Safety Performance of Autonomous Driving Systems Based on Artificial Intelligence

Fan Chen

School of Information Science and Engineering / School of Artificial Intelligence, Wuhan University of Science and Technology, Wuhan, China
202204416419@wust.edu.cn

Abstract. Safety incidents involving autonomous driving are on the rise with the increasing integration of artificial intelligence (AI) in this sector. Consequently, research on the safety performance of AI-based autonomous driving systems is gaining significance. This study delves into relevant security incidents, scrutinizes autonomous driving AI systems from the angles of adversarial attacks and dataset balance, and proposes a security performance assessment platform for adversarial attack algorithms to enhance the safety performance of autonomous driving AI systems. The analysis of safety accident cases elucidates various factors contributing to accidents, including erroneous decision-making due to misinterpretation of visual cues by AI systems, adversarial attacks introducing undetectable noise to input data, and system misidentification of objects. Proposed methods for improving safety performance encompass strategies such as proactive detection of data anomalies, development of robust data distribution frameworks, and implementation of defense mechanisms against adversarial attacks. Additionally, the study underscores the necessity for comprehensive evaluation platforms to assess the safety performance of AI systems in autonomous driving thoroughly. By addressing these issues, advancements in AI technology can be harnessed to ensure safer autonomous driving experiences, thereby mitigating risks and enhancing overall transportation safety. As AI technology continues to evolve, addressing safety concerns will remain paramount to realizing the full potential of autonomous driving for societal benefit.

Keywords: Artificial Intelligence, Autonomous Driving, Safety Performance, Adversarial Algorithms, Dataset Balance.

1 Introduction

Artificial intelligence has become ingrained in every aspect of life, including autonomous driving, medical diagnosis, industrial production, financial services, scientific research, and other fields, thanks to the quick development of machine learning and deep learning-based artificial intelligence technology. Among these, the application value of artificial intelligence in the field of autonomous driving has grown in prominence because of variables like traffic jams, frequent traffic accidents,

Y. Yue (ed.), *Proceedings of the 2024 International Conference on Mechanics, Electronics Engineering and Automation (ICMEEA 2024)*, Advances in Engineering Research 240,
https://doi.org/10.2991/978-94-6463-518-8_30

and human driving behavior brought on by the constant increase in traffic flow. Nevertheless, there are still issues with safety performance in autonomous driving systems that use artificial intelligence. These issues include unstable input and output, unsatisfactory model fitting, dynamic changes in the artificial intelligence technology environment, and the inability to ensure data privacy [1]. But as of late, the significance of artificial intelligence, since technology is now widely used and accepted worldwide, it is critical to find a solution to the artificial intelligence safety performance issue in the context of autonomous driving. This article primarily provides an overview of techniques for enhancing the safety performance of AI-powered autonomous vehicles by reviewing prior research and examining pertinent examples of AI-powered vehicles' safety performance.

2 Autonomous Driving Safety Accident Cases Caused by Artificial Intelligence

In spite of the fact that artificial intelligence has been included into the field of autonomous driving, a significant number of accidents using autonomous vehicles are still caused by AI. To offer just one example, red light was ran by a Tesla that was driving itself in December 2019, causing a collision that resulted in the instantaneous death of two passengers [1]. This tragedy took occurred in the state of California, in the United States of America. It was a pedestrian who was struck by a Tesla in Tokyo, Japan, in April of 2020, when the vehicle was utilizing its automatic driving assistance system [2]. The incident occurred in Tokyo. As a direct consequence of the accident, the pedestrian passed away completely and instantly. During the collision that took place on a highway in Taiwan in June of 2020, the automated driving system was mistaken for a cloud while it was attempting to identify the white truck that was in front of the car [1]. During the year 2022, more than seven hundred and fifty Tesla owners complained that their vehicles had unexpectedly slammed on the brakes or immediately stopped while they were on the road. The unstable input and output, inadequate model fitting, the inability to guarantee data privacy, and other issues related to the artificial intelligence technology environment in the field of autonomous driving are the primary causes of the numerous incidents of autonomous driving safety accidents that occur. The accident described above is just one example of the many incidents that occur.

3 Analysis of Autonomous Driving Safety Accident Cases

Regarding the 2019 incident in California, USA, where a Tesla ran a red light while operating an autonomous vehicle and caused an accident, the explanation was that the artificial intelligence system for autonomous driving mistook the red light for a green one. The autonomous driving visual model predicts distinct identification outcomes with noticeably different balances when the artificial intelligence autonomous driving system detects different hues. The autonomous driving artificial intelligence system's performance and safety are greatly impacted by the data set's balance [1]. This is because the system's ability to effectively detect and make decisions is dependent on the data set's balance. As a result, some research has discovered that the data set reduces the performance of a few data groups to balance the entire set after training because of the unequal proportions of various data groups in the data set. The artificial intelligence system for autonomous driving makes decisions based on the end outcome. The results vary from one another.

The 2020 Tesla incident in Tokyo, Japan involved an adversarial attack on the autonomous driving artificial intelligence system, which resulted in the vehicle striking and killing a pedestrian. It was specifically the victim of a black box attack in the real world. An adversarial attack occurs when a little amount of undetectable noise is introduced to the system's input data, leading to incorrect conclusions being drawn by the system. After adding noise, the sample is referred to as adversarial; the extra noise is known as adversarial perturbation. sex sample. Adversarial attacks include, for example, black box attacks in the real world. More precisely, this instance is vulnerable to a black box attack in the real world because of the high noise in the natural environment where the model structure and training data are unknown. Due to the adversarial attack, the system was unable to recognize speed limit road signs, which resulted in poor decision-making and accelerated driving, both of which contributed to the disaster. Relevant research from Beijing University of Aeronautics and Astronautics suggests that they developed an environmentally friendly physical assault technique for identifying road signs by creating adversarial stickers that resemble "graffiti" using the Generative Adversarial Network (GAN) and the mechanics of attention [3]. The artificial intelligence attack is predicated on effectiveness. As illustrated in Fig.1, "speed limit 20" is erroneously viewed as "speed limit 100" by the clever algorithm-based autonomous driving system in their interpretation.

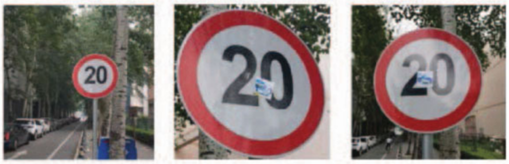

Fig. 1. Adversarial attacks schematic diagram on actual road signs [1].

Moreover, Zhejiang University of Science and Technology and Guangzhou University have conducted pertinent research and have put forth an algorithm for attacks that uses physical patches on QR codes [4], which creates feature maps with easily identifiable images of traffic signs using the road sign recognition system. This is used to identify the road sign's assault location for the QR code patch sticker, optimize the interference picture within the sticker, and create a camouflage example that works with the road sign recognition system, as illustrated in Fig. 2.

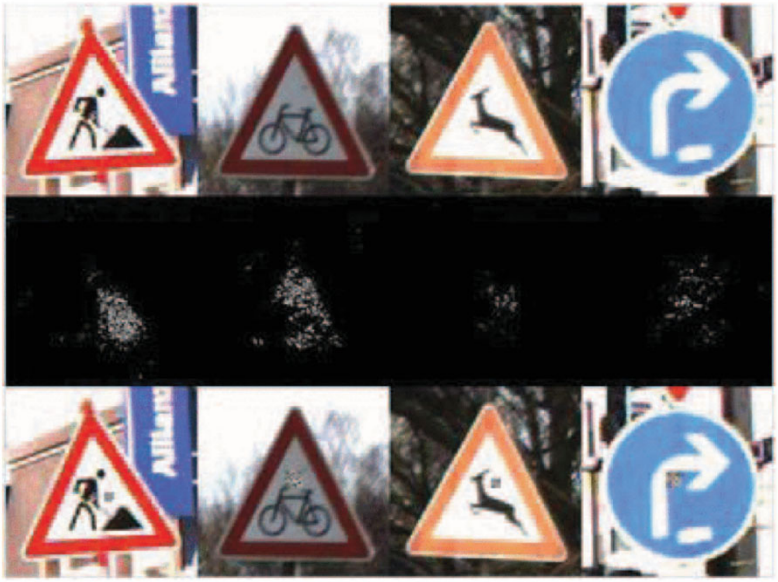

Fig. 2. Physical patch attacks against QR code-based traffic sign recognition [1].

A novel taint adversarial attack technique has also been presented by some academics for recognizing license plates and launching attacks against them [5]. This method uses an optimization algorithm to look for the best assault position on the license plate before smearing the characters. Subsequently, mud was used to replicate the organic stains on the license plate. The final study findings showed that the artificial stains had great hiding and denial capabilities and could successfully fool the neural network used to recognize license plates. The research also demonstrates how adversarial attacks cause systems to misidentify and produce incorrect decision outputs.

In June 2020, an automobile accident occurred on a Taiwanese highway when the autonomous driving system mistook the white truck ahead for a cloud. This error led to the system making incorrect decisions, which in turn caused the accident. An adversarial attack on the artificial intelligence autonomous driving system is another reason for this misidentification. It has specifically been the victim of a white-box attack. In reference to white-box attacks, pertinent literature suggests that the algorithm model of the artificial intelligence autonomous driving system may generate false prediction results and mistakenly identify pandas as gibbons by applying the Fast Gradient Symbol Method (FGSM) to add a little bit of noise to the picture, as illustrated in Fig. 3. System components are vulnerable to attacks on security and dependability [1].

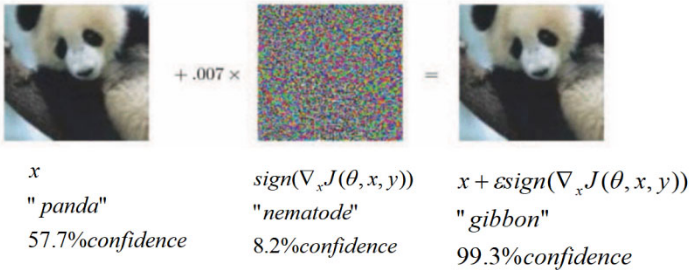

Fig. 3. An example of an adversarial assault diagram [1].

An additional piece of information concerning the DeepFool methodology was supplied by SM Moosavi-Dezfooli [6]. Through the application of this method, the L2 norm of the loss function is solved in order to determine the minimum amount of disturbance that is necessary in order to generate adversarial scenarios. In order to accomplish this, the loss function is taken into consideration. An additional piece of information concerning the Jacobian Saliency Map Attack (JSMA) strategy was supplied by N. Papernot [7]. For the purpose of delivering adversarial samples, this technique computes the derivatives throughout the forward propagation phase of the neural network. This is done in order to provide adversarial samples. In addition to this, the study team from Guangzhou University proposed the application of a technology known as adaptive gradient masking [8]. This was in addition to the one that came before it. Using this technique, sensitive sections of the image may be automatically detected, and the technology also has the capability to apply extremely modest perturbations to specific regions of the image. Su created a technique that is now known as one-pixel assault [9], which permits the adjustments of the results of the model classification as a consequence of a single pixel alteration. This method was published in the journal Computer Vision. Not only do these methods require a greater quantity of processing resources, but they also require a significantly more in-depth grasp of the structure of the model [1]. In addition, these methods require a significantly larger amount of processing resources. In every single one of them, there is not a single one that is not an algorithm consisting of artificial intelligence that is a white-box attack. Furthermore, these adversarial attack algorithms are able to demonstrate how an adversarial assault can result in inaccurate decision outputs and misidentifications in an artificial intelligence autonomous driving system. This is a significant accomplishment. One can say that this is a noteworthy achievement.

4 Suggestions or Methods for Improving the Safety Performance

The process of learning and training models is something that takes place during the design phase of autonomous navigation. When it comes to learning and training itself based on data sets, artificial intelligence systems often require a significant amount of data in order to be able to do so. It is vital to bear in mind that the system is particularly susceptible to attacks of all kinds throughout the whole of this procedure. These attacks can include poisoning, backdoors, and a variety of other types of attacks. This puts the data of the artificial intelligence system, which is responsible for giving power to autonomous vehicles, in peril of being vulnerable to model errors as well as abuses of privacy. This is because of the fact that the system is responsible for providing power to driverless vehicles. On the other hand, the distribution of data sets is skewed and uneven as a result of the limitations that are imposed on the environments in which data is collected. As a consequence of this, the likelihood that decisions may be made within the system that are not equitable to the parties involved is increased. The strategy of data poisoning is one that is utilized rather frequently in the context of attacks that entail the use of poisoning. The method of poisoning can be carried out in one of two ways: either by uploading the data directly to the model while acting as a user, or by placing the poisoned data on the internet in order for people to download it. These two approaches are both viable possibilities to consider. Afterwards, the model is trained by either including new data sets or modifying the ones that are already present in order to generate poisoning attacks. For the purpose of satisfying the requirements, this is carried out. The deployment of this form of attack is not only incredibly straightforward, but it also does not involve any adjustments to the infrastructure of the network. By leveraging training and retraining in order to accomplish this, it is possible to promptly manage the data collection that the organization is responsible for. Other important attack methods include the attack using a Trojan horse, the attack using quick poisoning, the attack using pure labels, the attack using clean labels, the attack using pruning aware, and so on. These are but a few of the many different options for assault that are accessible. There is a significant connection between the nature of the data collecting and the effectiveness and security of the artificial intelligence system that is utilized for autonomous driving. This connection is significant because it is a significant association. As a result of the fact that it is a substantial relationship, this relationship is of great significance. While the artificial intelligence system that is capable of autonomous driving is being created, it is of the utmost importance to conduct an evaluation of the quality of the data set that is being used by the system. This review should be carried out as soon as possible. It is absolutely necessary to carry out this review in order to achieve the ultimate goal of lowering the probability of poisoning attacks and biases in decision-making. In the course of carrying out this examination, it is of the utmost importance to pay particular attention to data abnormalities, data dispersion, data missing, and other issues of a similar nature. These are the kinds of concerns that require special attention for sure.

In order to discover a solution to the issue of data anomalies, research that is pertinent to this issue recommends a novel technique that makes use of an end-to-end data quality anomaly detection framework design. This is done in order to find a solution or solution to the problem. Utilizing predictive algorithms as a proactive strategy for discovering potential universal data quality issues of large data is made possible by this methodology, which makes it viable to utilize predictive algorithms. The reason for this is to ensure that the data is reliable and correct, and this is done.

Identifying quality anomalies that are connected with six quality qualities, namely accuracy, consistency, completeness, consistency, uniqueness, and readability, can be accomplished through the use of a technique known as anomaly detection. Furthermore, the utilization of predictive algorithms is made possible by this technique. As an additional point of interest, the "Quality Anomaly Score" is a brand new statistic that is given and developed through the utilization of this method. The significance of this score is referred to as a "quality anomaly score." The degree of abnormality and the low quality of the anomalies that were found can be determined by using this measure, which is relevant to each and every quality dimension as well as the entire dataset [10]. A number of distinct research investigations have proposed the All-to-All Comparison (ATAC) computer paradigm [11] as a potential solution for the distribution of data. This paradigm has been presented as a possible solution several times. This idea has been taken into consideration as a possible avenue for resolution. For those who are interested in improving the performance of computers, the application of this paradigm is an urgent must. MATLAB was utilized in order to develop the necessary model algorithms after the ATAC data distribution model had been designed. Following the initial design of the distribution model, this was carried out. This set of algorithms was constructed on top of tabu search, which served as the foundation for the algorithms. This study presents a solution framework for ATAC data distribution strategies that is based on tabu search. The framework is presented as potential solutions. The purpose of this is to set the stage for further exposition. The objective of this study was to investigate various methods of data distribution for ATAC computations, and it was carried out with that primary purpose in mind. To be more specific, the framework is made up of three separate components, each of which will be investigated in greater depth in the following paragraphs: A solution for data distribution that satisfies the load balancing and optimized storage policy is the responsibility of the storage optimization module, which is responsible for finding such a solution. This module is responsible for managing the data distribution process. Additionally, it is the responsibility of this module to find a solution on its own. 1) Managing interactions with users is the job of the driver module who is responsible for the system. 2) The load balancing module is the one that is responsible for coming up with a data distribution plan that is appropriate for the load balancing scenario so that it may be implemented. The storage optimization module is the one that is responsible for locating such a solution, which brings us to the third spot. This framework makes it feasible for the distributed system nodes that are participating in ATAC calculations to fulfill the requirements of load balancing and maintaining data integrity. This is made possible with the assistance of the data distribution strategy that is made possible by this framework. The framework is responsible for making this a reality. The architecture that is being explained here makes it possible to put into action the data distribution plan that has been developed. Furthermore, in contrast to the ATAC data distribution algorithm that is now being utilized, this framework data distribution solution has a higher processing performance and has the possibility to reduce the storage space capacity of the distributed system by between forty and fifty percent. This is a significant advantage. An important benefit is that this is the case. One of the most significant advantages is that this is the case. When it comes to dealing with missing data, there are specific circumstances in which the interpolation technique may prove to be a more effective strategy than other types of approaches that are comparable [12]. When it comes to making a classification of interpolation techniques, there are three basic categories that can be utilized. The statistical algorithms, the machine learning algorithms, and the deep learning algorithms are the categories that constitute these categories. The three categories are as follows, taking everything into consideration. The vast majority of interpolation algorithms produce the best results when used to numerical data sets; but, when applied to categorical data sets, they perform a very poor job with regard to the accuracy of their results. When the accuracy of both the interpolation method and the post-interpolation prediction task (which is based on the interpolation method) is higher, the capacity to forecast is raised. As the number of instances of missing data increases, the category of deep learning imputation methods displays an increased level of robustness. This is particularly true in circumstances in which there is a diverse assortment of data types, mechanisms that are absent, and post-imputation prediction tasks.

The discipline of autonomous driving makes substantial use of adversarial attack algorithms. This is due to the fact that these algorithms are vital to the accurate decision-making of artificial intelligence systems, which are utilized in autonomous driving applications that are utilized in the real world. The reason for this is that they are utilized in the field of autonomous driving, which is the reason for this. Therefore, as a result of this, they have a substantial influence on the increase of the safety performance of transportation systems that are able to operate independently and are powered by artificial intelligence. It is of the utmost importance that, over the course of the process of developing artificial intelligence systems, an evaluation of the system's levels of safety performance be carried out. This evaluation is an essential stage that must be finished in order to proceed. Because of this, a variety of different platforms have been developed with the purpose of evaluating the effectiveness and safety of artificial intelligence systems. These platforms have been established in order to accomplish this goal. FollBox is the name of the toolbox that is included in the ART tool bundle that IBM provides [13]. For further information, see the following. This is an illustration of the point. Additionally, there are extra scenarios to consider. Several different adversarial attack tactics have been incorporated into this toolkit in order to simplify the process of evaluating algorithms. This was done in order to maximize efficiency. Additionally, it is able to provide developers with suggestions for the creation of algorithms and defensive measures that are complementary to one another. This feature is a significant advantage. An important benefit is that this is the case. The DeepXplore platform, which was developed by Columbia University [14], provides a method that is popularly known as white-box testing. This method was developed by Columbia University. This technique is applied in the field of evaluating models that are based on deep learning. The application of this technology allows for the testing of deep learning models to be carried out successful. Furthermore, there is the DeepTest platform [15], which was developed by the University of Virginia. This platform will be discussed further below. By utilizing this platform, it is possible to do an automatic evaluation of the quality of artificial intelligence models that are used for autonomous driving. It is the platform that is providing this evaluation, and it is also the platform that is making it available. This evaluation is carried out with the purpose of verifying that these systems contain the required level of safety performance in order to successfully complete the task that is currently being performed. A good example of this would be the Deepsec adversarial algorithm testing platform, which was built jointly by Alibaba and the University of Illinois at Urbana-Champaign [16]. This platform was developed in collaboration with one another. Specifically, adversarial algorithms were exercised on this platform in order to evaluate their performance. This platform's principal assault method, which combines the employment of adversarial example technology, focuses primarily on original photographs as its major target. The original photographs are the major target of this attack strategy, which focuses exclusively on images. The results that were obtained provide the designers and developers of the system with direction on how to defend themselves against hostile assaults that are connected with the system. This is important because the system is associated with being vulnerable to attacks from adversaries. Furthermore, the findings will eventually lead to the development of associated algorithms, which will have the effect of improving the operational efficiency of the system's security. This will be implemented as a consequence of the findings. Additionally, research on adversarial attacks has been conducted at Tsinghua University, which is relevant to this topic matter. This study has been associated with this subject matter. In the past, investigations have been carried out on the aforementioned subjects. RealSafe is the name of the artificial intelligence security platform that was developed by the research team. RealSafe was developed by that team. The RealSafe software was created by that group. The very name of the company, RealSafe, is well-known. This platform is capable of performing a broad variety of other activities, such as the repair of algorithm models, in addition to the testing of adversarial assaults and the detection of vulnerabilities. Also included in this list is the capability to identify flaws. To add insult to injury, Guangzhou University and the Key Laboratory of Intelligent Product Testing and Reliability of the Ministry of Industry and Information Technology collaborated to build the F-lab, which is a neural network attack and defensive platform. This was done in order to add insult to injury. It was done in order to make matters worse than they already were. It was done with the intention of making things even more difficult than they already were. This action was taken with the purpose of making things even more challenging than they already were under the circumstances. To enhance the safety performance of autonomous driving artificial intelligence washing systems, the team developed an attack and defense module for the F-lab platform that is based on an artificial intelligence algorithm. This module was generated as part of their attempts to increase the safety performance of these systems. Within the context of the team's attempts to improve the safety performance of these systems, this module was built as part of those efforts. This action was carried out especially for the purpose of achieving the aforementioned objective of improving the outcome, which was the reason why it was carried out. The capabilities of this module include the ability to record voice, write text, and capture photographs, amongst a variety of other modes. This module integrates a number of various modes. In addition to including over twenty distinct adversarial sample production technologies, the platform in question has also developed fifteen distinct picture attacks and response tactics. This is a significant accomplishment. Additionally, the deployment of each of these technologies has already been finished within the organization. Through the use of the assault of neural networks and the platform of defense, it is possible to analyze a broad variety of well-known strategies for target identification, speech recognition methods, license plate recognition algorithms, and other approaches that are equivalent to these methodologies. Consequently, it is now possible to investigate a number of these different tactics. It is unfortunate that there is not yet a framework that is particularly thorough for the testing and assessment of artificial intelligence systems in the field of autonomous driving. This is a problem because autonomous driving is becoming increasingly popular. This is because hostile samples that are incorrect, data security detection methods that are inaccurate, and the complexity of building model security assessments are all aspects that contribute to the problem. This is the reason why this is the case. This is the reason why things are the way they are in the current situation. It will be determined that this is an essential topic that needs to be studied during the subsequent study that will be carried out on the safety performance of artificial intelligence systems that are capable of autonomous driving. This research will be carried out in the future.

5 Conclusion

To contribute to the safe application of artificial intelligence in the field of autonomous driving, we thoroughly analyze the safety concerns associated with AI-powered systems and provide a summary of various approaches or recommendations for enhancing system safety performance in this system review. Artificial intelligence technology is currently causing a global frenzy. The autonomous driving industry will see an increase in the demands placed on artificial intelligence technologies. In the future, there will be an increasing number due to technological advancements and the shifting of the times. Numerous safety concerns surface, and solutions to address and resolve these concerns will also. Additionally, autonomous driving will offer a stronger assurance for people's safety when traveling, reducing problems and losses.

References

1. X. Huang, L. Huang, G. Tong, X. Zhou, H. Xu and H. Shen, "Research on Safety Testing and Evaluation Technology of Artificial Intelligence System for Automatic Driving," 2023 International Conference on Mechatronics, IoT and Industrial Informatics (ICMIII), Melbourne, Australia, 2023, pp. 159-162
2. A. Liu, J. Wang and X. Liu, "Artificial intelligence security and evaluation", Artificial Intelligence, vol. 3, pp. 32-42, 2020.
3. A. Liu, X. Liu, J. Fan, Y. Ma and D. Tao, "Perceptual-sensitive GAN for generating adversarial patches", Proceedings of the AAAI Conference on Artificial Intelligence, pp. 1028-1035, July 2019.
4. Y. G. Qian, X. W. Liu, Z. Q. Gu, B. Wang, J. Pan and X. M. Zhang, "QR code based patch attacks in physical world", Journal of Cyber Security, vol. 5, pp. 75-86, December 2020.

5. Y. G. Qian, D. F. Ma, B. Wang, J. Pan, J. M. Wang, Z. Q. Gu, et al., "Spot evasion attacks: adversarial examples for license plate recognition systems with convolutional neural networks", Computers Security, vol. 95, August 2020.
6. SM. Moosavi-Dezfooli, A. Fawzi and P. Frossard, "DeepFool: a simple and accurate method to fool deep neural networks", arXiv preprint arXiv:1511.04599, 2016.
7. N. Papernot, P. Mcdaniel, S. Jha, M. Fredrikson and A. Swami, "The limitations of deep learning in adversarial settings", arXiv preprint arXiv:1511.07528, 2015.
8. Z. Q. Gu, W. X. Hu, C. J. Zhang, H. Lu and L. Wang, "Gradient shielding: towards understanding vulnerability of deep neural networks", IEEE Transactions on Network Science and Engineering, vol. 8, pp. 921-932, May 2020.
9. J. Su, D. V. Vargas and S. Kouichi, "One pixel attack for fooling deep neural networks", IEEE Transactions on Evolutionary Computation, vol. 23, pp. 828-841, October 2019
10. E. Widad, E. Saida and Y. Gahi, "Quality Anomaly Detection Using Predictive Techniques: An Extensive Big Data Quality Framework for Reliable Data Analysis," in IEEE Access, vol. 11, pp. 103306-103318, 2023,
11. D. Deng et al., "A Solution Framework for All-to-All Comparison Data Distribution Strategy Based on Tabu Search," 2021 IEEE 6th International Conference on Cloud Computing and Big Data Analytics (ICCCBDA), Chengdu, China, 2021, pp. 83-89,
12. X. Miao, Y. Wu, L. Chen, Y. Gao and J. Yin, "An Experimental Survey of Missing Data Imputation Algorithms," in IEEE Transactions on Knowledge and Data Engineering, vol. 35, no. 7, pp. 6630-6650, 1 July 2023,
13. J. Rauber, W. Brendel and M. Bethge, "Foolbox: A Python toolbox to benchmark the robustness of machine learning models", arXiv preprint arXiv:1707.04131, 2017.
14. K. X. Pei, Y. Z. Cao, J. F. Yang and S. Jana, "DeepXplore: Automated Whitebox Testing of Deep Learning Systems", arXiv preprint arXiv: 1705.06640, 2017.
15. Y. C. Tian, K. X. Pei, S Jana and B Ray, "DeepTest: Automated Testing of Deep-Neural-Network-driven Autonomous Cars", arXiv preprint arXiv: 1708.08559, 2017.
16. X. Ling, S. L. Ji, J. X. Zou, J. N. Wang, C. M. Wu, B. Li, et al., "Deepsec: A uniform platform for security analysis of deep learning model", 2019 IEEE Symposium on Security and Privacy(SP), pp. 673-690, May 2019.