



HAL
open science

Polarimetric image augmentation

Marc Blanchon, Olivier Morel, Fabrice Meriaudeau, Ralph Seulin, Désiré Sidibé

► **To cite this version:**

Marc Blanchon, Olivier Morel, Fabrice Meriaudeau, Ralph Seulin, Désiré Sidibé. Polarimetric image augmentation. 25th International Conference on Pattern Recognition (ICPR 2020), Jan 2021, Milan, Italy. hal-02977822

HAL Id: hal-02977822

<https://univ-evry.hal.science/hal-02977822>

Submitted on 26 Oct 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Polarimetric image augmentation

Marc Blanchon^{1†}, Olivier Morel¹, Fabrice Meriaudeau¹, Ralph Seulin¹ and Désiré Sidibé²

¹ERL VIBOT CNRS 6000, ImViA, Université Bourgogne Franche-Comté, 71200, Le Creusot, France

²IBISC, Univ Evry, Université Paris-Saclay, 91025, Evry, France

[†]Email: fr.marc.blanchon@gmail.com

Abstract—This paper deals with new augmentation methods for an unconventional imaging modality sensitive to the physics of the observed scene called polarimetry. In nature, polarized light is obtained by reflection or scattering. Robotics applications in urban environments are subject to many obstacles that can be specular and therefore provide polarized light. These areas are prone to segmentation errors using standard modalities but could be solved using information carried by the polarized light. Deep Convolutional Neural Networks (DCNNs) have shown excellent segmentation results, but require a significant amount of data to achieve best performances. The lack of data is usually overcome by using augmentation methods. However, unlike RGB images, polarization images are not only scalar (intensity) images and standard augmentation techniques cannot be applied straightforwardly. We propose enhancing deep learning models through a regularized augmentation procedure applied to polarimetric data in order to characterize scenes more effectively under challenging conditions. We subsequently observe an average of 18.1% improvement in IoU between not augmented and regularized training procedures on real world data.

I. INTRODUCTION

Navigation in urban environments can be prone to errors due to highly reflective areas while using RGB cameras. On the other hand, polarization imaging can cope with such environments. Accurate navigation requires dense depth map and relies on pixel-based information versus object detection methods like bounding boxes [1]. To achieve accurate segmentation using polarimetric images as input, one requires a good and large data-set illustrating hazardous areas such as cars, water spillage, windows, etc. Unfortunately, very few polarimetric datasets are available or they do not consider scenes from urban areas. Datasets represent a critical factor affecting learning methods and particularly deep learning ones. This need is sufficiently expressed by studying the size of different widely used datasets: ImageNet [2] contains 14 million images, CityScapes [3] 25,000 images, MNIST [4] 60,000 images etc.

Analyzing these data, one can deduce that the vast majority of frequently used datasets are provided with data acquired in a conventional mode (RGB, grayscale, etc.). These datasets, in addition to being significant, are therefore easily transformable and expandable using standard augmentation methods (rotation, symmetry, distortion, etc.). Such methods not only increase the amount of data, they also make learning more generic and therefore avoids overfitting [5]. However, unlike RGB images, polarization images are not only scalar (intensity) images and standard augmentation techniques cannot be applied straightforwardly.




Fig. 1: Illustration of the diverse segmentation results obtained with different augmentation methods. The top line shows results for DeepLabV3 network not pre-trained, and the bottom line the results with a pre-trained network. From left to right are presented predictions from networks trained with: un-augmented dataset, standardly augmented dataset and augmented dataset following our procedure. The four most representative segmented classes here are: *Road* (dark yellow/orange), *Cars* (red), *Sky* (light green) and *None* (light grey).

In this paper, we consider the problem of artificially augmenting polarization images that are in direct relation to the physics of the scene acquired by the camera. This acquisition/scene relationship drastically complicates the expansion of a dataset with standard techniques. However, it is necessary to address this issue, as the available data on polarimetric modality is extremely limited despite its increasingly popular uses [6]–[9].

We have explored augmentation operations applicable to polarimetry under any condition. Initially, using the interpolation proposed by Ratliff et al. [10] combined with the channel organization proposed by Wolff and Andreou [11], we propose an approach for applying rotation and/or symmetry to polarimetric images. In our case, for the experimental results, multiple trainings were performed with either a raw dataset (limited number of images), a standardly increased dataset (significant number of images without respecting physical properties) or an augmented dataset following our procedure (significant number of images with unaltered physical properties). To focus exclusively on the impact of the data, all trainings were performed using the DeepLabV3+ network [12]. Furthermore, we demonstrate the integrity of physical properties and the effectiveness of such a process on images

that should be used for training.

As shown in Fig. 1, a significant difference is then visible when training is performed successively with differently augmented datasets: raw, augmented by omitting physical properties and with a dataset using our augmentation method. Visually, it is possible to see a segmentation capability more adequate for the recognition of areas of interest for robotic applications.

The paper is organized as follows. Section II provides a brief overview of related works including pixel-wise segmentation and polarimetric image analysis. Then, in Section III we describe the proposed augmentation method in details. Section IV presents experimental results showing the effectiveness of our approach, and the paper ends with concluding remarks in Section V.

II. RELATED WORKS

A. Pixel-Wise Semantic Segmentation

Most of the research on pixel-wise semantic segmentation use conventional imaging as input data, either RGB or depth. We can observe a constant evolution of networks accuracy due on the one side to the improvement of networks and on the other side to the increase in datasets sizes. Also, some tasks are recurrent in the community: segmentation of urban scenes [5], [12]–[17], indoor scenes understanding [18]–[22] or medical images analysis [23]–[25]. The task addressed in this paper shares a common aspect with medical imaging. Indeed, one of the common disadvantages of medical images and urban scene acquired with polarimetric sensor is the lack of large annotated datasets for training. In particular, the specificity of the polarimetric information makes it rich but also rare. Semantic segmentation being a dense and valuable information, capable of characterizing a scene at any point, it can be exploited very extensively in robotics for navigation and detection.

B. Polarimetric Modality

Polarimetry has unique properties since it allows the acquisition of changes in the state of light [11]. It particularly characterizes reflective areas. The attractiveness of these images is their ability to describe both the diffuse and specular parts of a scene. Considering all these advantages, polarization imaging offers a comprehensive range of possibilities that could enhance more standard modality such as RGB.

As polarization images not exclusively carry intensity information but also information about the surface it is reflected upon, these extra informations makes it possible to extend the range of applications in computer vision, like Shape from Polarimetry, which reconstructs specular (or partially specular) objects [7], [26]–[28]. Other works have been carried out to improve certain approaches, such as attitude estimation [8], [29], water hazard detection [9], [30], catadioptric camera calibration [31], and depth estimation [6], [32].

Considering applications in urban environments, inferring extra knowledge about objects from the reflected light, would




Fig. 2: Illustration of the behaviour of HSL images when the scene has reflective areas. Here, four clearly visible puddles are strongly coloured green (green represents the angle of polarization, corresponding approximately to 90 degrees according to the HSL color wheel [33]).

result in better perception and therefore useful for unmanned vehicles and robots.

III. METHOD

A. Reminder on polarimetric image processing

The acquisition is made by rotating a polarizer in front of the camera, or via a sensor exploiting the division of focal plane (DoFP). In both cases, pre-processing is frequently necessary to acquire the informative part of the images. Indeed, raw images are grayscale images; they provide very low level description of the scenes. Commonly, many approaches require a transformation either to have specific information such as angle and degree of polarization or to interpolate the images. Ratliff et al. [10] proposed a widely used approach to merge these two information. Four images acquired with four distinct polarizer angles (0° , 45° , 90° , 135°) are required for this pre-processing step, via four unique acquisitions or using a DoFP sensor. From these 4 images, P_0 , P_{45} , P_{90} , P_{135} , one can reconstruct three informative images: the intensity I , the angle of polarization AoP and the degree of polarization DoP . The last two images correspond, from a physical point of view, to the angle and power of reflection, respectively.

$$\text{DoP} = \frac{\sqrt{(P_0 - P_{90})^2 + (P_{45} - P_{135})^2}}{(P_0 + P_{90})}, \quad (1)$$

$$\text{AoP} = \frac{1}{2} \arctan2 \left((P_{45} - P_{135}), (P_0 - P_{90}) \right), \quad (2)$$

$$I = \frac{P_0 + P_{45} + P_{90} + P_{135}}{2}. \quad (3)$$

These informative images are combined as proposed by Wolff and Andreou [11] for display:

$$H \rightarrow 2 * \text{AoP}, \quad S \rightarrow \text{DoP}, \quad L \rightarrow I/255. \quad (4)$$

H , S and L are the three channels of a standard HSL colorspace (Hue, Saturation, Luminance). As shown in Fig.2, the singular characteristic of these new images is that the more reflective the surface is the more colored it appears.





Fig. 3: Top row: rotation on real image and impact. From left to right, the initial image, then the $+90^\circ$ rotated image and then the regularized image using the equations 5 and 7. Bottom row: illustration of the rotation procedure on a DoFP polarimetric sensor. From left to right, the sensor’s pixel subset, then subset that has undergone physical rotation, and finally the regularized pixel subset using equations 5 and 7.

B. Applying transformations to polarimetric images

Once the images have been transformed it is possible to apply augmentation processes. However, care must be taken to consider the physical properties induced by the scene (and also by the sensor). Here, we use a camera with a division of focal plane, which makes it easier to explain the influence of rotation on images. Indeed, the pixel organization due to the sensor allows preserving the integrity of the transformation by validating their order. As shown in bottom row of Fig. 3, the images are composed of super-pixels of size 2×2 . Therefore, an isolated rotation without transformation is impossible since this would mix the order of the polarizer angles. To however apply a rotation to the image, we need an additional process to “realign” the pixels and keep the physical properties of the polarization angle.

In this case, we apply an anti-clockwise rotation to the image while at the same time applying a clockwise rotation to the polarization angle, which brings back both the physics of the scene and the sensor. Let θ be the rotation angle applied to the camera, R_θ the rotation operation and H the hue channel of the image (which as a reminder corresponds to the AoP):

$$H_{\text{rotated}} = R_\theta(H_{\text{prev}} - 2 * 1\theta). \quad (5)$$

As shown in the Figure 3, bottom row, we start from the sensor and more particularly from a set of 4 DoFP pixels. If we apply pure rotation, then the sensor will be disoriented and therefore the physical direction of each image will not be respected. The regularization applied by the equation 5 then allows the image




Fig. 4: Experimental validation of augmentation process. Images (a) and (b) are acquisitions with physical rotations applied to the camera. Image (c) is the result of augmentation applied to image (a). Note the correct recovery of angle information.

to be redirected and thus to keep the correct alignment of each polarizer.

The other channels are independent of the physical properties, which allows to neglect them in this transformation step. These two channels S and L are still subjected to the rotation operation. As shown in the top row of figure 3, a significant modification has been made to the hue of the image when applying a 90-degree rotation augmentation. As an experimental proof, it is equally possible to validate this procedure through different acquisitions. In the figure 4, two acquisitions of the same scene were taken by applying a physical rotation to the sensor. Consequently, we obtain two images with a real rotation of 0 and 90 degrees respectively. Following the procedure described above, an artificial rotation is possible by using the 0-degree corresponding image and applying a 90-degree rotation followed by regularization. Slight visible shade differences are explained by a parallax influenced by the acquisition system. In the optimal case, the acquisition should have been conducted with exact rotation around the optical axis. Nonetheless, the shade obtained is respected, which validates the integrity of the procedure presented.

Another possibility of augmentation proposed in this paper is symmetry. To ensure the integrity of the physical properties of the image, symmetry is applies as follows:

$$H_{\text{flipped}} = -H_{\text{prev}}. \quad (6)$$

As before for rotation, we can graphically show in Fig. 5 the impact of pure flipping and the action of regularization. As shown in bottom row of Fig.5 the right and left images seem to be the same, but unlike visual effect induced by the images particularity, the right image is a flipped version of the left image. This type of image is a grey level representation of the Hue channel. Each point of the circle in the image represents its corresponding angle with the center as reference. In conclusion, if only the top of the left circle is considered, from left to right, we have a circular gradient ranging from 180 degrees to 0 degrees. Therefore, this image represents the complete range of possible linear polarization angles in an image. Knowing that the Hue channel of images is 360 degrees periodic, flipping consists in reversing the selected axis for the transformation.

The inversion of the values of the Hue channel is made possible by the use of the periodic 360 degree property by




Fig. 5: Top row: flipping on real image and impact. From left to right, initial image, the flipped image and the regularized image using the equations 6 and 7. Bottom row: illustration of the flipping procedure. From left to right, the base image, a flipped image and the regularized image using the equations 6 and 7.

using consecutively the equations 6 and 7. This procedure is illustrated in Fig. 5.

In both augmentation operations previously presented, it is necessary to maintain a consistency with respect to the properties of the color space. The H channel being a value between 0 and 360, it is then necessary to normalize the intensities:

$$H_{final} = H_{transformed} \pmod{360} \quad (7)$$

where H_{final} is the final image and $H_{transformed}$ is the image resulting from a rotation, symmetry or a combination of both.

To verify the previous statements, it is possible to retrieve each polarizer intensity to verify the integrity of the information calculated as follows:

$$I_{\theta} = \frac{I}{2}(1 + \rho \cos(2\phi - 2\theta)), \quad (8)$$

where I_{θ} is the intensity image from the polarizer oriented at angle θ , I is the intensity calculated in equation 3, ρ the DoP calculated in equation 1 and ϕ the AoP deduced from equation 2. Then the reconstruction of the images allows to examine the physical properties by calculating the intensity obtained for each orientation of the polarizer and arranging them as on the original sensor. This procedure verifies that the addition of regularization is necessary and that without this step, the image is altered.

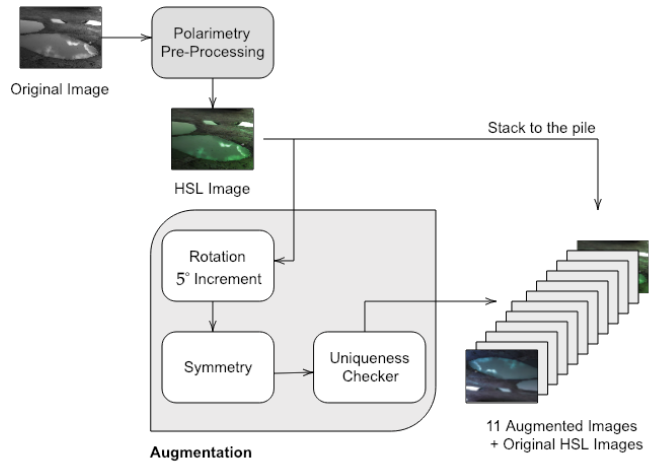


Fig. 6: Illustration of the augmentation per image procedure. This process is repeated for each image in the original dataset to obtain a consistent large dataset. Then, the entire set of augmented images is shuffled.

IV. EXPERIMENTS

A. Implementation details

In this section, different results from several augmentation procedures on an unique video sequence are presented. The idea here is to prove the augmentation reliability and show its contribution in improving segmentation results.

Datasets: The training dataset is unevenly divided into a part dedicated to training and a part dedicated to validation. The validation set is composed of 50 images selected for their characteristics allowing to have a balanced representation of the classes in the images and thus to have more precise validation metrics. As for the rest of the dataset, it is employed for training and will not be balanced so that it is closer to real world data.

Initially, the training data set is composed of 178 images that are then either kept as they are, or augmented in a standard way, without regularization, or augmented by following our approach. Whether the dataset is increased with our method or not, the transformations involved are identical and the set accumulates a total of 2136 images. As shown in Fig.6, the augmentation is performed by applying a random rotation in increments of 5° and/or symmetry with a probability of 20%. Note that each augmentation is unique.

The test dataset is a video sequence comprised of 8,049 images acquired at a frequency of 10Hz sharing many characteristics with the training dataset. It was acquired by mounting the Trioptics PolarCam 4D Technology V polarimetric camera on a remotely operated Robotnik Summit XL robot and using the ROS operating system.

Network training: The experiments are conducted using a server composed of an Nvidia Titan Xp (12GB memory)

GPU, 128GB of RAM and two CPU accumulating a total of 24 physical cores (48 threads). We use DeepLabV3+ [12] network, either with pre-trained parameters or without. Indeed, it is equally significant to compare the influence of pre-training. Since we use DeepLabV3+, it is possible to pre-train the xception subnetwork using the most recent provided model¹. The hyper-parameters and the loss function are kept identical for both models (pre-trained or not). We set empirical parameters like epoch number to 150, learning rate to $\eta = 10^{-2}$, batch size to 8 and use Adam algorithm as optimizer.

With regard to the loss, here an adapted Sørensen-Dice index is used to take into account the fact that some classes are under-represented in the dataset:

$$\lambda = \frac{\sum_c^N 1 - \frac{2|X_c \cap Y_c|}{|X_c| + |Y_c|}}{N}, \quad (9)$$

with X the label, Y the prediction, c the class and N the number of classes. Since classes are unequally represented in the dataset, this metric allows an equal valuation of each of them unlike other losses.

Metrics: Apart from qualitative evaluation such as Fig. 7, we have calculated a comprehensive range of metrics to quantitatively compare the different models according to augmentation processes. Results are shown in Tab. I and correspond to metrics computed on a sub-set composed of 15 different images where all classes are represented. This avoids redundancies and therefore allows for a more detailed case-by-case examination. The scenes correspond to urban areas with seven different classes: *Road* (dark yellow/orange), *Buildings* (grey), *Cars* (red), *Water* (blue), *Windows* (light yellow), *Sky* (green) and *None* (light grey).

B. Discussion

In this section, we will discuss the different impacts of the augmentation on the network. As shown in Fig. 7, a brief visual assessment makes it difficult to observe substantial differences between the models even if the regularized procedure seems more appropriate. It is noteworthy that, when using the proposed augmentation method, visual aberrations are reduced. A predominant visual defect of all models is the absence of building detection. This will be discussed further below.

As shown in Tab. I, a large panel of metrics has been calculated to evaluate the performance of each model. The two major metrics are mean intersection over union and recall.

For each of the major metrics, we have selected three classes representing reflective areas for comparison: *Water*, *Windows* and *Cars*. Class-specific metrics are, from a robotics application perspective, the critical points and also represent the core objective since they are derived from the detection

of danger zones. In addition, two more general metrics are proposed to evaluate the models from various angles: precision and specificity.

From a general point of view, the Tab. I points out that the regularized augmentation allows for better results in the vast majority of cases. In more detail, the IoU shows that models with an appropriate dataset perform better and mainly in class-specific metrics where we can observe substantial differences. On the contrary, the standard augmentation produces the worst results in terms of IoU. Metrics also emphasize it is better to use polarimetric data without augmentation than to augment them by neglecting physical properties of the scene. Indeed, with an unadapted augmentation method, the mean IoU is about twice lower than the result without augmentation. With the proposed augmentation approach, the mean IoU is improved by more than 40%.

The recall considers the model's ability to correctly classify zones independently of bad assignments. In our case, the recall ratio per class will indicate the ability to perceive hazardous areas in general. Once more, the majority of the high results are obtained by the model with regularized data while the unsatisfactory results are held by the trained model with standard augmentation. However, for the class *Water*, we observe that the model without augmentation achieves a better recall. This can be explained by the tendency for this model to overuse this class (this phenomenon is visible Fig. 7, row 2 columns 1 and 4).

As stated above, the common defect of the six models compared is the inability to detect buildings. This difficulty can be explained by the "physical similarity" of the components forming the *None* and *Buildings* classes. As the training was limited, it is likely that more training epochs would have benefited for this specific class. In our approach to detection for robotics, this defect is uncritical since the class is not drastically classified as a danger. Plus, it can be deduced by exploiting the other correctly segmented classes allowing the implementation of a system of rules or constraints.

We can conclude the discussion by arguing that polarimetry allows for better detection of areas prone to reflection. This capability could benefit robotic applications by improving existing algorithms. Indeed, since reflections are characterized upstream of the network, the new learnt features are specific to the reflectivity of the surfaces in the scene. This hypothesis is validated by the unsatisfactory results obtained while neglecting physical properties of the modality and moreover when analyzing the segmentation of reflective areas like cars and waters.

V. CONCLUSION

Highly reflective areas induce errors while navigating using RGB sensors. To overcome this issue, one solution is to rely on polarimetric images; however, this solution also suffers from the lack of images to train deep models and we therefore develop a successful data augmentation technique that take into

¹<https://data.lip6.fr/cadene/pretrainedmodels/>




Fig. 7: Examples of segmentation results according to the augmentation methods. From top to bottom are present the ground-truth, then consecutively the results from model with pre-training with: no augmentation, standard augmentation and regularized augmentation.

TABLE I: Impact of the augmentation procedure on DeepLabV3+ network. Specific classes have been highlighted in relation to the robotic application to witness the obstacle-wise performance. Due to the limited training, *Buildings* are almost undetected. For this reason, the averages denoted $\setminus B$ exclude the *Buildings* class from the calculation.

Augmentation	PreTraining	IoU (%)					Recall (%)					Precision (%)	Specificity (%)
		@water	@windows	@cars	Mean	Mean $\setminus B$	@water	@windows	@cars	Mean	Mean $\setminus B$		
None	No	40.0	20.6	20.8	30.5	32.2	35.2	15.8	22.5	50.9	50.0	50.0	89.6
	Yes	54.0	10.3	43.46	33.5	34.8	42.4	15.3	57.4	43.3	50.3	50.1	91.0
Standard	No	0.1	3.4	12.4	14.8	13.1	35.0	25.8	15.0	31.8	28.0	41.7	88.7
	Yes	10.2	3.0	19.7	21.8	20.0	35.2	22.9	23.4	37.0	33.4	41.2	91.2
Regularized	No	63.9	13.3	46.7	43.4	50.3	39.2	21.9	60.8	43.4	50.5	48.5	91.3
	Yes	70.0	26.6	47.1	37.8	38.5	35.0	26.0	48.0	42.0	38.5	53.7	90.7

account the vector aspect of polarimetric images. We proposed several regularization processes that maintain the integrity of the physical properties involved in polarization imaging. The experimental results show that the process of augmentation is useful and required when data amount is insufficient, but also that the results obtained at the end of the pipeline are improved in a significant way. To go further and increase the algorithm’s capabilities, other transformations such as distortion and noise addition can be investigated. Nonetheless, the use of a new modality could be beneficial for robotics and autonomous vehicles. It has been demonstrated that, with polarimetric modality, important areas for these domains can be segmented more effectively if physical properties are unaltered. Moreover, it is also possible to reconsider the structure of the CNNs used

to integrate the pre-processing step in order to avoid the HSL mapping of the polarized components.

ACKNOWLEDGMENTS

This work was supported by the French National Research Agency through ANR ICUB (ANR-17-CE22-0011). We gratefully acknowledge the support of NVIDIA Corporation with the donation of GPUs used for this research.

REFERENCES

- [1] J. Redmon and A. Farhadi, “Yolo9000: better, faster, stronger,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 7263–7271.
- [2] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “ImageNet: A Large-Scale Hierarchical Image Database,” in *CVPR09*, 2009.

- [3] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [4] Y. LeCun and C. Cortes, "MNIST handwritten digit database," 2010. [Online]. Available: <http://yann.lecun.com/exdb/mnist/>
- [5] L. Perez and J. Wang, "The effectiveness of data augmentation in image classification using deep learning," *CoRR*, vol. abs/1712.04621, 2017. [Online]. Available: <http://arxiv.org/abs/1712.04621>
- [6] K. Berger, R. Voorhies, and L. H. Matthies, "Depth from stereo polarization in specular scenes for urban robotics," in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*. IEEE, 2017, pp. 1966–1973.
- [7] Z. Cui, J. Gu, B. Shi, P. Tan, and J. Kautz, "Polarimetric multi-view stereo," in *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [8] M. Rastgoo, C. Démonceaux, R. Seulin, and O. Morel, "Attitude estimation from polarimetric cameras," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems-IROS*, 2018.
- [9] C. V. Nguyen, M. Milford, and R. Mahony, "3d tracking of water hazards with polarized stereo cameras," in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*. IEEE, 2017, pp. 5251–5257.
- [10] B. M. Ratliff, C. F. LaCasse, and J. S. Tyo, "Interpolation strategies for reducing ifov artifacts in microgrid polarimeter imagery," *Optics express*, vol. 17, no. 11, pp. 9112–9125, 2009.
- [11] L. B. Wolff and A. G. Andreou, "Polarization camera sensors," *Image and Vision Computing*, vol. 13, no. 6, pp. 497–510, 1995.
- [12] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-decoder with atrous separable convolution for semantic image segmentation," pp. 801–818, 2018.
- [13] A. Paszke, A. Chaurasia, S. Kim, and E. Culurciello, "Enet: A deep neural network architecture for real-time semantic segmentation," *arXiv preprint arXiv:1606.02147*, 2016.
- [14] G. Lin, A. Milan, C. Shen, and I. D. Reid, "Refinenet: Multi-path refinement networks for high-resolution semantic segmentation," in *Cvpr*, vol. 1, no. 2, 2017, p. 5.
- [15] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 4, pp. 834–848, 2018.
- [16] L. Li, B. Qian, J. Lian, W. Zheng, and Y. Zhou, "Traffic scene segmentation based on rgb-d image and deep learning," *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 5, pp. 1664–1669, 2018.
- [17] Z. Liu, X. Li, P. Luo, C. C. Loy, and X. Tang, "Deep learning markov random field for semantic segmentation," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 8, pp. 1814–1828, 2018.
- [18] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus, "Indoor segmentation and support inference from rgb-d images," in *European Conference on Computer Vision*. Springer, 2012, pp. 746–760.
- [19] M. Cimpoi, S. Maji, and A. Vedaldi, "Deep filter banks for texture recognition and segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3828–3836.
- [20] A. Handa, V. Patraucean, V. Badrinarayanan, S. Stent, and R. Cipolla, "Understanding real world indoor scenes with synthetic data," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [21] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, vol. 1, no. 2, p. 4, 2017.
- [22] A. Dai, A. X. Chang, M. Savva, M. Halber, T. A. Funkhouser, and M. Nießner, "ScanNet: Richly-annotated 3d reconstructions of indoor scenes." in *CVPR*, vol. 2, 2017, p. 10.
- [23] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [24] M. Havaei, A. Davy, D. Warde-Farley, A. Biard, A. Courville, Y. Bengio, C. Pal, P.-M. Jodoin, and H. Larochelle, "Brain tumor segmentation with deep neural networks," *Medical image analysis*, vol. 35, pp. 18–31, 2017.
- [25] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. Van Der Laak, B. Van Ginneken, and C. I. Sánchez, "A survey on deep learning in medical image analysis," *Medical image analysis*, vol. 42, pp. 60–88, 2017.
- [26] S. Rahmann and N. Canterakis, "Reconstruction of specular surfaces using polarization imaging," in *null*. IEEE, 2001, p. 149.
- [27] O. Morel, C. Stolz, F. Meriaudeau, and P. Gorria, "Active lighting applied to three-dimensional reconstruction of specular metallic surfaces by polarization imaging," *Applied optics*, vol. 45, no. 17, pp. 4062–4068, 2006.
- [28] O. Morel, F. Meriaudeau, C. Stolz, and P. Gorria, "Polarization imaging applied to 3d reconstruction of specular metallic surfaces," in *Machine Vision Applications in Industrial Inspection XIII*, vol. 5679. International Society for Optics and Photonics, 2005, pp. 178–187.
- [29] A. E. R. Shabayek, C. Démonceaux, O. Morel, and D. Fofi, "Vision based uav attitude estimation: Progress and insights," *Journal of Intelligent & Robotic Systems*, vol. 65, no. 1-4, pp. 295–308, 2012.
- [30] M. Iqbal, M. Morel, and F. Meriaudeau, "A survey on outdoor water hazard detection," *Skripsi Program Studi Siste Informatika*, 2009.
- [31] O. Morel, R. Seulin, and D. Fofi, "Catadioptric camera calibration by polarization imaging," in *Iberian Conference on Pattern Recognition and Image Analysis*. Springer, 2007, pp. 396–403.
- [32] A. Kadambi, V. Taamazyan, B. Shi, and R. Raskar, "Polarized 3d: High-quality depth sensing with polarization cues," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 3370–3378.
- [33] G. H. Joblove and D. Greenberg, "Color spaces for computer graphics," in *ACM siggraph computer graphics*, vol. 12, no. 3. ACM, 1978, pp. 20–25.