Planetary Rover Absolute Localization by Combining Visual Odometry with Orbital Image Measurements

M. Lourakis and E. Hourdakis

Institute of Computer Science
Foundation for Research and Technology – Hellas
Heraklion, Crete, Greece

ASTRA 2015, 11 - 13 May 2015, Noordwijk, The Netherlands

Introduction

- Planetary exploration increasingly depends on planetary rovers
- Accurate localization is essential to planetary rovers
- This talk is concerned with increasing the accuracy of rover localization by employing images collected on the ground and from orbit
- Reported work was pursued in *SEXTANT*, an ESA-funded activity aiming to develop visual navigation algorithms suitable for use by ExoMars Martian rovers

Background

- Visual odometry (VO) is a technology that has been used with considerable success by Martian rovers (e.g., MERs and MSL)
- VO works by tracking natural features in successive images
- VO provides only relative pose estimates and suffers from drift that accumulates without bound over time
- Countering drift with techniques such as visual SLAM or local bundle adjustment, challenges the limited computational capacity of current rovers
- Non-visual sensors (e.g., IMU) can only reduce the rate of drift but do not eliminate it completely

Using prior knowledge

- Localization accuracy can be increased by using prior knowledge regarding the environment (e.g., maps)
- Orthorectified aerial images of planetary surface (aka orbital images) are nowadays available, e.g. Mars HiRISE / MRO
- Several “map”-based localization techniques tailored to terrestrial environments exist, employing features such roads, line segments and building footprints or sensors such as GPS only available on Earth
- Here we choose to use as features boulders which abound on Mars

Sample images

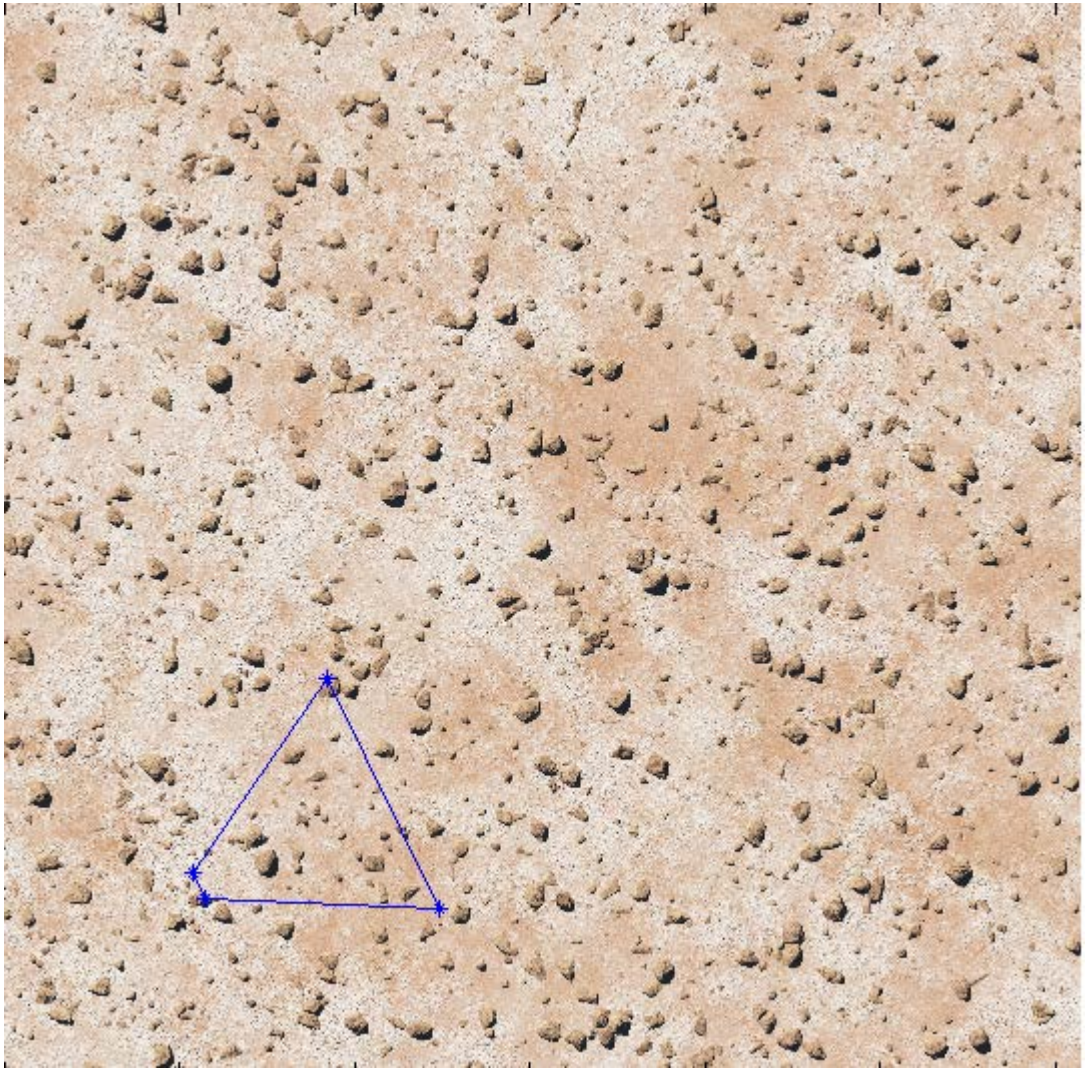

Orbital and ground images differ dramatically in appearance!

Using boulders for localization

- We seek to extract boulders visible in the orbital images and match them to those seen in the ground images
- Two boulder detection pipelines, i.e. orbital – ground
- Boulder matching is the primary challenge; cannot be based on appearance cues such as local patch descriptors
- The relative geometric arrangement of boulders remains the same, thus their matching is based on geometry and strives to align constellations of ground and orbital boulders
- This provides absolute localization as the orbital boulder locations are determined offline in a cartographic system

Overview of the approach

- Boulders are extracted offline from orbital images and their locations are recorded
- Boulders are detected online in the ground stereo images acquired by the rover
- Ground boulders are combined with VO to build a local ground map with relative boulder positions
- The rover periodically matches its local ground map against the map originating from the orbital images
- Successful matching permits absolute pose refinement, thus accounting for drift

Visual odometry

- VO estimates the inter-frame camera motion

- Binocular camera system
- Harris corner detection
- SIFT descriptor extraction and matching
- Sparse stereo triangulation
- Pose determination from two matched 3D point clouds

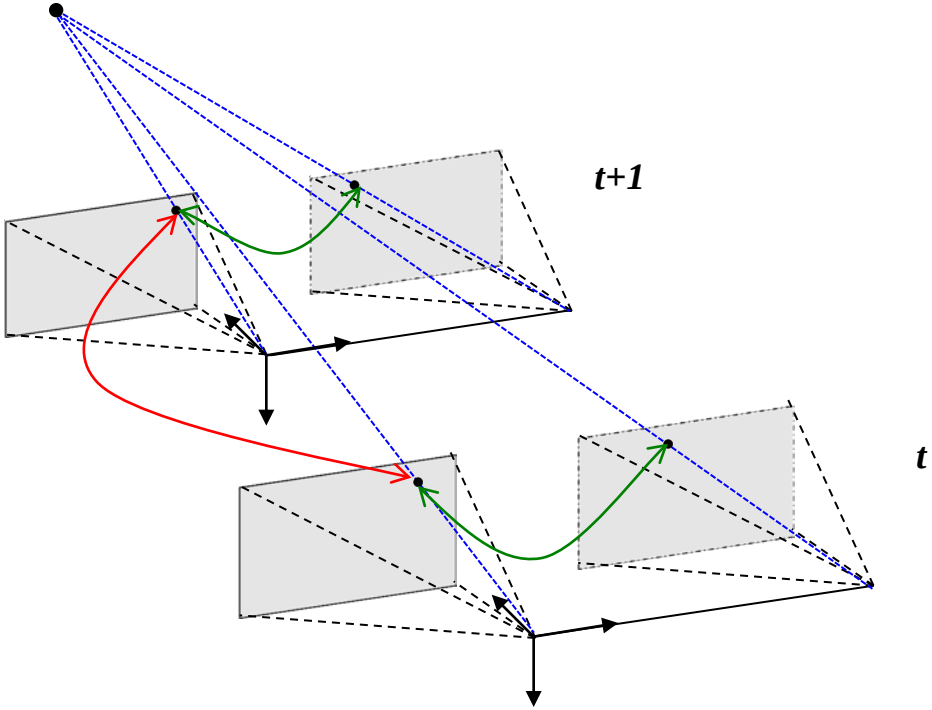

Blue: projection rays, Green: spatial matches, Red: temporal matches

Boulder detection

- Boulders have to be detected both in ground and orbital images
- Boulders differ with respect to their morphology, size, texture, color, etc.
- Shadows, changes in illumination and occlusions complicate things further
- Literature reports attempts to develop “boulder detectors” for ground images; these are not yet very reliable
- Boulder detection in orbital images is somewhat easier

Boulder detection: orbital images

- Image binarization with adaptive thresholding
- Noise reduction with median filtering
- Detection of boulders via connected components labeling
- Size filtering for eliminating very small components
- Boulder locations are computed from the centroids of detected regions; pixel coordinates then converted to cartographic coordinates
- Centroids are a coarse approximation of true boulder locations; this nevertheless suffices, as will be shown later

Boulder detection in orbital images example

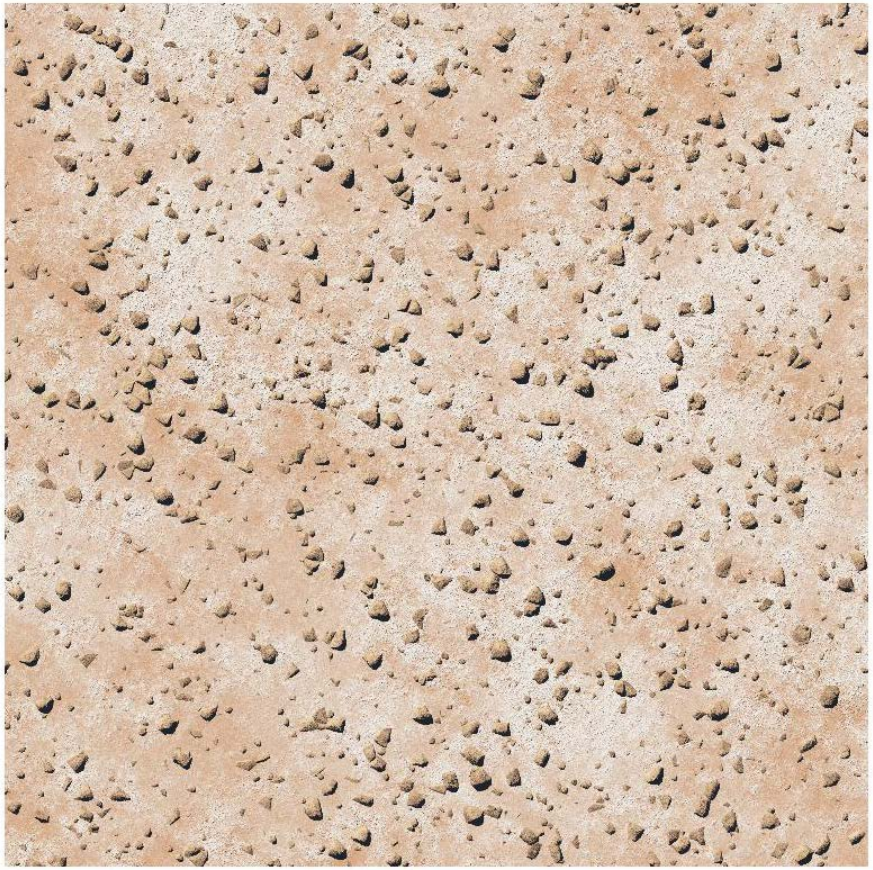

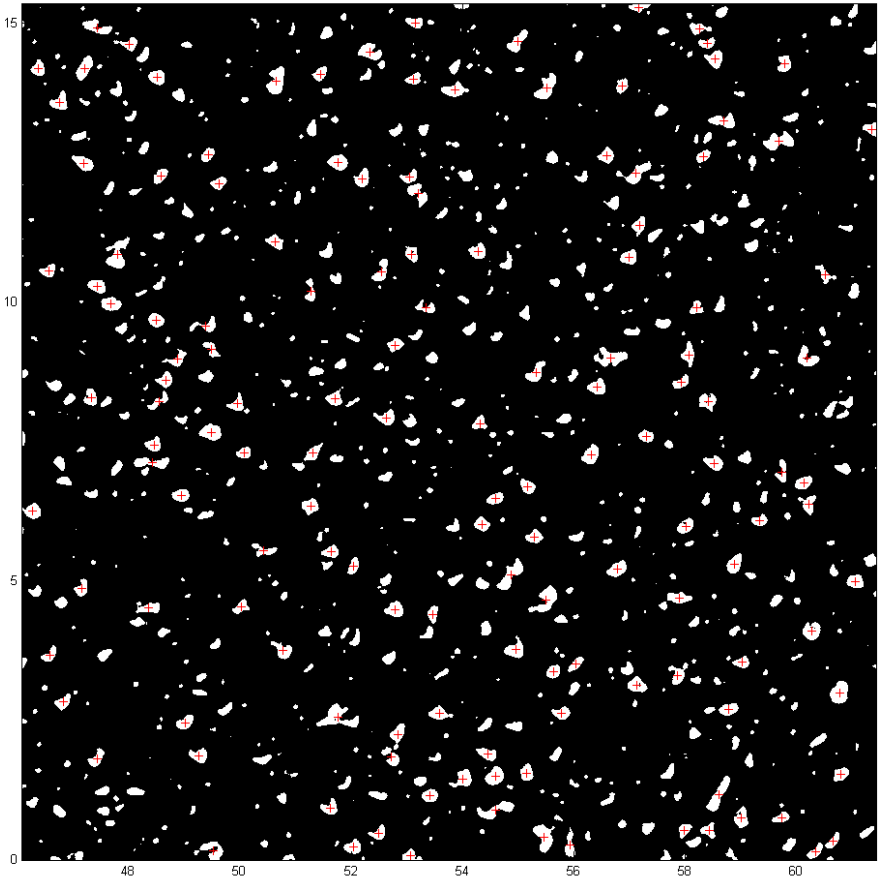

Part of an orbital image *(left)* and the boulders detected in it

Boulder detection: ground images

- Ground images segmentation into boulders & ground
- Formation of boulder point clouds via 3D reconstruction and temporal matching
- Refinement of boulder point clouds via clustering
- Computation of boulder point cloud centroids
- Transformation of centroids to world origin using VO estimates
- Down-projection by dropping Z coordinates
- Local map update with the centroids of detected boulders

Ground boulder detection: segmentation

- Segmentation is based on the mean shift algorithm

- Small regions are ignored
- Performs reliably, suffers from over-segmentation
- Post-processing: neighboring regions are merged when they have similar intensity properties
- Triangulated points are also merged in 3D (more later)

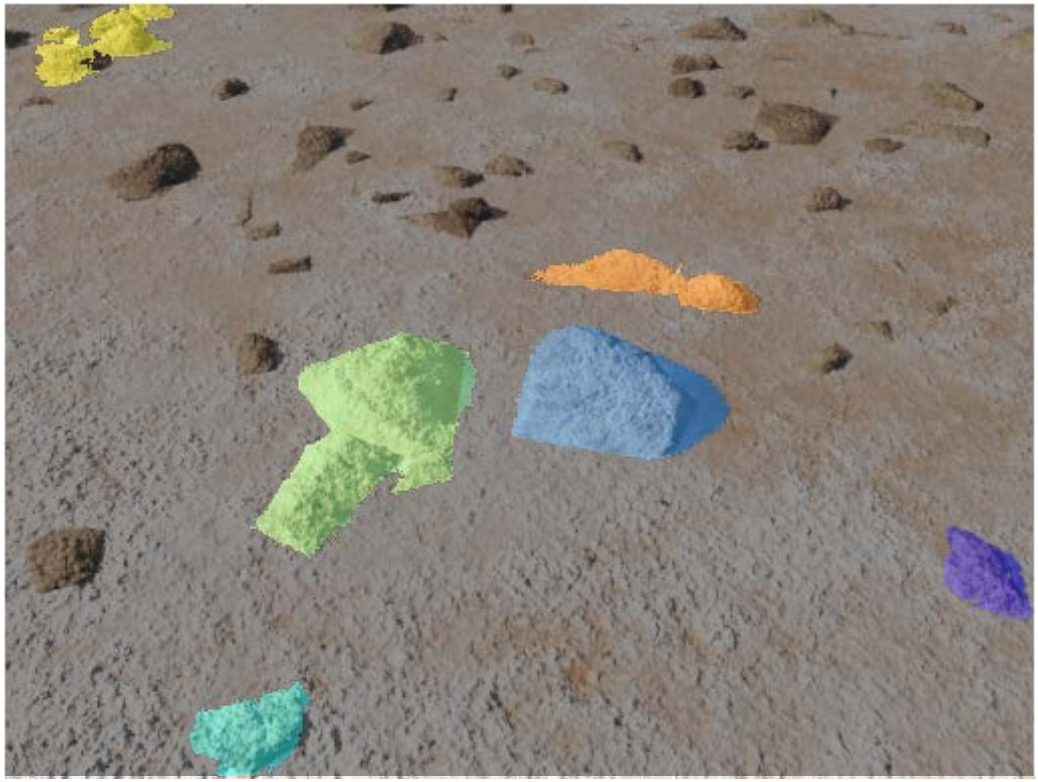

Ground boulder detection: sparse reconstruction

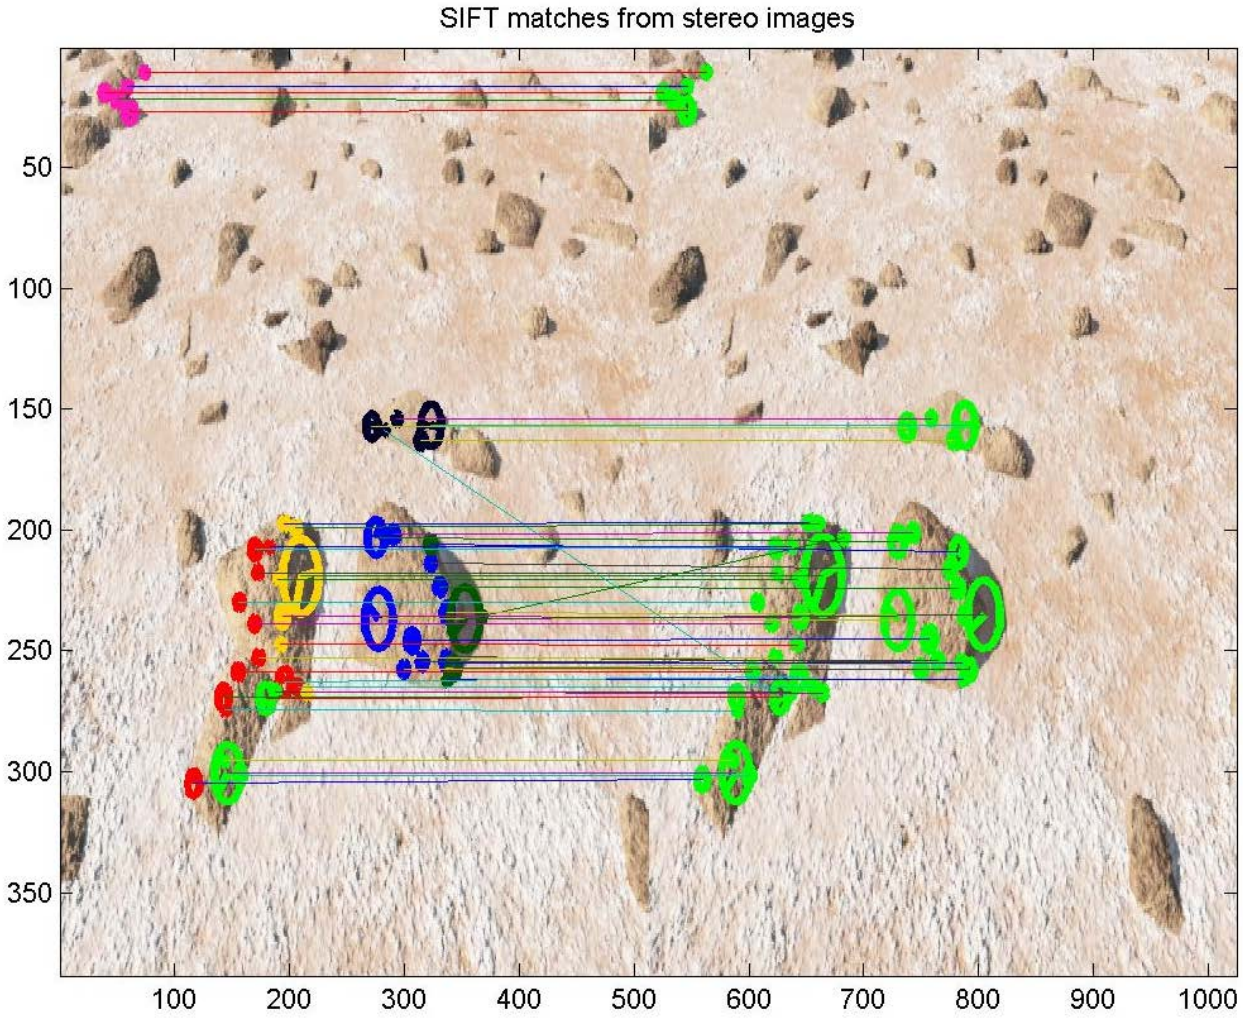

Features extracted from segmented regions and matched in stereo are reconstructed in 3D to form point clouds

Ground boulder detection: temporal matching

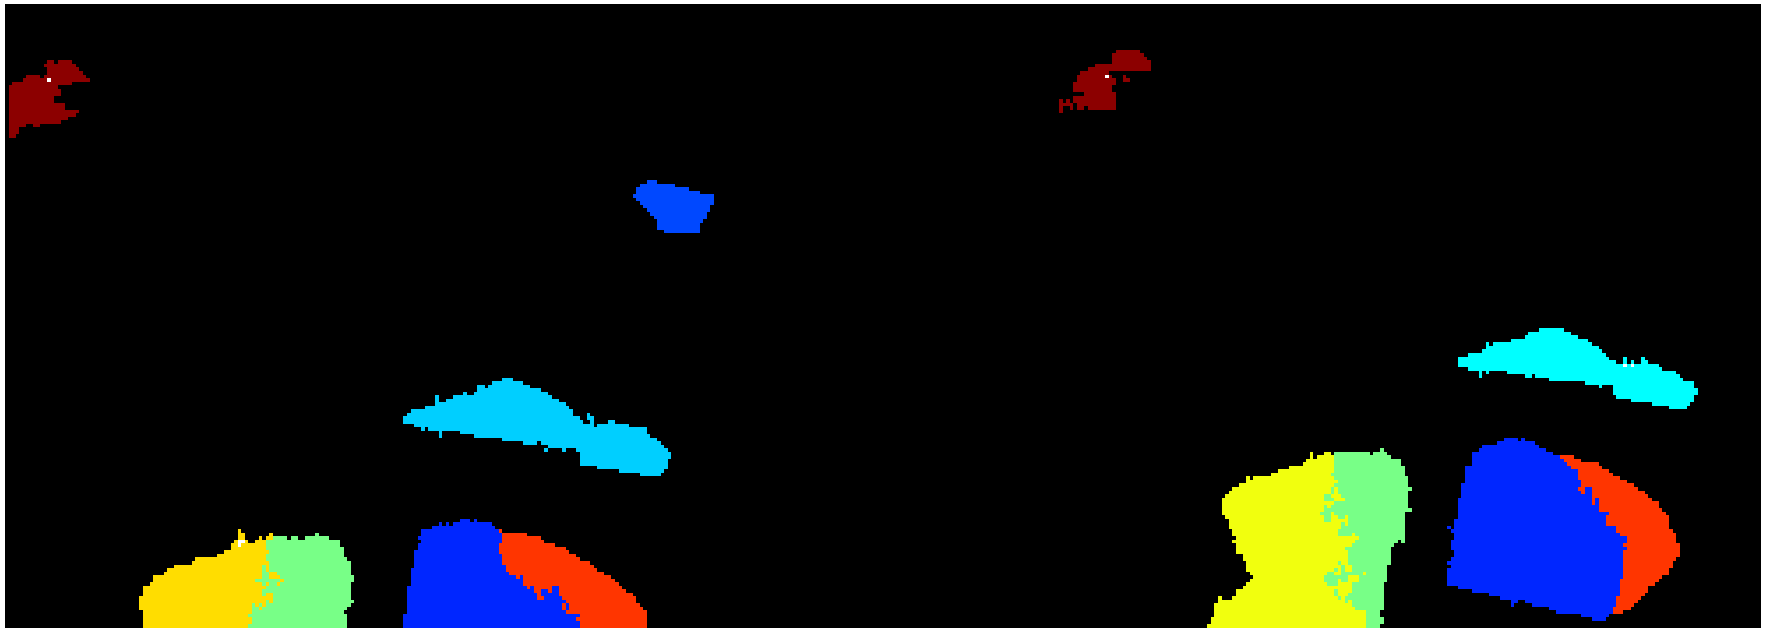

- Features extracted from segmented regions are also tracked over time and used to temporally match boulder regions
- Points reconstructed from matching regions are presumed to be on the same boulder and are assigned to the same point cloud

Ground boulder detection: clustering

Point clouds refined with agglomerative clustering

Ground – orbital boulder matching

- Key assumption: the rover knows its approximate location (i.e. the VO estimate can be trusted)
- Boulders extracted locally from the ground images are matched with the **ICP** algorithm against the overhead boulders located in the vicinity of the presumed rover location
- A byproduct of matching is the 2D transformation (Z-axis rotation and XY-axes translation) aligning the ground with the overhead boulders
- The VO estimate is refined with the 2D aligning transformation
- VO continues to estimate the relative displacement from the refined estimate
- The VO refinement is executed every few meters

Dataset for evaluation

- Graphically rendered images of a synthetic Mars-like environment *(courtesy GMV)*
- A stereo image sequence with accurate ground truth poses that corresponds to an S-shaped, 100m long traverse
    - Ground images are 512 × 384 pixels
- An orbital image with size 15360 × 15360, corresponding to an area of size 76.8 × 76.8 m^2 (resolution: 0.005 m/pixel)
    - Orbital image consists of 236 Mpixels
    - Coarser resolutions generated by sub-sampling

Results: boulders extracted & matched

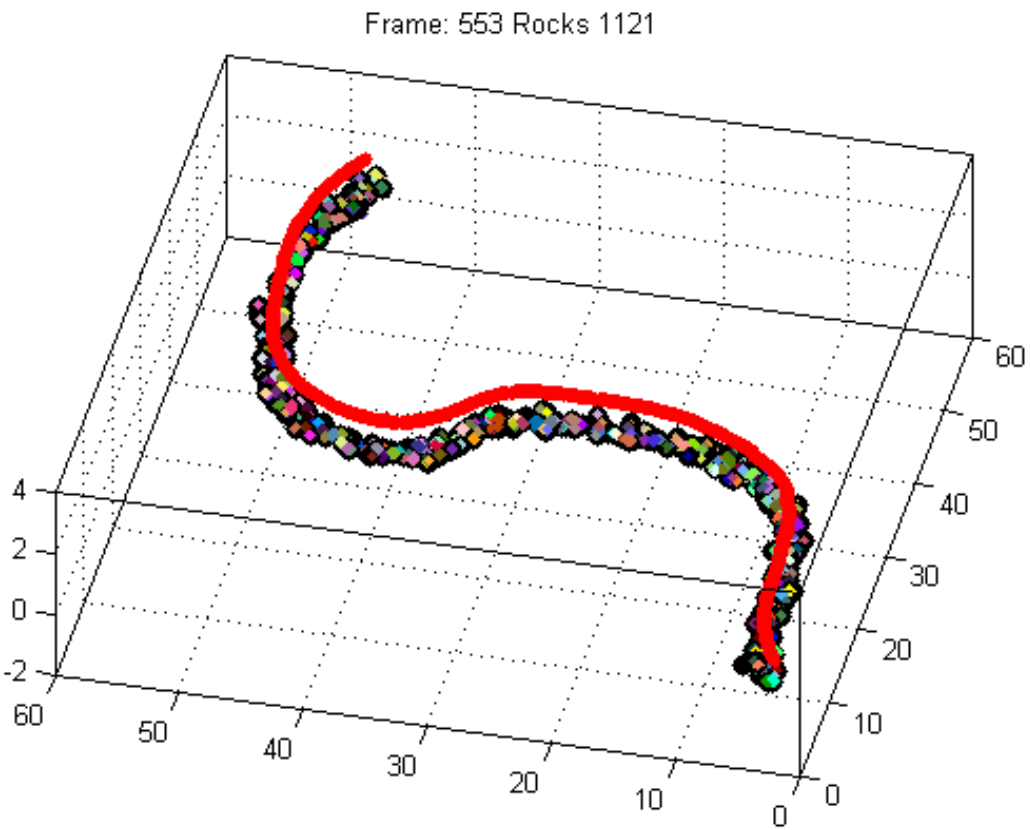

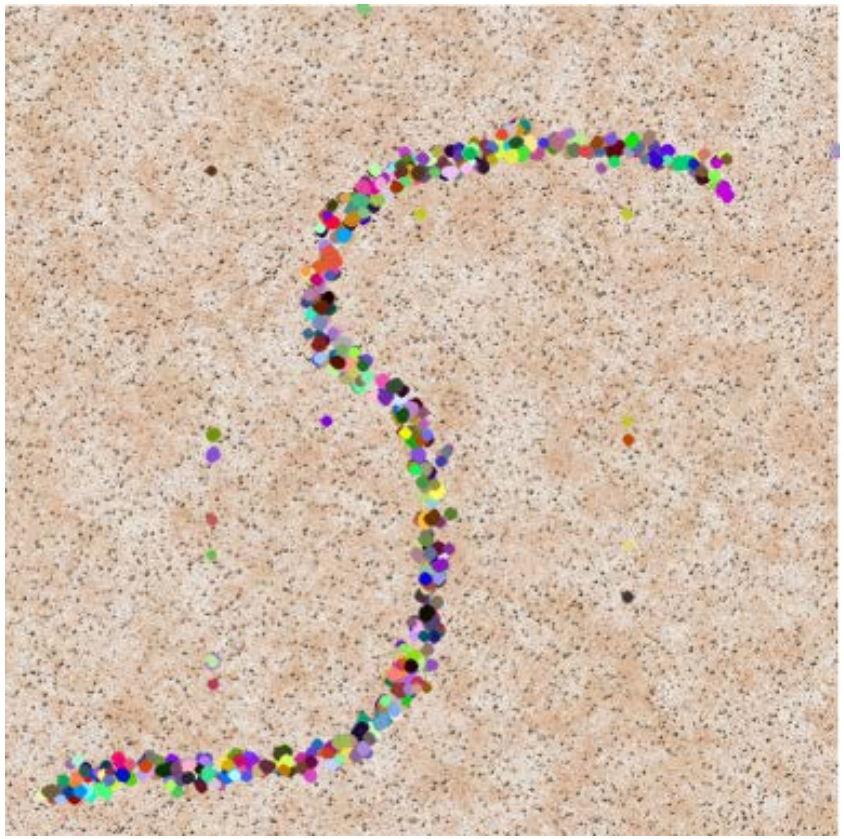

Ground boulders along the entire trajectory: 3D view *(left)* superimposed on orbital image *(right)*

Results: positional accuracy

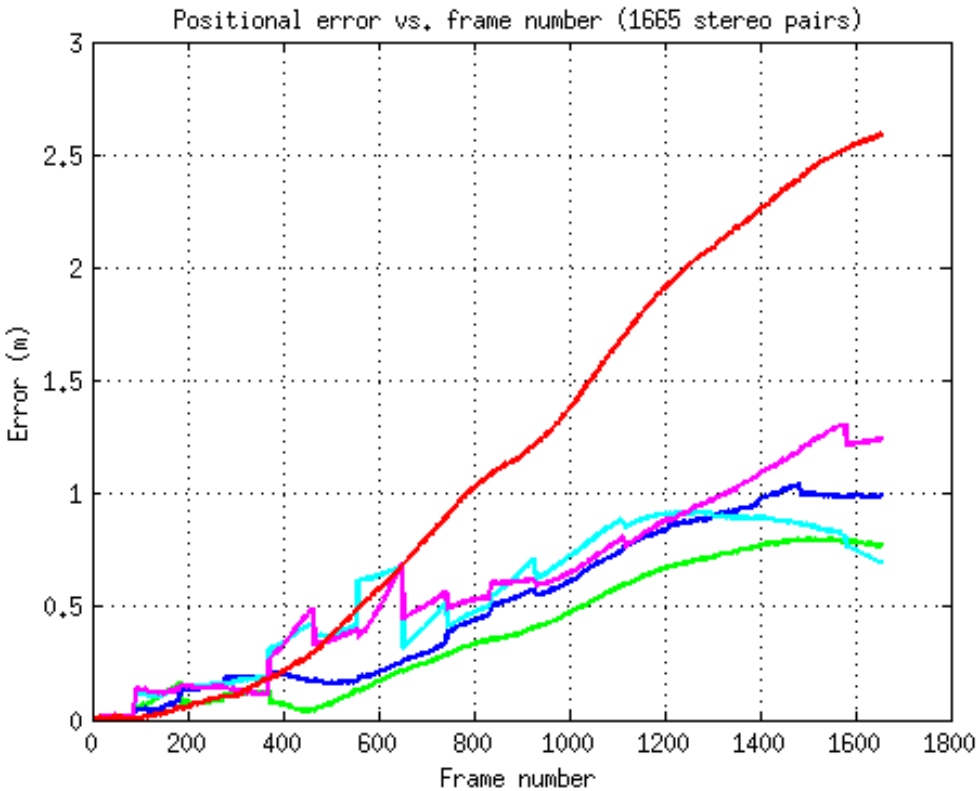

Resolutions
15360 x 15360
11520 x 11520
7680 x 7680
3840 x 3840

- Red is the plain VO positional error in all frames of a 100m traverse
- Green, blue, cyan and magenta are resp. the results of the proposed method in orbital images of the original, 25%, 50% and 75% lower resolutions

Results: rotational accuracy

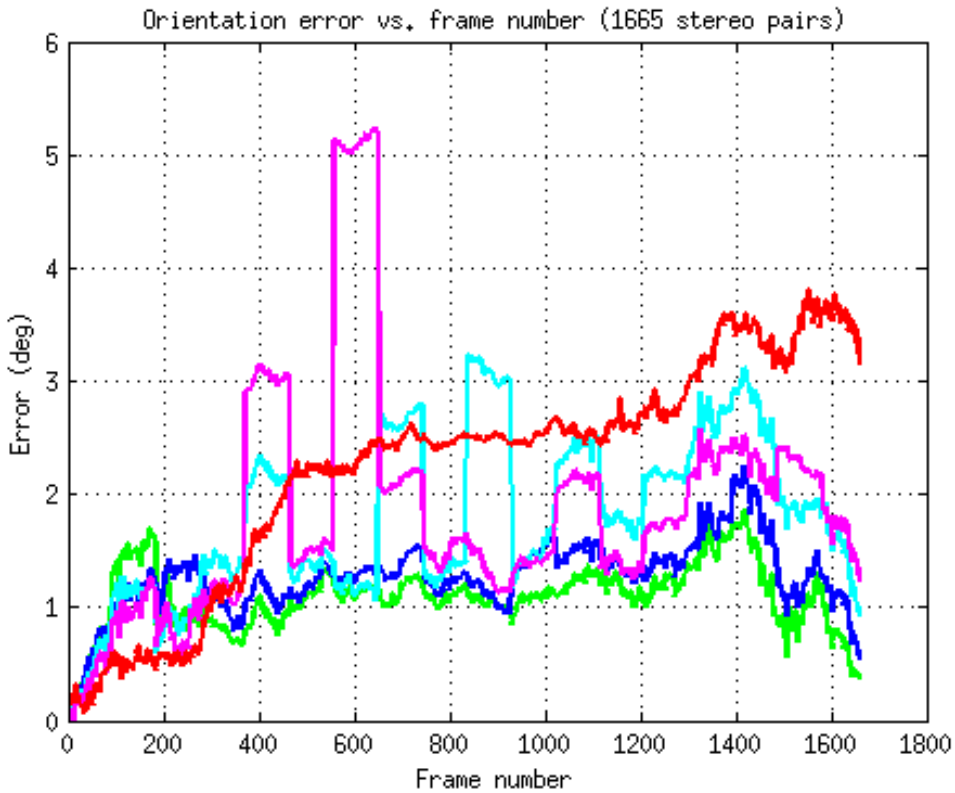

Resolutions
15360 x 15360
11520 x 11520
7680 x 7680
3840 x 3840

- Red is the plain VO rotational error in all frames of a 100m traverse
- Green, blue, cyan and magenta are resp. the results of the proposed method in orbital images of the original, 25%, 50% and 75% lower resolutions

Summary & conclusions

- Improving VO localization using boulders is a promising idea
  - A very terse representation consisting of only approximate 2D locations was shown to bring significant accuracy improvements
  - Same approach can potentially solve the “kidnapped robot” problem (substituting ICP with geometric hashing for matching)
  - More work on boulder detection and matching is needed
  - Testing on more datasets is also necessary
- Accuracy of localization corrections is limited by the resolution of the orbital image; they actually degrade performance at the beginning of the sequence
- Vertical displacements and out-of-plane rotations are not corrected; full 3D info necessary

Thank you

Any questions?

Schematic of the approach

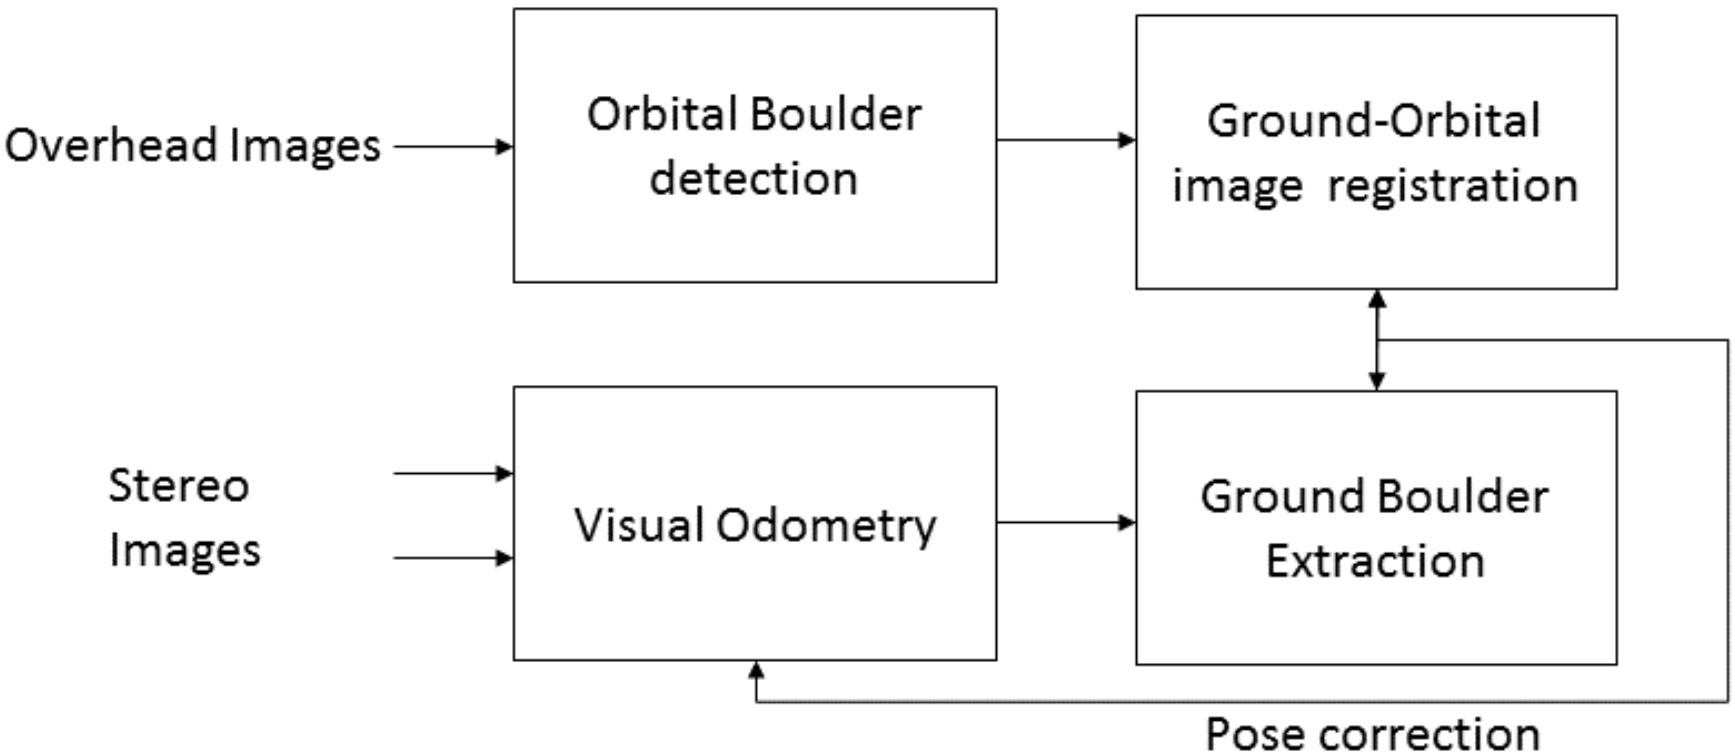

Ground – orbital boulder matching example

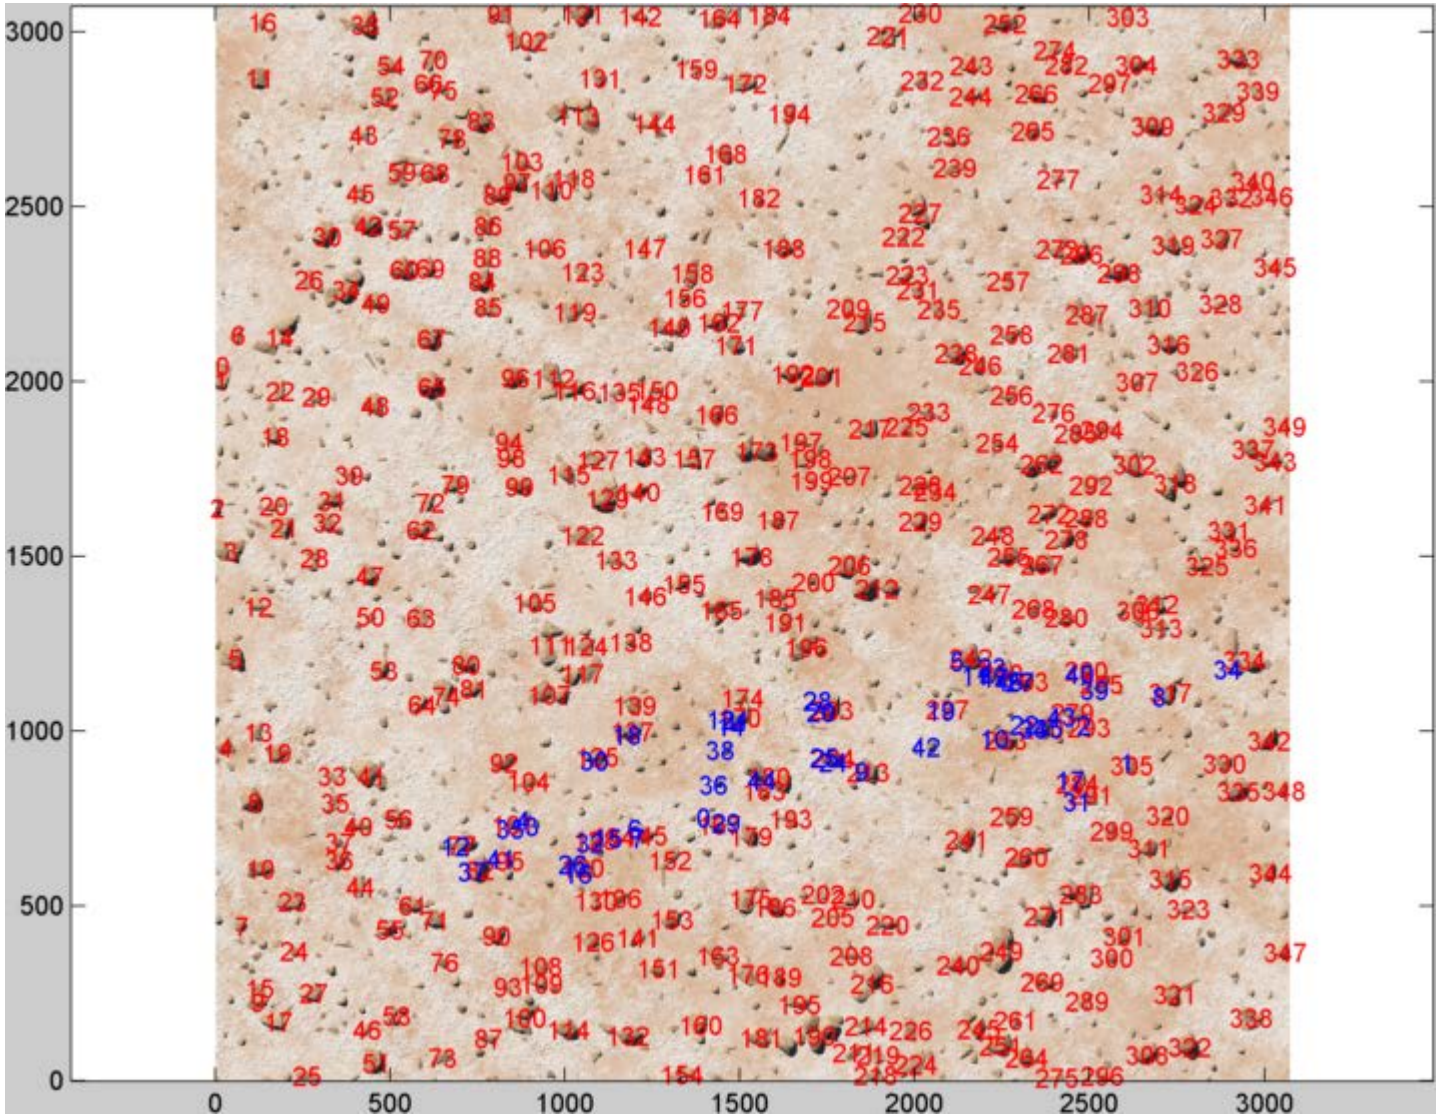

Ground boulders *(blue)* matched against ground ones *(red)*