Computational Materials Science 233 (2024) 112748

Contents lists available at ScienceDirect

Computational Materials Science

journal homepage: www.elsevier.com/locate/commatsci

Full length article

Improving uranium oxide pathway discernment and generalizability using contrastive self-supervised learning

Jakob Johnson [a,b], Luther McDonald [c], Tolga Tasdizen [b,d,*]

[a] University of Utah, Kahlert School of Computing, 50 S. Central Campus Drive, Room 3190, Salt Lake City, 84112, UT, United States
[b] University of Utah, Scientific Computing and Imaging Institute, 72 S Central Campus Drive, Room 3750, Salt Lake City, 84112, UT, United States
[c] University of Utah, Department of Civil and Environmental Engineering, Nuclear Engineering Program, 110 Central Campus Drive, Suite 2000, Salt Lake City, 84112, UT, United States
[d] University of Utah, Department of Electrical and Computer Engineering, 50 S. Central Campus Drive, Room 2110, Salt Lake City, 84112, UT, United States

ARTICLE INFO

Dataset link: https://doi.org/10.17632/74k85vybmk.1

Keywords:
Nuclear forensics
Uranium ore concentrates
Self-supervised learning
Contrastive learning
Machine learning
Image analysis

ABSTRACT

In the field of Nuclear Forensics, there exists a plethora of different tools to aid investigators when performing analysis of unknown nuclear materials. Many of these tools offer visual representations of the uranium ore concentrate (UOC) materials that include complimentary and contrasting information. In this paper, we present a novel technique drawing from state-of-the-art machine learning methods that allows information from scanning electron microscopy images (SEM) to be combined to create digital encodings of the material that can be used to determine the material's processing route. Our technique can classify UOC processing routes with greater than 96% accuracy in a fraction of a second and can be adapted to unseen samples at similarly high accuracy. The technique's high accuracy and speed allow forensic investigators to quickly get preliminary results, while generalization allows the model to be adapted to new materials or processing routes quickly without the need for complete retraining of the model.

1. Motivation

Commonly known as "yellowcake", Uranium Ore Concentrates (UOCs) are a class of materials that are a precursor to nuclear fuel and are the most common form of uranium shipped throughout the world. Should illicit material be intercepted, forensic investigators work to determine what the material is and where it originated. Recent focus has been on using image processing algorithms and machine learning to aid investigators in classifying the material's processing history based on surface morphology. Different chemical processing conditions result in unique morphological features which can be captured using SEM. Libraries of these SEM images are needed to improve image processing algorithms and machine learning, but the cost of acquiring additional images prohibits the creation of an exhaustive dataset.

Our goal is to utilize state-of-the-art self-supervised machine learning methods to generate an encoder that is able to create a robust encoding of scanning electron microscopy (SEM) images, even on previously unseen image classes.

1.1. Background and related work

After uranium ore is milled, dissolved, and purified, precipitating reagents are used to remove the uranium from the solution. This precipitate is then dried and calcinated at high temperatures, resulting in UO_2, UO_3, or U_3O_8. The combination of precipitating agent and calcination technique results in unique surface morphologies of the resulting uranium oxides, which investigators can use in determining the origin of an interdicted material as uranium processing facilities tend to use unique pathways to produce the desired UOCs, giving unique morphology "signatures".

Many techniques are available to forensic investigators when determining attribution, such as mass spectrometry or X-ray diffraction [1]. Recent work has demonstrated the power of image analysis algorithms using scanning electron microscope (SEM) images of uranium oxide. Fongaro et al. [2] used an angle measure technique (AMT) algorithm to characterize the surface texture captured in SEM images, using features such as particle size, homogeneity, and graininess. They performed statistical analyses on the characterization to classify the different types of UOCs. Heffernan et al. [3] used an automatic segmentation algorithm to identify and segment fully visible surface particles and then differentiated mixtures of U_3O_8 based on the area of the segmented particles.

During the recent growth in deep learning research, convolutional neural nets (CNNs) have been at the forefront of computer vision for

* Corresponding author.
E-mail address: tolga@sci.utah.edu (T. Tasdizen).

https://doi.org/10.1016/j.commatsci.2023.112748
Received 15 August 2023; Received in revised form 2 September 2023; Accepted 15 December 2023
Available online 3 January 2024

natural and scientific image tasks. Most common methods work in a supervised manner, assuming the labels for the training images are given ahead of time. The CNN optimizes its parameters to match the label output given an input image. Schwerdt et al. (2018) [4] and Hanson et al. [5] used supervised CNNs to directly differentiate materials by taking advantage of the powerful classification performance of these models. Most recently, Ly et al. [6] developed a multi-input single-output (MISO) method that would utilize different SEM magnifications at the same time. Their method beat a single input CNN by a large margin, showing the importance of magnification on surface morphology analysis. Supervised CNN models can be very accurate but are not generalizable to new data and require often expensive data labeling.

Unsupervised learning is a related branch of machine learning in which labels are not given for inputs. In computer vision, unsupervised learning usually involves extracting meaningful statistics or information from an image by condensing it into a vectorized encoding. A model is trained on a difficult task where it must learn to extract information from the image. This information can then be used with other machine learning methods as a lower-dimensional representation of an image. Girard et al. [7] show an unsupervised learning method called a vector-quantized variational autoencoder (VQ-VAE), where images are squeezed through an "information bottleneck" and reassembled to generate the same image. The low-dimensional representation in the bottleneck should then hold all of the relevant information about the image. When this encoding is generated, many methods, such as K-nearest neighbors or linear discrimination, can assign class labels to the vector-encoded image. Girard et al. [7] apply this method to the UOC pathway discernment task, achieve good accuracy, and demonstrate an ability to describe materials on which the encoder was not trained. Another technique for generating this encoding through unsupervised training is contrastive learning. In this technique, a starting image is augmented twice through transformations such as color shift, cropping, or rotation. The images are then passed through a neural net, generating a low-dimensional representation of each image. A contrastive loss function draws positive pairs (augmented versions of the same starting image) together and pushes negative pairs (other images) apart. This method is called self-supervised learning because the model generates both the prediction and the target, supervising itself. Recent work [8–10] has shown that these contrastive self-supervised models outperform autoencoders and are approaching or surpassing supervised accuracy. These results suggest that image reconstruction is not the best to learn when classification is the downstream task.

2. Data

Obtained from Schwerdt et al. (2019) [11], the SEM images used in this study represent four common UOCs and their calcination products. We use the same dataset as Girard et al. [7] and Ly et al. [6], allowing for a fair comparison between methods. The UOCs used are ammonium diuranate (ADU), sodium diuranate (SDU), ammonium uranyl carbonate (AUC), and magnesium diuranate, which commonly results in a uranium hydroxide (UH) as the final product. Each of these UOCs were calcined to reach the final products amorphous uranium trioxide (UO_3), triuranium octoxide (U_3O_8), and uranium dioxide (UO_2), except for UH, which was calcined to only UO_3 and U_3O_8, not UO_2. In addition, two processing routes were used between AUC and UO_2. In the AUC-direct (AUCd) pathway, AUC was directly calcined to UO_2. In the AUC-indirect (AUCi) pathway, AUC was first calcined to U_3O_8 before being further reduced to UO_2. We will refer to these processing routes in the form UOC → calcination product for the remainder of the paper. The dataset consists of 16,688 SEM images total, covering 4172 samples (see Table 1). Each sample was imaged at four magnification levels - 10,000×, 25,000×, 50,000×, and 100,000× - to capture differences in surface morphology details. The raw images are 1024 × 880 pixels. We cropped each image into four 512 × 440-pixel images to allow them to be used on our computational resources. Further details about the syntheses of the materials and SEM procedures can be found in Schwerdt et al. [11]. Example images can be found in Fig. 5 in the Appendix.

Table 1
Number of samples per processing route. Each sample was imaged at four different magnifications. Routes refer to the calcination or reduction of magnesium diuranate (UH), sodium diuranate (SDU), ammonium diuranate (ADU), and ammonium uranyl carbonate (AUC). AUCd → UO_2 refers to the direct single step process to convert AUC to UO_2, and AUCi → UO_2 refers to the indirect two-step process to form UO_2 from AUC. See Section 2 for more detail.

Processing route	Number of samples
ADU → U_3O_8	360
AUC → U_3O_8	360
UH → U_3O_8	360
SDU → U_3O_8	124
ADU → UO_2	352
AUCd → UO_2	360
AUCi → UO_2	396
SDU → UO_2	360
ADU → UO_3	360
AUC → UO_3	348
UH → UO_3	432
SDU → UO_3	360

3. Method

Contrastive learning requires comparisons between positive and negative examples. A logical way to keep track of these examples is in a memory bank or storage data structure; however, this has computational costs when using large datasets or high-dimensional data. First described in Chen et al. [9], SimCLR is a framework for contrastive learning without such a memory bank. This framework has shown to be a top performer in the self-supervised learning field. For a minibatch of size N, starting with a given image x, the method works by drawing two different compositions of data augmentations, creating a positive pair (x_i, x_j) (now $2N$ datapoints per minibatch), then generating an encoded representation through a standard model architecture such as a ResNet [12], (y_i, y_j). After the representation, a nonlinear projection head maps the representation to a 128-dimensional space. Instead of sampling negative examples from a memory bank, the remaining $2(N-1)$ augmented examples in the minibatch are treated as negative samples, removing the need to store examples between minibatches. The contrastive loss is applied across the positive pair (z_i, z_j) which aims to identify z_j in $\{z_k\}_{k\neq i}$ for a given z_i. The loss, termed the NT-Xent loss (normalized temperature-scaled cross-entropy loss [13]), is defined as

$$\ell_{i,j} = -\log \frac{\exp(\mathrm{sim}(z_i, z_j)/\tau)}{\sigma_{k=1}^{2n} \mathbb{1}_{[k\neq i]} \exp(\mathrm{sim}(z_i, z_k)/\tau)}, \tag{1}$$

for the positive pair (i, j), and where $\mathrm{sim}(u, v) = u^T v/\|u\|\|v\|$, $\mathbb{1}_{[k\neq i]} \in \{0, 1\}$ is an indicator function evaluating to 1 if $k \neq i$, and τ is a temperature parameter. This loss closes the distance between matching pairs while driving the negative examples away. The authors found that using a projection network $g(\cdot)$ before the contrastive loss is applied significantly increased the representation quality. This network $g(\cdot)$ is thrown away after training, and the representation y is used for evaluation.

Once the encoder is fully trained, it is frozen, and a small classification network is trained on the encoded representations. In [9], the authors used a single linear layer, we chose a two-layer nonlinear multilayer perceptron (MLP) for our implementation, a common adaptation of the original version.

Chen et al. [9] showed that the encoder requires training for a long time – upwards of 1000 epochs – and benefits from large minibatch sizes. Since the classification network is small, it is quick to train and works well with standard minibatch sizes. Fortunately, this works to our advantage as the encoder network learns powerful representations of the images, allowing us to quickly generalize to unseen classes without retraining the encoder, a benefit beyond accuracy over traditional supervised classification models.

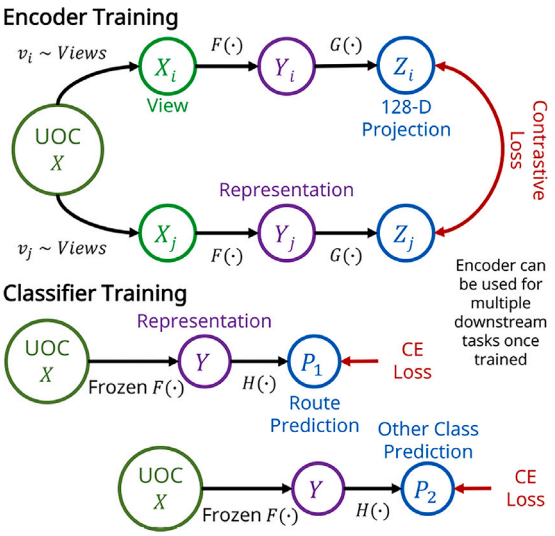

Fig. 1. Our model framework. X represents a UOC sample. *Views* refers to a family of transformations, including captured SEM views, where two distinct transformations v_i and v_j are drawn from. $f(\cdot)$ refers to a base encoder network such as ResNet50, $g(\cdot)$ refers to the projection head network, and $h(\cdot)$ refers to the classification network.

We build off the SimCLR framework and utilize the SEM magnification level as an additional "transformation". During training of the encoder, in addition to other transformations $\mathcal{T}$, two random magnifications of the same material sample are drawn and used as input images. In Ly et al. [6], the authors determined that exploiting multiple magnifications could significantly improve classification accuracy in a supervised learning context. We show that this holds true in a self-supervised context, where the contrastive method creates strong encodings. Nuclear forensics presents a unique opportunity; we can easily collect images at high magnification, showing different physical structures of the same material. We go beyond what is possible with simple transformations to the image and digital zoom. Instead, we add more information about the sample and make the task harder with each magnification or view added. The original x in our method is not simply a sample image but the uranium oxide sample itself. By applying the contrastive loss across multiple views, we force the encoder to map those views closely in encoding space, a more difficult task than transformations of single images. The method draws from co-training [14], where multiple conditionally independent views are shown to models to learn from unlabeled data and improve generalization. In co-training, models combine multiple weaker predictions to help each improve their predictions. We aim to accomplish the same, using a contrastive loss to combine two representations of different views, improving each of them. Incorporating multiple views helps the encoder learn to distinguish morphological structures more effectively. Fig. 1 shows a simple example of the framework.

When only one image is needed as input during evaluation time, we either use one magnification or do a majority vote on the final predicted class using all available magnifications.

3.1. Training procedure

We used a ResNet50 [12] backbone, pretrained on ImageNet [15], for the encoder model and a two-layer projection network for the contrastive loss. The encoder was trained for 1000 epochs (unless specified) using a minibatch size of 400 on two Nvidia RTX A6000 48 GB GPUs. We achieved the highest accuracy when using the LARS optimizer [16] and a cosine-annealing learning rate decay, starting at 1. A τ value of 0.5 was used in the NT-Xent loss.

Table 2
Validation accuracy on UOC pathway dataset, mean calculated with 5-fold cross-validation. For our method, we report both single-image and multi-magnification voting accuracy. See Section 4.1 for details.

Method	Mean accuracy
Supervised ResNet50	95.6%
Supervised MISO [6]	96.4%
Unsupervised VQ-VAE [7]	81.8%
Ours	90.2%
Ours, with multi-image voting	96.2%

For classification, a two-layer nonlinear MLP was attached to the end of the trained and frozen ResNet50 encoder, discarding the projection network. This network was trained for 30 epochs using the Adam optimizer [17] and an exponential learning rate decay, starting at 0.001 for a minibatch size of 64. The code for this project is published at github.com/jakobottar/bartali.

4. Results

Using the method described in Section 3, we demonstrate its strength through a few experiments. We show its overall classification accuracy through both single-image classification and multi-image voting. We demonstrate the robustness of the encoding by showing its generalizability to unseen classes through a series of tests. We use an additional dataset to boost classification accuracy further. Finally, we show that our magnification transformation is critical to the accuracy values seen throughout the other results.

4.1. Pathway classification

Using the encoder, we created 2048-dimensional feature vector representations for each image and classified them using a two-layer nonlinear fully-connected MLP.

We present two different approaches to classifying the images. Method A treats each sample and magnification separately, simply analyzing each image as distinct from the others. In method B, we classify each sample by analyzing each magnification of the sample, and then a final prediction is made using each view's predicted class to vote on the final class label. Method A is a more straightforward approach and is the method shown in Girard et al. [7]. Method B utilizes all available information and is most comparable to the MISO method proposed by Ly et al. [6].

Using five-fold cross-validation, the single magnification classifier (method A) achieved an accuracy of **90.2%**, surpassing the VQ-VAE method [7] by 9%, shown in Table 2.

Again using five-fold cross-validation, the multi magnification voting method (method B) scored an accuracy of **96.2%**, very close to the fully supervised MISO model, and higher than method A by 6%. This is also a reasonable scenario for practical applications where multiple magnifications of a sample can easily be acquired. The model performs very well, perfectly classifying many routes as shown for one example fold in Fig. 2. For the remainder of the paper, we will use the multi-magnification majority voting method to record model accuracy.

4.2. Generalization to unseen classes

A strength of self-supervised models is in transfer learning, or transferring the encodings to entirely new datasets. Transfer learning can also be used to generalize to classes not seen during encoder training. To test the model's performance in this context, we created two experiments: a leave-one-out experiment, and a more difficult drop-n experiment. For these experiments, we restricted the number of classes available during encoder training but added the classes back for classifier training.

Table 3
For the leave-one-out experiment, individual routes were dropped from the encoder training set then added back for classification, simulating a scenario where a new class is to be added during classification time. Accuracy is given on the overall validation set, as well as just on the dropped-out class. We include accuracy values from [7] for comparison.

Route left out	Our model		VQ-VAE [7]	
	Validation accuracy	Accuracy on left-out class	Validation accuracy	Accuracy on left-out class
None	96.2%	–	82.2%	–
ADU → U_3O_8	94.4%	89.3%	84.0%	84.2%
AUC → U_3O_8	96.2%	98.5%	84.6%	88.4%
UH → U_3O_8	96.9%	98.6%	81.1%	84.1%
SDU → U_3O_8	97.1%	80.0%	81.0%	76.3%
ADU → UO_2	96.4%	97.1%	82.9%	81.6%
AUCd → UO_2	95.6%	95.8%	81.2%	94.7%
AUCi → UO_2	95.8%	78.1%	84.5%	80.0%
SDU → UO_2	95.2%	100.0%	84.0%	89.5%
ADU → UO_3	94.1%	70.8%	84.5%	88.9%
AUC → UO_3	96.4%	78.9%	82.0%	67.8%
UH → UO_3	96.0%	71.8%	–	–
SDU → UO_3	95.7%	100.0%	79.6%	75.6%

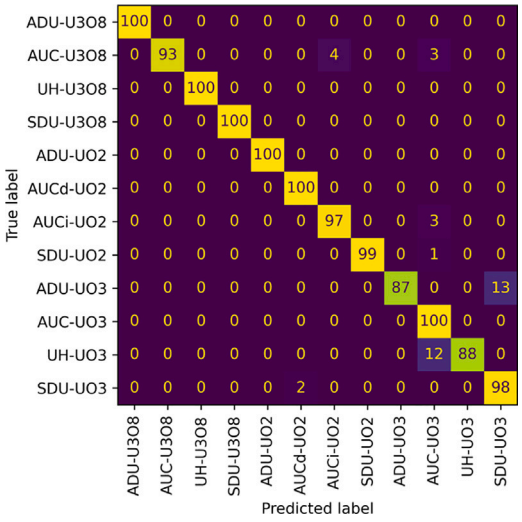

Fig. 2. A confusion matrix showing the per-class accuracy of multi-magnification method. The model performs very well, perfectly classifying many processing pathways.

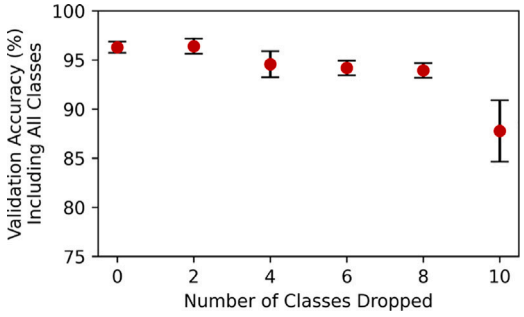

Fig. 3. In the drop-n experiment, n random routes were dropped during encoder training and added back during classifier training, simulating a scenario where many routes were missing from the training set. Accuracy is reported on the validation set, using the majority voting technique. Error bars represent ± one standard deviation across five runs.

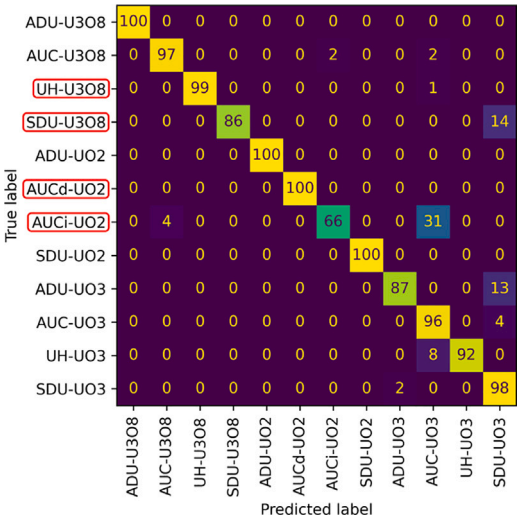

Fig. 4. A drop-four run confusion matrix, where the encoder was trained with four dropped classes (highlighted): UH → U_3O_8, SDU → U_3O_8, AUCd → UO_2, and AUCi → UO_2.

In the leave-one-out experiment, we dropped out each individual pathway to determine which pathways were most difficult to classify when left out or which were most important in training. Results from this experiment are shown in Table 3, where we see that validation accuracy does not significantly change when individual classes are dropped. The model performs well on those dropped classes, even though they were not seen during encoder training time.

Taking the experiment one step further, we dropped multiple of the 12 routes during encoder training, again adding them back during the classifier training. As shown in Fig. 3, the validation accuracy significantly decreases only when we drop over half of the routes. The confusion matrix of the drop-four run is shown in Fig. 4, with the dropped classes highlighted. It still achieves accuracy above 90%, even when trained without a third of the dataset. Even with very limited datasets, our model significantly outperforms existing methods.

4.3. Additional image sources

Going beyond magnifications, we incorporate images from a recent imaging campaign that contains images of the same material samples from two additional microscopes that can capture backscatter electron (BSE) and standard secondary electron (SE) images. The original dataset was captured on a Nova Nano SEM using SE detection. This new dataset contains images from the U_3O_8 and UO_2 pathways, imaged at the same four magnifications – 10,000×, 25,000×, 50,000×, and 100,000× – with both SE and BSE detection from Helios and Teneo SEMs. Images from different SEMs can have variations in how images appear due to differences in the beam voltages, working distances, and brightness and contrast settings. BSE detectors are more sensitive to the proton number (or Z number) of the materials being imaged. This mode of imaging can reveal the inclusion of lower Z impurities when

Table 4
Validation accuracy on UOC pathway datasets, mean calculated with 5-fold cross-validation. Accuracy reported using multi-image voting. Additional images from other SEMs were added to further improve contrastive learning performance.

Method	Mean accuracy
Supervised ResNet50	82.1%
Ours, trained on original dataset	96.2%
Ours, trained with new data	**99.0%**

Table 5
Transformation ablation study. Encoder models were trained for 150 epochs using different transformations. The accuracy is expected to be lower than those in Table 2 due to the shortened training time but the overall results are representative. In addition to the shown transformations, the images were randomly cropped and normalized.

Random flip	Color jitter	Magnification	Accuracy
✓	✓	✓	94.3%
	✓	✓	90.8%
✓		✓	**94.4%**
✓	✓		92.7%
✓			87.6%
	✓		91.1%
		✓	93.6%
			88.3%

imaging uranium oxides. We incorporate these images using the same strategy as with multiple magnifications — drawing random views that are used as input images. By using images from different SEMs and with different detection methods in addition to magnification, we aim to reduce the mutual information (MI) between views which has been shown to improve accuracy in contrastive learning [18].

We reduced the number of crops from each image from four to three in order to maintain a similar dataset size with the additional images. Adding these material views to our model significantly improved the model's accuracy to **99.0%**, shown in Table 4 using a five-fold cross-validation. Using images from other SEMs we see an increase of 3.5% over our previous model, and an increase of 16% over the supervised baseline model. This also shows we can use the variety of images that are collected to our advantage when training models. The baseline ResNet50 performs significantly worse with the new data, possibly indicating that the additional views are not as informative by themselves. As we collect more images from different SEMs over processing routes of interest, we continue to improve performance and generalization of contrastive learning-based models.

4.4. Effects of transformations

It has been shown that self-supervised contrastive models such as SimCLR depend on powerful transformations during encoder training. Our method utilizes a variety of these strong transformations as well as SEM magnification.

An ablation study was run comparing the effect of different transformations on accuracy. Results are shown in Table 5. Each combination of random horizontal and vertical flips, color jitter, and magnification were tried. Image cropping and normalization were applied to all experiments and omitted from the table. The encoder models were trained for 150 epochs using the transformations, then the evaluation model was trained using a consistent center-crop and normalize transformations. During runs with the magnification transformation disabled, all images of all samples were shown to the model so as not to lose information. The multi-magnification voting method was used to compute the final accuracy value.

While other strong transformations are important to representation quality and therefore model accuracy, this study shows that utilizing contrastive loss across magnifications is critical to high accuracy on the pathway discernment task and helps the model learn a high-quality representation. The magnifications used in this study are representative of those collected during realistic sample analysis and present views

Table 6
Images must be uniformly cropped for the model to process. This study compares the types of cropping, with and without magnification. The first column is accuracy with cropping only and the second column is accuracy with the SEM magnification transformation enabled. Encoder models were trained for 150 epochs using the transformations. Magnification significantly improved the model's accuracy on the validation set.

Crop type	Accuracy without magnification	Accuracy with magnification
Center crop	70.1%	86.5%
Random crop	81.1%	89.2%
Random resized crop	83.7%	88.6%

of the material that are not replicable by re-scaling and cropping. To validate the advantage of magnification compared to digital zoom, we ran a simple comparison experiment, whose results are shown in Table 6. Center crop crops a 256 × 256 region in the center of the image, random crop crops a random 256 × 256 square out of the image, and random resize crop randomly crops a region of the image (from 10% to 100% of the image height or width), then rescales it to 256 × 256. Random resized crop is a common image augmentation designed to induce scale invariability and replicates random digital zooms. We see that the models trained with magnification enabled significantly outperform the ones without magnification and that resizing or digital zoom is an insufficient replacement for magnification.

5. Conclusions and future work

In this paper, we introduced an application of self-supervised contrastive learning. We showed that using multiple views of material simultaneously in contrastive learning significantly improves the quality of the representation of uranium oxide SEM images in a way not replicable by digital manipulations alone. Our model is able to outperform competitive models in pathway discernment significantly and shows a strong ability to generalize to unseen classes, allowing our method to be quickly updated when new data is provided. We showed that we can improve the model's accuracy even more by using views from different SEMs and detectors. As more images of UOC material are collected and compiled into datasets, contrastive methods like the one shown here become even stronger. We hope our method can be at the forefront of nuclear forensics image processing.

CRediT authorship contribution statement

Jakob Johnson: Conceptualization, Methodology, Software, Validation, Writing – original draft, Data curation. **Luther McDonald:** Project administration, Funding acquisition. **Tolga Tasdizen:** Conceptualization, Resources, Writing – review & editing, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

A representative sample dataset has been made available with Mendeley Data, doi: https://doi.org/10.17632/74k85vybmk.1.

Acknowledgments

This work is supported by the Department of Energy's National Nuclear Security Administration, Office of Defense Nuclear Nonproliferation Research and Development.

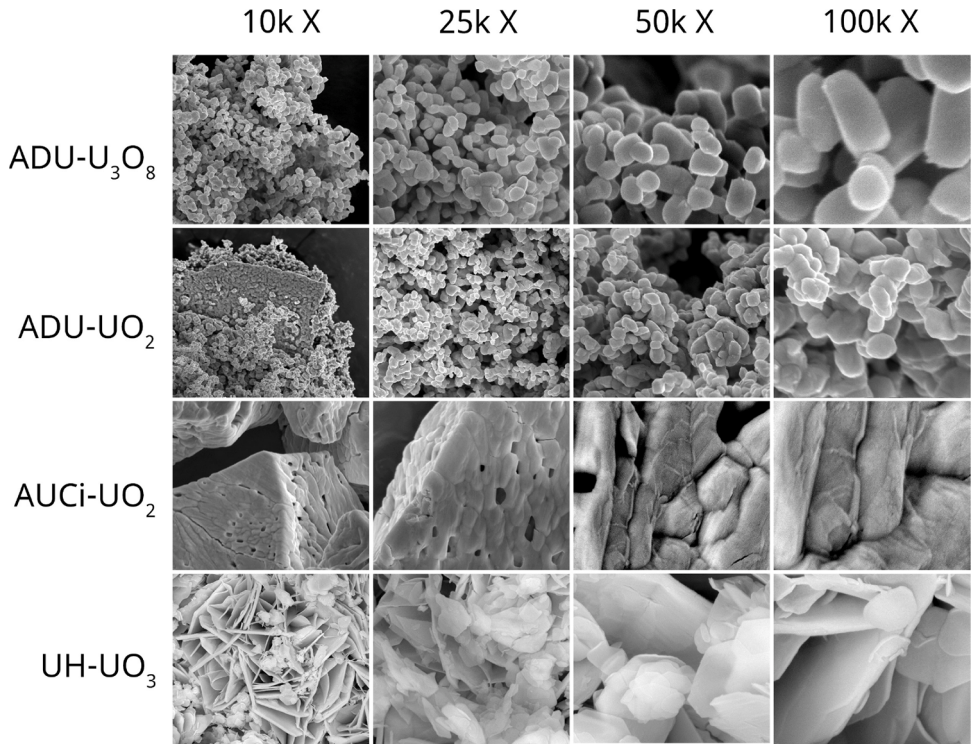

Fig. 5. Examples of Secondary Electron SEM images of four UOC processing pathways used in the study. From left to right, the magnification level increases from 10,000×, 25,000×, 50,000×, and 100,000×.

Appendix

See Fig. 5.

References

[1] K. Mayer, M. Wallenius, I. Ray, Nuclear forensics—a methodology providing clues on the origin of illicitly trafficked nuclear materials, Analyst 130 (4) (2005) 433–441.

[2] L. Fongaro, D.M. Lin Ho, K. Kvaal, K. Mayer, V.V. Rondinella, Application of the angle measure technique as image texture analysis method for the identification of uranium ore concentrate samples: New perspective in nuclear forensics, Talanta 152 (2016) 463–474.

[3] S.T. Heffernan, N.-C. Ly, B.J. Mower, C. Vachet, I.J. Schwerdt, T. Tasdizen, L.W. McDonald Iv, Identifying surface morphological characteristics to differentiate between mixtures of U_3O_8 synthesized from ammonium diuranate and uranyl peroxide, Radiochim. Acta 108 (1) (2019) 29–36.

[4] I.J. Schwerdt, A. Brenkmann, S. Martinson, B.D. Albrecht, S. Heffernan, M.R. Klosterman, T. Kirkham, T. Tasdizen, L.W. McDonald Iv, Nuclear proliferomics: A new field of study to identify signatures of nuclear materials as demonstrated on alpha-UO3, Talanta 186 (2018) 433–444.

[5] A.B. Hanson, R.N. Lee, C. Vachet, I.J. Schwerdt, T. Tasdizen, L.W. McDonald, Quantifying impurity effects on the surface morphology of α -U_3O_8, Anal. Chem. 91 (15) (2019) 10081–10087.

[6] C. Ly, C. Vachet, I. Schwerdt, E. Abbott, A. Brenkmann, L.W. McDonald, T. Tasdizen, Determining uranium ore concentrates and their calcination products via image classification of multiple magnifications, J. Nucl. Mater. 533 (2020) 152082.

[7] M. Girard, A. Hagen, I. Schwerdt, M. Gaumer, L. McDonald, N. Hodas, E. Jurrus, Uranium oxide synthetic pathway discernment through unsupervised morphological analysis, J. Nucl. Mater. 552 (2021) 152983.

[8] K. He, H. Fan, Y. Wu, S. Xie, R. Girshick, Momentum contrast for unsupervised visual representation learning, 2020, URL: http://arxiv.org/abs/1911.05722, arXiv:1911.05722 [cs].

[9] T. Chen, S. Kornblith, M. Norouzi, G. Hinton, A simple framework for contrastive learning of visual representations, 2020, URL: http://arxiv.org/abs/2002.05709, arXiv:2002.05709 [cs, stat].

[10] N. Zhao, Z. Wu, R.W.H. Lau, S. Lin, What makes instance discrimination good for transfer learning? 2021, URL: http://arxiv.org/abs/2006.06606, arXiv:2006.06606 [cs].

[11] I.J. Schwerdt, C.G. Hawkins, B. Taylor, A. Brenkmann, S. Martinson, L.W. McDonald Iv, Uranium oxide synthetic pathway discernment through thermal decomposition and morphological analysis, Radiochim. Acta 107 (3) (2019) 193–205.

[12] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, 2015, URL: http://arxiv.org/abs/1512.03385, arXiv:1512.03385 [cs].

[13] K. Sohn, Improved deep metric learning with multi-class N-pair loss objective, in: Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS '16, Curran Associates Inc., Red Hook, NY, USA, 2016, pp. 1857–1865, Event-place: Barcelona, Spain.

[14] A. Blum, T. Mitchell, Combining labeled and unlabeled data with co-training, in: Proceedings of the Eleventh Annual Conference on Computational Learning Theory, ACM, Madison Wisconsin USA, 1998, pp. 92–100, URL: https://dl.acm.org/doi/10.1145/279943.279962. http://dx.doi.org/10.1145/279943.279962.

[15] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A.C. Berg, L. Fei-Fei, ImageNet large scale visual recognition challenge, 2015, URL: http://arxiv.org/abs/1409.0575, arXiv:1409.0575 [cs].

[16] Y. You, I. Gitman, B. Ginsburg, Large batch training of convolutional networks, 2017, URL: http://arxiv.org/abs/1708.03888, arXiv:1708.03888 [cs].

[17] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, 2017, URL: http://arxiv.org/abs/1412.6980, arXiv:1412.6980 [cs].

[18] Y. Tian, C. Sun, B. Poole, D. Krishnan, C. Schmid, P. Isola, What makes for good views for contrastive learning? 2020, URL: http://arxiv.org/abs/2005.10243, arXiv:2005.10243 [cs].