

# Model editing for distribution shifts in uranium oxide morphological analysis

Davis Brown\*, Cody Nizinski\*, Madelyn Shapiro,  
Corey Fallon, Tianzhixi Yin, Henry Kvinge, and Jonathan H. Tu  
Pacific Northwest National Laboratory

davis.brown, cody.nizinski, jonathan.tu @pnnl.gov

## Abstract

*Deep learning still struggles with certain kinds of scientific data. Notably, pretraining data may not provide coverage of relevant distribution shifts (e.g., shifts induced via the use of different measurement instruments). We consider deep learning models trained to classify the synthesis conditions of uranium ore concentrates (UOCs) and show that **model editing** is particularly effective for improving generalization to distribution shifts common in this domain. In particular, model editing outperforms finetuning on two curated datasets comprising of micrographs taken of  $U_3O_8$  aged in humidity chambers and micrographs acquired with different scanning electron microscopes, respectively.*

## 1. Introduction

The morphological features of nuclear materials can provide information about the processing conditions used to produce the materials and the history of the materials after their production [18]. For uranium ore concentrates (UOCs) the precipitation chemistry and calcination conditions have been shown to impart certain morphological features that allow particles observed by scanning electron microscopy (SEM) to be correlated back to the synthesis conditions [27]. Aging studies have looked at how UOC properties and particle morphologies change over time during exposure to different storage conditions, particularly humidity and temperature [8, 9, 23, 25, 28, 29, 31, 33].

A number of analytical methods have been developed for morphological analysis of nuclear materials, including qualitative descriptions [30], quantitative measurement of particles by image segmentation [22] or other particle sizing methods, texture analysis [4, 5], and machine learning. Deep learning models for morphological analysis have been trained in supervised [17], self-supervised [13], and unsupervised settings [6]. The robustness of these models to out-of-distribution (OOD) data has been investigated [21],

along with methods for uncertainty quantification (UQ) or calibration of prediction confidence [7] and source-free domain adaptation (SFDA) for covariate shifts [16].

Computer vision models for predicting the provenance of nuclear materials must be robust to statistic-level variations (i.e., the appearance of images taken with various microscopes and in varying settings), and to some extent, feature-level variations (i.e., changes to the actual morphological characteristics). Whereas SFDA has only been demonstrated for the former, model editing methods [2, 20] could potentially be used for both types of distribution shifts and can be applied to existing validated models incrementally as new kinds of distribution shifts are discovered with lower computational needs than other methods.

## Contributions

Our contributions can be summarized as follows:

- Model editing methods generalize well on feature-level distribution shifts caused by aging uranium oxides under diel cycling humidity and temperature conditions.
- We compare editing methods and find that low-rank editing generally outperforms surgical finetuning.
- Both editing and surgical finetuning outperform full-model finetuning.

## 2. Description of data

The training and validation datasets consist of scanning electron microscope (SEM) images collected of UOCs representing five precipitation pathways — ammonium diuranate (ADU), ammonium uranyl carbonate (AUC), magnesium diuranate (MDU), sodium diuranate (SDU), and uranyl peroxide ( $UO_4 \cdot H_2O$ ) — that have been converted to three uranium oxides — uranium trioxide ( $UO_3$ ), triuranium octoxide ( $U_3O_8$ ), and uranium dioxide ( $UO_2$ ). The micrographs were collected using an FEI Nova NanoSEM 630 scanning electron microscope with the through lens detector (TLD) operating in secondary electron (SE) mode. More complete descriptions of the material synthesis and data collection can be found elsewhere [1, 11, 17, 27]. Previous work has shown that classification models trained on earlier

\*These authors contributed equally to this work.

versions of this dataset fail to generalize to new data that differs by process history and/or imaging parameters [21].

We consider two different concept distribution shifts, an **aging** shift and a **detector** shift:

- The dataset for the **aging** editing task uses micrographs of materials that were initially  $U_3O_8$  produced by calcining either AUC or  $UO_4 \cdot H_2O$ . The starting  $U_3O_8$  was then aged under conditions that cycled between high temperature/humidity and low temperature/humidity following a 24-hour daily cycle; samples were pulled at time steps of  $\{0, 14, 36, 43, 54, 60\}$  days to collect morphology data by SEM and measure crystallographic changes by X-ray diffraction (XRD) [8].
- The **detector** editing task uses micrographs of the same materials as the train and validation sets but collected on an FEI Teneo SEM with the “T2” SE detector [16].


### 3. Methods for model updating

Model editing makes tweaks to model weights to incorporate new facts [10, 19, 20], semantics [2], and behaviors [12]. Such methods have been used to patch model errors in the form of spurious correlations [26] and confusion [24] with only a single input example. We focus on the setting of *natural distribution shifts*, where we want to adapt models to inherent variations in the data — largely in the form of input/statistical-level variations (the **detector** editing task) and feature-level variations (the **aging** editing task).

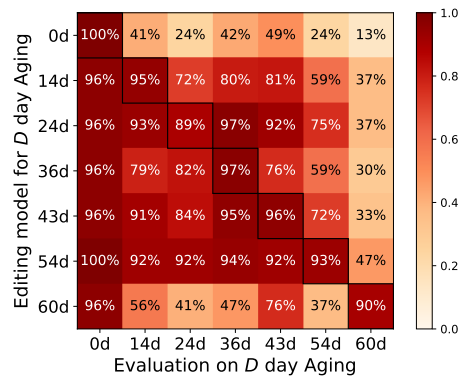
Let  $f$  be a neural network defined as a composition of layers  $f = f_L \circ f_{L-1} \circ \dots \circ f_2 \circ f_1$ , where  $f_l(x) = \sigma(W_l x)$  and  $W_l$  is an  $n_l \times n_{l-1}$  matrix for each  $l = 1, \dots, L$  with an activation function  $\sigma$ . For any layer, let  $f_{\leq l}$  denote the composition of the first  $l$  layers:  $f_{\leq l} \circ f_{l-1} \circ \dots \circ f_2 \circ f_1$ . In contrast with other work, e.g. [2, 19, 24], we do not use exemplar pairs  $(x, x', y)$  of transformed images  $x'$  with a fixed label  $y$  to update a model so that  $f_{\leq l}(x) \approx f_{\leq l}(x')$ . Instead, we update our model on new pairs of datapoints  $x'$  and labels  $y'$  from either the aging shift or the detector shift and use stochastic gradient descent (SGD) so that  $f(x') \approx y'$ . This is a fairly significant difference from some of the editing methods discussed above, however we keep the editing terminology because our methods modify single-layers with constrained updates and/or a limited number of examples.

We next describe the two model update methods we use in this work. We compare these update methods to a baseline of *full finetuning*, where all model weights  $W_1, \dots, W_L$ , are updated with stochastic gradient descent (SGD) to minimize the mean squared error (MSE) of  $f(x')$  and  $y'$ .


**Low-rank model editing:** the edit is an update  $W_l \leftarrow W_l + UV^T$  where  $U$  and  $V$  are low-rank matrices learned via SGD to minimize the MSE of  $f(x')$  and  $y'$ ; in our experiments we fix the rank to  $r = 2$ . The model weights  $W_1, \dots, W_L$ , are kept frozen. We use the implementation



(a) Surgical finetuning accuracy



(b) Low-rank editing accuracy



(c) Baseline model (unedited) accuracy

Figure 1. Comparison of how well models finetuned on an aging dataset  $D$  generalize to other ages. Both surgical finetuning ((a)) and low-rank editing ((b)) generalize to earlier ages. Also, all edited models outperform the baseline accuracy in (c). Note that because we drop runs that incur an accuracy drop on the baseline validation set of  $> 1.5\%$ , there are no successful full-model finetuning runs. See Fig. 2 for a more permissive threshold.

described in [2] where  $U$  and  $V$  are  $1 \times 1$  convolutions when applied to the weights of convolutional layers.

**Surgical finetuning:** we largely follow [14], updating only a subset of layers with SGD to minimize the MSE of  $f(x')$

and  $y'$ . In order to make a more precise comparison with low-rank model editing, we deviate from [14] and optimize only a single layer,  $W_l$ , rather than blocks of layers. This is also called *local finetuning* in [2].

## 4. Experiments

For both datasets, our goal is to adapt the model for the new domain while maintaining the model’s performance on the original dataset. On the aging dataset, this means updating a model to be accurate on the set of images and labels  $(x', y')$  that have undergone the  $D$  24-hour cycles of aging, while maintaining performance on the original unaged datapoints  $(x, y)$ . Similarly, for the detector dataset the model is updated to be accurate on datapoints taken from both the T2 SE detector and the original detector.

We use a ConvNeXt-Small model [15] trained with SGD on the original SEM image dataset with a validation accuracy of 97.9%. For each of the model updating methods, we perform a coarse grid search over learning rates for each convolutional layer, and then a finer-grained grid search over the two best performing layers for each method. To choose the best performing hyperparameters for an update, we employ the hyperparameter selection strategy from [2], where we select 50% of the dataset examples (for a given age  $D$  for the aging dataset or for the images taken from the T2 SE detector) to learn the update and to select the best hyperparameters and the remaining 50% to test performance. A hyperparameter run is dropped if it causes the model performance on the original SEM validation set to drop below a given threshold (e.g., a threshold of 1.5% drop for Fig. 1 and a substantially more permissive 7% threshold for Fig. 2).

### 4.1. Aging Experiments

The aging editing task comprises seven different material aging time steps. We use this dataset structure to evaluate how well models edited with images of age  $D$  generalize to other durations of particle aging. In other words, we ask how well a model update captures the relevant *morphology* of aging as opposed to the *particulars* of the collected dataset/aging timing. For example, we examine how well a material aged for  $D = 43$  days performs on unaged control samples (0 days), earlier ages (aging for 14, 24, and 36 days, respectively) and later ages (54 and 60 days).

Fig. 1 summarizes the results of the aging experiments. The diagonal elements capture the performance of a model on the targeted aging duration whereas off-diagonal elements capture the generalization of those models to other aging durations. The performance of the original model (without any updates) is plotted in Fig. 1c. Because all full-model finetuning runs incurred substantial accuracy drops on the base (non-aged) validation set, we also plot performance for a more permissive threshold in Fig. 2. Note that all of the model update methods — surgical finetuning, low-

rank editing, and even the baseline full-model finetuning for the permissive threshold — exhibit partial generalization across different ages. The generalization to shorter aging durations is stronger than for longer ones.

Comparing rows of Fig. 1a to those of Fig. 1b, we see that low-rank editing outperforms surgical finetuning on the targeted aging duration (e.g., editing a model on 43-day aging examples and testing on held-out 43-day aging examples) as well as on other aging datasets (e.g., editing a model on 43-day aging examples and testing on hold-out data for  $D \neq 43$ ). For example, in the aging tasks for 43 days and 54 days, low-rank editing has better editing performance for the target editing task by 15% for 43 days and 44% for 54 days of aging. For the off-target editing generalization to other aging sets, low-rank editing outperforms surgical fine-tuning as much as 27%. Both update methods increase model aging accuracy over the baseline model in Fig. 1c and over full finetuning in Fig. 2, and are targeted, i.e., do not pay substantial performance costs on the original SEM dataset. We expect that the better targeting of editing is due to regularization, where we update only a single-layer for surgical finetuning and single-layer with a low-rank restriction for low-rank editing.

Finally, we note that all update methods have worse generalization on the 60-day aging data. While the updates on the 60-day aging duration achieve reasonable performance on the held-out 60-day set, they do not generalize well to  $D \neq 60$ . Likewise, updating a model for  $D \neq 60$  leads to relatively minimal increases in accuracy for  $D = 60$ . We hypothesize that this poor performance is due to slightly different collection conditions for the 60-day micrographs, where a different scientist collected the SEM samples. This suggests that the success of model updates is sensitive to the image collection parameters; rather than editing for the changes in surface morphology seen in  $D = 60$  materials, edits were made on the style (e.g., brightness, contrast) of the image. We expect that editing performance will increase when using editing prototypes that better reflect the diversity in SEM images, however we leave this to future work.

### 4.2. SEM Detector Experiments

The detector editing task captures input-level and statistical changes in the distribution of SEM images due to using a different scanning electron microscope. In general, this task is more challenging for the editing methods under consideration. While editing does provide performance improvement over the baseline unedited model, the accuracy increases are modest. Box plots of model accuracy across five random seeds for each hyperparameter are plotted in Fig. 3. Like in the aging task, we learn model updates for the detector task with 50% of the SEM images, test on the remaining 50%, and drop runs that incur an accuracy drop on the original validation above a certain threshold (drop-

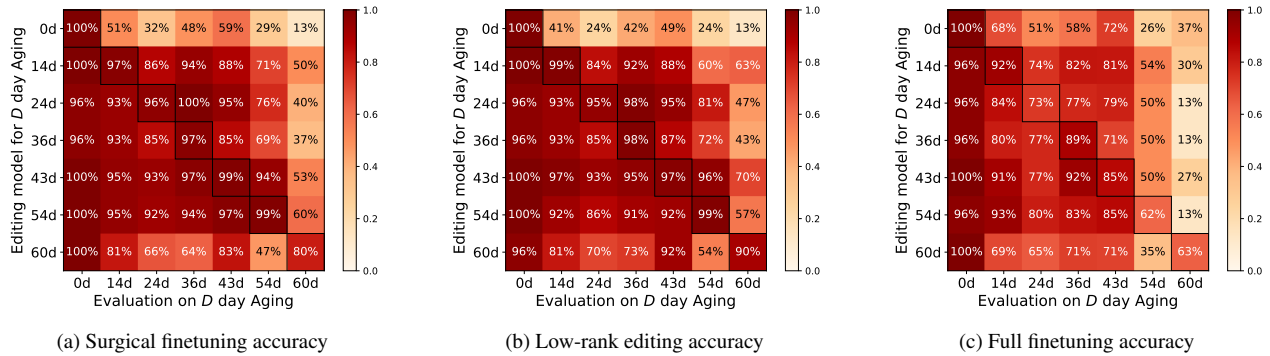


Figure 2. Comparison for how well models updated for an aging dataset  $D$  generalize to other ages, allowing up to a 7% drop in accuracy on the original validation set. Even for this more aggressive editing setting (as opposed to Fig. 1, which used a threshold of 1.5%), surgical finetuning and low-rank editing tend to outperform full finetuning. The aging accuracies for the unedited model are given in Fig. 1c.

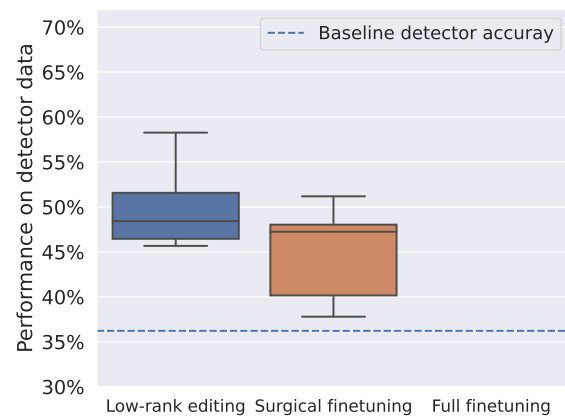
ping runs with performance drops greater than 1.5% for Fig. 3a and 7% for Fig. 3b).

For the more restrictive validation dropping threshold of 1.5% in Fig. 3a, low-rank editing tends to outperform surgical finetuning. Like in the aging editing task runs in Fig. 1, there were no training runs from full finetuning that met this threshold. With the more permissive threshold of 7% in Fig. 3b, surgical finetuning and low-rank editing again generally outperform full finetuning.


## 5. Conclusions and Future Work

We find that we can update models trained to predict the processing conditions of uranium ore concentrates to better handle a common domain-specific distribution shifts. Model editing methods perform particularly well on the set of aging distribution shifts, representing feature-level morphological changes in the SEM images. This has the potential to be particularly useful for designing future aging studies and incorporating their data, where model edits for a duration  $D$  can be reliably expected to generalize to shorter durations (and to a lesser but still notable extent, longer ones).

The detector distribution shift, reflecting lower-level variation in the statistics of input samples, proved more difficult. We are particularly optimistic about two directions for future work. First, we think that more targeted mixtures of image examples used for editing may better capture relevant directions of variation. For example, using exemplars from multiple detectors may allow for better generalization on a held-out detectors. Second, strong results in this domain have been achieved by using generative models to aide domain adaptation [16]. Related to domain adaptation, recent work has leveraged generative models to automatically discover failure modes [32] and debug them [3]. In future work, we would be excited to use comparable approaches to



(a) Detector accuracy, removing models with an original validation accuracy drop greater than 1.5%



(b) Detector accuracy, removing models with an original validation accuracy drop greater than 7%.

Figure 3. Model editing performance on the T2 SE detector dataset.

identify relevant failure modes with generative models and to edit models using generated exemplars. We expect such an approach to balance the domain coverage of generative methods with the targeting of model editing.

## 6. Acknowledgements

This research was supported by the Mathematics for Artificial Reasoning in Science (MARS) initiative at Pacific Northwest National Laboratory. It was conducted under the Laboratory Directed Research and Development (LDRD) Program at Pacific Northwest National Laboratory (PNNL), a multiprogram National Laboratory operated by Battelle Memorial Institute for the U.S. Department of Energy under Contract DE-AC05-76RL01830.

## References

- [1] Erik C Abbott, Hiram E O'Connor, Cody A Nizinski, Logan D Gibb, Elijah W Allen, and Luther W McDonald IV. Thermodynamic evaluation of the uranyl peroxide synthetic route on morphology. *Journal of Nuclear Materials*, 561: 153533, 2022. [1](#)
- [2] David Bau, Steven Liu, Tongzhou Wang, Jun-Yan Zhu, and Antonio Torralba. Rewriting a deep generative model. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 351–369. Springer, 2020. [1](#), [2](#), [3](#)
- [3] Atoosa Chegini and Soheil Feizi. Identifying and mitigating model failures through few-shot clip-aided diffusion generation. *arXiv preprint arXiv:2312.05464*, 2023. [4](#)
- [4] Lorenzo Fongaro, Doris Mer Lin Ho, Knut Kvaal, Klaus Mayer, and Vincenzo 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*, 152:463–474, 2016. [1](#)
- [5] L Fongaro, M Marchetti, IB Lande, M Wallenius, and K Mayer. Development of a new approach for a rapid identification and classification of uranium powders using colour, image texture and spectroscopy signatures. *Publications Office of the European Union, Luxembourg*, 2021. [1](#)
- [6] M Girard, A Hagen, I Schwerdt, M Gaumer, L McDonald, N Hodas, and E Jurrus. Uranium oxide synthetic pathway discernment through unsupervised morphological analysis. *Journal of Nuclear Materials*, 552:152983, 2021. [1](#)
- [7] Alex Hagen, Karl Pazdernik, Nicole LaHaye, and Marjolein Oostrom. Dbcal: Density based calibration of classifier predictions for uncertainty quantification. *arXiv preprint arXiv:2204.00150*, 2022. [1](#)
- [8] Alexa B Hanson, Cody A Nizinski, and Luther W McDonald IV. Effect of diel cycling temperature, relative humidity, and synthetic route on the surface morphology and hydrolysis of  $\alpha$ -u3o8. *ACS omega*, 6(28):18426–18433, 2021. [1](#), [2](#)
- [9] Alexa B Hanson, Ian J Schwerdt, Cody A Nizinski, Rachel Nicholls Lee, Nicholas J Mecham, Erik C Abbott, Sean Heffernan, Adam Olsen, Michael R Klosterman, Sean Martinson, et al. Impact of controlled storage conditions on the hydrolysis and surface morphology of amorphous-u3o8. *ACS omega*, 6(12):8605–8615, 2021. [1](#)
- [10] Peter Hase, Mona Diab, Asli Celikyilmaz, Xian Li, Zornitsa Kozareva, Veselin Stoyanov, Mohit Bansal, and Srinivasan Iyer. Methods for measuring, updating, and visualizing factual beliefs in language models. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2714–2731, Dubrovnik, Croatia, 2023. Association for Computational Linguistics. [2](#)
- [11] Sean T Heffernan, Nhat-Cuong Ly, Brock J Mower, Clement Vachet, Ian J Schwerdt, Tolga Tasdizen, and Luther W McDonald IV. Identifying surface morphological characteristics to differentiate between mixtures of u3o8 synthesized from ammonium diuranate and uranyl peroxide. *Radiochimica Acta*, 108(1):29–36, 2019. [1](#)
- [12] Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. *International Conference on Learning Representations*, 2022. [2](#)
- [13] Jakob Johnson, Luther McDonald, and Tolga Tasdizen. Improving uranium oxide pathway discernment and generalizability using contrastive self-supervised learning. *Computational Materials Science*, 233:112748, 2024. [1](#)
- [14] Yoonho Lee, Annie S Chen, Fahim Tajwar, Ananya Kumar, Huaxiu Yao, Percy Liang, and Chelsea Finn. Surgical fine-tuning improves adaptation to distribution shifts. In *NeurIPS 2022 Workshop on Distribution Shifts: Connecting Methods and Applications*, 2022. [2](#), [3](#)
- [15] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *CVPR*, 2022. [3](#)
- [16] Cuong Ly, Cody Nizinski, Alex Hagen, Luther McDonald IV, and Tolga Tasdizen. Improving robustness for model discerning synthesis process of uranium oxide with unsupervised domain adaptation. *Frontiers in Nuclear Engineering*, 2:1230052. [1](#), [2](#), [4](#)
- [17] Cuong Ly, Clement Vachet, Ian Schwerdt, Erik Abbott, Alexandria Brenkmann, Luther W McDonald, and Tolga Tasdizen. Determining uranium ore concentrates and their calcination products via image classification of multiple magnifications. *Journal of Nuclear Materials*, 533:152082, 2020. [1](#)
- [18] Luther W McDonald IV, Kari Sentz, Alex Hagen, Brandon W Chung, Cody A Nizinski, Ian J Schwerdt, Alexa Hanson, Scott Donald, Richard Clark, Glenn Sjoden, et al. Review of multi-faceted morphologic signatures of actinide process materials for nuclear forensic science. *Journal of Nuclear Materials*, page 154779, 2023. [1](#)
- [19] Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in gpt. *Advances in Neural Information Processing Systems*, 35: 17359–17372, 2022. [2](#)
- [20] Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D. Manning. Fast model editing at scale. *ICLR*, 2022. [1](#), [2](#)

- [21] Cody A Nizinski, Cuong Ly, Clement Vachet, Alex Hagen, Tolga Tasdizen, and Luther W McDonald IV. Characterization of uncertainties and model generalizability for convolutional neural network predictions of uranium ore concentrate morphology. *Chemometrics and Intelligent Laboratory Systems*, 225:104556, 2022. [1](#), [2](#)
- [22] Adam M Olsen, Bryony Richards, Ian Schwerdt, Sean Heffernan, Robert Lusk, Braxton Smith, Elizabeth Jurrus, Christy Ruggiero, and Luther W McDonald IV. Quantifying morphological features of  $\alpha$ -u3o8 with image analysis for nuclear forensics. *Analytical chemistry*, 89(5):3177–3183, 2017. [1](#)
- [23] Adam M Olsen, Ian J Schwerdt, Bryony Richards, and Luther W McDonald IV. Quantification of high temperature oxidation of u3o8 and uo2. *Journal of Nuclear Materials*, 508:574–582, 2018. [1](#)
- [24] Indu Panigrahi, Ryan Manzuk, Adam Maloof, and Ruth Fong. Improving data-efficient fossil segmentation via model editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 4829–4838, 2023. [2](#)
- [25] Kevin J Pastoor, Shane L Robinson, R Allan Greenwell, Camila V Quintero Hilsaca, Jenifer C Shafer, and Mark P Jensen. Understanding uranium oxide hardening during prolonged storage. *Radiochimica Acta*, 108(12):943–953, 2020. [1](#)
- [26] Shibani Santurkar, Dimitris Tsipras, Mahalaxmi Elango, David Bau, Antonio Torralba, and Aleksander Madry. Editing a classifier by rewriting its prediction rules. In *Advances in Neural Information Processing Systems*, pages 23359–23373. Curran Associates, Inc., 2021. [2](#)
- [27] Ian J Schwerdt, Casey G Hawkins, Bryan Taylor, Alexandria Brenkmann, Sean Martinson, and Luther W McDonald IV. Uranium oxide synthetic pathway discernment through thermal decomposition and morphological analysis. *Radiochimica Acta*, 107(3):193–205, 2019. [1](#)
- [28] Lucas E Sweet, Thomas A Blake, Charles H Henager, Shenyang Hu, Timothy J Johnson, David E Meier, Shane M Peper, and Jon M Schwantes. Investigation of the polymorphs and hydrolysis of uranium trioxide. *Journal of Radioanalytical and Nuclear Chemistry*, 296:105–110, 2013. [1](#)
- [29] Alison L Tamasi, Kevin S Boland, Kenneth Czerwinski, Jason K Ellis, Stosh A Kozimor, Richard L Martin, Alison L Pugmire, Dallas Reilly, Brian L Scott, Andrew D Sutton, et al. Oxidation and hydration of u3o8 materials following controlled exposure to temperature and humidity. *Analytical chemistry*, 87(8):4210–4217, 2015. [1](#)
- [30] Alison L Tamasi, Leigh J Cash, Christopher Eley, Reid B Porter, David L Pugmire, Amy R Ross, Christy E Ruggiero, Lav Tandon, Gregory L Wagner, Justin R Walensky, et al. A lexicon for consistent description of material images for nuclear forensics. *Journal of Radioanalytical and Nuclear Chemistry*, 307:1611–1619, 2016. [1](#)
- [31] Alison L Tamasi, Leigh J Cash, William Tyler Mullen, Alison L Pugmire, Amy R Ross, Christy E Ruggiero, Brian L Scott, Gregory L Wagner, Justin R Walensky, and Marianne P Wilkerson. Morphology of u 3 o 8 materials following storage under controlled conditions of temperature and relative humidity. *Journal of Radioanalytical and Nuclear Chemistry*, 311:35–42, 2017. [1](#)
- [32] Olivia Wiles, Isabela Albuquerque, and Sven Gowal. Discovering bugs in vision models using off-the-shelf image generation and captioning. *arXiv preprint arXiv:2208.08831*, 2022. [4](#)
- [33] Marianne P Wilkerson, Sarah C Hernandez, W Tyler Mullen, Andrew T Nelson, Alison L Pugmire, Brian L Scott, Elizabeth S Sooby, Alison L Tamasi, Gregory L Wagner, and Justin R Walensky. Hydration of  $\alpha$ -uo 3 following storage under controlled conditions of temperature and relative humidity. *Dalton Transactions*, 49(30):10452–10462, 2020. [1](#)