




Computer Vision and Machine Learning for Autonomous Characterization of AM Powder Feedstocks

BRIAN L. DECOST ¹ HARSHVARDHAN JAIN,¹
ANTHONY D. ROLLETT ¹ and ELIZABETH A. HOLM ^{1,2}

1.—Department of Materials Science and Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA. 2.—e-mail: eaholm@andrew.cmu.edu

By applying computer vision and machine learning methods, we develop a system to characterize powder feedstock materials for metal additive manufacturing (AM). Feature detection and description algorithms are applied to create a microstructural scale image representation that can be used to cluster, compare, and analyze powder micrographs. When applied to eight commercial feedstock powders, the system classifies powder images into the correct material systems with greater than 95% accuracy. The system also identifies both representative and atypical powder images. These results suggest the possibility of measuring variations in powders as a function of processing history, relating microstructural features of powders to properties relevant to their performance in AM processes, and defining objective material standards based on visual images. A significant advantage of the computer vision approach is that it is autonomous, objective, and repeatable.

INTRODUCTION

Computer vision is a branch of computer science that develops computational methods to extract information from visual images.¹ Familiar applications of computer vision include facial recognition on Facebook, Google image search, and autonomous vehicle navigation. Computer vision and its related field, machine vision, are also widely used in manufacturing, particularly in robotics, process control, and quality inspection.

Powder bed additive manufacturing (AM) is an emerging technology for three-dimensional (3D) metal printing.^{2,3} The process is conceptually simple. A layer of fine metal powder is spread on the build plate. By using a laser or electron beam, powder particles are melted and resolidified. Another layer of powder is spread, and the process repeats multiple times to build up the final part layer by layer.

Among the many challenges in deploying this new manufacturing system are numerous metal powder-related issues centered on understanding how the physical characteristics of the powder (size, shape, and surface character) affect processing parameters (flowability and spreadability) and build outcomes

(porosity and flaws). Fundamental to understanding these relationships is effectively characterizing the powders themselves.

To date, characterization of AM powder feedstocks has relied on direct measurements of powder properties of interest. For example, Strondl et al. used dynamic image analysis to capture photomicrographs of powders, segment them, and measure particle size and aspect ratio distributions. These quantities, along with powder rheology measurements, were found to correlate with powder flow and spreading characteristics.⁴ A group of case studies by Clayton et al. concluded that particle size distribution alone is insufficient to determine powder properties. Instead, they characterized powders by using rheological measurements, which they found to correlate with powder properties such as the degree of recycling, the manufacturing method, or the presence of additives.⁵ In perhaps the most comprehensive study of its kind, Slotwinski et al. systematically characterized virgin and recycled stainless steel and cobalt chrome powders in an effort to develop standards for AM feedstock materials. They measured particle size and shape with laser diffraction, x-ray computed tomography, and optical and scanning electron microscopy. In addition,

Table I. Metadata for the AM powder image dataset, including powder type, composition, source, sample labels, and number of images per sample

<u>Powder</u>	<u>Material system</u>	<u>Platform</u>	<u>Sample</u>	<u>Images</u>
Al-EOS	AlSi10Mg	EOS	1	9
			2	12
			Test	11
In-EOS	Inconel alloy 718	EOS	1	12
			2	11
			Test	12
MS-EOS	Maraging steel	EOS	1	12
			2	12
			Test	14
SS-EOS	Type 316 stainless steel	EOS	1	8
			2	17
			Test	17
Ti64-EOS	Ti-6Al-4V	EOS	1	12
			2	11
			Test	12
Ti64-#1	Ti-6Al-4V	ARCAM	1	11
			2	10
			Test	9
Ti64-#2	Ti-6Al-4V	ARCAM	1	8
			2	9
			3	9
Ti64-#2	Ti-6Al-4V	ARCAM	Test	10
			1	7
			2	7
			3	7
			4	5
Test	8			

they determined atomic structure and composition via energy-dispersive element x-ray analysis, x-ray photoelectron spectroscopy, and x-ray diffraction.⁶ Finally, Nandwana et al.⁷ studied particle size, flowability, and chemistry for two powders used in electron beam AM. Recycling caused significant changes to chemistry in one powder but minimal changes in the other; particle size and flowability were unaffected by recycling. Measurements such as these provide valuable insight into the factors that influence powder properties; nevertheless, data science offers a complementary approach that can extract information from a data stream directly, without reductive measurement.

In this article, we explore applications of computer vision for autonomously evaluating powder raw materials for metal AM. Instead of explicitly identifying and measuring individual particles, our method implicitly characterizes powder micrographs as a distribution of local image features.^{8,9} We demonstrate that the computer vision system is capable of classifying powders with different distributions of particle size, shape, and surface texture, as well as of identifying both representative and atypical powder images. AM applications include powder batch qualification, quantifying the effects of powder recycling, selecting build parameters based on powder characteristics, identifying

features that might be associated with powder spreading or build flaws, and defining objective material standards based on visual images.

Finally, we note that a significant advantage of the computer vision approach is that it is an autonomous and objective system that does not require a subjective human judgment about what to measure or how to measure it. It is not limited to powder micrographs and in fact can be extended to new image data sets, including bulk microstructural images, without customization.^{9,10}

METHODS

Materials

We collected micrographs from eight different gas-atomized powders, as listed in Table I. Five powders were procured from EOS and are intended for the EOS printer machine: AlSi10Mg (Al-EOS), Inconel 718 (In-EOS), Maraging Steel (MS-EOS), Stainless Steel 316 (SS-EOS), and Ti-6Al-4V (Ti64-EOS). The remaining three Ti64 powders are intended for the ARCAM machine and were procured from ARCAM and two additional suppliers.

We used a spatula to sample a small amount of powder after shaking the container to prevent sample biases from the settling of powder during transportation and storage. A thin layer of powder




Fig. 1. Eight powder micrographs for the Ti64-EOS material system, selected at random from 35 SEM images to illustrate typical microstructural variability for a given system. It is apparent that images vary in particle density as well as in individual particle particulars. Images have been processed to normalize image contrast and noise levels as described in the text.

was blown over double-sided carbon tape on a stub with an Air Duster. We cleaned the sample with pressurized air to remove any loose particles before performing scanning electron microscopy (SEM) on samples. We took images in backscattered imaging mode, adjusting the magnification for each powder system in an attempt to maintain a similar number density of powder particles across powder systems with differing average particle sizes. The motivation for this choice was to eliminate the influence of average particle size in classification and focus on particle morphology and the relative shape of the particle size distribution.

We found it necessary to preprocess the powder micrographs to mitigate differences in image contrast and noise levels introduced during image acquisition. We segment and suppress the background of each powder micrograph by applying Li's minimum cross-entropy threshold selection method,^{11,12} which yields good results for this dataset with no adjustable parameters. Following background suppression, we perform local histogram equalization¹³ to ensure that powder micrographs from different material systems and imaging conditions have similar intensity distributions. The processed images have more consistent contrast than do the raw images, which helps in better extraction of texture features that are relevant to the spatial and morphological characteristics of the powder particles, as opposed to high-frequency features related to SEM imaging conditions. Figure 1 shows several processed micrographs for the Ti64-EOS system to illustrate the typical variability in images within a powder type, and Fig. 2 shows a representative powder micrograph for each of the eight

powder systems. As shown in Figs. 1 and 2, the images used in the computer vision analysis do not contain scale bars to ensure that only features of the powder itself are used in the analysis; nevertheless, the absolute and relative scales of the images can be deduced from the powder size distributions in Fig. 3.

As detailed in Table I, we prepared between three and five independent powder samples for each powder system, with more samples for powders with larger-than-average particle sizes. From each powder sample, we collected between 5 and 17 backscattered electron micrographs, avoiding collecting images of intersecting regions on the sample. We verify that no images included in our analysis contain significant overlapping regions by applying to each pair of micrographs originating from the same physical sample a keypoint-matching algorithm with RANSAC filtering to identify micrographs related by a simple translation.¹⁴

Figure 3 shows cumulative particle size distributions obtained by applying threshold and watershed segmentation^{15,16} to the SEM micrographs after background suppression. For each powder system, we sample a total of 10,000 watershed particles from 20 powder micrographs. The five powders intended for the EOS machine use are smaller in size than those intended for ARCAM. Generally, the EOS powders have similarly shaped, approximately lognormal particle size distributions, albeit with different mean particle sizes. The powders intended for the ARCAM machine are much coarser and display a severe lower tail as a result of apparent sieving at 100 μm .




Fig. 2. Representative powder micrographs for each of the eight powder material systems: (a) Al-EOS, (b) In-EOS, (c) MS-EOS, (d) SS-EOS, (e) Ti64-#1, (f) Ti64-#2, (g) Ti64-#3, and (h) Ti64-EOS. Images have been processed to normalize image contrast and noise levels as described in the text.

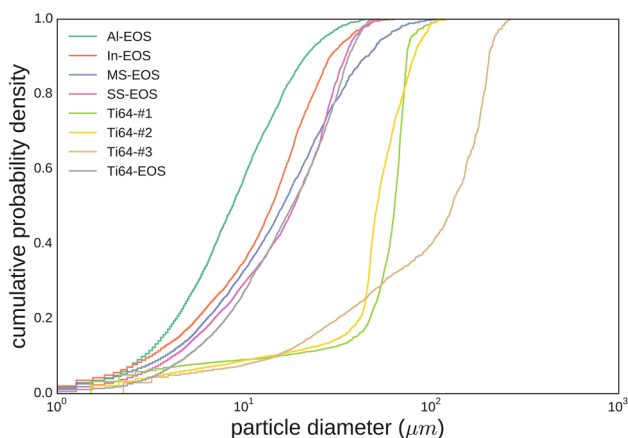


Fig. 3. Cumulative particle size distributions for the eight powder material systems, as measured by watershed segmentation for 10000 particles per system. Note that the systems differ in both mean particle size and overall shape of the distribution.

Computer Vision

The goal of our computer vision system is to produce a numerical representation that contains the relevant information necessary to characterize a visual image.⁹ In other words, we want to turn an image into a vector that summarizes its visual information content. This process has three steps: First, the system finds interest points, or visual features, in the image; that is, it decides what to look at. Next, the system encodes each feature as a vector descriptor; that is, it numerically characterizes the features. Finally, the system groups the

feature descriptors together to create a representation of the image as a whole; that is, it determines the microstructural fingerprint of the image. This process is illustrated schematically in Fig. 4.

To find the visual features in a powder micrograph, we use two complementary interest point localization methods common in computer vision: the difference of Gaussians¹⁷ and the Harris-LaPlace methods.¹⁸ Together, these methods yield a set of distinctive blob- and corner-like image features (respectively), each of which has a characteristic scale and orientation determined by the local contrast gradient in the image patch surrounding the interest point. A given micrograph may contain several thousand interest points, and they often correspond to the features that a materials scientist might identify as significant, i.e., spots, lines, corners, visual textures, etc. The yellow circles in Fig. 4a show the locations, scales, and orientations of 100 randomly selected interest points.

To characterize these individual features, we apply another widely used computer vision technique, scale-invariant feature transform (SIFT), to compute feature descriptors for each interest point patch.^{17,19} The SIFT descriptors describe the local image structure by encoding the orientation and magnitude of contrast gradients within 16 spatial bins surrounding the interest point. Basically, SIFT transforms each visual object into a characteristic 128-dimensional vector. The blue frame in Fig. 4b shows a schematic of the 16-bin SIFT descriptor corresponding to a large particle.

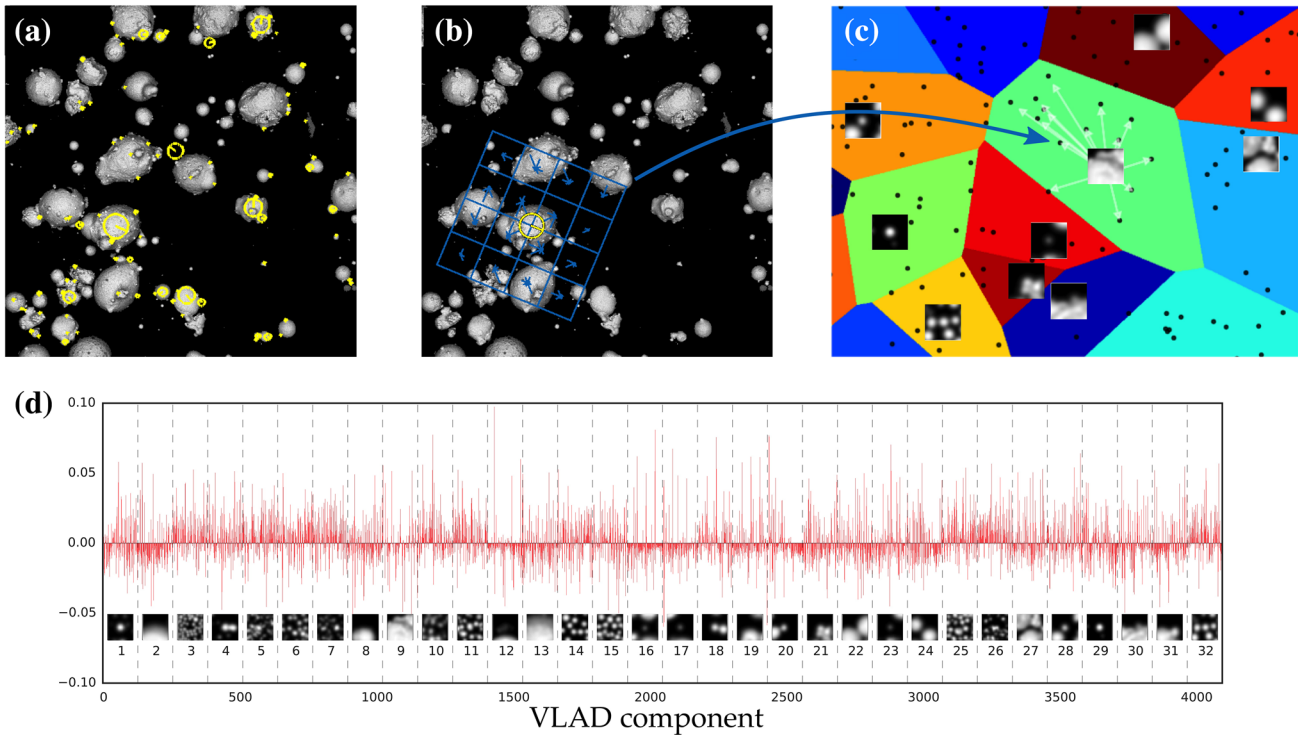


Fig. 4. Schematic diagram illustrating the construction of SIFT-VLAD microstructure representations. (a) Select oriented interest points (yellow markers) from a powder micrograph (100 randomly selected interest points shown). (b) Compute a SIFT descriptor (blue grid) for each interest point. (c) Cluster SIFT descriptors (colored regions) such that SIFT descriptors (black dots) are associated with their most similar visual word (image patches); compute a residual vector for each visual word (white arrows). (d) Concatenate the normalized residual vectors (red bars) of each visual word (image patches) to construct the VLAD representation, which serves as a microstructure fingerprint.

As materials scientists, we understand that microstructures often contain numerous examples of the “same” feature; thus, we group microstructural features into classes, such as precipitates, grain boundaries, or powder particles. The computer vision system performs a similar function by clustering the SIFT descriptors into groups of like features (termed “visual words”). In our system, we use k-means clustering²⁰ to partition 15% of the SIFT descriptors extracted from the training images into 32 visual words, as schematically indicated in Fig. 4c. The black markers in this 2D schematic indicate SIFT descriptors for individual interest points, and the colored cells indicate the clusters or partitions that demarcate the visual words; thus, each feature in the micrograph is associated with a particular visual word. In our powder micrographs, a visual word might represent a spherical particle, a neck between particles, a cluster of particles, a surface texture, or some other feature, as indicated by the image patches in Fig. 4c. It is important to note, however, that the visual words are determined by the computer vision system; there is no subjective (human) judgment involved.

To obtain an overall image representation, local feature methods often simply model an image as the histogram of its visual words [termed a bag of words (BoW) model.²¹ Nevertheless, this has the drawback that all features are given the same weight, even

though some unambiguously belong to a particular visual word, while others fall near the borders between visual words. We therefore apply a vector of locally aggregated descriptors (VLAD) encoding to model more effectively the overall distribution of local image features by summing up the difference between each local feature descriptor and its corresponding visual word.²² The white arrows in the central green cell in Fig. 4c illustrate this residual vector calculation for the SIFT descriptor shown in Fig. 4b. The result is a $128 \times 32 = 4096$ -dimensional vector that represents the image as a whole, as illustrated in Fig. 4d: The image patches illustrate the visual words, and the red bars show their corresponding residual SIFT vectors.

VLAD descriptors are often reduced in dimensionality by up to an order of magnitude without significant degradation of the image representation quality;²³ we apply principal component analysis (PCA) to reduce the dimensionality of the VLAD representations to 32. The first 32 principal components of the VLAD representations for the training images account for 76.4% of the variance of the high-dimensional representations. These first 32 principal components are the image representation or “microstructural fingerprint” we use to characterize each powder micrograph image.

We note that there are numerous image representation schemes in the computer vision literature.²⁴⁻³⁰ Our motivation for choosing the SIFT-VLAD system (as opposed to, for example, convolutional neural network-based texture representations³¹) is its strong rotation and scale invariance. Unlike photographic scenes or portraits, which are almost always oriented with gravity pointing down, powder micrographs do not have a natural orientation. The SIFT-VLAD system also captures mid-level features, such as particles, clusters, and particle intersections, with good fidelity, as shown in Fig. 4d.

RESULTS

Powder Classification

Recall that the computer vision system autonomously determines which features in a particular micrograph are visually important. Thus, our first question should be “Is the system looking at the right features, from a materials science point-of-view?” Moreover, the computer vision system does not perform any “conventional” feature measurements, such as particle size or surface area analysis. So our second question must be “Does the system capture the relevant microstructural quantities?”

To answer these questions, we challenge the system to sort powder micrographs according to powder type. To classify images correctly, the system will have to sense both the features that differentiate powders and the quantitative difference in their particle size distributions.

For each powder system, we held out the micrographs collected from one physical sample as an independent testing set, and we used the micrographs from the other samples as a training and validation set for machine learning via a support vector machine (SVM) classifier.³² The SVM classifier takes in the image representations from the SIFT-VLAD system and attempts to find a hyperplane that separates them into different classes. We used $10\times$ fivefold cross-validation to select the SVM regularization parameter, C , which represents the trade-off between the flexibility and generalizability of the classifier. For each of 50 cross-validation iterations, we randomly choose $1/5$ of the training images to evaluate the performance of a classifier trained on the remaining $4/5$ of the training images. Both training and validation accuracies converge to 100% (i.e., all images are assigned to their correct powder type) when $C \sim 0.25$, indicating that the SIFT-VLAD representation can reliably classify powder micrograph images without significant overfitting. Figure 5a shows the detailed validation results for all 50 validation folds when $C = 0.25$: The only misclassifications during cross-validation are a result of the similarity between the Maraging and Inconel powders, and especially between the

Ti64#1 and Ti63#2 powders. The overall cross-validation accuracy with these parameters is $96.5 \pm 2.5\%$.

Given a fully trained powder classifier, the next challenge is to demonstrate that it can recognize images taken from physical samples that were not included in its training dataset. To that end, we trained an SVM classifier by using the entire training set with the parameters selected via cross-validation, and we used it to classify the images in the independent testing set. As shown in Fig. 5b, the computer vision system classifies the previously unseen images with an overall accuracy of about 95%, which is comparable to the cross-validation accuracy. Half of the powder types achieved perfect classification, and no more than two images were misclassified in any powder system. Two Ti64-#1 images were misclassified as Ti64-#2, which has both similar particle morphologies (Fig. 2) and particle size distribution (Fig. 3); the same applies to the misclassification of In-EOS as MS-EOS. The other two misclassifications are a result of outlier (highly atypical) test images and emphasize the need for statistically representative image sets.

Overall, the computer vision system successfully identifies powder types from micrographs without requiring subjective judgments about what features to measure or how to measure them. This result demonstrates the promise for computer vision systems to provide autonomous microstructural analysis not only for AM powders but also for microstructural images in a more general sense.

Data Visualization

Although the confusion matrices in Fig. 5 indicate the classification accuracy for the powder images, they are not helpful in visualizing how the images vary within and between the powder types. This is, in fact, a significant challenge in data science generally because the data points (in this case, the image representations) occupy a high dimensional space. Thus, a variety of dimensionality reduction techniques have been developed to help humans interpret high-dimensional distributions.

PCA is a linear dimensionality-reduction technique that determines a set of n orthogonal axes (the principle components) along which an n -dimensional dataset varies, from most to least.³³ In essence, PCA reorients the natural axes of the dataset so as to maximize the spread in the data along each principle component, in decreasing order from the first to the n th. Along a given axis, PCA preserves pairwise high-dimensional distances between data points. Therefore, for a set of image representations, the distance between data points on a PCA plot should be related to the visual difference between the respective images. Nevertheless, because PCA plots are typically of lower

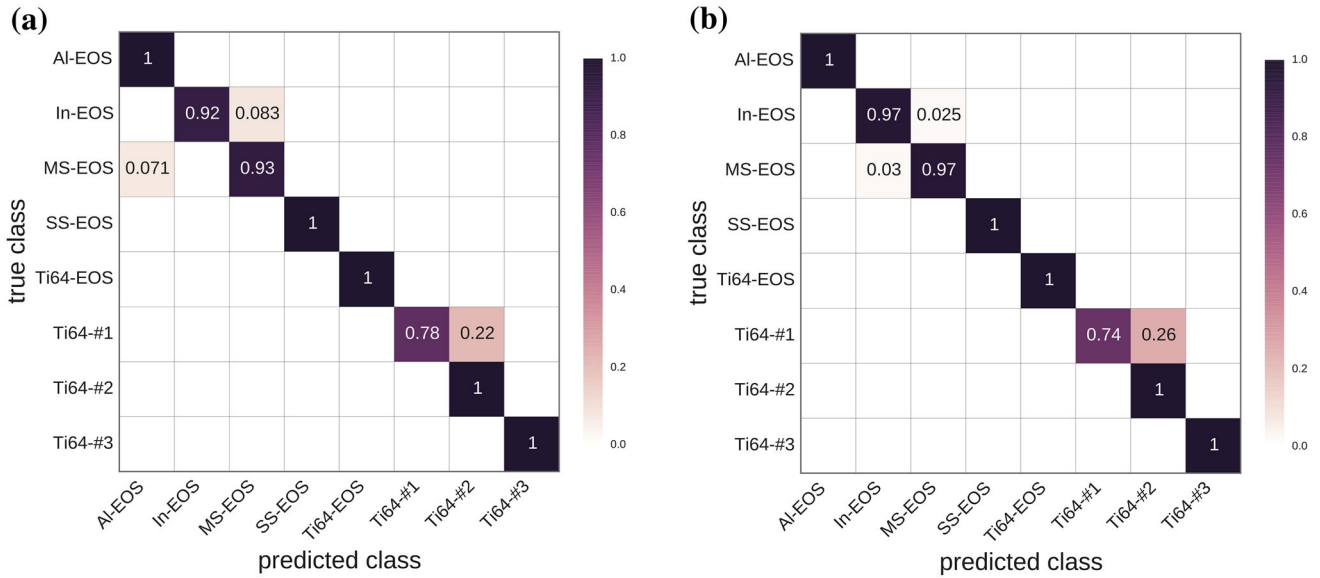


Fig. 5. Confusion matrices for powder classification in the (a) cross-validation and (b) independent test sets. The rows designate the actual powder system, and the columns indicate the powder system predicted by the computer vision. The numbers tally the fraction of images in each actual powder system that were classified as a given predicted powder system. Thus, the diagonal entries represent the fraction of accurate powder classifications for each powder system.

dimension than the data themselves, points may appear to overlap in a PCA plot but be distinct in a higher dimension.

Figure 6a shows a PCA plot for the first two principal components of the powder image dataset; the first two principal components account for approximately 42% of the total variance of the high-dimensional SIFT-VLAD representations. Notably, the first principal component effectively differentiates between the powders intended for the EOS machine (at the left of the PCA map) and the powders intended for the ARCAM machine (at the right of the PCA map). The Ti64-#3 micrographs have the most distinctive SIFT-VLAD representations in this dataset most likely because the Ti64-#3 powder was imaged at the highest magnification relative to its average particle size. This could potentially result in a different set of image texture features dominating these micrographs (finer-grained surface details or fewer higher-level groups of particles).

Although the EOS powders cluster together on the PCA map, we know that they are sufficiently distinct for classification. To visualize their differences, we must examine additional principal components or apply a different visualization method. t-Distributed Stochastic Neighbor Embedding (t-SNE) is a nonlinear technique that attempts to preserve the local neighborhood structure of the high-dimensional data in the low-dimensional map at the expense of retaining meaningful long-range correspondences.^{34,35} In other words, the pairwise distance between nearby points in a t-SNE map is meaningful, but the distance between clusters of points is not.

Figure 6b shows a t-SNE map for the powder image dataset. (Because t-SNE is a stochastic optimization technique, we selected the best map out of 10 independent t-SNE embeddings.) The t-SNE map reveals the finer cluster structure of the EOS powders, resolving them into five distinct clusters. The three ARCAM powders are also more differentiated in this mapping. It is apparent that the independent testing set images (shown as squares) are generally consistent with the training set representations (shown as circles) within each powder system and that the cluster overlap that causes misclassification is evident as well.

As Fig. 6 demonstrates, data visualization can be a helpful adjunct to computer vision outcomes, assisting in both understanding and interpretation of results. Nonetheless, no single, low-dimensional visualization contains all of the information in the high-dimensional dataset.

Representative Images

The microstructural fingerprint determined by the computer vision system is a numerical representation of a microstructural image. As such, its uses are not limited to classification tasks. For example, because visually similar images have numerically similar representations, the microstructural fingerprint can form the basis for a visual search, as we have reported previously.^{8,9} Another application addresses a classic problem in microstructural science: determination of a representative image.

Although individual micrographs are often presented in the literature as “representative” of the microstructure as a whole (cf. Ref. 36), there has

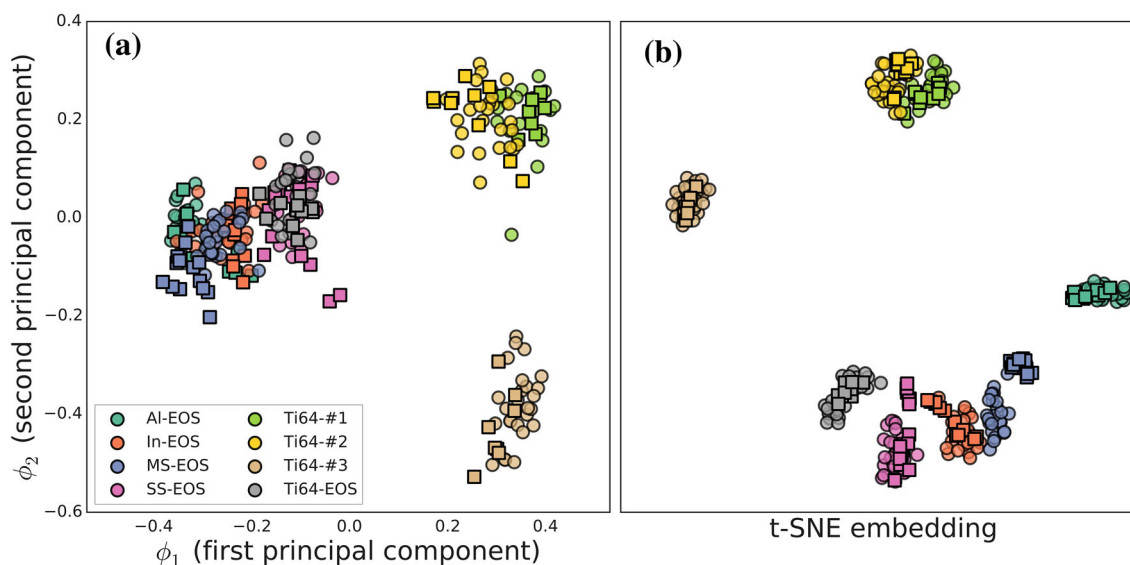


Fig. 6. (a) Data visualization showing the first two principal components of the powder micrograph SIFT-VLAD representations. The five EOS powders are clustered on the left side of the plot; the three ARCAM powders are on the right. (b) t-SNE visualization of the powder micrograph SIFT-VLAD representations. The eight powders occupy more distinct clusters. Marker colors indicate the material system of the corresponding map point; circular markers indicate training set micrographs; and square markers indicate the independent testing set.

been no rigorous test to confirm the validity of such assertions. Is a given image the most representative, or simply the most attractive, convenient, or well-prepared? By using the microstructural fingerprint, we can, for the first time, objectively and quantitatively determine which image is the most representative of a group of micrographs.

The procedure for selecting a representative micrograph involves computing the average high-dimensional SIFT-VLAD representation for each material system and choosing the micrograph with the minimum scalar distance between its representation and the corresponding material system average. The micrographs in Fig. 2 were selected with this procedure, and thus, they are rigorously and quantitatively representative of their material class. Similarly, Fig. 7 shows the “least representative” (or “most atypical”) processed micrograph for each powder system, obtained by selecting the image with the maximum distance between its representation and the corresponding material system average. Notably, the representative micrographs for the EOS-manufactured powders all have similar spatial and morphological characteristics (likewise for the Ti64 powders intended for the ARCAM machine). The atypical micrographs shown in Fig. 7 differ from their representative counterparts in Fig. 2 primarily in overall image intensity (panel a, Al-EOS) or in spatial distribution of powder particles (more dense for In-EOS, SS-EOS, Ti64-#1, Ti64-EOS, and Ti64-#1; less dense for Al-EOS and Ti64-#2).

The ability to quantify how well or how poorly an image represents a class of images enables a variety of applications. For instance, to qualify an AM

powder, an engineer could measure how closely a new batch of powder resembles previous batches or how significantly a recycled powder differs from its virgin condition. Similarly, one might compare a new powder to a library of known powders to select build parameters for the new system. Outlier images may contain valuable information about unusual microstructural features (e.g., atypical particle shapes and sizes) that might be associated with powder spreading or build flaws.

In a more general sense, the definition of a representative microstructure can form the basis of an objective standard for microstructural qualification. This is particularly significant for AM, where 3D printing makes it possible to build objects with the same composition and geometry as a conventionally manufactured part. Qualifying the AM microstructure (as a proxy for properties) is an important aspect of qualifying the part as a whole.

CONCLUSION

Computer vision and machine learning methods offer new possibilities for evaluating powder raw materials for metal AM. In place of identifying and measuring individual particles, this method implicitly characterizes powder micrographs as a distribution of local image features, termed the “microstructural fingerprint.” Operating on a set of powder micrographs with different distributions of particle size, shape, and surface texture, the computer vision system achieves a classification accuracy of more than 95% and a combination of data visualization techniques add further insight into powder characteristics. By representing a




Fig. 7. Least representative (most atypical) powder micrographs for each powder material system: (a) Al-EOS, (b) In-EOS, (c) MS-EOS, (d) SS-EOS, (e) Ti64-#1, (f) Ti64-#2, (g) Ti64-#3, and (h) Ti64-EOS. Images have been processed to normalize image contrast and noise levels as described in the text.

visual image with the microstructural fingerprint, both representative and atypical powder images can be identified and analyzed. As an autonomous and objective system, this method enables AM applications including powder batch qualification, quantifying the effects of powder recycling, selecting build parameters based on powder characteristics, identifying features that might be associated with powder spreading or build flaws, and defining objective material standards based on visual images. Finally, this approach is not limited to powder micrographs and in fact can be extended to new image data sets, including bulk microstructural images.

ACKNOWLEDGEMENTS

This work was funded in part by National Science Foundation Grant Numbers DMR-1307138 and DMR-1507830 and through the John and Claire Bertucci Foundation. We appreciate the authors and user communities of open-source image processing, computer vision, and machine learning algorithms via scikit-learn,³⁷ scikit-image,³⁸ VLFeat,³⁹ and the reference implementation of t-SNE. America Makes is acknowledged for the provision of the metal powders used in this work.

REFERENCES

1. R. Szeliski, *Computer Vision: Algorithms and Applications* (New York: Springer, 2010), pp. 1–812.
2. W.E. Frazier, *J. Mater. Eng. Perform.* 23, 1917 (2014).
3. D.D. Gu, W. Meiners, K. Wissenbach, and R. Poprawe, *Int. Mater. Rev.* 57, 133 (2012).
4. A. Strondl, O. Lyckfeldt, H. Brodin, and U. Ackelid, *JOM* 67, 549 (2015).
5. J. Clayton, D. Millington-Smith, and B. Armstrong, *JOM* 67, 544 (2015).
6. J.A. Slotwinski, E.J. Garboczi, P.E. Stutzman, C.F. Ferraris, S.S. Watson, and M.A. Peltz, *J. Res. Nat. Inst. Stand. Technol.* 119, 460 (2014).
7. P. Nandwana, W.H. Peter, R.R. Dehoff, L.E. Lowe, M.M. Kirka, F. Medina, and S.S. Babu, *Metall. Mater. Trans. B* 47, 754 (2016).
8. B.L. DeCost and E. Holm, *Comput. Mater. Sci.* 126, 438 (2017).
9. B.L. DeCost and E.A. Holm, *Comput. Mater. Sci.* 110, 126 (2015).
10. A. Chowdhury, E. Kautz, B.I. Yener, and D. Lewis, *Comput. Mater. Sci.* 123, 176 (2016).
11. C.H. Li and C.K. Lee, *Pattern Recognit.* 26, 617 (1993).
12. C.H. Li and P.K.-S. Tam, *Pattern Recognit. Lett.* 19, 771 (1998).
13. K. Zuiderveld, in *Proceedings of Graphics Gems IV*, p. 474 (1994).
14. M.A. Fischler and R.C. Bolles, *Commun. ACM* 24, 381 (1981).
15. C.A. Schneider, W.S. Rasband, and K.W. Eliceiri, *Nat. Methods* 9, 671 (2012).
16. P. Soille, and L.M. Vincent, in *Proceedings of Visual Communications and Image Processing*, p. 240 (1990).
17. D.G. Lowe, in *Proceedings of Seventh IEEE International Conference on Computer Vision*, p. 1150 (1999).
18. K. Mikolajczyk, and C. Schmid, in *Proceedings of Eighth IEEE International Conference on Computer Vision*, p. 525 (2001).
19. D.G. Lowe, *Int. J. Comput. Vis.* 60, 91 (2004).
20. S. Lloyd, *IEEE Trans. Inf. Theory* 28, 129 (1982).
21. G. Csurka, C. Dance, L. Fan, J. Willamowski, and C. Bray, in *Proceedings of Workshop on Statistical Learning in Computer Vision (ECCV)*, p. 1 (2004).
22. H. Jegou, M. Douze, C. Schmid, and P. Perez, in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, p. 3304 (2010).
23. R. Arandjelovic and A. Zisserman, in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, p. 1578 (2013).

24. Y. Guo, Y. Liu, A. Oerlemans, S. Lao, S. Wu, and M.S. Lew, *Neurocomputing* 187, 27 (2015).
25. Y. LeCun, Y. Bengio, and G. Hinton, *Nature* 521, 436 (2015).
26. J. Li and N.M. Allinson, *Neurocomputing* 71, 1771 (2008).
27. T.R. Reed and J.H. Dubuf, *CVGIP Image Underst.* 57, 359 (1993).
28. J.R. Schmidhuber, *Neural Netw.* 61, 85 (2015).
29. J. Zhang, M. Marszalek, S. Lazebnik, and C. Schmid, *Int. J. Comput. Vis.* 73, 213 (2007).
30. J. Zhang and T. Tan, *Pattern Recognit.* 35, 735 (2002).
31. A. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, p. 806 (2014).
32. C. Cortes and V. Vapnik, *Mach. Learn.* 20, 273 (1995).
33. H. Abdi and L.J. Williams, *Wiley Interdiscip. Rev. Comput. Stat.* 2, 433 (2010).
34. G.E. Hinton and S.T. Roweis, in *Proceedings of Advances in Neural Information Processing Systems*, p. 833 (2002).
35. L.V.D. Maaten and G. Hinton, *J. Mach. Learn. Res.* 9, 2579 (2008).
36. J.G. Kaufman, *Introduction to Aluminum Alloys and Tempers* (Materials Park: ASM International, 2000), p. 119.
37. F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, and V. Dubourg, *J. Mach. Learn. Res.* 12, 2825 (2011).
38. S. Van der Walt, J.L. Schonberger, J. Nunez-Iglesias, F. Boulogne, J.D. Warner, N. Yager, E. Gouillart, and T. Yu, *PeerJ* 2, 453 (2014).
39. A. Vedaldi and B. Fulkerson, in *Proceedings of 18th ACM international Conference on Multimedia*, p. 1469 (2010).