On the impact of pre-training datasets for matching dendritic identifiers using residual nets

Baaz Jhaj
Geometric Media Lab, Arizona State University
Tempe, Arizona, USA
bjhaj@asu.edu

Ankita Shukla
Computer Science and Engineering, University of Nevada, Reno
Reno, Nevada, USA
ankitas@unr.edu

Pavan Turaga
Geometric Media Lab, Arizona State University
Tempe, Arizona, USA
Pavan.Turaga@asu.edu

Michael Kozicki
School of Electrical, Computer, and Energy Engineering, Arizona State University
Tempe, Arizona, USA
michael.kozicki@asu.edu

Figure 1: Unique dendrites from the fine grained dendrite dataset

ABSTRACT

Dendrites are easy to synthesize branching structures that exhibit randomness; yet they are unique, non-repeatable, and identifiable with the right algorithmic innovations. This has created a novel application area where manufactured dendritic structures are being used as product identifiers - essentially "fingerprints for things". Unlike barcodes, which are linear structures, dendrites exhibit spatial randomness. This, coupled with a unique optical signal generated by light scattering from material inhomogeneities, ensures that each dendrite is unique and unclonable. While there have not yet been any established methods on reading dendritic patterns for verification using image data, identifying dendrites using computer vision techniques could have high potential. Due to limited data and low variance, dendrite identification can be considered to be a fine-grained classification task. In this paper, we examine how the selection of pre-trained models influences dendrite classification. The dendrites we work with share similarity to human fingerprints, thus we begin with a model trained for matching fingerprint data to extract features relevant to dendrites. Additionally, we explore broader pre-training approaches, using ImageNet-1K for our second model and ImageNet-21K for our third model. Surprisingly, our results indicate that even with the visual similarity with human fingerprints, more general pre-training with common image datasets achieves better performance on dendrite classification.

CCS CONCEPTS

• Computing methodologies → Machine learning; Bio-inspired approaches; Learning paradigms.

KEYWORDS

Fine-grained classification, pre-training, dendrites

ACM Reference Format:
Baaz Jhaj, Ankita Shukla, Pavan Turaga, and Michael Kozicki. 2024. On the impact of pre-training datasets for matching dendritic identifiers using residual nets. In Proceedings of International Workshop on Artificial Intelligence for Signal, Image Processing and Multimedia (AI-SIPM 2024). ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3643487.3662168

AI-SIPM 2024, June 10–14, 2024, Phuket, Thailand

ACM ISBN 979-8-4007-0548-9/24/06
https://doi.org/10.1145/3643487.3662168

1 INTRODUCTION

Dendrites are branching structures that look similar to fingerprints, and other naturally random structures found in nature. The use of dendrites in the supply chain is a new and emerging area of research [10, 15]. While blockchain technologies exist to create a digital record of product circulation within the supply chain, dendrites act as a physically unclonable function, providing another level of

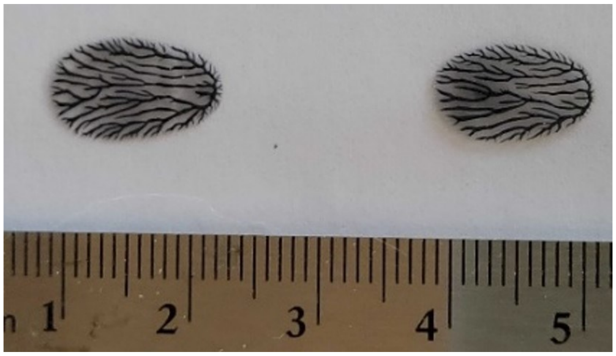

Figure 2: Dendrites with a scale comparison. Scale in cm.

security while linking the physical world to the cloud. The entropy of a dendrite's manufacturing process results in uniqueness among all dendrites, and due to the complexity of a dendrite's appearance, it is very dificult to replicate or physically remove without altering [9].

Current research on dendrite identification is ongoing, and published results are very limited. Most methodologies use 2D feature point detection and matching (c.f. [1]). As machine learning approaches provide a more scalable approach for learning image features, they offer greater potential for secure identification. In this paper, we analyze a few popular approaches for image matching using features learned from deep neural architectures, and apply them toward dendrite classification. We investigate different pre-training architectures and datasets for eficacy. We employ models pre-trained on ImageNet-1K and ImageNet-21K. We also explore more specific pre-training through a custom fingerprint dataset. We assess the results of these different pre-training choices on our target dataset.

In this paper, we analyze a few popular approaches for image matching using features learned from deep neural architectures, and apply them toward dendrite classification. We investigate different pre-training architectures and datasets for eficacy. We employ resnet models pre-trained on ImageNet-1K and ImageNet-21K. We also explore more specific pre-training through a custom fingerprint dataset. We assess the results of these different pre-training choices on our target dataset.

The main aim of this paper is to set a baseline on the dendrite classification problem with deep-learning methods and emphasizes the challenges of developing solutions for small yet demanding datasets commonly encountered in real-world applications.

2 DATASET CREATION

We have focused on a dendrites formation method that has shown great promise for high speed, high volume, low-cost manufacturing based on the Saffman-Taylor effect in viscous fluids [13]. In this scheme, we compress a few microliters of a bulk-produced low-cost fluid mixture between two surfaces and then separate the surfaces to form a dendritic shape in the fluid, which is then allowed to harden into a permanent trigger pattern. This is both simple and inexpensive and lends itself to roll-to-roll volume manufacturing using existing printing equipment. The patterns also have a subtle relief, with ridges typically in the order of 20 µm high in branches that are several hundred µm wide. This allows material inhomogeneities to generate a 3D optical signal that would be extremely dificult to replicate, ensuring that the pattern is not merely a photographic copy of the original trigger.

We developed two dendrite datasets that are described below.

(1) Fine Grain Dendrite Dataset. The dataset consists of 10 dendrites, with approximately 25 images per dendrite class. The images were captured from 5 different perspectives in order to fully encompass the dendrite's branches and the distinctive patterns or "constellations" created from the reflections of metal flakes in the substrate. The first dataset was created with a 20% test split, with 212 images in the training set, and 50 images in the test split.
(2) Ultra Fine Grain Dendrite Dataset. This dataset consists of 97 dendrites, without constellations. Like the first dataset, the images were captured from five different perspectives. However, there are only 5 images per class, making this dataset extremely fine grain. The second dataset was also created with a 20% test split, with 382 images in the training set, and 97 images in the test split. Both datasets were resized to 224 × 224 pixels.

Images in both the fingerprint and dendrite datasets were resized to be 224 × 224 pixels, as ResNet was trained on images of this size, and has been shown to perform better on these dimensions. A cutout augmentation technique was applied, randomly masking rectangular regions of the image, promoting better generalization by discouraging overfitting. Additionally, we applied a randomized sequence of image transformations, to enhance dataset diversity and encourage the model to learn more robust features. Finally, the images were converted to tensors for input into our model.

3 BACKGROUND AND RELATED WORK

Fine-grained identification is concerned with recognizing and distinguishing subordinate categories, such as dog species, monkey faces, or models of car. Generally, problems become increasingly "fine-grained" when there is less data and similar structure among subcategories [4]. There are several major challenges that are associated with fine grain identification problems [16], usually a lack of training data, and low class variance. Many fine-grained problems utilize a transfer learning paradigm, wherein a model is first trained on a large and diverse source dataset, in a process known as pre-training, then fine-tuned on a specific target dataset. The vast majority of transfer learning models are pre-trained using a supervised learning approach, where pre-training takes place on large labeled datasets such as ImageNet [3]. A frequently used version of ImageNet is ImageNet-1K [12], which contains 1281167 training images, 50000 validation images and 100000 test images, and is frequently the benchmark of most computer vision models. ImageNet-1K is sourced from the original ImageNet dataset, which contains more than 14 million images across more than 20000 categories, and is commonly referred to as ImageNet-21K. Despite being overlooked due to large size, complexity, and a perceived size-value disparity, ImageNet-21K has demonstrated enhanced model performance on various benchmarking datasets [11]. However, the effectiveness of this approach in a fine-grained classification setting remains relatively unexplored.

However, there are possible limitations to pre-training on datasets such as ImageNet for fine-grain identification problems. Often, much of the image features in these large datasets are not relevant to the target datasets [2]. Choosing data for model pre-training that is more aligned with the target dataset has shown to improve performance, as the features extracted from the source dataset are more likely to match that of the target dataset. We explore the impact of specific and general pre-training data on dendrite classification by utilizing transfer learning with fingerprint data, as well as with ImageNet-1K and ImageNet-21K weights.

3.1 Existing work on Dendrite Classification

3.1.1 Non Deep Learning Approaches. Current research on dendrite classification is ongoing, and published results are very limited. Most non-deep-learning approaches tend to utilize graph matching algorithms, where images of dendrites are encoded as graph structures, which are then utilized by graph-matching algorithms. Early approaches like [15] outline a conceptual graph-based matching algorithm approach to identify the closest dendrites from a sample after reduction techniques have been performed on the dataset. However, the paper lacks information regarding the eficacy of the approach. Similarly, [1] also developed a graph matching approach and represented test as well as reference images are graphs before matching. While this approach achieved a high identification rate, it is important to note that this matching rate refers to comparisons between dendrite image representations that had minimal noise or interference. The paper also explored some image modification results, such as rotation and scratch, but does not consider the alteration in shape that occurs when observing branching structures from various perspectives.

It is crucial to note that existing works are specifically focused on the matching of identical dendrite image (or image extracted data) pairs. Our paper is highly unique as it focuses on the challenging task of classifying unique dendrite images captured at various angles, which acknowledges the change in form when branching structures are viewed at different angles, an area where graph methods face limitations. Our training and test data consists entirely of unique images, as will be expanded on later in the paper.

3.1.2 Deep Learning Approaches. Machine learning approaches offer a scalable solution for learning image features, presenting significant potential for enhancing secure identification. As datasets increase in size, deep-learning approaches are more appropriate, following the same development in the areas of fingerprint biometrics, where deep-learning approaches have taken over simpler minutiae-matching approaches in the past. For example, [15] introduced a conceptual but untested method for training a neural network on 3D dendrite extractions using digital holography. Despite this, there is currently a lack of research concerning deep learning methodologies specifically focused on dendrite classification.

4 APPROACH

An overview of our approach is given in Algorithm 1. To establish baseline on the collected dendrites dataset, we develop pre-trained models based on ImageNet and Fingerprint datasets which are then finetuned on our datasets.

pre-training	Finetuning/test-data	Avg. accuracy
Fingerprint	Fine-grain dendrites	95.73%
ImageNet-1K	Fine-grain dendrites	100%
ImageNet-21K	Fine-grain dendrites	100%
Fingerprint	Ultra Fine-grain dendrites	n.a.
ImageNet-1K	Ultra Fine-grain dendrites	93.39%
ImageNet-21K	Ultra Fine-grain dendrites	96.13%

Table 1: Performance of ResNet50 architecture pre-trained with different choices, fine-tuned on our dendrites dataset, and final test performance on a held-out test-set from our dendrites dataset. Accuracies are reported as an average over 5 runs. The fingerprint pre-training approach did not converge when finetuned on the ultra fine-grain dendrites data.

Model Pre-Training. We pretrained our models using two datasets. ImageNet-1K and ImageNet-21K are large image datasets used for computer vision, featuring 1000 and 21000 classes, respectively. Carefully curated through crowd-sourcing and expert curation, they offer diverse and high-quality images. The datasets were constructed with standardized 224 × 224 resolution images, facilitating compatibility with popular deep learning architectures like ResNet. ImageNet-21K contains 14197122 images across 21841 classes. ImageNet-1K, which was created from ImageNet-21K, contains 1281167 training images, 50000 validation images and 100000 test images, across 1000 classes.

Additionally we also pretrained the model with The NIST SD 302b Fingerprint dataset [5]. This dataset was collected with a rolling technique which aimed to capture the entire fingerprint surface from nail to nail as well as a segmented 4-4-2 slap configuration, wherein the target fingerprints were simultaneously captured with imaging. Both methods of collection are visually distinctive. For this reason, the source dataset is split between these two capture methods, with 8002 images in the roll subset and 2921 images in the segmented slap subset. The dataset was created with a 20% test split, with 10920 images in the training set, and 2184 images in the test split.

Model details. We use a ResNet50 architecture [6] trained using the Adam optimizer [8] with a learning rate of 3e-4. The loss function employed is the standard cross-entropy loss with label smoothing [14]. The data is loaded in batches of 64 with 8 worker threads for data loading. Mixed precision training is utilized for improved performance and reduced memory usage, through PyTorch's Automatic Mixed Precision (AMP) feature and the GradScaler.

5 EXPERIMENTAL SETUP AND RESULTS

ResNet50 pre-trained with NIST Fingerprint SD 302b finetuned with Fine Grain Dendrite Dataset. The ResNet50 architecture trained from scratch on the fingerprint source dataset for 20 epochs, with a batch size of 64. The model reached a test accuracy of 100% at the eleventh epoch, and this performance was maintained through the rest of training. The model achieved a perfect ROC curve, zero

Algorithm 1 Evaluate Pre-training Datasets for Dendrite Classification

```
Initialize an empty list to store dataset evaluation scores
for each pre-training dataset □ in the list of candidate datasets do
    // Step 1: Pre-train a model on dataset □
    Initialize a new model □ with architecture suitable for dendrite classification (e.g., ResNet)
    Pre-train model □ on dataset □
    // Step 2: Fine-tune and evaluate the model on the dendrite dataset
    Fine-tune model □ on the dendrite target dataset
    Evaluate the fine-tuned model on a held-out validation set of the dendrite dataset
    Compute the evaluation metric (e.g., accuracy, F1 score)
    // Step 3: Store the evaluation score
    Add the evaluation metric score to the list of dataset evaluation scores
    // Step 4: Clean up for the next iteration
    Reset model □ for the next dataset
end for
// Step 5: Analyze the evaluation scores
Identify the dataset □max with the highest evaluation score
Select □max as the optimal pre-training dataset

return □max along with its evaluation score
```

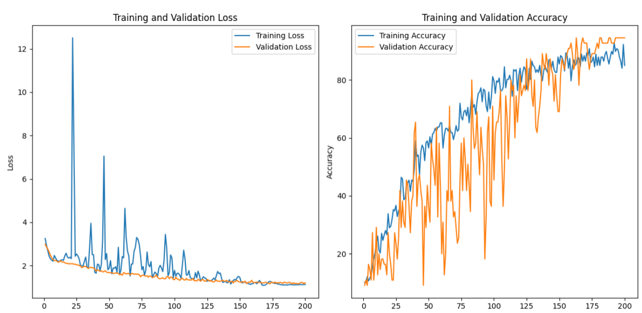

Figure 3: ResNet50 pre-trained on Fingerprint dataset, model training history on Fine-grain dendrite dataset.

false positives, and a threshold of 1, indicating extremely high model performance on the source dataset.

This pre-trained ResNet50 model was fine-tuned on the fine grain dendrite dataset for 200 epochs. The model reached a highest test accuracy of 96.36% at epoch 174 and one vs all precision, recall, and f1 scores (averaged across classes) of 0.9731, 0.8233, 0.8754, respectively, at the end of training. The model achieved a nearly perfect one-vs-all ROC curve. The average test accuracy reached a test accuracy of 95.73% across five individual training sessions. The overall model training and convergence history is displayed in figure 3.

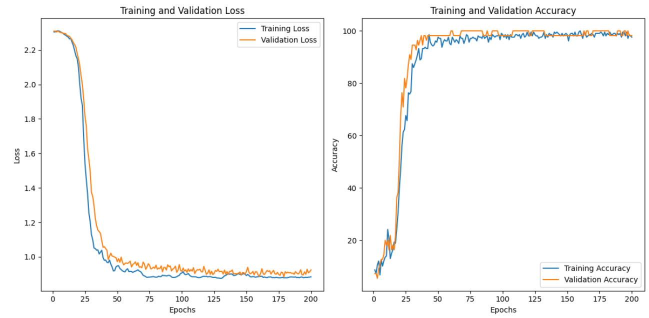

Figure 4: ResNet50 pre-trained on ImageNet-1K, model training history on Fine-Grain dendrite dataset.

Figure 5: ResNet50 pre-trained on ImageNet-21K, model training history on Fine-Grain dendrite dataset.

ResNet50 pre-trained with ImageNet-1K finetuned with Fine Grain Dendrite Dataset. The ResNet50 model loaded with ImageNet-1K weights was trained on the fine grain dendrite dataset for 200 epochs with a batch size of 64. The model reached a test accuracy of 100% at epoch 108, and perfect one-vs-all precision, recall, and F1 scores at the 111th epoch. However, at the final epoch, the test accuracy declined to 98.18%, while the recall and F1 score for one of the dendrite classes had declined to 0.8 and 0.8889 respectively. The model achieved a perfect one-vs-all ROC curve. The overall performance history is displayed in figure 4. The average test accuracy reached 100% across five training sessions.

ResNet50 pre-trained with ImageNet-21K finetuned on Fine Grain Dendrite Dataset. The Resnet50 model loaded with ImageNet-21K weights was trained on the fine grain dendrite dataset for 200 epochs with a batch size of 64. The model reached a test accuracy of 100% at epoch 30, and perfect one vs all percussion, recall, and f1 scores at the 35th epoch. The model achieved a perfect one vs all ROC curve. The overall training convergence history is displayed in figure 4. The average test accuracy reached 100% across five training sessions.

ResNet50 pre-trained with NIST Fingerprint SD 302b finetuned with Ultra Fine Grain Dendrite Dataset. We were not able to successfully fine tune the model on the Ultra Fine Grain dataset. We also found, as we describe above, that the pre-training with

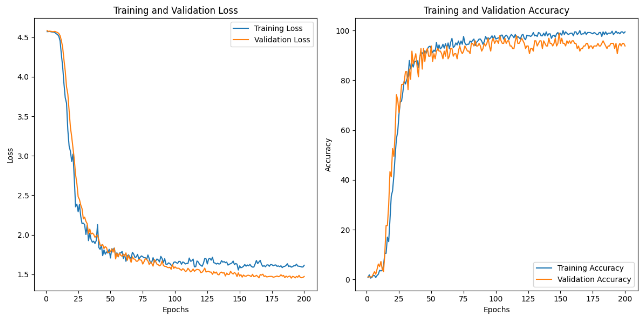

Figure 6: ResNet50 pre-trained on ImageNet-1K, model training history on Ultra Fine-Grain dendrite dataset.

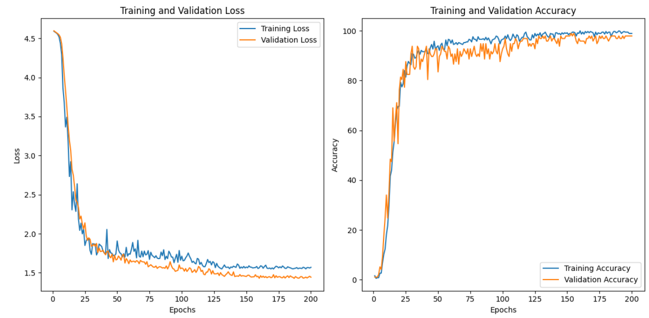

Figure 7: ResNet50 pre-trained on ImageNet-21K, model training history on Ultra Fine-Grain dendrite dataset.

ImageNet already leads to better performance on the Fine Grain dataset. Thus, we limit our comparisons below to ImageNet pre-training only.

ResNet50 pre-trained with ImageNet-1K finetuned with Ultra Fine Grain Dendrite Dataset. The ResNet50 model loaded with ImageNet-1K weights was trained on the ultra fine grain dendrite dataset for 200 epochs with a batch size of 8. The best performing model reached a test accuracy of 94.85% at the end of training. The model achieved a nearly perfect one-vs-all ROC curve. The overall training history is displayed in figure 6. The average test accuracy reached 93.39% test accuracy across five individual training sessions.

ResNet50 pre-trained with ImageNet-21K finetuned on Ultra Fine Grain Dendrite Dataset. The ResNet50 model loaded with ImageNet-21K weights was trained on the ultra fine grain dendrite dataset for 200 epochs with a batch size of 8. The best performing model reached a test accuracy of 97.94% at the end of training. The model achieved a nearly perfect one-vs-all ROC curve. The overall performance history is displayed in figure 7. The average test accuracy reached 96.13% test accuracy across four individual training sessions.

Summary of findings. All of our results are summarized in table 1. Our findings underscore that, for dendrite classification using the ResNet50 architecture, pre-training on ImageNet-21K stands out as the most effective approach by a considerable margin. This holds true when compared to ResNet50 models pre-trained on ImageNet-1K and even fingerprint data. While our exploration of the fingerprint pre-training approach did not yield the expected outcomes, it nonetheless provided valuable insights. The unsatisfactory performance can be attributed to several potential factors. Our fingerprint source dataset encompassed only two classes, an insuficient setup for effective fine-tuning on a multi-class dataset. While not mentioned in this paper, preliminary experiments involving a multiclass fingerprint source dataset displayed even poorer dendrite classification performance, strongly suggesting that relying solely on pre-training with a single object/type (in this case, fingerprints) is not an effective approach. An intriguing extension of this study could involve pre-training on a diverse multiclass dataset featuring objects akin to dendrites, such as sea shells, coils, fingerprints, and other similarly structured entities. In conclusion, our findings highlight the substantial promise that computer vision methods hold for dendrite identification, provided the right pre-training data is utilized.

6 CONCLUSION

Dendrite classification poses complexities due to their intricate branching patterns. While computer vision has the potential to aid this task, challenges arise from the intricate nature of dendrites and the scarcity of data. In this paper, our focus lies in investigating the influence of different pre-training datasets on dendrite classification. We evaluate three models, each trained on distinct data sources: ImageNet-21K, ImageNet-1K, and NIST 302b fingerprint data. These models are then applied to our Dendrite dataset, with their layers remaining fully adaptable. We aim to understand how the choice of pre-training data impacts the classification process.

Intriguingly, as the target data becomes more 'particular', the effectiveness of highly 'general' pre-training becomes more pronounced in dendrite classification. This finding runs counter to previous studies that have suggested the eficacy of pre-training on source data closely resembling the target dataset. Ultimately, our research reveals that computer vision, particularly when rooted in hyper-general pre-training datasets like ImageNet-21K, holds significant promise for dendrite classification. Recent results in the broader computer vision literature have also explored pre-training with synthetically generated images, including fractals, which can provide competitive performance on natural image classification tasks [7]. We would also explore these ideas in future work

ACKNOWLEDGMENTS

Work supported in part by USDA NIFA grant 2020-67017-33078 and NSF grant 2323086, as well as the ASU STC-MADE program and Densec ID LLC.

REFERENCES

[1] Zaoyi Chi, Ali Valehi, Han Peng, Michael Kozicki, and Abolfazl Razi. 2020. Consistency Penalized Graph Matching for Image-Based Identification of Dendritic Patterns. IEEE Access 8 (2020), 118623–118637. https://doi.org/10.1109/ACCESS.2020.3005184
[2] Yin Cui, Yang Song, Chen Sun, Andrew Howard, and Serge J. Belongie. 2018. Large Scale Fine-Grained Categorization and Domain-Specific Transfer Learning. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Computer Vision Foundation / IEEE Computer Society, 4109–4118.

[3] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In IEEE Conference on Computer Vision and Pattern Recognition. 248–255.

[4] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In International Conference on Learning Representations, (ICLR) 2021. OpenReview.net.

[5] Gregory Fiumara, Patricia Flanagan, Kenneth Ko, Karen Marshall, El-ham Tabassi, and Bryan Woodgate. 2019. NIST Special Database 302. https://www.nist.gov/itl/iad/image-group/nist-special-database-302. https://doi.org/10.6028/NIST.TN.2007

[6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 770–778. https://doi.org/10.1109/CVPR.2016.90

[7] Hirokatsu Kataoka, Kazushige Okayasu, Asato Matsumoto, Eisuke Yamagata, Ryosuke Yamada, Nakamasa Inoue, Akio Nakamura, and Yutaka Satoh. 2022. Pre-Training Without Natural Images. International Journal of Computer Vision (IJCV) 130, 4 (Apr 2022), 990–1007.

[8] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In International Conference on Learning Representations, (ICLR), Yoshua Bengio and Yann LeCun (Eds.).

[9] Michael N. Kozicki. 2021. Information in electrodeposited dendrites. Advances in Physics: X 6, 1 (2021), 1.

[10] Michael N. Kozicki. 2023. Fabrication of dendritic structures and tags. Patent US-11598015-B2.

[11] Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik. 2021. ImageNet-21K Pretraining for the Masses. In Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks, Vol. 1. Curran.

[12] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. International Journal on Computer Vision 115, 3 (Dec 2015), 211–252.

[13] P. G. Saffman and F. R. S. Sir Geofrrey Taylor. 1988. The penetration of a fluid into a porous medium or Hele-Shaw cell containing a more viscous liquid. In Dynamics of Curved Fronts, Pierre Pelcé (Ed.). Academic Press, San Diego, 155–174.

[14] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. 2016. Rethinking the Inception Architecture for Computer Vision. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 2818–2826.

[15] H. Wang, X. Chen, A. Razi, M. Kozicki, R. Amin, and M. Manfredo. 2022. Nano-Resolution Visual Identifiers Enable Secure Monitoring in Next-Generation Cyber-Physical Systems. In International Conference on Computational Science and Computational Intelligence (CSCI). 856–861.

[16] Ze Yang, Tiange Luo, Dong Wang, Zhiqiang Hu, Jun Gao, and Liwei Wang. 2018. Learning to Navigate for Fine-Grained Classification. In 15th European Conference on Computer Vision (Munich, Germany). Springer-Verlag, Berlin, Heidelberg, 438–454.

Received 31 March 2024; revised 11 April 2024; accepted 25 April 2024