RMI Analyst

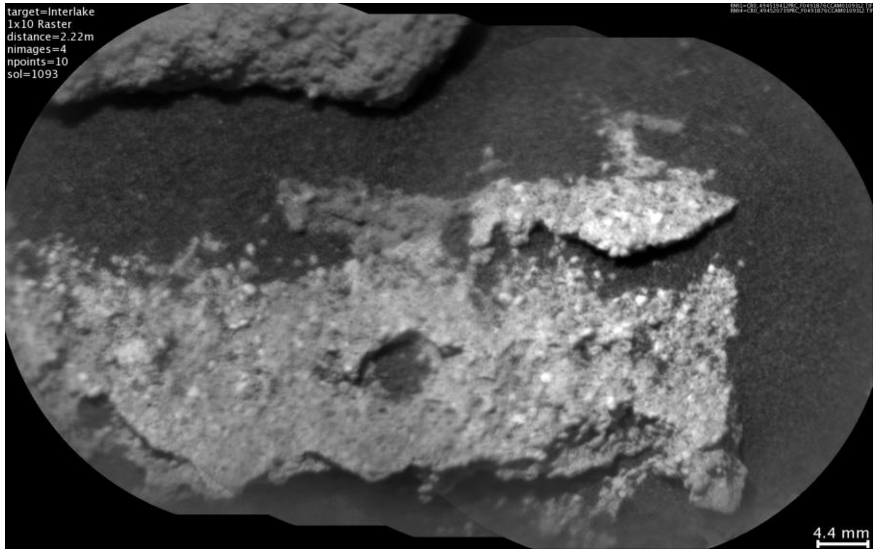

F. Arrignon, W. Rapin, O. Gasnault

IA@IRAP, 15/10/2025

ESA/MPS/UPD/LAM/IAA/RSSD/INTA/UPM/DASP/IDA

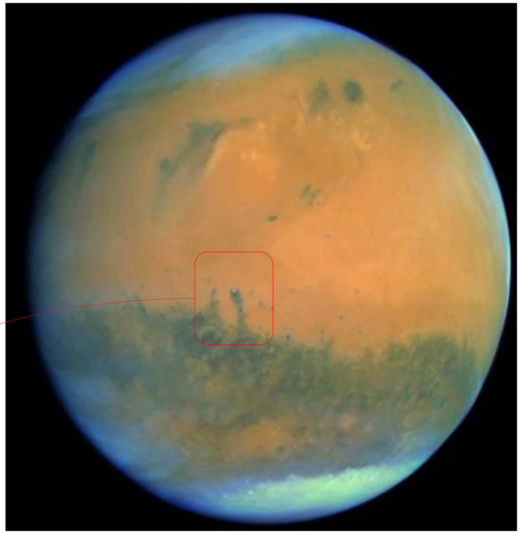

Gale crater, **Mars**
Curiosity rover

ChemCam

and its Remote Micro-Imager (RMI) camera

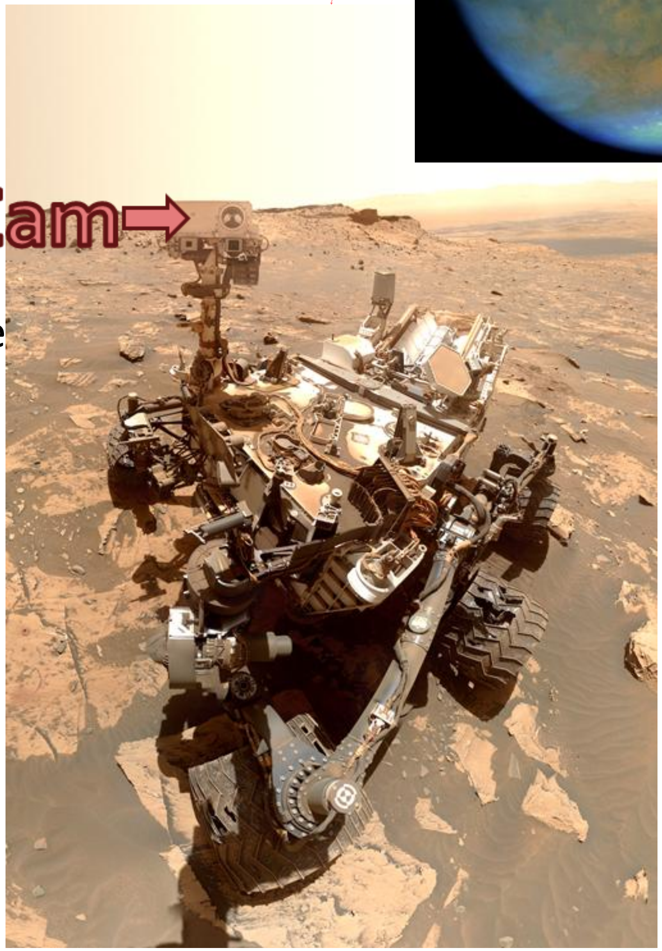

RMI mosaic of images

NASA/JPL-Caltech/
LANL/CNES/CNRS/ IRAP/IAS

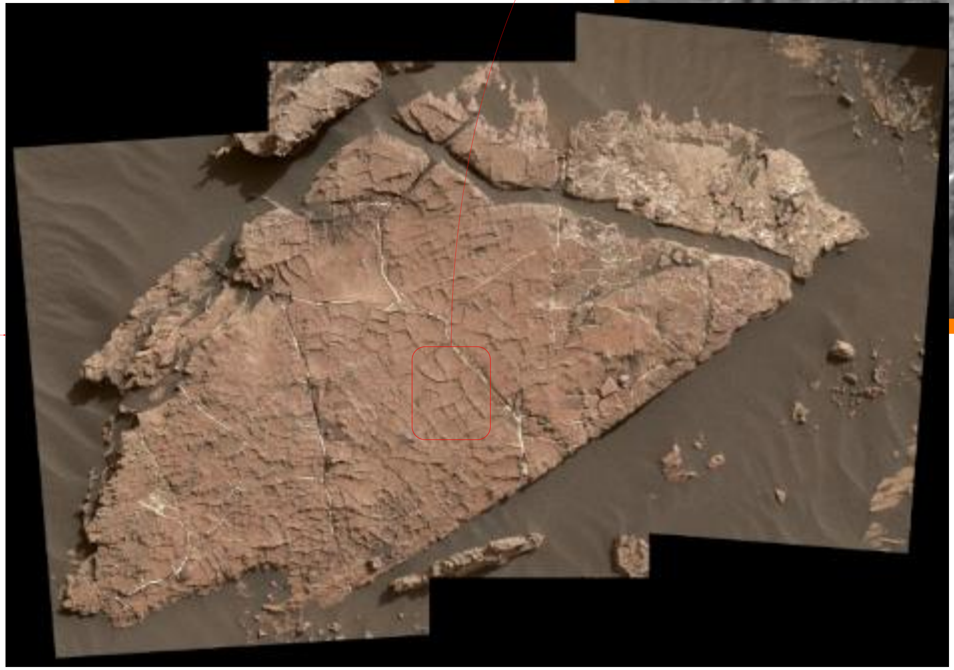

NASA/JPL-Caltech/MSSS

RMI Analyst

main goal:

Create a tool able to propose similar images encountered in the past

secondary:

Propose a geological unsupervised classification of rocks encountered by Curiosity

methods:

unsupervised approaches
test the whole computer vision methods portfolio
mix of « classic » and Deep Learning methods

→ **Project launched early 2025**

Data available

From a regular mosaic to a standardized input

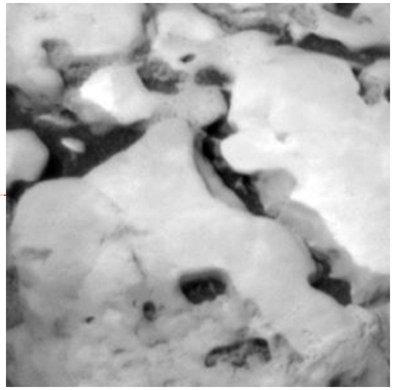

- RMI Mosaics, scaled to the same physical resolution, sliced into 640x640 and centered-normalized on themselves.
- 3002 mosaics pre-laser & 3002 post-laser mosaics
- No redundancy between data, each mosaic being one site (before or after laser shot)
- Geological expertise on part of the dataset:
  - Sulfur blocks (11 images)
  - Iron meteorites (13 images)
  - Wave-ripples rocks (26 images)
  - Broad classification: bedrock, mixed-bedrock, mixed-soil, soil (1825 images)

Computer Vision : methods 1/2

- **Portilla-Simoncelli Statistics** = Steerable Pyramid Decomposition and Reconstruction
  - applying several filters at different scales to different parts of the grayscale spectrum
  - projecting the resulting statistical constraints into a white gaussian noise

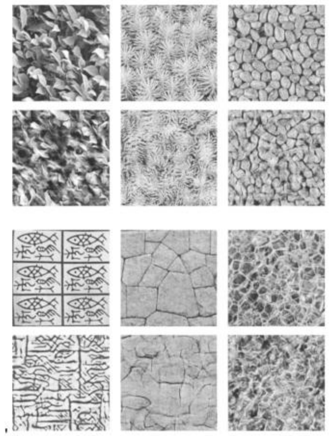

- **Local Binary Pattern** = looks for differences in the neighborhood and compute the resulting histogram

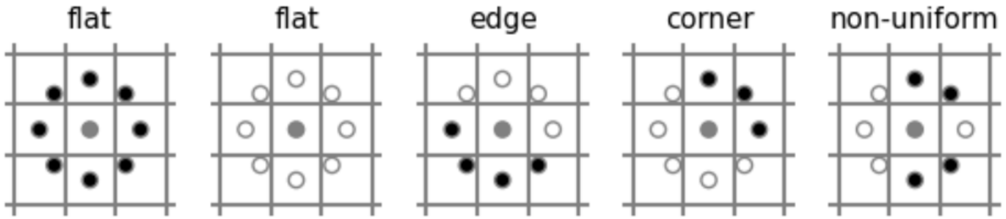

- **SSIM** and derivatives = looks for statistics (subset of PSS) in slices of the image

- **SIFT**, **ORB** and derivatives = search for « key features » in the image
  - search for common forms in the dataset
  - find similarities in the distribution and nature of key features

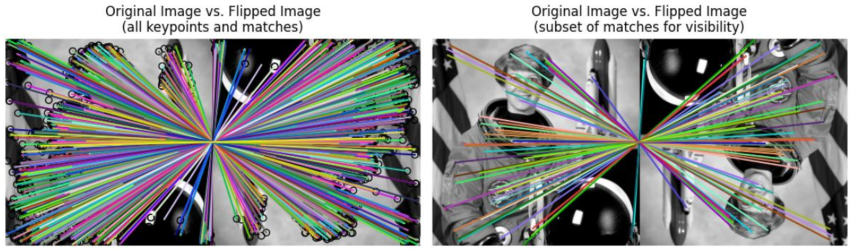

- **Histogram of Oriented Gradients** (HOG)

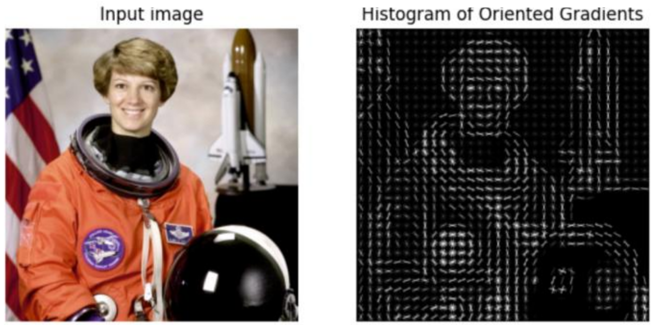

Deep Learning: methods 2/2

- **CNN** = explore Region of Interest in a CNN already fitted (transfer learning) without classifier part
  - fitted with terrestrial data in RGB
  - provides a synthetic feature map in 2D

- **AutoEncoders** = reconstruction of input image with an encoded representation (1-D latent vector)
  - convolutional and/or variational
  - specific to RMI dataset
  - similarity computation using the latent vector
  - search for links between gradients and latent vector ?

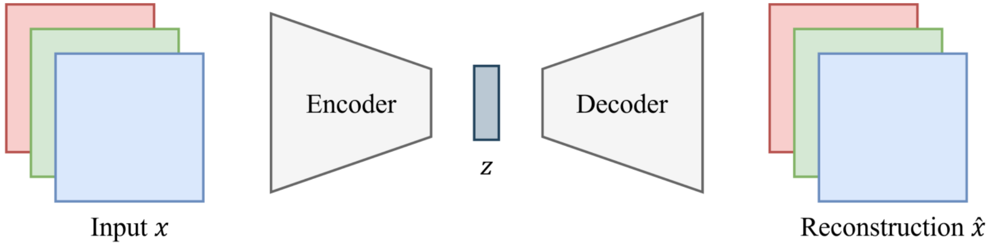

- **VisionTransformer** = same encoding process but with slices of image (tokens) and attention mechanism

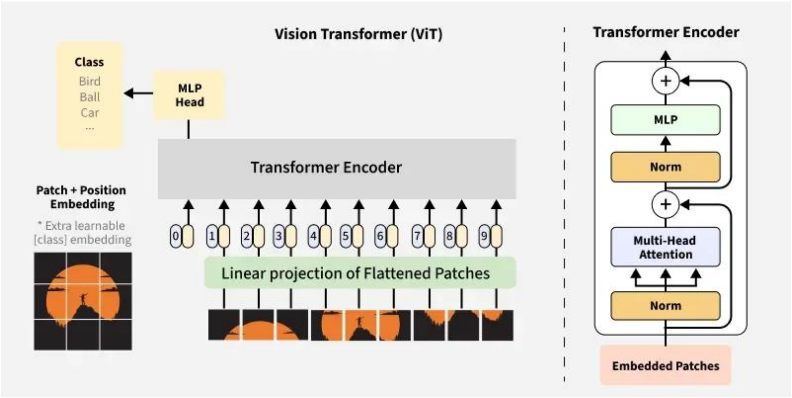

Exploration philosophy

Hamano et al., 2023, *Exploring the role of texture in deep-CNN: Insights from PS statistics, Neural Networks, 168, 300-312*

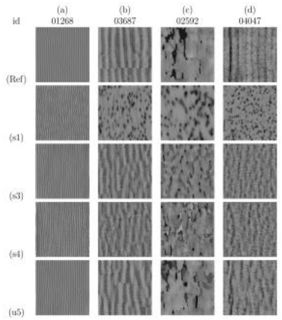

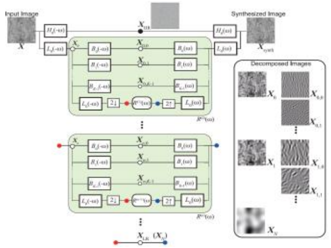

Table 1
Number of features for each group of PSS and minPS. The second and third columns, respectively shows the original number of features and the number after dimensionality reduction.

Group	Original PSS	minPS
Spectral statistics (SS)	18	4
Marginal statistics (MS)	13	2
Linear cross position (LCP)	125	4
Linear cross scale (LCS)	96	4
Energy cross position (ECP)	400	6
Energy cross orientation (ECO)	40	6
Energy cross scale (ECS)	48	4
Total	740	30

Deep CNN (Vgg19) trained to mimic input image

Portilla-Simoncelli process

No black box !

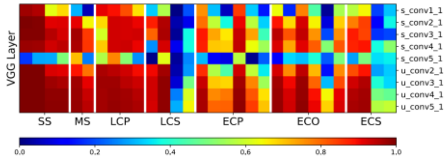

Correlation between PSS simplified and CNN feature maps

Arrignon et al., 2025, *Artificial Intelligence-Enhanced Detection of Biogenicity Using… Experimental MISS , Astrobiology, 6, 414-436*

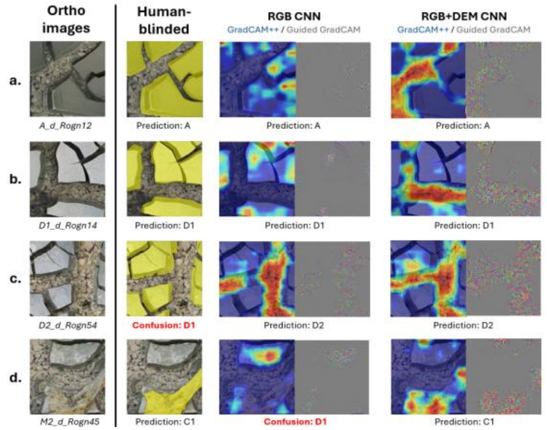

→ Class Activation Maps to explore feature maps and gradients in CNN

Assessment & metrics

- visual assessment when models allow it

- assessment by mean ranking of similarity obtained with the feature vectors :
  - Absolute (all images of each type are classified):
    - Iron meteorites
    - Sulfur blocks
    - Wave ripples
    - etc...
  - Relative (only a subset is documented):
    - Rough classification: rock / rock+sand / sand

- use of features vectors in unsupervised clustering:
  - measurement of accuracy after projecting real labels
  - explore the mistakes of each method

- results differ if clustering maths are different from similarity computation maths = different insights

Software specs

- able to explore all images with regular IDs, before and after laser
- propose a customized experience to navigate in the data
- can integrate many computer vision & DL models
- provide global stats and graphs for the different methods
- expose the models results as much as possible, especially for DL models

→ Python : pyside6 / openCV / sk-learn / pytorch

Current version

Mission timeline

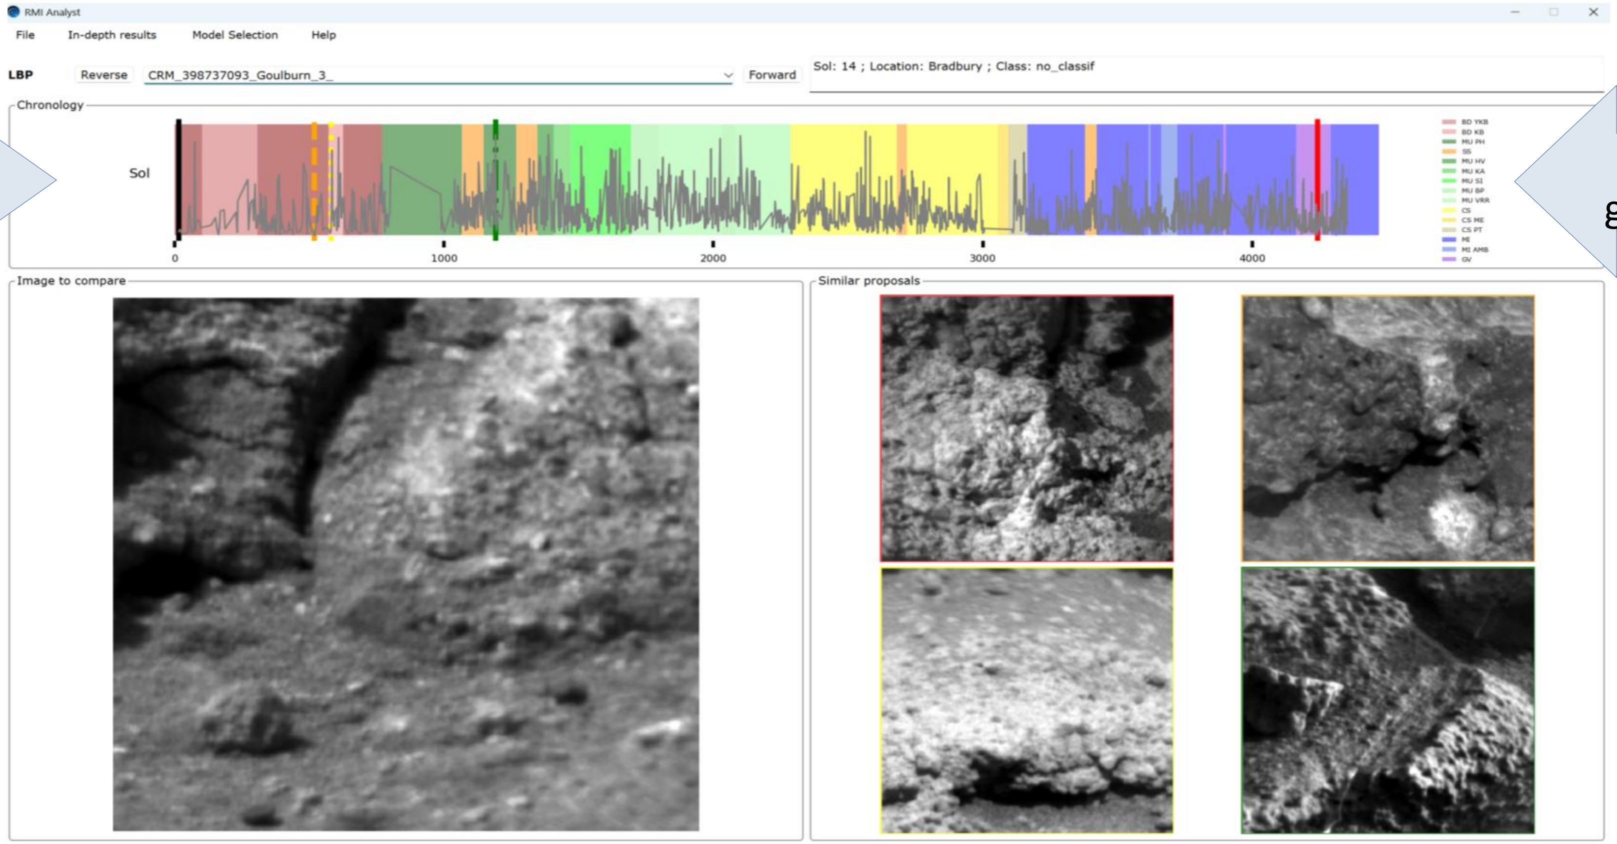

Color-coded geologic regions

Reference image

Best matches

Methods implemented with full similarity computation:

- Local Binary Pattern (LBP)
- Structural Textural Similarity index (STSIM)
- Histogram of Oriented Gradients
- Autoencoders (simple convolutional architecture)
- LBP on Autoencoders output

Methods yet to implement:

- ORB
- CNN
- Variational Autoencoders
- Vision Transformers

Autoencoders

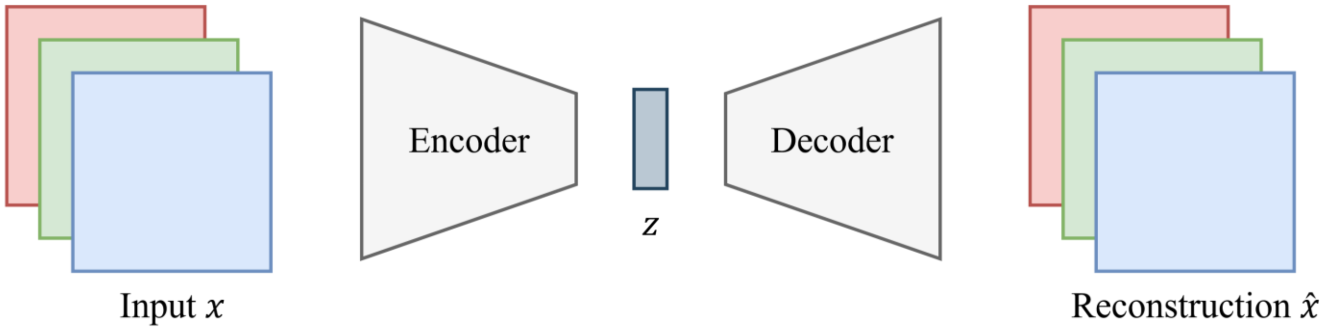

First attempts with very simple convolutional architecture:

- 10 or 14 convolutional layers with 5- or 7-dimensions reductions
- 100 , 250- or 1000-variables latent spaces
- sigmoid as activation function
- trained from scratch
- assessment at 130 epochs

→ **Best trade-off « visual performance / computation time » with 14 layers and 250 latents**

Results autoencoders

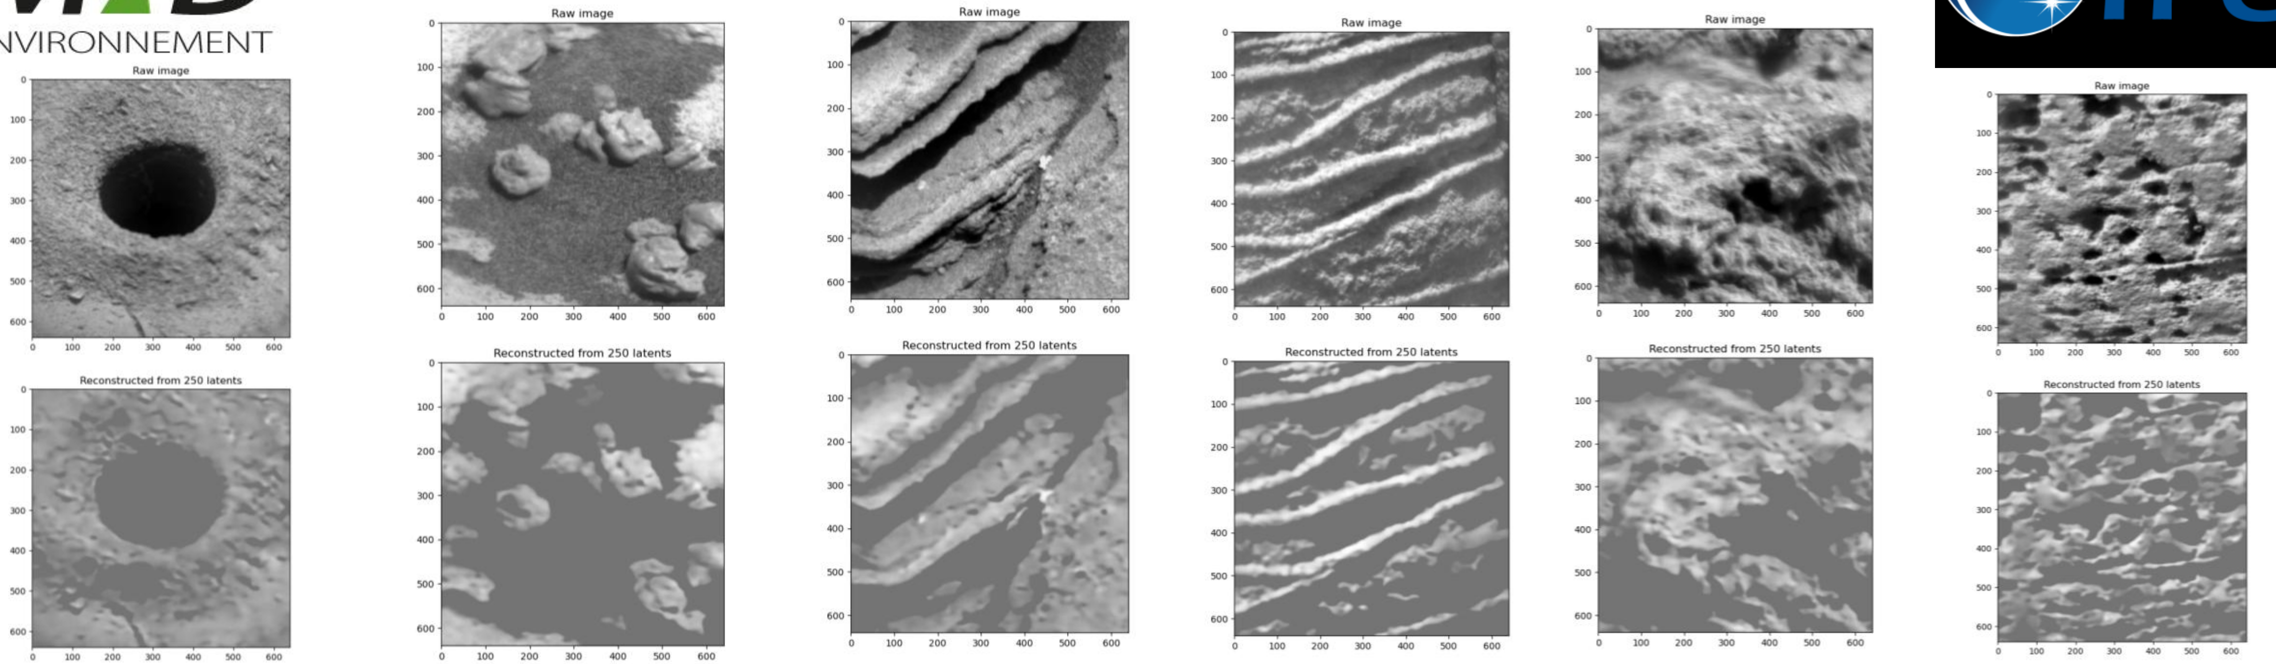

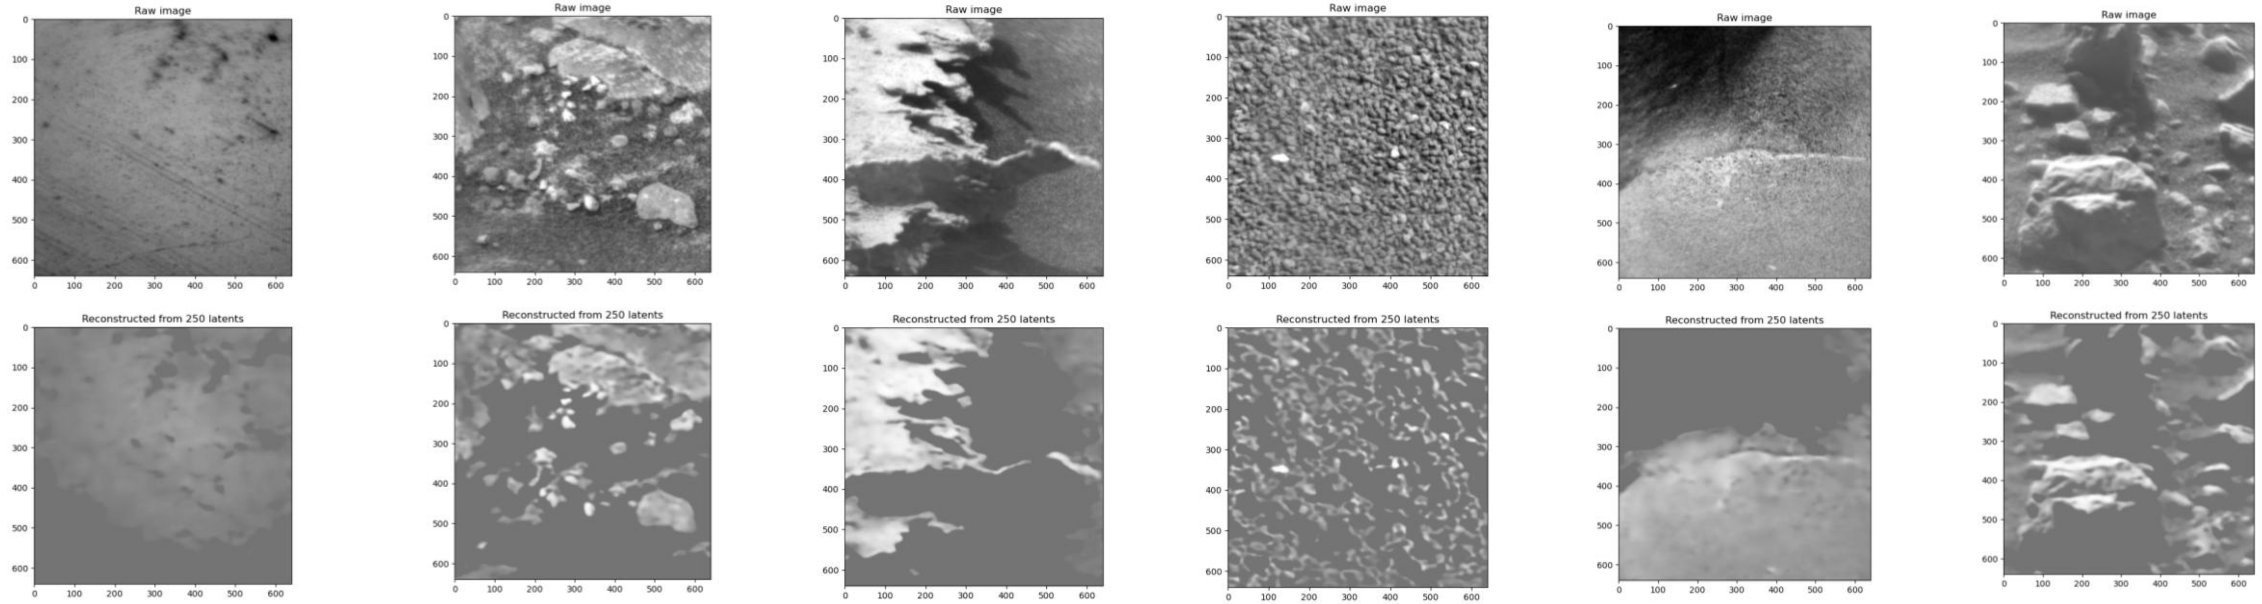

Results autoencoders

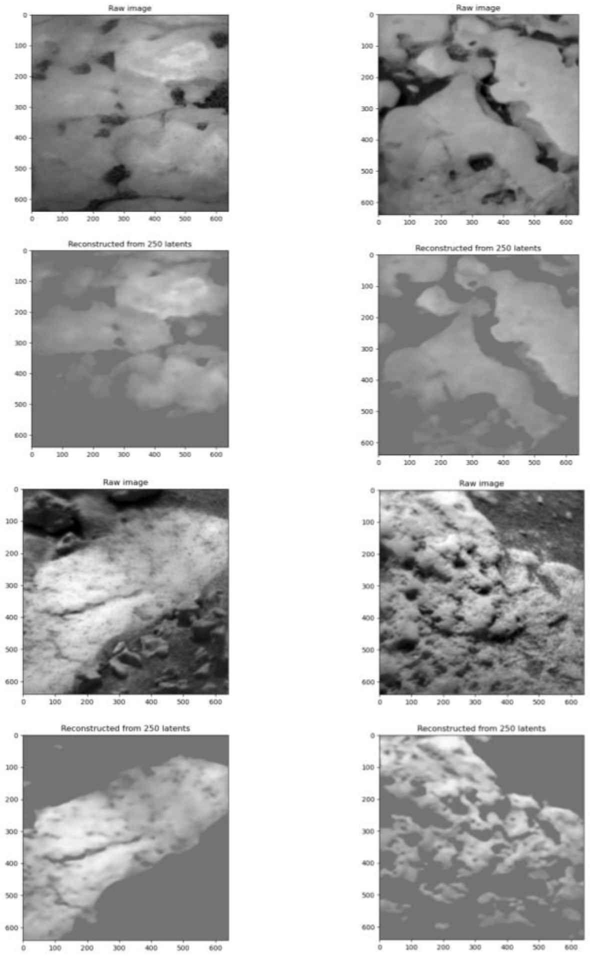

sulfur blocks

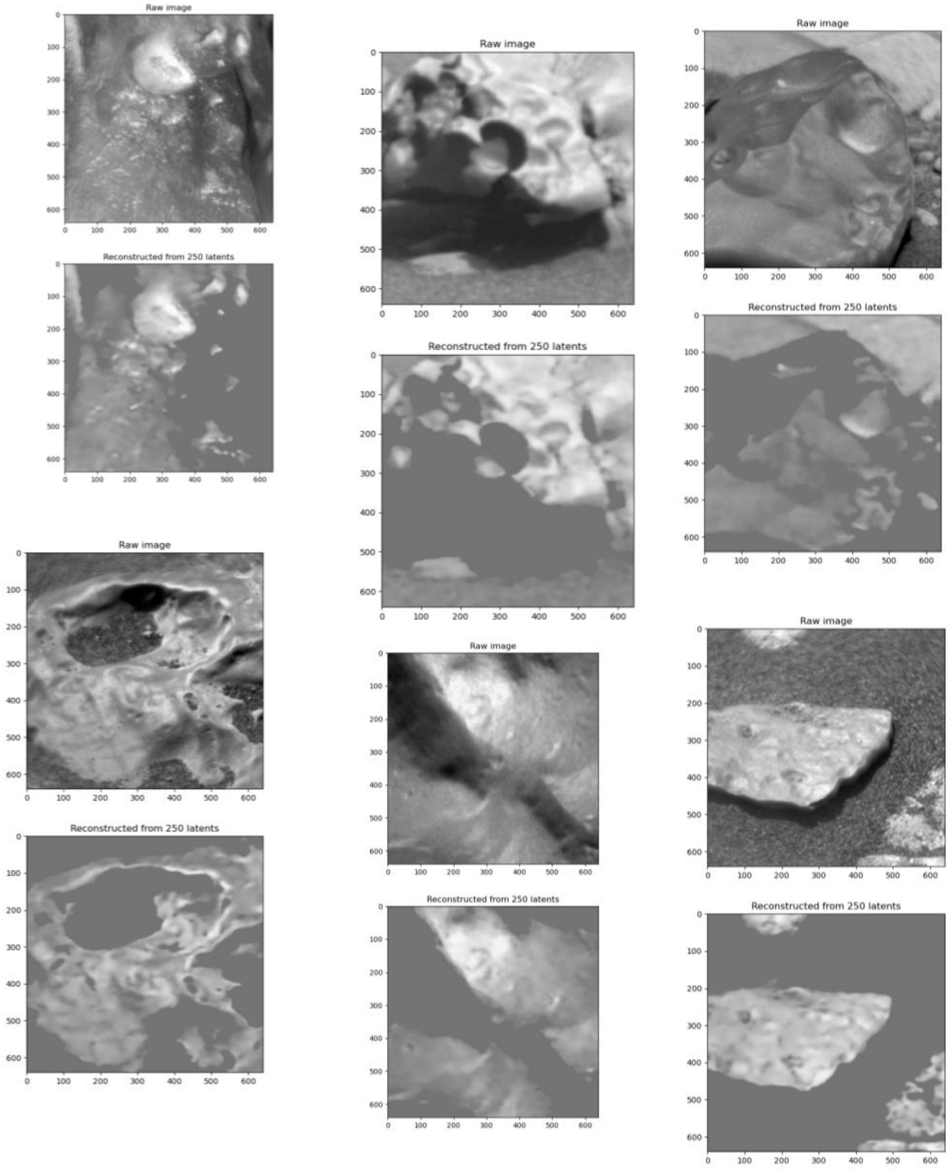

Iron meteorites

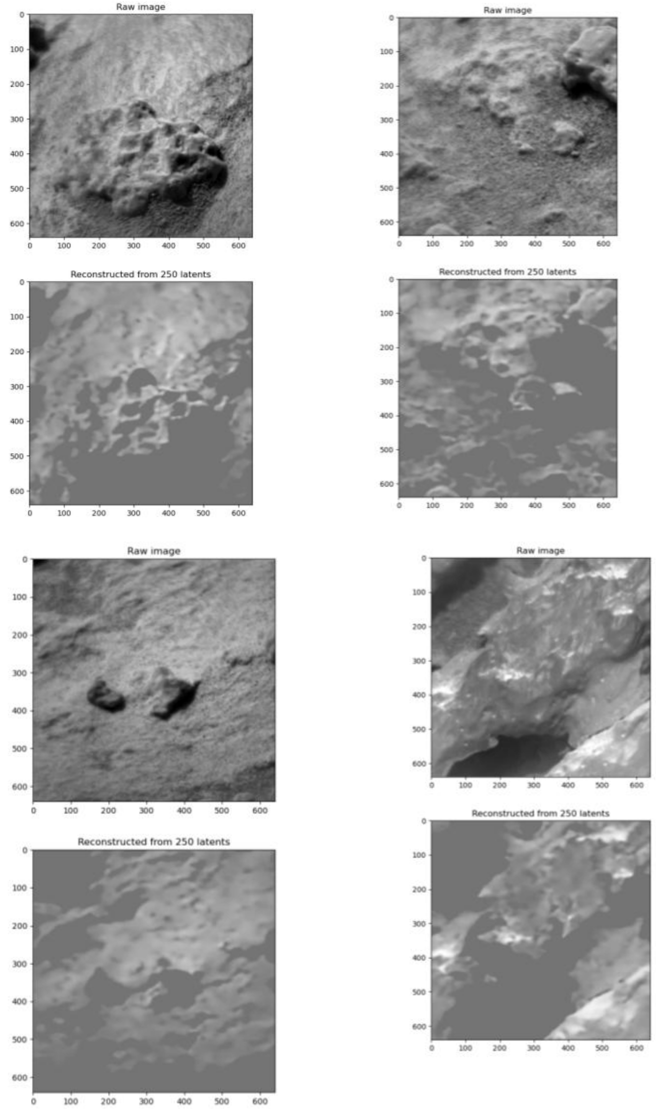

Wave ripples

Mean Ranking Metric results

Sulfur	mean	*variance*
STSIM-M	0,518	*0,005*
LBP	0,497	*0,029*
HOG	0,587	*0,002*
AE	0,745	*0,003*
LBP_on_AE	0,817	*0,018*

N = 11

IronMeteor	mean	*variance*
STSIM-M	0,516	*0,026*
LBP	0,449	*0,013*
HOG	0,478	*0,009*
AE	0,635	*0,001*
LBP_on_AE	0,610	*0,014*

N = 13

WavesRipple	mean	*variance*
STSIM-M	0,392	*0,002*
LBP	0,437	*0,003*
HOG	0,339	*0,001*
AE	0,538	*0,001*
LBP_on_AE	0,592	*0,002*

N=26

Bedrock	mean	*variance*
STSIM-M	0,475	*0,001*
LBP	0,467	*0,004*
HOG	0,407	*0,001*
AE	0,438	*0,001*
LBP_on_AE	0,549	*0,001*

N = 1108

Mixed_bedrock	mean	*variance*
STSIM-M	0,430	*0,001*
LBP	0,423	*0,006*
HOG	0,641	*0,001*
AE	0,586	*0,001*
LBP_on_AE	0,442	*0,001*

N = 530

Mixed_soil	mean	*variance*
STSIM-M	0,437	*0,005*
LBP	0,422	*0,021*
HOG	0,378	*0,001*
AE	0,682	*0,001*
LBP_on_AE	0,563	*0,006*

N = 13

Soil	mean	*variance*
STSIM-M	0,647	*0,007*
LBP	0,371	*0,060*
HOG	0,406	*0,069*
AE	0,419	*0,040*
LBP_on_AE	0,451	*0,004*

N = 174

Results LBP, good matches…

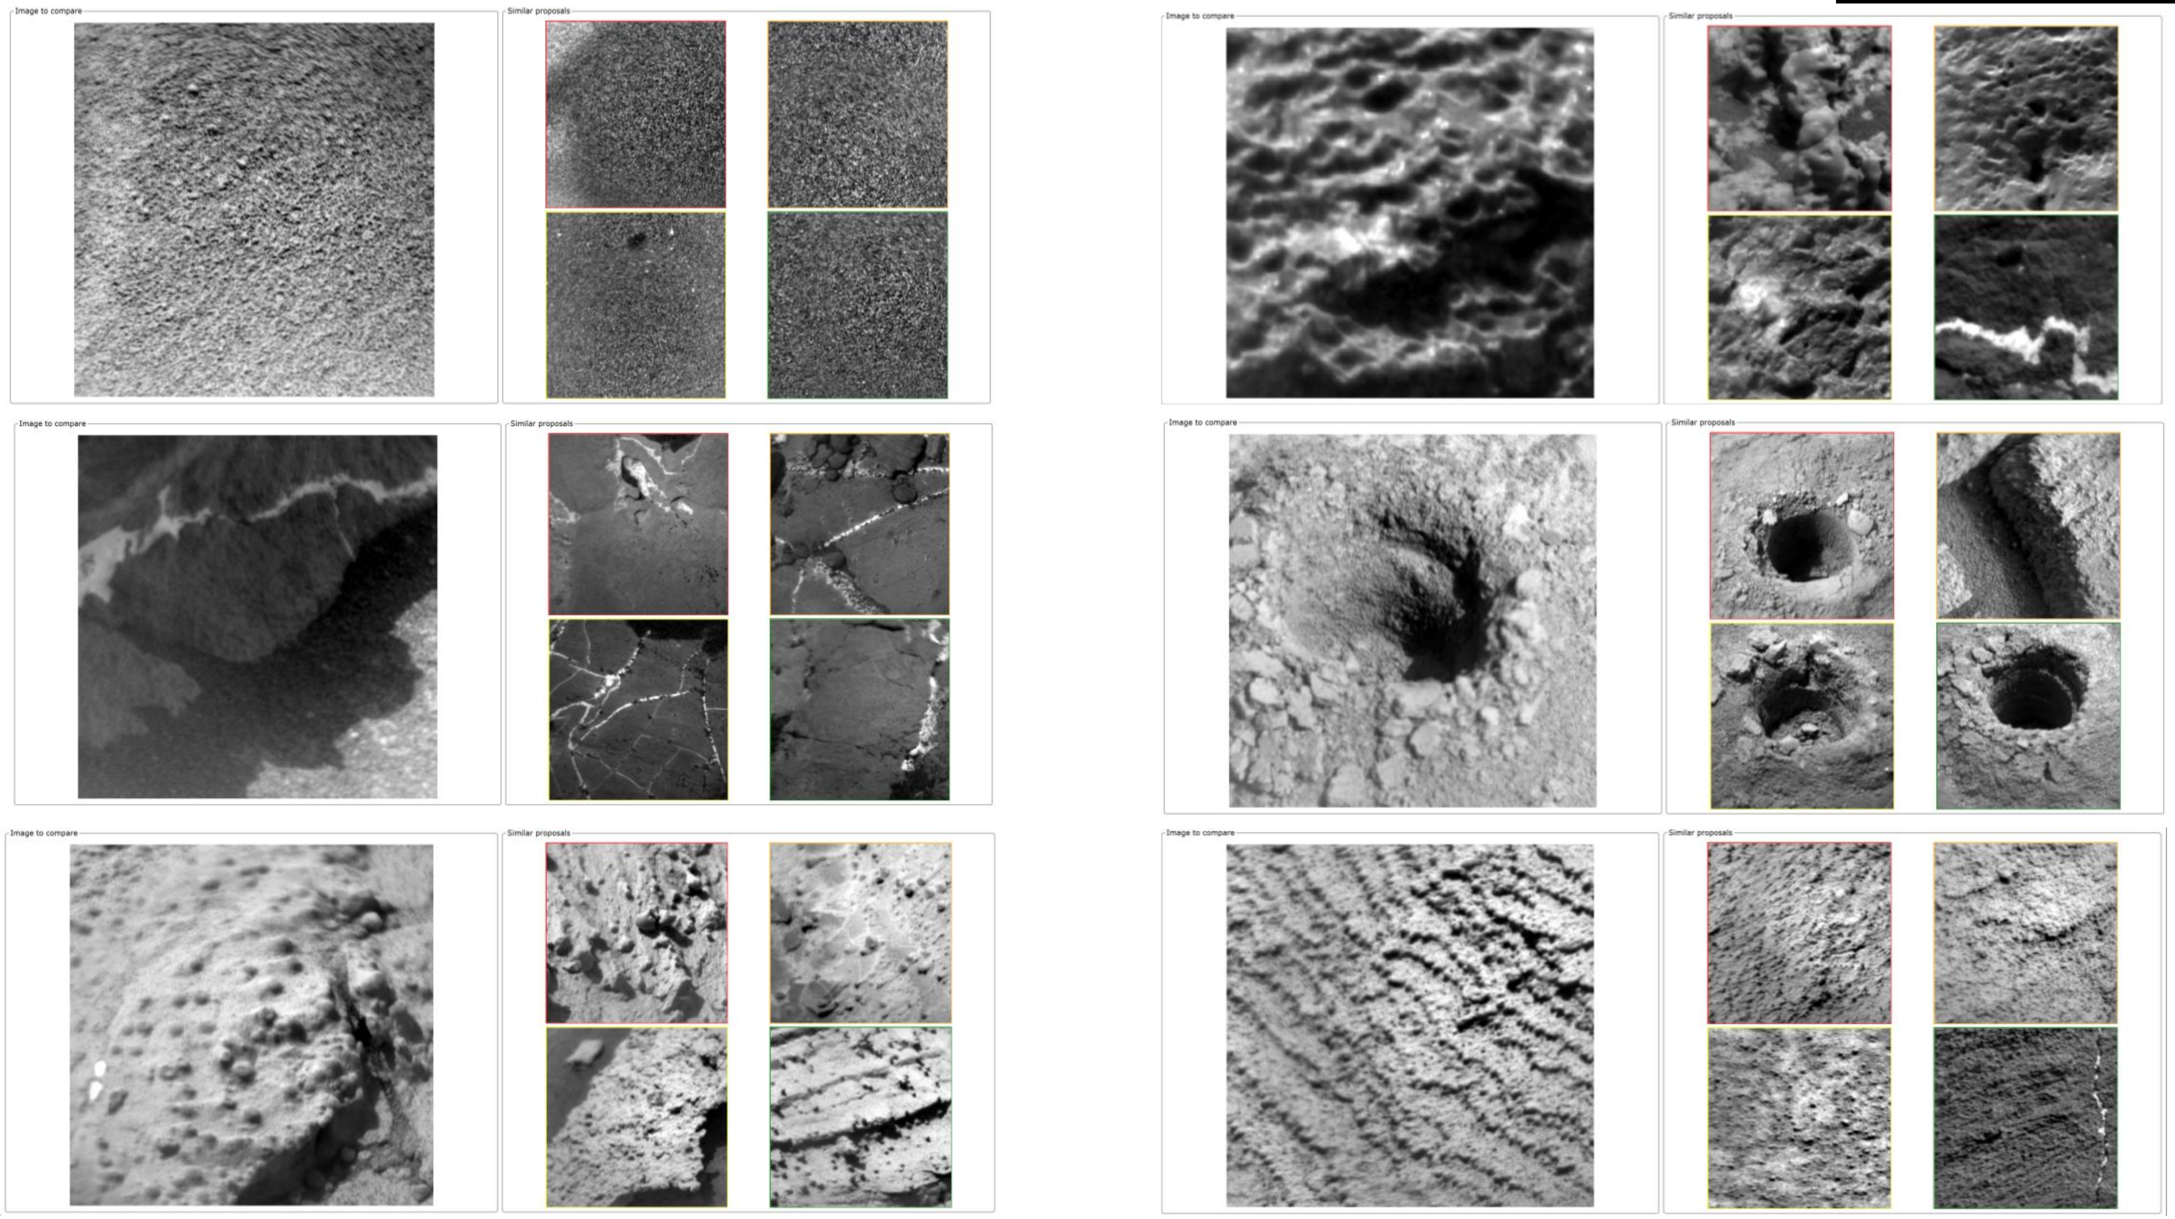

Results LBP, bad cases…

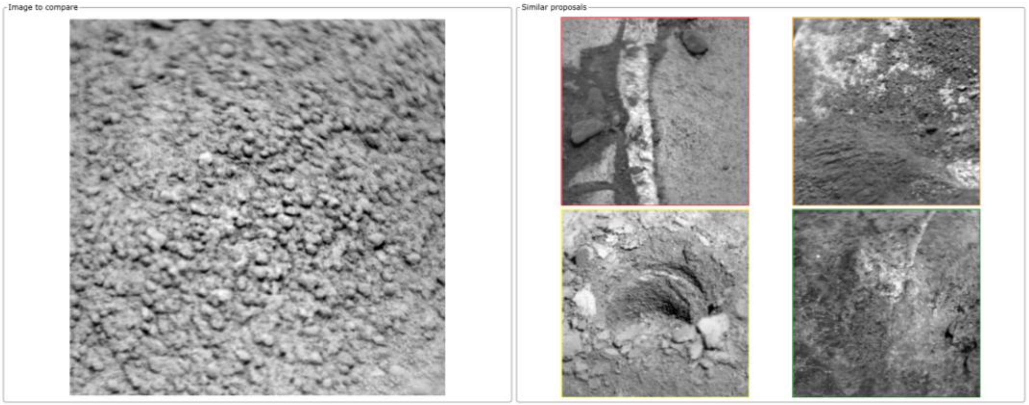

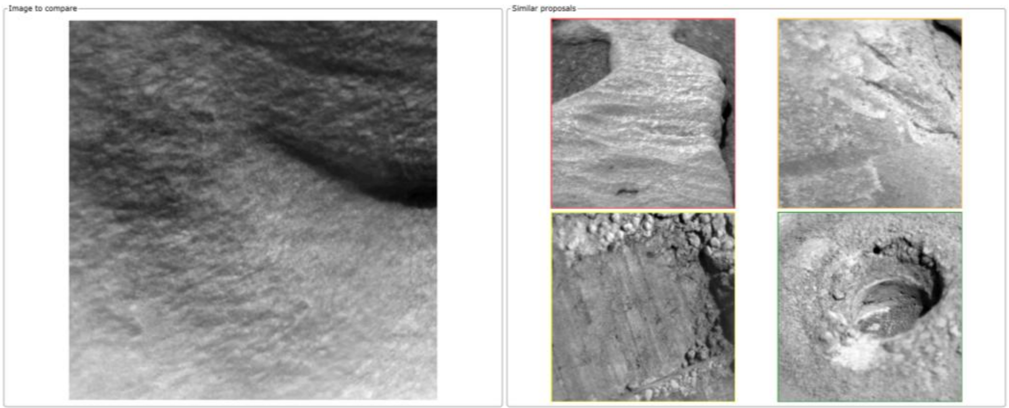

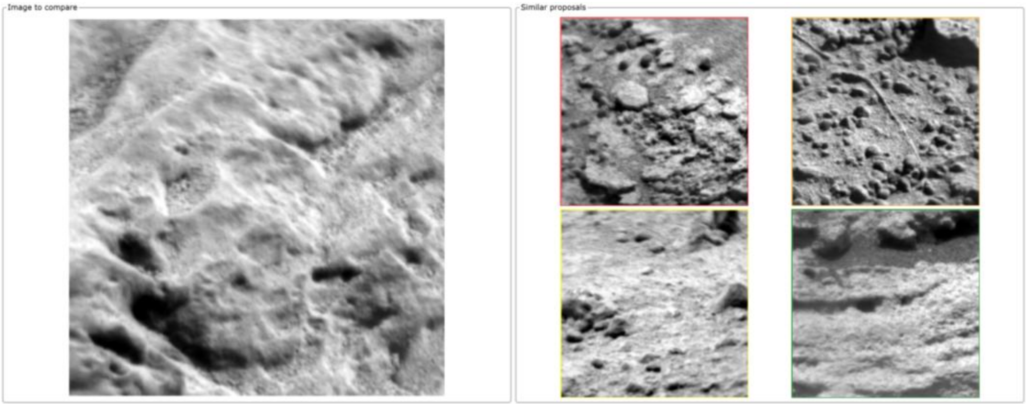

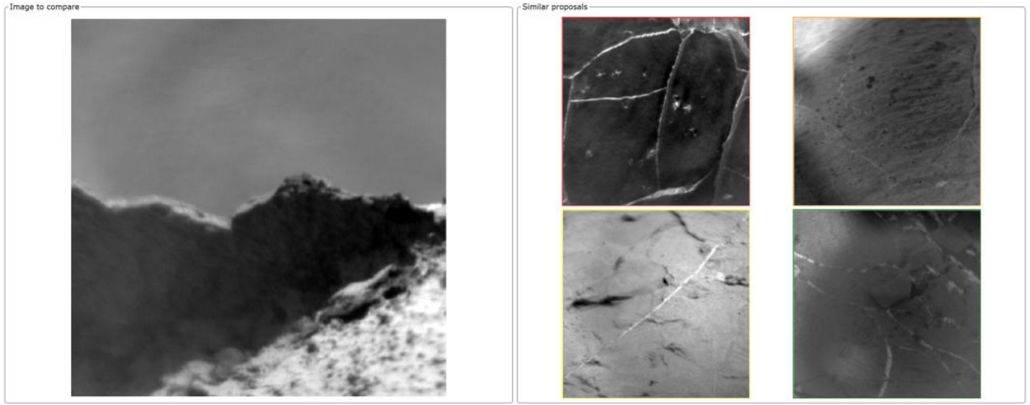

Results summary

- Not great overall statistics until now, but some trends are emerging.
- Local Binary Pattern is consistent both visually & metrically.
- Auto-Encoders provide decent reconstructs, but are not the best at finding similarities.
- Classifications can be too fuzzy depending on what the models take into account.

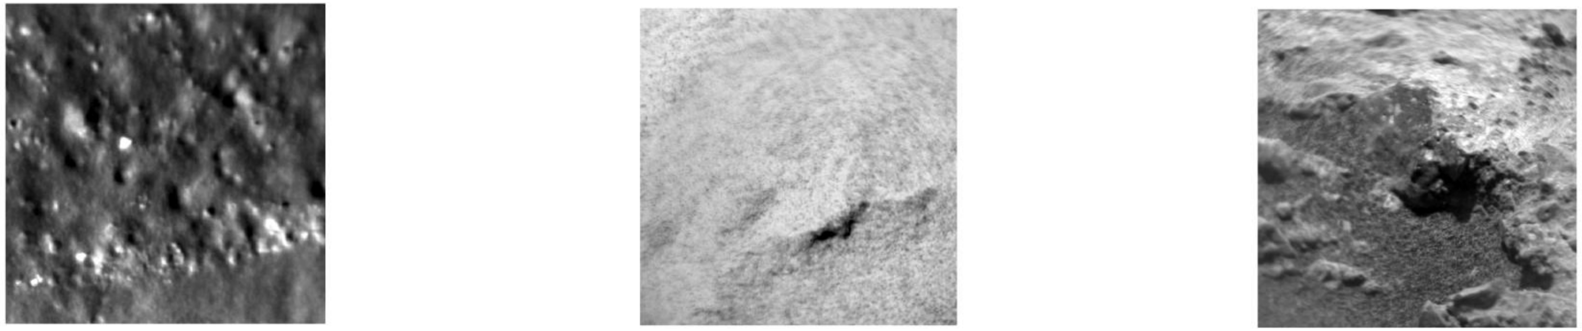

bedrock

Further developments 2026

Two parallel workflows

Models improvement

- Autoencoders architectures
- Vision Transformers (DINO, Mars ViT-mae)
- LBP & HOG adjustments
- Unsupervised clustering of feature vectors
- Similarity computation

mixed AI approach

- using expertise to build decision rules & define global patterns
- sequential model with AutoEncoders then HOGs / LBPs

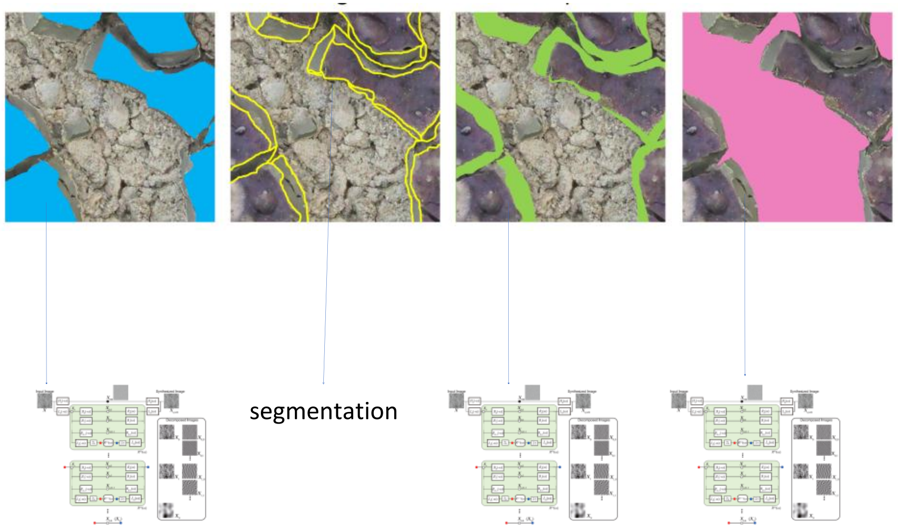

→ **Still at early stage of the project!**

Thank you for your attention !