

Fine-Grained Visual Categorization of Fasteners in Overhaul Processes

Sajjad Taheritajani
 Department of Informatics
 Technical University of Munich
 Garching, Germany
 e-mail: sajjad.taheri@tum.de

Juan Haladjian
 Department of Informatics
 Technical University of Munich
 Garching, Germany
 e-mail: haladjia@in.tum.de

Bernd Bruegge
 Department of Informatics
 Technical University of Munich
 Garching, Germany
 e-mail: bruegge@in.tum.de

Abstract—Commercial aircraft engines must be overhauled approximately every six years, during which hundreds of different parts must be disassembled, checked, and then reassembled. This includes undoing up to thousands of fasteners, cleaning, checking, refitting, and tightening them. Prior to refitting the fasteners, they must be classified and packaged. In this paper, we describe a system for classifying fasteners automatically, by use of computer vision and machine learning. Using the proposed system, we created sample datasets and performed a fine-grained visual categorization of the fasteners. Our trained model classifies 20 bolts and 14 washers with an accuracy of 99.4%. Our work is the first step towards an automated fastener classification system in overhaul processes.

Keywords—Industrial Automation, Fine-grained Visual Categorization of Fasteners, Small Parts Categorization, FGVC, Overhaul Processes

I. INTRODUCTION

Airplane maintenance includes an overhaul process in which the aircraft is disassembled to examine all the components thoroughly. One challenge during this overhaul process is number of fasteners and small parts that must be manually classified and packaged in their respective boxes [1] [2]. The disassembled fasteners, as shown in Figure 1, are brought to special workstations in maintenance checking plants with access to digital manuals, pictures of each part, magnifiers, and measurement devices with which the technicians can classify them.

The technicians pick the similar fasteners, put them on the table and classify them with their part numbers, using the magnifiers. Bolts, screws and nuts have a part number carved on them. These parts are classified using the part number. In the event that the fastener does not have any part number, for instance washers, or that the part number is not recognizable because of scratches or erosion, the technicians must compare them with their comparison-sketch in a catalog in an 1:1 scale with regards to the fastener structure - see Figure 2 - to be able to place them in the correct package. The technicians might



Figure 1: Fasteners that were taken from part of airplane engine and must be classified during an overhaul process.

also compare the fasteners with the fastener samples on their workstations. Most of the aero-engine fasteners are made of resistant alloys, like steel or titanium, and are manufactured under specific considerations [3]. Hence, they can be expensive and the process of 1:1 scale comparison must be repeated for all fasteners without part number or with unrecognizable part number, which makes the classification time consuming and costly [2]. As an example, for classifying a bolt with unrecognizable part number, the technicians might need up to 30 seconds to measure the sizes and compare them with the comparison-sketch for the classification. Therefore, an automatic classification and sorting process would save time and reduce costs during overhaul processes.

The main contributions of this paper are: 1) we propose a system that employs polarized backlighting to create datasets for fasteners; 2) we train a model using state of the art convolutional neural networks to classify the fasteners. To the best of our knowledge, there is no reported study on the classification of fasteners in overhaul processes. Our system




Figure 2: The structure of a sample bolt and washer displaying the length, width, shaft, pitch and diameters.

can be also applicable to other small shiny objects and other industrial fields.

The rest of this paper is structured as follows: Section II describes the related work, Section III the challenges and the requirements of the system, Section IV the system, including data creation, preprocessing, and training, Section V the results and discussion, and Section VI the contributions.

II. RELATED WORK

Classification of fasteners can be defined as a Fine-grained Visual Categorization (FGVC) problem. The definition of FGVC originates from a Nilsback and Zisserman paper in 2006 [4] in which they introduced an object classification dataset for flower species. This problem was particularly challenging given the high intra-class and low inter-class variance nature among flowers, which made the older classification methods inefficient. Over the last decade, FGVC has gained popularity and is now used in different domains including flowers, birds, dogs, aircrafts, and many more [4] [5] [6] [8]. FGVC has been also used for industrial purposes. In 2013, Maji et al. introduced a large dataset of aircraft images for fine-grained visual categorization [8]. They obtained the images from aircraft spotter collections, maximizing internal diversity in order to reduce unwanted correlation between images taken by a limited number of photographers.

Aside from FGVC, work has also been carried out to detect and classify fasteners for industrial purposes, in particular in railway maintenance. Aytekin et al. analyzed the specific case of hexagonal headed fastener detection from depth images that were acquired using a high speed 3-D laser range finder [9]. They proposed a fused approach with an

appearance based and a histogram peak checking method. In 2014, Feng et al. proposed a railway inspection system, which assess the damage of multiple types of fasteners [7]. Their proposed system is insensitive to the illumination used and can model different types of fasteners using unlabeled data, and ranked the statuses of fasteners. Gibert et al. proposed another method for railway fastener detection by aligning training data, reducing intra-class variation, and bootstrapping difficult samples to improve the classification margin [10]. They used histogram of oriented gradients features and a combination of linear support vector machine (SVM) classifiers to inspect ties for missing or defective rail fastener problems.

In 2016, Laptev et al. designed a framework to incorporate expert knowledge on nuisance variations in the data during training deep neural networks [11]. This framework handled prior knowledge on nuisance variations in the data, such as rotation or scale changes. Influenced by [11], Xuan et al. designed a pearl classification machine, which automatically collected multi-view images of pearls [12]. Using their machine, the pearls could be classified with a multi-view CNN. In 2017, Usman and Rajpoot presented an algorithm to hierarchically classify the tumor into three regions: whole tumor, core tumor and enhancing tumor [13]. They extracted Intensity, intensity difference, neighborhood information and wavelet features from MRI scans with various classifiers.

These studies inspire the current work and we propose a system with backlighting and light polarization which records images of fasteners as the input. Our work is similar to recent study performed by Usman and Rajpoot, but the characteristic of the fasteners - see III - make this study significantly different from tumor classification. Moreover, we created a system to capture data for classifying the fasteners, while [13] used the ready MRI scans for classifying the tumors.

III. BACKGROUND

In FGVC the training dataset must contain the main discriminative features [14]. Therefore, there are three challenges in comparison with most other image datasets [15] [16] [17]. Firstly, unlike the cases in which the background of the image can contribute to the classification, here the background is not informative, rather it is a source of noise, and results in a computation overhead in FGVC problems. Secondly, many classes in FGVC can only be separated by subtle details, for example small length differences or difference in the thread pitch. Figure 3 is an example of such cases. Thirdly, the FGVC of fasteners faces challenges due to the nature of the fasteners themselves:

- **Perspective:** in the graphic arts, perspective is an approximate representation of an image on a flat surface as it is seen by the eye. Perspective is characterized by its correlation between object's size and its distance to the observer. The objects are smaller as their distance from the observer increases [18]. Unlike most of the objects where scaling the size, width or height does not affect their actual category, the size of the fasteners must be preserved in all the instances of the dataset. Any scaling transformation of the image of the fastener can result in misclassification. Figure 4




Figure 3: An example of similar fasteners. A is identical to C except that its non-threaded shaft is one millimeter shorter. B has a two millimeter longer non-threaded shaft and a one millimeter smaller diameter than C.




Figure 4: An example of two bolts that each one is a scale version of the other one. Scaling one's image is an example of the perspective issue and can result in a misclassification.

shows an example of two bolts which are scale versions of each other, considering their length and width. Their distance to the camera lens must be known and always scaled to a fixed size in all the data samples.

- **Reflection:** the fasteners used in airplane engines are metallic and primarily silver color. Their surface reflects the light, which might introduce extra noise and hides important details of the fastener in the image such as the threads. Hence, a specific light polarization can help to control and reduce the light reflection.

- **Shadow:** laying the fasteners on a surface with a normal source of light such as room lighting results in them being surrounded by shadow. Considering the size of the fasteners and the generated shadows, shadows could be a source of noise which makes the data augmentation process less useful.

Aside from the challenges, the shape, color, and texture have previously been studied as the most common visual image features [19] [20] [21]. However, the aero-engine fasteners used in this study have metallic silver or dark color and a metallic texture, irrespective of their constituent alloy.

As the texture and color are identical for all of the fasteners in the study, we focus on shape, threads, and head.

Shape: the shape of a fastener can identify whether the part is a screw/bolt, nut, washer or pin and how large, small, narrow or thick it is - length and width.

Threads: for threaded parts such as bolts, screws, and nuts, the threads can be used to distinguish between different similar-shaped items. Differences in the thread pitch, size, and the length of the threaded shaft of the part indicate a different fastener.

Head: there are many screws and bolts which have the same shape, length and width and also the same threads—pitch, length and location of threads on the shaft. However, they have different heads such as hex socket, hex cap, flat head, or curved head bolts and screws. The screw or bolt head indicates what kind of wrench or screw driver is needed to drive it and different head types are appropriate for different threaded screw holes.

IV. SYSTEM

Considering the challenges faced in FGVC fastener identification - perspective, reflection and shadow with normal ambient lighting - as well as the shape, threads and head features on which we want to focus, we propose a system for recording images of the small parts for both creating the datasets, training and prediction-inference.

A. Data Creation

To reduce the shadows around the parts and the reflection on their surface, we use a glass table covered with an opal foil and dimmable backlighting. Light polarization achieved using the opal foils allows us to control the reflection [24]. In addition, shining a controlled light behind the object directly at the camera lens darkens the object, which again reduces reflection [24]. A darker object image helps in focusing on the shape, threads, and form of the head. Moreover, a fixed camera is used to overcome the perspective issue. Using fixed cameras, the distance between the fastener and the camera lens is fixed - neglecting minor variation in distance resulting from the different positions of the fasteners on the glass table surface. We can also employ a non-fixed camera. However, we must compute the camera distance to the fastener, using fiducial markers, and scale the image to a 70 cm - for our Genie Nano GigE industrial camera. Figure 5 shows the sketch of the system design and its realization with a rotating glass table. Figure 6 shows an image of two bolts and one washer obtained with normal ambient lighting (6a) and with our system (6b).

As a single camera setup is used, it cannot capture images of the screw and bolt recesses or the inner threads of nuts. Therefore, nuts and similar bolts with different head recesses are excluded from data creation.

We created the data in a number of iterations with different light intensities and a differing number of images. For the first iteration, we started with a smaller dataset, trained a model, and evaluated how well the created model performed. If the performance was below our expectations - less than 95% - we




Figure 5: The sketch of the system (left) and its realization (right) with a rotating glass table, covered with an opal foil. The backlighting comes from below the table.



(a)



(b)

Figure 6: Image of two bolts and one washer obtained using normal ambient lighting (a) and using the system with reduced shadows and reflection (b).




Figure 7: The camera view is divided into five regions shown in blue and aqua. The red dots indicate the differences between the bolt head when the bolt is moved vertically along its shaft. The yellow dots indicate the differences in light reflection

added additional examples to our training data and tested the trained model again. After several iterations it was realized that we can divide the camera view into five regions. Placing the fastener in each of these regions can result in capturing a slightly different image in compare with other ones. Figure 7 shows five regions in the camera view and the differences between the sample captured bolt in each of them. Considering the vertical or horizontal position of the bolt along with its shaft, the bolt head and the reflections on the sides of its shaft are shown differently. Therefore, we must make sure that our dataset has sample images of the bolt in each of these regions.

We recorded 10 different images per region of each bolt - 50 images in total - and augmented 950 more images from these original images. Since the aero-engine washers are thin plates with a hole in the middle, without any threads and different sections such as a shaft and head, the quantity of actual washer images required to train a classifier may be less than with bolts and screws. Therefore, we only captured 10 images of washers without considering the region divisions and augmented 290 more images from them. This dataset could yield an acceptable accuracy - over 97% - see Figure 8.

B. Data Preprocessing

Figure 9 shows an overview of the preprocessing pipeline prior to training. After recording the images with our system, we converted the images to gray scale to filter out the color feature.

In addition, we created a square bounding box around the object and cropped the image to reduce both the background and the computation overhead during the training phase. The




Figure 8: Twenty bolts and fourteen washers which are used for the classification.




Figure 9: Overview of steps executed before feeding the data to the train process.

size of the bounding box was a fixed 500×500 square, so that the largest object in our dataset - bolts 11 and 19 - can fit in.

As the light reflection and shadows around the parts was minimized using light polarization and backlighting, it was possible to use rotation and translation augmentation techniques to enhance the dataset. In different runs, we realized that using augmentation without minimizing the shadows and light reflection is not helpful because the network learns these noises as part of the classification. For example, using a bolt image with shadows on its right side results in having shadows on the right side of the bolt after rotation and translation augmentations and the network may

learn that this bolt should always have shadows on its right side. In our study, mirroring and flipping augmentations do not function well for threaded fasteners - right hand versus left hand fasteners.

In short, for the preprocessing, we applied the following methods: width and height shift of 10%, a range of 0.7 to 1.3 light intensity augmentation, to make the models more insensitive to slight light changes [22], and from 0 to 270° rotation - just for the bolts. Using these augmentation techniques, we created 950 augmented data samples for each bolt and 290 for each washer, and added them to the original images for the training.

C. Training and Validation

We trained 23 different models for the fasteners. Table I summarizes the most important runs and their hyper-parameters - number of epochs, steps per epochs, and batch size. Apart from experimental runs to tune hyper-parameters of the convolutional neural network, we performed 9 different runs for training the classifier, each with different datasets and configurations. We started with AlexNet and VGG Net 19 using a Nvidia Titan Xp graphic card for training the models and achieved less than 75% accuracy. Therefore, we decided to use the InceptionV3 Model with Keras on top of TensorFlow. We also used stochastic gradient descent (SGD) optimizer with the learning rate of 0.0001 and the momentum of 0.9 for all the trainings. 80% of the data was used for the train, 10% for validation, and 10% for the test in all of our experiments.

Before using the dataset obtained by the system we conducted several experimental runs using images with normal ambient lighting to have a baseline for comparison. We started training using four similar bolts - bolts 1, 2, 3 and 4 in Figure 8 where bolt 1 and 2 were accordingly 5 and 2 millimeters smaller on their non-threaded shaft in compare with bolt 3 and

Table 1: Overview of the training results, using InceptionV3 and Keras on top of TensorFlow, optimizer=SGD (Stochastic Gradient Descent), learning rate=0.0001 and momentum=0.9. The accuracy is reported on our test dataset which is 10% of the whole data.

#	Number of Classes (Bolts/ Washers)	Images per Class (Real/ Augmented)	Epochs	Steps per Epoch	Batch Size	Accuracy
1-4	4/ 0	5/ 95	100	20	20	51.1%
		20/ 180		40		79.5%
		25/ 470		100		89.3%
		50/ 950		200		97.7%
5	10/ 0	50/ 950	100	500	20	96.2%
6	10/ 3	10-50 / 290-950	100	545	20	96.7%
7	10/ 6	10-50 / 290-950	150	590	20	97.8%
8	10/ 9	10-50 / 290-950	200	635	20	99.4%
9	20/ 14	10-50 / 290-950	250	1210	20	99.4%

Table 2: Comparison between the manual and the automatic classifications for 20 bolts and 14 washers.

	Manual Classification	Automatic Classification
Accuracy	100%	99.4%
Precision	100%	99.7%
Recall	100%	99.7%
Inference Time per Item	5 to 30 seconds	Less than half a second

4. However, the obtained trained model for bolts 1–4 using the normal ambient lighting never reached beyond 49% accuracy over the test dataset.

Subsequently we studied the performance of the following training datasets: 4 bolts (in four runs), 10 bolts, 10 bolts and 3 washers, 10 bolts and 6 washers, 10 bolts and 9 washers and 20 bolts and 14 washers - see Table I. For 4 bolts we used the data obtained by the system and applied the preprocessing steps explained in Section IV-B. By applying grid search and a process of trial and error in respect of the number of augmented images, number of epochs, number of steps per epoch, batch size and early stopping we finally obtained 98.9% accuracy on our test dataset, where the model was trained with 100 steps per epoch, and early stopping being applied after 50 epochs. We generated random data augmentations on the fly - see IV-B - which allowed us to have an infinite generated training set. Therefore, we changed the steps per epoch to limit the number of training data during experiments.

For 10 bolts, we continued using the same configurations, although the number of validation steps during the training was increased from 3 to 30 resulting in overall 96.2% accuracy. Subsequently we continued adding washers to our dataset and increased the number of epochs in each experiment - runs 6 to 8 in Table I. The accuracy of classification increased in each experiment with the addition of the washers. The surprising phenomenon during our experiments was the improvement in the classification results of the bolts after simply adding washers to our training dataset. We consider that this is because the addition of the objects with different shapes activates filters in convolution steps that were not active before. This could help to boost the filters to achieve superior results in comparison with previous results [23].

Finally, we added another 5 washers and 10 bolts to our dataset and used 1210 steps per epoch for a total number of 250 epochs, a batch size of 20 and a validation step size of 100 and applied early stopping after 50 epochs - run 9 in Table I.

V. RESULTS AND DISCUSSION

Table II summarizes the results of the classification. For training with all the 34 fasteners in our dataset, we ultimately achieved 99.4% accuracy over our test dataset which contains 50 images of each fastener.

There is no public available reference to compare the results with the current manual classification process. During our visit at MTU Maintenance (Hannover) ¹ we performed the manual classification and collected the data for comparison in two consecutive days. During the first day we manually classified a box full of 307 bolts and nuts, all from

the 34 fasteners in Figure 8, and the day after we double-checked them to measure the accuracy. Only one fastener was misclassified. Nevertheless, the double-checking process is done for all the fasteners when they are checked for damages. Therefore, we decided to use 100% accuracy, precision and recall for the manual classification. In the automatic classification approach, the inference time was less than half of a second - including taking an image of the fastener and sending it to the classifier. Comparing this with the time used in the manual classification shows that the automatic classification achieves up to 60 times faster results.

The dataset creation and training time overheads are excluded from this comparison. The time spent for dataset creation was about 5 seconds per real image - 6 hours and 31 minutes for the whole dataset. The training time was 29 hours and 14 minutes for the last run with 20 bolts and 14 washers. On the other hand, in the manual approach, technicians must create 1:1 comparison-sketch for each fastener. The time to create the sketches is unknown to us. Technicians must also pass special training courses to prepare for job. This time could be varied per trainee. For the automatic classification in a real environment, data creation and training phases must be performed every couple of months, when new unknown fasteners are detected. Therefore, in a real environment, the average data creation and training time might be more than the one we presented in this experiment.

Only a single camera is used in this study. A single camera limits the approach to a single view and a single view cannot extract all the discriminative features from some of the fasteners - namely bolts and nuts. Using multiple cameras, we can use multi-view CNNs to generalize the fasteners FGVC to all bolt types and also nuts [25]. We also used a polarized backlighting supported glass table to create our datasets and our fasteners FGVC model can only be employed using this system. However, we assume aero-maintenance companies can simply provide such a system with reusing the existing materials and tools they have on their workstations - dimmable lighting and glass tables.

VI. CONCLUSION

In this paper, we have proposed a system to minimize the shadows and light reflections arising during the creation of datasets for fasteners. Using a fixed camera, backlighting and light polarization, we were able to capture shape and the threads of fasteners with less noise than under normal room lighting conditions.

In addition, we trained a convolutional neural network to classify the fasteners. We achieved 99.4% accuracy over our test dataset of 20 bolts and 14 washers. Our trials showed that adding new classes to the dataset tends to improve the classification accuracy of other classes as well because of the application of different convolutional filters during convolutional steps.

To the best of our knowledge, this is the first reported study to apply fine-grained visual classification to small parts and fasteners. Our system can classify fasteners without part numbers or with unrecognizable part numbers with an inference time of less than half of a second. Enhancing our approach with damage detection and combining the system with automatic machineries - like a robotic arm to pick the fasteners and place them on their own package - can help aero-maintenance companies to reduce time and expense in sorting and classifying fasteners during the overhaul of airplane engines. This can also be used for classification of other small shiny objects or in other fields such as in the automotive and shipbuilding industries.

ACKNOWLEDGMENT

The Titan Xp used for this research was donated by the NVIDIA Corporation. We would like to thank NeuPro Solutions for helping us creating the system. We are also grateful to Mr. Christoph Haas and Mr. Karsten Bunde at MTU Maintenance (Hannover) for their kind invitation to visit the company and their comments on the manuscript.

REFERENCES

- [1] A. K. Jardine and A. H. Tsang. "Maintenance, replacement, and reliability: theory and applications." CRC press, 2013, page 71. D. A. Marx and R.C. Graeber. "Human error in aircraft maintenance." *Aviation psychology in practice*, 87-104, 1994.
- [2] D. G. Teer and F. B. Salem. "The formation of low friction wear-resistant surfaces on titanium by ion plating." *Thin Solid Films* 45, no. 3 (1977): 583-589.
- [3] M. E. Nilsback and A. Zisserman. "A visual vocabulary for flower classification." *CVPR*, 2006.
- [4] O. M. Parkhi, A. Vedaldi, A. Zisserman, and C. V. Jawahar. "Cats and dogs." *CVPR*, 2012.
- [5] S. Yang, L. Bo, J. Wang, and L. G. Shapiro. "Unsupervised template learning for fine-grained object recognition." *NIPS*, 2012.
- [6] H. Feng, Z. Jiang, F. Xie, P. Yang, J. Shi, and L. Chen. "Automatic fastener classification and defect detection in vision-based railway inspection systems." *IEEE transactions on instrumentation and measurement*, 63(4), pp.877-888, 2014.
- [7] S. Maji, E. Rahtu, J. Kannala, M. Blaschko, and A. Vedaldi. "Fine-grained Visual Classification of Aircraft." *arXiv preprint arXiv:1306.515*, 2013.
- [8] C. Aytekin, Y. Rezaeitabar, S. Dogru, and I. Ulusoy. "Railway fastener inspection by real-time machine vision." *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 45, no. 7 (2015): 1101-1107.
- [9] X. Gibert, V. M. Patel, and R. Chellappa. "Robust Fastener Detection for Autonomous Visual Railway Track Inspection." In *Applications of Computer Vision (WACV)*, 2015 IEEE Winter Conference on, pp. 694-701. IEEE, 2015.
- [10] D. Laptev, N. Savinov, J. M. Buhmann, and M. Pollefeys. "TI-POOLING: transformation-invariant pooling for feature learning in convolutional neural networks." In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 289- 297. 2016.
- [11] Q. Xuan, B. Fang, Y. Liu, J. Wang, J. Zhang, Y. Zheng, and G. Bao. "Automatic Pearl Classification Machine Based on a Multistream Convolutional Neural Network." *IEEE Transactions on Industrial Electronics* 65, no. 8 (2018): 6538-6547.
- [12] K. Usman and K. Rajpoot. "Brain tumor classification from multi-modality MRI using wavelets and machine learning." *Pattern Analysis and Applications* 20, no. 3 (2017): 871-881.
- [13] Y. Chai. "Advances in fine-grained visual categorization." *Doctoral dissertation*, University of Oxford, 2015.
- [14] J. Deng, W. Dong, R. Socher, L. J. Li, K. Li, and F. F. Li. "Imagenet: A large-scale hierarchical image database." In *Computer Vision and Pattern Recognition*, 2009. *CVPR* 2009. IEEE Conference on (pp. 248-255). IEEE.
- [15] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller. "Labeled faces in the wild: A database for studying face recognition in unconstrained environments." *Technical Report 07-49*, University of Massachusetts, Amherst, 2007, Vol. 1, No. 2, p. 3.
- [16] A. Khosla, N. Jayadevaprakash, B. Yao, and F. F. Li. "Novel dataset for fine-grained image categorization: Stanford dogs." In *Proc. CVPR Workshop on Fine-Grained Visual Categorization (FGVC) 2011*, June, Vol. 2, p. 1.
- [17] O. Grau. "Virtual Art: from illusion to immersion." *MIT press*, 2003.
- [18] D. ping Tian. "A review on image feature extraction and representation techniques." *International Journal of Multimedia and Ubiquitous Engineering* 8, no. 4 (2013): 385-396.
- [19] P. L. Stanchev, D. Green Jr. and B. Dimitrov. "High level colour similarity retrieval." *International Journal of Information Theories and Applications*, vol. 10, no. 3, (2003), pp. 363-369.
- [20] D. Zhang and G. Lu. "Review of shape representation and description techniques." *Pattern Recognition*, vol. 37, no. 1, (2004), pp. 1-19.
- [21] C. McCarthy, D. Feng, and N. Barnes. "Augmenting intensity to enhance scene structure in prosthetic vision." In *Multimedia and Expo Workshops (ICMEW)*, 2013 IEEE International Conference on (pp. 1-6). IEEE.
- [22] J. G. Zilly, J. M. Buhmann, and D. Mahapatra. "Boosting convolutional filters with entropy sampling for optic cup and disc image segmentation from fundus images." In *International Workshop on Machine Learning in Medical Imaging*, pp. 136- 143. Springer, Cham, 2015.

- [24] F. Hunter, S. Biver, and P. Fuqua. "Light Science & Magic: An Introduction to Photographic Lighting." Focal Press, 2015.
- [25] H. Su, S. Maji, E. Kalogerakis, and E. Learned-Miller. "Multi-view convolutional neural networks for 3d shape recognition." In

Proceedings of the IEEE international conference on computer vision, pp. 945-953. 2015