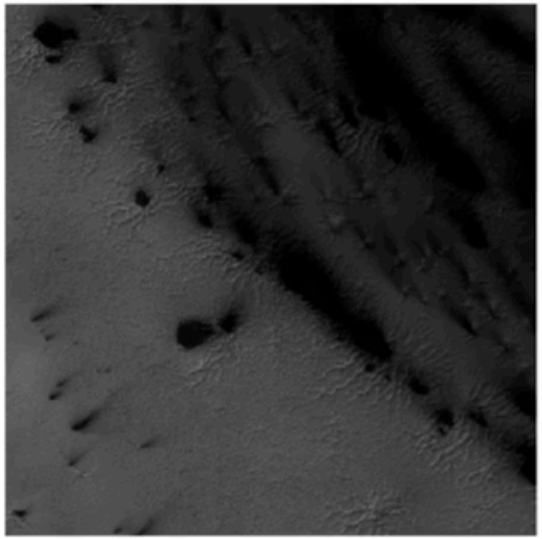

Input Image Stream

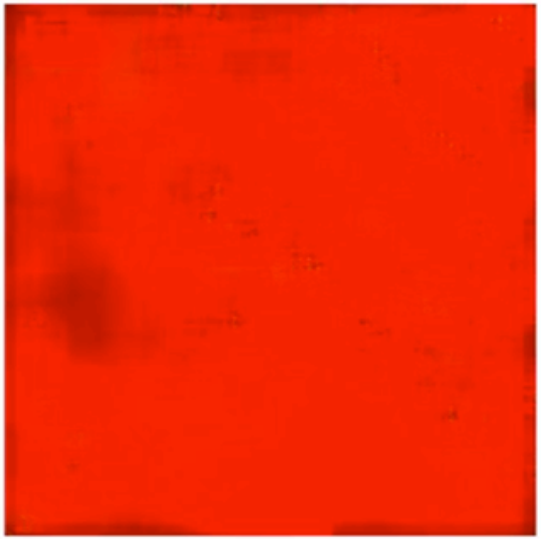

Network Heatmap

Binary Segmentation

HiRISE image exploration with the FCN

Found small, potentially fresh impacts (from 10 HiRISE browse images)

Possible fresh impacts (~5%):

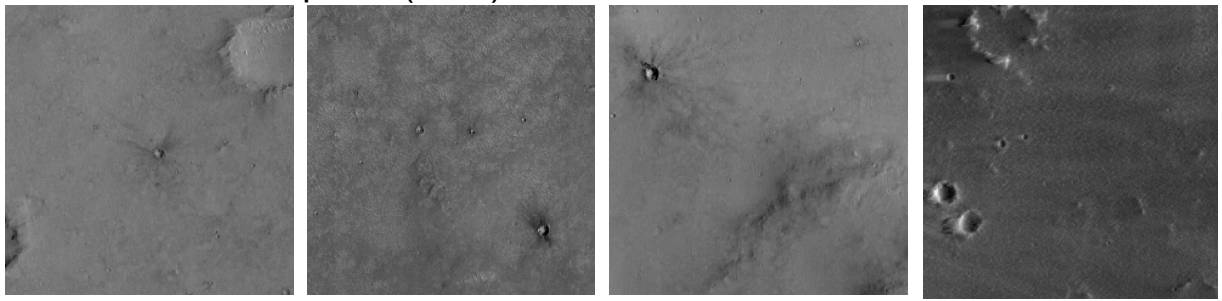

Not fresh impacts (~95%):

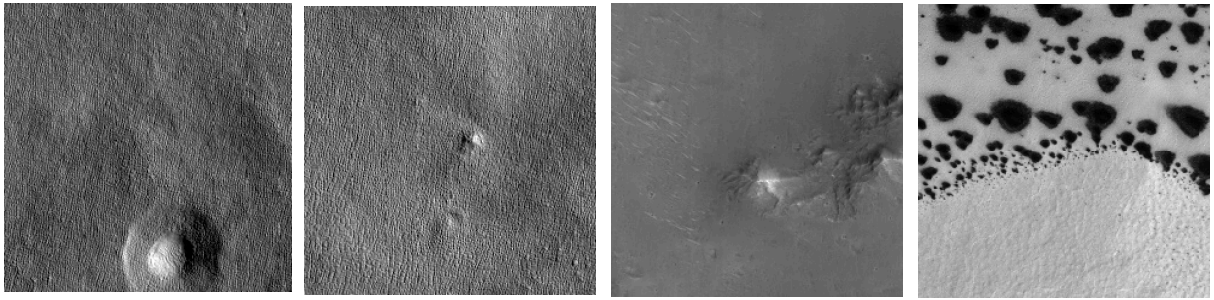

FCN selected **~5%** of considered patches.

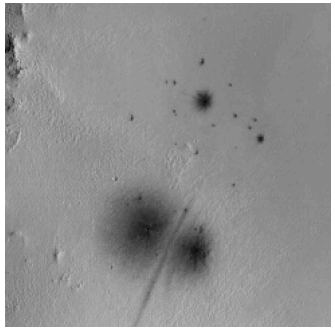

“Re-discovered” fresh impact from training set

Next steps

Working on COSMIC in the coming semester

- Cross validation (running now!)
- Comparison of other FCN architectures, e.g. AlexNet, GoogLeNet, VGG-16
  - VGG-16 already implemented
- Cleaning up code (constantly working on this)
- Model compression

Thanks!

First step of the curriculum
for cross validation (5 folds pictured)
before I ran out of memory

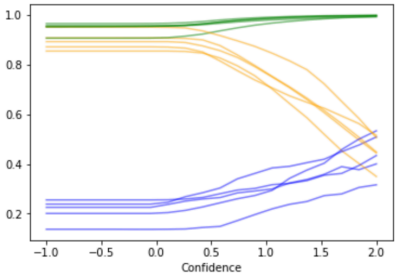


JPL
Jet Propulsion Laboratory
California Institute of Technology