Deep learning map of fresh crater ejecta on Mercury

Michele Lissoni[1], Alain Doressoundiram[1], Sebastien Besse[2]

[1]LIRA, Observatoire de Paris, CNRS, Université PSL, Sorbonne Université, Université Paris Cité, CY Cergy Paris Université, 92190 Meudon, France

[2]ESAC, European Space Agency, Madrid, Spain

Introduction: fresh crater ejecta

- Recent impact craters are usually surrounded by rays of high-albedo material extending outwards to many times the crater radius.
- Crater rays are formed by sub-surface material ejected by the primary impact and by secondary impacts (caused by fragments ejected from main crater).
- The high-albedo is usually caused by the low maturity of the material. Over time, space weathering (SpWe) causes rays to disappear. Older craters have fainter or absent rays.
- Ejecta rays are a prominent feature on the surface of Mercury.
    - Craters with visible ejecta rays may be < 300 Ma old (Banks+ 2017).
    - Rays from one crater, Hokusai, stand out: they extend over most of the planetary surface.


Very recent (2010 – 2012) impact crater on Mars (HiRISE, MRO, NASA/JPL-Caltech/Univ. of Arizona)

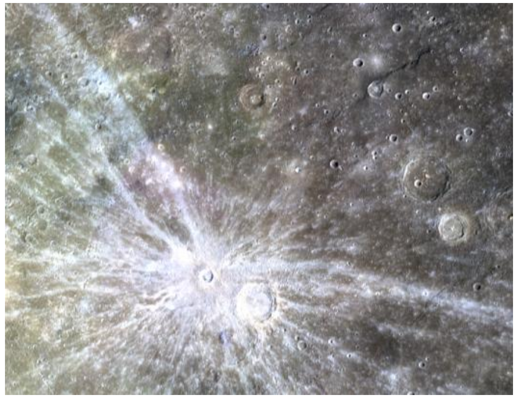
Xiao Zhao crater (Mercury): its ejecta are very bright and evident, suggesting a very young age.

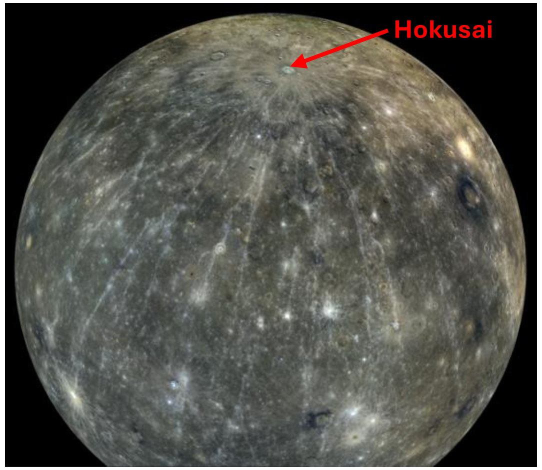

Hokusai crater rays (Mercury ACT-REACT Quickmap)

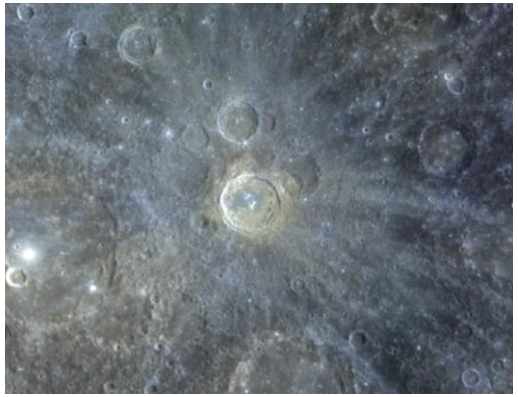
Motonobu crater (Mercury): its ejecta are visible, but faint. It is probably older than Xiao Zhao, its ejecta partially erased by SpWe.

Crater ejecta maps

- Mapping fresh crater ejecta would make it possible to:
  - systematically study ejecta spectra;
  - compare the degree of SpWe between ejecta belonging to different craters or to the same craters: information on crater age and SpWe spatial pattern;
  - reconstruct the dynamics of the impacts that created the craters.
- Bright crater ejecta have already been mapped manually in Mercury geological maps (Galluzzi+ 2016, …), but these are not yet complete and do not distinguish ejecta of different craters.
  - Distinguishing ejecta of multiple craters controls for the age variable (ejecta rays from one crater all formed at the same time) and isolates the ejecta pattern of an impact and thus its dynamics.
  - A finer spatial resolution is also desirable.

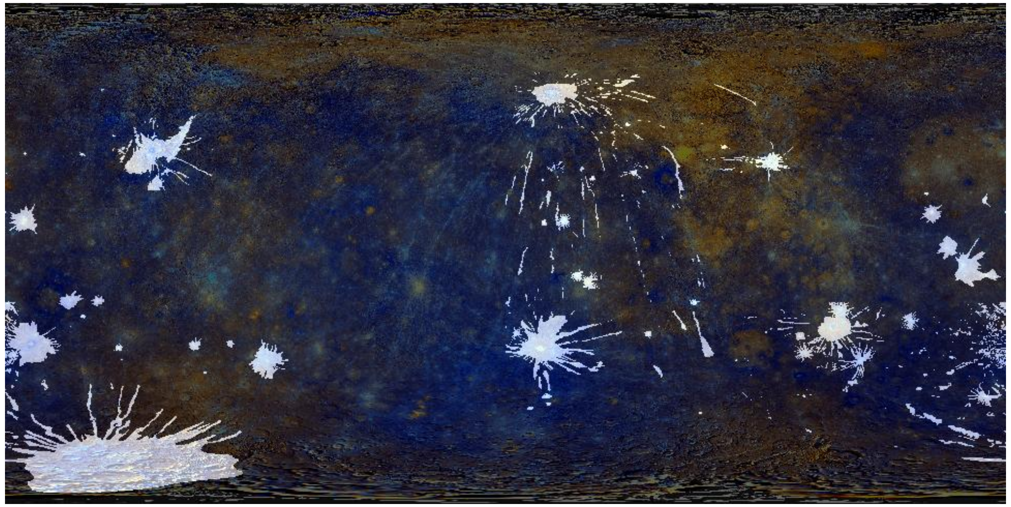

Existing ejecta map (March 2025) (Galluzzi+ 2016, …)

Deep learning approaches

- Idea: train semantic segmentation model to produce binary ejecta masks from sub-images of a Mercury global mosaic. Then merge the masks together into a global mask.
- Previous approaches:
  1. A model segmenting ejecta without distinguishing the craters of origin.
     - produced reasonably complete maps of brighter ejecta;
     - struggled with fainter ejecta and on certain terrains. Sometimes misinterpreted brightness due to viewing angle.
  2. A model trained to segment ejecta from given craters with only the sub-images and their distance from the crater as inputs.
     - no information on what was between the crater and the sub-image was provided;
     - segmented ejecta blankets in the vicinity of craters but failed with rays.
- New approach: two models
  1. Connection model: evaluates whether a sub-image is “connected” to the crater by an ejecta ray.
  2. Segmentation model: segments the sub-image.

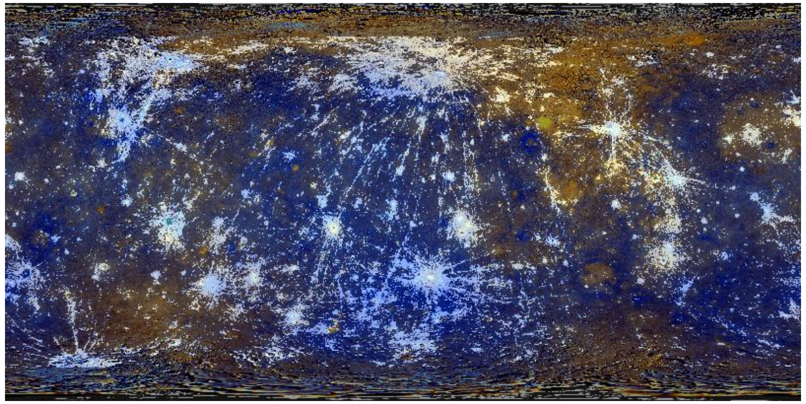

Ejecta deep learning map, no crater distinctions.

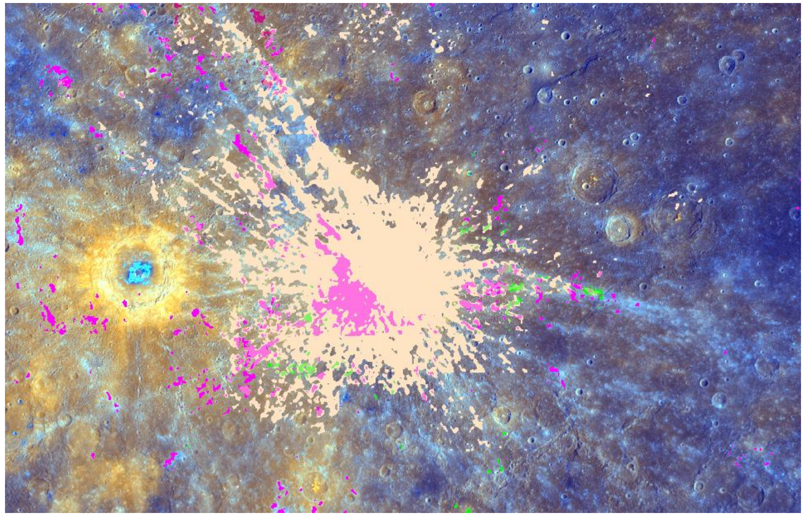

Xiao Zhao ejecta, mapped according to the 2nd approach. The ejecta rays are not recognized and there are spurious overlaps from other craters.

Training data

- 119 fresh craters
- 215 sub-images of the Enhanced Color mosaic, manually segmented in Equidistant Cylindrical projection.
  - Sub-image size: 224x224 px. Spatial resolution: 665 m.
  - Separate ejecta binary masks for each crater.

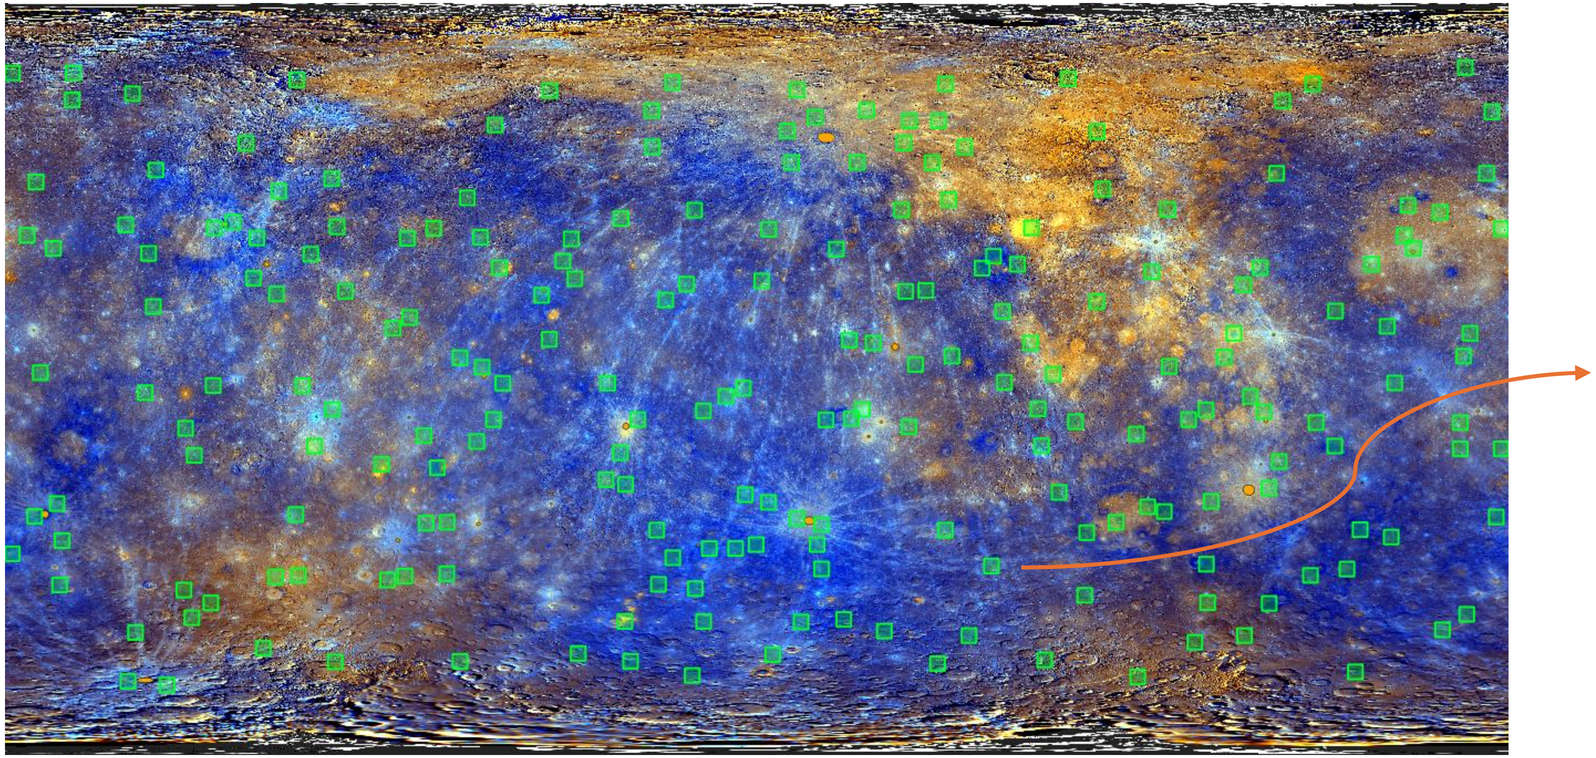

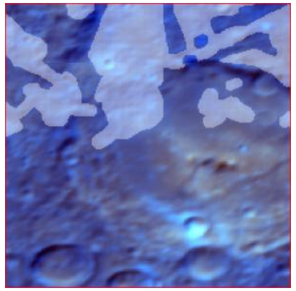

Sub-image position 215
Crater 12 (Debussy)

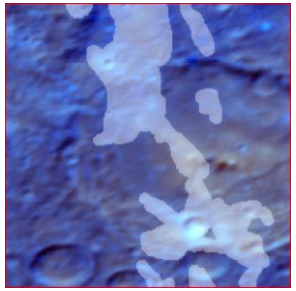

Sub-image position 215
Crater 0 (Hokusai)

Training data

- For each crater, mosaic and mask are reprojected into an Azimuthal Equidistant projection centered about the crater.
  - For each cylindrical sub-image, the largest square contained within its reprojected boundaries (zoomed to 224x224 px) is used to train the segmentation model.
  - The area between the crater and the azimuthal sub-image (4:1 rectangle, reprojected to 448x112 px) is used to train the connection model.

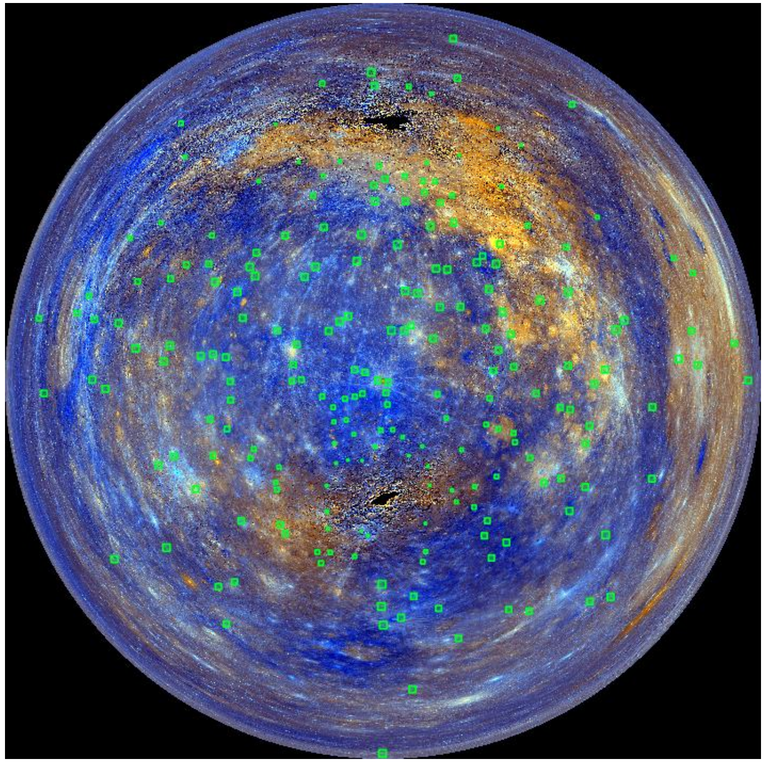

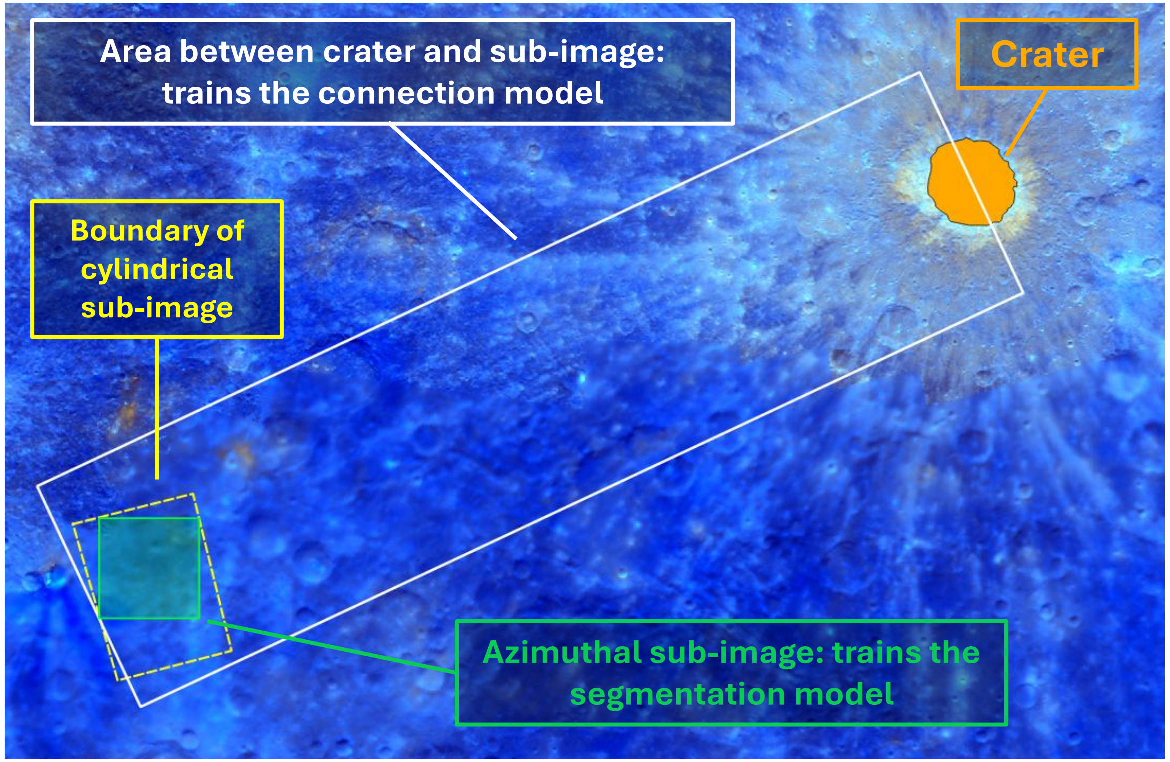

Position 93
Crater 12 (Debussy)

Connection model

- Input: 4-channel 448x112 px image.
    - 3 bands of the Enhanced Color mosaic, showing the area between the crater and the sub-image.
    - 1 channel showing the positions of the crater and of the sub-image.
- Output: binary value, 0 (not connected) or 1 (connected).
    - Connected: the sub-image contains some ejecta from that crater.
- Trained, for every crater, on every sub-image (connected or not) within 0.25 x $\pi R_{Mercury}$ from the crater (3868 samples).
    - 0.5 for Debussy, Degas, Copley, Han Kan; 0.75 for Hokusai.
    - Data split into 20% test, 8% validation (for early stopping), 72% training. Stratified random split according to distance from crater and connectedness.
- Predicts, out of a set of sub-images within the same distances, which ones are connected to the crater.
    - The model is trained 10 times over the same data set (fixed train-test split, different train-val split). The connected sub-images are those classified as 1 in at least half the iterations. This improves robustness but lengthens training.

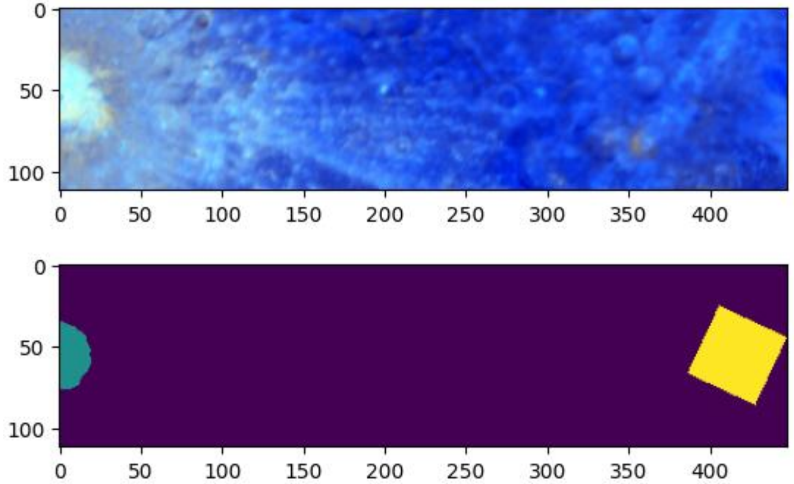

Sub-image position 93
Crater 12 (Debussy)

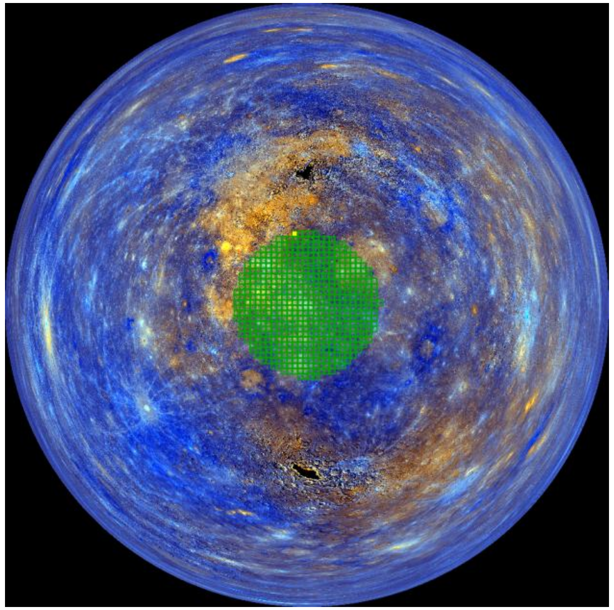
Prediction sub-images
Crater 34 (Xiao Zhao)

Connection model: hyperparameters

- Architecture: res2net50_26w_4s
  - Pretrained on ImageNet-1k
  - Any other classification model from the *timm* library could alternatively be implemented; no customization.
- Loss: Focal (alpha = 0.1)
  - Alternatives: Dice, Binary Cross-Entropy, Jaccard.
- Optimizer: AdamW
  - Learning rate: 0.0001; scheduler StepLR (step size: 10; gamma: 0.9)
- Early stopping patience: 20
- Batch size: 64
  - Choice influenced primarily by dataset size
- Sample weights: none
  - Weighting the samples can influence significantly the results.
  - Attempted weights:
    - Weights for connected and non-connected sub-images inversely proportional to class size.
    - Larger weights for more distant sub-images, to improve performance over distant rays.
  - No apparent benefit from weighting.
- Scaling: images divided by 255 (maximum possible pixel value).
  - No apparent benefit from standardization based on training data.

Segmentation model

- Input:
  - 4-channel 224x224 px sub-image.
    - 3 bands of the Enhanced Color mosaic.
    - 1 channel containing the log-distance (in m) of each pixel from the crater center.
  - Metadata: mean values of the 3 bands in an annulus around the crater (2 to 3 crater radii), log crater radius (in m).
- Output: ejecta binary mask.
  - When assembling the predicted sub-masks into the planet-wide mask, the model predicts probabilities between 0 and 1. Where the sub-masks overlap, the probabilities are averaged. The mask is set to 1 when the probability is > 0.5.
- Trained only on the sub-images containing at least some ejecta from craters (261 samples).
  - Same train-test split as the connection model.
  - Training data split into 10% validation (for early stopping), 90% training. Stratified random split according to distance from crater.
- Segments the sub-images that the connection model has classified as connected to the crater.

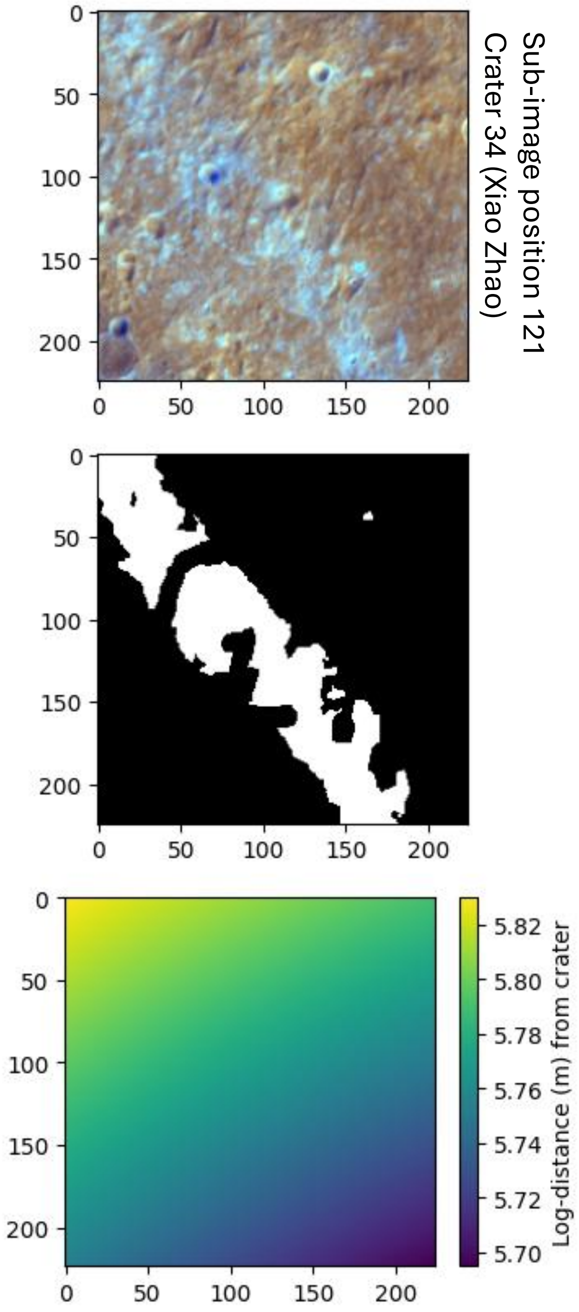

Segmentation model: hyperparameters

- Architecture: Unet++ encoder and decoder
  - Pretrained on ImageNet.
  - No customization on encoder and decoder; metadata injected into bottleneck and skip connections through fully-connected layer.
  - Encoders and decoders from the *smp* library could alternatively be implemented.
- Loss: Dice
  - Alternatives: Focal, Binary Cross-Entropy, Jaccard.
- Optimizer: AdamW
  - Learning rate: 0.0001; scheduler StepLR (step size: 10; gamma: 0.9)
- Early stopping patience: 20
- Batch size: 4
  - Choice influenced primarily by dataset size
- Sample weights: larger weights on sub-images where the fraction of ejecta from the crater in question wrt the total ejecta is small.
  - Done in an attempt to boost the model's ability to tell apart ejecta from different craters.
- Scaling: images and annulus values divided by 255 (maximum possible pixel value), log-distances and log-radius unaltered.
  - No apparent benefit from standardization based on training data.

Crater subsets

- Should the model be trained and used on all the craters, or would it be better to divide them into families?
- The number of connected sub-images, according to the connection model, is a rough estimate of the ejecta area. Plotting it vs the crater radius reveals some possible families.

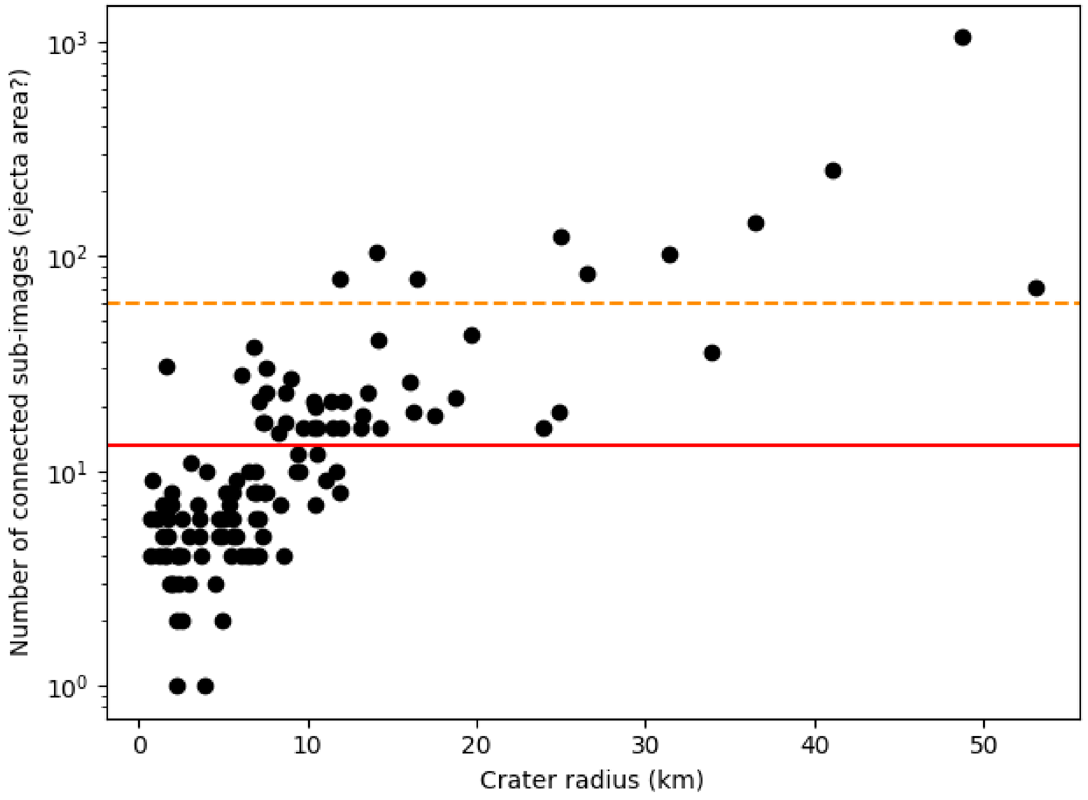

Threshold: 60

Worse results

<u>Small ejecta</u>
Craters: 109

<u>Large ejecta</u>
Craters: 10

Threshold: 13

Slight improvement: <u>**BEST MAP**</u>

<u>Small ejecta</u>
Craters: 75
Training samples: 2186
Connected training samples: 59

<u>Large ejecta</u>
Craters: 44
Training samples: 1682
Connected training samples: 202

Best map

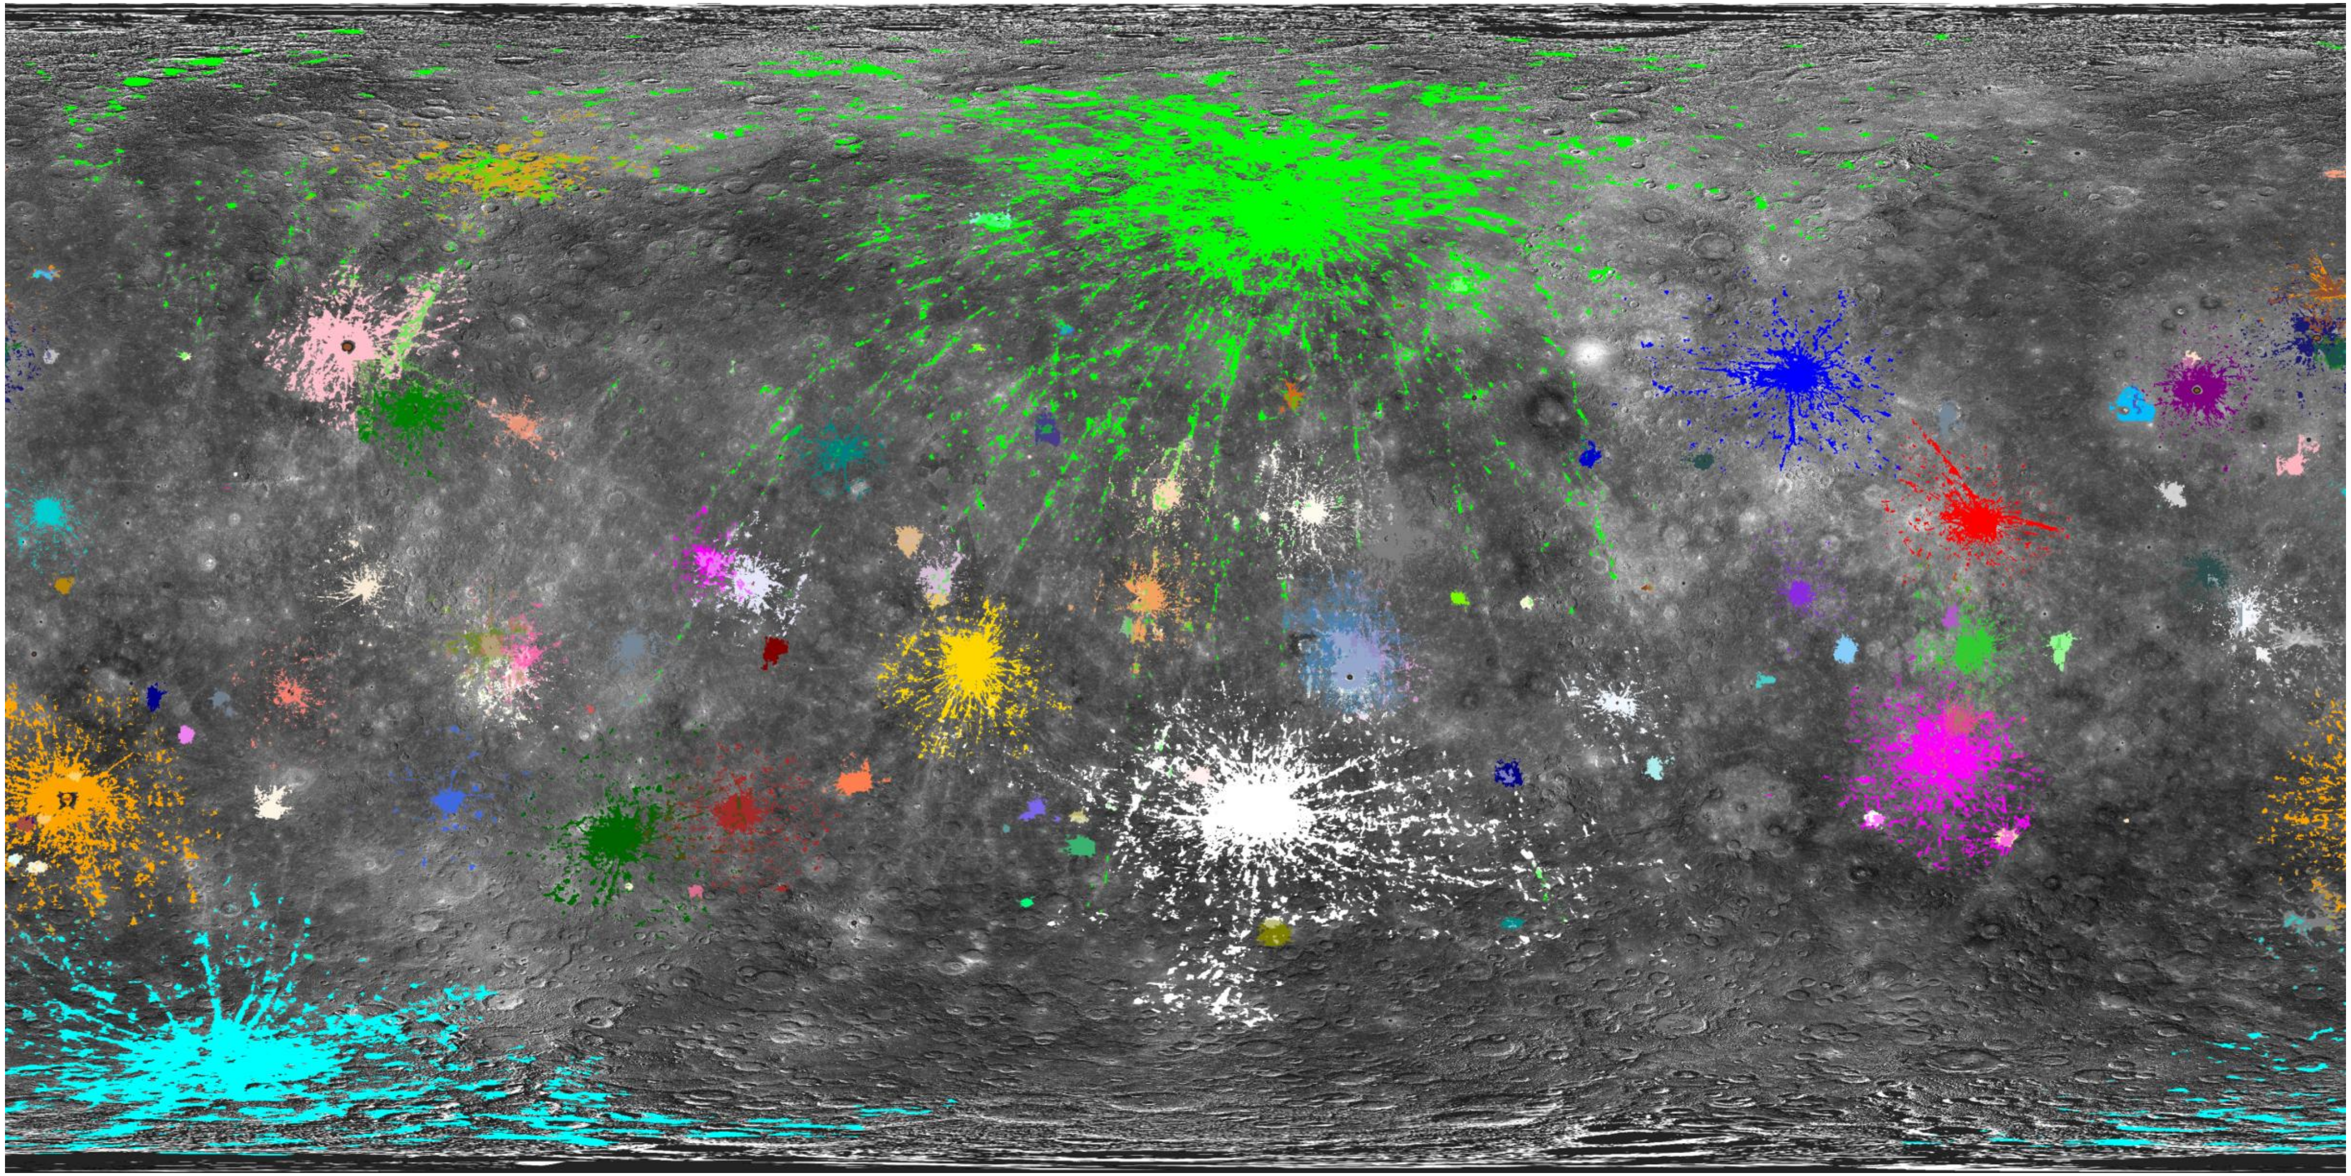

Best map

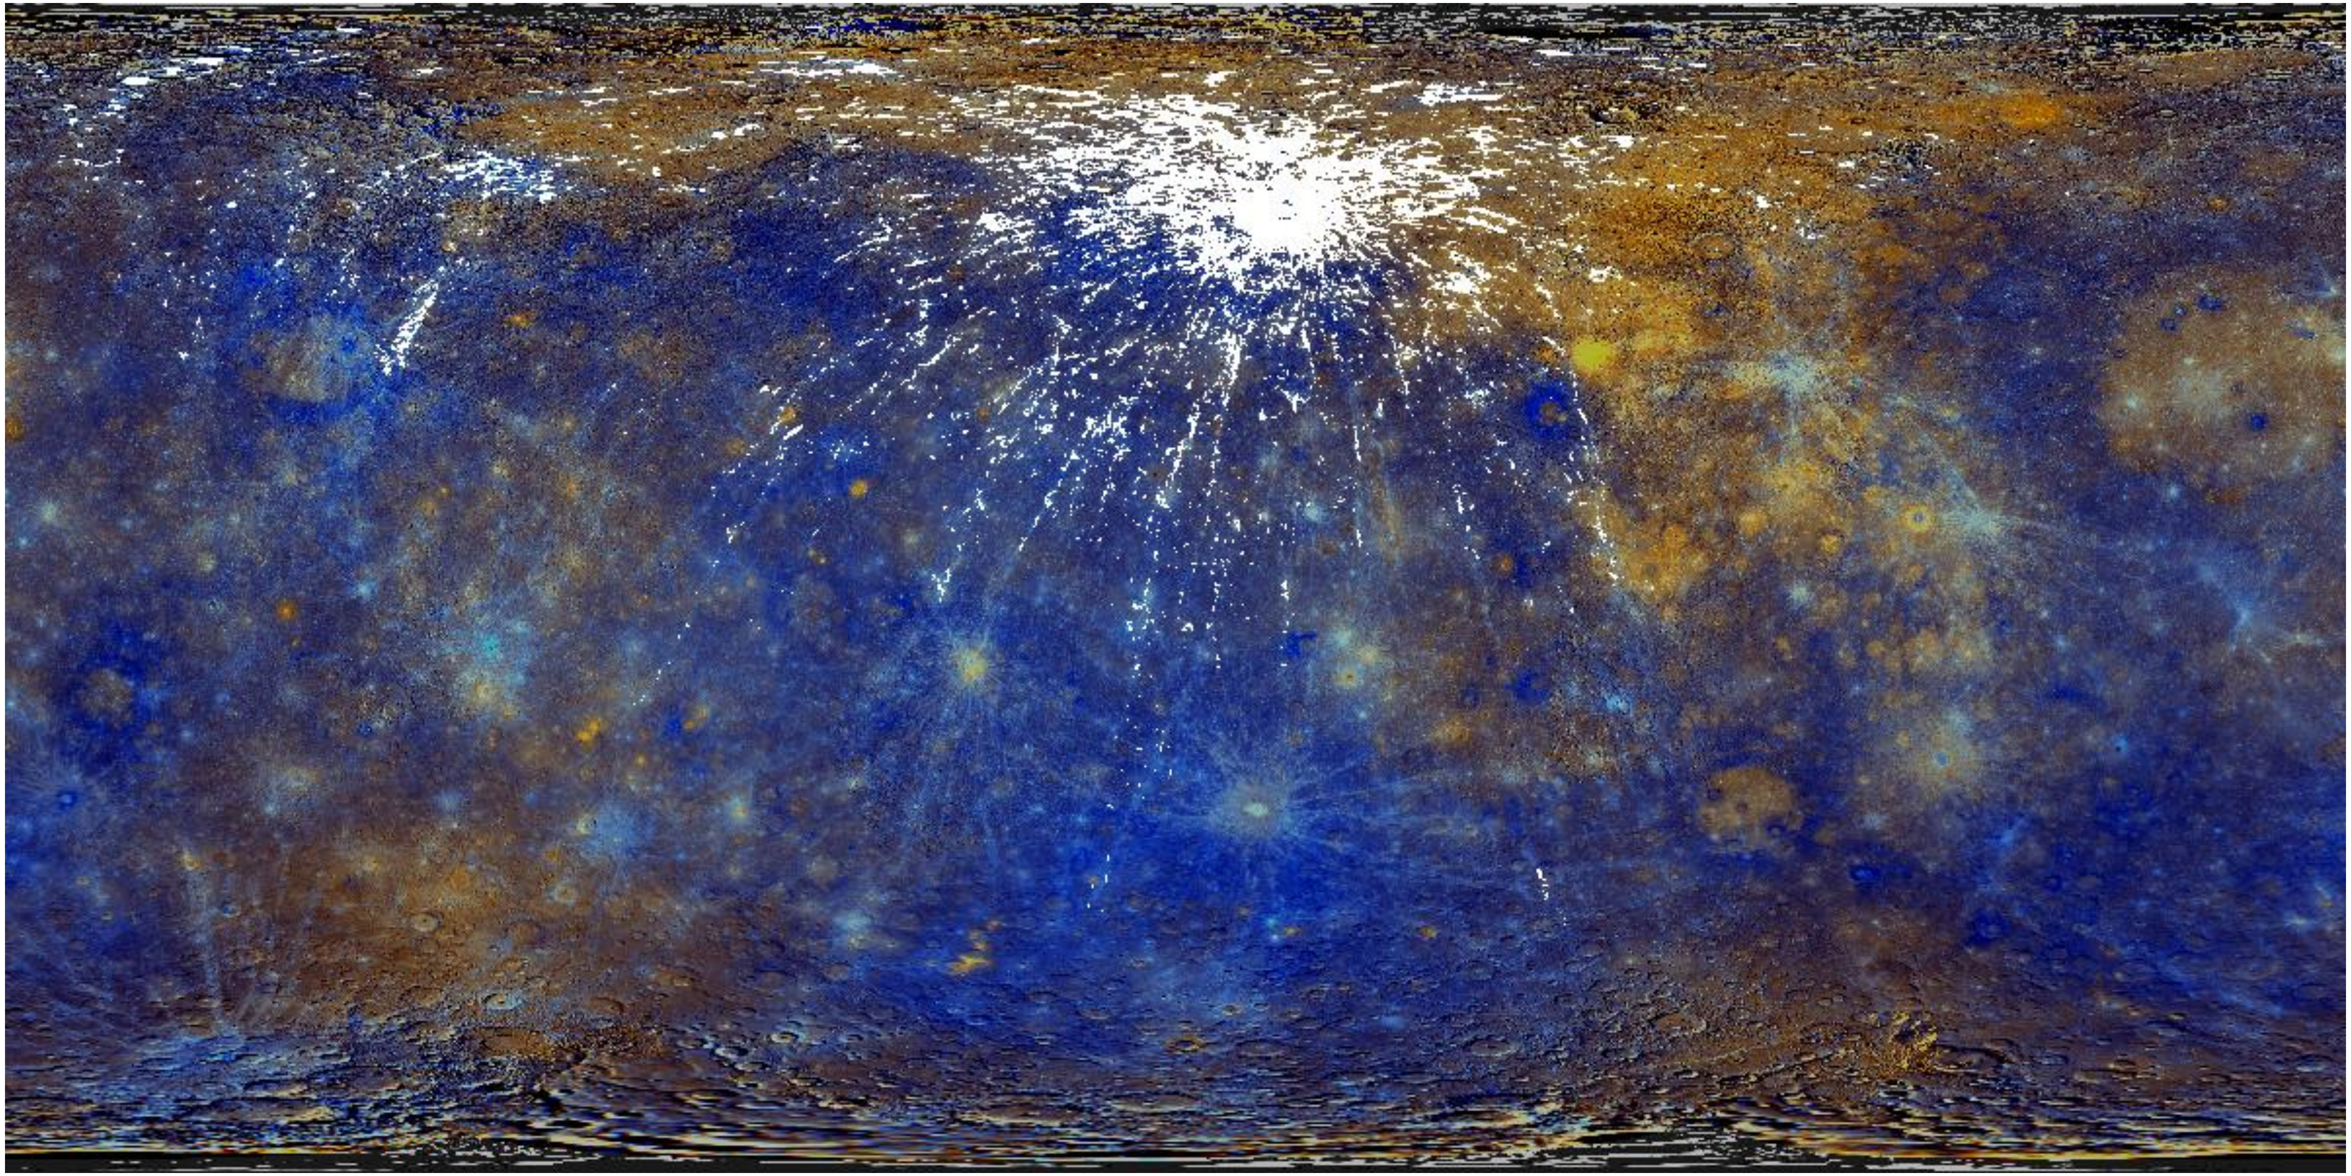

Crater 0 (Hokusai)

Best map

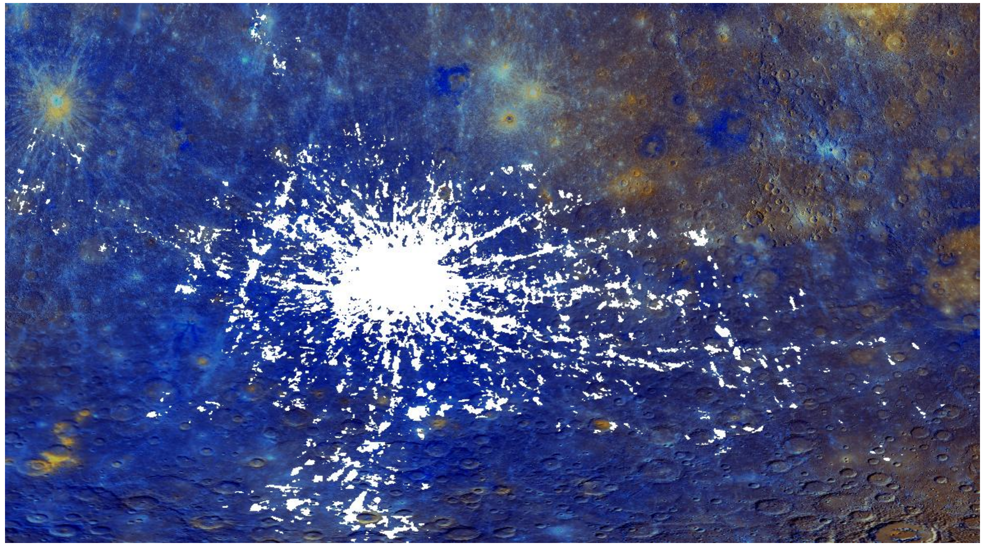

Crater 12 (Debussy)

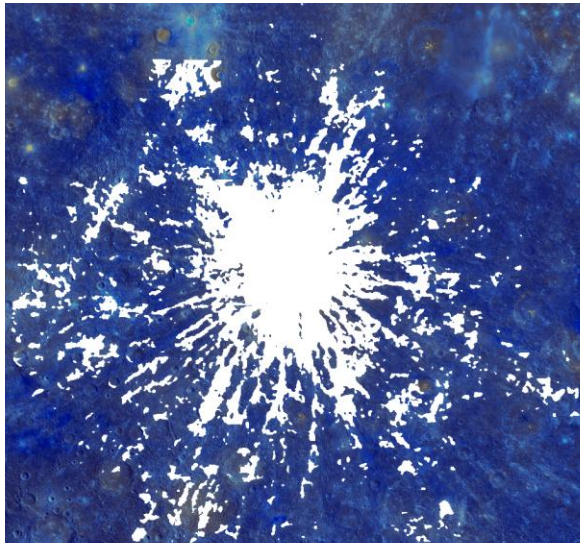

Crater 17 (Kuiper)

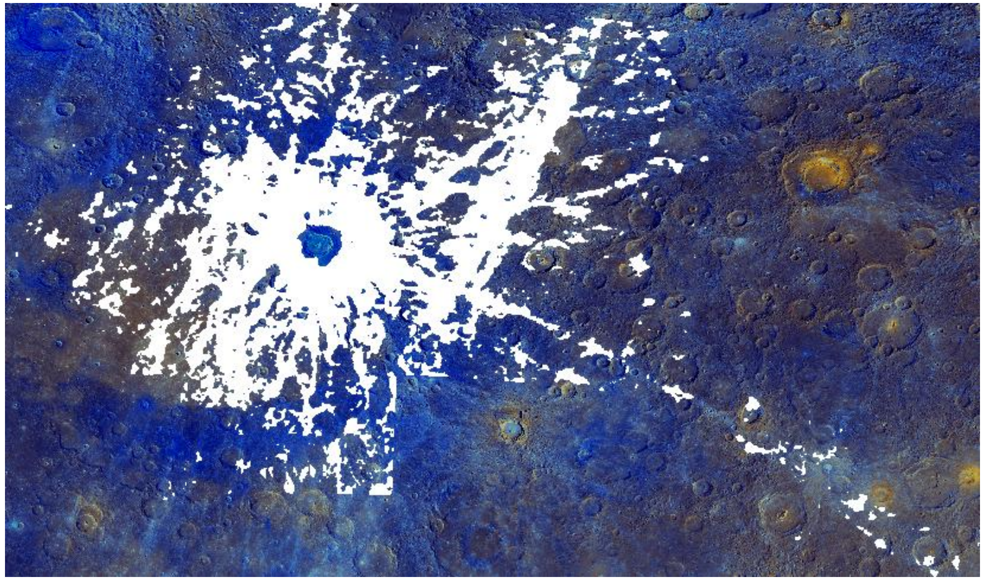

Crater 13 (Degas)

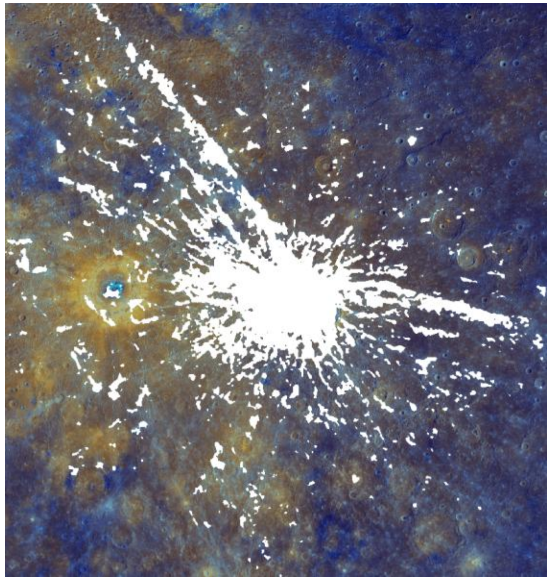

Crater 34 (Xiao Zhao)

Compared to the approach without the connection model, correct segmentation of rays further away from craters has significantly improved.

Flaws

- Model performs badly where ejecta from multiple craters overlap and sometimes confuses ejecta from different craters.
  - The model cannot always trace good boundaries when ejecta from multiple craters overlap (some of these boundaries are very difficult for a human too) or are close to each other.
  - Ejecta from craters close to a large crater are often assumed to belong to the large crater.
  - Rays from large craters are sometimes confused as ejecta from smaller craters.

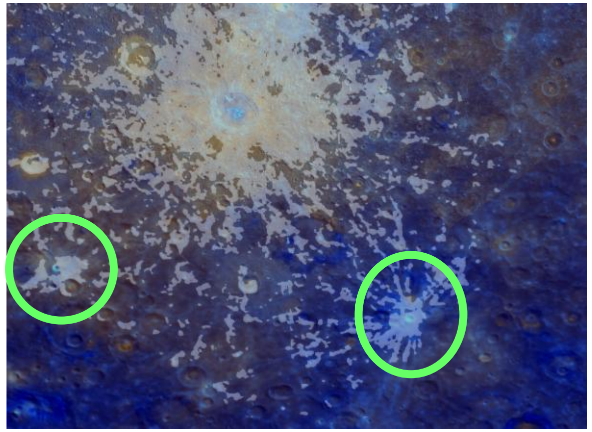

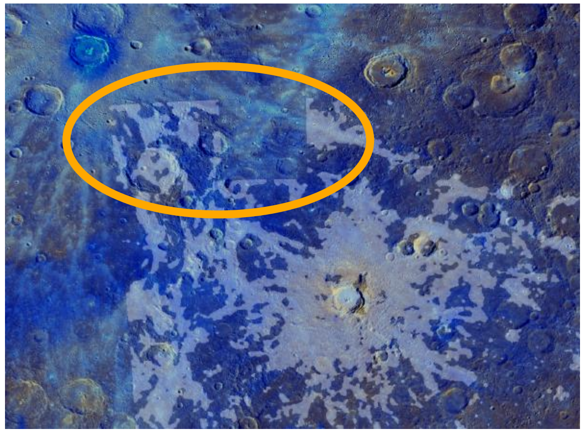

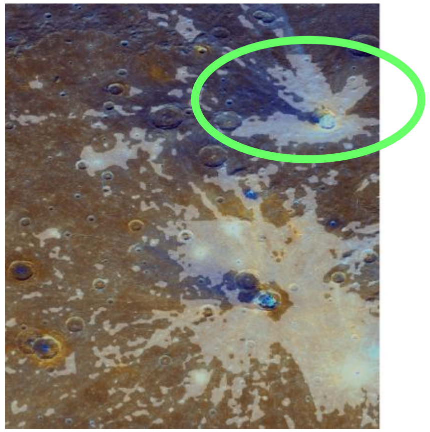

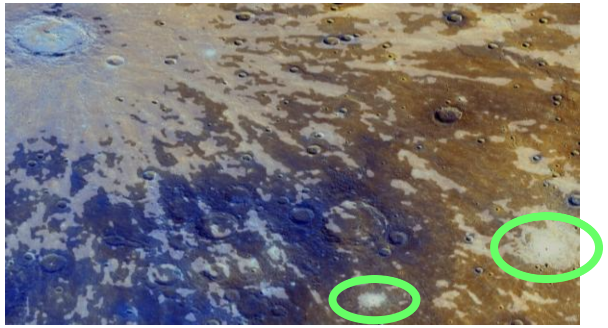

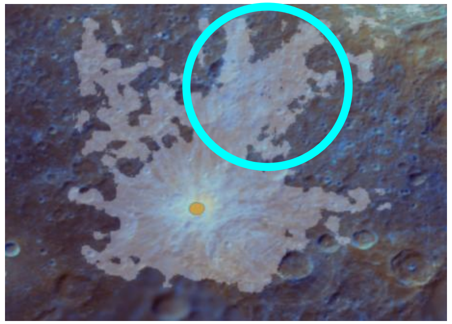

Flaws

- Ability to track rays far from the crater has improved, but can still sometimes fail.
  - Rays far from the crater may be neglected or covered in a fragmented way.
  - Rays from smaller craters are often neglected or they are not segmented finely enough.

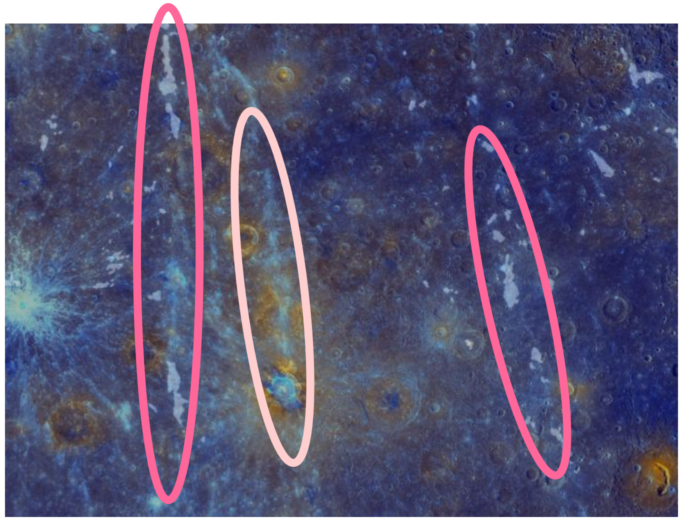

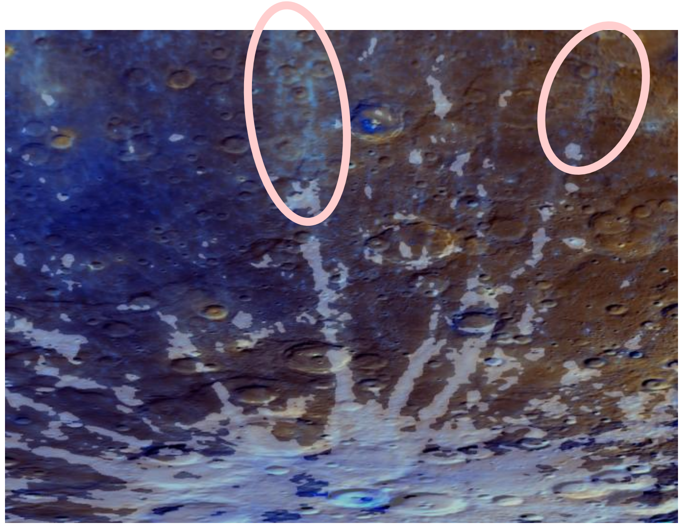

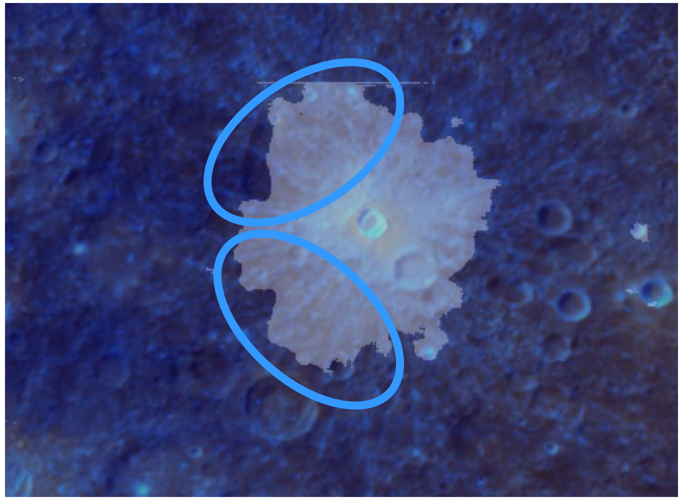

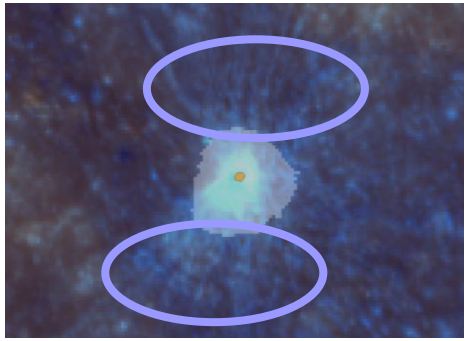

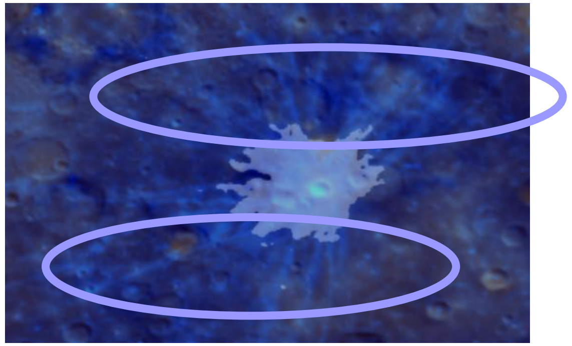

Evaluation

- We can evaluate the models by comparing, on the global masks for each crater, the sub-images that were manually segmented.
  - Comparison in cylindrical projection (since the manual segmentation was done in this projection).
  - Segmentation model: comparison per pixel. Is each pixel in the mask covered in ejecta in the manual segmentation? And in the neural network segmentation?
  - Connection model: comparison per sub-image. Does the sub-image contain manually-segmented ejecta from a given crater? What about ejecta mapped by the model?

Train & Test set

	Segmentation	Connection
F1 Score	0.691	0.855
Precision	0.672	0.819
Recall	0.711	0.894
IoU	0.528	0.747

Test set only

	Segmentation	Connection
F1 Score	0.638	0.724
Precision	0.622	0.731
Recall	0.656	0.717
IoU	0.469	0.567

We're more interested in the quality of the global map than in the model's ability to generalize (the major craters are all in the training data anyway), so computing metrics over the entire dataset (left table), more representative of Mercury ejecta, may be more appropriate.

Evaluation

- If isolating a test set is not important, we can use all the data for training.

All data used for training and evaluation

	Segmentation	Connection
F1 Score	0.691	0.902
Precision	0.662	0.895
Recall	0.723	0.909
IoU	0.528	0.822

- Segmentation performance is almost the same, improvement in connection model performance.

Optimization and evaluation problem

- “Best” map was found through visual, qualitative examination.
  - Not replicable.
  - Not useful for hyperparameter optimization. We cannot evaluate visually the map produced by every single combination of hyperparameters.
- Do the quantitative metrics described earlier properly represent the quality of an ejecta map?
  - With previous approach (ejecta maps regardless of crater of origin) we tried to compare quantitative metrics with quality rankings by the authors. No correlation.
  - Optimization is only feasible with a good quantitative metric to evaluate each hyperparameter combination.
- Do better evaluation metrics exist?
  - Structure-aware metrics? Hausdorff distance, boundary F1-score…