https://doi.org/10.2352/EI.2022.34.6.IRIACV-275

Leveraging Gradient Weighted Class Activation Mapping to Improve Classification Effectiveness: Case Study in Transportation Infrastructure Characterization

Thomas Karnowski, Deniz Aykac, Regina Ferrell, Christy Gambrell, Zach Langford, Lauren Torkelson
Oak Ridge National Laboratory, Oak Ridge, TN 37831

Abstract

Roadway "corners" are common for pedestrian use, whether designated with markings or not. Different types of markings have been deployed, ranging from simple parallel lines to more complex designs. Understanding the impact of different types of crosswalks is important for public safety. In this work we explore methods to improve the logging of marked crosswalk types. We used the Roadway Information Database from the Second Strategic Highway Research Project and used active learning methods with transfer learning to identify the crosswalk types (marked or unmarked). Upon completion we found our classifiers were unable to perform above roughly 90% correct classifications. To improve their efficacy, we separated the crosswalks into their "fine grained" types and used Gradient-Weighted Class Activation Mapping to isolate and study the features that classified the crosswalks. We compared this with sampled manually marked crosswalks and present findings. We believe this use case can represent a process to improve the active learning method for some visual machine learning applications.

Keywords: *Active learning, automated labeling, explainability, image analysis*

** Corresponding author's E-mail: karnowskitp@ornl.*[*]

Introduction

Roughly 35,000 people die on roadways in the United States of American annually [1]. In 2019, as an example, there were 36096 motor vehicle deaths, including 6205 pedestrian deaths [2]. A crosswalk is defined as "…the extension of the sidewalk or the shoulder across the intersection, regardless of whether it is marked or not. … Most jurisdictions have crosswalk laws that make it legal for pedestrians to cross the street at any intersection, whether marked or not, unless the pedestrian crossing is specifically prohibited"[3]. There have been various studies in how to better protect pedestrians and a large focus has been on signage and markings to create safer conditions; as an example, in [4], the visibility of different crosswalk marking patterns was investigated.

Additional insights can be obtained from actual data collected in the field under real driving conditions. Naturalistic Driving Studies (NDS) collect data by recording driving information using a variety of instrumentation [6]. One particular NDS, the Second Strategic Highway Research Project (SHRP2) [1], was conducted with approximately 3000 drivers in 6 data collection sites in the US between 2010 and 2013 [7]. The SHRP2 data is a valuable resource for researchers to use in their analyses of traffic and transportation conditions. An accompanying dataset, the Roadway Information Database (RID), was created [8] by driving an instrumented van on the main expected roadways of the SHRP2 study. The RID is a geospatial database of roadway features and high-resolution imagery, and contains other information including weather information, crash histories, etc. depending on the collection site. Several researchers have used these and similar resources to explore the relationships between roadway features and actual driving events particularly for pedestrian / driver interactions [10][11][12].

In our case, we had a scientific support goal of simply categorizing all the intersections in the RID and determining if they were marked crosswalks or unmarked crosswalks. We also had a secondary goal of identifying the type of markings used, with some examples shown in Figure 1. We explored a simple but effective method of performing this ground truth, leveraging the binary nature of our original goal (marked vs unmarked crosswalk) and then separating the subsequent classification of crosswalk types into additional binary classes (i.e., standard vs not-standard, then taking the not-standard set and separating into continental vs not-continental, etc.)

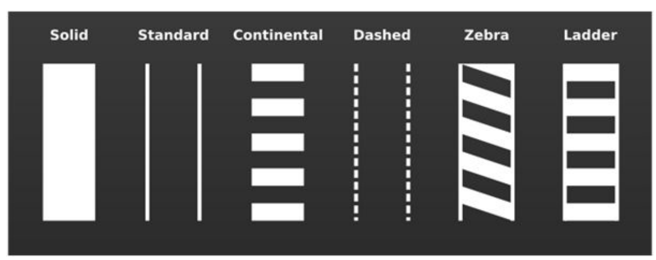

Figure 1. Examples of different crosswalk markings [13].

We used a machine-learning assisted method which we devised ad-hoc to complete this task. We obtained overhead imagery of each intersection, then sampled several cases to separate them into marked and unmarked examples. We used transfer learning to train a convolutional neural network (CNN), then sorted the unlabeled

[*] *This manuscript has been authored in part by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (http://energy.gov/downloads/doe-public-access-plan).*

examples and manually corrected them. We repeated this process until all the imagery were classified. The marked crosswalk types were classified completely manually but in a binary fashion. Upon completion, we had some outstanding questions about the labeling process, as well as the effectiveness of the transfer learning. We explore these concepts here by using this pragmatic labeled data set as a case study in automated labeling and active learning methods, as well as trying to understand the transfer learning feature extraction process. We also explore whether explainable Artificial Intelligence (AI) methods such as gradient class activation mapping (GradCAM) [14] could be used to help inform and assist classification problems which move from coarse detections to more fine-grained classifications. We note that [15] pursues a similar objective of labeling marked crosswalks, particularly for a single marking pattern type; we believe our main difference is a focus on methods to improve annotation and transition from coarse to fine-grain classification, as opposed to developing a network for crosswalk classification. In the rest of this paper, we discuss the data set, followed by the process for the automated labeling effort in separating marked / unmarked crosswalks with different strategies on sample selection. We extend this approach to marked crosswalk types. We compare using explainable methods with more brute-force methods for the fine-grained classification of the marked crosswalk types. We conclude with a discussion on findings and future work.

Data Set

The geospatial database of the RID features 40,387 intersections which were exported to comma separated value files. Each intersection has a latitude-longitude coordinate which was used to retrieve an overhead image using the Google Static Map API [15] with a function developed for the MATLAB environment [17]. Some intersections were not retrievable, so we were limited to roughly 30,007 intersections of interest. An example of a set of downloaded images is shown in Figure 2. The images ranged in size from 448x448 pixels to 640x640 pixels in size. The types and number of crosswalks are shown in Table 1 below.

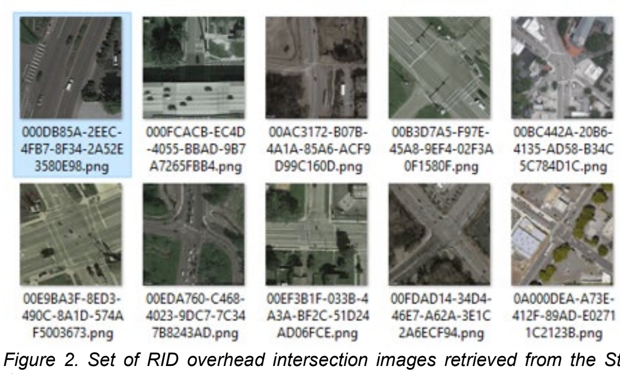

Figure 2. Set of RID overhead intersection images retrieved from the Static Google Map API [15].

Table 1. Types of Crosswalks and Quantities

Type	Quantity
Unmarked	26528
Standard	3805
Continental	3755
Ladder	2411
Zebra	61
Dash	35
Multiple/Unknown	412

Methods

Image Review Statistics

Our first method involves estimating the time to mark errors in a classified / labeled set of images. Given the size of the images and the nature of the labeling process (i.e., pick out unmarked vs marked crosswalks), we executed our hand labeling processing by displaying 16 images in a 4x4 array. We leveraged the nature of the labeling process and created an interface where the labeler could pick out the incorrect entries; for example isolate the unmarked crosswalks in an array of images from the marked crosswalks. If all entries are marked crosswalks, then no interfacing is needed (the user simply confirms all labels are correct). The approach we took has some similarities to active learning [18][19]. We also note that there are existing assisted annotation tools, but we sought to mimic our initial ad-hoc approach. In Figure 3 we show the mean time to mark the incorrectly classified images in a 4x4 grid. This was generated by experimental trials using a graphical-user interface with simple user instructions to mark errors, then indicate when the review was complete. There is a clear performance gain by allowing access to multiple images, as we note we also conducted an experiment with a single image marked correct or incorrect which averaged 1.5 seconds (which is consistent with the result below for a 1-error case.) Evaluating a worst-case 16 errors takes roughly 8 seconds on average and thus there is an approximately 3x time savings on a per-image basis. We used these statistics in our subsequent tasks described below.

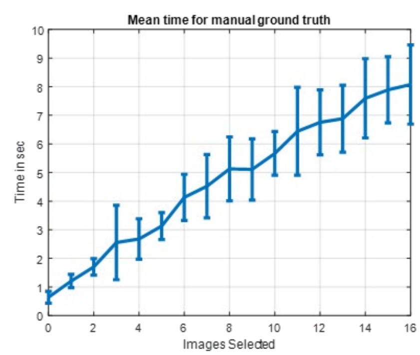

Figure 3. Mean time to confirm/correct in 4x4 image array for a 2 class problem.

Labeling Process

Our second method concerns how we execute the labeling process. Our overall approach is to (1) create a ground-truth data set, (2) train a classifier on the set, (3) machine label the remaining entries in the data set, (4) present the results in a specific manner based on the machine predictions to allow a user to correct / confirm the decisions, and (5) use the corrections / confirmations to repeat the process with a new ground truth data set (Figure 4). The main variable in this process which we investigated was the presentation manner of step 4. We used convolutional neural networks for the machine learning method, so for each image, we received a value between 0 and 1 for the membership in each of the two candidate classes. Thus, we used one of these measures as our score. We used two evaluation methods for presenting the results to the user, shown in the graph at step 4 in Figure 4. In the first method (dubbed "extremes"), we focused on the extreme results where the score was high for one of the classes. In the second method (dubbed "middle"), we focused on the cases where the scores were most ambiguous (choosing the middle images from the region around the "cross over" point where the scores are nearly equal). Our rationale was that the extreme cases would be easier to confirm/correct with fewer errors and thus decrease the time as indicated in Figure 3. We also hoped that capturing gross errors would help direct the classifier in a more nuanced manner. On the other hand, the middle approach would capture more problematic, ambiguous cases and perhaps increase the variation the classifiers used. As opposed to active learning, we are not seeking to choose the best candidates for a training set, but rather we want to reduce the overall time needed for evaluation by exploring how to select training set samples. We targeted confirmation / correction on a fixed number of samples for both approaches and doubled this from the seed set which was chosen arbitrarily as 1000 data points (thus we used training sets of 1000, 3000, 7000, etc with the new number of points doubling each time). Our experiments consisted of using these two approaches on the marked vs unmarked crosswalk problem, and we then repeated them on a 'fine grained' classification of the marked crosswalks as well, using standard vs ladder/continental and then ladder vs continental. (We omitted the underrepresented classes for illustrative purposes; generally, this would be more difficult in an actual application of this method.)

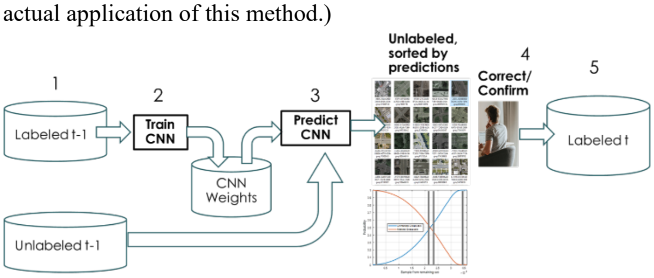

Figure 4. Flow diagram of method. (1) start with set of labeled images; (2) train CNN; (3) use trained CNN to perform prediction on unlabeled set; (4) user interface to correct / confirm results, based on strategies depicted in graph either using extreme cases or cases in the middle of the distribution (fold-over point); (5) new labeled set is fed back to step 1 and process is repeated.

Fine-Grained Classification

For the final methods, we investigated how the results of the coarse marked / unmarked crosswalk problem may be leveraged to perform a more "fine grained" classification, specifically, into the different types of crosswalks shown in Figure 1. We sought to leverage explainable methods, which seek to identify important regions in images where CNNs concentrate learning. Previous methods for pixel-space gradient visualizations such as Guided Backpropagation [22] and Deconvolutional networks [20] are high-resolution and highlight fine-grained details in the image, but are not class-discriminative [14]. Deconvolutional networks aim to approximately reconstruct the input of each layer from its output by understanding neuron activations in feature maps [20]. Guided Backpropagation leverages which elements are positive in the preceding layer with Deconvolutional networks by setting the gradient and negative gradients to zero to highlight the pixels that are important in the image [22].

In contrast, Class Activation Mapping (CAM) are class-discriminative, which localizes the category or class of the image by using global average pooling in CNNs [23]. CAM works by global average pooling on the convolutional feature maps just before the final output layer and use those as features for a fully-connected layer that produces the importance of the image regions by projecting back the weights of the output layer on to the convolutional feature maps. However, CAM requires feature maps to directly precede the prediction layer and is only applicable to CNN architectures performing global average pooling over convolutional maps immediately prior to the prediction layer.

Instead, Gradient-weighted Class Activation Mapping (Grad-CAM) highlights important regions of an input image for CNN's prediction using the gradients of any target concept, flowing into the final convolutional layer to produce a coarse localization map highlighting the important regions in the image for predicting the concept [14]. Grad-CAM may be combined with existing pixel-space visualizations to create a high-resolution class discriminative visualization (Guided Grad-CAM) [14]. Grad-CAM forward propagates an input image to obtain a raw score for a class of interest setting gradients to zero and the desired class to 1, then the signal is backpropagated to the rectified convolutional feature maps of interest to create heatmaps of where the model looks for the class of interest. Guided Grad-CAM visualizations are created by performing a pointwise multiply on the heatmap with guided backpropagation. Grad-CAM is further improved by other works, such as Grad-CAM++ [21] and Ablation-CAM [24].

We specifically applied GradCAM to our final labeled dataset and used the heatmaps produced as proxies for bounding box data, by using a simple threshold of 0.80 and computing the bounding box around the threshold area. We compared these with a sample of hand-drawn oriented bounding boxes on the image features of interest for the marked crosswalks. Our hope was that these bounding boxes could be leveraged for an improvement on the fine-grained classifier; thus we tested how effective the GradCAM regions were at localizing those features that were unique to each fine-grained class.

Results

The estimated times to correct / confirm a 4 x 4 array of images are shown in Figure 3, so we simply note that this method allows an operator to perform a more rapid per-image confirmation / correction than the review of a single image would allow. As a comparison, we estimated a single image takes 1.5 s to estimate, so the entire dataset of 37007 images would require 55,510 seconds to label (15.4 hours). Moving to our "divide and conquer" experiments, the times to complete each step of the assisted annotation are shown in Figure 5 through Figure 7. We summarize these results in Table 2. Each row represents the labeling objective (Marked vs Unmarked for coarse, then Standard vs Cont/Ladd and Continental vs Ladder for fine), with the two approaches shown and the total time for each method and objective presented.

Table 2. Timing for Different Approaches and Test Classes

Test Case	Approach	
	Extreme	Middle
Marked vs Unmarked	4150 s	4696 s
Standard v Cont/Ladd	981 s	1367 s
Continental v Ladder	373 s	560 s

Overall, the extreme approach saves time, in all cases. However, the time savings occur at the initial phases where the training set is smaller and likely has less variation. The intermediate CNNs used here (based on Alexnet [25]) were created by splitting the available set into 4 pieces. The unlabeled data was set aside for prediction only for the next iteration. The labeled data was split by taking the minimum number of entries for the two classes, then dividing that into 50% for training, 25% for validation, and the remaining was used for testing. This created balanced training and validation sets, but resulted in an unbalanced testing set (and obviously an unbalanced unlabeled set). Since the overall data set was unbalanced, this meant the testing set was more unbalanced than the overall set. We note that while we were not focused on the performance of the CNN, the CNNs used in the intermediate steps here (based on Alexnet) were typically better performing on the validation data for the extreme approach than the middle approach – which makes intuitive sense because they likely had an easier data set for the initial iterations, but eventually ran into the more difficult sets later in the process and thus the operator had to work harder in latter stages than in earlier stages. Meanwhile, the middle approach took more time initially because there were more difficult training data at the start, which was harder to correct and confirm.

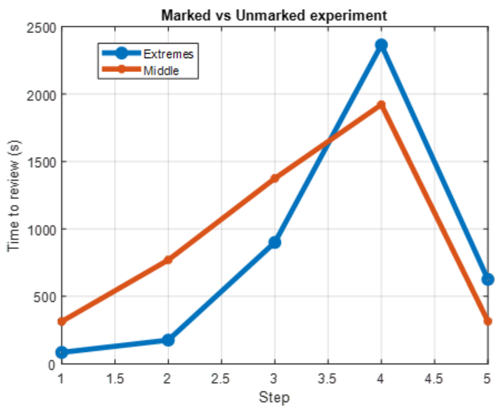

Figure 5. Timing for marked vs unmarked experiments

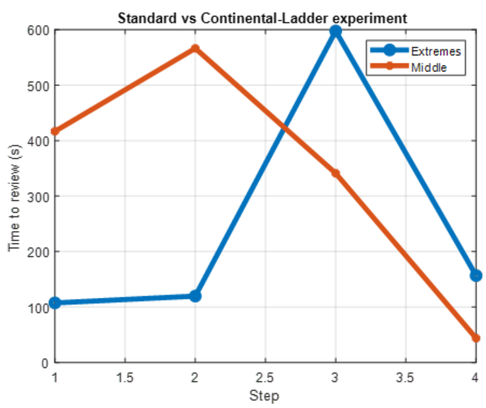

Figure 6. Timing for Standard vs Continental/Ladder fine-grain experiments

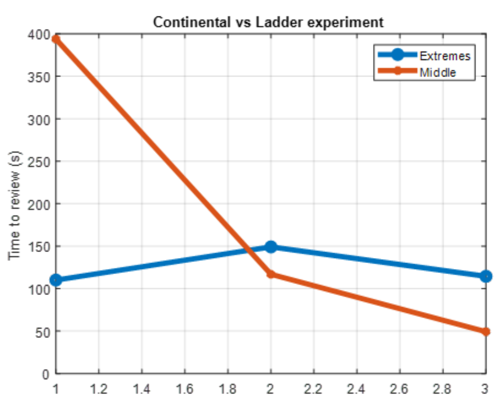

Figure 7. Timing for Continental vs Ladder experiments

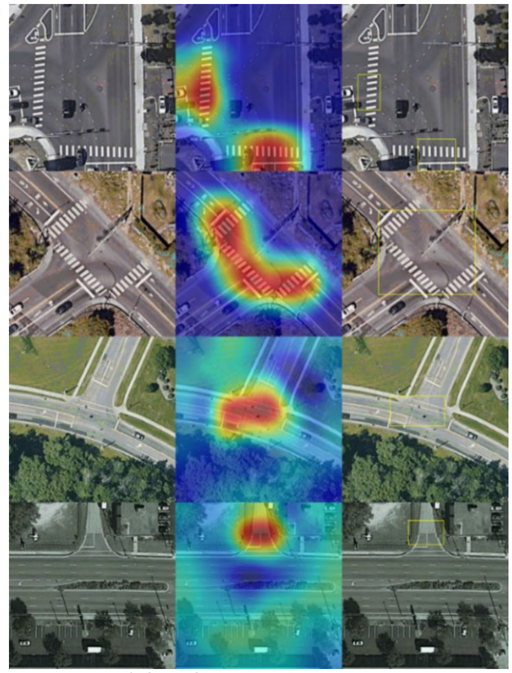

Figure 8. Examples of Grad-CAM based heatmaps and localization. Left column are the original images, with the heatmaps in the middle and the isolated bounding boxes on the right. Top row shows a good localization. The second row shows a good localization for the entire intersection, but does not distinguish the individual crosswalks. The third and fourth rows scored at 0.63 and 0.50 for the marked crosswalk class, and in both cases do not localize the key feature well although they do focus on the intersection itself.

We performed several experiments to understand the potential of using GradCAM to direct from the coarse classification problem to the fine-grained problem. These experiments were conducted more to determine the efficacy of the approach rather than as a full-scale implementation, but we note that regardless our goal was to classify / review the data, not generate the best classifier. We first trained a final coarse classifier (this time using the InceptionV3 topology [26]) and used its weights to generate GradCAM heatmaps. Some examples are shown in Figure 8, along with the bounding box drawn by simple thresholding the heat map at 0.8. Generally we noted that when the class score for marked crosswalk was high, the GradCAM heatmap showed a good mapping to the fine-grained crosswalk. But, when the score was lower the result was not as promising, and even when the scores were high the heatmap did not always separate the features into separated crosswalks. We quantified this by performing hand-segmentation of a sample of 1979 crosswalks in the data set. We created a custom interface that allowed crosswalk segmentation by clicking the four corners of the crosswalk. The time to generate these bounding boxes is shown in Figure 9. We compared the hand boxes with the automated ones using the Intersection over Union (IoU) score, and Figure 10 shows a plot of the IoU score as a function of the GradCAM score. This reveals that many IoU are low (which is not unexpected, especially since the hand-drawn results are oriented boxes and the GradCAM are not), but significantly the higher IoU scores occur more frequently when the GradCAM score is high, suggesting that high GradCAM does indeed localize the image features well. A visual review of the GradCAM bounding boxes found that roughly 75% of the GradCAM boxes gave good localization, and thus the localization property for this particular application is reasonably good.

As a final check, we compared the classification of the GradCAM localized data set for fine-grained classification to the non-localized set. We used an InceptionV3 classifier model, with transfer learning, and trained a single instance on the GradCAM focused images (made by extracting the bounding box) and the entire image. Our rationale was that the GradCAM focused cases may allow better control of undesired variation in the data set, such as the surrounding imagery and other roadway infrastructure and markings. However, the results are comparable (Table 3 and Table 4), suggesting that the GradCAM focused imagery has too much clutter as well; while 75% of the images localize well, that leaves 25% that do not.

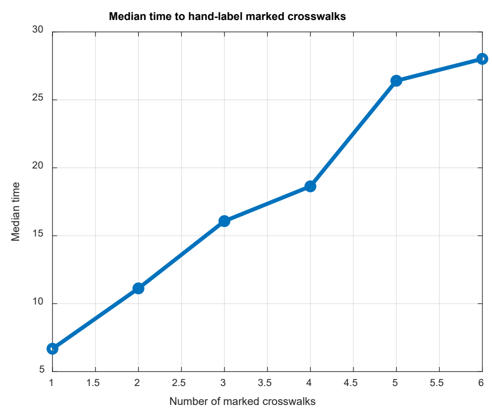

Figure 9. Median time to draw marked crosswalk features, based on the number of marked crosswalks in each image.

Table 3. Fine Grained Using GradCAM Focused Images

Actual	Predicted		
	Continental	Ladder	Standard
Continental	0.92	0.02	0.05
Ladder	0.02	0.90	0.08
Standard	0.06	0.07	0.86

Table 4. Fine Grained Using Uncropped Images

Actual	Predicted		
	Continental	Ladder	Standard
Continental	0.91	0.03	0.06
Ladder	0.03	0.99	0.08
Standard	0.05	0.06	0.89

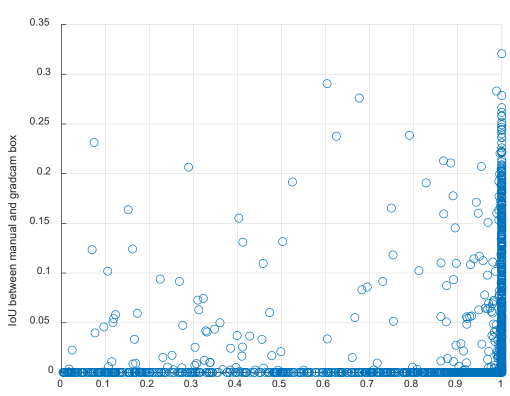

Figure 10. IoU score for hand drawn crosswalks plotted as a function of the GradCAM score. While many IoU are zero, we see that non-zero IoU are more common when the GradCAM score is near unity on the right side of the plot.

Conclusions

There are many factors to consider when executing hand-labeling or confirmation/ correction of a data set. Our experiments showed that the extreme method seems to save time at the beginning of the process, while the middle method takes longer initially but improves over time. The extreme method did take less overall time in our experiments, and both were superior to hand-labeling with no automation assistance. There are open-source and commercial tools (for example [27][28]) that can be used for image annotation as well, and they may be well suited for this but we focused on the particulars of this application.

The GradCAM results showed that high scoring results were often good at picking out fine-grained features of interest – and thus, the coarse classifier truly learned important features of interest that defined the fine-grained problem. However, we were not able to leverage this to effectively assist a fine-grained classification system. One promising idea is the concept of using the GradCAM boxes as an initial approach to hand-drawn boxes. This method would work by having a user manually review GradCAM boxes for confirmation / correction, then using an object-based CNN method to iteratively generate additional examples. There would likely be considerable time savings with this approach if it could be implemented. Finally, there remain many unanswered questions; a more thorough literature review is likely needed. The training time needed for each step is a consideration that we did not address. Other ideas include sampling from both the middle and the extreme ends; leveraging the GradCAM analysis for sample selection; and studies on the impact of more flexibility in the number of samples selected in each step.

Acknowledgments

We would like to acknowledge the support and assistance of Virginia Tech Transportation Institute, research staff at FHWA Turner Fairbank, and Omar Smadi of the Center for Transportation Research and Education at Iowa State University.

References

[1] "Early Estimate of Motor Vehicle Traffic Fatalities for the First Half of 2021," https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/813199

[2] https://www.iihs.org/topics/fatality-statistics/detail/pedestrians

[3] *Uniform Vehicle Code and Model Traffic Ordinance*, Millennium Edition, National Committee on Uniform Traffic Laws and Ordinances, Evanston, IL, 2000.

[4] Fitzpatrick, Kay, et al. Crosswalk marking field visibility study. No. FHWA-HRT-10-068. United States. Federal Highway Administration. Office of Safety Research and Development, 2010.

[5] Zegeer, Charles V., et al. Safety effects of marked versus unmarked crosswalks at uncontrolled locations final report and recommended guidelines. No. FHWA-HRT-04-100. United States. Federal Highway Administration. Office of Safety Research and Development, 2005.

[6] Guo, Feng. "Statistical methods for naturalistic driving studies." Annual review of statistics and its application 6 (2019): 309-328.

[7] Hankey, Jonathan M., Miguel A. Perez, and Julie A. McClafferty. *Description of the SHRP 2 naturalistic database and the crash, near-crash, and baseline data sets*. Virginia Tech Transportation Institute, 2016.

[8] Smadi, O., Hawkins, N., Hans, Z., Bektaş, B., Knickerbocker, S., Nlenanya, I., & Hallmark, S. (2015). Naturalistic driving study: development of the Roadway Information Database. SHRP 2 Report S2-S04A-RW-1 (https://onlinepubs.trb.org/onlinepubs/shrp2/SHRP2prepubS04ARepo rt.pdf)

[9] Ghasemzadeh, Ali, et al. "Complementary methodologies to identify weather conditions in naturalistic driving study trips: Lessons learned from the SHRP2 naturalistic driving study & roadway information database." Safety Science 119 (2019): 21-28.

[10] Pantangi, Sarvani Sonduru, et al. "Do high visibility crosswalks improve pedestrian safety? A correlated grouped random parameters approach using naturalistic driving study data." Analytic methods in accident research 30 (2021): 100155.

[11] Sheykhfard, Abbas, et al. "Analysis of the occurrence and severity of vehicle-pedestrian conflicts in marked and unmarked crosswalks through naturalistic driving study." Transportation research part F: traffic psychology and behaviour 76 (2021): 178-192.

[12] Du, Eliza Yingzi, et al. "Pedestrian behavior analysis using 110-car naturalistic driving data in USA." 23rd International Technical Conference on the Enhanced Safety of Vehicles (ESV). 2013.

[13] https://sf.streetsblog.org/2013/11/01/noticed-more-continental-crosswalks-theyre-now-standard-on-sf-streets/

[14] Selvaraju, Ramprasaath R., et al. "Grad-cam: Visual explanations from deep networks via gradient-based localization." Proceedings of the IEEE international conference on computer vision. 2017.

[15] Berriel, Rodrigo F., et al. "Deep learning-based large-scale automatic satellite crosswalk classification." IEEE Geoscience and Remote Sensing Letters 14.9 (2017): 1513-1517.

[16] http://code.google.com/apis/maps

[17] Val Schmidt (2022). get_google_map, MATLAB Central File Exchange. Retrieved January 2, 2022. https://www.mathworks.com/matlabcentral/fileexchange/24113-get_google_map

[18] Settles, 2010] B. Settles. Active learning literature survey. Computer Science Technical Report 1648, University of Wisconsin-Madison, January 2010

[19] Ren, P., Xiao, Y., Chang, X., Huang, P. Y., Li, Z., Gupta, B. B., ... & Wang, X. (2021). A survey of deep active learning. ACM Computing Surveys (CSUR), 54(9), 1-40.

[20] Mahendran, Aravindh, and Andrea Vedaldi. "Salient deconvolutional networks." European Conference on Computer Vision. Springer, Cham, 2016.

[21] Chattopadhay, A, et al. "Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks." Winter Conference on Applications of Computer Vision. IEEE, 2018.

[22] Springenberg, Jost Tobias, et al. "Striving for simplicity: The all convolutional net." arXiv preprint arXiv:1412.6806 (2014).

[23] Zhou, Bolei, et al. "Learning deep features for discriminative localization." Proceedings of the IEEE conference on computer vision and pattern recognition. 2016.

[24] Ramaswamy, Harish Guruprasad. "Ablation-cam: Visual explanations for deep convolutional network via gradient-free localization." Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. 2020.

[25] Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. "Imagenet classification with deep convolutional neural networks." Advances in neural information processing systems 25 (2012): 1097-1105.

[26] Szegedy, Christian, et al. "Rethinking the inception architecture for computer vision." Proceedings of the IEEE conference on computer vision and pattern recognition. 2016.

[27] https://openvinotoolkit.github.io/cvat/docs/

[28] https://www.mathworks.com/help/vision/ug/create-automation-algorithm-for-labeling.html

Author Biography

Thomas P Karnowski received a BSEE from the University of Tennessee (1988), a MS from NC State University (1990), and a PhD from Tennessee (2010). He has worked at ORNL since 1990 on a variety of research projects in applications of signal and image processing

Deniz Aykac received a BS in Physics from the Bogazici University (1994) and an MS in Biomedical Engineering from The University of Iowa (2000). She has worked at Oak Ridge National Laboratory since 2002 extensively on 3D medical image processing, image, and video analysis.

Regina K. Ferrell received a BS in Electrical Engineering (1984) and a MS in Electrical Engineering (1994) from the University of Tennessee, Knoxville. She has worked for Oak Ridge National Laboratory since 1992 on a variety of projects with her primary focus on applications in image processing and big data.

Christy Gambrell received at BS in Psychology from the University of Phoenix (2014) and is currently pursuing her MPS from Tennessee Technical University (2022). She has worked at Oak Ridge National Laboratory since 2019 as an administrative assistant and has assisted on a variety of research projects in signal and image processing.

Zach Langford completed his PhD in Energy Science and Engineering from the University of Tennessee. His PhD research applied data mining and machine learning methods for understanding environmental processes. He received his M.Sc. in Civil Engineering from Penn State University focusing on object-based imagery analysis and artificial neural networks. His research interests include developing new data analytics methods in scientific disciplines such as cyber-physical systems, image/signal processing, and environmental sciences.

Lauren Torkelson received a BS in Cyber Operations from Dakota State University (2019) and a MS in Cyber Defense from Dakota State University (2020). She has worked at Oak Ridge National Laboratory since 2021, and has worked on several research projects involving machine learning.