



Improved segmentation of meteorite micro-CT images using local histograms

L.D. Griffin ^{a,*}, P. Elangovan ^{b,1}, A. Mundell ^a, D.C. Hezel ^{b,2}

^a Computer Science, University College London, London WC1E 6BT, UK

^b Natural History Museum, Department of Mineralogy, Cromwell Road, SW7 5BD London, UK

ARTICLE INFO

Article history:

Received 4 March 2011

Received in revised form

24 June 2011

Accepted 4 July 2011

Available online 22 July 2011

Keywords:

Meteorites

Micro-CT

Ground truth

Image segmentation

Local histograms

Pixel classification

ABSTRACT

In micro-CT images of meteorites individual components such as matrix, chondrules, Ca,Al-rich inclusions (CAIs), and opaque phases (metal and sulfide) are visually distinguishable. Automated classification of the components is desirable to deal with the large amount of data in a 3-D CT image. Classification by pixel intensity achieves a performance only 25% of the way from baseline to perfect. The poor performance is explained by an overlap in the range of intensities present in the different components. An improved method of semiautomated classification is presented, based on local histograms of the intensity. This achieves a performance 60% of the way from baseline to perfect.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

X-ray computerized tomography (CT) is a radiographic imaging modality that produces 3-D digital images of volumetric objects. The 3-D images are reconstructed from multiple 2-D images, each acquired by projecting X-rays through the object at a different orientation. X-ray CT images are composed of scalar-valued pixels (called voxels when the 3-D nature is important to stress) each of which records the reconstructed X-ray transmissibility of a location within the object. X-ray CT is widely used in medical and biological fields. In medicine the images encompass entire human organs, and have voxel dimensions around 1 mm. Micro-CT is the term used for smaller scale imaging setups, typically with a field-of-view dimension around 1 cm, and a voxel dimension of down to 1 μ m. Micro-CT has already been extensively applied in medicine, biological sciences, materials science, engineering, and bio-engineering and has been recently extended to the field of geosciences and meteoritics (Van Geet et al., 2001; Ketcham and Carlson, 2001; Ebel and Rivers, 2007; Spinsby et al., 2008, Friedrich, 2008; Friedrich et al., 2008a, 2008b). When applied to rock samples the pixel values are determined by the

proportions and distribution of different mineral types, and by porosity.

Asteroids are the building blocks of planets; hence, they predate the formation of Earth. Meteorites are fragments from asteroids that orbit the sun between Mars and Jupiter. One group of meteorites, called “primitive chondrites,” contain micrometer- to millimeter-sized components that formed in the protoplanetary disk of our Solar System. Study of these components is informative about the formation and evolution of the early Solar System before planet formation. Primitive chondrites contain the following: (i) ~20–80% by volume of micrometer- to millimeter-sized, roundish “chondrules” primarily consisting of silicate minerals; (ii) ~5–80% by volume of matrix primarily formed of micrometer-sized silicates; (iii) 0.1–3% by volume of Ca,Al-rich inclusions consisting of silicates and oxides; (iv) sulfides and metals (collectively opaques) that occur as minor abundances in chondrules and matrix (Brearley and Jones, 1998; McSween, 1999; Hezel et al., 2008, 2010). These components formed separately under different conditions in the nebula, and were then aggregated in the meteorite parent bodies, some of which survived as asteroids.

Information on meteorite internal structure, such as proportions, spatial distribution, and relative orientation of components, is pertinent to a number of important problems in meteoritics (e.g., Ebel and Rivers, 2007). This information was for a long time accessible only through slow, difficult, and destructive methods. Since micro-CT can image the interior of a meteorite it has the potential to nondestructively access the required 3-D information (e.g., Ebel et al., 2008; Friedrich et al., 2008a, 2008b; Uesugi et al., 2010). Whether this

* Corresponding author.

E-mail address: L.Griffin@cs.ucl.ac.uk (L.D. Griffin).

¹ Current address: University of Surrey, Department of Electronic Engineering, Guildford, GU2 7XH, Surrey, UK.

² Current address: University of Cologne, Department of Geology and Mineralogy Zulpicher Str. 49b, 50674 Köln, Germany.

potential can be realized depends on whether the different components of the meteorite can be accurately distinguished in the X-ray tomographic image. The experience of observers is that the components can be visually distinguished by careful inspection. An example is shown in Fig. 1.

Since the components are visually distinguishable, manual segmentations, such as that shown in Fig. 1, are possible but they are laborious. A single slice will take between 1 and 10 min even with bespoke software, a graphics tablet, and an experienced observer. A 3-D micro-CT image will typically have thousands of slices, making full manual segmentation of 3-D micro-CT images impractical.

The fact that manual segmentation is possible implies that there is information available in the images that could be used to automate the segmentation task. Many methods of automated image segmentation have been developed. Different methods are

effective for different tasks. Edge-based methods focus on detecting particular image structures, for example, intensity discontinuities, that signal the boundaries of regions; while region-based methods focus on delineating regions with homogeneous textural properties (Haralick and Shapiro, 1985). Model-based approaches that make use of a strong prior expectation of the geometry of the desired segmentation (McInerney and Terzopoulos, 1996), while pixel classification approaches impose no geometric constraints on the produced segmentation, but simply consider each pixel in turn and independently classify it according to features computed from the local neighborhood (Harsanyi and Chang, 1994; Phung et al., 2005).

Pixel classification approaches are the simplest and fastest, and so are preferred if adequate to the task. The simplest feature that can be used as the basis for the classification is pixel intensity. There are three reasons why this feature would be

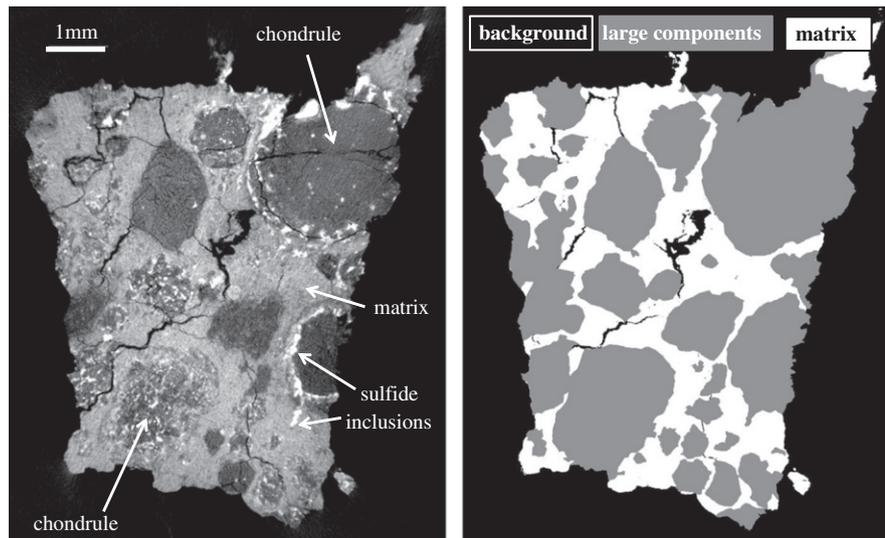


Fig. 1. On the left is a slice (approximately 1200^2 pixels) from a micro-CT image of the Mokoia meteorite. The original range of pixels values was [1883, 35,611] but this has been clipped to [6500, 19,000] to improve visibility. The meteorite can be seen to consist of a collection of dark convex large components bound together by a lighter matrix. The bright spots, blobs, and lines are the (X-ray opaque) sulfide component of this meteorite. They are found within both large components and matrix. On the right is a ground truth segmentation produced by visual inspection and manual pixel painting. The segmentation is into the classes background, matrix, and large components.

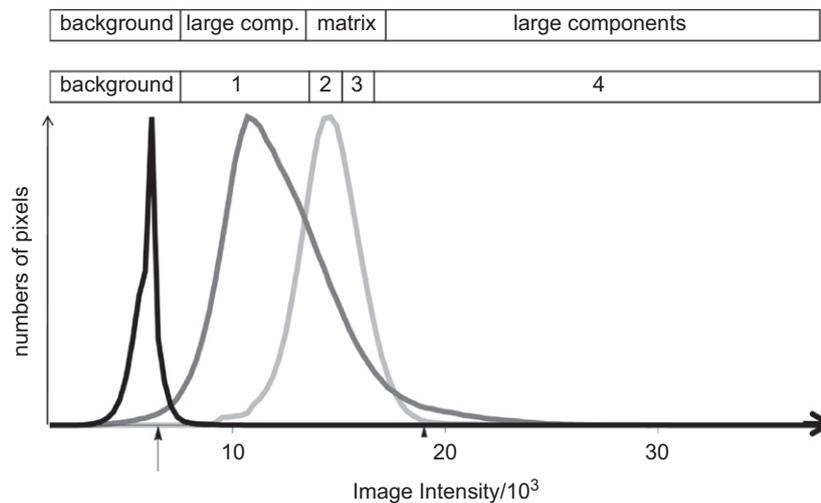


Fig. 2. At the bottom are shown the distribution of intensities that occur in each of the three components in the two slices of Mokoia meteorite data that we have analyzed. Reading the peaks left to right, the distributions correspond to background, large component, and matrix. The arrows on the axis indicate the clipping values that we have used when reproducing the image data in Figs. 1 and 3. The bars at the top of the figure illustrate some of the key parameters in the segmentation methods we have explored. The top bar shows how different intensities are classified in the optimum pixel-intensity-based classification scheme. The lower bar shows the system of four bins that are used in the optimum local histogram-based classification scheme.

suspected to be of little effectiveness for segmenting meteorite components in X-ray CT. First, in a simplified idealized meteorite, each component would exhibit a unique X-ray absorbance, and would therefore give rise to a unique image intensity. However, even with such an idealized meteorite, the image would exhibit a greater variety of pixel values than one per component due to the *partial volume effect*, which occurs when the spatial extent of a voxel contains a mixture of different components. Second, each component of a real meteorite exhibits a range of X-ray absorbances due to variations in mineral content, composition, and porosity. Crucially, the ranges for the different components overlap significantly, as shown in Fig. 2. Third, errors (known as imaging artifacts) in reconstructing the 3-D X-ray image from the 2-D X-ray projection images may lessen the accuracy with which image voxel intensities record the X-ray absorptions of locations in the sample (Remeysen and Swennen, 2006).

It is the aim of this study to present a new technique that in the long run and in combination with other techniques will allow accurate semiautomated component segmentation. As a first step we focus on segmenting matrix from what we will call “large components.” These comprise chondrules, CAIs, and dark inclusions. Chondrules are by far the most abundant of the large components. Large components are very similar in their densities and, hence, other techniques than we described here are required to further segment the large components. Opaques are not considered, as these are easily thresholded out from silicate and oxid material. We refer to the method as semiautomated, rather than automated, as the algorithms parameters need to be tuned on a subset of the data for any given meteorite.

This article is structured as follows. In Section 2 we introduce the image and ground truth data we use to assess different methods of semiautomated classification. In Section 3 we describe these methods. In Section 4 we report results on their effectiveness.

In Section 5 we give a summary description of the winning method (local histogram-based classification) in a form suitable for reimplementation. In Section 6 conclusions are drawn.

2. Materials

2.1. Image data

The micro-CT images were taken with a Metris X-Tek HMX ST 225 scanner at the Natural History Museum, London. The polychromatic X-ray beam was produced by a Pb target. Operating conditions of the X-ray beam were around 200 kV and 160 mA. Images were recorded on 2024 × 2024 wide panel with a 16-bit gray scale. A shading correction was applied to correct for streak and ring artifacts. The X-Tek software was used to correct for beam hardening. For this study we scanned a piece of the Mokoia (CV3) meteorite. The resolution of the resulting 3-D image is 7 μm. We choose two slices from the 3-D image, shown in Figs. 1 and 3. The slices were chosen for the wide variation in shape, size, and texture of large components, matrix, and opaques that they exhibit. The slices are far apart in the 3-D stack, and so represent effectively independent samples from this meteorite.

2.2. Ground truth data

Ground truth data were generated for the two slices by visual inspection and manual painting using the software tool ImageJ (Abramoff et al., 2004). Ground truth data are always subjective to a certain extent, but it has been demonstrated that differences in ground truthing in this problem domain are usually not extensive (Ebel et al., 2009). The data were segmented into three classes: background, matrix, and large components. Thirty-five percent of

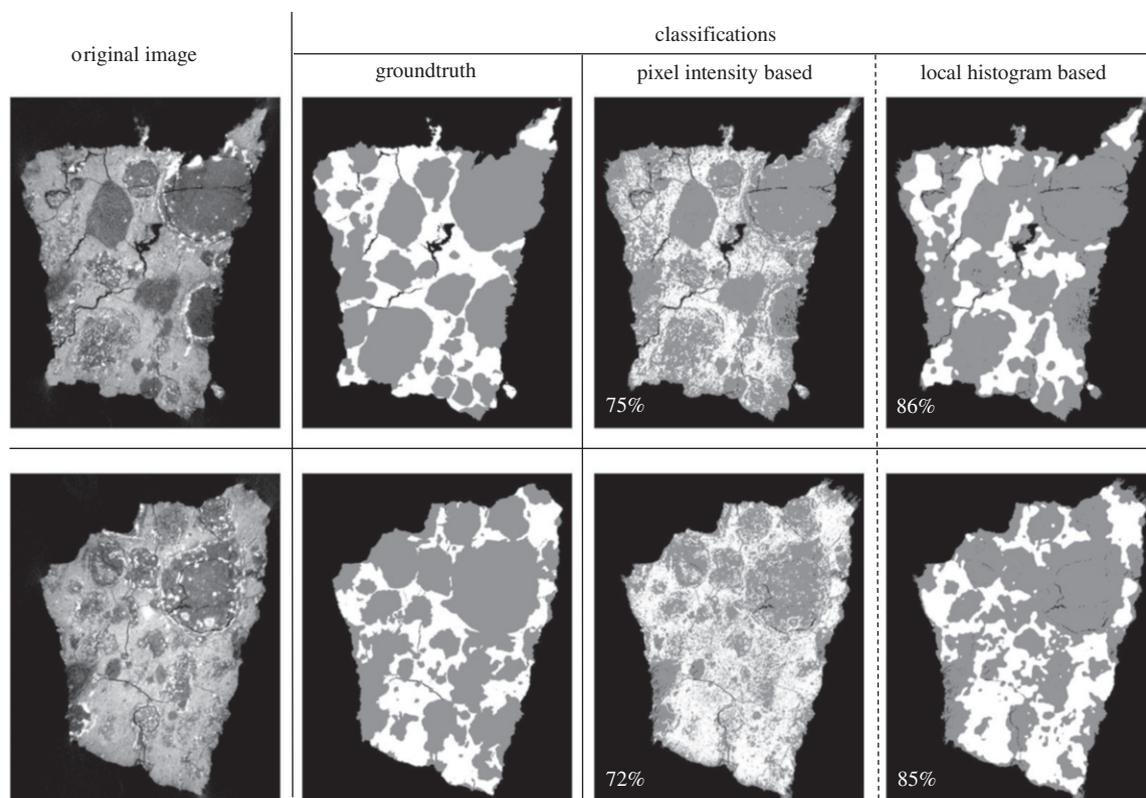


Fig. 3. The two rows correspond to the two slices used in this study; the top row is the same slice as in Fig. 1. The original slices, using an intensity clipped range as in Fig. 1, are shown in the left column. The second column shows the ground truth labeling. The third and fourth columns show labeling by algorithm.

the area of the meteorites was classified as matrix, 65% as large components. Sulfide inclusions in the meteorite are evident as bright lines, specks, and blobs, and are seen to occur in matrix and in large components. As pointed out above, sulfides were not separately labeled, but were instead grouped in with the component in which they were found. Based on the ground truth labelings we were able to compute histograms of the intensities found in each component—these are shown in Fig. 2. The mean intensities of background, meteorite, matrix, and large components were 5897, 13,074, 14,546, and 12,285, respectively.

3. Methods

3.1. Segmentation of meteorite from background

Meteorite was distinguished from background in a two-step procedure. First, pixels with intensity above a threshold were marked as candidate meteorite. Next, the candidate meteorite pixels were grouped into 4-connected clusters. Clusters below 1000 pixels in area were discarded as due to artifacts in the background. The accuracy of the resulting segmentation was determined by comparison with the ground truth, and quantified by computation of overlap. Overlap is the area of the intersection of the algorithmic and ground truth segmentations as a fraction of the area of their union (Crum et al., 2006). The threshold used in the segmentation was optimized to maximize this overlap score.

3.2. Segmentation of components using pixel intensity

As can be seen from the histograms in Fig. 2, mean large component intensity is less than mean matrix intensity. So the simplest way of using the pixel intensity feature would be to classify pixels with intensity below a threshold as large component, and above as matrix. We assessed this approach. Further inspection of the histograms in Fig. 2 shows that the range of matrix intensities is a subset of the range of large component intensities. This suggests the application of a two-threshold approach—pixels with intensities between the two thresholds are classified as matrix, otherwise as large component. We assessed this approach, and some systems based on three or more thresholds. In all cases we tuned the values of the thresholds by maximizing the accuracy of the resulting classifications as compared to the ground truth. Accuracy was quantified using a *percentage-correct score*. This score considers only pixels that were identified as meteorite in the ground truth labeling and in the algorithmic identification of the meteorite described in 3.1. Of those pixels, the score expresses the fraction of them for which the algorithmic classification agrees with the ground truth classification.

3.3. Segmentation of components using local histograms

We hypothesize that a local histogram of pixel intensities (Griffin, 1997; Koenderink, 1999; Shen, 2007) would be a better feature for classification into components than simple pixel intensity. In particular we expect coarsely binned local histograms to perform the best.

We compute local histograms as follows. First we make a $\{0,1\}$ -valued *bin-image* for each intensity bin in the histograms we aim to construct. A '1' in these maps indicates that (i) the location was identified as within the meteorite in Section 3.1 and (ii) the intensity of the corresponding pixel within the image lies in the range of the bin. We then blur each of the bin-images with an *aperture function*, which is nonnegative with unit integral. The list of blurred bin-image values at a location will sum to unity if the location is sufficiently within the interior of the meteorite, to less

than unity if near its border, in which case we normalize the numbers so that they do sum to unity. The list of blurred bin-image values (possibly normalized) at a location is then the local histogram of that location—where *local* is that part of the meteorite picked out by the aperture function. Following previous practice, we use a fuzzy aperture: a sampled Gaussian kernel of standard deviation σ_H (Griffin, 1997; Koenderink, 1999).

Let b be the number of histogram bins. Each raw histogram can be regarded as a point in b -D space. Since the bin values always have unit sum, the points will lie in a $(b-1)$ -D linear subspace. However, following previous practice (Griffin, 2006; Crosier and Griffin, 2010) we work, not with raw histograms, but with *rooted histograms*, obtained by computing the square root of each bin count of a raw histogram. The effects of the square-rooting transformation are (i) the distribution of values for each bin tends to become more normally distributed, and (ii) the corresponding b -D points no longer lie in a degenerate linear subspace, instead they lie on the unit b -dimensional sphere. Both effects are convenient as we proceed by approximating the distribution of rooted histograms for a meteorite component with a multivariate normal distribution $(N_{\vec{\mu}, \Sigma})$. The mean $(\vec{\mu})$ and covariance (Σ) of the best approximation are easily found by computing the mean and covariance of the rooted histograms.

We use one “training” image slice together with its ground truth labels to establish the parameters $\langle \vec{\mu}_M, \Sigma_M \rangle, \langle \vec{\mu}_C, \Sigma_C \rangle$ of the approximating distributions of the matrix and large component classes, and to count the number of pixels in each class n_M, n_C . We then use these approximations to classify the pixels of the training slice. For each location we use the rooted local histogram (\vec{h}) to classify according to which is the larger of $n_M \cdot N_{\vec{\mu}_M, \Sigma_M}(\vec{h})$

and $n_C \cdot N_{\vec{\mu}_C, \Sigma_C}(\vec{h})$ (Jain et al., 2000). By maximizing the percentage-correct score, derived by comparing the classifications to the ground truth labels, we tune the aperture size σ_H and the bin boundaries. Finally, those parameters together with n_M, n_C and $\langle \vec{\mu}_M, \Sigma_M \rangle, \langle \vec{\mu}_C, \Sigma_C \rangle$ are used to classify the other unseen “test” slice in order to obtain unbiased estimates of performance.

Superior approximations of the distribution of rooted histograms for a component can be constructed using a weighted sum of multiple multivariate normal distributions with different means and covariances, rather than a single distribution as described above. Such Gaussian mixture models (Permuter et al., 2006) can be fitted using the EM algorithm (Dempster et al., 1977; Petrou and Petrou, 2010).

We have compared performance using between three and five bins in the local histograms, and one or two multivariate normal distributions in the approximation of the distributions of rooted histograms. For each combination we optimize the aperture size σ_H and the bin boundaries.

3.4. Segmentation of components using local intensity

We wanted to assess whether any advantage that we found for using local histograms was actually due to the use of histograms, or simply due to the use of a local-based rather than a pixel-based feature. The simplest local feature is local image intensity.

Local image intensity can be measured using a localized linear filter. Following previous practice we use a Gaussian filter of standard deviation σ_L for such measurement (Koenderink, 1984). Pixel intensity is thus a limiting case of local intensity, where σ_L approaches zero. Near the edge of the meteorite, simple local intensity will be strongly influenced by the background. To prevent this we (i) set background pixels to value 0, (ii) convolve the result with a Gaussian filter of standard deviation σ ,

(iii) create a {0.1}-valued image indicating meteorite pixels, (iv) convolve the binary image with the same Gaussian filter, and (v) divide the output of (ii) by the output of (iv). The result of this is an image of local image intensity within and restricted to the meteorite.

We have used local image intensity as a feature for classification. As in Section 2.2 we experimented with different numbers of thresholds, and their values, but also with different values of σ_L . We optimized these parameters by maximizing the percentage-correct scores.

4. Results

When tuning algorithms, performance on the data used to tune will overestimate general performance. To prevent this error, we have used the standard methodology of tuning on one dataset and evaluating on a second. In particular we tune each algorithm on slice 1, and test on slice 2—and vice versa. The performance scores we will give are the average of these two test scores. When optimum values of parameters are indicated these are the mean of the tuned values for slice 1 and slice 2.

The results are presented visually in Fig. 3 and in overview in Table 1; extra detail is given in the following subsections, whose numbering corresponds with that in Section 3. The table gives scores for each separately and an overall mean. The differences between the per slice score are always modest, indicating stability of the algorithms evaluated. The table includes, as baseline, the approach to classification of labeling everything as large component; since there is more large component than matrix, this approach scores more than 50%. Using this baseline, and the assumption that a perfect score would be 100% agreement with ground truth, we can place the performance of each algorithm along the spectrum from baseline to perfection, as shown in the third column. The figures in the table show that histogram-based classification is more than twice as far along this spectrum than pixel-based classification. This is reinforced by Fig. 3, which moreover shows that histogram-based classification produces more large component-like structures than pixel-based.

4.1. Segmentation of meteorite from background

Using simple thresholding alone to separate the meteorite from the background achieves an overlap score of 97.2%. When additionally small regions are eliminated using connected component analysis, the score rises to 97.6%. The optimal threshold found is indicated in both bars of the top part of Fig. 2.

4.2. Segmentation of components using pixel intensity

Using a single threshold to classify pixels as matrix or large component achieves a percentage-correct score of 71.8%. Using a pair of thresholds the score rises to 73.4%. No further increase

comes from using more thresholds. The optimal pair of thresholds found is indicated in the top bar of the top part of Fig. 2.

4.3. Segmentation of components using local histograms

Using three-bin histograms achieves a percentage-correct score of 83.7%; this increases to 85.8% for four-bin histograms; no further increase comes from using more bins. The optimal system of four histogram bins is indicated in the second row of the upper part of Fig. 2. The optimum size for the histogram aperture was a standard deviation of $\sigma_H = 7.5$ pixels, which is 0.6% of the linear dimension of the slices shown in Figs. 1 and 3. Representing the distribution of rooted histograms for either component using a mixture of two multivariate normal distributions did not produce a worthwhile improvement over and above using just a single multivariate normal distribution.

The score of 85.8% for local histogram-based classification is 60% of the way from baseline to perfection, and represents a substantial improvement over the pixel intensity-based approach, which scored 73.4%, which is only 25% of the way. The results obtained using histograms are also qualitatively different from those using pixels: as can be seen in Fig. 3, the segmentations are lower frequency in their articulations.

4.4. Segmentation of components using local intensity

Classification based on local intensity works best when local intensity is based on a Gaussian filter with standard deviation $\sigma_L = 2.0$ pixels. As with pixel intensity, the best results are obtained using a pair of thresholds and the optimum values are very similar in the two cases. The best score obtained using local intensity was 74.7%, which is only a very slight improvement over the score for pixel intensity. We conclude that the majority of the improvement of histogram, over and above pixel-based methods, can be attributed to the use of histograms rather than simply to the use of a local-based rather than a pixel-based features.

5. Summary of the local histogram algorithm

In this section we summarize the local histogram algorithm that the analysis of Sections 2 and 3 found to be the best of the algorithms assessed, in a form that will facilitate its reimplementation.

The following parameters of the algorithm have been optimized on our training data. We provide the optimum values we determined, but note that these may change for a different meteorite or imaging setup:

- meteorite/nonmeteorite threshold $T_m = 7500$;
- minimum foreground component size $n_f = 10^3$;
- intensity bin boundaries $\{b_0, \dots, b_n\} = \{0, 13600, 15700, 36000\}$;
- scale of local histograms $\sigma = 7.5$ pixels.

The following parameters have been determined from our training data when using the optimum values above. These values may need to be redetermined for a different meteorite or imaging setup:

- fraction of matrix pixels and large component pixels in training data $f_M = 0.35, f_C = 0.65$
- mean rooted local histograms for matrix and for large

$$\text{component } \vec{\mu}_M = \begin{pmatrix} 0.50 \\ 0.68 \\ 0.36 \end{pmatrix}, \vec{\mu}_C = \begin{pmatrix} 0.84 \\ 0.34 \\ 0.25 \end{pmatrix}$$

Table 1

Performance of the main methods of matrix/large component classification.

Classification basis	% correct			% to perfect
	Slice 1	Slice 2	Mean	
All pixels=Large component	63.4	65.6	64.5	0
Pixel intensity	75.2	71.7	73.4	25
Local intensity	76.4	73.0	74.7	29
Local histogram	86.4	85.3	85.8	60

- covariance matrices for the rooted local histograms of matrix and large component

$$\Sigma_M = \begin{pmatrix} 0.051 & -0.014 & -0.038 \\ -0.014 & 0.021 & -0.001 \\ -0.038 & -0.001 & 0.050 \end{pmatrix},$$

$$\Sigma_C = \begin{pmatrix} 0.032 & -0.031 & -0.031 \\ -0.031 & 0.049 & 0.027 \\ -0.031 & 0.027 & 0.043 \end{pmatrix}.$$

The steps of the local histogram-based classification algorithm are as follows:

1. Produce a *meteorite mask image* by thresholding the image at T_m and keeping 4-connected components of pixels above the threshold only if they are of size at least n_f . Pixels flagged in this image will be called *meteorite pixels*.
2. For each of n histogram bins, form a *bin image*, which has the value unity for *meteorite pixels* whose intensity lies in the interval $[b_r, b_{r+1})$, and zero elsewhere.
3. Blur each *bin image* with a Gaussian filter of standard deviation σ to form the *unnormalized local bin images*.
4. For each *meteorite pixel*, multiply the corresponding n values in the *unnormalized local bin images* by the same factor so that they sum to unity. The result is *local bin images*.
5. Square root all values in the *local bin images* to form the *rooted local bin images*.
6. For each *meteorite pixel*, assemble the corresponding n numbers from the *rooted local bin images* into a *rooted local histogram* \vec{h} .
7. Compute and compare $f_M \cdot N_{\vec{\mu}_M, \Sigma_M}(\vec{h})$ and $f_C \cdot N_{\vec{\mu}_C, \Sigma_C}(\vec{h})$, where $N_{\vec{\mu}, \Sigma}$ is the pdf of a multinormal distribution with the specified mean and covariance. Classify the pixel as matrix if the former is larger, as large component otherwise.

6. Conclusions

Micro-CT is an emerging technology that is now widely used to study the internal features of rock specimens. To be of use in the analysis of meteorites it is essential that the constituent components can be distinguished with semiautomated methods, and preferable that this should be possible with fully automated methods. We have confirmed that the simplest semiautomated approach—classification based on pixel intensity—does not perform well, achieving a score, which is only 25% of the way from baseline to perfect performance. We have presented an alternative approach—classification based on the local histogram of intensities—that performs better: 60% of the way from baseline to perfect. A control experiment confirmed that most of this increase in performance should be attributed to the use of a local histogram rather than merely to the use of a local-based rather than a pixel-based feature. Whether this level of performance is adequate or not is application dependent. Some minor improvement could be expected by shifting to 3-D histograms, but this would require 3-D ground truth data to tune the algorithm and to determine the means and covariances of the different classes. Avenues that could yield substantial increases in algorithm performance are local texture assessment, and enforcing expectations about large component shape. This might then as well allow distinction among chondrules, CAIs, and dark inclusions.

References

- Abramoff, M.D., Magelhaes, P.J., Ram, S.J., 2004. Image processing with ImageJ. *Biophotonics International* 11 (7), 36–42.
- Brearely, A.J., Jones, R.H., 1998. Chondritic meteorites. *Reviews in Mineralogy and Geochemistry* 36 (1), 1–398.
- Crosier, M., Griffin, L.D., 2010. Using basic image features for texture classification. *International Journal of Computer Vision* 88 (3), 447–460.
- Crum, W.R., Camara, O., Hill, D.L.G., 2006. Generalized overlap measures for evaluation and validation in medical image analysis. *IEEE Transactions on Medical Imaging* 25 (11), 1451–1461.
- Dempster, A.P., Laird, N.M., Rubin, D.B., 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society* 39 (1), 1–38.
- Ebel, D.S., Leftwich, K., Brunner, C.E., Weisberg, M.K., 2009. Abundance and size distribution of inclusions in CV3 chondrites by X-ray image analysis. *Lunar and Planetary Science Conference* 40(Abstract), 2065.
- Ebel, D.S., Rivers, M.L., 2007. Meteorite 3-D synchrotron microtomography: methods and applications. *Meteoritics and Planetary Science* 42 (9), 1627–1646.
- Ebel, D.S., Weisberg, M.K., Hertz, J., Campbell, A.J., 2008. Shape, metal abundance, chemistry, and origin of chondrules in the Renazzo (CR) chondrite. *Meteoritics and Planetary Science* 43 (10), 1725–1740.
- Friedrich, J.M., 2008. Quantitative methods for three-dimensional comparison and petrographic description of chondrites. *Computers and Geosciences* 34 (10), 1926–1935.
- Friedrich, J.M., Macke, R.J., Wignarajah, D.P., Rivers, M.L., Britt, D.T., Ebel, D.S., 2008a. Pore size distribution in an uncompacted equilibrated ordinary chondrite. *Planetary and Space Science* 56 (7), 895–900.
- Friedrich, J.M., Wignarajah, D.P., Chaudhary, S., Rivers, M.L., Nehru, C.E., Ebel, D.S., 2008b. Three-dimensional petrography of metal phases in equilibrated L chondrites—effects of shock loading and dynamic compaction. *Earth and Planetary Science Letters* 275 (1–2), 172–180.
- Griffin, L.D., 1997. Scale-imprecision space. *Image and Vision Computing* 15 (5), 369–398.
- Griffin, L.D., 2006. Optimality of the basic colour categories for classification. *Journal of the Royal Society: Interface* 3 (6), 71–85.
- Haralick, R.M., Shapiro, L.G., 1985. Image segmentation techniques. *Computer Vision, Graphics and Image Processing* 29 (1), 100–132.
- Harsanyi, J.C., Chang, C.-I., 1994. Hyperspectral image classification and dimensionality reduction: an orthogonal subspace projection approach. *IEEE Transactions on Geoscience and Remote Sensing* 32 (4), 779–785.
- Hezel, D.C., Needham, A.W., Armytage, R., George, B., Abel, R., Kurahashi, E., Coles, B.J., Rehkammer, M., Russell, S.S., 2010. A nebula setting as the origin for bulk chondrule Fe isotope variations in CV chondrites. *Earth and Planetary Science Letters* 296, 423–433.
- Hezel, D.C., Russell, S.S., Ross, A.J., Kearsley, A.T., 2008. Modal abundances of CAIs: implications for bulk chondrite element abundances and fractionations. *Meteoritics and Planetary Science* 43 (11), 1879–1894.
- Jain, A.K., Duin, R.P.W., Mao, J., 2000. Statistical pattern recognition: a Review. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22 (1), 4–37.
- Ketcham, R.A., Carlson, W.D., 2001. Acquisition, optimization and interpretation of X-ray computed tomographic imagery: applications to geosciences. *Computers and Geosciences* 27, 381–400.
- Koenderink, J.J., 1984. The structure of images. *Biological Cybernetics* 50 (5), 363–370.
- Koenderink, J.J., 1999. The structure of locally orderless images. *International Journal of Computer Vision* 31 (2–3), 159–168.
- McInerney, T., Terzopoulos, D., 1996. Deformable models in medical image analysis: a survey. *Medical Image Analysis* 1 (2), 91–108.
- McSween, H.Y., 1999. *Meteorites and their Parent Planets*. Cambridge University Press, Cambridge, UK.
- Petrou, M., Petrou, C., 2010. *Image Processing: The Fundamentals*. John Wiley & Sons, UK, Chichester.
- Permuter, H., Francos, J., Jermyn, I., 2006. A study of gaussian mixture models of color and texture features for image classification and segmentation. *Pattern Recognition* 39 (4), 695–706.
- Phung, S.L., Bouzerdoum, A., Chai, D., 2005. Skin segmentation using color pixel classification: analysis and comparison. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 27 (1), 148–154.
- Remeysen, K., Swennen, R., 2006. Beam hardening artifact reduction in microfocus computed tomography for improved quantitative coal characterization. *International Journal of Coal Geology* 67 (1–2), 101–111.
- Shen, D., 2007. Image registration by local histogram matching. *Pattern Recognition* 40 (4), 1161–1172.
- Spinsby, J., Friedrich, H., Buseck, P.R., 2008. Volume and surface-area measurement using tomography, with an example from the Brenham pallasite meteorite. *Computers and Geosciences* 34 (1), 1–7.
- Uesugi, M., Uesugi, K., Oka, M., 2010. Non-destructive observation of meteorite chips using quantitative analysis of optimized X-ray micro-computed tomography. *Earth and Planetary Science Letters* 299 (3–4), 359–367.
- Van Geet, M., Swennen, R., Wevers, M., 2001. Towards 3-D petrography: application of microfocus computer tomography in geological science. *Computers and Geosciences* 27 (9), 1091–1099.