

CNN-based Grain Size Determination of Microscopy Images

Jingchao Zhang

Stanford Center for Professional Development
Stanford University

jingchao@stanford.edu

Abstract

Traditional grain size determination in microscopy images is a laborious process involving lots of manual input and human expertise. In recent years, the development of computer vision (CV) has provided an alternative approach to this problem, and preliminary work has been done to simplify the grain size determination process. This work aims to build an end-to-end workflow to measure the grain size in microscopy images without any manual input. Results from the traditional line intercept (Heyn method) and planimetric (Saltykov method) approaches are used as baseline. A pre-trained holistically-nested edge detection (HED) model is used for CV-based edge detection, and the results are further compared to the classic Canny edge detection method. Post-processing was performed using open-source image processing packages to extract the grain size. It was reported that for grain size determinations in optical and scanning electron microscope (SEM) images, the pre-trained HED model achieves much higher accuracy compared to the Canny edge detection method, meanwhile reducing the image processing time from 15 minutes per image to under 10 seconds compared to the traditional methods. Effects of morphological operations on the predicted grain size accuracy are also explored. Overall, the proposed end-to-end CNN-based workflow can significantly reduce the processing time while maintaining the same accuracy as the traditional method.

1. Introduction

Grain size plays a vital role in the microstructure analysis of metallurgical and ceramic materials.[1] It is well known that microscopic morphology, such as grain size, is greatly responsible for macroscopic mechanical characteristics.[2] Grain sizes in ceramics and metals are historically determined via the American Society for Testing and Materials (ASTM)[3] and International Organization for Standardization (ISO)[4] standards. Knowledge of rules and exceptions is needed. Still, the primary flaw in conventional grain size

determination is the time required to generate accurate results.

Artificial intelligence (AI) has quickly risen to the center of image processing due to its adaptability in recognizing objects in diverse image sets.[5] While being used in some research, AI remains vastly untapped by the materials science community compared to its uses in medical image segmentation.[6] ResNet is an AI-powered image segmentation method designed for use on the diversified image dataset ImageNet.[7] The same pretrained model was used in hand with U-Net models by Liu et al.[8] for the segmentation of breast and liver lesions with success with diverse anisotropic resolution. In recent years, research in materials has started to implement computer vision (CV) and image processing to produce more precise grain determinations in less time. Pretrained models are becoming increasingly important by requesting less computational power, time, and data than conventional neural networks.[9] The implementation of a pretrained model instead of outright training a new model allows for fast, reduced computation without overfitting to specific material types. Given the limited dataset in this work, we will implement pre-trained models for edge detection and grain-size extraction.

2. Related Work

Grain size determination is traditionally a very laborious work. There are two traditional grain size determination methods that are used by the American Society for Testing and Materials organization, i.e., line intercept (Heyn method)[10] and planimetric (Saltykov method)[11]. Both methods require lots of manual input and human expertise, and are very time expensive. These methods are considered as golden standards in microscopy image grain size determinations. The CNN predicted grain sizes are compared to these two methods to evaluate the model performance.

As shown in Fig. 1, grain size determination is primarily an issue of edge detection. Image processing techniques like Canny edge detection facilitate grain size determination by converting original images to masked outlines of grain borders.[12] Polarized micrographs were processed




Figure 1. Edge detection is vitally important for accurate grain size determination.

via Lazy Grain Boundary produced grain boundary maps of steel microstructures.[13] Peregrina-Barreto et al[14] used an advanced thresholding technique to discern grains and a planimetric grain count to determine the ASTM grain size. Past research has also employed Canny edge detection to create closed contours and similarly determine planimetric grain size, with small improvements in accuracy.[15] Zhang’s thinning algorithm has been shown to reduce processed micrographs to its main edge components.[16]

CV has been implemented in grain size analysis, but relies on the inclusion of large datasets seldomly available to the typical laboratory[17]. The efficacy of a neural network is dependent on the amount of data that it is given, even in transfer learning.[18] Pretrained models can directly apply complex computer vision algorithms without the input of any training data. U-Net, for instance, is a deep learning model widely used in medical image segmentation.[8] Edge detection is a fundamental and important topic in computer vision. Various approaches have been proposed to tackle this problem, such as Canny detector[12], Statistical Edges[19], Structured Edges[20] and DeepEdge[21]. Pre-trained edge detection algorithms with high accuracy already exist and may have impacts outside of their original field. These methods, however, have some key drawbacks. Manual input is often still needed to connect grain boundaries. Current methods work especially well with specific materials but have not been applied to diverse sets of micrographs. More specifically, a large focus is placed on metallurgical optical images, with little consideration for ceramic or scanning electron microscope (SEM) micrographs.

In this paper, a novel method of average grain diameter determination using the pretrained holistically-nested edge detection algorithm is explored.[22] This method aims to use the sophistication of a model trained on non-micrograph data to process and characterize the grain size of various materials in an automated fashion.




Figure 2. Illustration of the traditional line intercept method.

3. Methods

The traditional line intercept method is used as the baseline to evaluate the performance of CNN-based models. In the line intercept method, a set of 10 randomly drawn straight lines are drawn over the microscopy images, cumulatively spanning the entirety of the image height and width. The intercepts, including all exceptions, are noted along with line length and converted to μm using the pixel to μm ratio derived from the scale bar within the image. Intercepts are determined manually with visual detection, while the ImageJ software is used to determine individual line lengths.[23] Total line length is divided by total intercepts to give an average grain diameter length. All samples exhibited relatively equiaxed structures, meaning no further specially oriented test lines are needed. An illustration of the line intercept method is shown in Fig. 2.

In the planimetric method, the entire rectangular area of the image is selected. Each grain within the image is counted as N_{inside} , with edge touching grains counting as 0.5 a grain $N_{intercepted}$, and corner grains as 0.25 a grain. The following equation determines the number of grains per μm , N_A , using magnification, M , and total image area, A :

$$N_A = (M^2/\bar{A})(N_{inside} + 1/2N_{intercepted} + 1) \quad (1)$$

The average grain area, $\bar{A}$, is the reciprocal of $\bar{A}$. Each area is approximated as a circle with diameter, d , as the square root of $\bar{A}$. Using a rectangular area allows for a higher grain count and minimized scatter, which increases accuracy compared to a circular area. This is especially true in the micrographs with larger grains, such as the ones encompassed in the test dataset. The planimetric method also alleviates any impact of un-equiaxed grain orientations, which is minimal in the test dataset.

Although proper sample preparation and imaging practices may be used, the resulting images are often not opti-




Figure 3. traditional workflows vs. CNN based workflow

mal for direct image processing. As shown in Fig 1., large scale edges must be distinguished against smaller material imperfections to discern grain boundaries from crystalline material. Typical microscopy images have noise: fluctuations in intensity as a byproduct of exposure and deformities on the sample surface. Gaussian smoothing is commonly employed to address this issue but fails to preserve edges. Local Laplacian filters (LLF) alleviate this by smoothing noise and keeping large edges intact. Initially formulated by Paris et al[24], this image processing technique computes the Laplacian pyramid in neighborhoods, locally enhancing details of one area while neglecting the rest of the image. Small Gaussian convolutions and simple point-wise nonlinearities reduce the complexity of the algorithm, allowing for many images to be processed with less computational power. Simple pixel value thresholding discerns large edges over noise in each neighborhood.[25] Neighborhoods are then stitched together, forming partial pyramids and eventually a fully processed image. LLF does not introduce halos, an otherwise common occurrence in traditional Laplacian pyramids and frequency boosting. Three parameters impact the performance of an LLF: sigma, σ , defines minimum detail size, alpha, α , determines the degree of smoothing, and beta, β , controls dynamic range compression. Images in the test dataset were processed in MATLAB using LLF with parameters $\sigma = 0.4$ and $\alpha = 0.5$. No dynamic range compression is needed thus β is neglected, with an equivalent value of 1. Comparison of the traditional workflows vs. the CNN-based workflow is shown in Fig. 3.

The pretrained HED model is used to extract the grain boundaries in composite microscopy images. This model is able to learn the hierarchical features which is crucial in detecting natural image edges and object boundaries. The model is inspired by fully convolutional neural networks with trimmed VGG nets for deep supervised learning. The model is pretrained on Berkeley Segmentation Dataset and Benchmark (BSDS 500), which has 200 training, 100 validation and 200 testing images. The neural network architecture of the HED model is shown in Fig. 4. Side output layers is inserted after the convolutional layers. Deep supervision is applied to each side output layer to direct the side output to edge prediction with the features we want. The output of HED is multiscale and multilevel, with the side output planes becoming smaller in size and the perceptual field size becoming larger. A weighted fusion layer is added to automatically learn how to combine the outputs from multiple scales. The entire network is trained using multiple error propagation paths (dashed lines).[22]

4. Dataset and Features

A total of 30 images were collected using two different approaches, i.e., optical microscope and scanning electron microscopy. Multiple material samples were used to obtain a diverse dataset. Sample materials were chosen to have various microstructures and grain sizes while possessing clear, equiaxed grain boundaries.

All optical microscopy images were of Ceria (CeO_2) pellets which were unaltered via etching or polishing be-




Figure 4. Illustration of the HED architecture.[22]




Figure 5. (a) Optical image. (b) SEM image

fore being imaged. A Nikon Eclipse LV100 optical microscope with a Lumenera Infinity 1 microscope camera attachment took 1280 by 1024 resolution TIF images, each 9.01 megabytes in size. Lumenera’s Infinity Analyze software was used to place scale bars in each bottom right corner. Polyvinylpyrrolidone, Magnesium-oxide, Ceria, and Barium-titanate samples were imaged using SEM. Gold and carbon coatings facilitated imaging for these non-conducting materials. A Joel Neoscope JCM-5000 bench-top scanning electron microscope produced 1280 by 1120 JPGs ranging from 231.42 to 563.21 kilobytes in size. Sample images from optical and SEM measurements are shown in Fig. 5.

All TIF files are converted to JPG format before preprocessing. Original metadata is conserved for further use in the later pixel-to-micrometer ratio as a separate image. Image samples are then converted to RGB format, if not already, for use in holistically-nested edge detection.

5. Experiments/Results/Discussion

5.1. Edge map prediction

The newly proposed method contains two primary parts: grain boundary amplification and grain size determination. In the former, an image sourced from an SEM or optical microscope is processed in an image editing software for surface adjustments in contrast. This ensures grain boundaries




Figure 6. (a) Original image (b) HED output binary image (c) post-processed image with edge enhancement. (d) color mapped image




Figure 7. (a) Optical and (b) SEM micrographs after HED and initial thresholding, 8 iterations of erosion, and 8 iterations of dilation.

are visible and noise is reduced. Metadata is also assigned in this step, with a pixel-to-micrometer ratio inserted under the Exif user comment.

In this work, ImageJ was used for contrast filtering and a python script for metadata assignment. The pre-trained HED model transforms the edited JPG image into a grain boundary prediction map. The original image is conserved for its metadata use in the following step.

In grain size determination, the PNG output of HED is converted to a grayscale and thresholded to distinct grain boundaries from crystalline material. The binary image undergoes eight iterations of erosion and dilation using a disk kernel of radius 1 pixel. A 3×3 structural element then distinguishes cells as the regions between grain boundaries. Cells touching the border of the image are cleared so that only full inside cells are considered. Scikit’s regionprops attribute calculates the area of each region.[26]

In polycrystalline single-phase materials, the most energy favorable grain structure is such that the intersection of three grains forms at dihedral angles of 120. In a two-dimensional plane, an idealized polycrystalline material has




Figure 8. Mean diameter determination of honeycomb structure (a) with an expected value of 1.866 units. Alterations on the initial figure include (b) sidewall thickening, (c) low-intensity low contrast, and (d) high-intensity low contrast.




Figure 9. Grain size measured by HED and Canny edge detection algorithms compared to line intercept method and reported as % error.

a normal hexagon honeycomb structure. The area of each region (A_r) is thus taken to be that of a hexagon with an inner diameter (d) resulting from the smallest possible inscribed circle and an outer diameter (D) from the largest possible inscribed circle. The mean of the two diameters is used as the average diameter of the grain.

The pixel-to-micrometer ratio is retrieved from the metadata to convert pixel units to micrometers. Grain size is reported as a distribution or as a mean of the grains detected. No input other than an image and pixel to micrometer ratio is needed. The goal of this algorithm is not necessar-

ily to follow ASTM or ISO standards, but to curate results which are accurate to the standards which it is based on, quickly, and without great input from the researcher. The grain boundary detection process from the original image to the color mapped image is shown in Fig. 6.

5.2. Morphological operations

Morphological operands remove noise while largening or shrinking the background of an image.[27] Erosion preserves foreground pixels only when a structural element entirely matches that of the pixels it covers. The affect is the thinning of foreground images. Dilation works similarly by comparing the structural element to that of the background, resulting in a thickening of foreground images. Two inputs are needed: dimensions of the kernel, and a binary image. The process may be repeated, in iterations, to achieve desired effects.

In grain size determination, morphological operations alleviate some issues caused by an edge detection methodology of grain size determination. Micrographs do not typically possess perfectly distinct grain boundaries. The HED thus interprets grain boundaries with defects such as broken edges, double edges on thick boundaries, and undesired particles in locations of noise. In morphological erosion, broken edges are united and double edges are joined together by a disk (radius of 1 pixel) structural element, given the distance between interpreted edges is small enough. The resulting image contains large grain boundaries and miscellaneous bits, which is misrepresentative of the material




Figure 10. A comparison of four grain determination methods' error to the average grain Haywood diameter. For reference, the micrographs are located under, with areas of less resolution circled in red.

sample. Morphological dilation using the same structural element reduces edge thickness and size of unwanted particles to yield an image much closer to the material sample. The effects on optical and SEM images are depicted in Figure 7.a and Figure 7.b.

5.3. Model validation with honeycomb testing

To analyze the effectiveness of the newly proposed grain size determination algorithm, a series of tests in increasing complexity were carried out.

To begin, a representative test on a computer-generated honeycomb structure was executed. Each image was processed through HED, thresholding, morphological operations, and grain size determination as discussed above. By using a standard hexagon, the diameter is definitive, with no interpretation or error needed from a standardized method like Heyn's or Saltykov's. The side length of each hexagon is normalized to 1 arbitrary unit, defining the minimum and maximum diameters as 1.732 and 2.000 units, respectively. The results of the algorithm are shown in Figure 8 with images at each processing step. An optimal construct (Figure 8.a) and an accurate pixel-to-micrometer ratio result in accurate diameter determination. Thickening the hexagon borders (Figure 8.b) leads to a decrease in grain area and a smaller resulting diameter. Underexposed and overexposed micrographs are mimicked in Figure 8.c and Figure 8.d. The HED allows for edge detection at a range of luminosities, although high intensity, low contrast images perform marginally better than equivalent images at low intensity. The largest deviation from the expected value (Figure 8.b) produced a 2.84% error.

The honeycomb test highlights some benefits of the

CNN-based HED method. Firstly, the omission of edge grains allows for a more accurate representation of the material at hand because only whole grains are considered. Secondly, low errors are experienced in a wide range of image types. Thirdly, the use of Exif metadata to convey a pixel-to-micrometer ratio allows for easy batch operations of similar images. Finally, area approximation is shown to accurately represent the mean diameter of a material's crystalline regions.

The proposed algorithm must, however, concur accurate results on real micrographs to be considered a viable grain size determination method. A set of 25 optical microscopy images were analyzed using conventional and new methods. To measure the consistency of the algorithm, all optical micrographs exhibited similar grain structures and image dimensions. The newly proposed algorithm, though, is not limited to these constraints.

Heyn's method set the benchmark by using the following set of basic parameters. All lines were randomly selected before being placed on the image as a mask, but care was taken to ensure all regions of the micrograph were represented. The test lines were also oriented so that a wide range of angles was considered. A set of ten lines was found to be sufficient with an average of 127.47 intersections, greater than the 50 minimum stated by ASTM E112-13 13.1. The average grain diameter of 61.62 μm and a standard deviation of 5.149 μm was reported. The average grain count in each image was about 167 grains.

5.4. Model performance comparison

To determine the robustness of HED in comparison with other edge detection types, the same images are processed

with Canny edge detection. Thresholding and the number of eroding iterations were changed so that contained cells were formed in the Canny method, but the rest of the methodology remained the same. Canny edge detection requires two values of edge thresholding: 0 and 70 were determined based on visual inspection. The HED method was untouched. The Canny method exhibited an average error of 11.867% among optical images and 2.342% among the hexagon tests as shown in Fig. 9. In comparison, the average accuracy of 5.128% was reported for the CNN-based HED method on the optical images.

Only some past research account for error in grain size determination.[28] In practice, the majority of uncertainty is caused by the variation in measurements from one field of view to another. A 95% confidence interval (CI) and percent relative accuracy (%*RA*) can help determine the accuracy of the measurement. The standard ASTM E122-13 15.6 states a %*RA* less than 10% is generally acceptable, which is followed in this paper. Fields (images) for each sample are collected until the %*RA* reached an acceptable value: this equated to 5 samples for each ceria pellet. It is important to note, however, that high local precision may only yield a slight increase in overall precision (ASTM E112-13 15.1.1).

All grain size determination methods inevitably possess errors. A comparison to a more accurate benchmark than Heyn's may thus be made. To obtain the ground truth average diameter, the Heywood diameter of each cell is taken. A closed contour is manually drawn using ImageJ software around each interior grain. The area is calculated and equated to the area of a circle with a constant diameter, *d*. The average of all Heywood diameters is taken to represent the true average diameter of the micrograph. A set of 10 micrographs from the original 25 were selected and measured using Heyn's, Saltykov's, Canny edge detection, and HED methods as shown in Fig. 10.

The HED method had an error percentage lower than Heyn's and Canny edge detection for 8 of the 10 micrographs. The two micrographs where the opposite is true are the two images in which a portion of the image is not fully resolved. Saltykov's method reliably had the most accurate grain size, except for two images (1 and 4) where HED outperformed. It is important to note the accuracy of Heyn's method is dependent on the number of intersected grain boundaries (ASTM E112-13 13.1), meaning the error of Heyn's is not set at the values obtained. An increase in the number of test lines and intersections undoubtedly would increase Heyn's accuracy.

On the other hand, the proposed CNN-based HED method reduced the grain boundary detection time from 15 minutes to under 10 seconds. The CNN-based method eliminate the needs of manually selecting the grain areas, deciding pixel dimensions, and connecting missing edges while

maintaining a high accuracy comparable to the golden standards.

6. Conclusion/Future Work

In this work, a CNN-based grain boundary detection approach was proposed for microscopy images in the material science domain. Traditional methods such as line intercepting and planimetric methods are used as baseline to evaluate the performance of the pre-trained HED model. The HED model is trained on the Berkeley Segmentation Dataset and Benchmark (BSDS 500) dataset. The classic Canny edge detection algorithm is also explored for comparison. The HED model outperforms the Canny approach in the hexagonal test and microscopy image test. The Canny approaches worked well on some special cases but cannot provide consistent performances over different microscopy image under varying lighting conditions and background noises. On the other hand, the HED model achieves a grain size accuracy comparable to the golden standard provided by the line intercepting method.

It is worth noting that morphological operations play an important role in refining the microscopy images to help achieve better grain size accuracy. Multiple iterations of dilation and erosion are performed to reduce the noises in the detected edge map and enhance the edge contour. We report that morphological operations needs to be applied after the HED prediction to achieve the best accuracy result, while the operations need to be applied to the original images to achieve the optimal results for Canny method. The morphological operations are integrated as part of the workflow thus no manual inputs is needed for different set of images. As a result, the proposed CNN-based method could reduce the image processing time from 15 minutes to under 10 seconds per image.

Given the constraint of time and computational resources, the primary focus of this work is on optical images, which has achieved high accuracy and time-efficiency. While the HED model works on some cases, there are still corner cases of scanning electron microscope images that the HED model cannot predict. The SEM images are radically different from the optical images with different scale and background color. The grain boundaries in the SEM images also have a "glowing effect" due to the measurement methods, which created difficulties for the current edge detection model. In the future, we will explore the possibilities of adding pre-processing steps to remove the glowing edges, or using different pre-trained models for the SEM images.

References

- [1] Robert L Folk. A review of grain-size parameters. *Sedimentology*, 6(2):73–93, 1966. 1

- [2] Henrik Jensen. Determination of macroscopic electro-mechanical characteristics of 1-3 piezoceramic/polymer composites by a concentric tube model. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 38(6):591–594, 1991. 1
- [3] HT Hahn and RY Kim. Proof testing of composite materials. *Journal of Composite Materials*, 9(3):297–311, 1975. 1
- [4] Marcel Heires. The international organization for standardization (iso). *New Political Economy*, 13(3):357–367, 2008. 1
- [5] Stephanie Robertson, Hossein Azizpour, Kevin Smith, and Johan Hartman. Digital image analysis in breast pathology—from image processing techniques to artificial intelligence. *Translational Research*, 194:19–35, 2018. 1
- [6] Joseph R England and Phillip M Cheng. Artificial intelligence for medical image analysis: a guide for authors and reviewers. *American journal of roentgenology*, 212(3):513–519, 2019. 1
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1
- [8] Zhenqing Liu, Yiwen Cao, Yize Wang, and Wei Wang. Computer vision-based concrete crack detection using u-net fully convolutional networks. *Automation in Construction*, 104:129–139, 2019. 1, 2
- [9] Xipeng Qiu, Tianxiang Sun, Yige Xu, Yunfan Shao, Ning Dai, and Xuanjing Huang. Pre-trained models for natural language processing: A survey. *Science China Technological Sciences*, 63(10):1872–1897, 2020. 1
- [10] George F Vander Voort. *Grain size measurement*. ASTM International, 1984. 1
- [11] GF Vander Voort. Grain size measurements using circular or rectangular test grids. *Practical Metallography*, 50(1):17–31, 2013. 1
- [12] Lijun Ding and Ardeshtir Goshtasby. On the canny edge detector. *Pattern recognition*, 34(3):721–725, 2001. 1, 2
- [13] Anirban Naskar, Madhumanti Bhattacharyya, Krishnan S Raja, Indrajit Charit, Jens Darsell, and Saumyadeep Jana. Pitting behavior of friction stir repair-welded 304l stainless steel in 3.5% nacl solution at room temperature: role of grain and defect structures. *SN Applied Sciences*, 2(12):1–21, 2020. 2
- [14] H Peregrina-Barreto, IR Terol-Villalobos, JJ Rangel-Magdaleno, AM Herrera-Navarro, LA Morales-Hernández, and F Manríquez-Guerrero. Automatic grain size determination in microstructures using image processing. *Measurement*, 46(1):249–258, 2013. 2
- [15] Lixin Zhang, Zhengguang Xu, Shuailing Wei, Xuechong Ren, and Meiling Wang. Grain size automatic determination for 7050 al alloy based on a fuzzy logic method. *Rare Metal Materials and Engineering*, 45(3):548–554, 2016. 2
- [16] Tongjie Y Zhang and Ching Y. Suen. A fast parallel algorithm for thinning digital patterns. *Communications of the ACM*, 27(3):236–239, 1984. 2
- [17] Akshansh Mishra and Tarushi Pathak. Estimation of grain size distribution of friction stir welded joint by using machine learning approach. *ADCAIJ: Advances in Distributed Computing and Artificial Intelligence Journal*, 10(1):99–110, 2021. 2
- [18] Karl Weiss, Taghi M Khoshgoftaar, and DingDing Wang. A survey of transfer learning. *Journal of Big data*, 3(1):1–40, 2016. 2
- [19] Roger W Minich, Christopher A Schuh, and Mukul Kumar. Role of topological constraints on the statistical properties of grain boundary networks. *Physical Review B*, 66(5):052101, 2002. 2
- [20] Kuan-Chia Chen, Wen-Wei Wu, Chien-Neng Liao, Lih-Juann Chen, and King-Ning Tu. Observation of atomic diffusion at twin-modified grain boundaries in copper. *Science*, 321(5892):1066–1069, 2008. 2
- [21] Gedas Bertasius, Jianbo Shi, and Lorenzo Torresani. Deepedge: A multi-scale bifurcated deep network for top-down contour detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4380–4389, 2015. 2
- [22] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. In *Proceedings of the IEEE international conference on computer vision*, pages 1395–1403, 2015. 2, 3, 4
- [23] Tony J Collins. Imagej for microscopy. *Biotechniques*, 43(S1):S25–S30, 2007. 2
- [24] Sylvain Paris, Samuel W Hasinoff, and Jan Kautz. Local laplacian filters: edge-aware image processing with a laplacian pyramid. *ACM Trans. Graph.*, 30(4):68, 2011. 3
- [25] Derek Bradley and Gerhard Roth. Adaptive thresholding using the integral image. *Journal of graphics tools*, 12(2):13–21, 2007. 3
- [26] Oliver Kramer. Scikit-learn. In *Machine learning for evolution strategies*, pages 45–53. Springer, 2016. 4
- [27] Meysam Siyah Mansoori, Hojjat Tajik, Gelareh Mohamadi, and Mohsen Pashna. Edge defect detection in ceramic tile based on boundary analysis using fuzzy thresholding and radon transform. In *2008 IEEE International Symposium on Signal Processing and Information Technology*, pages 58–62. IEEE, 2008. 5
- [28] JP Le Roux and EM Rojas. Sediment transport patterns determined from grain size parameters: Overview and state of the art. *Sedimentary Geology*, 202(3):473–488, 2007. 7