Data Augmentation using Generative Adversarial Networks to Reduce Data Imbalance with Application in Car Damage Detection

Mohamed Mahyoub
Tutor Reach
United Kingdom
m.mahyoub@tutorreach.com

Friska Natalia
Faculty of Engineering and Informatics
Universitas Multimedia Nusantara
Tangerang, Indonesia
friska.natalia@umn.ac.id

Sud Sudirman
School of Computer Science and Mathematics
Liverpool John Moores University
Liverpool, United Kingdom
s.sudirman@ljmu.ac.uk

Panos Liatsis
Khalifa University
Abu Dhabi, United Arab Emirates
panos.liatsis@ku.ac.ae

Abdulmajeed Hammadi Jasim Al-Jumaily
Universiti Putra Malaysia
Serdang, Malaysia
a.aljumaily@ieee.org

***Abstract*— Automatic car damage detection and assessment are very useful in alleviating the burden of manual inspection associated with car insurance claims. This will help filter out any frivolous claims that can take up time and money to process. This problem falls into the image classification category and there has been significant progress in this field using deep learning. However, deep learning models require a large number of images for training and oftentimes this is hampered because of the lack of datasets of suitable images. This research investigates data augmentation techniques using Generative Adversarial Networks to increase the size and improve the class balance of a dataset used for training deep learning models for car damage detection and classification. We compare the performance of such an approach with one that uses a conventional data augmentation technique and with another that does not use any data augmentation. Our experiment shows that this approach has a significant improvement compared to another that does not use data augmentation and has a slight improvement compared to one that uses conventional data augmentation.**

Keywords—Image Classification; Deep Learning; Generative Adversarial Networks; Data Augmentation; Car Insurance Claim

I. INTRODUCTION

The success of a car insurance company is determined by how quickly and accurately it can settle a claim when a car accident occurs [1]. It can lose a lot of money in claim leakages because of manual inspection of cars which are often biased and inaccurate. There has been a considerable increase in the number of claim requests for car insurance and to speed up the claim settlement accurately, there is a need for an automatic car damage detection and car damage severity estimation system. This will help filter out any frivolous claims that can take up time and money to process manually. Research to find solutions to these computing problems fall into the computer vision field, which is a field of artificial intelligence that uses computer algorithms to derive high-level understanding from digital images or videos. Computer vision methods traditionally use machine learning algorithms to extract and process information from images. Machine learning is a well-studied area of research and has been been used to solve, not just computer vision problems, but also a wide range of computing tasks including non-destructive testing in engineering [2], digital watermarking for copyright protection [3], and tourism data analytics [4].

More recently, progress in machine learning is eclipsed by deep learning, a subset of machine learning that uses artificial neural networks in which a high number of processing layers are used to extract progressively higher-level features from the input data. With the increase in computing resources and power, solving computer vision problems using deep learning algorithms is becoming more accurate and efficient. Specific to solving computer vision problems, there has been tremendous progress in one type of deep learning model namely Convolutional Neural Network (CNN). Many CNN models have been proposed to solve a wide range of computer vision problems ranging from image classification and organ measurement in medical images [5][6], to pest detection and classification in agriculture [7][8], and to sign language recognition [9]. Many of the proposed methods achieve their goals through the use of readily available pre-trained deep learning models. Many of these pre-trained models have been trained using ImageNet [10]. Although the dataset itself contains over 14 million images with more than 20,000 categories, most pre-trained models were trained using only a subset of the dataset, for example, the pre-trained VGG16 and VGG19 models were trained using 1.2 million images with 1000 categories [11].

To utilize a pre-trained model for different applications, such as for car damage detection, a transfer learning approach needs to be used whereby the weights of the model are used as initial weight values. The classification layers of the model are replaced with a new classification layer specifically targeted to the new task. The resultant neural network (after replacing the classification layers) needs to be trained on the new dataset, which in our case, is the damaged car dataset. However, the availability of car damage datasets is limited and the available ones suffer from small dataset size and data imbalance [12]. This research will attempt to address those challenges by generating more images by utilizing data augmentation techniques using Generative Adversarial Networks (GANs).

The main aim of this research is to propose a data augmentation technique to improve the suitability of a damaged car dataset by generating synthetic damaged car images using GANs. With the help of data augmentation using GANs the issue of data insufficiency and data imbalance in the damaged car can be resolved. The model trained on this dataset will be able to learn more features and observations from the additional images and is expected to perform better on unseen images. More specifically, this research has the following objectives; to suggest data augmentation using GANs by generating the synthetic images, to evaluate the performance of the model trained on datasets appended with synthetic images generated by GANs against the model trained on original datasets, and to evaluate both models' training time and speed of convergence. A pre-trained VGG19 model will be used for transfer learning and the model's accuracy will be compared against the same model which is trained on the unmodified dataset and the dataset augmented using another traditional data augmentation technique. The paper will provide proof that the proposed method can produce a model that performs better on unseen car damage images.

II. Literature Review

Traditionally, the approach for image classification using machine learning is based on a two-layered system where the first layer was composed of a feature extractor while the second layer consisted of a classifier. In this case, the feature extractor is manually designed. One example of such a feature extractor is the Scale Invariant Feature Transform (SIFT) [13], which the authors claimed to be able to extract features that are invariant to image scaling, rotation, translation, and illumination. The classification layer can be any or a combination of machine learning classification methods such as Support Vector Machine, K-Nearest Neighbor, or Neural Networks.

A deep learning approach is also a two-layered system consisting of a feature extractor and a classification layer. However, the feature extractor in this case is a deep neural network composed of many convolutional layers. Compared to traditional feature extractors, deep neural networks can automatically learn the feature to be extracted. Currently, there are several deep neural network models that have been pre-trained and are readily available to be used for a wide range of applications including medical image classification [6], object detection and recognition [7], and semantic segmentation [14].

A. *Pre-trained CNNs for Car Damage Detection*

Kyu and Woraratpanya used pre-trained VGG16 and VGG19 models for car damage detection [15]. They used transfer learning and L2-regularization to reduce the training time and overfitting problems. The experiment used images that were downloaded using the Selenium tool [16] and were labeled manually as damaged and undamaged. Despite reporting a good result, the authors conceded that they still face the overfitting problem and in future work, they cite the need for more data to overcome the problem of overfitting.

Dwivendi et. al., use transfer learning techniques on pre-trained AlexNet, VGG19, Inception V3, MobileNet, and ResNet50 models [17]. The paper reported the best performance of 94.9% and 93.2% accuracy using the VGG19 model with and without data augmentation, respectively. The authors downloaded data manually from the Web and categorized the images into seven types of damages. The authors applied data augmentation techniques by applying rotation and horizontal flip transforms. The authors cite the need for more data to improve the study.

Patil et al. employed AlexNet, VGG-19, VGG-16, and ResNet models to achieve similar goals [18]. The authors compared the result of training them using an autoencoder with the transfer learning methods. The dataset was downloaded from the web and image annotation was done manually. In the end, they found that transfer learning combined with ensemble learning was the method of choice. Interestingly, the authors found that data augmentation marginally reduces the accuracy of the classification results. The paper reported a maximum accuracy of 87.92% and 88.24%, with and without data augmentation using ResNet. The data augmentation was done using a horizontal flip and appending it with random rotations between -20 to 20 degrees.

In a more recent paper [19], the authors performed car damage localization on a dataset containing 4,343 front-view car images from traffic scenes. Each image contains the bounding boxes of different car parts e.g., left and right lights, mirrors, license plates, bumper, and grill. And used these bounding boxes to manually crop the images into several parts containing the different car parts. A localization accuracy of 87% with false positive and false negative rates of 0.63% and 0.97%, respectively. The authors also mentioned that their dataset can be used for vision-based applications and for identifying car makes and models.

Pachón-Suescún et al. used a transfer learning technique on a pre-trained AlexNet model to detect scratches on car bodies [20]. The authors used 862 images without scratches and 424 images with scratches and augmented the dataset using the tool developed in [21] to generalize the model. The paper reported a validation accuracy of 88.29% can be attained. The authors also mentioned an issue with some false positive cases where surface dirt is categorized as scratches.

In related work, Murillo et al. found that detection accuracy is affected differently depending on the type and extent of the image augmentation process used to increase the dataset size [21]. The paper reported the maximum accuracy for image recognition is achieved if the dataset is augmented using transformation on original images where the rotation is 30 degrees and background color exists. However, if the rotation is limited to only 10 degrees, the accuracy is reduced by 20%. This is because the increase in rotation allows the network to learn the object in a different orientation. The study thus echoes the main argument for data augmentation via transformation which states that image transformation is an important factor to enrich a database by adding more diversity where the images exist with different backgrounds and at different orientations.

Bandi et al. used pre-trained VGG16 and transfer learning methods to achieve 87.9% and 68% accuracy in damage identification and location identification tasks, respectively [22]. The authors argued that low accuracy is the result of the lack of available images for training in the dataset which was from the Stanford car image dataset supplemented with additional images from scrapping the web.

B. *Image Augmentation using GANs*

From the study and analysis of the above research work, it is clear that the most common problem faced by several authors is the unavailability of proper training datasets and the importance of data image transformation and augmentation.

Most authors were able to compensate for the lack of datasets by image transformation. In other studies, which are not dealing with car damage detection and classification, the problem of imbalanced datasets is tackled using Generative Adversarial Networks (GANs). This approach is very common and is being utilized extensively in many areas such as medical imaging for example, but there is only very limited study done for car damage detection and classification using car damage datasets.

The authors in [23] have demonstrated that with the help of conditional GANs image-to-image translation can be done. The networks learn the mapping from the input image to the output image and also learn the loss functions which makes it an effective approach to synthesizing images from label maps, reconstructing objects from edge maps, and colorizing images.

Wei et al. proposed the use of Multi-Angled Generative Adversarial Networks (MAGAN) which generates realistic images of the same objects from different view angles [24]. These images help in mitigating the issue of imbalanced datasets. Grocery product images were used in the study whereby after image augmentation there was a significant improvement in training accuracy. The author has compared the results of using MAGAN against various other state-of-the-art GANs including Deep Convolutional GANs, Auxiliary Classifier GANs, and Wasserstein GANs.

Frid-Adar et al. applied GANs to generate synthetic images to improve the performance of CNN in medical image classification [25]. The author has used GAN for synthesizing high-quality liver lesion regions of interest. The paper reported that after the GAN augmentation, the CNN model can achieve a classification sensitivity of 85.7% and specificity of 92.4% up from 78.6% sensitivity and 88.4% specificity using a traditional image augmentation. The authors argued that the greatest challenge in the medical imaging domain is the limited number of datasets and synthetic images created using GANs have the potential to help solve the problem.

Arruda et. al. [26] used GANs to resolve the issue of lighting effects in car detection by generating a nighttime image of a car from a daytime image of the same car, and vice versa. The authors reported a consistent improvement of 10% when the synthetic images were introduced into the dataset. The authors used CycleGan where the model is trained in an unsupervised way to generate the CNN-based model responsible for the day-to-night translation.

III. Material and Method

A. Material

The dataset used in this study consists of 182 images of damaged cars, 80 of which were taken from Kaggle [12] with the remaining scrapped from the Web. The images have annotations that are stored in COCO format. The images have different sizes and show cars at different scales or magnification levels, as illustrated in Figure 1.

Fig. 1. Two examples of damaged car images in the dataset.

The dataset has been annotated in two ways. The first is the location of the damage in the image and the second is the severity of the damage. The damage location is categorized into one of nine car parts, namely Front Bumper, Left Head Lamp, Right Head Lamp, Hood, Fender, Left Front Door, Left Rear Door, Right Front Door, and Right Rear Door. The damage severity is categorized into one of three classes, namely Low, Medium, and High. The number of images that belong to each category is depicted as bar charts in Figures 2 and 3.

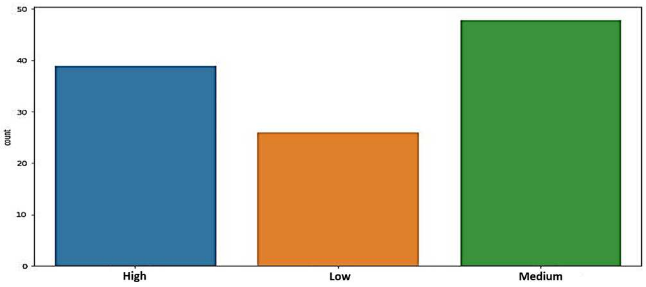

Fig. 2. Image counts per damage category.

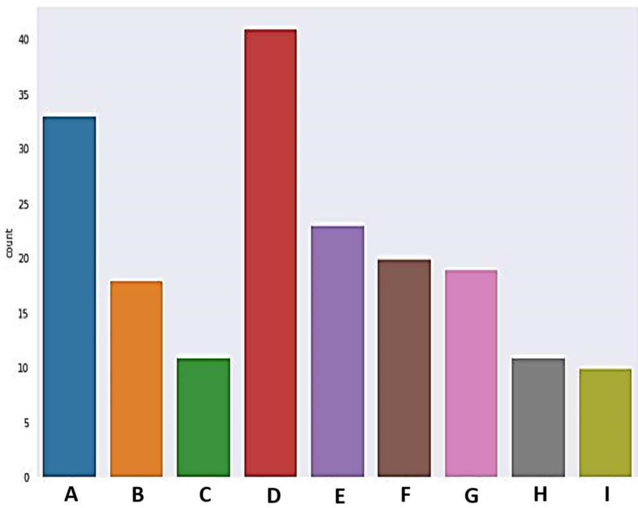

Fig. 3. Image counts per damage location; Front Bumper (A), Left Head Lamp (B), Right Head Lamp (C), Hood (D), Fender (E), Left Front Door (F), Left Rear Door (G), Right Front Door (H), and Right Rear Door (I).

B. Method

As evident from our review of the literature and observation of several existing datasets [12], many car damage detection models suffer from the problem of using insufficient and imbalanced datasets. Therefore, data augmentation is needed to make the trained models to be better generalized.

The network needs to learn different features in the images in different orientations and conditions. We will consider two data augmentation approaches namely the classical method through transformation and image synthesis by using GANs.

Our proposed research method is broadly divided into two parts as illustrated in Figure 4. In the first part, data augmentation will be applied using GANs which will generate the synthetic images and supplement them with the existing datasets. In the second part, transfer learning techniques will be employed on a pre-trained VGG19 model. The pre-trained VGG19 model was trained on a subset of the ImageNet dataset consisting of over a million images having 1000 different labels [27]. The advantage of using this transfer learning technique is that there is no need to train the model from scratch requiring millions of training images which would take a very long time to train with mid-range CPU and GPU.

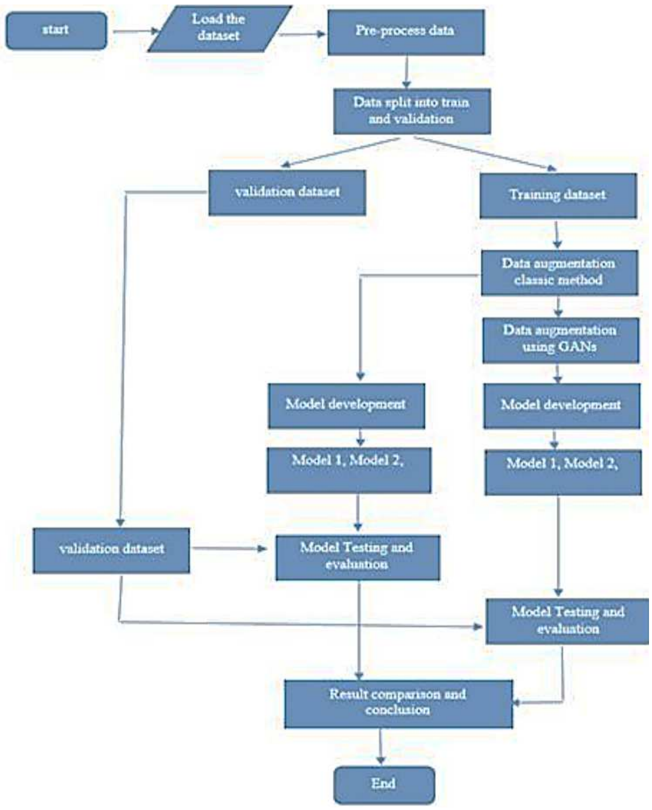

Fig. 4. A flow chart illustrating the proposed research methodology.

Transfer learning is done by replacing the classification layer with a new classification layer that has randomly initialized weights. This layer will be retargeted to perform car damage location and severity classifications. Two models will be trained, the first will detect the damage location and the second model will identify the severity of the damage. The two models will be trained using different labeled images based on damaged parts of the cars. The activation function in the first and last classification layers will be the Rectified Linear Activation (ReLU) function and the Sigmoid function, respectively. The learning rate will be subject to experiments at a later stage. For the loss function, categorical cross-entropy will be used.

When using the standard data augmentation method, the datasets will be appended by using random rotations between -20 to 20 degrees and horizontal flip transformations as was done in [18]. The images will also be subject to resizing where the image width will be changed. For some of the images cropping and padding will be done randomly and will be appended to the existing datasets. Also, color augmentation such as changes in brightness, contrast, saturation, and hue will be done on the randomly chosen images.

One of the disadvantages of using the classical method of image augmentation is that there is no new data being produced. The model has already observed the images but in a different rotation, and the model has already learned from the same images but in a different state. GANs are a technique where the data is augmented using new synthesized images. GAN is a relatively recent invention to generate artificial data from random noise as input. In their seminal paper, [28] the authors proposed a new framework for estimating a generative model via an adversarial process. There are two models which are trained simultaneously through the backpropagation weights and biases. The first one is a generative model, called the generator, that tries to produce fake images and the other is a discriminatory model, called the discriminator, that tries to detect fake images. At its core, the two neural networks compete with each other to produce more realistic synthetic images. Both models try to improve themselves until the synthetic images are no more distinguishable from the real images. The discriminator learns what best differentiates the synthetic images from the real images, and relays the information to the generator. The training process for both the generator and discriminator happens at the same time and training weights are updated/adjusted together. For successful GANs training, the generator should ultimately be able to fool the discriminator by producing realistic-looking synthetic images where the discriminator is not able to differentiate from the real images.

This approach has been attempted in different applications and contexts. For example, in [29] Fang et. al. use Deep Convolutional Generative Neural Networks (DCGAN) architecture [30] to generate images that improved a CNN-based classifier's accuracy. In our work, we also opted to use DCGAN for data augmentation. This is an improved version of GANs because as reported in [29], it can give a strong modeling performance and is more stable across various architectures, and helps in generating high quality images.

IV. Experiment and Result Analysis

The proposed methodology is implemented using the following software/hardware setup:

- Jupyter Notebook, Python, and Google Colab.
- Pytorch, Keras, Tensorflow, Numpy, Matplotlib, Seaborn, Sklearn, CV2, OS, and Random.
- NVIDIA P5000 GPUs with 500GB Hard Disk space

To provide a better understanding of the sort of images used in our experiment, we show in Figure 5 some examples of the images used for the training of the VGG19 and DCGAN models.

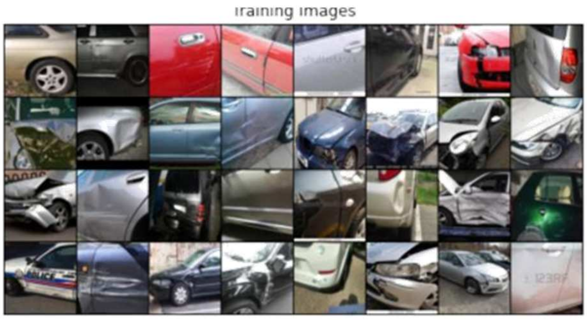

Fig. 5. Examples of car damage images used for training from the dataset.

Evaluation of any model is an important step that measures the model's quality. In the paper, the accuracy metrics will be used to measure the performance of each model. More specifically, the following data will be presented and analyzed:

- Training and validation accuracy using no data augmentation.
- Training and validation accuracy using standard data augmentation.
- Training and validation accuracy using DCGAN data augmentation.

The training log consisting of the training and validation accuracy for the best model after hyperparameter optimization is recorded. Figures 6, 7, and 8 show the training and validation accuracy of the damage location and damage severity classifiers. Figure 6 shows the training log when no data augmentation is performed, Figure 7 when standard data augmentation is performed, and lastly Figure 8 shows the training log when DCGAN is used for data augmentation. We define an accurate aggregate prediction as the prediction in which both car damage location and car damage severity are predicted accurately.

The comparison of the accuracy of the models between the models with data augmentation and without data augmentation is summarized in Table I. The result shows that the validation accuracy of the model when no data augmentation is performed is very low (44.1%). The result shows that both augmentation techniques increase the classifier performance by a significant margin. We observe some improvements when comparing the results of the two types of data augmentation, from 84.6% when using the standard data augmentation to 85.3% when using DCGAN's augmentation.

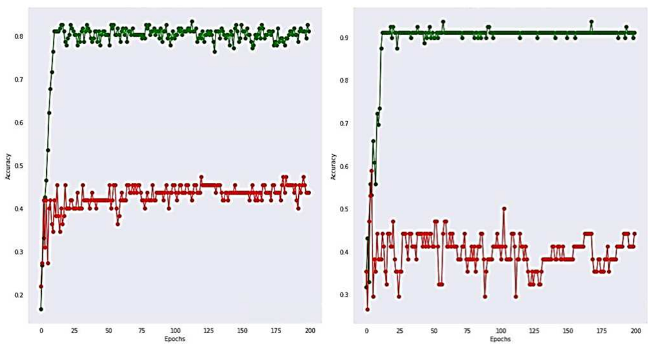

Fig. 6. Training accuracy (green) and validation accuracy (red) of damage location (left) and damage severity (right) classifiers without image augmentation.

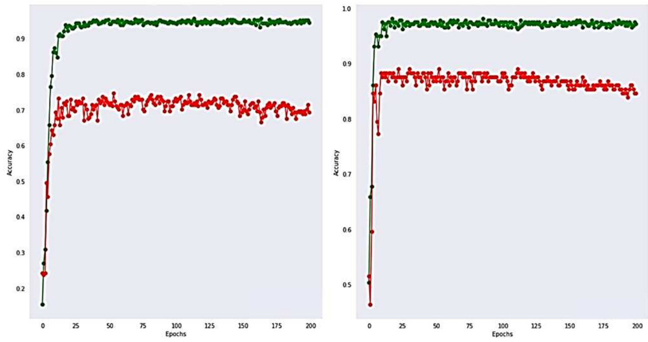

Fig. 7. Training accuracy (green) and validation accuracy (red) of damage location (left) and damage severity (right) classifiers with standard image augmentation.

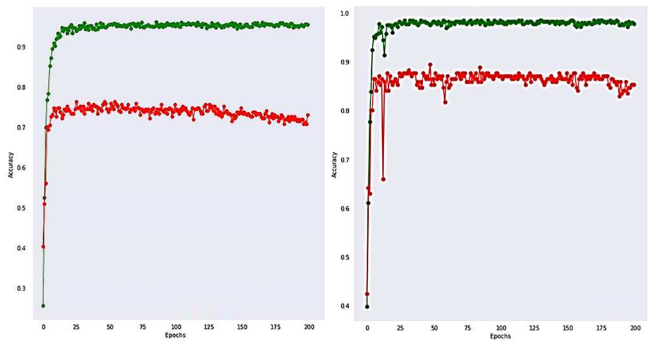

Fig. 8. Training accuracy (green) and validation accuracy (red) of damage location (left) and damage severity (right) classifiers with DCGAN image augmentation.

TABLE I. AGGREGATE PERFORMANCE OF MODELS

Dataset	Training Accuracy	Validation Accuracy
Datasets with no Augmentation	91.1	44.1
Dataset with Standard Augmentation	97.2	84.6
Dataset with DCGAN Augmentation	97.7	85.3

The analysis of the results suggests that after all the avenues of data augmentation have been exhausted, both standard data augmentation and generative data augmentation methods can increase classification performance and improve overall accuracy. Both augmentation methods inflate the dataset size artificially, but the standard augmentation technique can become counterproductive when the augmented images are not as per the requirements. On the other hand, generative data augmentation re-arranges the pixels just enough so that the classification model can learn the new features. When the datasets appended with augmented images are used for the VGG19 network, it produces better results than the original datasets because the model learns the new features from the augmented images as these images do not have the same structure.

V. CONCLUSION

We have explored the potential of extending a car damage dataset by using synthetic images generated using a DCGAN to improve the performance of an image classifier. We compare the classification performance of a pre-trained VGG19 on the original dataset, the dataset augmented using standard data augmentation technique, and the dataset augmented using DCGAN. It was observed that the classifier accuracy was improved by image augmentation with the generative augmentation method producing the best performance. The major challenge in our research was in training the GAN architecture because it is computationally intensive and needs a large number of training images. The model is quite sensitive to the hyperparameters values hence a slight change in their values could lead to mode collapse. Having said that, we conclude that the results we obtained support the primary objective of our research that the proposed method can be employed to generate synthetic images which can improve the accuracy and performance of a classifier.

REFERENCES

[1] Team Acko, "Common Problems Faced By Vehicle Insurance Claims Team," 2021. [Online]. Available: https://www.acko.com/articles/general-info/5-common-problems-faced-by-vehicle-insurance-claims-team. [Accessed: 23-Oct-2022].

[2] S. Sudirman, F. Natalia, A. Sophian, and A. Ashraf, "Pulsed Eddy Current signal processing using wavelet scattering and Gaussian process regression for fast and accurate ferromagnetic material thickness measurement," *Alexandria Eng. J.*, vol. 61, no. 12, pp. 11239–11250, 2022.

[3] C. Song, S. Sudirman, M. Merabti, P. Xiao, S. Sudirman, and M. Merabti, "Region adaptive digital image watermarking system using DWT-SVD algorithm," in *Proceeding of NASA/ESA Conference on Adaptive Hardware and Systems*, 2014, pp. 196–201.

[4] S. Monica, F. Natalia, and S. Sudirman, "Clustering Tourism Object in Bali Province Using K-Means and X-Means Clustering Algorithm," in *2018 IEEE 20th International Conference on High Performance Computing and Communications; IEEE 16th International Conference on Smart City; IEEE 4th International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*, 2018, pp. 1462–1467.

[5] F. Natalia, H. Meidia, N. Afriliana, J. C. Young, R. E. Yunus, M. Al-Jumaily, A. Al-Kafri, and S. Sudirman, "Automated measurement of anteroposterior diameter and foraminal widths in MRI images for lumbar spinal stenosis diagnosis," *PLoS One*, vol. 15, no. 11, pp. 1–27, 2020.

[6] F. Natalia, J. C. Young, N. Afriliana, H. Meidia, R. E. Yunus, and S. Sudirman, "Automated selection of mid-height intervertebral disc slice in traverse lumbar spine MRI using a combination of deep learning feature and machine learning classifier," *PLoS One*, vol. 17, no. 1, p. e0261659, 2022.

[7] L. Liu, R. Wang, C. Xie, P. Yang, F. Wang, S. Sudirman, and W. Liu, "PestNet: An End-to-End Deep Learning Approach for Large-Scale Multi-Class Pest Detection and Classification," *IEEE Access*, vol. 7, pp. 45301–45312, 2019.

[8] L. Liu, C. Xie, R. Wang, P. Yang, S. Sudirman, J. Zhang, R. Li, and F. Wang, "Deep Learning based Automatic Multi-Class Wild Pest Monitoring Approach using Hybrid Global and Local Activated Features," *IEEE Trans. Ind. Informatics*, p. 1, 2020.

[9] R. Cui, H. Liu, and C. Zhang, "A deep neural framework for continuous sign language recognition by iterative training," *IEEE Trans. Multimed.*, vol. 21, no. 7, pp. 1880–1891, 2019.

[10] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, 2009, pp. 248–255.

[11] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in *International Conference on Learning Representations*, 2015.

[12] L. P. Lenka, "Coco Car Damage Detection Dataset," *Kaggle*, 2020. [Online]. Available: https://www.kaggle.com/datasets/lplenka/coco-car-damage-detection-dataset. [Accessed: 23-Oct-2022].

[13] D. G. Lowe, "Object recognition from local scale-invariant features," in *Proceedings of the seventh IEEE international conference on computer vision*, 1999, vol. 2, pp. 1150–1157.

[14] A. S. Al-Kafri, S. Sudirman, A. Hussain, D. Al-Jumeily, F. Natalia, H. Meidia, N. Afriliana, W. Al-Rashdan, M. Bashtawi, and M. Al-Jumaily, "Boundary Delineation of MRI Images for Lumbar Spinal Stenosis Detection Through Semantic Segmentation Using Deep Neural Networks," *IEEE Access*, vol. 7, pp. 43487–43501, 2019.

[15] P. M. Kyu and K. Woraratpanya, "Car damage detection and classification," in *Proceedings of the 11th international conference on advances in information technology*, 2020, pp. 1–6.

[16] Selenium, "Selenium automates browsers. That's it!," *Software Freedom Conservancy*, 2022. [Online]. Available: https://www.selenium.dev/. [Accessed: 26-Oct-2022].

[17] M. Dwivedi, H. S. Malik, S. N. Omkar, E. B. Monis, B. Khanna, S. R. Samal, A. Tiwari, and A. Rathi, "Deep learning-based car damage classification and detection," in *Advances in Artificial Intelligence and Data Engineering*, Springer, 2021, pp. 207–221.

[18] K. Patil, M. Kulkarni, A. Sriraman, and S. Karande, "Deep learning based car damage classification," in *2017 16th IEEE international conference on machine learning and applications (ICMLA)*, 2017, pp. 50–54.

[19] P. Rostami, A. Taimori, S. Sattarzadeh, H. O. Shahreza, and F. Marvasti, "An Image Dataset of Vehicles Front Views and Parts for Vehicle Detection, Localization and Alignment Applications," in *2020 10th International Symposium on Telecommunications (IST)*, 2020, pp. 25–30.

[20] C. G. Pachón-Suescún, P. C. U. Murillo, and R. Jimenez-Moreno, "Scratch Detection in Cars Using a Convolutional Neural Network by Means of Transfer Learning," *Int. J. Appl. Eng. Res.*, vol. 13, no. 16, pp. 12976–12982, 2018.

[21] P. C. U. Murillo, J. O. P. Arenas, and R. J. Moreno, "Implementation of a data augmentation algorithm validated by means of the accuracy of a convolutional neural network," *J. Eng. Appl. Sci.*, vol. 12, no. 20, pp. 5323–5331, 2017.

[22] H. Bandi, S. Joshi, S. Bhagat, and A. Deshpande, "Assessing car damage with convolutional neural networks," in *2021 International Conference on Communication information and Computing Technology (ICCICT)*, 2021, pp. 1–5.

[23] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1125–1134.

[24] X.-S. Wei, Q. Cui, L. Yang, P. Wang, and L. Liu, "RPC: A large-scale retail product checkout dataset," *arXiv Prepr. arXiv1901.07249*, 2019.

[25] M. Frid-Adar, I. Diamant, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan, "GAN-based synthetic medical image augmentation for increased CNN performance in liver lesion classification," *Neurocomputing*, vol. 321, pp. 321–331, 2018.

[26] V. F. Arruda, T. M. Paixão, R. F. Berriel, A. F. De Souza, C. Badue, N. Sebe, and T. Oliveira-Santos, "Cross-domain car detection using unsupervised image-to-image translation: From day to night," in *2019 International Joint Conference on Neural Networks (IJCNN)*, 2019, pp. 1–8.

[27] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems*, 2012, pp. 1097–1105.

[28] J. Goodfellow Ian, P.-A. Jean, M. Mehdi, X. Bing, W.-F. David, O. Sherjil, and C. Courville Aaron, "Generative adversarial nets," in *Proceedings of the 27th international conference on neural information processing systems*, 2014, vol. 2, pp. 2672–2680.

[29] W. Fang, F. Zhang, V. S. Sheng, and Y. Ding, "A method for improving CNN-based image recognition using DCGAN," *Comput. Mater. Contin.*, vol. 57, no. 1, pp. 167–178, 2018.

[30] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *4th Int. Conf. Learn. Represent. ICLR 2016 - Conf. Track Proc.*, pp. 1–16, 2016.