2024

Automated Flood Depth Estimation on Roadways

Kwame Ampofo
Old Dominion University, kampo001@odu.edu

Megan A. Witherow
Old Dominion University, mwith010@odu.edu

Alex Glandon
Old Dominion University, aglandon@odu.edu

Monibor Rahman
Old Dominion University, mrahm006@odu.edu

Ahmed Temtam
Old Dominion University, atemtam@odu.edu

See next page for additional authors

Original Publication Citation

Ampofo, K., Witherow, M., Glandon, A., Rahman, M., Temtam, A., Cetin, M., & Iftekharuddin, K. M. Automated flood depth estimation on roadways, in Optics and Photonics for Information Processing XVIII, edited by Khan M. Iftekharuddin, Abdul A. S. Awwal, Victor Hugo Diaz-Ramirez, Andrés Márquez, Proc. of SPIE 13136, 1313603 [30/09/2024]; https://doi.org/10.1117/12.3028126

Automated Flood Depth Estimation on Roadways

Kwame Ampofo[a], Megan Witherow[a], Alex Glandon[a], Monibor Rahman[a], Ahmed Temtam[a], Mecit Cetin[b], Khan M. Iftekharuddin*[a]
[a]Vision Lab, Dept. of Electrical and Computer Engineering, Old Dominion University, Norfolk, VA, USA 23529; [b]Dept. of Civil and Environmental Engineering, Old Dominion University, Norfolk, VA, USA 23529

ABSTRACT

Recurrent nuisance flooding is common across many parts of the globe and causes extensive challenges for drivers on the roadways. The prevailing monitoring methods for roadway flooding are costly and not automated or effective. The ubiquity of visual data from cameras and advancements in computing such as deep learning may offer cost-effective methods for automated flood depth estimation on roadways based on reference objects such as cars. However, flood depth estimation faces challenges due to the limited amount of data annotated with water levels and diverse scenes showing reference objects at various scales and perspectives. This study proposes a novel deep learning approach to automated flood depth estimation on roadways. Our proposed pipeline addresses variations in object perspective and scale. We have developed an innovative approach to generate and annotate flood images by manipulating existing image datasets of cars in various orientations and scales to simulate four floodwater levels for augmenting real flood images. Furthermore, we propose object scale normalization for our reference objects (cars) to improve water level predictions. The proposed model achieves an accuracy of 74.85% and F1 score of 74.32% for four water levels when tested with real flood data. The proposed approach substantially reduces the time and labor required for labeling datasets while addressing challenges in perspective/scale, offering a promising solution for image-based flood depth estimation.

Keywords: flood depth estimation, deep learning, data augmentation, object scale normalization, roadway safety

1. INTRODUCTION

Frequent nuisance flooding has become more prevalent in most parts of the world, posing significant challenges for drivers on roadways[1]. Accurate and timely monitoring of flood depth is crucial for ensuring public safety, managing traffic flow, and mitigating the impacts of flooding on transportation infrastructure[2]. However, traditional flood monitoring methods are often costly and not fully automated[3,4].

Recent advancements in deep learning techniques have shown great promise for various computer vision tasks, including image segmentation and object detection[5, 6]. These techniques have the potential to revolutionize flood monitoring by enabling automated analysis of visual data from cameras and other imaging devices. Most methods[7–9] estimate floodwater height relative to segmented reference objects such as cars. However, few incorporate learning into the flood level prediction step, instead relying on geometric calculations[8, 9] that may fail to generalize to challenging scenes involving variations in the segmentation quality, perspective, and scale of reference objects. To address these challenges, this study proposes a novel deep learning approach for automated flood depth estimation on roadways. Our specific contributions are as follows:

- We introduce an automatic data generation and annotation pipeline that simulates varying floodwater levels by manipulating image datasets of cars shown at various scales and perspectives. This approach alleviates the need for extensive manual annotation and facilitates the training of deep neural networks for flood analysis.
- We propose and systematically investigate the impact of training data augmentation and object scale normalization to improve deep learning-based flood depth estimation. To understand the individual contributions of data augmentation and object scale normalization, we perform an ablation study.
- We evaluate the proposed deep learning-based floodwater level prediction approach using a curated data set of real annotated flood images and assess its effectiveness in estimating flood depth.

*kiftekha@odu.edu

Optics and Photonics for Information Processing XVIII, edited by Khan M. Iftekharuddin, Abdul A. S. Awwal, Victor Hugo Diaz-Ramirez, Andrés Márquez, Proc. of SPIE Vol. 13136, 1313603
 · 0277-786X · doi: 10.1117/12.3028126

This study contributes to the advancement of automated flood monitoring techniques and provides valuable insights for flood management and disaster response efforts. The remainder of this paper is organized as follows.: Section 2 provides an overview of related work in flood monitoring and deep learning techniques. Section 3 describes the proposed methodology, including the data generation pipeline and the deep learning architecture. Section 4 presents the experimental results and discussion of the proposed methods. Finally, Section 5 concludes and describes future research directions.

2. RELATED WORK

Effective flood management and assessment play a vital role in disaster response and public safety. Recent advancements in deep learning techniques have paved the way for various methodologies to address the challenges associated with estimating flood extent and depth.

In Ref. 8., Park et al. propose a method that compares flooded cars with a database of synthetic cars in various orientations to estimate flood water depth. However, their approach has limitations in terms of prediction time and scalability, as it requires finding the most similar car to the flooded car before computation. Moreover, the absence of a learning-based component in the water level prediction part of their pipeline may hinder its ability to generalize to challenges such as varying segmentation quality of cars from floodwater, which can be affected by occlusions and orientation issues. Additionally, the use of synthetic data may lead to domain adaptation problems.

Similarly, in Ref. 9, Sazara et al. introduce a deep learning method to predict the depth of floodwater on streets based on side-view images of vehicles. Their approach involves four main steps: semantic segmentation, object detection, segmentation refinement, and water depth calculation. The water depth is estimated relative to the wheel size without considering the actual wheel and tire dimensions, which may vary across different vehicles. However, their method is constrained by the requirement of specific side-view orientations of vehicles for accurate flood level prediction. The flood level prediction part of their pipeline relies on geometric calculations based on the detected wheels and water edges rather than a learning-based approach. Consequently, their method may encounter difficulties in generalizing to diverse vehicle orientations and challenging scenarios where the segmentation quality of cars and floodwater is compromised.

Despite the progress made, problems involving scale and perspective issues, including vehicle orientations, persist. These challenges hinder the development of accurate and efficient flood depth estimation methods that can handle diverse scenes and orientations. Furthermore, the limited availability of annotated data poses a challenge for learning-based flood level prediction. To address these limitations, we propose a deep learning pipeline that automatically generates annotated data of various scales and perspectives to augment training, incorporates object scaling techniques to normalize the scale of reference objects (cars) across different scenes, and introduces a learning-based approach for flood depth estimation.

Unlike others[8, 9], our approach does not rely on specific vehicle orientations or geometric calculations for flood level prediction. Instead, we leverage data augmentation and the power of deep learning to learn the relationship between the visible portions of scaled cars and the corresponding flood depths, enabling our model to adapt to challenging scenarios with limited annotated data and diverse vehicle orientations.

3. METHOLODOGY

3.1 Data Preparation and Generation

To train and validate our deep learning models for flood depth estimation, we curate a dataset consisting of real flood images collected from various sources. The images are carefully annotated with multiple water levels as shown in Table 1, providing ground truth labels for training and evaluation purposes.

One of the primary challenges in applying deep learning to flood depth estimation is limited labeled data to train deep networks[7]. To address this issue, we develop a data generation method for augmenting the annotated real floodwater image dataset. Our data generation pipeline leverages the COCO dataset[10] which contains a large number of car images with diverse orientations. We extract all cars from the COCO images using their masks and place them against a blank background. To simulate the appearance of cars at various levels of submersion in floodwater, we horizontally cut portions of the cars using their masks. We assume the average height of a car is approximately 150 cm. We randomly assign a flood severity level to each car image and segment the car at a heights Level 0 (0-10cm), Level 1 (11-30), Level

2 (31-50) and Level 3 (51-70) according to information from the national weather service (NWS)[11] as seen in Table 1. In this paper, we will refer to this augmented data as the "generated" dataset. We manually review the generated images to remove any instances with excessive occlusions or unrealistic water levels, ensuring the quality and consistency of the augmented database.

Table 1. Flood severity table.

Water Level	Range (cm)	Severity
Level 0	0-10	no flood
Level 1	11-30	low water
Level 2	31-50	medium water
Level 3	51-70	high water

3.2 Object Scaling

Due to the inherent variations in image resolutions and car sizes within the generated dataset, we normalize the dimensions of every car within each image. We scale all cars to a uniform size of 224x224 pixels. For the scaling process, we identify each car using an object detector, draw a bounding box around it, and transform it using the bounding box information as follows: Equation (1) computes the scale factor for the height (or width) by dividing the target height (or width) of 224 by the original height (or width) of the segmented car. The height or width scale factor is then used to scale the height and width of the segmented car in equations (2) and (3), respectively.

$$scale = \frac{Target\ Width\ or\ Height}{Original\ Width\ or\ Height} \tag{1}$$

$$New\ Width = Original\ Width\ \times\ Scale \tag{2}$$

$$New\ Height = Original\ height \times\ Scale \tag{3}$$

Applying this scaling process, we ensure that all cars in the generated dataset have consistent dimensions, facilitating the training of our deep learning model.

3.3 Deep Learning Model for Flood Depth Estimation

Figure 1 illustrates the proposed depth estimation pipeline. For training, the input images are segmented to extract the cars. Then, we further segment the upper portion of the cars to simulate different flood levels occluding the lower portion of the vehicle. These generated partial cars are scaled to a consistent size and used to train the modified VGG-16 architecture. The trained model is then evaluated using a separate testing set of real flood images. For testing, the input image is segmented to extract cars using a pre-trained Mask R-CNN[5] model. The segmented cars are then scaled and fed into the trained CNN model to estimate the flood level.

For our flood depth estimation model, we employ a modified version of the VGG-16 architecture[12] as the backbone. VGG-16 is a well-established convolutional neural network that has demonstrated excellent performance in various computer vision tasks. We make several modifications to the architecture to adapt it to our specific requirements for flood depth estimation, including the incorporation of batch normalization layers, dropout layers, and L1 and L2 regularization at various points in the network. We replace the last layer of VGG-16 with an output layer of four units for our four-class classification task. The model is trained using the Adam optimizer with a learning rate of 0.00001, a batch size of 32, and the categorical cross-entropy loss function. We also employ early stopping with a patience of 10 epochs to prevent overfitting and restore the best weights. Our model is capable of estimating flood depths and classifying water levels into four categories: no water, low water, medium water, and high water. This information is crucial for providing timely and reliable flood monitoring and management, ultimately contributing to enhanced public safety and resilience to flood-related hazards.

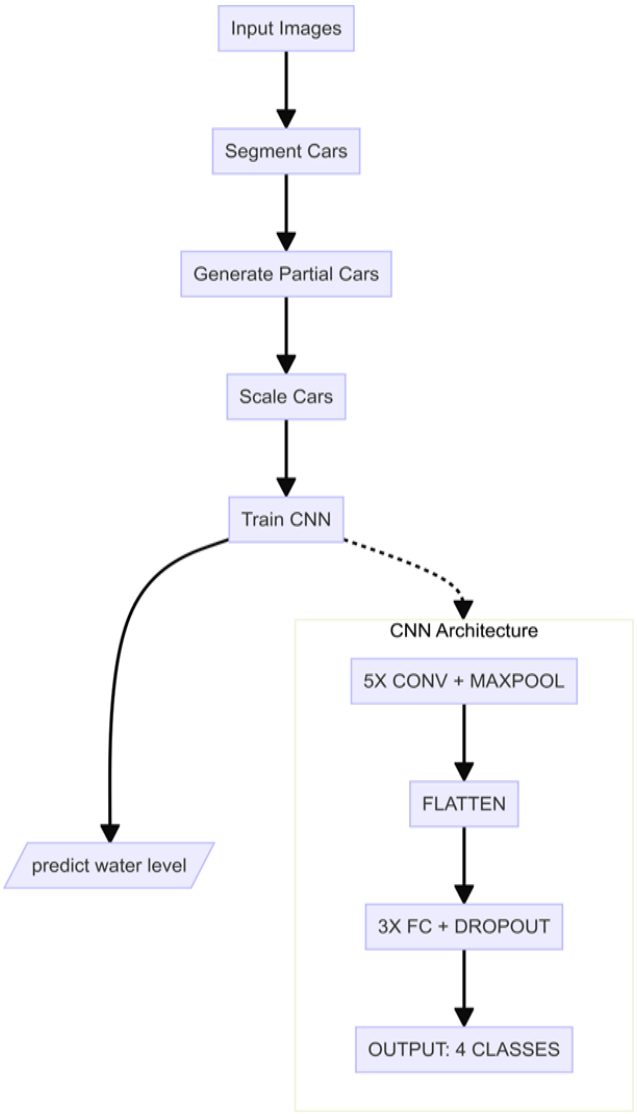

Figure 1. Automated flood depth estimation pipeline.

4. RESULTS AND DISCUSSION

4.1 Model Training and Validation

To train and validate our flood depth estimation model, we utilize a dataset consisting of 2988 generated flood images and 447 real flood images. We employ a 10-fold cross-validation approach to assess the model's performance across different data splits and obtain a reliable estimate of its generalization ability.

During the cross-validation process, our model achieves a mean validation accuracy of 73.31% with a standard deviation of 1.8% as shown in Table 2. This indicates that the model consistently performs well across different folds, with minimal variation in its predictions. To evaluate the model's performance on unseen data, we test it on a separate set of 334 real flood images. These images are carefully selected to represent diverse flood scenarios and car orientations from various locations worldwide.

Table 2. 10-fold validation of training on the mixed dataset.

Fold	Loss	Validation Accuracy (%)
1	2.5428	74.13
2	2.5752	74.42
3	2.5028	70.64
4	2.8717	73.26
5	2.6943	72.67
6	2.5706	71.51
7	2.5858	76.38
8	2.4515	75.51
9	2.4984	70.85
10	2.6499	73.76
Average	2.5943 ± 0.082	73.31 ± 1.82

Our model demonstrates reasonable results on the test set, achieving an accuracy of 74.85% and an F1 score of 74.32% as shown in Table 3. These metrics suggest that the model can effectively estimate flood depths and classify water levels in real-world scenarios, despite being trained on a combination of generated and real images.

Table 3. Performance of mixed dataset model.

	Training	Validation	Test
Number of Images	3000	344	334
Accuracy (%)	85	75.51	74.85
F1 Score (%)	88	75.49	74.32

4.2 Ablation Study

To evaluate the impact of training data augmentation and object scaling on flood depth estimation, we conduct an ablation study by comparing the performance of our model with and without the scaling pipeline on three different data sets: real images, generated images, and the mixed data set of real and generated images (data augmentation).

In the first setup, we train and test our model exclusively on real flood images. We evaluate the model's performance using 10-fold cross-validation and report the test F1 score. When object scaling is not applied (Table 4, row 1), the model achieves a test F1 score of 54.8% and a validation F1 score of 55.4%. However, by incorporating object scaling into the pipeline (Table 4, row 2), the model's performance improves, resulting in a test F1 score of 65.57% and a validation F1 score of 66.10%. This improvement highlights the importance of normalizing the scale of reference objects (cars) across different scenes, as it enables the model to learn more robust and generalizable features for flood depth estimation.

In the second setup, we explore the impact of object scaling in the context of data augmentation. We train our model on a dataset of generated images, which are created by manipulating car images to simulate various flood levels. We then evaluate the model's performance on a separate test set of 334 real flood images. When the generated images are used without scale normalization (Table 4, row 3), the model achieves a high validation F1 score of 91.2% on the generated data. However, when tested on real flood images, the model's performance significantly drops, resulting in a test F1 score of only 33.3%. This discrepancy suggests that the model overfits to the generated data and struggles to generalize to real-

world scenarios when object scales are not normalized. In contrast, when object scaling is applied to the generated images during training (Table 4, row 4), the model's performance on the real flood image test set improves substantially. The model achieves a validation F1 score of 72.91% on the generated data and a test F1 score of 67.66% on the real flood images. This improvement demonstrates the effectiveness of object scaling in bridging the domain gap between generated and real data, enabling the model to learn scale-invariant features that generalize well to real-world flood scenes.

In the third setup, we investigate the impact of object scaling when training on a mixed dataset of real and generated images (data augmentation). When object scaling is not applied (Table 4, row 5), the model achieves a validation F1 score of 82.01% and a test F1 score of 54.69%. However, by incorporating object scaling into the pipeline (Table 4, row 6), the model's performance improves significantly, resulting in a validation F1 score of 75.49% and a test F1 score of 74.32%. This improvement underscores the effectiveness of object scaling in enhancing estimation.

Table 4. Performance showing the impact of scaling.

Row	Data Description	Validation F1 Score (%)	Test F1 Score (%)
1	Real Images (without scaling)	55.4	54.8
2	Real Images (with scaling)	66.10	65.57
3	Generated Images (without scaling)	91.2	33.3
4	Generated Images (with scaling)	72.91	67.66
5	Mixed Images (without scaling)	82.01	54.69
6	Mixed Images (with scaling)	75.49	74.32

These ablation study results highlight the importance of object scaling in enhancing the accuracy and robustness of our flood depth estimation model. By normalizing the scale of reference objects across different scenes, our approach mitigates the challenges posed by variations in object sizes and orientations, leading to improved performance on real, generated, and mixed datasets. The incorporation of object scaling into our deep learning pipeline proves to be a crucial component in addressing the limitations of existing methods and enabling accurate flood depth estimation in diverse real-world scenarios.

4.3 Qualitative Results

Figure 2 showcases representative examples that highlight the model's performance across various scenes. The scaled images in the second column illustrate how the object scaling process normalizes the car sizes, enabling the model to focus on the relevant features for flood depth estimation. While the model accurately classifies the several scenes (rows 1-4) into their correct flood categories, the last scene (bottom row), which has a groundtruth of level 3, is misclassified as level 0 (no flooding). This particular example may be challenging due to the orientation of the SUV. It also appears that the water depth at the front of the SUV may be greater than the rear which is closest to the camera. This misclassification highlights the challenges and limitations of the current approach in handling certain complex scenarios. Despite this misclassification, the overall qualitative results demonstrate the capability of the proposed deep learning pipeline in handling diverse scenes and object scales. By incorporating object scaling techniques, the model can effectively normalize the size of reference objects (cars) across different scenes, improving its ability to learn meaningful patterns and estimate flood depths accurately.

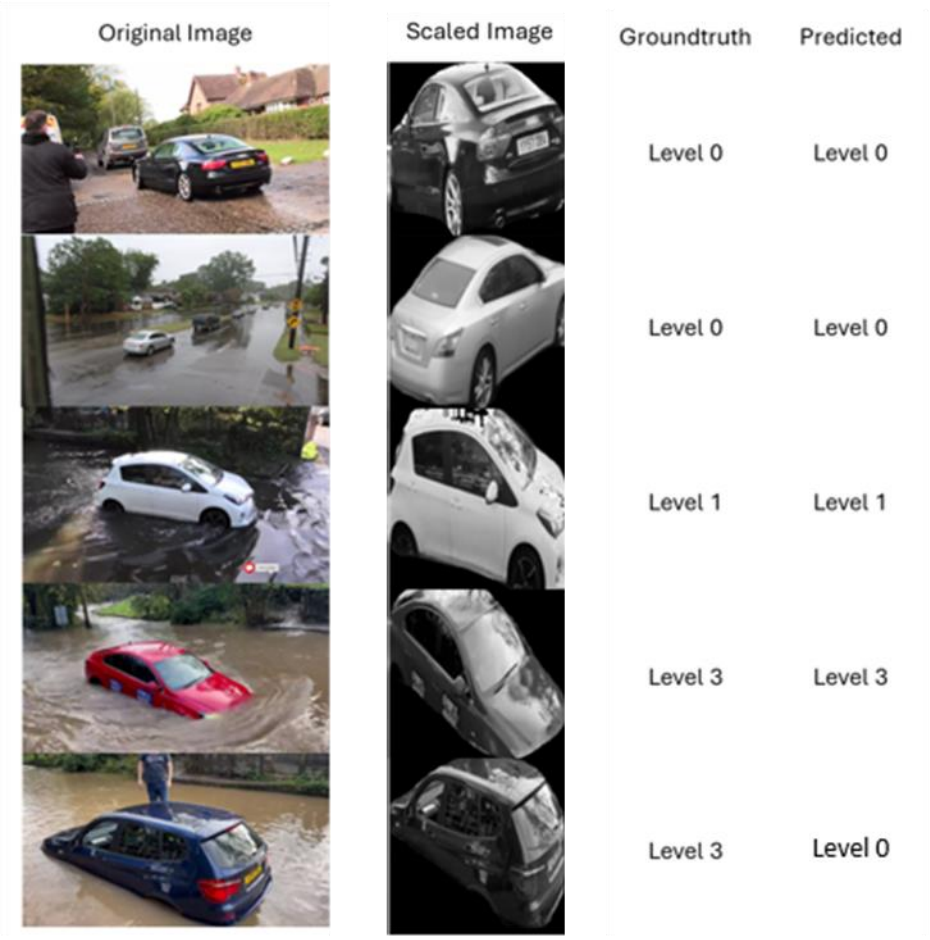

Figure 2. Representative examples of real flood data and water level predictions.

5. CONCLUSION

In this study, we propose a deep learning pipeline for automated flood depth estimation on roadways. Our method addresses the challenges of limited annotated data and variations in object perspectives/scales by introducing a data generation approach that mimics flood images of various perspectives and scales by manipulating car image datasets. When the dimensions of cars are a consistent size, the model achieves improved performance. Furthermore, we propose object scale normalization. Our ablation study reveals the synergistic effect of object scale normalization and augmenting real flood images with generated flood images for training, resulting in the best overall performance. The experimental results demonstrate the effectiveness of the proposed pipeline in classifying water levels into four categories, obtaining an overall F1 score of 74.32%. Future research could focus on addressing these challenges by incorporating additional contextual information, exploring advanced data augmentation techniques, and investigating the integration of 3D information to enhance the model's ability to handle complex flood scenes. Our study contributes to the field of automated flood monitoring and depth estimation. We hope this work may facilitate further research on timely and reliable flood management and disaster response efforts.

ACKNOWLEDGEMENTS

This research has been partially supported by National Science Foundation grants 1951745 and 1828593 at Old Dominion University.

REFERENCES

[1] National Oceanic and Atmospheric Administration, "NOAA: 'Nuisance flooding' an increasing problem as coastal sea levels rise." https://www.noaa.gov/news/ noaa-nuisance-flooding-increasing-problem-as-coastal-sea-levels-rise (July 2014).

[2] S. Van Ackere et al. "A review of the internet of floods: Near real-time detection of a flood event and its impact," *Water* **11**(11), 2275 (2019).

[3] X. Sun, R. G. Mein, T. D. Keenan, and J. Elliott, "Flood estimation using radar and raingauge data," Journal of Hydrology, vol. 239, no. 1-4, pp. 4-18, 2000.

[4] S. Cohen et al., "Sensitivity of Remote Sensing Floodwater Depth Calculation to Boundary Filtering and Digital Elevation Model Selections," Remote Sensing, vol. 14, no. 21, p. 5313, 2022.

[5] He, K., Gkioxari, G., Dollár, P., and Girshick, R. B., "Mask r-cnn," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2961–2969 (2017).

[6] Tan, M., Pang, R., and Le, Q. V., "Efficientdet: Scalable and efficient object detection," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10778–10787 (2020).

[7] Chaudhary, P., D'Aronco, S., de Vitry, M. M., Leit˜ao, J. P., and Wegner, J. D., "Flood-water level estimation from social media images," ISPRS Annals of the Photogrammetry, *Remote Sensing and Spatial Information Sciences* **4**(2/W5), 5–12 (2019).

[8] Park, S., Baek, F., Sohn, J., and Kim, H., "Computer vision–based estimation of flood depth in flooded- vehicle images," *Journal of Computing in Civil Engineering* **35**(2), 04020072 (2021).

[9] Sazara, C., Salahshour, B., Cetin, M., and Iftekharuddin, K., "A deep learning method for floodwater depth prediction on roadways from side-view real and synthetic images of vehicles," *Journal of Big Data Analytics in Transportation* **4**(1), 85–101 (2022).

[10] Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L., "Microsoft COCO: Common objects in context," in *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, 740–755, Springer (2014).

[11] "Turn around, don't drown." https://www.weather.gov/safety/flood-turn-around-dont-drown, (accessed 29/08/, 2023).

[12] Simonyan, K. and Zisserman, A., "Very deep convolutional networks for large-scale image recognition," *arXiv preprint* arXiv:1409.1556 (2014).