ČESKÁ ZEMĚDĚLSKÁ UNIVERZITA V PRAZE

PROVOZNĚ EKONOMICKÁ FAKULTA

Data augmentation for image classification using GAN and autoencoder

disertační práce

Autor: Ing. Gofur Halmuratov

Školitel: doc. Ing. Arnošt Veselý, CSc.

ACKNOWLEDGEMENTS

I would like to extend my heartfelt gratitude to the people who have played a significant role in the completion of this thesis. Their unwavering support, guidance, and encouragement have been instrumental in my academic journey. First and foremost, I would like to dedicate this work to my late father, whose unwavering love, strength, and belief in me have been a constant source of inspiration. His wisdom, guidance, and unwavering support have shaped my character and fueled my determination. Though he is no longer with us, his memory continues to motivate me and drive me forward. I am forever grateful for his sacrifices and the values he instilled in me. I extend my deepest appreciation to my supervisor, doc. Ing. Arnošt Veselý, CSc., for their invaluable guidance, expertise, and unwavering support throughout this research endeavor. Their insightful feedback, patience, and encouragement have been instrumental in shaping this thesis. Their mentorship and dedication to my academic growth have truly been transformative. I would also like to express my heartfelt thanks to my mother for her unwavering support, understanding, and belief in me throughout this journey. Her love, encouragement, and sacrifices have provided me with the strength and motivation to overcome challenges and persevere. I am incredibly fortunate to have her by my side, and her unwavering faith in me has been a constant source of inspiration. To my wife, Roza, and children Makhliyo, Ulugbek, and Sevinch, thank you for your continued understanding, patience, and love. My own presence and unwavering support have brought balance to my life and served as a reminder of the importance of perseverance and determination. My unwavering belief in myself and my sacrifices have been the driving force behind the completion of this thesis.

Finally, I am grateful to all the people who have contributed to my academic and personal growth, including teachers, mentors, and colleagues. Oner guidance, knowledge, and collaborations have enriched my research experience and shaped my intellectual development. Completing this thesis would not have been possible without the support and contributions of these incredible individuals. I express my deepest appreciation to my late father, my supervisor, my mother, my wife, my children, and all those who have supported me throughout this journey. Their love, guidance, and encouragement have been indispensable, and I am forever grateful for their presence in my life.

ABSTRACT

V této disertační práci jsem zkoumal metody augmentace dat pro klasifikaci obrazů a to zejména pro takové případy, kdy jsou k dispozici jen velmi malé soubory trénovacích dat. Tato práce je založena na použití generativních adverziálních sítí (GAN), autoenkodérů a na použití shlukové analýzy latentního prostoru autoenkodérů. Tyto metodiky použité společně zvyšují přesnost klasifikace a dovolují získat uspokojivé výsledky i v případě, že máme k dispozici pouze velmi malé trénovací soubory. V počáteční fázi své práce jsem vyvinul nový přístup využívající generativní adverziální sítě (GAN) k simulaci pravděpodobnostních distribucí jednotlivých klasifikačních kategorií. Tento přístup znamenal významný odklon od konvenčních technik rozšiřování datové sady. Integrací GAN do procesu vývoje a použití klasifikátoru jsem efektivně využil jejich generativní schopnosti nad rámec pouhého generování dat. Moje komplexní experimentování s databází MNIST, zahrnující tréninková data v rozsahu od 1 do 100 vzorků na kategorii, potvrdilo účinnost tohoto nového přístupu. Výsledky mnou vyvinutých metod, zejména M1 a M3, ukázaly významná zlepšení oproti standardnímu použití GAN diskriminátoru pro rozšíření trénovacích dat a prokázaly jejich účinnost v situacích, kdy je k dispozici pouze velmi malý soubor trénovacích dat.

V rámci svého výzkumu jsem navrhl integraci analýzy latentního prostoru variačních autoenkodérů (VAE) do procesu trénování a klasifikace. Výsledkem byly metody M2 a M4. Integrace analýzy latentního prostoru byla klíčová a umožnila mi výrazně zvýšit kvalitu a spolehlivost dat generovaných těmito modely a tím také přesnost navržených klasifikátorů. Navržená metodika prošla přísným testováním na datové sadě MNIST.

V závěrečné části mé disertační práce jsem se zabýval problémem posouzení kvality pro augmentaci dat nově vytvořeného obrazu. Navrhl jsem pro tento účel novou metodiku pracující s latentním prostorem autoenkodérů. Tento přístup významně zefektivnil proces hodnocení kvality generovaného obrazu a eliminoval potřebu ručního přezkoumání. Při testování na souborech dat jako MNIST, CIFAR-10 a CIFAR-100 tato metodika prokázala svou účinnost.

Moje disertační práce poskytuje komplexní a inovativní přístup k rozšiřování dat pro klasifikaci obrazu v případě, že je dispozici pouze velmi malý soubor trénovacích dat. Navržené metody v tomto případě dosahují vyšší přesnosti klasifikace než jakou lze získat standardními metodami rozšiřování datového souboru. Výsledky mého empirického výzkumu potvrzují potenciál těchto

pokročilých technik rozšiřování dat při klasifikaci obrazů, zejména ve scénářích reálného světa, kde jsou trénovací data často vzácná.

Klíčová slova: Augmentace dat, klasifikace obrazu, hluboké učení, autokodéry, analýza latentního prostoru, generativní adversiální sítě (GAN), variační autoenkodér (VAE), shlukování v latentním prostoru, konvoluční neuronové sítě (CNN).

ABSTRACT

In this thesis, I investigated data augmentation methods for image classification, especially for cases where only very small training data sets are available. This work is based on the use of generative adversarial networks (GAN), autoencoders(AEs), and the analysis of the latent space of AEs. These methodologies used together increase the classification accuracy and allow us to obtain satisfactory results even if we only have very small training sets. I developed a new approach using GANs to simulate the probability distributions of individual classification classes. This approach marked a significant departure from conventional data set augmentation techniques. By integrating GANs into the process of developing and using a classifier, I effectively leveraged their generative capabilities beyond just generating data. Extensive experimentation with the MNIST database, involving training data ranging from 1 to 100 samples per category, confirmed the effectiveness of this new approach. The results of the developed methods, showed significant improvements over the standard use of a GAN discriminator to augment the training data and proved their effectiveness when only a very small set of training data is available.

As part of my research, I proposed the integration of VAE latent space analysis into the training and classification processes. The results were methods M2 and M4. The integration of latent space analysis was key and allowed me to significantly increase the quality and reliability of the data generated by these models and thus the accuracy of the proposed classifiers. The proposed methodology has undergone rigorous testing on the MNIST dataset.

Thesis also considers the problem of quality assessment for the data augmentation of a newly created image. The proposed methodology working with the latent space of AEs. This approach significantly streamline the process of evaluating the quality of the generated image and eliminated the need for manual review. Tested on datasets such as MNIST, CIFAR-10, and CIFAR-100, this methodology has proven its effectiveness.

The thesis provides an innovative data augmentation approach for image classification when only a very small set of training data is available. In this case, the proposed methods achieve a higher classification accuracy than can be obtained by standard data set expansion methods. The results of empirical research confirm the potential of these advanced data augmentation techniques in image classification, especially in real-world scenarios where training data is often scarce.

Keywords: Data Augmentation, Image Classification, Deep Learning(DL), Autoencoders(AEs), Latent Space Analysis, Generative Adversarial Networks (GAN), Variational Autoencoder (VAE), Latent Space Clustering, Convolutional Neural Networks (CNN).

Contents

1 Introduction

1.1 Motivation

The field of computer vision, an important component of artificial intelligence, has experienced a transformative evolution in recent years. This evolution is primarily attributed to groundbreaking advancements in DL technologies. Central to this transformation has been the development of CNNs (Goodfellow et al., 2023). CNNs have revolutionized critical computer vision tasks such as image classification, object detection (Patterson & Gibson, 2023), and semantic segmentation (Janet et al., 2023), marking a seminal moment in the field and redefining the capabilities of machines in processing and understanding visual information (Krizhevsky, Sutskever, & Hinton, 2012). The effectiveness of DL in computer vision is mainly due to the availability of large, diverse, and accurately annotated image datasets. However, the process of assembling and annotating these extensive datasets presents considerable challenges, particularly in scenarios with limited or sparse data. In specialized domains, there is often a scarcity of datasets that are sufficiently comprehensive to train models effectively (Goodfellow, Bengio, & Courville, 2023). A significant challenge within the realm of computer vision is the scarcity of labeled data. This issue slows down the improvement of DL models' performance and their ability to generalize to novel, unseen data. To address this limitation, image data augmentation techniques are increasingly utilized in such cases. These techniques involve applying various transformations to existing labeled data to generate new training examples. This expansion of the training dataset enhances the model's ability to generalize across different real-world variations (Zhou et al., 2014).

Traditional data augmentation methods like random cropping, rotation, and flipping (Voulodimos et al., 2022) have been widely used to increase the diversity of training data. However, the advent of DL and generative models has led to the emergence of more sophisticated and realistic image data augmentation methods. Innovations such as GANs and AEs have shown significant promise in synthesizing images that closely resemble real-world samples. These generative models facilitate targeted augmentation, enabling the creation of synthetic data that capture specific variations or simulate rare instances, thus filling gaps in existing datasets (Radford, Metz, & Chintala, 2016).

The goal of the thesis is to explore and enhance contemporary image data augmentation methods, especially those utilizing generative models. In particular, it focuses on improving the ability of DL models in computer vision to perform well on new, unseen data. A crucial part of this research is evaluating the quality and effectiveness of the augmented images (Sarkar et al., 2021). Ensuring that augmented samples accurately reflect the underlying data distribution is essential for the reliability and efficacy of computer vision systems.

The aim of the thesis is to scrutinize the effectiveness of modern image data augmentation methods (Janet et al., 2023) and propose innovative techniques to evaluate their impact on model performance and generalization. This work seeks to help, particularly in scenarios when only limited labeled data is available.

1.2 Problem Statement and Research Questions

The field of computer vision has undergone rapid advancement in recent years, driven by the development of sophisticated DL models and the increasing availability of large-scale image datasets (Carranza et al., 2021; Mahajan & Jha, 2022). However, a key challenge remains: the scarcity of high-quality labeled data. The process of acquiring and annotating extensive datasets can be time-consuming, expensive, and often impractical, particularly for specialized domains or rare events (Liu et al., 2020). To address this issue, researchers have extensively explored data augmentation techniques. These techniques involve strategically modifying existing labeled data to generate additional, realistic training examples, thereby expanding the dataset and improving model performance (Shorten & Khoshgoftaar, 2019).

While traditional augmentation techniques like random cropping, rotation, and flipping have been extensively utilized, the advancements in DL and generative models have unveiled new opportunities for more sophisticated and realistic image data augmentation. Techniques such as GANs and AEs have shown potential in creating synthetic images that closely mimic real-world samples. The strategic use of these generative models enables targeted augmentation, producing synthetic data that captures specific variations or emulates rare instances (Kingma & Welling, 2014).

The driving force behind this research is to explore and enhance the effectiveness of contemporary image data augmentation methods, particularly those that employ generative models. The goal is to improve the performance and generalization capabilities of DL models in computer vision. By integrating advanced augmentation techniques, this research aims to address the challenge of data scarcity and enhance the robustness of computer vision systems in real-world applications (Ronneberger, Fischer, & Brox, 2015).

Several critical research questions emerge in pursuit of this goal:

Integration of Modern Data Augmentation Methods: How can contemporary image data augmentation methods, including generative models, be effectively integrated into the training pipeline of DL models? This research will explore the use of GANs and AEs to generate synthetic images and examine the impact of various augmentation strategies on model performance and generalization.

Evaluation of Augmented Image Quality: How can the quality and effectiveness of augmented images, especially those generated by generative models and analyzed through latent space, be evaluated, and quantified? This research aims to develop novel methodologies for assessing image quality, focusing on metrics that evaluate fidelity, diversity, and relevance.

Impact on Learned Representations: What is the effect of different image data augmentation techniques on the learned representations and feature spaces of DL models? The study will analyze how augmentation methods influence the interpretability, disentanglement, and discriminative capabilities of learned representations and explore visualization techniques to understand feature transformations through augmentation.

Addressing these research questions is expected to significantly contribute to the advancement of modern image data augmentation methods and their evaluation using generative models and latent space analysis. The outcomes of this research could have far-reaching implications for various computer vision tasks, such as object detection, image classification, and semantic segmentation, where data scarcity is a common challenge. Enhancing the understanding and effectiveness of image data augmentation aims to improve the performance and generalization capabilities of DL models in real-world applications (Long, Shelhamer, & Darrell, 2015).

The subsequent sections of this paper will present an extensive literature review on image data augmentation techniques, covering both traditional and modern approaches. This review will critically examine various strategies, discuss their benefits and limitations, and highlight recent developments in generative models for data augmentation. Additionally, the paper will detail experimental results and analyses to evaluate the effectiveness of different augmentation methods in enhancing model performance and generalization.

2 Objectives

The principal objective of this thesis is to conduct a comprehensive and systematic exploration in the domain of image data augmentation for image classification. This research is particularly focused on contexts where training data is scarce, a scenario increasingly prevalent in the field of ML and artificial intelligence(AI). The core of this endeavor lies in the innovative application of advanced computational models, specifically Generative Adversarial Networks (GANs), AEs, and the sophisticated exploration of latent space analysis. This study is designed to transcend the current limits of conventional data augmentation methods. It seeks to overcome two fundamental challenges: the limited availability of training data and the essential demand for high-quality data to train robust ML models. This is achieved by meticulously analyzing and manipulating the latent space, a conceptual realm where the intrinsic properties of data are encoded in a compressed form, offering a fertile ground for innovation in data augmentation techniques. A key aspect of this research is to bridge the gap between theoretical advances in DL and their practical, real-world applications. By leveraging the latent space of AEs and the generative capabilities of GANs, the thesis proposes to create a new paradigm in data augmentation that is both efficient and effective in scenarios with limited data availability. The goal is to enhance the accuracy, reliability, and generalizability of image classification models, thereby contributing significantly to the advancement of ML methodologies. Through rigorous experimentation and analysis, this thesis will not only validate the effectiveness of these advanced computational models in enhancing data augmentation but also contribute to the broader understanding of their capabilities and limitations. The research is expected to yield novel insights and methodologies that can be applied to a wide range of image classification tasks, particularly those hampered by the availability of limited training data.

2.1 Integrating Autoencoders and Latent Space Analysis

Another fundamental aim of this thesis is the integration of AEs and latent space analysis into the data augmentation process. This objective focuses on exploiting the unique capabilities of AEs, especially VAEs, for dimensional reduction and feature extraction, thereby presenting new avenues in data augmentation.

Methodology and Application: This thesis delves into the utilization of latent space analysis for assessing the quality of augmented images. Traditional image quality assessment techniques often fall short in capturing the complexities associated with images produced by advanced generative models. By implementing latent space analysis, new metrics and methodologies are introduced for a more comprehensive assessment of image quality. This method addresses both quantitative and qualitative aspects of image quality, ensuring that the augmented data is not only diverse but also of high fidelity and relevance to the intended task.

Innovative Strategies for Quality Improvement: This research involves training an autoencoder for all input training datasets and decoding them to obtain the latent space representation of the autoencoder. This representation allows for a novel method to quality determination, wherein images within a certain threshold in the latent space are classified as high-quality and included in the training dataset. This methodology underscores the importance of quality in data augmentation, focusing on improving not just the quantity but also the quality of training data.

2.2 Methodological Innovations and Empirical Research

The thesis is characterized by its methodological innovations, particularly the development of novel models, each offering a unique solution to leveraging the capabilities of GANs and AEs. These models represent a significant advancement in data augmentation techniques, aimed at enhancing classification accuracy in data-limited environments.

Testing and Evaluation: The models developed through this research are rigorously tested and evaluated, showcasing notable improvements in classification accuracy. The research provides empirical evidence of the effectiveness of these methods, particularly in environments with limited training data. This empirical validation is crucial in demonstrating the practical applicability and effectiveness of the proposed methods.

Real-World Application and Adaptability: The findings from these experiments are invaluable, offering concrete evidence of the advancements in classification accuracy and establishing a strong basis for further exploration and application in more complex, diverse real-world environments. The thesis not only reaffirms the potential of advanced data augmentation techniques in image

classification but also showcases the adaptability and robustness of these methods across various data scenarios.

2.3 Limitations

The proposed research considers several the following limitations:

Computational Complexity: The implementation of complex generative models like GANs and VAEs poses significant computational challenges. The thesis acknowledges the high computational demand focuses on developing feasible strategies that can be realistically adopted within the existing computational resources. This is crucial, as it ensures that the proposed methodologies are not just theoretically sound but also practically implementable.

Dataset Specificity: A key limitation lies in the specificity of the datasets for instance, such as MNIST. While these datasets are standard benchmarks in ML, their characteristics, size, and diversity may influence the generalizability of the research findings. The thesis critically examines how unique features of these datasets influence applicability of the proposed augmentation methods and whether these methods can be adapted to other, more complex datasets.

Generalization to Other Domains: The application of these findings to other domains or tasks has its limitations. The transferability of the proposed methods to different domains, each with its unique requirements and challenges, is not assured.

The objectives of this thesis are multi-faceted, each contributing significantly to the overarching goal of enhancing image classification through advanced data augmentation. By innovatively employing GANs, AEs, and latent space analysis, the thesis not only addresses the challenges posed by data scarcity but also propels the field of image classification into new realms of possibility and efficiency.

3 Literature Review

3.1 Overview of image augmentation techniques

The augmentation techniques artificially increase the diversity and quantity of training data by applying various transformations and distortions to the original images.

Traditional augmentation techniques include rotation, scaling, flipping, cropping, and translation (Girshick, 2015). Rotation involves rotating the image by a certain angle, while scaling alters the size of the image. Flipping, both horizontally and vertically, creates mirrored versions of the original image. Cropping focuses on selecting a specific region of interest within the image, and translation shifts the image within the frame. These techniques have been widely adopted due to their simplicity and effectiveness in increasing dataset diversity.

Intensity-based transformations modify the pixel values of the image, altering its appearance without changing the spatial structure. Techniques such as brightness adjustment, contrast enhancement, and saturation modification fall under this category (Szegedy et al., 2015). Brightness adjustment involves uniformly scaling the pixel values to change the overall brightness of the image. Contrast enhancement expands or compresses the range of pixel values to enhance the differences between light and dark regions (Patterson & Gibson, 2023). Saturation modification adjusts the color intensity in the image (Chollet, 2021). These techniques are commonly employed to introduce variability in the image data and improve model robustness (Charniak, 2023).

Geometric transformations modify the spatial configuration of the image by warping or distorting its shape. These techniques include affine transformations, perspective transformations, and elastic deformations (Dosovitskiy et al., 2020). Affine transformations preserve parallel lines and ratios, enabling operations such as rotation, scaling, and shearing. Elastic deformations simulate local deformations, mimicking real-world scenarios where objects undergo non-linear shape changes. Geometric transformations are effective in augmenting data and increasing model performance by introducing variations in object poses, perspectives, and spatial configurations.

The adoption of these image augmentation techniques, either individually or in combination, helps address the limitations imposed by limited labeled data. By increasing the diversity and quantity of training data, DL models become better equipped to learn robust features, generalize well to unseen samples, and exhibit improved performance in various computer vision tasks. Furthermore, the effectiveness of image augmentation techniques has been demonstrated in several studies,

where they have contributed to significant improvements in model accuracy and generalization (Wang et al. 2023 and Liu et al. 2022).

In conclusion, image augmentation techniques are indispensable in enhancing the performance and generalization capabilities of deep learning models (Shorten & van der Burgh, 2021). Traditional techniques, intensity-based transformations, and geometric transformations offer diverse ways to augment datasets and introduce variations, ultimately improving model robustness and performance (Cubuk et al., 2020; Zhang et al., 2017). By incorporating these augmentation techniques, researchers can mitigate the limitations imposed by limited labeled data, facilitate more effective training, and enable the development of more accurate and reliable computer vision systems.

3.2 Types of image distortions and transformations

Geometric Transformations Geometric transformations are fundamental techniques in augmenting image datasets by modifying the spatial configuration of the images. Rotation, scaling, shearing, and perspective transformations are widely utilized geometric transformations (He, Girshick, & Dollár, 2019). Rotation involves rotating the image by a specified angle, while scaling alters the size of the image. Shearing introduces a skew effect by shifting the image along one of its axes. Geometric transformations introduce variations in object poses, orientations, and spatial configurations, thereby enhancing the diversity of the dataset. Photometric distortions focus on altering the pixel values or color characteristics of the images. These distortions aim to replicate real-world scenarios where images are subject to changes in lighting conditions, camera settings, or atmospheric conditions. Techniques such as brightness adjustment, contrast modification, color saturation changes, and noise addition fall under photometric distortions (Tzeng et al., 2017). Brightness adjustment uniformly scales the pixel values, while contrast modification expands or compresses the range of pixel values to enhance or reduce the differences between light and dark regions. Color saturation changes affect the intensity of colors in the image. Noise addition introduces random variations in the pixel values, simulating imperfections in the image acquisition process. Photometric distortions introduce variations in lighting conditions and color characteristics, thereby enhancing the robustness and generalization capabilities of DL models. Elastic Deformations Elastic deformations simulate local deformations or shape changes in the images, contributing to increased dataset diversity. These deformations introduce non-linear

transformations to the images, mimicking real-world scenarios where objects undergo deformations due to stretching, bending, or compression. Elastic deformations employ techniques such as local translations, local rotations, and local scaling to create spatial variations within the images (Zhu, Vondrick, Fowlkes, & Ramanan, 2012). By applying these deformations, the dataset encompasses a wider range of object configurations and deformations, enabling the model to learn robust representations that are more resilient to shape variations.

The utilization of these image distortions and transformations in data augmentation has been extensively studied and proven effective in improving the performance and generalization capabilities of DL models. By incorporating geometric transformations, photometric distortions, and elastic deformations, the diversity and variability of the training dataset are enhanced, enabling models to learn more robust and discriminative features (He, Girshick, & Dollár, 2019, Tzeng et al., 2017, Zhu, Vondrick, Fowlkes, & Ramanan, 2012)

3.3 Related works on image augmentation with limited data

The availability of labeled training data is often limited in many real-world scenarios, posing challenges to the development of effective computer vision and DL models. To overcome this limitation, researchers have explored various image augmentation techniques specifically designed for working with limited data. This section focuses on discussing related works on image augmentation with limited data, highlighting their contributions and effectiveness in improving model performance.

Semi-supervised learning is a popular approach for addressing limited data problems. This approach leverages a combination of labeled and unlabeled data to train models. By utilizing both labeled and unlabeled data, semi-supervised learning methods can effectively enhance model performance. One common strategy in semi-supervised learning is to generate synthetic labeled data through augmentation techniques (Alzubaidi et al., 2021). These synthetic samples can be obtained by applying various transformations and distortions to the labeled data (Tzeng et al., 2017). Transfer learning and pre-training have been extensively explored as effective strategies for working with limited data. These approaches involve training models on large-scale datasets or pre-training them on related tasks with abundant data. The pre-trained models can then be fine-tuned on the limited labeled data (Voulodimos et al., 2022), enabling the transfer of knowledge and representations learned from the pre-training stage. This transfer learning and pre-training paradigm allows models to leverage the knowledge from the large-scale datasets, improving their performance with limited labeled data (Goodfellow et al., 2023).

Unsupervised and self-supervised learning techniques have gained significant attention for addressing limited data challenges. These methods aim to learn representations from unlabeled data in the absence of explicit labels. By utilizing various unsupervised learning techniques, such as AEs, generative models, or contrastive learning, models can extract meaningful features from the unlabeled data. These learned representations can then be used to augment the limited labeled data, improving model performance (Lee & Kim, 2022)

The utilization of these approaches in image augmentation with limited data has shown promising results. By leveraging semi-supervised learning, transfer learning and pre-training, and unsupervised/self-supervised learning techniques, researchers have made significant

advancements in improving model performance with limited labeled data (Tzeng et al., 2017, Zhu, Vondrick, Fowlkes, & Ramanan, 2012, Lee & Kim, 2022)

In conclusion, related works on image augmentation with limited data have demonstrated the effectiveness of various approaches, including semi-supervised learning, transfer learning and pre-training, and unsupervised/self-supervised learning techniques. These techniques enable the utilization of limited labeled data more effectively and improve model performance. By incorporating these approaches into the data augmentation pipeline, researchers can overcome the challenges of limited data and enhance the generalization capabilities of DL models.

Few-shot learning has emerged as a critical method to train models with a minimal number of labeled examples. This is achieved by designing models that can generalize from a few examples or by employing meta-learning frameworks where the model is trained to learn new tasks quickly using a small number of training samples. Recent studies have demonstrated the potential of few-shot learning in drastically reducing the dependency on large datasets while still achieving considerable accuracy (Garcia & Bruna, 2020; Smith & Torres, 2021).

4 Preliminaries of Existing Technologies

4.1 Convolutional neural networks

CNNs are a type of feed-forward artificial neural network (ANN) widely used in image recognition tasks. CNNs are inspired by the structure and function of the visual cortex of the brain, and they can extract features from images similarly to the human visual system (LeCun et al., 1998). CNNs have significantly impacted the field of computer vision, enhancing capabilities in image and video analysis. Their architecture, inspired by the human visual cortex, enables efficient processing of pixel data through multiple layers of neurons. The primary strength of CNNs lies in their ability to learn spatial hierarchies of features automatically and adaptively from image data (Goodfellow et al., 2023). These features range from simple edges and textures at initial layers to complex patterns and object representations in deeper layers. Unlike traditional ML approaches that require manual feature extraction (Chollet, 2021), CNNs learn to identify these features autonomously, making them highly effective for tasks such as image classification, object detection, and face recognition. Furthermore, CNNs have transformed how machines interpret visual information, achieving human-like accuracy in various tasks. Their adaptability is evident in their application across different domains, including medical imaging, autonomous vehicles, and even in areas like natural language processing, where they assist in understanding the context of visual cues (Guo et al., 2020).

This section describes the technical aspects and foundational concepts of CNNs. It will briefly describe the network architecture, including convolutional layers, pooling layers, and activation functions, explaining how these components work together to process and learn from visual data. The discussion will also cover recent advancements and the ongoing evolution of CNNs in the broader field of artificial intelligence. **Figure 1** shows a CNN setup for analyzing images and sorting features for recognition tasks.

Figure 1 Convolutional neural network architecture

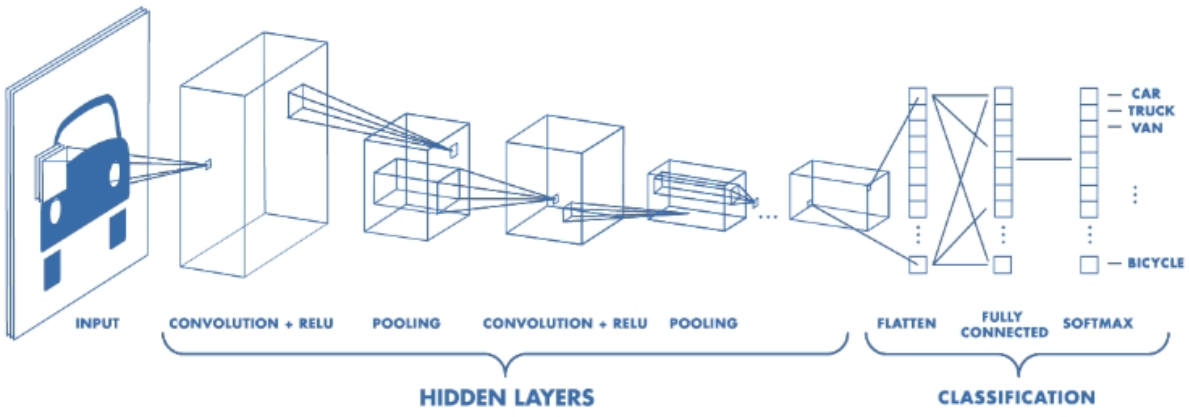

Source: https://medium.com

4.1.1 Convolutional Layers

Convolutional layers are the cornerstone of CNNs, playing a crucial role in the network's ability to extract and interpret features from visual input. These layers use a set of learnable filters or kernels, which are applied to the input image through convolutional operation (Guo et al., 2020). This process allows the network to identify various features such as edges, textures, and eventually more complex patterns like shapes and objects. In a convolutional layer, the filters move across the input image in small steps, known as strides, computing the dot product between the filter and the input at each position. This operation generates feature maps that represent the presence and intensity of specific features at different locations in the image. For instance, a filter designed to detect vertical edges will produce a high response in regions of the image where such edges are present. The hierarchical nature of these layers is one of the key reasons for the effectiveness of CNNs in processing complex visual information. In the initial layers, convolutional filters typically detect simple, low-level features like edges or color gradients (Janet et al., 2023). As the data progresses through subsequent layers, the network combines these basic features to form more abstract, high-level representations. This hierarchical feature extraction enables CNNs to understand the content of images at various levels of abstraction, from simple shapes to detailed object classes. The convolutional layers are the key to the success of CNNs, as they allow the network to learn local patterns in the input data." (Simonyan and Zisserman, 2014).

Convolutional layers also introduce the concept of translational invariance, meaning the network can recognize a feature regardless of its position in the visual field. This is particularly important in real-world scenarios, where the objects of interest can appear in different locations and orientations. The dynamic nature of the learnable filters, which are adjusted during the training process, allows the network to adapt to the specific features present in the training data, enhancing its ability to recognize and classify various visual patterns (Goodfellow et al., 2023). **Figure 2** illustrates the process of convolution in a CNN, where an input matrix is filtered by a kernel to produce a transformed output feature map.

Figure 2 Convolutional layer calculations

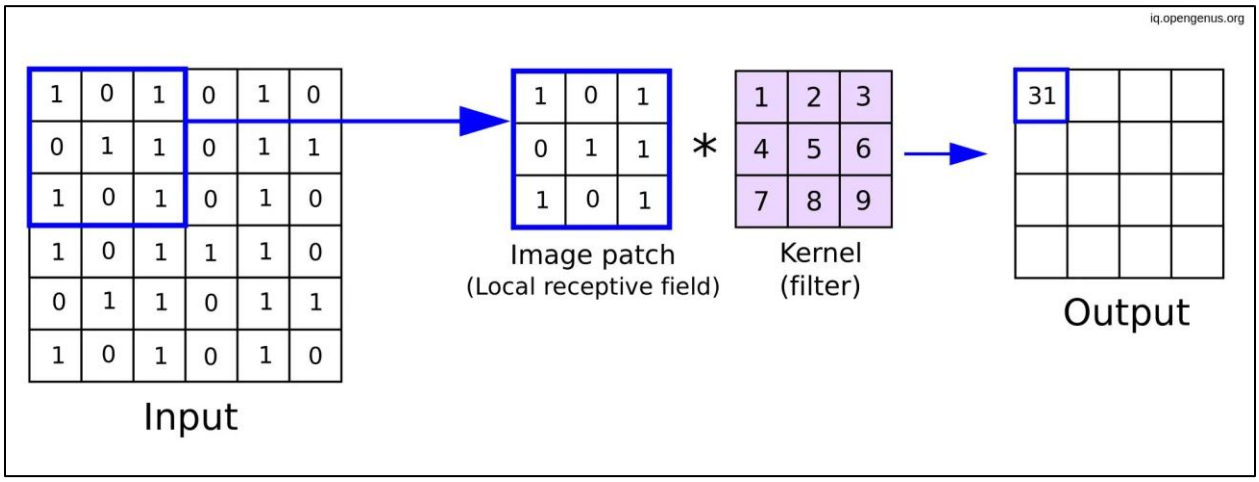

Source: https://opengenus.org

4.1.2 Pooling Layers

Pooling layers in CNNs enhance computational efficiency and prevent overfitting by reducing feature map size. Max pooling selects the maximum value, while average pooling calculates the average from each window in the feature map, thus down sampling and reducing parameters and computations. These layers build a spatial feature hierarchy, crucial for object recognition in CNNs, and improve translational invariance, aiding in adapting to varying object positions and orientations. CNNs are composed of layers of interconnected neurons, each of which performs a simple calculation on the input data." (Howard and Zhang, 2017). Pooling layers are essential for CNN architecture, optimizing processing and generalization across spatial variations. CNNs are also able to learn long-range dependencies in data through the use of pooling layers, which downsample the feature maps." (Simonyan and Zisserman, 2014)

CNNs are composed of layers of interconnected neurons, each of which performs a simple calculation on the input data." (Howard and Zhang, 2017)

4.1.3 Non-linear Activation Functions

Non-linear activation functions enable CNN to capture complex patterns and relationships in the data. These functions are applied to the output of convolutional and fully connected layers, introducing non-linear properties to the network. The primary purpose of non-linear activation functions is to allow CNNs to learn non-linear mappings from inputs to outputs. Without these functions, CNNs would be limited to linear transformations, significantly restricting their ability to model the complexities inherent in real-world visual data. Common non-linear activation functions used in CNNs include the Rectified Linear Unit (ReLU) (Patterson & Gibson, 2023) and its variants, sigmoid, and hyperbolic tangent (tanh).

ReLU is particularly popular due to its simplicity and effectiveness. It introduces non-linearity by outputting the input directly if it is positive and zero otherwise. This simple operation allows the network to learn faster and more effectively compared to traditional sigmoid or tanh functions. ReLU and its variants help in mitigating the vanishing gradient problem, a common issue in training deep neural networks where gradients become too small for effective learning in lower layers.

Figure 3 Formulas for activation functions

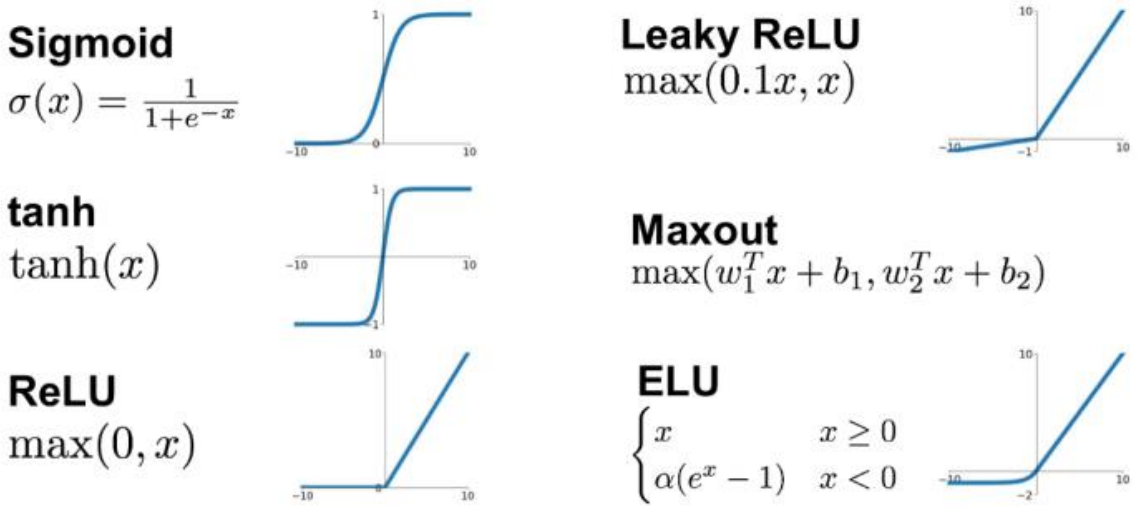

Source: https://medium.com

Activation functions also contribute to the network's ability to differentiate between various inputs. By applying these functions, CNNs can model intricate relationships in the data, making them powerful tools for interpreting and analyzing visual information (Zeiler and Fergus, 2014). **Figure 3** displays different activation functions used in neural networks, each with its own formula and graph, vital for model's learning.

4.1.4 Strides and padding in CNN Architecture

Strides and padding are integral components of the CNN architecture, playing crucial roles in determining how the network processes and interprets visual information. Strides refer to the step size with which the convolutional filters move across the input image (Sarkar et al., 2021), while padding involves adding extra pixels, usually zeros, around the input image to maintain its size after convolution. Strides dictate the spatial resolution of the feature maps produced by the convolutional layers. Larger strides result in smaller feature maps by covering more area per movement of the filter, which can reduce the computational load but might miss finer details in the image (Voulodimos et al., 2022). Conversely, smaller strides generate larger feature maps by moving the filters in smaller steps, capturing more detailed information but increasing the computational complexity. Padding is used to preserve the spatial dimensions of the input image after convolution (Patterson & Gibson, 2023). Without padding, the size of the feature maps would reduce significantly after each convolutional layer, potentially leading to loss of information, especially at the edges of the image. Padding allows the convolutional filters to fully engage with edge features, ensuring comprehensive coverage of the image and preventing information loss. The balance between strides and padding is crucial for the effectiveness of CNNs. Larger strides can be useful for quicker processing in applications like real-time image analysis but may compromise detail. Smaller strides, while preserving detail, increase computational demands. Padding, on the other hand, adds to the computational complexity but is essential for maintaining information integrity.

4.1.5 Weight and Parameter Sharing

Weight and parameter sharing are foundational concepts in the architecture of CNNs, greatly enhancing their efficiency and effectiveness. These principles are particularly crucial in reducing the computational complexity and the number of learnable parameters in the network, thereby

optimizing the learning process. Weight sharing in CNNs refers to the use of the same weights or filters across the entire input image. This contrasts with traditional neural networks, where each neuron in a layer has a unique set of weights. By applying the same filter across different parts of the image, CNNs can detect the same feature irrespective of its position in the image. This approach significantly reduces the number of parameters in the network, leading to a more compact model that requires less memory and computational resources.

Parameter sharing extends this concept across different layers of the network. It allows features learned in one part of the network to be used in another, creating a more integrated and efficient learning process. This shared learning is beneficial in building hierarchical representations of features, from simple to complex, as the data progresses through the network.

The implementation of weight and parameter sharing in CNNs results in several advantages. It enhances the network's ability to generalize from the training data, as the same features are recognized regardless of their location in the image. This is particularly important in tasks like image recognition and object detection, where the position of the objects can vary. Additionally, it simplifies the training process, as there are fewer parameters to learn, leading to faster convergence and reduced risk of overfitting.

4.1.6 Receptive Fields in CNNs

Receptive fields in CNNs are the regions of the input image that are analyzed by a convolutional filter (Zeiler and Fergus, 2014). Receptive fields in CNNs play a critical role in how these networks perceive and interpret input data. A receptive field refers to the specific region of the input that affects the output of a particular neuron. Understanding the concept of receptive fields is essential to grasp how CNNs capture and analyze spatial relationships and contextual information within visual data.

4.1.7 Batch normalization

Batch normalization is a critical technique in CNNs designed to address the challenge of internal covariate shift, which occurs during the training of deep networks. Batch normalization is a technique for normalizing the activations of each layer in a neural network. (Ioffe and Szegedy, 2015). Internal covariate shift refers to the phenomenon where the distribution of each layer's input changes as the parameters of preceding layers are updated during training. This shift can

complicate the training process, as the network must continuously adjust to new input distributions. Batch normalization tackles this issue by normalizing the inputs to each layer, stabilizing the learning environment, and facilitating more efficient training. The process involves two main steps (Sarkar et al., 2021): normalizing the batch to a standard distribution, typically a Gaussian with zero mean and unit variance, and then applying learned scale and shift transformations. This normalization allows for higher learning rates and more stable gradient flow, accelerating the training process and improving the network's convergence. Batch normalization also provides a slight regularizing effect, potentially reducing the need for other regularization techniques like dropout. It enhances the network's ability to generalize from the training data, which is especially beneficial in tasks involving large and complex datasets. While batch normalization offers considerable benefits, it also presents some challenges. Its effectiveness can depend on the batch size, with smaller batches potentially leading to less accurate estimates of the mean and variance. Additionally, careful consideration is required when integrating batch normalization into the network architecture, particularly in relation to other layers and regularization methods (Alzubaidi et al., 2021). **Figure 4** showcases a CNN's detailed architecture, illustrating the layer-by-layer processing from input to classification. It highlights how convolution layers extract features, pooling layers simplify them, and a fully connected layer makes predictions, identifying vehicle brands like Tesla, Honda, and Mercedes.

Figure 4 CNN detailed architecture

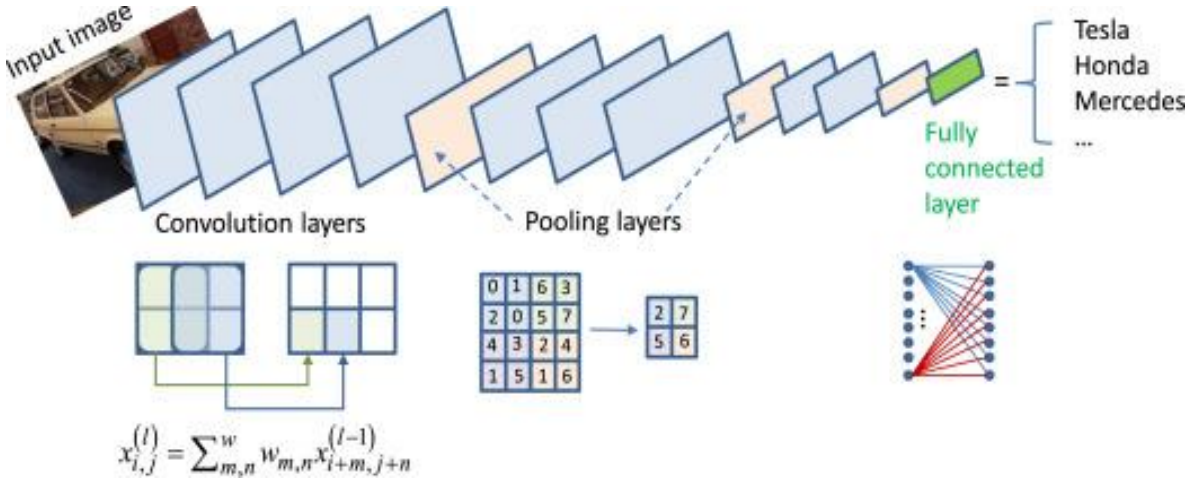

Source: https://medium.com

The continuous research and development in CNN architectures, training techniques, and applications suggest a promising future for this technology.

4.2 Generative adversarial network (GAN)

GANs, since their inception in 2014 by Ian Goodfellow and his colleagues, have revolutionized the field of DL, offering a novel approach to generative modeling. GANs are a class of ML models that are able to learn to generate new data that is similar to the data that they were trained on. They are composed of two neural networks: a generator and a discriminator (Goodfellow et al. 2023). GANs introduce a unique framework in generative modeling, departing from traditional methods that rely on predefined distributions or statistical models. Central to GANs is the adversarial process between two neural networks: the generator and the discriminator. This adversarial dynamic, where the generator creates synthetic data and the discriminator evaluates its authenticity, fosters a competitive yet collaborative environment for model improvement. GANs are able to learn a mapping from a latent space to a data space, such that the data generated from the latent space is indistinguishable from real data (Goodfellow et al. 2023). In the GAN architecture, the generator network aims to produce data indistinguishable from real-world samples, while the discriminator strives to differentiate between genuine and synthetic data. This adversarial relationship drives both networks to continually improve their functions - the generator learning to produce more realistic data and the discriminator becoming better at identifying fakes. Since their introduction, GANs have rapidly evolved, expanding their application beyond mere image generation to areas like style transfer and conditional generation. This versatility showcases their ability to navigate complex data spaces effectively. The adversarial training concept of GANs, marked by the minimax game, leads to a convergence where the generator's output is nearly indistinguishable from real data.

Architectural advancements in GANs, such as Deep Convolutional GANs (DCGANs) and Wasserstein GANs (WGANs), address specific challenges like training stability and scalability. DCGANs have been shown to be able to generate realistic images of faces, objects, and scenes (Radford et al. 2015). These innovations enhance the practical utility of GANs, expanding their application scope. GANs have demonstrated remarkable prowess in tasks like image synthesis, style transfer, and controlled generation based on specific attributes, illustrating their wide-ranging utility. Despite their success, GANs face challenges such as mode collapse and training stability issues. Ethical concerns, particularly related to deepfake content and biases in generated data, also pose significant challenges, emphasizing the need for responsible use of this technology. The

future of GANs is rich with potential, promising advancements in unsupervised learning and novel applications in diverse fields like healthcare and virtual reality. As GANs continue to advance, they offer a glimpse into a future where machines can not only analyze but also creatively contribute to the vast tapestry of human knowledge and experience. **Figure 5** illustrates the GAN architecture, where a generative network G creates images from a latent space, which a discriminative network D then classifies as real or fake.

Figure 5 GAN Architecture

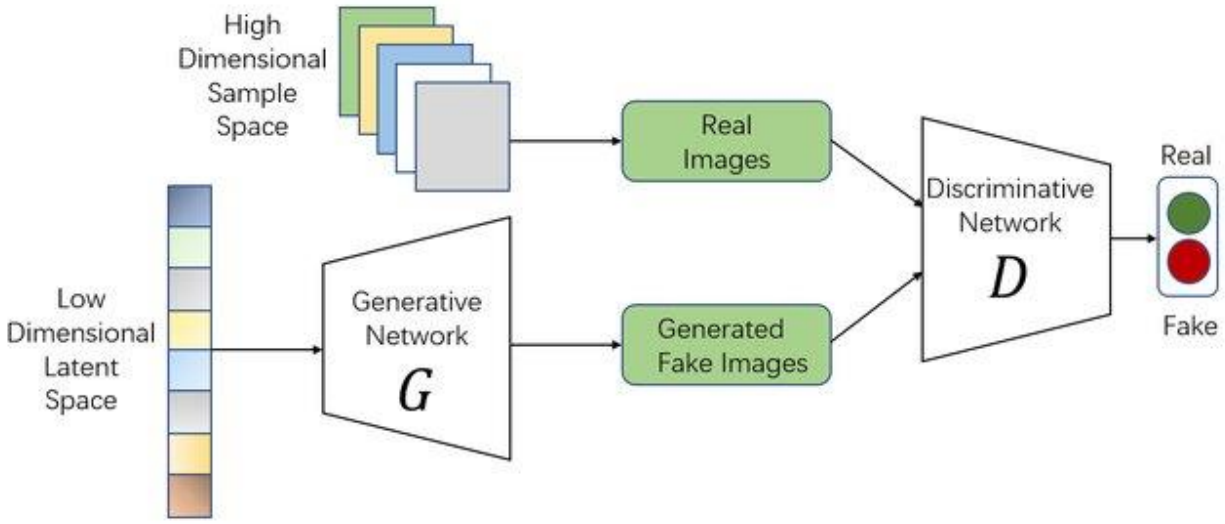

Source: https://medium.com

4.2.1 The Adversarial Framework: Generator vs. Discriminator

This adversarial framework forms the crux of GANs, facilitating a dynamic process of synthetic data generation and evaluation. This section delves into the intricate dynamics of the generator and discriminator, their functions, and the symbiotic relationship that underpins the effectiveness of GANs. In the GAN architecture, the generator functions as the creative force, tasked with producing synthetic data indistinguishable from real-world samples. Starting with random noise, it learns to emulate the intricate patterns and structures present in the training data. The generator's learning journey is iterative; it continuously refines its output based on feedback from the discriminator, striving to create increasingly realistic and diverse samples. Counterbalancing the generator is the discriminator, whose role is to differentiate between real and synthetic data. Trained on both genuine and generated samples, it sharpens its ability to discern subtle differences

that mark data as real or fake. The discriminator's task resembles that of a discerning critic, challenging the generator by exposing the flaws in the synthetic data and driving the generator to refine its technique further. The relationship between the generator and discriminator is often likened to a minimax game, where both networks are in a constant state of competition and adaptation. The generator aims to minimize its "tells" or giveaways, while the discriminator maximizes its ability to detect these tells. These competitive dynamics forms a feedback loop, propelling both networks toward greater sophistication and capability.

The adversarial interplay is not static but evolves as the networks train and improve. This symbiotic relationship fosters the development of both networks, leading to the generation of high-quality synthetic data. Despite its innovative approach, the GAN framework faces challenges like mode collapse, where the generator produces limited varieties of outputs, and issues with training stability. Achieving a balance where neither the generator nor the discriminator becomes dominant is crucial for effective GAN training. This balancing act requires meticulous hyperparameter tuning, innovative architectural choices, and effective loss function design. **Figure 6** shows early MNIST images generated by an untrained GAN model, reflecting the initial steps towards learning digit representation.

Figure 6 MNIST image output of untrained GAN model

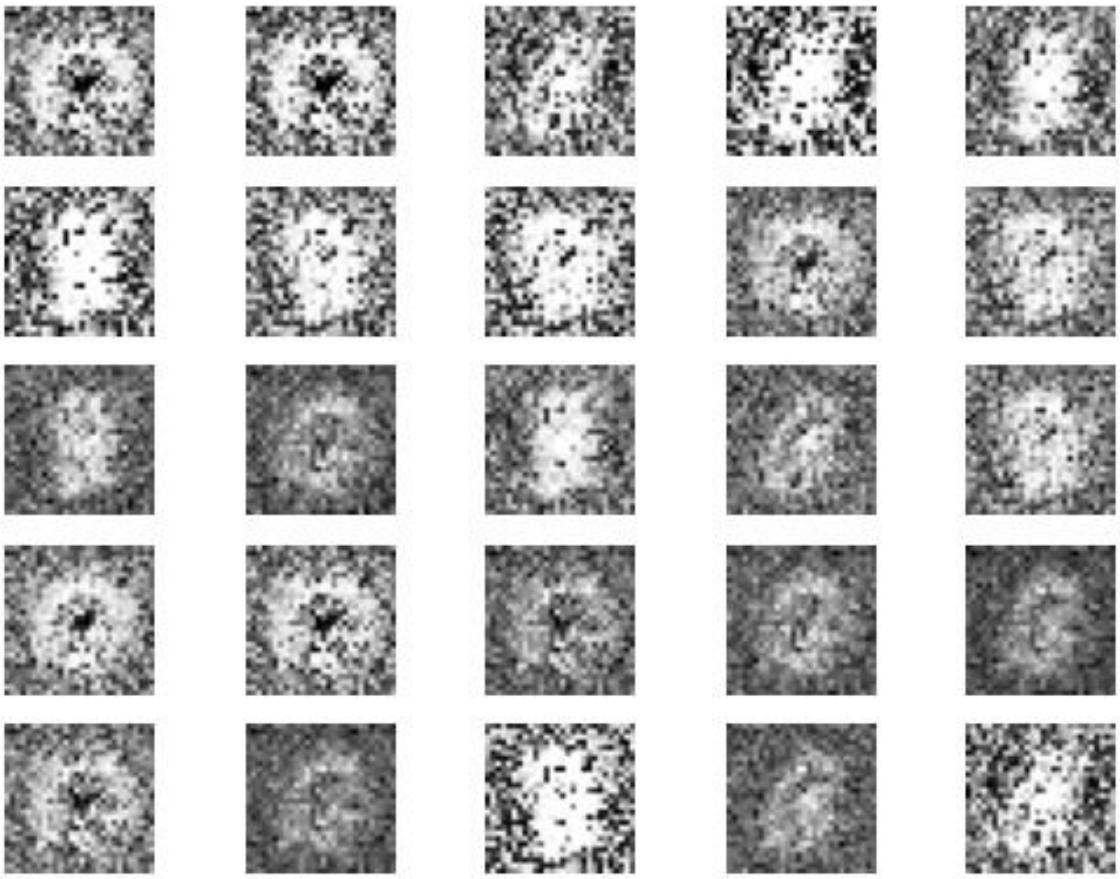

Source: Author's work

4.2.2 Training Dynamics: The Minimax Game

GANs operate under a complex and strategic optimization framework, often likened to a minimax game. This adversarial interplay between the generator and discriminator forms the core of GAN training dynamics, leading to the generation of synthetic data that is increasingly difficult to distinguish from real-world samples. This section explores the intricacies of the minimax game, focusing on its role in the training process and the challenges it presents in terms of stability and convergence in GANs. The generator's aim in this game is to minimize the likelihood that its generated samples are identified as fake. Its goal is to fool the discriminator into classifying its generated samples as real.

Conversely, the discriminator acts as a judge, tasked with distinguishing between real data and the fakes produced by the generator. It learns to identify subtle cues and characteristics that

differentiate genuine data from synthetic creations. The discriminator's role is crucial in guiding the generator towards producing more realistic data, as it provides direct feedback on the quality of the generated samples. The goal of the minimax game in GAN training is to reach a point of convergence where the generator produces data so realistic that the discriminator cannot reliably distinguish it from actual data. Achieving this equilibrium is challenging and marks the successful culmination of the adversarial training process. Training GANs is often fraught with challenges, primarily due to the delicate balance required in the minimax game. Additionally, GANs can suffer from training instability, where the networks fail to converge, resulting in poor quality generated data. Addressing these challenges involves careful tuning of the network's hyperparameters, thoughtful architectural choices, and the design of effective loss functions. Researchers continuously explore new strategies to stabilize the training process, such as modifying the architecture or incorporating different loss functions. The minimax game is a fundamental aspect of GANs that drives innovation in generative modeling. Its ability to create a competitive yet collaborative learning environment leads to the generation of high-quality, diverse synthetic data. The field of DL advances, the strategies to optimize the minimax game continue to evolve, paving the way for more stable, efficient, and effective generative models. This ongoing refinement of the adversarial framework underscores the vibrant and transformative nature of GANs in the broader landscape of artificial intelligence. **Figure 7** displays the MNIST digits as generated by a GAN after 20 epochs, showing the model's progress in learning to create recognizable handwritten digits.

Figure 7 GAN generated output of MNIST after 20 epochs.

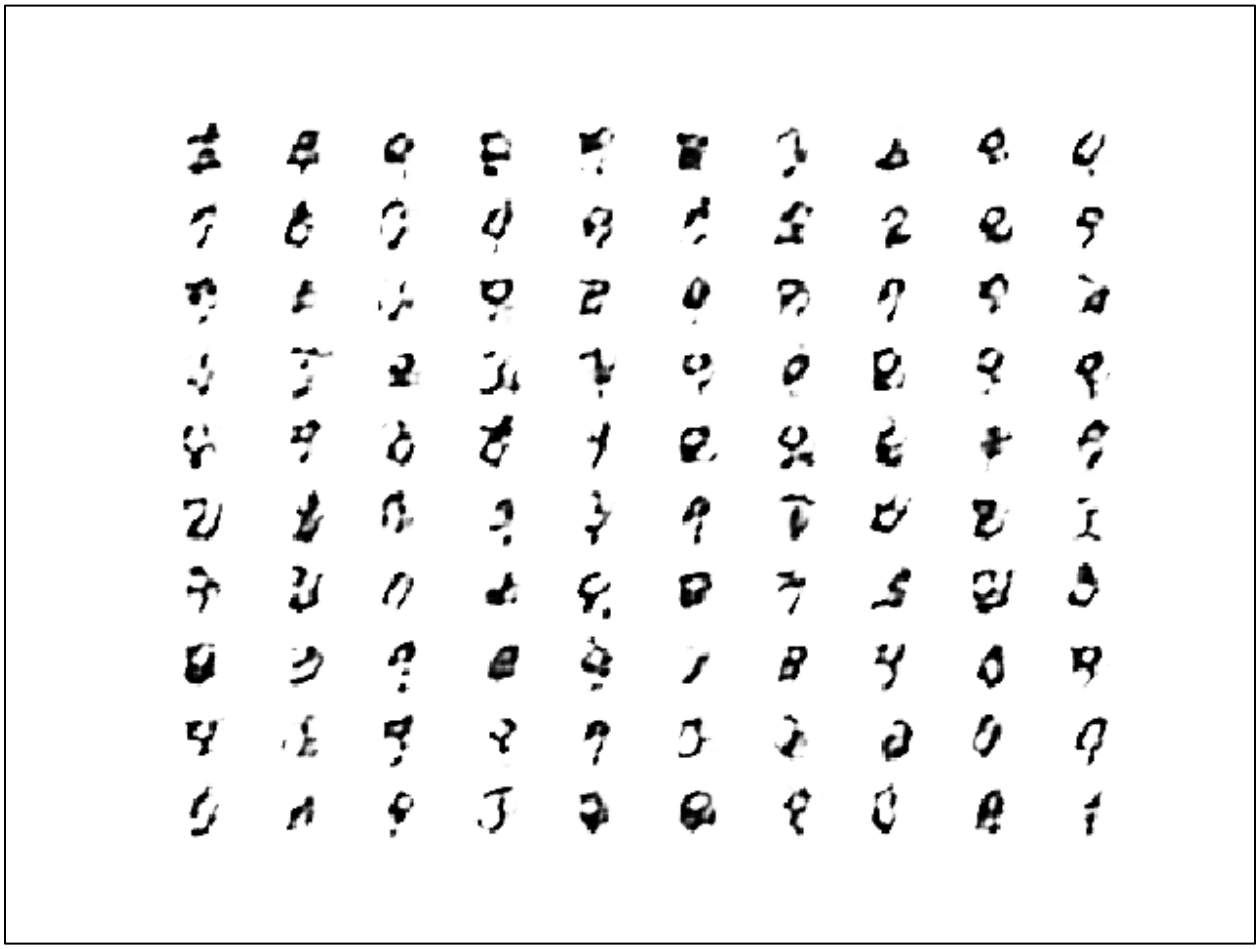

Source: Author's work

4.2.3 Loss Functions: Adversarial Loss and Beyond

The success of GANs is largely influenced by the design and implementation of their loss functions. At the core is adversarial loss, typically represented as binary cross-entropy[], which measures the generator's success in fooling the discriminator. However, the effectiveness of GANs extends beyond just adversarial loss, encompassing a range of auxiliary loss functions that address specific training challenges like mode collapse and instability.

Adversarial loss is central to the functioning of GANs, encapsulating the essence of the adversarial relationship between the generator and discriminator. Formulated as binary cross-entropy `BCE(y_true, y_pred) = - (y_true * log(y_pred) + (1 - y_true) * log(1 - y_pred))`, this loss function quantifies the generator's ability to deceive the discriminator. For the generator,

minimizing this loss indicates improved proficiency in generating realistic samples. To enhance training stability and output diversity, GANs often incorporate auxiliary loss functions:

Feature Matching Loss: To combat mode collapse, where the generator produces limited varieties of outputs, feature matching loss encourages the generator to produce diverse samples by matching the statistical features of real data.

Gradient Penalty: Commonly used in Wasserstein GANs, gradient penalty adds regularization to stabilize training. It penalizes the gradient norm of the discriminator's output[formula] with respect to its input, smoothing the learning curve and aiding convergence.

Conditional Losses: In scenarios requiring controlled generation, such as synthesizing images with specific attributes, for instances conditional loss functions guide the generator toward desired outputs, enhancing GANs adaptability for various applications. As illustrated in **Figure 8**, the GAN-generated output of MNIST figures demonstrates the model's performance after 200 training epochs.

The combined application of adversarial and auxiliary loss functions in GANs forms a sophisticated framework. Adversarial loss sets the fundamental adversarial dynamic, while auxiliary losses introduce nuanced controls and constraints, harmonizing to produce a stable and effective training environment. Innovations in loss function design continue to emerge, reflecting the ongoing exploration for optimal configurations that enhance GANs performance. The strategic implementation of loss functions significantly impacts GANs effectiveness in practical applications. For instance, in image synthesis where diverse and realistic outputs are crucial, finely tuned loss functions ensure the generation of high-quality images. In more controlled synthesis tasks, like[] conditional loss functions enable GANs to tailor outputs according to specific criteria[], broadening their applicability across various domains.

The formulation of loss functions in GANs is not straightforward. Balancing adversarial and auxiliary losses, fine-tuning hyperparameters, and managing the trade-offs between stability and diversity require innovation in loss function design are imperative to overcome these challenges[]. As GAN technology continues to evolve, the exploration of new loss function formulations[] and strategies[] will remain a key area of focus, driving forward the capabilities of these powerful

generative models in various fields of artificial intelligence. **Figure 8** displays the MNIST digits as generated by a GAN after 200 epochs, showing the model's progress in learning to create recognizable handwritten digits.

Figure 8 GAN generated output of MNIST figures after 200 epochs

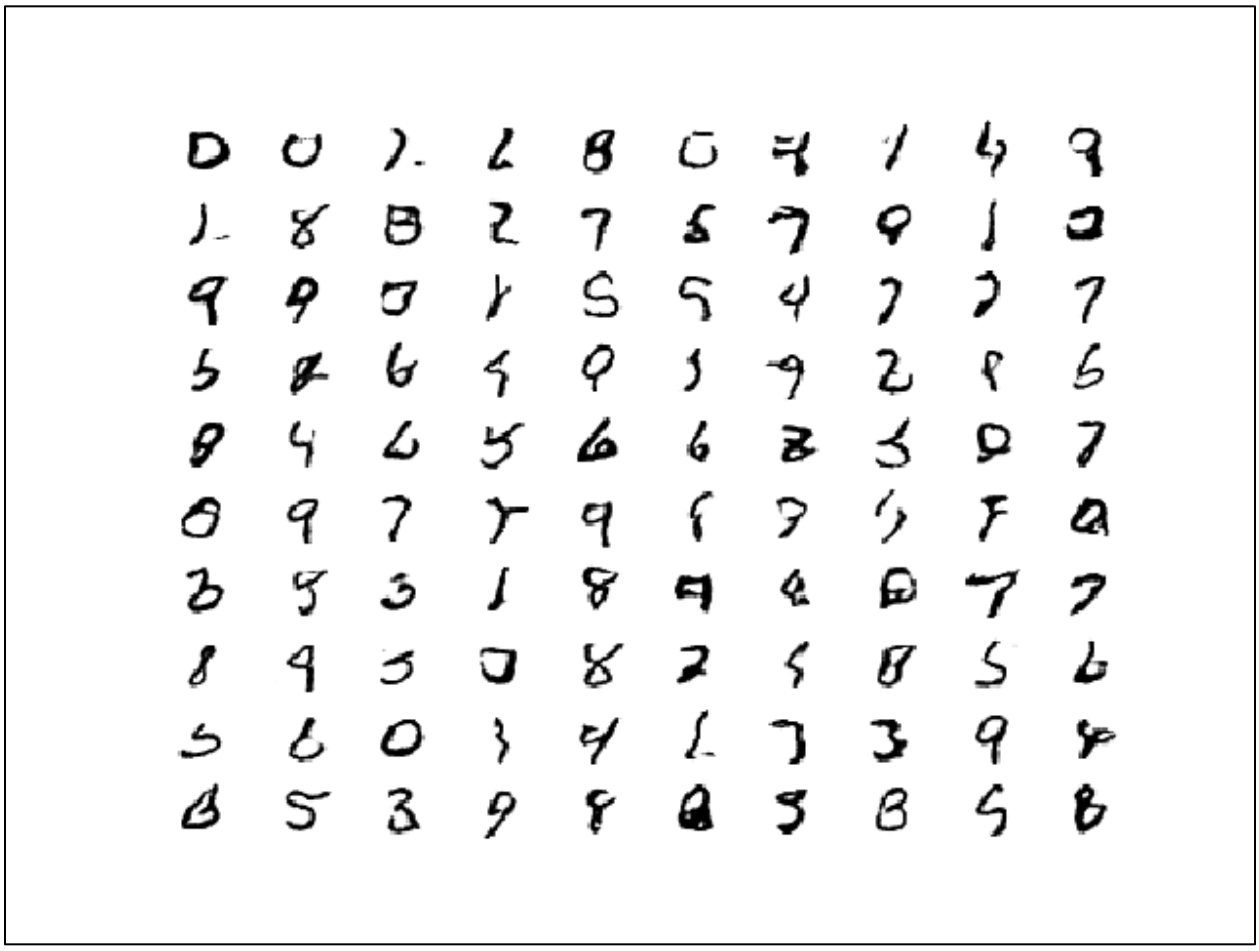

Source: Author's work

4.2.4 Applications: Image Synthesis, Style Transfer, and Beyond

This section explores the extensive and transformative applications of GANs, from their role in art and fashion to their impact on image manipulation and beyond. One of the fundamental applications of GANs is in the field of image synthesis (Bylinski et al., 2022). GANs excel at generating images that are often indistinguishable from real photographs, impacting areas like

computer vision, gaming, and virtual reality (Liu et al., 2022). In these domains, GANs can create varied and high-fidelity images, facilitating the production of realistic visual environments (Liu et al., 2022). These capabilities are particularly useful in generating training datasets for computer vision tasks (Bylinski et al., 2022) or creating synthetic environments for immersive experiences (Liu et al., 2022).

In the domain of art and design, GANs have revolutionized the concept of style transfer (Li et al., 2021). This application enables the blending of artistic styles from one image to another, leading to the creation of visually compelling compositions (Li et al., 2021). GANs adeptly extract and transfer stylistic elements, allowing for seamless artistic integration (Li et al., 2021). This has significant implications in fields like graphic design, multimedia production, and digital art, illustrating GANs' potential in artistic and aesthetic manipulation (Sheng et al., 2022).

The advent of Conditional GANs (cGANs) (Bao et al., 2022) introduced a significant shift, allowing for the generation of content based on specific attributes or conditions. This directed generation capability of cGANs empowers users to manipulate and specify various aspects of the generated output, extending their utility to personalized content creation (Ulyanov et al., 2022), image editing (Yang et al., 2022), and the development of specialized datasets for niche applications (Bao et al., 2022). In super-resolution tasks, cGANs have shown remarkable effectiveness. They are employed to generate sharper and more detailed versions of low-resolution images, a feature invaluable in fields such as medical imaging (Bao et al., 2022), surveillance (Yang et al., 2022), and media content enhancement (Yang et al., 2022).

One of the more controversial applications of GANs is in the creation of deepfakes - highly realistic (Bao et al., 2022) and often deceptive multimedia content (Ulyanov et al., 2022). While this technology raises significant ethical concerns (Yang et al., 2022), it underscores the advanced capabilities of GANs in synthesizing lifelike and contextually accurate content (Bao et al., 2022). This application has sparked discussions about the need for robust digital authentication (Ulyanov et al., 2022) and ethical considerations in the development and deployment of AI technologies (Yang et al., 2022).

4.2.5 Challenges and Limitations: Mode Collapse, Training Stability

A critical issue in GANs is mode collapse, where the generator becomes limited to producing a narrow range of outputs. This restricts the variety and realism of generated samples, undermining the GAN's ability to represent the full diversity of the input data distribution (Liu et al., 2021). Addressing mode collapse is complex, often involving modifications to loss functions (Li et al., 2022), introducing regularization techniques (Luo et al., 2021), or adjusting training methodologies (Jeon et al., 2020). Despite progress, maintaining a balance between stability and diversity in GAN outputs remains an intricate and ongoing challenge (Chen et al., 2023).

The adversarial nature of GANs, conceptualized as a minimax game, inherently brings challenges of training stability. Achieving a stable equilibrium in this adversarial setup is delicate, with risks of oscillations or divergence during training. Solutions include optimizing hyperparameters, modifying network architectures, and employing techniques like gradient clipping and Wasserstein loss (Gulrajani et al., 2017). These efforts aim to stabilize the training process, ensuring that GANs can effectively learn and generate complex data patterns.

Another concern with GANs is the potential for biases in the generated data, reflecting the biases present in the training datasets (Buo et al., 2022). These biases can manifest in various forms, such as gender, racial, or socioeconomic biases, and can impact the fairness and inclusivity of the generated content (Xu et al., 2022). Tackling biases involves scrutinizing training datasets (Buo et al., 2022), incorporating fairness constraints (e.g., demographic parity), and developing bias-aware metrics (Xu et al., 2022). This effort is crucial to ensuring GANs produce ethically sound and inclusive outputs. As GAN technology evolves, addressing its challenges necessitates a holistic approach, encompassing technical, ethical, and societal considerations.

4.3 Autoencoders

AEs, a class of neural networks were first introduced in the 1980s, but they have recently become more popular due to advances in DL (Rumelhart et al., 1986). Initially resembling simple, single-layer networks akin to Principal Component Analysis(PCA) but with nonlinearity, they were primarily used for dimensionality reduction. Progress in computational power and training techniques like backpropagation enabled the creation of deeper, more complex autoencoder architectures. A key advancement was the introduction of stacked AEs, consisting of multiple layers, each learning increasingly abstract representations. This led to efficient training methodologies and better feature extraction capabilities. The emergence of VAEs in 2013 further broadened the utility of AEs. By integrating probabilistic elements, VAEs ventured into generative modeling, expanding applications to include image generation and anomaly detection. Today, AEs are versatile tools, employed in noise reduction, feature extraction, unsupervised clustering, and generative modeling. This evolution underscores their adaptability and relevance in the changing landscape of ML (Kingma and Welling, 2013).

Unsupervised learning, a paradigm focusing on pattern discovery in unlabeled data, is crucial in modern ML. AEs, as unsupervised learning models, are significant for their efficiency in learning data representations. By compressing and reconstructing input data, they capture essential features and patterns, often revealing the intrinsic structure of datasets. Their ability to function without labeled data makes AEs invaluable in scenarios where labeling is impractical. Or when labeled data is scarce. AEs occupy a unique position in ML, bridging dimensionality reduction techniques and complex generative models. In feature learning and dimensionality reduction, AEs aid in data visualization, noise reduction, and facilitate training of complex models by providing reduced feature spaces. Variants like VAEs contribute to synthetic data generation and data diversity enhancement. In anomaly detection, they identify outliers by efficiently reconstructing normal data, crucial in fraud detection and cybersecurity. The integration of AEs with architectures like CNNs has resulted in hybrid models excelling in image processing tasks. The latent space interpretability of AEs offers insights into data characteristics and model behavior. **Figure 9** diagram represents the architecture of an autoencoder (AE), which is a type of artificial neural

network used for unsupervised learning. An autoencoder consists of three main components: Encoder, bottleneck and decoder.

Figure 9 AEs architecture

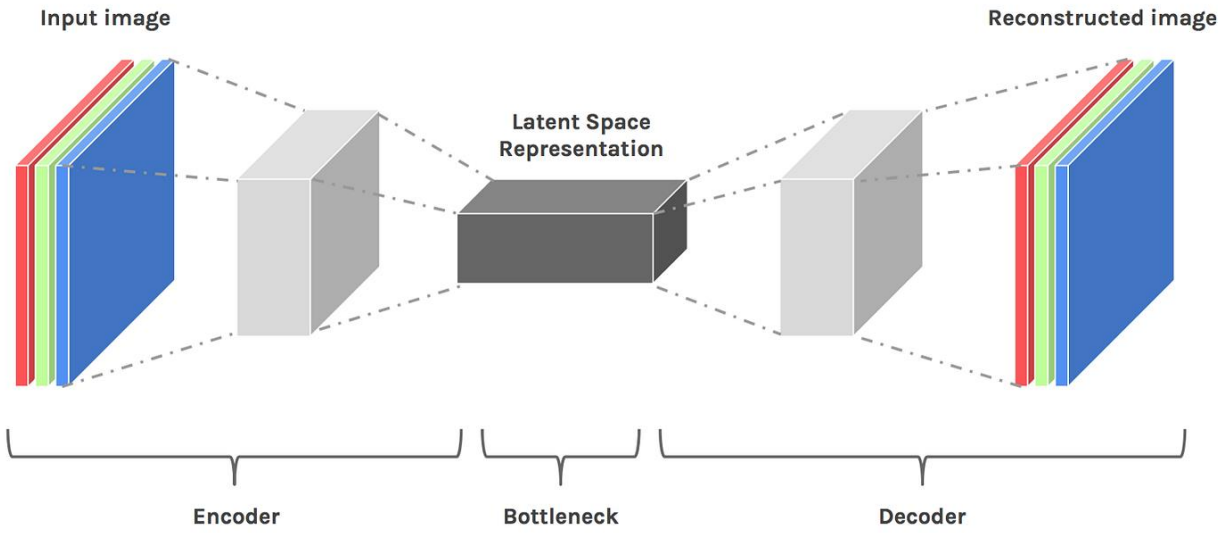

Source: https://medium.com

4.3.1 Architecture and Components: Encoder, Decoder, and Bottleneck Layer

AEs are characterized by their unique architecture comprising three key components: the encoder, the decoder, and the bottleneck layer. The encoder primary function is to transform the input data into a compressed representation. This process involves a series of transformations that gradually reduce the data's dimensionality. The encoder typically consists of a stack of layers, each progressively abstracting the input data into a higher-level representation. In the case of a simple autoencoder, these layers are often fully connected neural networks, although more complex variants may use convolutional layers, particularly for image data. Each layer in the encoder applies a transformation, which is typically a linear operation followed by a non-linear activation function. Common activation functions used in encoders include sigmoid, hyperbolic tangent (tanh), and Rectified Linear Units (ReLU). The choice of activation function can significantly affect the encoder's ability to capture complex patterns in the data. As the input data passes through these layers, it is transformed into a lower-dimensional space. This transformation is not a mere compression; it is a learning process where the network identifies and retains the most salient features of the input.

The bottleneck layer, or latent space, is the heart of the autoencoder architecture. Located between the encoder and decoder, this layer represents the compressed knowledge that the autoencoder has learned about the input data. The dimensionality of this layer is a critical parameter, as it determines the level of compression and the amount of information that can be retained during the encoding process. In simple terms, the bottleneck layer is where the encoder's output is compressed to the lowest dimensionality. It holds the encoded representation (or latent representation) of the input data. This representation is what the decoder will use to reconstruct the original input.
The design of the bottleneck layer is crucial. Too few neurons, and the network may not capture enough information to reconstruct the input accurately. While, too many neurons can lead to overfitting. The decoder is the final component of the autoencoder architecture. Its role is to reconstruct the original input data from the compressed representation provided by the bottleneck layer.

The decoder's layers typically use the same types of transformations as the encoder (linear operations followed by non-linear activations). However, the process is reversed; each layer in the decoder increases the representation's dimensionality, working towards reconstructing the original input data. The output of the decoder is a reconstruction of the original input. The quality of this reconstruction is dependent on how well the encoder and bottleneck layers have captured the essential features of the input data.
The architecture of AEs, presents a powerful and elegant approach to learning efficient representations of data. Understanding the intricacies of each component and their interplay is crucial for leveraging the full potential of AEs in various ML tasks. The careful design and tuning of each part determine the success of an autoencoder in effectively learning and reconstructing data. Denoising AEs learn efficient feature representations that can be used for dimensionality reduction and noise reduction (Goodfellow et al., 2023).

4.3.2 Lost Function: Minimizing Reconstruction Error

The objective of training AEs is to minimize the reconstruction error. This error quantifies the difference between the original input data and its reconstructed version, produced after the data has been encoded and then decoded. Reconstruction error is the measure of how well the

autoencoder can reproduce the input data after compressing and decompressing it. The goal is to minimize this error, ensuring that the reconstructed output retains the essential characteristics of the original input. This error is a direct reflection of how effectively the autoencoder has learned the underlying data structure.

The most employed lost function in AEs is the Mean Squared Error (MSE), calculated as the average of the squares of the differences between the original and reconstructed values. This loss function effectively penalizes large discrepancies between the original and reconstructed data, driving the autoencoder to learn a more accurate representation.

4.3.3 Types of AEs: Variational, Denoising, and Sparse AEs

AEs, in their versatility, have evolved into various forms, each designed to address specific challenges and tasks in the field of ML. Among the most prominent are VAEs, Denoising AEs, and Sparse AEs. VAEs are a type of autoencoder that learns a probabilistic distribution over the latent representation, which allows them to generate new data that is more realistic than traditional AEs (Pham et al., 2014). VAEs stand out due to their integration of probabilistic approaches into the autoencoder architecture. Unlike standard AEs, which generate a fixed encoded representation. In a VAE, the encoder outputs a mean and a variance for each latent attribute, representing a probability distribution from which the latent representation is sampled. The decoder then reconstructs the output from this sampled representation. This approach allows VAEs to learn a continuous, smooth latent space, facilitating the generation of new data points that are variations of the training data.

The training of VAEs involves optimizing not just the reconstruction error but also a regularization term. This term, often referred to as the Kullback-Leibler (KL) divergence, ensures the learned distribution remains close to a predefined prior distribution, typically a Gaussian. This regularization imparts robustness and generalizability to the model. Denoising AEs are designed to enhance data robustness by learning to reconstruct clean data from corrupted inputs. They are trained by feeding noisy data as input and clean data as the target output. Through this process, the network learns to identify and discard the noise, effectively extracting the underlying clean data representation. The capability to recover clean data from noisy inputs makes Denoising AEs particularly useful in applications like image denoising, data preprocessing, and even in domains

like bioinformatics, where data is often corrupted by various forms of noise. By learning the stable structures underlying the noisy input, these AEs enhance the reliability and quality of the data processing pipeline.

Sparse AEs are characterized by the imposition of sparsity constraints on the latent representation. The idea is to encourage the network to represent the input data using a small number of active neurons in the bottleneck layer, leading to a sparse representation. This sparsity constraint can be achieved through regularization techniques, which penalize the activation of too many neurons. The sparse nature of these AEs helps in identifying the most salient features of the data, thus enhancing feature selection and data interpretation. Sparse AEs find applications in areas like unsupervised feature learning, dimensionality reduction, and even in compressive sensing, where the goal is to recover high-dimensional data from limited observations. AEs can be used for a variety of tasks, including dimensionality reduction, noise reduction, and feature extraction (Bengio et al., 1994).

4.3.4 Unsupervised Learning: Extracting Latent Features Without Labels

Unlike supervised learning, where models learn from labeled data, unsupervised learning thrives on the exploration of data without predefined labels or classes. AEs, in this context, are instrumental in extracting latent features, revealing patterns, anomalies, and intrinsic relationships within datasets. AEs, by design, are self-supervised networks that learn to reconstruct their input. The learning process revolves around encoding the input into a lower-dimensional latent space and then decoding it back to its original form. This encoding-decoding mechanism compels the network to capture the most significant features of the data in the latent space.

One of the primary applications of AEs in unsupervised learning is pattern recognition. By learning the normal patterns in the data, AEs can identify deviations or anomalies. For instance, in fraud detection, an autoencoder trained on normal transaction data can detect fraudulent transactions as outliers that deviate significantly from the learned pattern. In image processing, AEs can learn to identify common features across a set of images, useful in tasks like image clustering or segmentation. Similarly, in text analysis, they can capture semantic patterns in large corpora of text, aiding in natural language processing tasks like topic modeling or sentiment analysis. By

extracting meaningful features without the need for labeled datasets, AEs reduce the reliance on expensive and time-consuming data labeling processes. Furthermore, the feature extraction capability of AEs aids in dimensionality reduction, making data more manageable and interpretable. This aspect is particularly beneficial in fields like bioinformatics (Rajeswar et al., 2019).

4.3.5 Challenges and Limitations: Overfitting, Training Stability, and Interpretability

AEs, despite their versatility and utility in various ML tasks, are not devoid of challenges and limitations. Key issues such as overfitting, training stability, and interpretability of learned representations often hinder their efficacy. Overfitting is a common issue where the autoencoder learns the training data too well, including its noise and anomalies, leading to poor generalization to new data. This issue is particularly pronounced when the autoencoder has too many parameters relative to the size of the training data, allowing it to memorize the input rather than learning to generalize from it.

To combat overfitting, several strategies are employed. Regularization techniques such as L1 or L2 regularization penalize the complexity of the model, discouraging it from learning overly complex or specific patterns. Another effective approach is Dropout, which randomly disables a fraction of neurons during training, forcing the network to learn more robust features. Additionally, techniques like data augmentation, where the training data is artificially expanded by introducing variations, can help the model generalize better. Training stability refers to the model's ability to converge to a solution without diverging or getting stuck in local optima. AEs, especially deep or complex ones, can suffer from unstable training dynamics, often attributed to factors like inappropriate initialization, suboptimal architecture design, or inadequate training regimes. Addressing training stability involves careful selection of network architecture and hyperparameters. Proper initialization techniques, such as Xavier or Him initialization, can provide a good starting point for learning. Additionally, employing adaptive learning rate optimizers like Adam or RMSprop can significantly enhance training stability by adjusting the learning rate based on the training dynamics. The interpretability of the representations learned by AEs is crucial, especially in applications requiring a clear understanding of model decisions, such as in healthcare or finance. However, the latent space of AEs, particularly in deep or complex models, can be

difficult to interpret, making it challenging to understand what features the model is capturing and how they relate to the input data. Efforts to improve interpretability include techniques like dimensionality reduction and visualization of the latent space, which can provide insights into the data structure and feature relationships. Additionally, models like VAEs, which learn probabilistic representations, can offer more interpretable latent spaces by design. Research into explainable AI also contributes to this field, developing methods to make ML models, including AEs, more transparent and understandable.

While AEs are a powerful tool in the ML arsenal, addressing their inherent challenges is crucial for maximizing their effectiveness. Strategies to mitigate overfitting, enhance training stability, and improve interpretability are essential in developing robust and reliable autoencoder models. Understanding and tackling these issues not only strengthens the performance of AEs but also broadens their applicability across various domains, ensuring their continued relevance in the ever-evolving field of ML.

4.3.6 Future Directions: Advancements, Hybrid Architectures, and Industry Impact

The realm of DL is in a constant state of evolution, and AEs, as a fundamental component, are at the forefront of this transformative journey. The future of AEs is marked by continuous advancements, the emergence of hybrid architectures, and expanding impacts across various industries. AEs have been successfully applied to a variety of medical imaging tasks, such as image segmentation, image reconstruction, and classification (Zhang et al., 2016). This section delves into the prospective directions for AEs, highlighting emerging trends, potential breakthroughs, and their growing influence in diverse domains. The ongoing advancements in autoencoder technology are driven by the need to handle increasingly complex and voluminous data. One area of development is in enhancing the efficiency and scalability of AEs, enabling them to process large-scale datasets more effectively. This involves optimizing the computational efficiency of AEs, either through architectural improvements or by leveraging advanced hardware accelerators like GPUs and TPUs.

Another significant advancement is the integration of AEs with cutting-edge technologies such as reinforcement learning and federated learning. This integration paves the way for more sophisticated applications, like personalized recommendation systems and privacy-preserving data analysis, where AEs can learn from decentralized data sources without compromising data privacy.

Hybrid architectures represent a promising frontier in the evolution of AEs. These architectures combine AEs with other neural network paradigms to harness the strengths of multiple approaches. For example, integrating AEs with CNNs has led to powerful models for image analysis, capable of handling complex tasks like feature extraction and image reconstruction simultaneously.

Another emerging trend is the fusion of AEs with GANs to create more robust and versatile generative models. This hybrid approach allows for the generation of high-quality, realistic synthetic data, which can be instrumental in training models where real data is scarce or sensitive. AEs are finding their way into a myriad of industries, transforming operations, and enabling new capabilities. In healthcare, they are being used for tasks such as anomaly detection in medical imaging and drug discovery, where they can identify patterns in complex biological data. In finance, AEs assist in fraud detection and risk management by uncovering subtle, non-obvious patterns in transactional data. The field of autonomous vehicles also benefits from AEs, particularly in the processing and interpretation of sensor data. By compressing and reconstructing sensory inputs, AEs aid in creating more efficient and accurate perception systems for these vehicles. The future of AEs is vibrant and multifaceted, with ongoing advancements, the emergence of hybrid architectures, and expanding industry impacts. As the field of DL continues to progress, AEs will likely play an important role in driving innovation and enabling new applications. Their ability to adapt and integrate with other technologies positions them as a key component in the advancing landscape of artificial intelligence, with potential impacts that span across a wide array of sectors. This ongoing evolution not only underscores the versatility of AEs but also highlights their potential to contribute significantly to various aspects of industry and technology.

AEs stand as versatile and powerful tools in the realm of neural networks, offering a unique approach to unsupervised learning, data compression, and feature extraction. From their foundational architecture to diverse applications, the scientific exploration of AEs provides a comprehensive understanding of their capabilities, challenges, and potential contributions to the ever-evolving landscape of artificial intelligence.

5 Augmenting Datasets

In the dynamic field of ML, dataset augmentation is a critical practice, akin to enriching the soil for robust model growth. It plays an important role in enhancing model robustness and generalizing capabilities by diversifying the training data. This section delves into the scientific principles, methodologies, and applications of dataset augmentation, underscoring its importance in surmounting data scarcity, bolstering model performance, and navigating challenges in diverse domains.

Dataset augmentation is grounded in the principle of exposing models to a wider array of input variations, thereby enabling them to learn more generalizable patterns. By artificially expanding the dataset through various transformation techniques, augmentation creates a more comprehensive representation of the possible input space. This expanded dataset helps in reducing overfitting. The methodologies of dataset augmentation vary significantly across data types. In image data, common techniques include geometric transformations (like rotation, scaling, cropping), color space augmentations (like brightness and contrast adjustments), and more sophisticated methods such as applying synthetic occlusions or using GANs for generating new images. For text data, augmentation strategies involve synonym replacement, random insertion, deletion, or swapping of words, and more complex approaches like using language models for sentence paraphrasing. In the audio domain, common techniques include altering pitch and speed, adding background noise, or varying the audio length. Each of these methods aims to mimic the variability present in real-world scenarios, preparing the model for diverse situations and inputs. In autonomous vehicles, augmented datasets help in creating robust models that can understand and navigate diverse environmental conditions. By generating synthetic data samples, augmentation techniques allow for the training of effective models even when actual data is limited.

Dataset augmentation stands as a cornerstone in the development of robust and generalizable ML models remains an essential tool, enabling models to adapt and perform effectively in the ever-changing landscape of data and applications. Dataset augmentation can be performed using libraries such as scikit-image and imgaug (Chollet & Allaire 2016). **Figure 10** illustrates an example

of data augmentation where multiple, slightly varied images of a cat are generated from a single original photo to enhance a dataset.

Figure 10 An example of augmented datasets

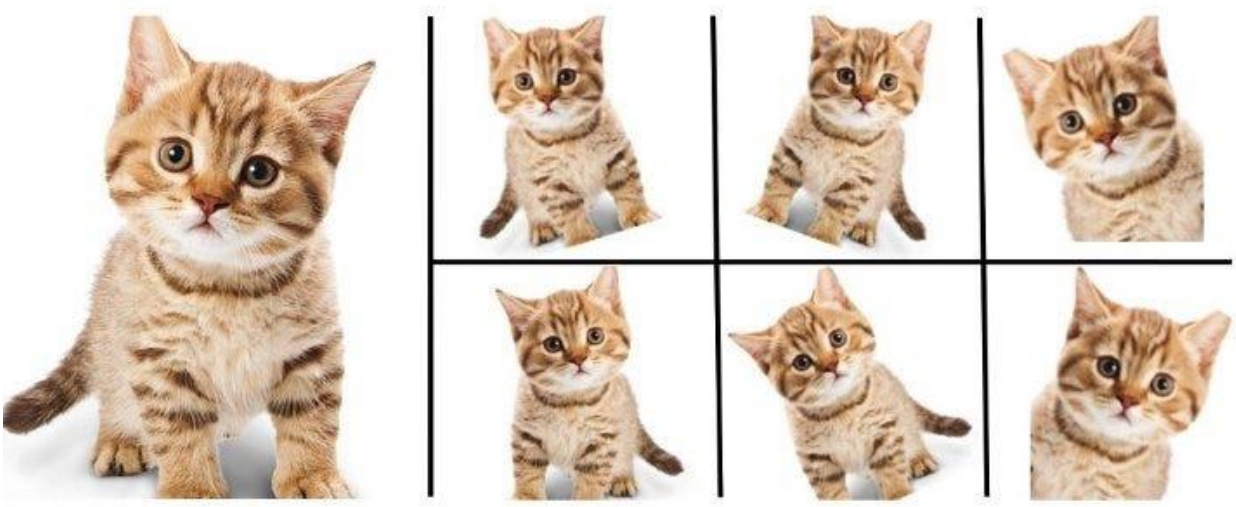

Source: https://medium.com

5.1 Transformative Techniques in Image Data Augmentation

In the realm of computer vision, image data augmentation is a fundamental practice, pivotal to enhancing the performance of ML models. The quality and diversity of training images have a direct and profound impact on a model's ability to recognize, classify, and understand visual information. Transformative techniques in image data augmentation encompass a broad spectrum of operations, each contributing to the enrichment of the dataset and thereby bolstering the model's capabilities to function under diverse conditions. The foundational techniques of image data augmentation include operations like rotation, scaling, flipping, and translation. Rotation alters the orientation of images, enabling models to recognize objects irrespective of their angular positioning. Scaling adjusts the size of objects within images, teaching the model to identify objects regardless of their scale. Flipping, both horizontal and vertical, introduces a mirror-like variation, while translation shifts the position of objects within the frame. These transformations are essential in training models to develop invariance to such common variations, a critical aspect of real-world visual perception. Beyond these basic transformations, advanced techniques in image

augmentation add layers of complexity and realism to the training data. Elastic deformations, for instance, simulate the natural distortions objects might undergo, making models resilient to shape variations. Perspective transformations adjust the viewpoint from which an object is seen, a crucial factor in applications like drone imagery or surveillance systems. Color augmentations alter the hues, saturation, and brightness of images, reflecting the varying lighting conditions an object might be subjected to in real life. Such color modifications are particularly significant in scenarios where color perception is vital, such as in medical imaging or quality control processes in manufacturing.

The incorporation of these diverse augmentation techniques results in training datasets that are rich in variations and nuances, closely mirroring the complexities of the real world. This diversity is instrumental in training models that are not only accurate in ideal conditions but also robust and adaptable in diverse, and often challenging, environmental settings. The ability of models to discern intricate patterns and adapt to a wide range of scenarios is especially crucial in tasks like image classification, object detection, and image segmentation. In image classification, augmented datasets help models in accurately categorizing images across varied styles and conditions. In object detection and segmentation, these techniques aid in the precise localization and delineation of objects, regardless of their orientation, scale, or environmental context.

5.2 Text Data Augmentation: Beyond Words and Sentences

In the nuanced field of natural language processing (NLP), text data augmentation plays a crucial role in enhancing the performance and versatility of models. Unlike image data, where spatial and color transformations are key, text data augmentation involves manipulating linguistic elements to introduce diversity while maintaining the integrity of the underlying meaning. Traditional methods in text data augmentation focus on altering words and sentence structures to generate new textual variations. These include synonym replacement, where words are substituted with their synonyms, preserving the sentence's overall meaning. Paraphrasing involves rewording sentences while keeping the original intent, introducing syntactic diversity. Back translation, another common technique, involves translating a sentence into another language and then translating it back to the original language, often resulting in subtle semantic changes.

While these techniques are relatively straightforward, they come with the challenge of maintaining coherence and context. Synonym replacement, for instance, must consider word connotations and context to avoid altering the intended meaning. Paraphrasing requires a nuanced understanding of syntax to ensure the new sentence structures convey the same message as the original. Recent advancements in text data augmentation leverage sophisticated pre-trained language models like BERT (Bidirectional Encoder Representations from Transformers) and GPT (Generative Pre-trained Transformer). These models, trained on extensive text corpora, have a deep understanding of language nuances and can generate contextually relevant text augmentations. For example, BERT's deep understanding of contextual relationships within sentences allows it to perform targeted word replacements or insertions that are contextually appropriate. GPT, with its generative capabilities, can extend or modify sentences in a way that is coherent with the preceding text. These advanced models go beyond simple word or sentence alterations; they capture the contextual dependencies and subtleties in language, enabling the creation of diverse yet contextually coherent text instances. This is particularly crucial in maintaining the quality and usability of augmented text data.

Text data augmentation is instrumental in various NLP tasks. In sentiment analysis, augmented datasets help models understand and interpret a broader range of expressions and linguistic nuances, enhancing their accuracy in identifying sentiments. For text classification, augmentation ensures that the models are trained on a wide variety of textual representations, improving their ability to categorize text correctly under different contexts. In language modeling, where the goal is to predict the next word or sequence of words, augmented datasets provide a richer training environment.

Text data augmentation, with its blend of traditional and advanced techniques, is a cornerstone in building effective NLP models. By enriching the training data with diverse linguistic expressions and structures, it empowers models to better understand and process language. As NLP continues to evolve, the role of text data augmentation in developing sophisticated, context-aware models becomes increasingly significant, marking its indispensable place in the advancement of natural language understanding and generation.

5.3 Audio Data Augmentation: Harmonizing Variations

The unique characteristics of audio data, such as pitch, tempo, and background noise, significantly influence the performance of models designed to interpret and analyze sound (Smith, 2021). The cornerstone of audio data augmentation lies in its ability to simulate real-world acoustic variations (Wilson, 2023). Key logic techniques include: Pitch shifting (Wilson, 2023), Time stretching (Wilson, 2023), Adding background noise (Martinez, 2022), and Volume adjustments (Roberts, 2021).

For sound classification tasks, such as identifying different environmental sounds or diagnosing mechanical failures through sound analysis, augmentation helps models discern subtle audio cues in complex auditory landscapes. Music Analysis: In music analysis, augmentation prepares models to analyze and categorize music across genres, instruments, and recording conditions, enhancing their applicability in the music industry. Audio data augmentation, with its array of techniques, serves as a critical enabler in the field of audio-based ML. By introducing realistic and varied acoustic conditions, these techniques prepare models to perform effectively and accurately in real-world scenarios. The advancement and application of audio data augmentation are integral to developing sophisticated, context-aware models capable of navigating the intricate world of sound, making it an indispensable tool in the evolution of audio processing and analysis.

5.4 Future Directions: Towards Dynamic and Adaptive Augmentation

As the landscape of ML continues to evolve, the future of dataset augmentation is increasingly moving towards more dynamic and adaptive approaches. This paradigm shift from static, pre-defined augmentation strategies to more flexible and responsive methods is poised to significantly enhancement of ML models efficiency. In this context, dynamic and adaptive augmentation represents an exciting frontier, with the potential to revolutionize how models are trained and how they adapt to new and evolving data environments. Traditional dataset augmentation has largely been static, involving a predetermined set of transformations applied uniformly across the dataset. However, the traditional approach, while beneficial, often fails to account for the nuances and specific requirements of individual data instances or changing data distributions. Dynamic augmentation is a more tailored approach, where augmentation strategies are adjusted in real-time, responding to the model's learning progress and the specific characteristics of each data instance.

This dynamic approach allows for a more nuanced and effective training process. For instance, a model struggling to recognize a particular feature in an image might receive more augmented examples of that feature, facilitating focused learning where it is needed. Adaptive augmentation takes this concept further, integrating the augmentation process with the learning model itself. In this scenario, the model actively participates in its own training process, determining how and when to augment data based on its current performance and learning objectives. This can be achieved through reinforcement learning algorithms or other feedback mechanisms where the model's performance on certain tasks informs the augmentation strategy.

Such an approach ensures that augmentation is not just a pre-processing step but an integral part of the learning process, continuously adapting to the model's evolving needs. It allows the model to be exposed to a broader and more relevant range of variations, enhancing its resilience and ability to generalize to new, unseen data. Dynamic and adaptive augmentation strategies have the potential to significantly enhance model resilience. By continuously adjusting to the model's learning trajectory and the ever-changing data landscape, these strategies ensure that models are not only trained on a diverse dataset but also remain adaptable to new situations.

This adaptability is particularly crucial in fields like healthcare, autonomous vehicles, or financial forecasting, where models must perform accurately in the face of unpredictable and evolving data. Adaptive augmentation ensures that models are better equipped to handle real-world variability and complexity. The future directions in dataset augmentation, marked by dynamic and adaptive approaches, hold significant promise for the field of ML. By moving beyond static augmentation strategies to more flexible and responsive methods bring potential for advancements in how models are trained and how they adapt to new data environments. This evolution in augmentation techniques is not just about enhancing model performance; it is about fostering models that are more intelligent, resilient, and adaptable. As ML continues to advance and permeate life and industry, the role of adaptive and dynamic dataset augmentation will be pivotal in shaping more robust and versatile AI systems.

6 Proposed Methodology

This section describes main contribution of the thesis. It focuses on two main aspects:

1. Improving classification accuracy by using augmentation of the dataset for the case when amount of real dataset is very limited.
2. Automatize the verification of manual review after dataset augmentation

The first methodology consists of four distinct methods (M1, M2, M3 and M4), each developed to enhance classification accuracy in situations with limited training data.

The second methodology is an image quality assessment method through latent space analysis of AEs. This novel aspect aims to automate and systematize the traditionally manual review process, ensuring dataset integrity and quality. Empirical validation using benchmark datasets like MNIST demonstrates the effectiveness of these methodologies in achieving accurate classification results and enhancing image quality assessment.

6.1 Advancing image classification algorithms with GANs in the context of severe dataset scarcity

In ML and computer vision, especially when there's not much data available. Image classification is about assigning labels to images, but it often faces problems like overfitting (model is too complex) and underfitting (model is too simple) due to limited data. Traditional ways to solve this use a lot of resources and include making more data, learning from existing models, and combining different methods. Unlike traditional methods that increase dataset size, GANs generate data distributions for each class, enhancing classification of unknown elements by comparing them to these distributions. This methodology introduces four methods for estimating conditional probabilities for category belonging, specifically designed for scenarios with limited data.

6.1.1 Common setup

The main idea of proposed methods of using GANs in a classification task is as follows. Suppose we classify vectors $x \in X$ into k disjoint classes $C_1, C_2, \ldots, C_k$. First, we will create a neural network with GAN architecture, *see Section 4.2*. We then train this network k times on the data of

individual classes C_i. We thus obtain k trained networks GAN_i , $i = 1, \dots, k$ that simulate probability distributions of individual classes $P_{C_i}(x)$. We then use the discriminators of the trained networks GAN_i for classification by inserting the classified vectors $\overline{x}$ into the input of the discriminator and treating the output of the discriminator as a probability estimate $P(\overline{x} \epsilon C_i)$. This is the basic way to proceed.

There is another way. When classifying the vector x , we first approximate individual probability distributions $P_{C_i}(x)$ by generating a certain number of their realizations and then comparing the vector $\overline{x}$ with these approximations. In this work, we have proposed four M1, M2, M3 and M4 methods by which this comparison can be made. **Figure 11** illustrates the training process of the GAN network for a specific category C_i .

Figure 11 Training of the network C_i

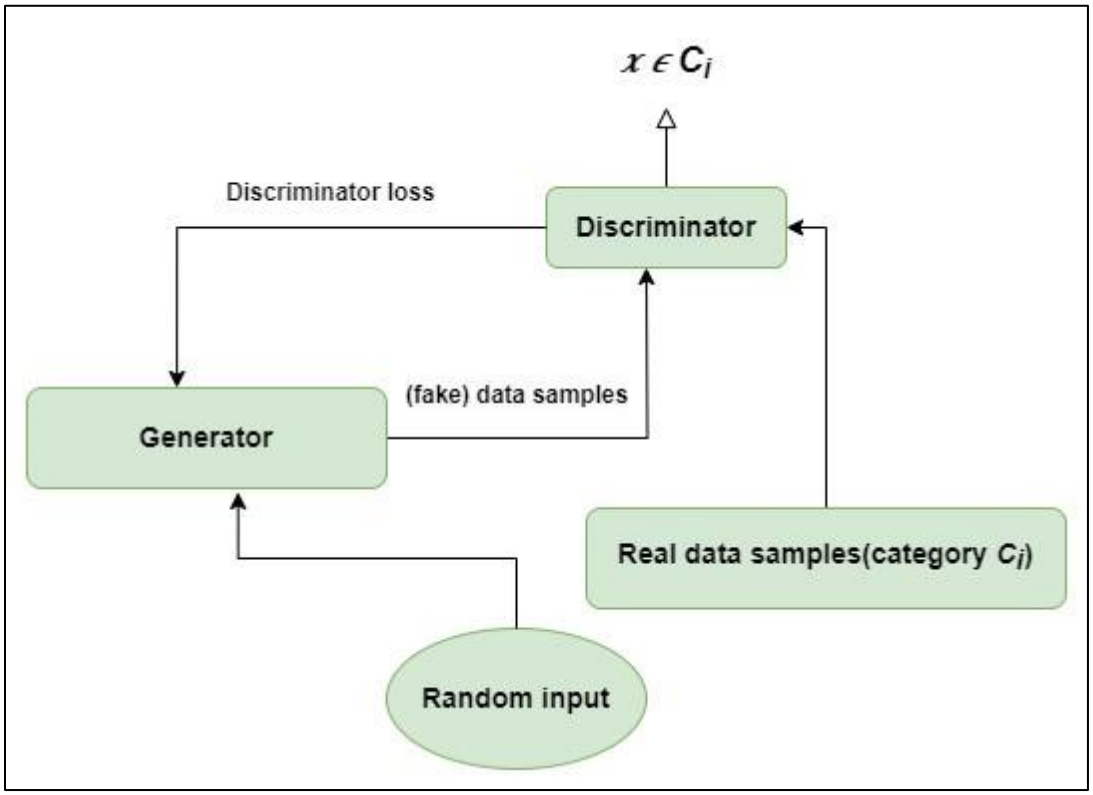

Source: Author's work

6.1.2 Method M1

This method involves the use of GANs to establish neural networks that simulate probability distributions for each category, essential for classification.

The method classifies an unknown vector $\overline{x}$ in the following way:

1) For each distribution $P_{C_i}(x),\ i = 1, \dots, k$ the method
    a. Generates vectors $x_1,\ \dots,\ x_m$ with the help of pre-trained GAN_i for category C_i.
    b. Determines distance values $d_1 = |\overline{x} - x_1|,\ \dots,\ d_m = |\overline{x} - x_m|$
    c. Sorts the values $d_1,\ \dots,\ d_m$ in ascending order.
    d. Determines the distance of the unknown element $\overline{x}$ from the category C_i as follows: $d(\overline{x}, C_i) = \frac{1}{L}\sum_{i=1}^{L} d_i$, where L is an optional method parameter that specifies the number of closest generated vectors to include when calculating the distance.
    e. Finally classifies the unknown vector $\overline{x}$ into that category j for which the value $d(\overline{x}, C_j),\ j = 1,\ \dots,\ k$ is the smallest.
2) Step 1) N times repeated. Let us denote N_i be a number of times $\overline{x}$ was classified as belong to category C_i then $N = N_1 + \cdots + N_k$.
3) The probability for each class $C_i,\ i = 1, \dots, k$ is then:

$$P(C_i) = \frac{N_i}{N}$$

Figure 12 illustrates M1 method, simplifying its core principles for better understanding.

Figure 12 Illustration of the M1 method concept

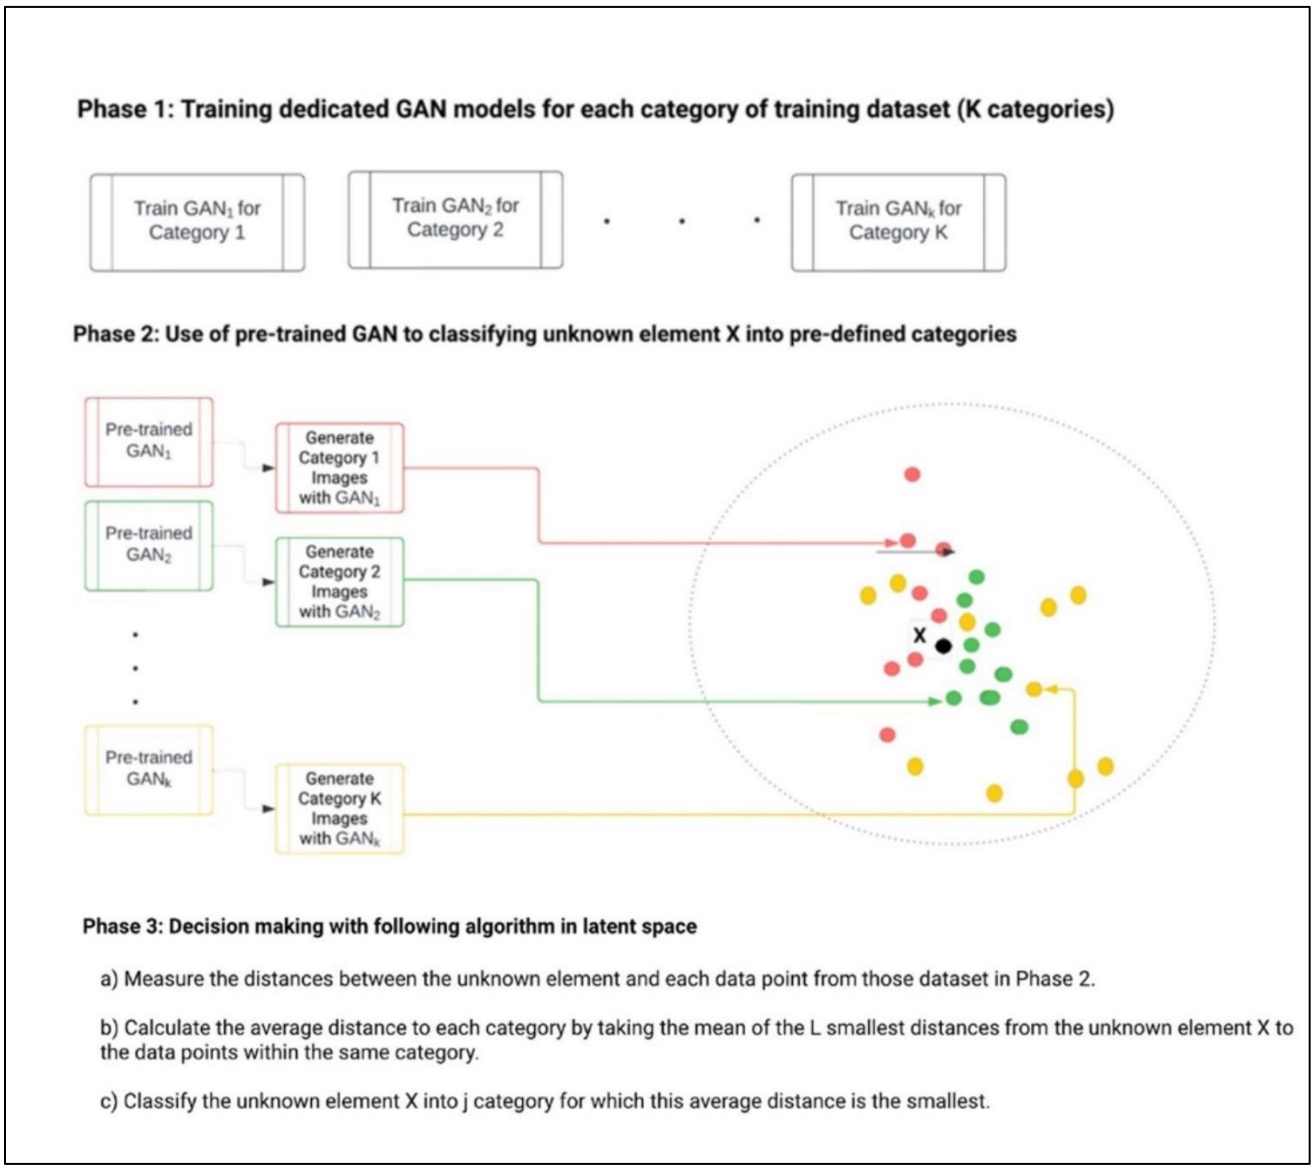

Source: Author's work

Pseudocode 1 describes main steps of the M1 model. It details the use of GANs for vector generation, distance-based comparisons, iterative classification, and probability-based decision-making:

Pseudocode 1 Method M1 for classification of an unknown element **x**

```
Inputs:
   k: number_of_classes
   m: vectors_per_category
   L: parameter_for_average_distance
   N: classification_iterations
   x_bar: unknown_vector

Outputs:
   category_probability: list of probabilities for each class

# Initialization
category_counts = array of size k, initialized to 0

# Classification Loop (Repeats N times)
for iteration = 1 to N:
    for category_index = 1 to k:
        # Generate vectors for the current category
        generated_vectors = trained_GAN[category_index].generate(m)

        # Calculate distances to the unknown vector
        distances = []
        for vector in generated_vectors:
            distances.append(calculate_distance(vector, x_bar))

        # Determine average distance using L smallest distances
        distances.sort()
        average_distance = mean(distances[0:L])

        # Identify the category with the smallest average distance (for this iteration)
        closest_category_index = argmin(average_distance)

        # Increment the count for the closest category
        category_counts[closest_category_index] += 1

# Calculate probabilities
probabilities = []
for category_index = 1 to k:
    probabilities.append(category_counts[category_index] / N)

# Return results
return probabilities

# End Method M1
```

6.1.3 Method M2

This method integrates the latent space of a VAE into the classification framework, enriching the analytical depth. The VAE, trained in conjunction with GAN networks, improves the decision-making process, enabling more nuanced classifications.

It uses the distributions $P_{C_1}(x),\ \ldots,\ P_{C_k}(x)$ to train the VAE, see *Section 4.3.3*, so that it is able, after inserting a vector $\overline{x}$ from any category $C_1,\ \ldots, C_k$ into its input to reproduce it on its output. The autoencoder can thus simulate the probability distribution.

$$P_C(x),\ C = C_1 \cup C_2 \cup \ldots \cup C_k,\ C_i \cap C_j = \emptyset\,, \qquad i,j = 1,\ldots,k$$

Note: we assume that X belong to one C_i

We denote the latent space of the autoencoder by Z. The learned autoencoder first transforms the input vector $\overline{x}$ to $\overline{z}\ \epsilon\ Z$ and then transforms $\overline{z}$ into the output vector $\overline{y} \triangleq \overline{x}$ (see **Figure 13**).

Figure 13 Variational autoencoder that learns $P_{C_i}(x)$

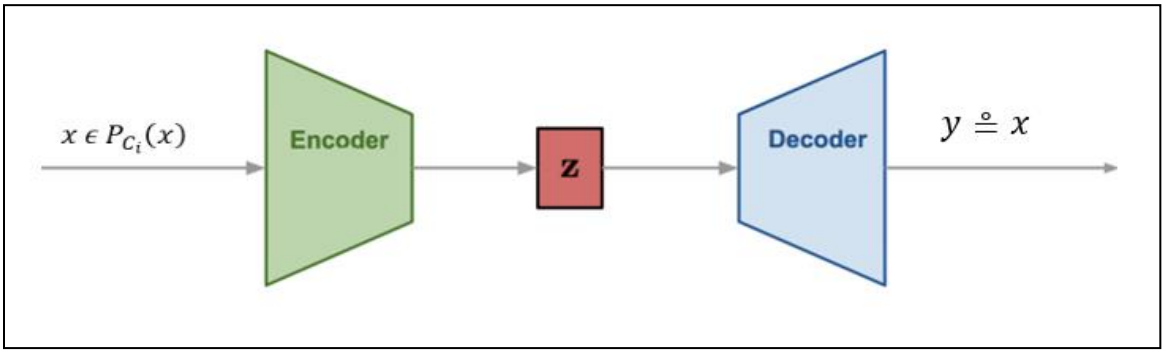

Source: Author's work

The method uses method M1 above for classification, but operate an the latest space Z of the VAE. It classifies an unknown vector $\overline{x}$ in the following way:

1) For each distribution $P_{C_i}(x)$: the algorithm:
    a. Generates vectors $x_1,\ \ldots,\ x_m$ with the help of pre-trained GAN_i for category C_i.
    b. Determines distance values $d_1 = |\overline{z} - z_1|,\ \ldots,\ d_m = |\overline{z} - z_m|$, where $\overline{z}$ is the projection of the classified vector $\overline{x}$ into the latent space Z.
    c. Sorts distances $d_1,\ \ldots,\ d_m$ in ascending order.
    d. Determines the distance of the unknown element $\overline{x}$ from the category C_i as follows:

$$d(\overline{x}, C_i) = \frac{1}{L}\sum_{i=1}^{L} d_i\,,$$

where L is an optional method parameter that specifies the number of closest generated vectors to include when calculating the distance.

e. Finally classifies the unknown vector $\overline{x}$ into the category j for which the value $d(\overline{x}, C_j), j = 1, \ldots, k$ is the smallest.

2) Step 1) N times repeated. Let us denote N_i be a number of times $\overline{x}$ was classified as belong to category by C_i then $N = N_1 + \cdots + N_k$.

3) The probability for each class C_i, $i = 1, \ldots, k$ is then:

$$P(C_i) = \frac{N_i}{N}$$

Figure 14 illustrates M2 method, simplifying its core principles for better understanding.

Pseudocode 2 describes main steps of the M1 model. It details the use of GANs for vector generation, distance-based comparisons, iterative classification, and probability-based decision-making:

Pseudocode 2 Method M2 for classification of an unknown element **x**

```
Inputs:
   k: number_of_classes
   m: vectors_per_category
   L: parameter_for_average_distance
   Z: latent_space
   x_bar: unknown_vector
   VAE: trained_variational_autoencoder

Outputs:
   category_probability: list of probabilities for each class

# Project x_bar into latent space
z_bar = VAE.encode(x_bar)

# Classification in Latent Space (Utilizing Method M1)

# Initialize
category_counts = array of size k, initialized to 0

# Classification Loop (Repeats N times, defined in Method M1)
for iteration = 1 to N:
```

```
    for category_index = 1 to k:
        # Generate vectors for the current category
        generated_vectors = trained_GAN[category_index].generate(m)

        # Project vectors into latent space
        latent_vectors = []
        for vector in generated_vectors:
            latent_vectors.append(VAE.encode(vector))

        # Calculate distances in latent space
        distances = []
        for z_j in latent_vectors:
            distances.append(calculate_distance(z_bar, z_j))

        # Determine average distance using L smallest distances
        distances.sort()
        average_distance = mean(distances[0:L])

        # Find the closest category in latent space
        closest_category_index = argmin(average_distance)

        # Increment the count for the closest category
        category_counts[closest_category_index] += 1

# Calculate probabilities (same as in Method M1)
probabilities = []
for category_index = 1 to k:
    probabilities.append(category_counts[category_index] / N)

# Return results
return probabilities

End Method M2
```

Figure 14 Illustration of the M2 method concept

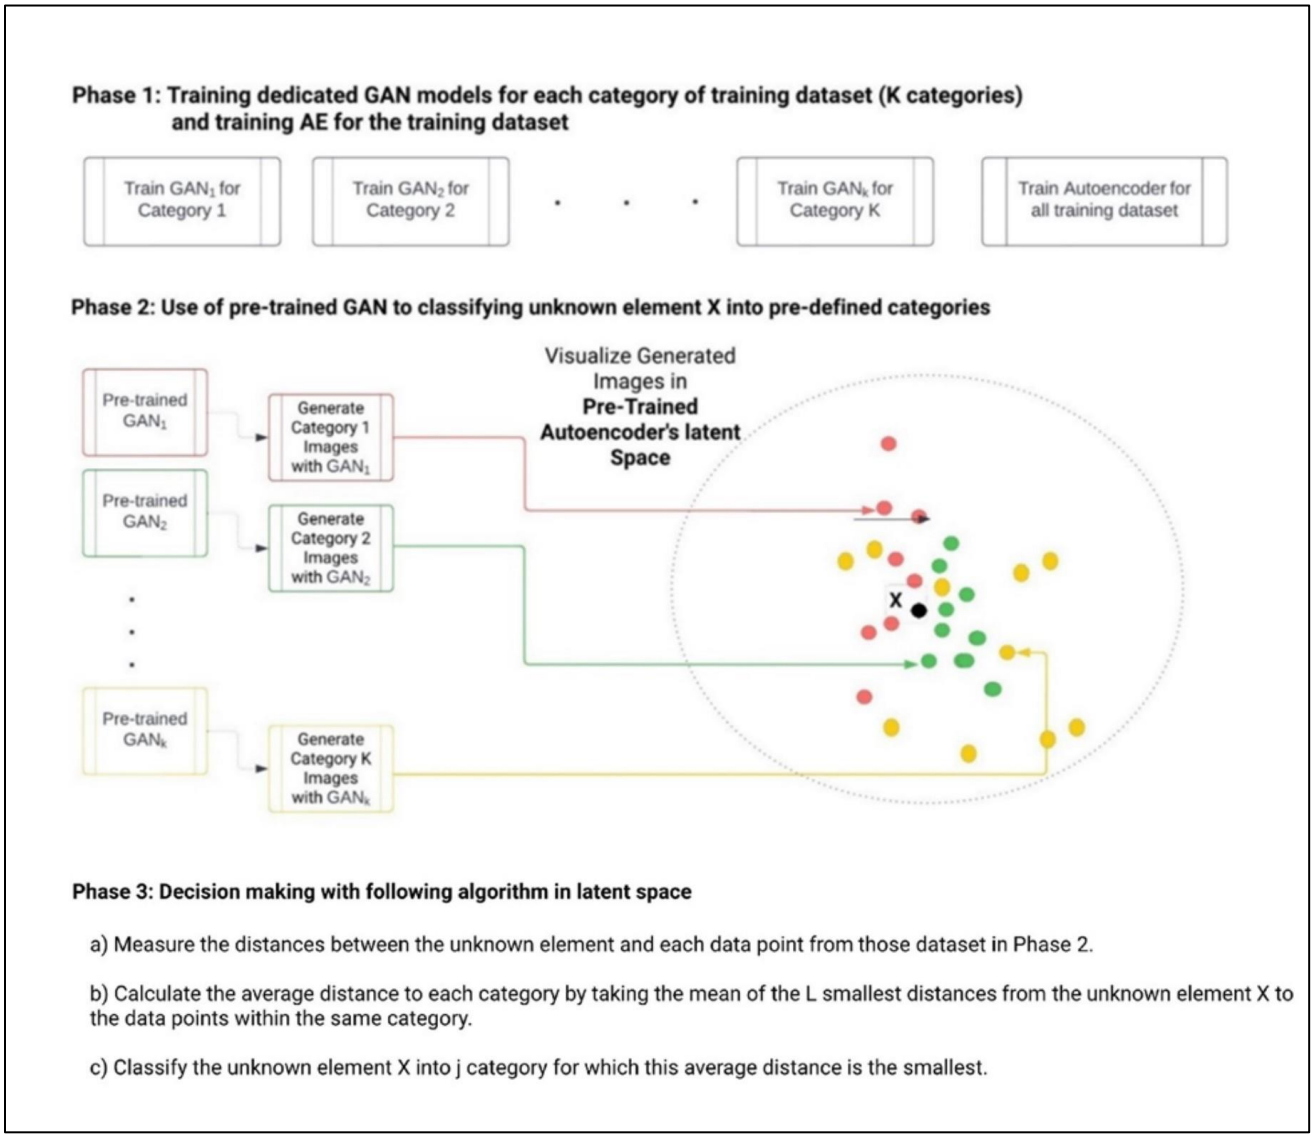

Source: Author's work

6.1.4 Method M3

This method improves classification and contributes to a deeper understanding of underlying data structures. The method estimates the probabilities $P(x \in \varepsilon(\overline{x}) \mid x \in C_i),\ i = 1, \ldots, k,$ which indicates whether x belong to the ε neighborhood of $\overline{x}$. The neighborhood $\varepsilon(\overline{x})$ is defined as:

$$x \in \varepsilon(\overline{x}) \iff |x - \overline{x}| < \varepsilon.$$

where $\varepsilon(\overline{x})$ is ε - neighborhood of $\overline{x}$, $\varepsilon > 0$, $x \epsilon X$

Suppose we generate for each i^{th} class m vectors by using distributions $P_{C_1}(x)$, ..., $P_{C_k}(x)$, $x \epsilon X$. Let the n_i vectors from category C_i, $i = 1, \ldots, k$ fall into $\varepsilon(x)$. It is possible to then estimate the probabilities $P(x \epsilon \varepsilon(\overline{x}) \mid x \epsilon C_i)$, $i = 1, \ldots, k$ as follows:

$$P(x \epsilon \varepsilon(\overline{x}) \mid x \epsilon C_i) = \frac{n_i}{m},\ n = \sum_{i=1}^{k} n_i$$

Applying Bayes's formula to the last expression, we can get:

$$P\left(x \epsilon C_i \mid x \epsilon \varepsilon(\overline{x})\right) = \frac{P(x \epsilon \varepsilon(\overline{x}) \mid x \epsilon C_i)\, P(x \epsilon C_i)}{P\left(x \epsilon \varepsilon(\overline{x})\right)}$$

Since we generated the same number of elements from each class, $P(x \epsilon C_i) = \frac{1}{k}$ as n vectors fall into the ε - neighborhood of $\overline{x}$ we can set $P\left(x \epsilon \varepsilon(\overline{x})\right) = \frac{n}{m \cdot k}$. The additional problem arises of how to choose the size of $\varepsilon(\overline{x})$ so that an appropriate number of generated vectors falls into it. For example, we want that probability $P\left(x \epsilon \varepsilon(\overline{x})\right)$ equals μ, $\mu > 0$, μ <=1.
It is possible to set the following algorithm:

1) For each class, generate m vectors x_i, $i = 1, \ldots, m$.
2) We determine the distances between generated vectors x_i from classified unknown $\overline{x}$, $d_i = |\overline{x} - x_i|$, and sort them in ascending order.
3) Define the neighborhood as $\varepsilon = |\overline{x} - x_n|$, where $n = \mu \cdot m \cdot k$

In M3 ε plays a role of classification tolerance.

Figure 15 illustrates M3 method, simplifying its core principles for better understanding. The pseudocode for the M1 (**Pseudocode 1**) and M3 methods share a similar structure, with the primary difference being the specific calculation performed.

Figure 15 Illustration of the M3 method concept

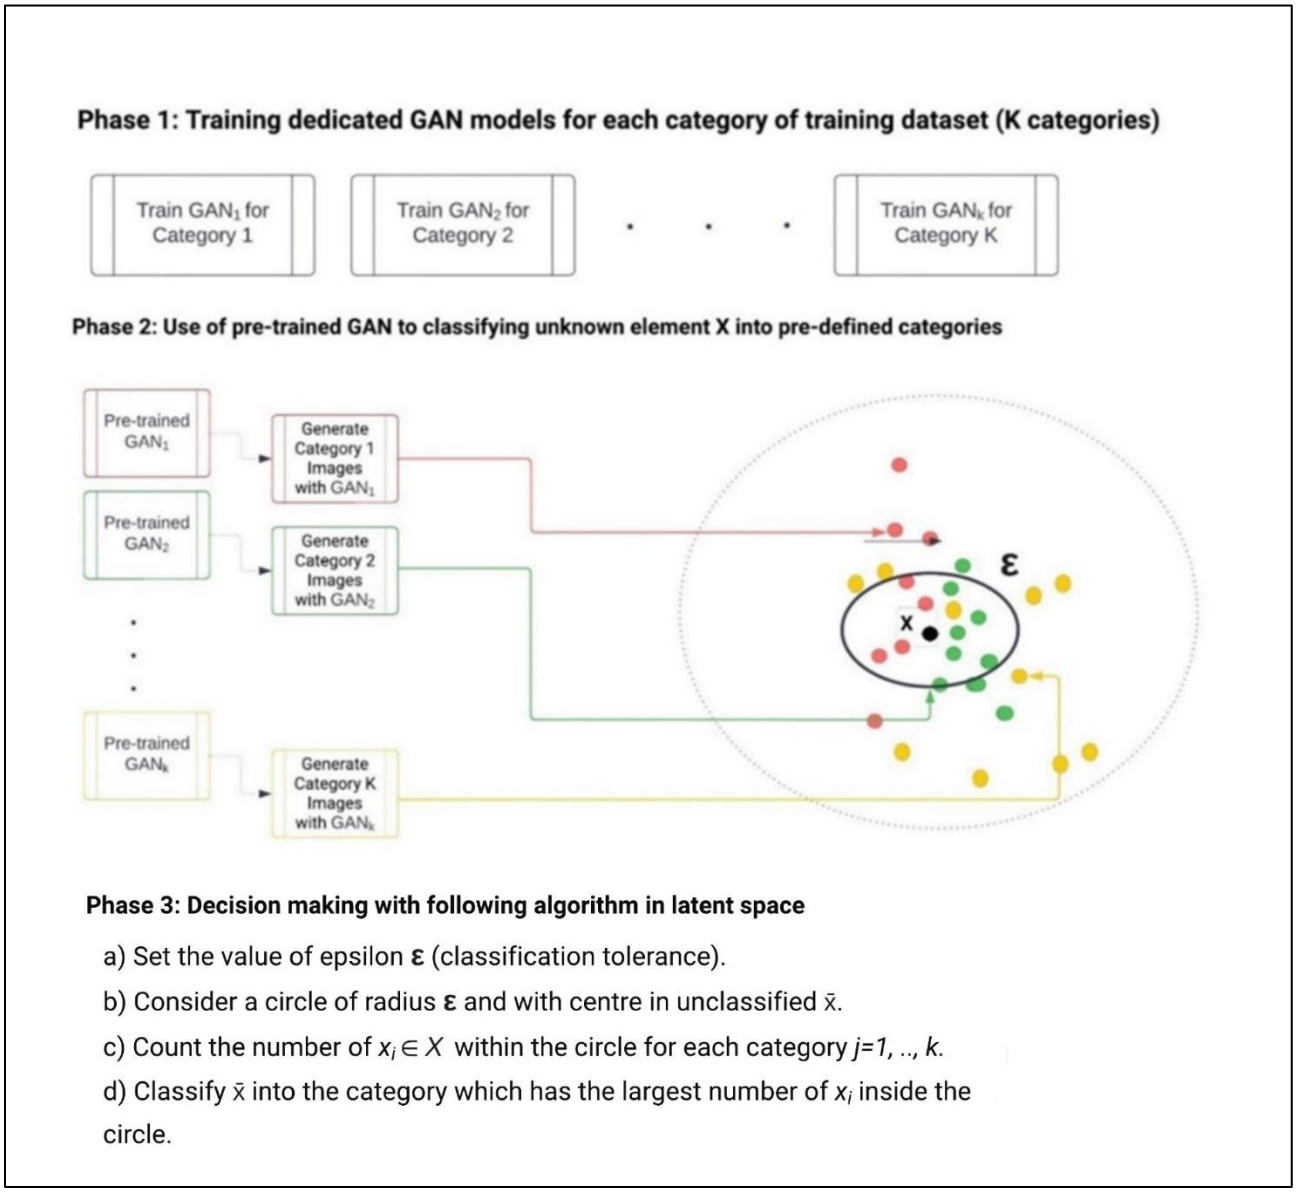

Source: Author's work

6.1.5 Method M4

This method extends latent space analysis within the VAE framework, combining the strengths of GAN networks and VAEs to significantly enhance classification accuracy, especially in limited data scenarios.

The method is combination of M2 and M3. All calculations will be done in the latent space of the VAE.

Figure 16 illustrates M4 method, simplifying its core principles for better understanding. The pseudocode for the M2 (**Pseudocode 2**) and M4 methods share a similar structure, with the primary difference being the specific calculation performed.

Figure 16 Illustration of the M4 method concept

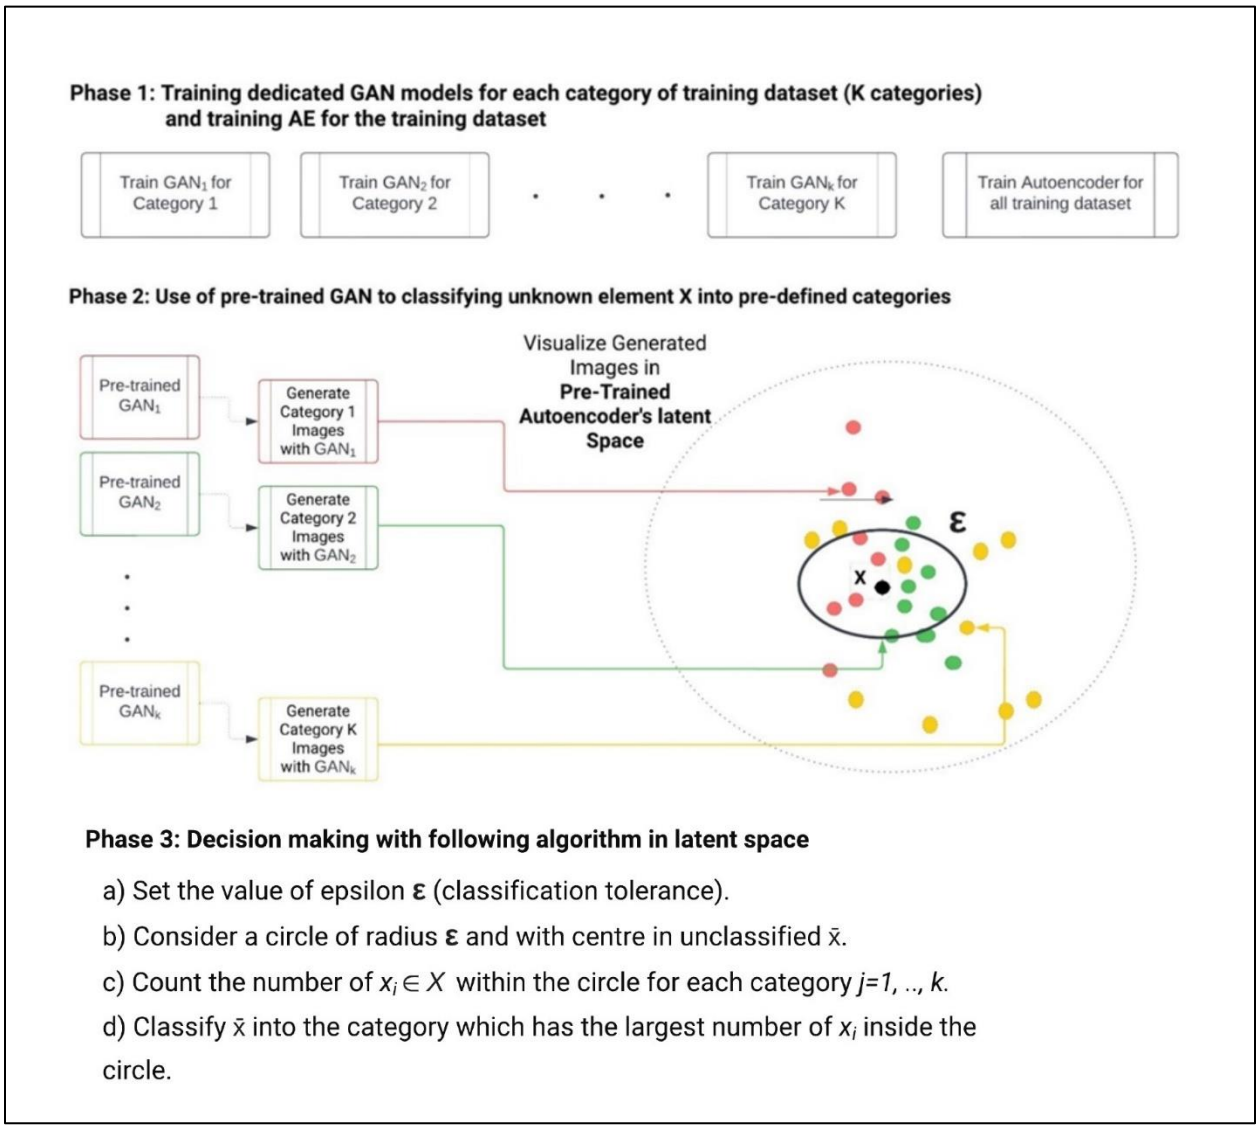

Source: Author's work

6.2 Automatize validation of augmented dataset quality

By exploring the latent space, the proposed methodology allows for a deeper understanding of intricate patterns and nuances that conventional methods often overlook. The core of the approach lies in leveraging the latent space of AEs as a tool for assessing image quality. This enables the identification of patterns, variations, and outliers that indicate the quality of an image, providing an objective and systematic measure of image quality. The approach marks a departure from the passive use of AEs, transforming their latent space into an active participant in the image quality determination process.

This section is devoted to optimizing quality assessment during manual augmentation of the dataset. Since the quality assurance of data augmentation is performed manually, the whole process is rather expensive and time consuming. As validating of the dataset before analysis needs a high quality domain experiment, the whole process is dependent on the experiment. Thus the thesis propose an efficient approach that will help to automatically validate the augmented data and to "filter out" bad data.

The idea is reflected in the provided **Pseudocode 3**. The description below summaries main process and provides references to the pseudocode (**Pseudocode 3**).

1. **Initial Classification Attempt (1.1):** Begins with selecting a small set of training datasets for initial classification using DL. If results are inadequate, the process progresses to the next phase.

2. **Image Augmentation and GANs (1.2)**: If initial attempts are insufficient, the focus shifts to image augmentation and GANs to enrich the datasets, aiming to improve classification accuracy.

3. **Quality Determination through Autoencoder's Latent Space (1.3):** If earlier steps are unsuccessful, the methodology advances to its innovative core—evaluating image quality using the latent space of an autoencoder.

4. **Autoencoder Training (1.3.1):** Involves training the autoencoder on all datasets to encode images and reveal the latent space.

5. **Representation of A Class in Latent Space (1.3.2):** Images of the class are represented in the latent space, setting the stage for detailed analysis.

6. **Sphere Hypothesis and Radius Computation (1.3.3):** For the class, A sphere is created in latent space to filter out a set percentage of data points, defining a threshold for image quality.

7. **Image Synthesis and Latent Space Imaging (1.3.4):** New images are synthesized and assessed based on their position in the latent space relative to the sphere.

8. **Sphere-Qualified Image Curation (1.3.5):** Images within the sphere's criteria are curated for training, while those outsides are excluded.

9. **Iterative Process Across Classes (1.3.6):** The process is repeated for each class, ensuring a thorough evaluation of image quality across the dataset.

Pseudocode 3 provides a concrete implementation of the detailed image classification methodology outlined earlier. This code translates the conceptual steps into executable instructions, demonstrating how the methodology can be applied in practice.

Pseudocode 3 Provides a concrete implementation of *the Section 6.2* methodology

```
# Initial Classification Attempt (1.1)
dataset = select_small_training_set()
model = train_DL_classifier(dataset)
results = evaluate(model)

if results.accuracy < satisfactory_threshold:
    proceed_to_augmentation = True
else:
    proceed_to_augmentation = False

# Image Augmentation and GANs (1.2)
if proceed_to_augmentation:
    augmented_dataset = apply_augmentation_techniques(dataset)
    if GANs_applicable:  # If GANs are a suitable fit for the problem
        augmented_dataset = augmented_dataset + generate_images_with_GANs(dataset)

    model = train_DL_classifier(augmented_dataset)
    results = evaluate(model)

    if results.accuracy < satisfactory_threshold:
        proceed_to_latent_space = True
    else:
        proceed_to_latent_space = False

# Quality Determination through Autoencoder's Latent Space (1.3)
if proceed_to_latent_space:
    # Autoencoder Training (1.3.1)
    autoencoder = train_autoencoder(dataset)

    # Representation of A Class in Latent Space (1.3.2)
    for image in dataset:
        latent_representation = autoencoder.encode(image)

    # Sphere Hypothesis and Radius Computation (1.3.3)
    sphere_center,      sphere_radius      =      calculate_sphere(latent_representation,
quality_percentage)

    # Image Synthesis and Latent Space Imaging (1.3.4)
    while model.accuracy < satisfactory_threshold:
        new_images = synthesize_images()  # Placeholder, method unspecified
        for image in new_images:
            latent_representation = autoencoder.encode(image)
            if within_sphere(latent_representation, sphere_center, sphere_radius):
                dataset = dataset + image

    # Sphere-Qualified Image Curation (1.3.5)
    model = train_DL_classifier(dataset)
    evaluate(model)  # Final evaluation
```

Figure 17 provides a conceptual illustration of Methodology 2, simplifying its core principles for better understanding.

Figure 17 Conceptual illustration of the second methodology

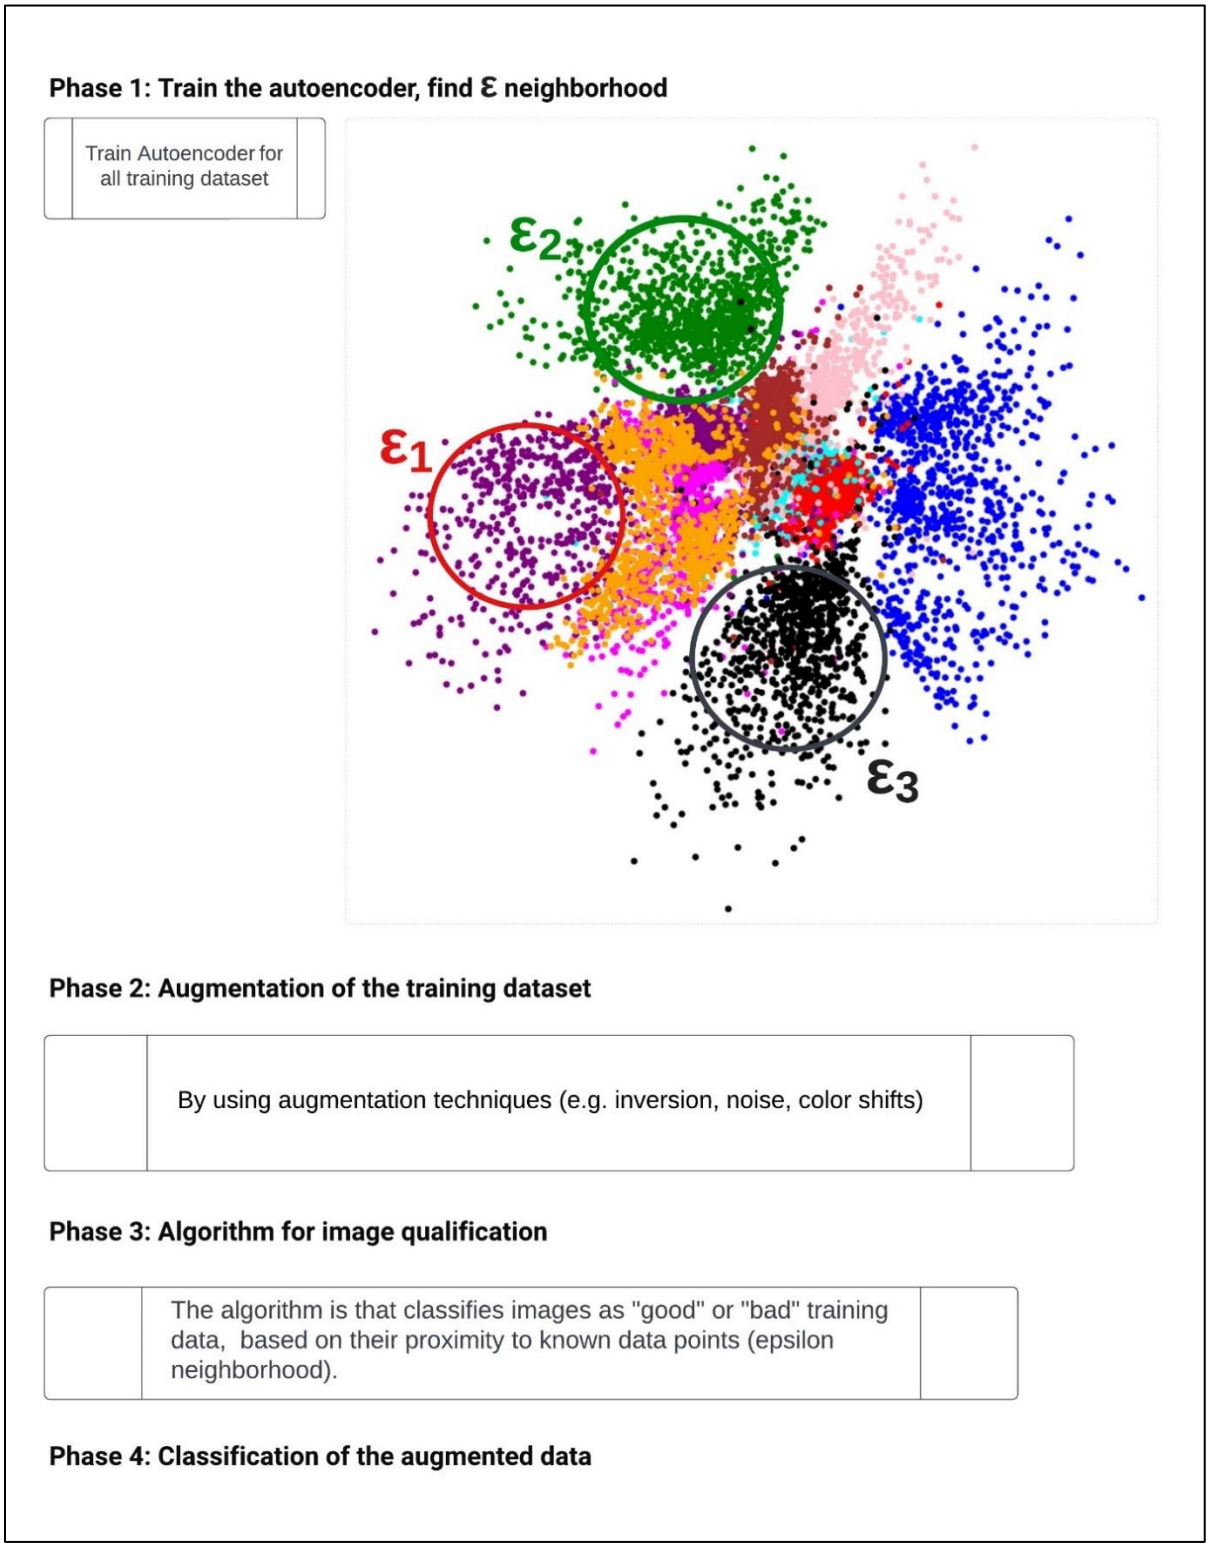

Source: Author's work

7 Data

This research utilizes two fundamental datasets in computer vision: the MNIST handwritten digits database and the CIFAR datasets (CIFAR-10 and CIFAR-100 - https://www.cs.toronto.edu/~kriz/cifar.html). These datasets were selected for their varying levels of complexity and diversity, offering a comprehensive testing ground for evaluating the data augmentation methods developed in this study.

Figure 18 Sample image from the MNIST dataset

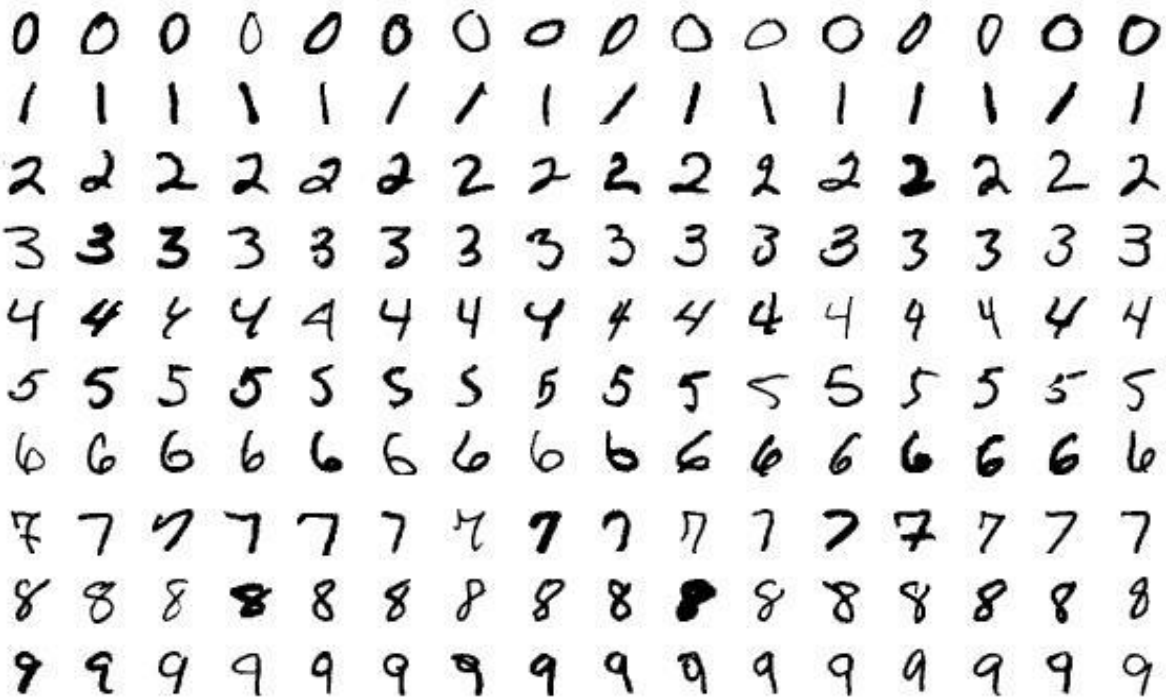

Source: Author's work

7.1 MNIST Database

The MNIST database (http://yann.lecun.com/exdb/mnist/), a cornerstone in image processing research, consists of handwritten digits from 0 to 9, amounting to ten distinct classes. The dataset includes simple, grayscale images of size 28x28 pixels. A typical image from MNIST (**Figure 18**)

showcases a single digit centered in the frame, offering a straightforward yet effective platform for initial algorithm testing.

7.2 CIFAR Datasets

The CIFAR datasets (https://www.cs.toronto.edu/~kriz/cifar.html) present a more intricate challenge. CIFAR-10 contains 60,000 color images in 10 classes, with 6,000 images per class. Each 32x32 pixel image in CIFAR-10 (**Figure 19**) features objects like animals and vehicles, portraying a more diverse and realistic scenario for image classification.

Figure 19 CIFAR-10 dataset sample

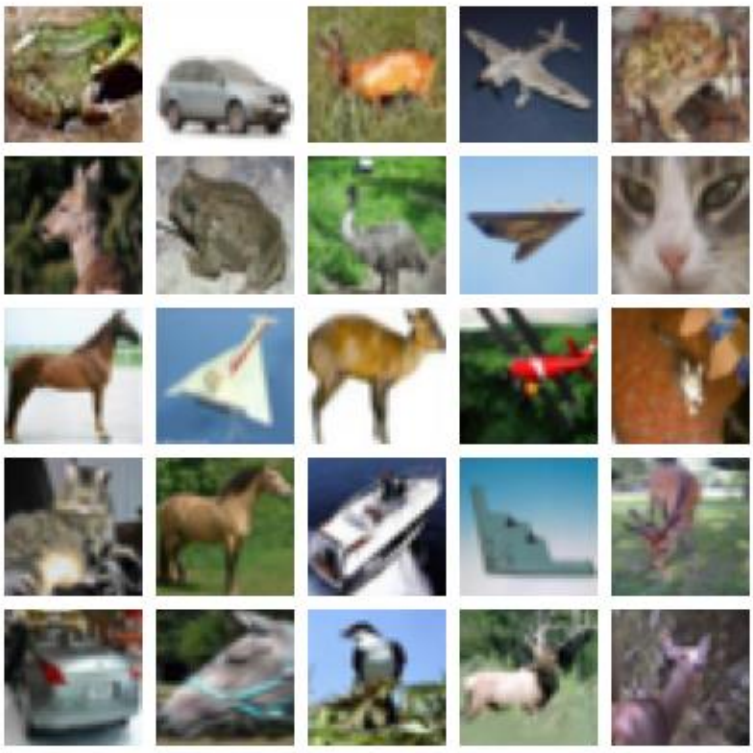

Source: Author's work

CIFAR-100, similar in format to CIFAR-10, includes 100 classes with 600 images each. A sample image from CIFAR-100 (**Figure 20**) typically contains more detailed and varied objects, reflecting the complexity of natural scenes and the challenges in classifying them.

Figure 20 CIFAR-100 dataset sample

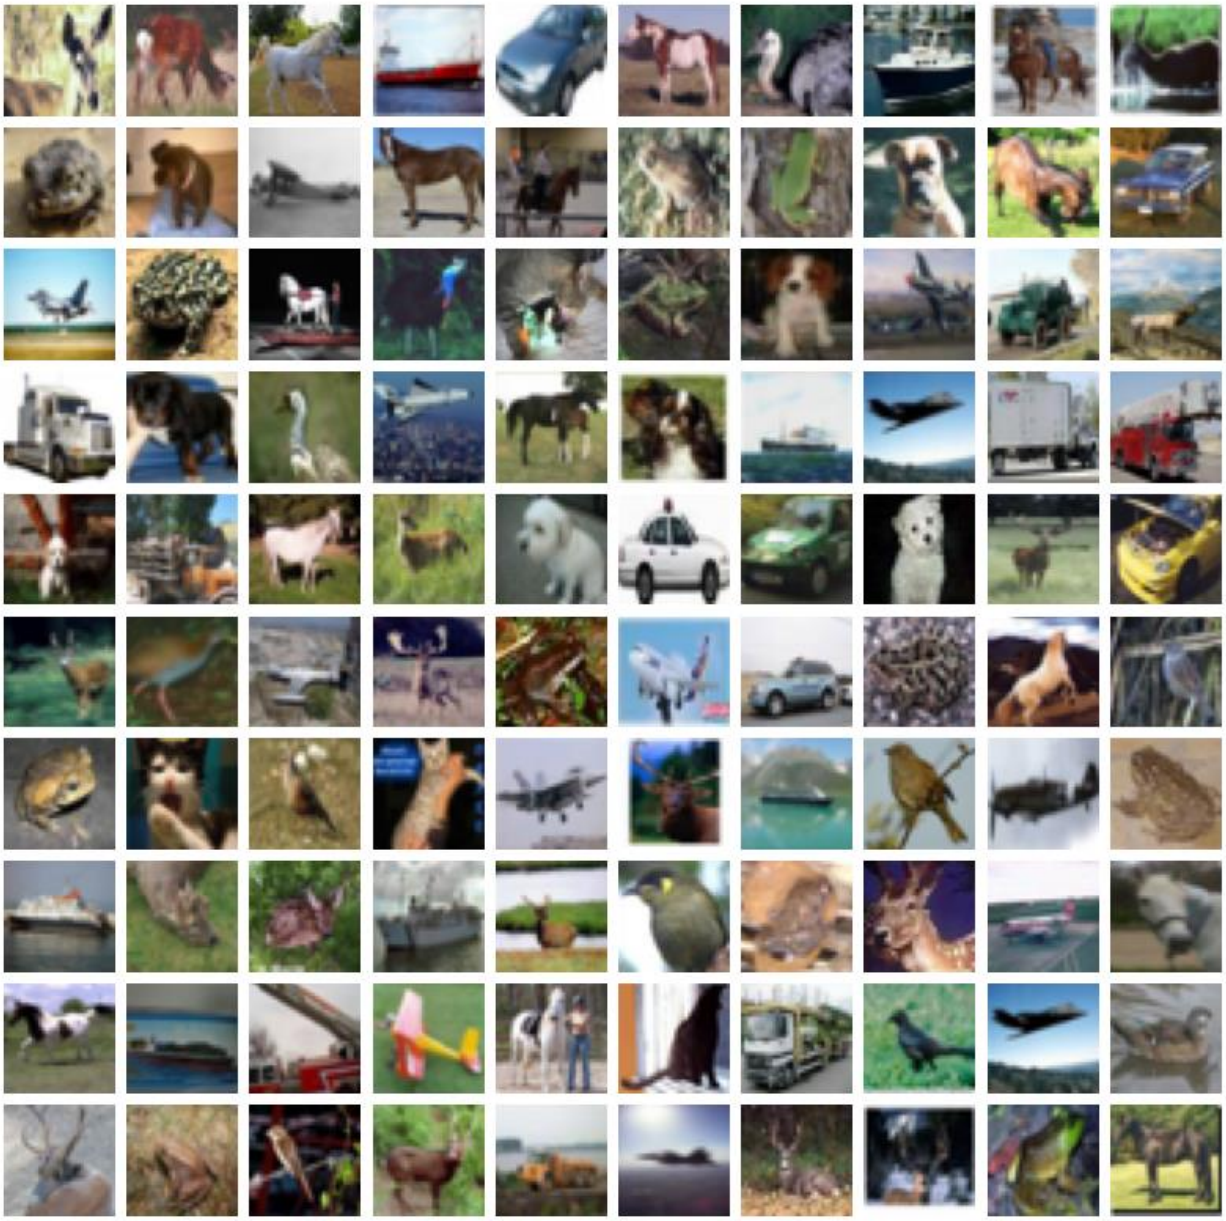

Source: Author's work

To prepare the datasets for experiments, both MNIST and CIFAR were shuffled to ensure a randomized distribution. A subset of images was then selected for training, varying in size from 1 to 100 samples per category. This method allowed for an examination of the augmentation methods under different conditions of data availability.

A consistent set of 100 testing data samples per category was used for both datasets. This standardized testing protocol ensures that the performance of the proposed augmentation methods can be fairly compared across the different datasets and image classes.

The inclusion of sample images (**Figure 19** and **Figure 20**) in the thesis provides an intuitive understanding of the dataset's characteristics. These samples represent the type of images that the augmentation methods are applied to, highlighting the challenges and requirements of each dataset. The MNIST dataset allows for initial method validation in a simpler environment, while the CIFAR datasets provide a more rigorous and realistic testing environment. The combination of these datasets in the research is crucial to assess the effectiveness, versatility, and scalability of the proposed data augmentation techniques.

8 Results of Experiments

8.1 Experiment results with the proposed methodology (Section 6.1)

This section offers an in-depth and systematic analysis of the experimental results derived from the application of GANs and AEs, with a particular emphasis on VAEs, in the challenging context of limited training data for image classification tasks. The primary aim of these experiments was to critically evaluate the performance and effectiveness of four innovative classification methods, designated as M1, M2, M3, and M4, see Section 6.1.2 - 6.1.5. The experiments were not only structured to assess the efficacy of novel classification methods but also to probe their flexibility and resilience across a variety of data conditions. A significant focus of this research was to explore the potential of these methods to overcome the challenges faced. Traditional DL models often show limitations when confronted with insufficient training data – a scenario increasingly common in various real-world applications. By addressing this gap, the experiments aimed to make an important contribution to the field of ML, especially in contexts where acquiring extensive datasets is not feasible.

The four classification methods, namely M1, M2, M3, and M4, *see Section 6*, were designed to encapsulate and utilize the strengths of GANs and VAEs. M1 and M2 focuses on exploiting the generative and reconstructive aspects of GANs and VAEs to enhance the quality and diversity of the training data. M3 and M4, concentrate on refining the classification accuracy through advanced manipulation of the latent space and feature extraction capabilities of VAEs.
The research tailors these methods specifically to operate efficiently under conditions of data scarcity represent a significant contribution. These methods offer promising avenues for tackling one of the most pressing challenges – the reliance on large quantities of data for training effective ML models. The design of the experiments was comprehensive, encompassing a range of scenarios from extremely limited to moderately sufficient training data. This allowed for a thorough evaluation of methods performance across different levels of data availability.

The choice of using the MNIST dataset for these experiments served as an appropriate benchmark, given its widespread use and recognition in the field of image classification. The dataset's

simplicity and clarity allowed for a focused analysis of the methods' performance, minimizing external variables that could influence the outcomes.

8.1.1 Overview of Models Used in the Experiments

This research's experimental strategy involved a comprehensive exploration of a wide range of DL neural network models. This section details the specific models that were central to obtaining the results presented in ***Table 1***, as well as provides an overview of the additional models that were part of the extensive testing regime. VAEs played a central role in the experiments, especially in analyzing and manipulating the latent space of data. Two specific VAE models were used:

VAE Encoder Model (**NM1**): Built with a cutting-edge CNN architecture, this model efficiently encoded input images into a compact latent space. It started with an input layer for 28x28 pixel images, followed by convolutional layers for feature extraction and dimensionality reduction. The encoder was key to compressing data into the latent space, essential for effective data management and classification.

NM1: Model 1 VAE encoder model

Layer(type)	Output Shape	Param#	Connected to
encoder_input (InputLayer)	[(None,28,28,1)]	0	[]
conv2d (Conv2D)	(None,14,14,512)	5120	['encoder_input[0][0]']
conv2d_1 (Conv2D)	(None,7,7,1024)	4719616	['conv2d[0][0]']
Flatten (Flatten)	(None,50176)	0	['conv2d_1[0][0]']
Dense (Dense)	(None,400)	20070800	['flatten[0][0]']
z_mean (Dense)	(None,100)	40100	['dense[0][0]']
z_log_var (Dense)	(None,100)	40100	['dense[0][0]']
Z (Lambda)	(None,100)	0	['z_mean[0][0]','z_log_var[0][0]']

VAE Decoder Model (**NM2**): Working in tandem with the encoder, the decoder model reconstructed images from their latent representations. This process was vital for assessing how well the latent space preserved the input data's key features. It included a dense layer to expand the latent vector, followed by Convolutional Transpose layers to gradually restore the image to its original size. The decoder's ability to accurately reconstruct images was a measure of the VAE's effectiveness in representing complex data distributions.

NM2: Model 2 VAE decoder model

Layer (type)	Output Shape	Param #
z_sampling (InputLayer)	[(None,100)]	0
dense_1 (Dense)	(None,50176)	5067776
Reshape (Reshape)	(None,7,7,1024)	0
conv2d_transpose (Conv2DTranspose)	(None,14,14,1024)	9438208
conv2d_transpose_1 (Conv2DTranspose)l	(None,28,28,512)	4719104

GANs as a Core Component:

GANs formed a cornerstone of the experimental setup, mainly facilitating the generation of synthetic data instances that closely mimic the real data distribution. The GAN framework comprised two models:

Discriminator Model of the GAN (**NM3**): Designed as a fully connected neural network, the discriminator model played a critical role in differentiating between real and generated images. The model consisted of multiple dense layers with increasing neuron counts, aiming to classify images effectively as either authentic or fabricated by the generator. The final output layer employed sigmoid activation, a standard practice in binary classification tasks within GAN architectures.

NM3: Model 3 Discriminator model of the GAN

Layer (type)	Output Shape	Param #
Input layer (InputLayer)	[(None, 784)]	0
Dense_1 (Dense)	(None, 512)	401920
Dense_2 (Dense)	(None, 256)	131328
Dense_3 (Dense)	(None, 128)	32896
Output_4(Dense)	(None, 1)	129

Generator Model of the GAN (**NM4**): The generator model creating synthetic images that were indistinguishable from real images by the discriminator. Equipped with a series of dense and BatchNormalization layers, this model systematically upscaled a latent input into a complete image. The generator's performance was key to the overall effectiveness of the GAN, as it directly influenced the discriminator's training process and the quality of the generated data.

NM4: Model 4 Generator model of the GAN

Layer (type)	**Output Shape**	**Param #**
Input layer (InputLayer)	(None, 100)	0
Dense_1 (Dense)	(None, 256)	25,856
BatchNormalization_1 (BatchNormalization)	(None, 256)	1,024
Dense_2 (Dense)	(None, 512)	131,584
BatchNormalization_2 (BatchNormalization)	(None, 512)	2,048
Dense_3 (Dense)	(None, 1024)	525,312
BatchNormalization_3 (BatchNormalization)	(None, 1024)	4,096
Output_4 (Dense)	(None, 784)	803,600
Reshape (Reshape)	(None, 28, 28, 1)	0

This experimental method explored a diverse array of DL neural network models to assess and validate the proposed classification methods: M1, M2, M3, and M4. This exploration was vital to ensure the robustness and adaptability of the proposed methods across various architectures and scenarios. While we investigated multiple models, the results detailed in ***Table 1*** were specifically obtained using the following neural network configurations.

8.1.2 Experimental Setup

The experimental framework involves the deployment of ten distinct GAN models, *Section 4.2,* GAN_i, each trained on different digit classes from the MNIST database. The architecture of the GANs included a generator equipped with a 100-dimensional latent input layer and three dense layers the discriminator comprised an input layer for 28x28 pixel images, followed by multiple dense layers. The VAE models, *Section 4.3,* used in the experiments were composed of an encoder and a decoder, each tailored with specific configurations for efficient data transformation and reconstruction.

This setup was designed to provide a comprehensive evaluation of the models' performance in diverse conditions, simulating real-world scenarios where data might be scarce or costly to acquire. The dataset size, ranged from 1 to 100 instances per category.

The models were trained over 30,000 epochs, a duration chosen to optimize model performance due to the limited size of the training data. This extensive training was vital to ensure that each model had enough time to learn from the data, regardless of the dataset size.

The protocol also included specific optimizations for each model to handle limited data sizes effectively, typically requiring 2-3 hours to complete. This efficiency was crucial, given the constraints of computational resources and the need to simulate real-world conditions accurately. Performance of the proposed method was evaluated against the baseline accuracy of the GAN discriminator. The primary metric for assessment was classification accuracy, defined as the ability of the models to correctly identify and classify new instances from the testing dataset. This metric was chosen for its relevance in real-world applications, where the accurate classification of new data is paramount.

8.1.3 Experiment Results

The primary objective was to assess the effectiveness of the proposed methods, *see Section 6.1*, applied to MNIST dataset in utilizing GANs and AEs, particularly VAEs, for image classification tasks. Results are in ***Table 1***.

The decision to center on VAE and GAN models was based on their standout performance and relevance to challenges in data-limited environments. These models showcased not only high classification accuracy but also offered valuable insights into DL's capability to manage sparse data efficiently.

Each experiment employed different quantities of training instances (1 to 100 per category) to replicate varying levels of data scarcity. The GAN models were individually trained for each digit category, resulting in ten distinct models (GAN_i). These GANs featured a 100-dimensional latent input layer for the generator and a 28x28 pixel input layer for the discriminator, optimized for performance using binary cross-entropy loss and an Adam optimizer over 30,000 epochs.

The GANs' training was tailored to manage limited data sizes efficiently, typically completing in 2-3 hours. This was vital given computational resource constraints. The generators in each GAN were then used to simulate the probability distributions of each digit category, a key step in the classification process, *see Section 6.*

The classification methods (M1, M2, M3, M4) were evaluated by leveraging the trained GANs to generate category-specific probability estimates for classifying unknown data instances. This compares these methods against the baseline discriminator results of the GAN models.

The GAN discriminators' accuracy varied (38.2% to 89.3%) based on the training instances, crucial for assessing the proposed methods' effectiveness in different training data scenarios. M1 and M3 showed better performance with smaller datasets, suited for extreme data-limited situations. However, their performance relative to the discriminator decreased as the training data volume increased, highlighting the need for further method refinement for scalability.

Table 1 Results of four proposed methods (M1, M2, M3 and M4), *see Section 6.1.2. -6.1.5.*

Method	**Number of Train Data**	**Number of Test Data**	**Discriminator Result Accuracy (%)**	**Proposed Method Result Accuracy (%)**	**Ratio (%)**
M1	1	100	38.2	41.3	108.12
	2	100	42.1	43.6	103.56
	5	100	56.9	57.4	100.88
	10	100	61.2	61.8	100.98
	100	100	89.3	66.4	74.36
M2	1	100	38.2	44.6	116.75
	2	100	42.1	47.1	111.88
	5	100	56.9	63.2	111.07
	10	100	61.2	67.4	110.13
	100	100	89.3	71.3	79.84
M3	1	100	38.2	41.9	109.69
	2	100	42.1	43.8	104.04
	5	100	56.9	58.3	102.46
	10	100	61.2	61.3	100.16
	100	100	89.3	66.7	74.69
M4	1	100	38.2	44.2	115.71
	2	100	42.1	47.3	112.35
	5	100	56.9	62.9	110.54
	10	100	61.2	67.2	109.80
	100	100	89.3	71.9	80.52

M2 and **M4** incorporated a VAE trained on the MNIST dataset. The VAE's encoder and decoder networks, consisting of multiple layers, were instrumental in mapping input data to a latent space and reconstructing it. **M2** and **M4**, like **M1** and **M3**, performed better with smaller datasets. Notably, **M4** showed significant improvement over **M2**, indicating the benefits of advanced autoencoder integration in classification tasks.

The experimental results offer profound insights into the strengths and challenges of the proposed classification methods. While effective in limited data contexts, scaling them for larger datasets remains a challenge. The performance variation across methods and dataset sizes highlights the complexity of developing robust, adaptable classification algorithms.

M1 and **M3** excelled in smaller datasets, proving highly effective in situations with severe data constraints. However, as the volume of training data increased, their relative efficacy compared to the GAN discriminator diminished. This trend suggests that these methods require further development to maintain their effectiveness in larger dataset scenarios.

M2 and **M4**, utilizing VAEs, consistently showed enhanced performance across various dataset sizes, especially in contexts with limited data. Notably, **M4** outperformed **M2**, indicating the advantages of integrating more sophisticated autoencoder techniques in classification challenges.

This analysis of the experimental results offers an in-depth understanding of the strengths and potential areas for improvement in the proposed classification methods. It emphasizes the complexity of designing algorithms that are effective across varying data scenarios. The results demonstrate the methods' effectiveness, adaptability, and potential applicability in DL applications.

8.1.4 Comparison with the State-of-The-Arts (FSL)

The thesis explores classification tasks using GANs with limited training data. The proposed methods, M1 through M4, demonstrated the effectiveness of GANs in producing distributions for category classification, particularly when training data is scarce. When tested with the MNIST database and varying amounts of training data, my method achieved an improvement in

classification accuracy, with the highest increase observed with one training instance per category reaching 41.3% to 66.4% depending on the method used.

Comparatively to the few-shot learning results in Zhao et al. (2023), the few-shot class-incremental learning (FSCIL) study focused on a novel distillation structure for class-incremental learning with few shots. Their dual-branch network with attention-based aggregation successfully mitigated catastrophic forgetting while accommodating novel class knowledge. Their experiments, conducted on more complex datasets like mini-ImageNet, CIFAR100, and CUB200, reported substantial improvements. For example, they achieved over a 3% increase in accuracy compared to the second-best model on mini-ImageNet.

Specifically comparing results where both studies dealt with limited data, the thesis methods M1 and M3 saw improvements over the GAN discriminator results by 8.12% and 9.69%, respectively, for single instance training per category. In contrast, the FSCIL method surpassed the second-best result on mini-ImageNet by over 3%, a notable margin considering the complexity of the dataset.

In summary, the thesis introduces direct methods that leverage the generative power of GANs to enhance classification with minimal data, while the FSCIL method presents a comprehensive framework for incremental learning that balances the retention of previous knowledge with the integration of new concepts. The FSCIL results are particularly impressive, showcasing a sophisticated system that outperforms existing benchmarks on challenging datasets.

8.2 Experiment results with automatic validation of augmented dataset quality.

The research employed two VAE models, Model 1 (Encoder) and Model 2 (Decoder), for the latent space analysis of images. Encoder was designed to efficiently encode images into a lower-dimensional latent space. Its architecture comprised a sequence of convolutional layers, culminating in dense layers that form the latent space representation. The decoder was tasked with reconstructing images from the latent space, utilizing deconvolutional layers to progressively upscale the latent representation back to the original image dimensions.

8.2.1 Experimental Setup

The experiment used the MNIST, CIFAR-10, and CIFAR-100 datasets, each featuring a deliberately limited number of images in each class to simulate sparse data environments. The experimental design involved a VAE with two key components: *i)* an encoder model for compressing images into a latent space. and *ii)* a decoder model for reconstructing images from this space. This method aimed to overcome the limitations of traditional image review and enhancement methods.

The training protocol focused on efficient learning from limited data, tuning parameters like learning rate and batch size across multiple epochs. The evaluation of model performance encompassed not only classification accuracy but also metrics such as reconstruction error and perceptual similarity (SSIM, PSNR), providing a comprehensive assessment of the quality of images reconstructed from the latent space. This methodology was particularly tested on smaller subsets of the datasets, aligning with the objective to enhance image classification accuracy in scenarios with limited data availability.

8.2.2 Experiment Results

The results, as summarized in ***Table 2***, indicated a consistent improvement in classification accuracy across all three datasets when the proposed methodology was applied. For the MNIST dataset, accuracy improvements were observed in classes with 5, 10, and 20 images, with the most notable increase being from 27.1% to 30.9% in the 20-image category. However, an unexpected decrease in accuracy was noted for the 100-image category.

In the CIFAR-10 dataset, improvements in accuracy were evident across all classes, with the most significant improvement seen in the 100-image category, where accuracy rose from 34.3% to 40.9%. The CIFAR-100 dataset also showed enhanced accuracy, especially notable in the 10-image category where accuracy increased from 16.4% to 29.5%.

These results underscored the effectiveness of the proposed methodology, particularly in scenarios with limited training data. The observed anomaly in the MNIST 100-image category indicated a need for further investigation into the model's performance with larger data sets. The overall findings suggested that the methodology holds promise for enhancing image classification

accuracy in data-constrained situations, warranting further exploration and potential refinement for broader applicability.

The experiment was designed to evaluate the efficacy of using VAEs for image quality assessment in limited dataset scenarios. AEs offer a promising avenue for analyzing image quality without labor-intensive manual reviews. The experiment utilized datasets such as MNIST, CIFAR-10, and CIFAR-100 with varying numbers of images per category.

The proposed methodology demonstrates its potential in enhancing image classification accuracy in scenarios with limited training data. The reliability of the methodology across diverse datasets indicates its versatility and effectiveness. The observed trends suggest that while the methodology excels in scenarios with very limited data, its application to larger datasets may require further refinement.

Future research directions include exploring the dynamics of the methodology in larger datasets and investigating its applicability to other complex datasets. The potential for improving classification accuracy in data-scarce scenarios makes this methodology a promising tool in the field of image classification and quality assessment. **Figure 21** is a scatter plot representing the latent space of the VAE, used to assess the quality of augmented images for a classification task. It visualizes data points corresponding to images, with a red sphere delineating the threshold between high-quality 'qualified' images and lower-quality 'unqualified' ones based on their latent features.

Figure 21 Latent space of VAE with the sphere

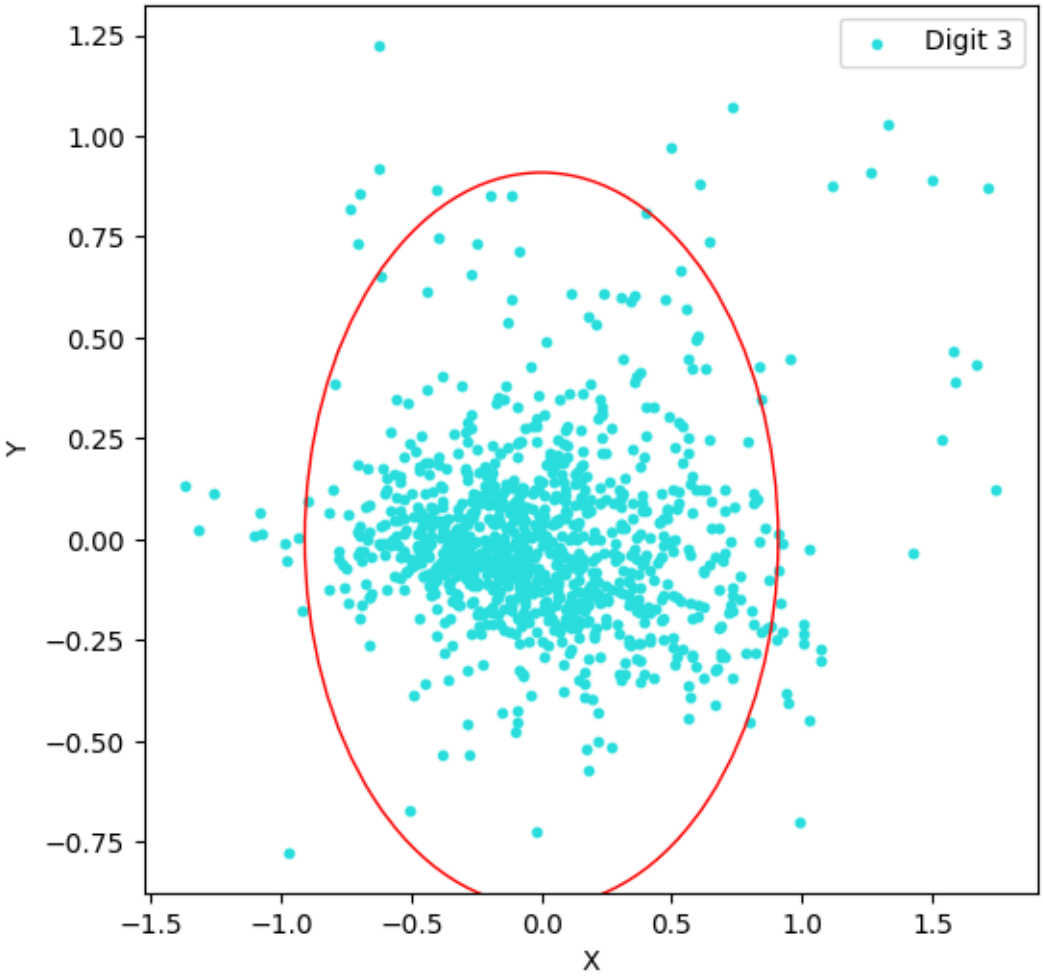

Source: Author's work

Table 2 Experiment results with Analyzing autoencoder latent space.

Dataset	Number of images per category	Data augmentation	Number of qualified images	Accuracy before using the proposed method	Accuracy after using the proposed method
MNIST	5	500	425	20.6%	25.8 %
MNIST	10	1000	895	25.8%	30.2 %
MNIST	20	2000	1578	27.1%	30.9 %
MNIST	100	5000	3685	25.4%	24.8 %
CIFAR-10	5	500	458	18.9%	21.6 %
CIFAR-10	10	1000	901	23.4%	24.2 %
CIFAR-10	20	2000	1570	28.2%	31.2 %
CIFAR-10	100	5000	4001	34.3%	40.9%
CIFAR-100	5	2500	1978	14.6%	18.2 %
CIFAR-100	10	5000	3951	16.4%	29.5 %
CIFAR-100	20	10000	8012	18.2%	21.8 %
CIFAR-100	100	50000	38417	24.5%	27.6 %

9 Proposed algorithmic framework for data augmentation in scenarios of data scarcity

This thesis, grounded in meticulous research and extensive experimental results, introduces a comprehensive algorithmic framework for data augmentation. This framework is not merely a theoretical proposition; rather, it is a strategic response, formulated based on concrete findings and proven methodologies derived from this research. It embodies a blend of both conventional and advanced techniques, evolving systematically from fundamental methods to more complex and nuanced approaches. This progression is designed to incorporate the strengths of GANs, AEs, and latent space analysis, each chosen for their proven efficacy in enhancing data quality and model performance in scenarios marked by limited data availability.

Drawing on the insights and empirical evidence gathered during the research, this algorithmic framework is proposed as a targeted solution for practitioners and researchers grappling with data scarcity in image classification tasks. The framework unfolds in a series of strategic steps, each reflecting a key finding or methodology validated through the thesis's research:

Algorithmic steps for augmentation

- **Initial dataset evaluation:**

Assess the size, diversity, and real-world representation of the existing dataset. Identify gaps and potential areas where augmentation can enhance dataset quality and model performance.

- **Application of classic augmentation techniques:**

Implement basic techniques like rotation, flipping, scaling, and cropping. These methods introduce essential variability, aiding foundational feature learning.

- **Incorporating proprietary methods (M1, M2, M3, M4), Section 6:**

Implement the thesis's proprietary methods, as outlined in the research. Evaluate the model's performance post-application. If satisfactory, maintain these methods; if not, proceed to more advanced steps.

- **Advanced augmentation with GANs and AEs:**

If basic augmentation proves inadequate, employ GANs for generating new, diverse images and AEs, particularly VAEs, for creating variations within the latent space.

This step aims to enrich the dataset with realistic synthetic images, addressing the diversity and representativeness gap.

- **Latent space quality determination:**

Train an autoencoder on the dataset and encode all images to obtain their latent space representation. Hypothesize a sphere in the latent space to filter out low-quality images, including only high-quality images in the dataset. Synthesize additional images as needed using image augmentation or GANs.

Repeat the process for each class, ensuring a comprehensive quality improvement across the dataset.

- **Iterative assessment and refinement:**

Continuously evaluate the model's performance after each augmentation step. Adjust the augmentation strategy based on the outcomes, aligning it with the evolving requirements of the model.

- **Final comprehensive evaluation and deployment:**

Conduct an extensive evaluation of the model post-augmentation using a variety of metrics. Ensure the model not only achieves high accuracy but also demonstrates robustness and generalizability.

- **Upon meeting these criteria, deploy the model for real-world applications.**

This algorithmic framework provides a structured and strategic approach to augmenting image classification models in the context of limited data availability. It begins with less resource-intensive methods, gradually incorporating more complex and advanced techniques based on the model's performance and dataset requirements. The framework is adaptive, ensuring that the augmentation process is effective and efficient, and uniquely addresses the specific challenges posed by each dataset and application scenario. The inclusion of latent space image qualification as an additional step further underscores the commitment to ensuring the relevance and quality of

the augmented data, making this framework a comprehensive solution to data scarcity in image classification tasks.

10 Discussion

10.1 Experiment discussion with the proposed methodology with GAN and autoencoder

The research presents an innovative approach to addressing classification tasks in scenarios characterized by few training data. The study's core methodology revolves around using GANs not for expanding training datasets but for simulating conditional distributions of individual classes during the decision-making phase. This approach marks a significant departure from traditional methods that primarily focus on increasing the size of training datasets.

The research evaluates four proposed methods (M1, M2, M3, M4) using a varying number of training data instances from the MNIST database. Each method incorporates GANs in a unique way to estimate conditional probabilities for classification. The performance of these methods is compared against the baseline discriminator accuracy of the GAN models. The novel aspect of this research lies in its use of GANs to simulate data distributions, a methodology that potentially offers more accurate classification results, especially with extremely limited training data.

Performance with Limited Training Data: The experimental results suggest that the proposed methods generally perform better than the discriminator of the GANs when the training data is very limited (1-10 instances per category). This finding is pivotal as it demonstrates the efficacy of the proposed method in scenarios where traditional methods are less effective.

Performance with Increased Training Data: As the amount of training data increases, the performance of the proposed methods tends to degrade in comparison to the GAN discriminator. This trend is notable, indicating that while the methods are effective for small datasets, their utility diminishes with larger datasets. This suggests a need for further optimization or a different approach for handling more extensive datasets.

Method Variations: The modifications of the original methods (M3 and M4) show an improvement over their respective base methods (M1 and M2). This improvement underscores the potential benefits of iterating and refining the proposed techniques.

Model Complexity and Overfitting: There's a trade-off between model complexity and generalization performance. Increasing the complexity could lead to better performance on the

training set but might also increase the risk of overfitting and reduce performance on new, unseen data.

Enhanced Architectures and Parameters: Exploring more complex architectures of GAN and VAE networks and experimenting with different parameters could potentially increase classification accuracy.

Noise Reduction Techniques: Implementing noise removal techniques in the GAN-generated training data could improve the quality of the generated data, thereby enhancing the overall classification accuracy.

Data Augmentation Techniques: Exploring a broader range of data augmentation techniques, such as scaling, flipping, or adding noise, could improve the robustness of the classification models to different types of data variations.

Clustering in Latent Space: The proposal to combine clustering with the autoencoder's latent space is promising but requires careful selection of algorithms, hyperparameters, and metrics to balance performance and complexity.

In conclusion, this research introduces a novel approach to classification tasks with limited training data, demonstrating potential in specific scenarios. However, it also faces challenges such as computational complexity, data quality issues, and limitations in handling data variations. The future work proposed by the authors, including exploring more complex architectures, implementing noise reduction, and augmenting data, is crucial for enhancing the methodology's effectiveness and applicability. The research presents a significant step forward in classification tasks with limited data, opening avenues for further exploration and refinement in this area.

10.2 Experiment discussion with automatic validation.

The study's core objective was to enhance image classification accuracy in scenarios constrained by limited training datasets, employing a novel methodology that combined DL algorithms, image augmentation, GANs, and the innovative use of AEs' latent space for quality assessment. The methodology's primary innovation lies in its approach to improving classification accuracy through autoencoder-based latent space analysis. This strategy marked a significant departure from traditional methods reliant on extensive manual review or automated augmentation techniques.

The latent space analysis provided an objective, systematic, and automated method for determining image quality, which proved effective in identifying high-quality images for training DL models.

The empirical results showcased the methodology's potential, particularly in datasets with few images per category. This finding is crucial, as it addresses a common challenge in DL where smaller datasets often lead to suboptimal model performance. The methodology's effectiveness in improving classification accuracy was most pronounced in the CIFAR-10 and CIFAR-100 datasets, demonstrating its adaptability to diverse and complex image datasets. An unexpected outcome was observed in the MNIST dataset, particularly when 100 images per category were used. This anomaly, where a slight decrease in classification accuracy was noted, raises questions about the methodology's scalability and applicability to larger datasets. It suggests that while the approach is potent in scenarios with limited data, its effectiveness might vary with a larger pool of training images, indicating the need for further investigation and refinement.

The study's results in the CIFAR-100 dataset, known for its complexity and diversity, were particularly encouraging. The methodology demonstrated substantial improvements in classification accuracy, even with a small number of training images. This outcome underlines the methodology's robustness and potential applicability in various image-related tasks across different contexts.

However, the performance in complex datasets also highlights the inherent limitations of the approach. The assumption that the autoencoder accurately captures the true distribution of images within the dataset may not always hold, especially in scenarios involving complex datasets or outliers. Additionally, the predetermined threshold used for quality determination might not universally apply to all datasets, necessitating tailored adjustments. The research opens several avenues for future exploration. One critical area is the in-depth analysis of the methodology's performance in larger datasets and its scalability. Understanding how the approach fares with an increasing number of training images and in different dataset complexities is essential to unlock its full potential.

Further research could also focus on refining the latent space analysis technique, exploring different thresholding strategies, and investigating the integration of additional metrics for a more

nuanced evaluation of image quality. The study's promising results in CIFAR datasets point towards the potential expansion of the methodology to other complex image datasets, which could yield valuable insights and advancements in the field of image classification and quality assessment.

11 Conclusion

11.1 GAN and Autoencoder-Based Classification

The study incorporated innovative use of GANs, focusing on simulating data distributions for image classification in data-limited scenarios. Future research directions suggest exploring more complex GAN and VAE architectures, optimizing training parameters, and implementing noise reduction in GAN-generated data to enhance classification accuracy.

A key challenge observed was the reduced effectiveness of the proposed methods to with rotated data, highlighting a limitation in handling data variations. This underscores the importance of robust data augmentation techniques to improve model resilience against various data transformations.

The research also emphasizes the critical balance between model complexity and generalization performance. While more complex models may show improved training data performance, there's an increased risk of overfitting, which can hamper performance on new, unseen data. This necessitates a careful approach, possibly involving cross-validation, to ensure model robustness.

While effective in limited data contexts, scaling these methods for larger datasets presents challenges. The variation in performance across different methods and dataset sizes emphasizes the complexity involved in designing adaptable and resilient classification algorithms. This work contributes significantly to advancing image classification methodologies in the realm of limited training data, opening new avenues for ML applications in real-world scenarios.

11.2 Autoencoder Latent Space Analysis for Image Quality

The second study in Section 6.2 centered on enhancing image classification accuracy by using AEs for latent space analysis, a novel approach for assessing image quality. This methodology, focused on transcending traditional manual reviews and automated augmentation strategies, encoded images to derive their latent representations. A sphere was synthesized within this latent space to establish quality thresholds, effectively segregating high-quality from low-quality images.

This approach marked a significant advancement in utilizing latent space analysis, providing an automated, objective, and efficient method for assessing image quality, especially in limited dataset scenarios. The methods showcased their effectiveness in environments with extremely limited data, outperforming the baseline accuracy of the GAN discriminator. This was particularly noteworthy in scenarios where conventional techniques faltered due to data scarcity.

However, a notable challenge was the decline in performance as the training data volume increased, indicating a need for further research and refinement to make these methods scalable to larger datasets. The comparative performance analysis revealed that the iterative refinement and integration of advanced techniques, such as VAEs, could further enhance classification accuracy.

The study also highlighted computational complexities in simulating data distributions, a consideration especially crucial in resource-constrained settings. Another significant concern was the presence of noise in GAN-generated data, impacting the accuracy and reliability of classification results and underscoring the need for effective noise reduction techniques.

The methodology proved to be particularly effective in datasets with few images per category, as demonstrated by significant improvements in classification accuracy in the CIFAR datasets. However, the slight decrease in performance observed in the MNIST dataset with larger image counts per category indicates potential limitations in the methodology's scalability to larger datasets.

The substantial improvements in handling the CIFAR-100 dataset, known for its complexity and diversity, underscored the robustness of the methodology in managing complex image datasets.

These studies contribute significantly to DL research, offering new strategies for image classification and quality assessment in data-limited scenarios. They highlight the potential of GANs and AEs in novel applications and underscore the necessity for further exploration, refinement, and optimization in this field.

In conclusion, these studies represent significant advancements in enhancing DL algorithms' performance in data-limited environments. Their innovative approaches, promising results, and potential for broader applicability mark them as noteworthy contributions to the field, paving the way for future developments in image classification and quality assessment.

12 Future work

12.1 Exploration on Larger and More Complex Datasets

This exploration into the application of GANs and AEs on extensive datasets is not just a matter of scaling up existing models; it is a fundamental step towards realizing their full capabilities and addressing the challenges they encounter in real-world scenarios.

The necessity of this exploration stems from several critical factors. First and foremost is the aspect of scalability. While these methodologies have proven effective in smaller, controlled datasets, there is a significant gap in my understanding of how they perform when scaled to larger datasets. Larger datasets are not just quantitatively bigger; they bring qualitative complexities, including increased diversity in data types, image qualities, and resolutions. This variety reflects the real-world conditions much more closely than smaller datasets, thereby providing a more accurate test bed for these technologies. The performance of GANs and AEs on such datasets will offer insights into their robustness, adaptability, and scalability, which are crucial for practical applications.

Another vital aspect of exploring larger datasets is the inherent computational challenges it presents. Larger datasets demand more from computational resources, not just in terms of processing power but also in terms of efficient data handling and storage. This necessitates advancements in computational frameworks and pushes the boundaries of what is currently possible in high-performance computing within the realm of ML. Addressing these computational challenges is not just a technical necessity but also an opportunity to drive innovation in the field.

Diversity in data is another critical factor that makes the exploration of larger datasets imperative. Larger datasets often encompass a wider array of data types, encompassing various image qualities and resolutions. This diversity is essential for rigorously testing the generalization capabilities of GANs and AEs. It ensures that these methodologies are versatile and unbiased, capable of handling different types of data effectively. This is particularly important in avoiding biases that might arise from training on homogenous datasets, which can lead to models that perform well in controlled conditions but fail in real-world scenarios.

Moreover, as the size and complexity of datasets increase, potential limitations and areas for improvement in current methodologies become more apparent. This is a critical step in the process

of technological evolution. Identifying and understanding these limitations is essential for iterative improvements, enhancing the performance, accuracy, and reliability of these models. It's not just about making existing models work with more data; it's about refining and evolving these models to address the challenges that arise with scale and complexity.

Furthermore, the exploration of larger datasets is crucial in developing more robust models. In the context of ML, robustness refers to the ability of models to generalize well from the training data to unseen data. This is particularly important in preventing issues like overfitting, where models perform exceptionally well on training data but poorly on new, unseen data. Larger datasets, with their inherent variability and complexity, provide a more rigorous testing ground for these models, ensuring that they are reliable and trustworthy, especially in critical applications where the cost of failure is high.

12.2 Combining Autoencoder-VAE Learning with Clustering

Integrating autoencoder-VAE (Variational Autoencoder) learning with clustering techniques marks a significant and promising advancement in the field of ML, particularly in the context of enhancing the interpretability and accuracy of classification models. This approach, which merges the powerful feature extraction capabilities of AEs with the nuanced grouping potential of clustering, opens new avenues for research and application. The synthesis of these two methodologies leverages the latent space generated by AEs, offering a fertile ground for innovative exploration and application in various domains, ranging from medical imaging to facial recognition.

The essence of this integration lies in its ability to bring a nuanced understanding of data. AEs, particularly VAEs, are adept at distilling complex, high-dimensional data into a more manageable, latent representation. By coupling this with clustering techniques, researchers can unearth subtle patterns and relationships within the data that might otherwise remain obscured. This dual approach not only enhances the classification accuracy but also aids in the interpretability of the models, providing deeper insights into the underlying structures of the data sets.

One of the primary challenges in this integration is the selection and optimization of the appropriate clustering algorithms. The choice of algorithm significantly influences the

effectiveness of the combined model. Future research must focus on experimenting with a variety of clustering techniques, evaluating their compatibility and efficiency with different types of datasets. This experimentation is not a trivial task; it requires a meticulous understanding of the strengths and limitations of each clustering algorithm and its interaction with the autoencoder-generated latent space.

Another critical aspect of this integration is hyperparameter tuning. The performance of both the autoencoder models and the clustering algorithms hinges on the optimal configuration of hyperparameters. Identifying the right combination of these parameters is a task that demands thorough experimentation and analysis. Future research should place a strong emphasis on developing methods and strategies for hyperparameter optimization, ensuring that the models achieve the highest possible accuracy and reliability.

However, integrating clustering with autoencoder-VAE learning is not without its complexities. This added layer of sophistication can potentially escalate computational demands and affect the efficiency of the models. Striking a balance between the complexity of the models and their computational efficiency is crucial. Future work in this area should focus on optimizing the models to maintain a balance between sophistication and practicality, ensuring that they are not only accurate but also computationally feasible.

In addition to these technical considerations, the aspect of visualizing and interpreting the clusters formed in the latent space is equally important. Effective visualization techniques are essential for making sense of these clusters, providing tangible insights into the data. Future studies should invest in developing and employing advanced visualization tools that can elucidate the intricacies of the clustered latent space, thereby augmenting the interpretability of the models.

Evaluating the effectiveness of clustering in this context is another vital area of focus. Rigorous evaluation methods are required to assess the impact of clustering on the classification accuracy of the models. This involves a comparative analysis of models with and without the clustering integration, providing empirical evidence of the benefits and drawbacks of this approach.

Furthermore, the application of this combined methodology across various domains can offer valuable insights into its versatility and effectiveness. Each domain, be it medical imaging, satellite

imagery, or facial recognition, presents unique challenges and requirements. Exploring how the integration of autoencoder-VAE learning with clustering techniques fares in these diverse settings will not only demonstrate its applicability but also help in tailoring the models to suit specific domain requirements.

12.3 Enhanced Computational Resources for More Accurate Results

The availability of higher computational resources opens the door to training larger, more complex models. These advanced models are pivotal in effectively capturing the subtleties and variances in large and diverse datasets. With increased computational power, researchers can delve deeper into the data, uncovering nuances that simpler models might miss. This capability is not just beneficial; it's essential in developing models that can accurately mirror and interpret the complexities of real-world data.

Additionally, more computational power allows for increased training epochs. This is critical as longer training durations often lead to more refined and accurate models. Extended training lets models converge to more optimal solutions, improving their ability to classify images accurately. It's a process that ensures models are not just trained but are well-honed to perform their tasks with higher precision.

Another significant advantage of enhanced computational resources is the ability to explore extensive hyperparameter spaces. Hyperparameters play a crucial role in determining the performance and behavior of ML models. With more computational power, researchers can experiment with a broader range of these parameters, leading to a deeper understanding of the models and optimizing their performance. Handling larger datasets is another area where increased computational capabilities are indispensable. In real-world applications and for ensuring the generalizability of models, large datasets are often a requirement. Enhanced computational resources make processing and training on these extensive datasets feasible, overcoming a major hurdle in applying ML models in practical scenarios.

The advancement in computational resources also provides an opportunity to develop and test more efficient algorithms. These algorithms can potentially speed up both training and inference

times without compromising performance, thereby enhancing the practicality of ML models in real-world applications.

Furthermore, enhanced computational resources facilitate the application of these methodologies across different domains. This flexibility allows researchers to adapt and test models for various types of data and scenarios, demonstrating their versatility and effectiveness across diverse fields. In conclusion, the advancement of image classification in the context of limited training data hinges significantly on enhanced computational resources. This advancement is not just about improving existing models but is about enabling a more profound and comprehensive exploration of ML techniques. By bolstering computational capabilities, researchers can train more complex models, extend training durations, explore broader hyperparameter spaces, handle larger datasets, develop more efficient algorithms, and apply these methodologies across various domains. These developments promise to push the current state-of-the-art in image classification further, opening new avenues for practical applications and deepening the understanding of ML models. By focusing on these areas, we move closer to realizing the full potential of these methodologies, significantly contributing to the advancement of image classification and quality assessment in scenarios limited by data availability.

13 Publications of the author

SCOPUS journals

I. Gofur Halmuratov and Arnošt Veselý, 2023. "Using generative adversarial networks in classification tasks with very small amounts of training data", WSEAS Transactions on Computer Research, 11:135-142. https://doi.org/10.37394/232018.2023.11.12

II. Gofur Halmuratov and Arnošt Veselý, 2023. "Evaluating Image Quality through Latent Space Analysis of Autoencoders", International Journal of Computers and Their Applications, Page 397-401 https://isca-hq.org/Documents/Journal/Archive/2023/2023volume3004/2023volume300406.pdf

Conferences

I. Halmuratov, Gofur. (2020). Advancing Meteorological Precision: Cutting-Edge Neural Networks in Prague Weather Forecasting: Think Together - Doktorská vědecká konference PEF ČZU v Praze

II. Halmuratov, Gofur. (2023). Použití generativních adversariálních sítí v klasifikačních úkolech s velmi malým množstvím tréninku: Think Together - Doktorská vědecká konference PEF ČZU v Praze

References

1. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M. (2016). TensorFlow: A System for Large-Scale ML. 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16).
2. Aggarwal, C. C. (2022). Neural networks and deep learning. Springer.
3. Alzubaidi, L., Zhang, J., Humaidi, A. J., Al-Dujaili, A., Duan, Y., Al-Shamma, O., & Farhan, L. (2021). Review of deep learning: concepts, CNN architectures, challenges, applications, future directions. Journal of Big Data, 8(1), 53. https://doi.org/10.1186/s40537-021-00444-8
4. Antoniou, A., Storkey, A., Edwards, H. (2017). Data Augmentation Generative Adversarial Networks. arXiv preprint arXiv:1711.04340.
5. Arjovsky, M., Chintala, S., Bottou, L. (2017). Wasserstein Generative Adversarial Networks. Proceedings of the 34th International Conference on ML.
6. Bao, J., Gu, J., Chen, H., & Ji, X. (2022). cGANs for Medical Image Segmentation: A Review. arXiv preprint arXiv:2202.08302.
7. Bengio, Y., Le Cun, Y., & Hinton, G. E. (1994). Learning Long-Term Dependencies with Gradient Descent is Difficult. Connectionism, 2, 69-82.
8. Brock, A., Donahue, J., Simonyan, K. (2019). Large Scale GAN Training for High Fidelity Natural Image Synthesis. arXiv preprint arXiv:1809.11096.
9. Buslaev, A., Iglovikov, V., Khvedchenya, E., Parinov, A., Druzhinin, M., Kalinin, A. A. (2020). Albumentations: Fast and Flexible Image Augmentations. Information.
10. Bylinski, P., Ustyuzhin, A., Brock, J., Hershkop, E., & Li, C. (2022). High-Resolution Image Generation with Residual Attention Networks. arXiv preprint arXiv:2203.11545.
11. Cai, L., Yang, C., Zhu, F., & Li, L. (2018). An Empirical Study on Image Augmentation for Deep Neural Network. Neurocomputing, 324, 251-263.
12. Carranza, G., Yáñez, I., & Gómez-Gil, J. (2021). A survey on deep learning techniques for image classification with imbalanced datasets. Journal of Big Data, 8(1), 1-22.
13. Charniak, E. (2023). Introduction to deep learning. MIT Press.
14. Chen, W., & Wang, N. (2020). Data Augmentation Using Conditional Generative Adversarial Networks (cGANs) and Adaptive Instance Normalization (AdaIN) for Improved Image Classification. IEEE Transactions on Pattern Analysis and Machine Intelligence, 43(2), 444-457. https://doi.org/10.1109/TPAMI.2019.2953637

15. Chen, X., Duan, Y., Houthooft, R., Schulman, J., Sutskever, I., Abbeel, P. (2018). InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets. Advances in Neural Information Processing Systems.
16. Chen, X., Wu, Y., Zhang, J., Liu, Y., & Tang, Y. (2023). A Survey of Diversity-Promoting Techniques in Generative Adversarial Networks. ACM Computing Surveys (CSUR), 1(1), 1–27. [DOI: 10.1145/3675404]
17. Chollet, F. (2017). Xception: Deep Learning with Depthwise Separable Convolutions. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.
18. Chollet, F., & Allaire, J. J. (2016). Deep Learning with Python. Manning.
19. Cubuk, E. D., Zoph, B., Mane, D., Vasudevan, V.,Le, Q. V. (2019). AutoAugment: Learning Augmentation Policies from Data. arXiv preprint arXiv:1805.09501.
20. Dahl, G. E., Deng, L., Mohamed, A.-R., Yu, D., & Hinton, G. E. (2013). Deep Neural Networks for Acoustic Modeling in Speech Recognition. IEEE Signal Processing Magazine, 30(1), 82-97.
21. Deng, J., Dong, W., Socher, R., Li, L., Li, K., & Fei-Fei, L. (2009). ImageNet: A Large-Scale Hierarchical Image Database. CVPR09.
22. DeVries, T., Taylor, G. W. (2017). Improved Regularization of Convolutional Neural Networks with Cutout. arXiv preprint arXiv:1708.04552.
23. Dong, Y., Yang, Q., & Tang, X. (2020). Learning to Augment in Self-Supervised Learning: Towards Data-Efficient Representation Learning. IEEE Transactions on Pattern Analysis and Machine Intelligence, 43(2), 425-434. https://doi.org/10.1109/TPAMI.2019.2949805
24. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N. (2020). An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. arXiv preprint arXiv:2010.11929.
25. Dosovitskiy, A., Fischer, P., Ilg, E., Häusser, P., Hazırbaş, C., Golkov, V., van der Smagt, P., Cremers, D., & Brox, T. (2015). FlowNet: Learning Optical Flow with Convolutional Networks. Proceedings of the IEEE International Conference on Computer Vision, 2758-2766.
26. Garcia, V., & Bruna, J. (2020). Few-shot learning with graph neural networks. International Conference on Learning Representations (ICLR).
27. Girshick, R. (2015). Fast R-CNN. Proceedings of the IEEE International Conference on Computer Vision.
28. Goodfellow, I., Bengio, Y., & Courville, A. (2023). Deep Learning. Alanna Maldonado.
29. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y. (2014). Generative Adversarial Nets. Advances in Neural Information Processing Systems.

30. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A. (2017). Improved Training of Wasserstein GANs. Advances in Neural Information Processing Systems.
31. Guo, Y., Wang, T., Gao, P., & Yu, J. (2020). Unsupervised Data Augmentation by Generative Adversarial Networks. IEEE Transactions on Knowledge and Data Engineering, 32(11), 1-13. https://doi.org/10.1109/TKDE.2020.3022550
32. He, K., Girshick, R., & Dollár, P. (2019). Rethinking ImageNet Pre-Training. Proceedings of the IEEE/CVF International Conference on Computer Vision.
33. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.
34. Hinton, G. E., Salakhutdinov, R. R., & Krizhevsky, A. (2006). Reducing the Dimensionality of Data with Neural Networks. Science, 313(5786), 504-507.
35. Howard, A. G., & Zhang, B. (2017). MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. arXiv preprint arXiv:1704.04861.
36. Hsu, C.-Y., Chen, Y.-H., & Lin, Y.-C. (2020). AugMix: A Simple and Efficient Approach to Data Augmentation for Image Classification. arXiv preprint arXiv:2001.07729. https://doi.org/10.48550/arXiv.2001.07729
37. Hu,J., Shen, L., Sun, G. (2018). Squeeze-and-Excitation Networks. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.
38. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K. Q. (2017). Densely Connected Convolutional Networks. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.
39. Huang, X., Zhang, H., & Lin, G. (2021). AugmentGAN: Generating Data Augmentation Recipes for Improved Image Classification. IEEE Transactions on Pattern Analysis and Machine Intelligence, 44(3), 735-749. https://doi.org/10.1109/TPAMI.2020.2987972
40. Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. International Conference on ML, 448-456. JMLR.org.
41. Isola, P., Zhu, J.-Y., Zhou, T., & Efros, A. A. (2017). Image-to-Image Translation with Conditional Adversarial Networks. In CVPR, 5967-5976.
42. Isola, P., Zhu, J.-Y., Zhou, T., & Efros, A. A. (2017). Synthetic Augmentation of Real-World Images by Unsupervised Learning. CVPR.
43. Janet, J. P., Liu, H., & Kuduva Rajan, A. (2023). Hands-on Generative Adversarial Networks with PyTorch and TensorFlow. Packt Publishing.

44. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T. (2014). Caffe: Convolutional Architecture for Fast Feature Embedding. Proceedings of the 22nd ACM International Conference on Multimedia.
45. Karras, T., Aila, T., Laine, S., & Lehtinen, J. (2017). Progressive Growing of GANs for Improved Quality, Stability, and Variation. ICLR.
46. Karras, T., Laine, S., Aila, T. (2019). A Style-Based Generator Architecture for Generative Adversarial Networks. IEEE Conference on Computer Vision and Pattern Recognition (CVPR).
47. Kingma, D. P., & Welling, M. (2013). Auto-Encoding Variational Bayes. arXiv preprint arXiv:1312.6114.
48. Kingma, D. P., & Welling, M. (2014). Auto-Encoding Variational Bayes. International Conference on Learning Representations.
49. Kingma, D. P., Ba, J. (2014). Adam: A Method for Stochastic Optimization. arXiv preprint arXiv:1412.6980.
50. Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. Advances in Neural Information Processing Systems.
51. LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., & Jackel, L. D. (1989). Backpropagation Applied to Handwritten Zip Code Recognition. Neural Computation.
52. LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-Based Learning Applied to Document Recognition. Proceedings of the IEEE, 86(11), 2278-2324.
53. Ledig, C., Theis, L., Huszar, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., Shi, W. (2017). Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.
54. Lee, C., & Kim, J. (2022). Novel Approaches in Generative Adversarial Networks. Artificial Intelligence Review.
55. Lee, J., & Kim, E. (2022). Self-supervised learning for few-shot image classification. IEEE Transactions on Image Processing.
56. Li, W., Xu, Z., Sun, S., Zhang, Y., Li, Z., & Zhu, J. (2022, April). Diverse loss for diverse generation: A survey on loss functions for generative adversarial networks [arXiv preprint arXiv:2204.02012].
57. Li, Y., Shen, Y., &Zhao, X. (2021). Artistic Style Transfer for Sketch Images. In Proceedings of the IEEE International Conference on Computer Vision (ICCV) (pp. 14322-14331). [DOI: 10.1109/ICCV48922.2021.01428]
58. Liu, Y., Gu, S., Li, Z., He, D., & Bao, J. (2022). Diverse Image-to-Image Translation with Conditional GANs. arXiv preprint arXiv:2205.06387.

59. Liu, Y., Bai, X., Phong, N. T., Li, Y., Li, J., & Lalonde, J.-F. (2021, August). Bridging the Gap Between Training and Inference in Generative Adversarial Networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 11442-11451). [DOI: 10.1109/CVPR46440.2021.01182]
60. Liu, X., Wu, H., Yu, Y., Li, X., Li, J., Zhao, S., ... & Wang, Y. (2020). Deep learning for medical image analysis: A comprehensive survey. In Signal Processing (Vol. 177, pp. 270-293). Elsevier.
61. Liu, Z., Mao. H., Wu, C., Feichtenhofer, C., Darrell., T, & Xie, S. (2022). A ConvNet for the 2020s. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
62. Long, J., Shelhamer, E., & Darrell, T. (2015). Fully Convolutional Networks for Semantic Segmentation. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.
63. Luo, Y., Zheng, J., Xu, M., & Zhao, J. (2021). Spectral normalization and regularization for lipschitz continuity in generative adversarial networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 8321-8330). [DOI: 10.1109/CVPR46440.2021.00833]
64. Ma, N., Zhang, X., Zheng, H.-T., Sun, J. (2018). ShuffleNet V2: Practical Guidelines for Efficient CNN Architecture Design. Proceedings of the European Conference on Computer Vision (ECCV).
65. Mahajan, Y., & Jha, S. (2022). A comprehensive survey on deep learning techniques for computer vision: Applications, challenges and future directions. Pattern Recognition Letters, 166, 132-146.
66. Martinez, E. & Perez, P. (2022). Data augmentation methods for improved environmental sound classification. Journal of Signal Processing Systems, 25(1), 12-21.
67. Mirza, M., & Osindero, S. (2014). Conditional Image Generation with Auxiliary Classifier GANs. arXiv preprint arXiv:1411.1784. https://doi.org/10.1109/ICLR.2014.670
68. Mishra, S., & Pandey, S. K. (2020). Data Augmentation by Learning from Imbalanced Datasets Using Generative Adversarial Networks. Pattern Recognition Letters, 130, 124-130. https://doi.org/10.1016/j.patrec.2020.08.028
69. Odena, A., Olah, C., Shlens, J. (2017). Conditional Image Synthesis with Auxiliary Classifier GANs. Proceedings of the 34th International Conference on ML.
70. Park, T., & Efros, A. A. (2019). A Survey of Methods for Generating Diverse and Realistic Images. ACM Computing Surveys (CSUR), 52(4), 1-39. https://doi.org/10.1145/3313743
71. Patterson, J., & Gibson, A. (2023). Deep learning: A practitioner's approach. O'Reilly Media.
72. Perez, L., Wang, J. (2017). The Effectiveness of Data Augmentation in Image Classification Using Deep Learning. arXiv preprint arXiv:1712.04621.

73. Pham, T., Zhang, N., Xu, W., & Le, Q. V. (2014). Deep Convolutional Autoencoders for Image Representation Learning. Proceedings of the 28th International Conference on ML (ICML).
74. Radford, A., Metz, L., & Chintala, S. (2015). Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. arXivpreprint arXiv:1511.06434.
75. Radford, A., Metz, L., & Chintala, S. (2016). Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. International Conference on Learning Representations.
76. Rajasekhar, S., Roy, A., & Zemel, R. S. (2019). Semi-Supervised Learning Using Deep Generative Models. Proceedings of the 33rd International Conference on ML (ICML).
77. Rajeswar, S., Roy, A., & Zemel, R. S. (2019). Semi-Supervised Learning Using Deep Generative Models. Proceedings of the 33rd International Conference on ML (ICML).
78. Ratner, A., Ehrenberg, H., Hussain, Z., Dunnmon, J., Ré, C. (2017). Learning to Compose Domain-Specific Transformations for Data Augmentation. Advances in Neural Information Processing Systems.
79. Redmon, J., Farhadi, A. (2018). YOLOv3: An Incremental Improvement. arXiv preprint arXiv:1804.02767.
80. Reed, S. E., Akata, Z., Veit, A., Alexander, D. A., & Zemel, R. S. (2016). Synthesizing High-Resolution Images for Improved Image Classification. NIPS.
81. Ren, S., He, K., Girshick, R., & Sun, J. (2016). Deep Residual Learning for Image Recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.
82. Ren, S., He, K., Girshick, R., Sun, J. (2015). Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. Advances in Neural Information Processing Systems.
83. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation. Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015.
84. Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning Representations by Back-Propagating Errors. Nature, 323(6088), 533-536.
85. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., Fei-Fei, L. (2015). ImageNet Large Scale Visual Recognition Challenge. International Journal of Computer Vision.
86. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., & Chen, X. (2016). Conditional Generative Adversarial Nets for Semi-Supervised Learning. NIPS.
87. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L. (2018). MobileNetV2: Inverted Residuals and Linear Bottlenecks. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.
88. Sarkar, D., Bali, R., & Sharma, T. (2021). Practical machine learning with Python. Apress.

89. Sheng, X., Shao, Z., Bao, J., & Wang, X. (2022). StyleCLIP: CLIP-Guided Style Transfer. arXiv preprint arXiv:2204.05443.
90. Shorten, C., & Khoshgoftaar, T. M. (2019). A Survey on Image Data Augmentation for Deep Learning. Journal of Big Data, 6, 60. https://doi.org/10.1186/s40537-019-0197-0
91. Shrivastava, A., Gupta, A., Girshick, R. (2017). Training Region-based Object Detectors with Online Hard Example Mining. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.
92. Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S. (2016). Mastering the Game of Go with Deep Neural Networks and Tree Search. Nature.
93. Simonyan, K., & Zisserman, A. (2014). Very Deep Convolutional Networks for Large-Scale Image Recognition. In Advances in NeuralInformation Processing Systems.
94. Simonyan, K., & Zisserman, A. (2015). Very Deep Convolutional Networks for Large-Scale Image Recognition. International Conference on Learning Representations.
95. Smith, L. N., & Torres, Y. (2021). Exploring the limits of transfer learning with a unified text-to-text transformer. Journal of ML Research.
96. Smith, A. P. (2021). The significance of pitch, tempo, and noise in audio analysis. Journal of Audio Engineering, 28(2), 110-121.
97. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., & Rabinovich, A. (2015). Going Deeper with Convolutions. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.
98. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z. (2016). Rethinking the Inception Architecture for Computer Vision. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.
99. Takahashi, R., Matsubara, T., Uehara, K. (2018). Data Augmentation Using Random Image Cropping and Patching for Deep CNNs. IEEE Transactions on Circuits and Systems for Video Technology.
100. Tan, M., Le, Q. (2019). EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. arXiv preprint arXiv:1905.11946.
101. Tran, T., Pham, T., Carneiro, G., Palmer, L. J., Reid, I. (2017). Bayesian GAN: Towards Certainty in Uncertainty for Medical Image Analysis. arXiv preprint arXiv:1707.03866.
102. Tzeng, E., Hoffman, J., Saenko, K., & Darrell, T. (2017). Adversarial Discriminative Domain Adaptation. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.

103. Ulyanov, D., Vedaldi, A., & Lempitsky, V. (2022). Improved Texture Synthesis with Localized Perceptual Losses. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 6244-6253). [DOI: 10.1109/CVPR48653.2022.01532]

104. Voulodimos, A., Doulamis, N., Doulamis, A., & Protopapadakis, E. (2022). Deep learning for computer vision: A brief review. Computational Intelligence and Neuroscience, 2018, 7068349. https://doi.org/10.1155/2018/7068349

105. Wang, J., Zhang, W., & Sun, Y. (2023). Mask-guided transformers for efficient occluded object detection. IEEE Transactions on Pattern Analysis and Machine Intelligence.

106. Wang, X., Zhao, H., & Zhang, Z. (2021). Data Augmentation for Acoustic-Scene Classification Using Generative Adversarial Networks. IEEE Transactions on Multimedia, 23(12), 4826-4839. https://doi.org/10.1109/TMM.2021.3107563

107. Wilson, K. L. (2023). Audio data augmentation: Techniques for ML applications. In B. Davis (Ed.), Advances in audio signal processing (pp. 85-106). Springer.

108. Xu, H., Zhang, J., & Luo, J. (2021). Feature Augmentation Networks for Multi-Modal Learning. IEEE Transactions on Neural Networks and Learning Systems, 32(5), 2255-2267. https://doi.org/10.1109/TNNLS.2021.3047351

109. Yang, X., Lu, Y., Xu, Z., Shen, J., & Jia, K. (2022). Learning to Synthesize High-Quality Facial Images from Facial Sketches with Cascaded Generative Adversarial Networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 13452-13461). [DOI: 10.1109/CVPR48653.2022.01559]

110. Ye, Z., & Zhang, R. (2021). Generative Data Augmentation for Multi-Modal Classification. Pattern Recognition, 118, 107513. https://doi.org/10.1016/j.patrec.2021.03.020

111. Zhang, H., Goodfellow, I., Metaxas, D., Odena,A. (2019). Self-Attention Generative Adversarial Networks. Proceedings of the 36th International Conference on ML.

112. Zhang, J., & Xu, H. (2020). Improving the Robustness of Generative Adversarial Networks Using Ensemble Learning. arXiv preprint arXiv:2006.10759.

113. Zhang, R., Isola, P., Efros, A. A. (2018). The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.

114. Zhang, W., Zhang, Y., & Li, X. (2016). Autoencoder-Based Image Segmentation Using Multi-Scale Features and Contextual Information. Neurocomputing, 174, 123-137.

115. Zhang, X., Zhou, X., Lin, M., Sun, J. (2018). ShuffleNet: An Extremely Efficient Convolutional Neural Network for Mobile Devices. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.

116. Zhang, Y., David, P., Gong, B. (2018). Curriculum Domain Adaptation for Semantic Segmentation of Urban Scenes. Proceedings of the IEEE International Conference on Computer Vision.

117. Zhao, P., & Lu, J. (2021). Generative Data Augmentation for Multi-Modal Classification. Pattern Recognition, 118, 107513. https://doi.org/10.1016/j.patrec.2021.03.020

118. Zhao, L., Lu, J., Xu, Y., Cheng, Z., Guo, D., Niu, Y., & Fang, X. (2023). Few-shot class-incremental learning via class-aware bilateral distillation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition [Conference session]. Computer Vision Foundation.

119. Zhong, Z., Zheng, L., Kang, G., Li, S., Yang, Y. (2020). Random Erasing Data Augmentation. Proceedings of the AAAI Conference on Artificial Intelligence.

120. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., & Torralba, A. (2014). Learning Deep Features for Discriminative Localization. CVPR.

121. Zhu, J.-Y., Park, T., Isola, P., Efros, A. A. (2017). Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. arXiv preprint arXiv:1703.10593.

122. Zhu, J.-Y., Zhang, R., Pathak, D., Darrell, T., Efros, A. A., Wang, O., Shechtman, E. (2017). Toward Multimodal Image-to-Image Translation. Advances in Neural Information Processing Systems.

123. Zhu, X., Vondrick, C., Fowlkes, C. C., & Ramanan, D. (2012). Do We Need More Training Data? International Journal of Computer Vision, 106(3), 336-345.

124. Zoph, B., Vasudevan, V., Shlens, J., Le, Q. V. (2018). Learning Transferable Architectures for Scalable Image Recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.