

Acquisition and Online Display of High-Resolution Backscattered Electron and X-Ray Maps of Meteorite Sections

Ryan C. Ogliore^{1,2}

¹Department of Physics, Washington University in St. Louis

²McDonnell Center for the Space Sciences, Washington University in St. Louis

Key Points:

- Large scanning electron microscopy datasets have been difficult to visualize and share
- New techniques described here allow for the acquisition, stitching, and display of electron and x-ray mosaics
- Multi-gigapixel, multi-channel maps displayed online can be used to compare meteorites to returned asteroid samples

Corresponding author: Ryan Ogliore, rogliore@physics.wustl.edu

Abstract

Analyses of meteorites in thin or thick section begins with a detailed mineralogic/petrologic study of the sample. Backscattered electron and x-ray imaging in a secondary electron microscope is critical for the characterization and study of the meteorite sections at sub- μm to cm size scales. Here, I describe techniques to acquire backscattered electron and x-ray images of an entire one-inch thin or thick section at high resolution, assemble large mosaic maps of the data, and display the maps conveniently online in a web browser. The code to acquire, stitch, and display the maps is made available as an open-source project.

Plain Language Summary

I describe techniques to create a “virtual scanning electron microscope” for scientifically important meteorite samples. Scientists all over the world can view high-resolution images easily and conduct studies without having either the meteorite sample in hand, or access to an electron microscope.

1 Introduction

In recent years there have been increased efforts to make raw scientific data publicly available online. Many astronomy and planetary science data sets are available for any interested scientist to analyze. These open data sets allow for greater transparency of published work and exploration of data in novel ways by people outside of the main community. Open data sets also facilitate a “first look” into an interesting scientific question, with more in-depth studies that follow. Open data sets in cosmochemistry, meteoritics, and astromaterials are more scarce. The different types of laboratory instrumentation (electron beam, ion probe, synchrotron, . . .), techniques, protocols, standardization/calibration, and samples make it difficult to usefully share data in a standard way with other researchers. However, there are a few open data sets that have proven very useful to the community. For example, the presolar grain database (Stephan et al., 2020), has been used extensively by cosmochemists and astronomers to explore a variety of problems.

The “first look” data that is often critical for cosmochemists is a detailed mineralogic and petrographic description of the meteorite sample. For a meteorite prepared as a thin or thick section, backscattered electron (BSE) and x-ray elemental maps acquired at the effective resolution limit for these two modalities (~ 50 nm/pixel and ~ 2 μm /pixel, respectively) in a scanning electron microscope (SEM) are critical for determining if a given sample can answer a given scientific question. However, sections of precious extraterrestrial samples are rare and often difficult to acquire. Each section is unique, and only one scientist can analyze a given sample at one time. In addition, it is expensive to analyze samples using SEM techniques (typical rates are \$25/hour), and many scientists do not have immediate access to an SEM.

BSE and x-ray maps of meteorite section are typically limited to an area of interest less than 1 mm wide. High BSE resolution-imaging over an entire meteoritic section, much more than 1 mm, would facilitate mineralogic/petrologic studies from sub- μm to cm size scales. However, even a well-polished thin section may have a slight tilt, and may not be mounted perfectly flat in the SEM. The depth of focus for high-resolution BSE imaging is tens of μm , and the working distance (the distance between the sample and objective lens) across a 1-inch section varies on scales much larger than this. Therefore, to acquire high-resolution BSE images over the entire sample, the SEM must adjust its focus. Autofocus routines on modern SEMs are slow and must adjust the brightness and contrast to find the focus. Sometimes the brightness/contrast is not set back to the original value accurately. For these reasons, techniques to reliably and efficiently collect high-

resolution BSE and x-ray maps over the entire sample (e.g. a 1-inch round) at high resolution are not readily available.

Once individual BSE images are captured, they must be stitched together to make an image with minimal artifacts. Stitching algorithms using feature recognition, matching with neighboring images, and image warping are commonly used in many applications (Brown & Lowe, 2007; Wang et al., 2017). Image tiles in the mosaic are linked back to some origin image. These types of algorithms work well for dozens of images, but will fail on a large number of images (thousands to tens of thousands) if some images are matched poorly, or if some images lack features. A robust algorithm that can handle outliers is required to reliably stitch together multi-gigapixel images with thousands of tiles.

A multi-gigapixel BSE image, with registered elemental map overlays, is very difficult to view on a computer (the entire image often cannot be loaded into memory). Sharing such data with remote colleagues is even more difficult. Image pyramids, where the single large image is broken down into tiles and resampled at different sizes, is a useful way to view large scientific images (Hayashi et al., 2016). With image pyramids, the computer only needs to load and display the current view at the appropriate resolution.

Here I describe techniques to solve these three problems: 1) acquire high-resolution BSE images (~ 50 nm/pixel) and the associated x-ray maps (~ 2 μ m/pixel) of an entire ~ 1 -inch section, 2) stitch these images into one multi-gigapixel image, and 3) display these maps conveniently online using an image pyramid viewer with overlays.

2 Methods

For the technique described here, I am using a Tescan Mira3 FEG-SEM and EDAX Octane Plus (30 mm² silicon-drift detector) energy-dispersive x-ray system. However, the techniques and code I describe here can be easily converted for use on a different SEM and EDS system. The required SEM and EDS software are protocols that can automatically acquire electron images and x-ray images from a user-defined input text file (where the working distance can be changed). For the Tescan SEM, this requirement is fulfilled with the ImageSnapper function (Mira3 control software, version 4.2.27.0), or with the SharkSEM Python scripting protocol. For the EDAX x-ray system (TEAM Enhanced software, version 4.5.1), this requirement is fulfilled with the Multifield Analysis function.

The code used here is available on GitHub: <https://github.com/ogliore/DeepZoomSEM>. I will refer to the name of the script or function in parentheses in describing the algorithm below.

The meteorites analyzed using this technique include DOM 14305,5 (CO3), DaG 749 (CO3), Aguas Zarcas (CM2), Acfer 094 (C2-ung), Acfer 182 (CH3), Orgueil (CI), and Tarda (CY). These maps can be viewed at: <https://presolar.physics.wustl.edu/meteorite-deep-zoom/>.

2.1 SEM acquisition of electron and x-ray images

The meteorite section is mounted on a large SEM stub with clips or a set screw to ensure the sample does not move during the long acquisition. The SEM is tuned for optimal BSE image acquisition at high magnification and 15–30 kV accelerating voltage. The electron beam current is chosen for a beam spot size to match the pixel size in our final mosaic, typically ~ 50 nm. A working distance of ~ 12 mm is typically used as a balance between backscattered electron signal (which is higher for shorter working distance) and depth-of-focus (which is smaller for shorter working distance). The BSE brightness/contrast and look-up-table gamma value is changed depending on the sample analyzed and the phases of interest.

First, we acquire a “focus map” before the high-resolution BSE acquisition. The user selects ~ 100 points including the perimeter of the sample using the Image Snapper with autofocus enabled. The acquired images are not used, but the header files associated with each image records the optimal focus (working distance) determined by the autofocus method. These working distances are used in a Matlab script (`TescanImageSnapperPoints.m`) to build a focus map (Figure 1). Outliers are removed, then the remaining points are fit to a two-dimensional, second-order polynomial (to account for curvature and tilt of the sample from polishing and mounting) or a two-dimensional interpolated surface.




Figure 1. Working distance, x , and y values for images in the focus map of meteorite thin section Acfer 182 (solid blue points), with second-order polynomial surface curve fit.

The x , y , and working distance values are then interpolated from this surface fit to calculate the coordinates of the full-resolution BSE scan (including a user-defined overlap fraction, $\sim 20\%$). These coordinates are fed into a Matlab function that writes an Image Snapper acquisition file (`writeImageSnapper.m`) for collection of the full-resolution BSE scan. Images are only acquired over the actual sample (defined as the perimeter of points that the user defined in the focus map), minimizing wasted acquisition time. Each individual tile is a 16-bit BSE image in png format, 2048×2048 pixels, $100\text{--}200\text{ }\mu\text{m}$ field of view, and $2\text{--}4\text{ }\mu\text{s}/\text{pixel}$ dwell time.

Next, the BSE brightness, contrast, and gamma is optimized for the particular sample. With auto-working-distance and auto-brightness-contrast disabled, BSE images are acquired over the entire sample using Image Snapper and the acquisition file written in the previous step (acquisition takes ~ 3 days). After acquisition, images are renamed to their locations in the scan grid using a bash script (`MoveTescanImages.sh`).

Following BSE acquisition, the SEM is re-optimized for X-ray acquisition (higher beam current, 15 mm working distance). We acquire a new focus map and write a multifield acquisition file using Matlab for the EDAX TEAM software (`writeEDAXMultifieldMaps.m`). We acquire a 512×400 pixel images over a $1024 \times 800\text{ }\mu\text{m}$

field of view ($2\ \mu\text{m}/\text{pixel}$). We use an amp time of $0.48\ \mu\text{s}$ (which is relatively short) to maximize the x-ray count rate (at the expense of larger sum peaks). We tune the primary beam current to achieve a deadtime of 20%. A typical x-ray spectrum summed over one field of view is shown in Figure 2. It takes 5–10 days to acquire X-ray maps over the entire thin section using the 30-mm^2 SDD Octane x-ray detector.




Figure 2. Typical x-ray spectrum summed over one $512\ \mu\text{m}$ field of view in DOM 14305,5. Major element peaks are labeled, sum peaks are visible at 3–4 keV.

2.2 Stitching of electron image tiles into mosaic

The positions of the individual images in the final mosaic are calculated using a Matlab script (`mosaic_maker.m`). Identifying features in the overlapped regions of neighboring images are found using Matlab’s `detectBRISKFeatures` which uses the Binary Robust Invariant Scalable Keypoints algorithm to detect multi-scale corner features (Figure 3).




Figure 3. Left) Backscattered electron image of Acfer 182. Field of view is $100\ \mu\text{m}$. Right) Same image but with identified BRISK features circled in red, in the region of overlap with the neighboring image.

Identifying features in each image are extracted from only the overlap regions (typically 20%) with the image in the previous row and the image in the previous column. Typically, a hundred features are identified for each image. This function (`computeMatchedPointsMosaicMaker.m`) is run in parallel for each set of images in a row using a Matlab `parfor` loop, as it is computationally intensive. Features are extracted and matched to the previous-row image and previous-column image. The following four geometric transforms are calculated: current image to previous-row image $(i, j) \rightarrow (i-1, j)$ which is defined as $T_{(i,j) \rightarrow (i-1,j)}$, previous-row image to current image $(i-1, j) \rightarrow (i, j)$, current image to previous-column image $(i, j) \rightarrow (i, j-1)$, and previous-column image to current image $(i, j-1) \rightarrow (i, j)$. The geometric transforms are 3×3 matrices. The transform is assumed to be a similarity transform, which allows for translation, rotation, and scaling (straight lines remain straight and parallel lines remain parallel). Similarity was chosen (instead of affine or projective) to minimize distortion of feature shapes, though with the compromise of decreasing the accuracy of the transform. The accuracy of the transform is quantified by the Euclidean distance between the features in the original image and the mapped previous image. If the transform is perfect, this distance is zero for all of the features. The *transform error* is calculated as the mean of the squares of these distances for all identified features. Since we know that the stage should have moved a certain distance between neighboring row and column images (given by the image overlap that we set when acquiring the images on the SEM), we can compute another type of error—the difference between the similarity transform and the expected translational shift. This is the *overlap error*, and may arise from either inaccurate feature matching (for example, between two images covering a single, featureless crystal) or from inconsistent movements of the SEM stage.




Figure 4. A representative BSE image (red) and its previous-column neighbor (turquoise) with their matched features linked by yellow lines. Field of view is $150 \mu\text{m}$.

This image matching calculation for all tiles in the mosaic yields transformation matrices from a given tile to its neighbor, and two estimates of the error of these transformations. To construct the final mosaic, it is necessary to link all images back to an

origin tile that will define the origin of the mosaic coordinate system. Each image in the mosaic is calculated relative to the origin tile, so a path from the origin tile to every tile in the mosaic is needed to assemble the final mosaic. The path from $(0,0)$ to (a,b) may be:

$$(0,0) \rightarrow (0,1) \rightarrow (0,2) \rightarrow (1,2) \rightarrow (1,3) \rightarrow \cdots (a-1,b) \rightarrow (a,b) \quad (1)$$

and the transformation matrix to place the tile image (a,b) into the final mosaic is the matrix product of the individual transformation matrices:

$$T_{a,b} = T_{(a-1,b) \rightarrow (a,b)} \cdots T_{(1,2) \rightarrow (0,1)} T_{(1,3) \rightarrow (1,2)} T_{(0,1) \rightarrow (0,2)} T_{(0,0) \rightarrow (0,1)} \quad (2)$$

The path, however, is not unique. To find the optimal path, a bidirectional graph is constructed. Each node on the graph represents the location of a tile image. Each node is connected to its neighboring node by an edge. Matlab's `shortestpath` function will calculate the shortest path between one node (e.g., the origin tile image) and another node by minimizing the sum of the edge distances between the nodes. The edge distance is defined as a weighted sum of the overlap error and transform error. This will penalize the steps between image tiles that have large errors by increasing the distance. The relative weight between the overlap and the transform error is a user-defined quantity, though equal weight usually works well. Outlier images with large transform or overlap errors, or too few matched features, are assigned infinite weight, so paths through these images are avoided. An example of the paths back to the origin image is shown in Figure 5.



Figure 5. Shortest paths between tile images and the origin tile image (filled black circle) with bidirectional graph edges weighted by the sum of the overlap and transform errors (for DOM 14305,5).

The entire set of $\sim 10,000$ individual tiles are mapped back to an origin tile near the middle of the image (one that has a large number of matched features with its neighbors) via individual paths. The tile positions are then refined for each horizontal and tile pair using an additional affine transform on the mapped matched feature locations. If the refined position between an image pair improves the overall transform error, the new position is kept, else it is discarded. The entire image set is refined ~ 8 times which improves the overall transform error by $\sim 25\%$. The transforms for the outlier images are calculated via two-dimensional interpolation and extrapolation (`inpaint_nans.m`) with the surrounding image transforms.

Each vertical and horizontal neighboring images will have some overlapping areas that can be computed from the locations of their matched features. The average brightness in these overlapping regions should be the same, but changes in the primary electron beam current and surface geometry of the sample may cause the brightness to change. Because an image needs to match brightness in both its vertical and horizontal neighbor images, we employ an iterative algorithm to adjust the brightness in each image. We iterate the brightness adjustments for each image until the summed differences in the mean brightness of the overlapping images changes below some threshold. This brightness correction is calculated with a Matlab function (`globalbalance2.m`) that is called after the similarity transform calculation.

The image is subdivided into row strips with fixed vertical boundaries, to facilitate faster stitching. The distance between these boundaries is set equal to the image height. The Matlab script (`mosaic_maker.m`) calculates which images have overlap into the boundaries for each row.

A comma-separated text file is written where each line is a tile image filename, the components of the transformation matrix for that image, and the brightness correction. The text file is separated into the rows described previously, so that each row can be stitched independently. This is the last step of the Matlab script `mosaic_maker.m`. The script will take a few hours to run for a ~ 150 gigapixel mosaic.




Figure 6. Mapping of x-ray BSE image (turquoise) to high-resolution BSE image (red) of DaG749, with their matched features linked by yellow curves.

The image transformation text file written by Matlab is then fed to a Python script to actually stitch the final 100+ gigapixel mosaic. The similarity transform, brightness correction, and compositing of images is done with pyvips—the Python implementation of vips (Martinez & Cupitt, 2005), a demand-driven, horizontally threaded image processing library, to apply the transformation and brightness corrections, and assemble the individual image tiles.

A row of images is stitched using the Python script `affinetransform6r.py` which employs the vips composite function. Since each row is independent of the other rows, the rows can be stitched simultaneously as parallel processes. From a bash script (`allrows.sh`), a Python process is spawned to stitch one row. The row is written to disk in the vips image format. The number of parallel processes spawned is equal to the number of logical cores on the machine (16 or more on modern CPUs). (Vips does not run efficiently on the GPU). When all rows have been written to disk, the final image is assembled from the row images using the vips command `arrayjoin` (`affinetransform6f.py`). The final image is written as a 10 megapixel thumbnail image, and as a full-resolution image pyramid in the dzi (deep zoom image) format. The dzi image pyramid is ~ 150 gigabytes in size. The stitching process takes about a day, but depends on the size of the mosaic and the speed and number of logical cores.

The resulting stitched BSE image mosaic has very few stitching artifacts. However, the choice of using a similarity transform instead of affine or projective transform preserves shapes in the BSE images at the expense of transform errors which can result in some stitching artifacts. This choice was made to preserve the fidelity of the BSE images, where shapes have important scientific meaning. Stitching artifacts are most apparent far from the origin image, at the edges of the sample. Shading artifacts are also minimized by the brightness correction algorithm.

2.3 Assembly and Registration of X-ray Maps

X-ray maps are saved by the EDAX TEAM software as data cubes in the spd file format. The intensities at each pixel are extracted for elements of interest using the Matlab function `processEDAXmaps.m` and script `assemble_xray_maps_AguasZarcas.m`. The widths and locations of the peaks for each element are estimated by interpolation using the most abundant elements in the summed spectrum (Figure 2). For each pixel, the intensity of each element is calculated from the sum of all x-ray counts at the peak energy plus or minus two estimated peak widths. The background counts are estimated from a neighborhood between two and three widths from the peak, and are subtracted from the total intensity. Each element map is normalized by the total summed x-ray counts in the data cube. Backscattered electron images are acquired simultaneously and saved as separate image files during collection of the x-ray maps.

X-ray maps are acquired at $2\ \mu\text{m}/\text{pixel}$, which is 40 times larger than the electron images. This results in many fewer stage movements and less distortion of the final assembled image. For this reason, the BSE image that is acquired along with the x-ray maps is assembled into a mosaic based on the stage position. The x-ray maps and associated BSE image are not stitched using feature-matching. This (x-ray collected) BSE image is then registered to the thumbnail of the high-resolution BSE image described previously using a projective transformation (Figure 6, `transform_edx_to_bse_AguasZarcas.m`). The x-ray BSE image was acquired with different SEM conditions, including a much higher beam current, so it can be distorted compared to the high-resolution BSE image. The same projective transformation is then applied to the assembled X-ray maps for each element so that they are warped to align with the high-resolution BSE map. Pre-compiled RGB maps, such as Ca-Al-Si or Fe-S-O, are also created and saved. Histogram equalization for each channel or other image adjustments are applied as needed. The magma colormap is applied to each element map (Figure 7) to facilitate a larger visual dynamic




Figure 7. Magnesium x-ray map of DOM 14305,5 with magma colormap.

range (Nuñez et al., 2018). A red, green, and blue color map is also applied to each element so that users can create their own RGB colormaps in the web-based viewer, described below. All of these maps are saved as image pyramids in dzi format.

2.4 Online Display of Electron and X-ray Maps

The dzi image pyramids are uploaded to a web server. For each sample, a web page is created to view the BSE and x-ray dzi pyramids. The webpage template needs only minimal customization for each sample: nm/pixel of the full-resolution BSE image, ratio of the BSE and x-ray image sizes, and the sample name. The javascript library OpenSeaDragon is used to display the images with seamless panning and zooming. Buttons are added to the OpenSeaDragon viewer to allow the user to switch between BSE and x-ray maps. A scalebar is overlaid in the lower-left corner. The viewer can be made fullscreen, flipped, or rotated. Each field of view has a unique URL that can be shared with collaborators. The user can save a high-resolution screenshot of the current field-of-view.

An important feature in the OpenSeaDragon viewer is the ability for the user to create a custom x-ray RGB map. The user clicks the “RGB” button then the maps to be in the red, green, and blue channels (the images are composited to the canvas using the OpenSeaDragon composite operation “lighten”). This map is then saved as a button, and the user can toggle between this map, the individual element maps, and the BSE image.




Figure 8. BSE image (top) and Fe-Si-Mg RGB x-ray image (bottom) of DOM 14305,5.

3 Future Improvements

The software presented here is mature, but some improvements are planned for the future:

- **Simultaneous acquisition of secondary and backscattered electron images.** Some samples polish poorly, or have surface contamination, so it is helpful to view secondary electron images alongside backscattered electron images. The Orgueil meteorite is one such sample. In this map: https://presolar.physics.wustl.edu/maps/Orgueil_X.html we acquired secondary electrons after backscattered electrons. However, they can be acquired simultaneously in batch mode using the SharkSEM Python scripting protocol.
- **Image adjustments** Standard image adjustments (such as brightness, contrast, and gamma) using sliders would be useful to bring out shadow details in electron and x-ray images. The NIST OpenSeaDragon image adjustments plugin can be used for this purpose.
- **Faster x-ray map acquisition** X-ray image acquisition takes about a week, but can be reduced to two days with a large 100 mm² SDD x-ray system.
- **Feature-matching and stitching of the x-ray images** Currently the x-ray images are laid out in a grid, not stitched, which results in some artifacts at the edges of the x-ray image tiles. The x-ray images could be stitched using the simultaneously acquired BSE images to compute the transforms.
- **Other imaging modalities** Other imaging modalities can be added to the OpenSeaDragon viewer, as long as a BSE map can be acquired simultaneously for registration purposes. For example, cathodoluminescence would be important for certain samples.
- **Quantitative x-ray analyses** The x-ray intensity maps are useful for qualitative studies but cannot be used to determine, for example, the Fe/Mg ratio in a chondrule olivine. Energy dispersive spectroscopy can be made quantitative with appropriate standard analyses (Newbury & Ritchie, 2013). For the maps presented here, the user could click a button that would look up the original x-ray data cube file for that field of view, and associated standard spectra. The user could then select to download these files. Alternatively, the user could circle an area, and the spectrum from the encircled pixels would be summed and displayed as a spectrum. This would be relatively straightforward in javascript, and would require more space on the webserver.

4 Conclusions

I have presented techniques and software for an online “virtual SEM” of meteorite thin sections. Data acquisition is lengthy, 1–2 weeks (determined by the x-ray acquisition time), but it is unsupervised and requires only minimal setup. Data processing is also mostly unsupervised. The investment in effort and SEM time is worthwhile for important samples where mm- to μ m-scale mineralogic context is critical for targeting subsequent *in-situ* micro-analyses such as FIB-TEM and SIMS. The ability to easily share the detailed mineralogy of a sample with colleagues possessing different expertise is invaluable.

My lab has used these maps in a number of projects. We have identified cosmic symplectite in Acfer 094, searched for cosmic symplectite in other meteorites, located carbonaceous-chondrite-like clasts in the howardite Kapoeta (Liu et al., 2020), and identified an unusual Ti-rich sulfide mineral in Acfer 182 (CH3) that would have been nearly impossible to find without these maps.

Many scientists all over the world have limited access to an SEM, and cannot easily acquire interesting meteorite samples. A catalog of these high-resolution maps can

allow for any interested scientist to perform basic mineralogy/petrology of meteorite samples at various size scales.

The algorithms and software presented here can play an important curation role for the next generation of returned samples. Hayabusa2 returned mm-sized and larger stones from Ryugu that may possibly be prepared in thin section. OSIRIS-REx will likely return stones from Bennu at least this large that may also allow for thin sections. The characterization of Ryugu and Bennu samples with this technique will allow for researchers all over the world to simultaneously analyze thin sections and target particular samples for more detailed analysis.

One of the immediate questions concerning the returned Bennu and Ryugu samples will be: is this material similar to any known meteorites? Many of the most important studies of the comet Wild 2 samples returned by NASA's Stardust mission have been comparative studies with meteorites (Frank et al., 2014). Comparative asteroid mineralogy is most efficiently done with open meteorite data sets. The closest analogs to Bennu and Ryugu will likely be the CI, CM, and CY chondrites (Hamilton et al., 2019). With a collection of 20–30 publicly accessible, high-resolution BSE/EDS maps of appropriate analogous meteorites, covering a range of petrologic type, comparative mineralogy can be done by scientists around the world without needing all these samples in hand.

Acknowledgments

The author thanks Lionel Vacher (Washington University in St. Louis) for helpful discussions. Datasets for this research are available at <https://presolar.physics.wustl.edu/meteorite-deep-zoom/> under Creative Commons v 4.0 CC BY-SA license.

References

- Brown, M., & Lowe, D. G. (2007). Automatic panoramic image stitching using invariant features. *International Journal of Computer Vision*, 74(1), 59–73.
- Frank, D. R., Zolensky, M. E., & Le, L. (2014). Olivine in terminal particles of stardust aerogel tracks and analogous grains in chondrite matrix. *Geochimica et Cosmochimica Acta*, 142, 240–259.
- Hamilton, V., Simon, A., Christensen, P., Reuter, D., Clark, B., Barucci, M., . . . others (2019). Evidence for widespread hydrated minerals on asteroid (101955) bennu. *Nature Astronomy*, 3(4), 332–340.
- Hayashi, S., Gopu, A., Kotulla, R., & Young, M. D. (2016). Imagex: New and improved image explorer for astronomical images and beyond. In *Software and cyberinfrastructure for astronomy iv* (Vol. 9913, p. 99134E).
- Liu, N., Ogliore, R. C., & Vacher, L. G. (2020). Nanosims isotopic investigation of xenolithic carbonaceous clasts from the kapoeta howardite. *Geochimica et Cosmochimica Acta*, 283, 243–264.
- Martinez, K., & Cupitt, J. (2005). Vips-a highly tuned image processing software architecture. In *Ieee international conference on image processing 2005* (Vol. 2, pp. II–574).
- Newbury, D. E., & Ritchie, N. W. (2013). Is scanning electron microscopy/energy dispersive x-ray spectrometry (sem/eds) quantitative? *Scanning*, 35(3), 141–168.
- Núñez, J. R., Anderton, C. R., & Renslow, R. S. (2018). Optimizing colormaps with consideration for color vision deficiency to enable accurate interpretation of scientific data. *PloS one*, 13(7), e0199239.
- Stephan, T., Bose, M., Boujibar, A., Davis, A., Dory, C., Gyngard, F., . . . others (2020). The presolar grain database reloaded-silicon carbide. *LPI(2326)*, 2140.
- Wang, M., Niu, S., & Yang, X. (2017). A novel panoramic image stitching algorithm based on orb. In *2017 international conference on applied system inno-*

vation (icasi) (pp. 818–821).