Journal of Current Research in Engineering and Science
JCRES
Bi-Annual Online Journal (ISSN : 2581 - 611X)

Volume 7- Issue 1, January 2024
Paper : 93

Bridging the Gap in COVID-19 X-ray Datasets:Stable Diffusion-based Synthetic Image Generationfor Enhanced Diagnosis

Bhawna Jain
Computer Science Engineering
IGDTUW
Delhi, India
bhawnajain@igdtuw.ac.in
ORCID: 0009-0006-7455-2854

Saumya Mehra
Computer Science Engineering
IGDTUW
Delhi, India
saumya.mehra23@gmail.com

Nishchal Gaba
Director
Prodigal AI Technologies
Delhi, India
nishchal@prodigalai.com

Almas Ansari
ML Researcher
Prodigal AI Technologies
Delhi, India
almas@prodigalai.com

Pranav Bansal
ML Engineer
Prodigal AI Technologies
Delhi, India
gunubansal129@gmail.com

Dheeraj Verma
COO
Prodigal AI Technologies
Delhi, India
dheeraj@prodigalai.com

Gauri Maheshwari
Mechanical Engineering
IGDTUW
Delhi, India
gaurimaheshwari009@gmail.com

Ishrita Saxena
Mechanical Engineering
IGDTUW
Delhi, India
Ishrita004btmae21@igdtuw.ac.in

***Abstract*— Stable diffusion, a robust generative modeling technique, is employed to tackle the issue of limited availability of COVID-19 patient X-ray images for accurate diagnostic tools. This approach utilizes deep neural networks to learn the data distribution and generate synthetic images. By applying the Stable diffusion model, researchers aim to overcome the scarcity of annotated COVID-19 patient X-ray pictures. Through fine-tuning the model with a Kaggle dataset, they create synthetic X-ray images that closely mimic genuine patient X-rays. This synthetic data generation technique addresses the data scarcity challenge, enabling more robust model training and enhancing the accuracy of detecting and categorizing COVID-19 instances. Utilizing Stable diffusion for data augmentation involves implementing key steps and processes to effectively generate synthetic X-ray images and enhance the overall performance of machine learning models for COVID-19 diagnosis and categorization. This dataset forms the foundation for training the Stable diffusion model. Fine-tune the model to capture essential characteristics and patterns in real patient X-rays. Once fine-tuned, the model generates synthetic X-ray images resembling real patient X-rays by sampling from the learned data distribution. Synthetic images are integrated into the original dataset, creating a larger and more diverse collection of X-ray images. Finally, machine learning models are trained, like deep neural networks, using the augmented dataset. This leverages information from real and synthetic images to enhance accuracy and robustness in detecting and classifying COVID-19 cases.**

***Index Terms*—Stable Diffusion,COVID 19, Generative AI**

I. INTRODUCTION

The COVID-19 pandemic has posed unprecedented challenges to healthcare systems globally, necessitating the development of accurate and efficient diagnostic tools for the identification and monitoring of infected individuals. Among the various medical imaging modalities, X-ray imaging has emerged as a valuable technique for detecting and evaluating respiratory diseases, including COVID-19. Nevertheless, the limited availability of annotated X-ray images from COVID-19 patients presents a significant barrier in the development of reliable machine learning models for automated diagnosis. In this context, Stable diffusion, a robust generative modelling technique, offers a promising solution as stated in [5]. Stable diffusion refers to a computational approach that allows for the generation of high-quality synthetic images by fine-tuning a diffusion model with an existing dataset. The technique has garnered attention for its ability to generate realistic and diverse synthetic images across different domains. Researchers can solve the shortage of annotated X-ray pictures from COVID-19 patients by using Stable diffusion. The procedure entails fine-tuning the Stable diffusion model with a dataset gathered from sites such as Kaggle, which comprises X-ray pictures of COVID-19 patients. The model learns the underlying patterns and traits inherent in genuine patient X-rays through this fine-tuning process, allowing it to create synthetic pictures that closely match actual X-rays [7]. These simulated X-ray pictures are a great addition to the current collection, considerably increasing its size and diversity. The enhanced dataset, which includes both real patient X-rays and synthetic pictures, is subsequently used to train machine learning algorithms. Researchers can increase their accuracy and resilience in detecting and categorising COVID-19 instances by training these models on the expanded dataset [10].

To address the aforementioned issue, researchers investigated the use of Stable diffusion, a strong generative modelling tool, to produce synthetic X-ray pictures that can supplement the restricted dataset. Stable diffusion has received a lot of interest in recent years because of its capacity to produce high-quality synthetic pictures in a variety of areas. The researchers

Journal of Current Research in Engineering and Science
Bi-Annual Online Journal (ISSN : 2581 - 611X)
Volume 7- Issue 1, January 2024
Paper : 93

want to generate a larger and more diversified dataset that properly represents the main properties of genuine patient X-rays by fine-tuning the Stable diffusion model using a dataset supplied from Kaggle that includes X-ray pictures from COVID-19 patients. This enriched dataset may then be used to train and improve machine learning models built for automated COVID-19 diagnosis.

The use of stable diffusion to generate synthetic X-ray pictures has various advantages. For starters, it tackles the issue of data scarcity, as having a bigger dataset allows for more robust and accurate model training. Second, it enables the investigation of a broader spectrum of illness presentations and changes, allowing for a better knowledge of the disease's radiological properties. Furthermore, the produced dataset allows you to test the generalizability of machine learning models by introducing variables that were not included in the original dataset [2].

The quality and realism of the generated synthetic X-ray images are evaluated through visual inspection and quantitative metrics. Furthermore, the impact of augmenting the dataset with synthetic images on the performance of machine-learning models for COVID-19 diagnosis is investigated. The study aims to contribute to the field of medical imaging and AI-assisted diagnosis by showcasing the potential of Stable diffusion in addressing data scarcity and enhancing the accuracy and reliability of automated COVID-19 detection systems.

Clinical Importance: The development of accurate and efficient diagnostic tools for COVID-19 is of paramount importance in effectively managing the ongoing pandemic. Medical imaging, such as X-ray imaging, plays a crucial role in the detection and monitoring of respiratory diseases, including COVID-19. However, the limited availability of annotated X-ray images from COVID-19 patients hampers the development of reliable AI-based diagnostic systems.

The clinical importance of this research lies in its potential to address the scarcity of annotated data by leveraging Stable diffusion, a generative modeling technique [4]. By generating synthetic X-ray images that closely resemble real patient X-rays, we can significantly augment the available dataset. This augmented dataset enables the training of machine learning models on a larger and more diverse set of COVID-19 cases, thereby improving their accuracy and reliability in automated diagnosis.

The use of Stable diffusion to generate synthetic X-ray images offers valuable clinical implications [1]. Firstly, it enhances the robustness and generalizability of AI models by providing a more comprehensive representation of COVID-19 manifestations. This, in turn, improves the accuracy of diagnosis, enabling healthcare professionals to make informed decisions regarding patient management and treatment. Moreover, the availability of a larger dataset through synthetic image generation allows for the exploration of rare or subtle radiological features that may be challenging to encounter in the original dataset. This expanded understanding of disease variations contributes to the refinement of diagnostic criteria. [9]

The remainder of the paper is organised as follows:

Section II delves into the existing literature and previous studies on the use of medical imaging and the challenges associated. Section III is Implementation which contains the details,methodology and steps involved in utilizing Stable diffusion for generating synthetic X-ray images. Section IV presents the findings of the research. It includes the evaluation of the quality andrealism of the generated synthetic X-ray images. Section V discusses the limitations and obstacles encountered during the research process. It addresses the potential difficulties and complexities in implementing Stable diffusion for data augmentation. Section VI concludes the paper by summarizing the findings on stable diffusion in generating synthetic X- ray and its effect on detecting COVID-19 cases. Section VII outlines potential avenues for further work and development.

II. Background and Related Work

Stable diffusion operates by iteratively applying a series of diffusion steps to an initial noise vector, gradually transforming it into a realistic image. Each diffusion step involves a controlled blending of the current image with a noise signal, guided by a diffusion model. This diffusion model, typically based on deep neural networks, captures the intricate mapping between the noisy image and the target distribution, effectively learning the underlying data distribution from the available training data. It discusses the opportunities for extending the application of Stable diffusion beyond COVID-19 to other respiratory diseases and imaging modalities [6].

Unlike traditional generative models such as Variational Autoencoders (VAEs) [2] or Generative Adversarial Networks (GANs), Stable diffusion avoids the mode collapse issue and produces diverse and high-fidelity samples. It achieves this by employing a diffusion process that enables fine-grained control over the level of noise in the generated images. By gradually reducing the noise level during the diffusion process, Stable diffusion ensures that the generated images exhibit increasingly realistic and precise features. The original stable diffusion model is shown in Figure 1.

The Stable diffusion framework is employed in this research to generate synthetic X-ray images that closely resemble real patient X-rays. By fine-tuning the Stable diffusion model on a dataset of X-ray images from COVID-19 patients, the model captures the underlying patterns and characteristics of the disease [3]. This allows for the generation of synthetic images that exhibit realistic radiological features specific to COVID-19, including opacities, consolidations, and ground-glass opacities. By utilizing the Stable diffusion framework, the researchers effectively expand the dataset and create a more comprehensive representation of COVID-19 manifestations, thereby enhancing the training of machine learning models for automated diagnosis. [8]

By leveraging the capabilities of Stable diffusion, the limitations imposed by the available dataset are effectively overcome, and a significant number of additional images capturing the diverse variations of COVID-19 in X-ray imaging are generated. This innovative approach proves instrumental in

Journal of Current Research in Engineering and Science
Bi-Annual Online Journal (ISSN : 2581 - 611X)
Volume 7- Issue 1, January 2024
Paper : 93

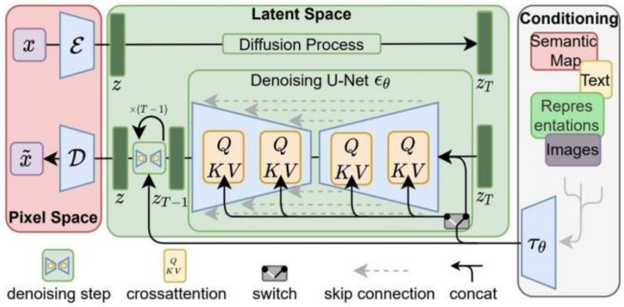

Fig. 1. Original Stable Diffusion Model

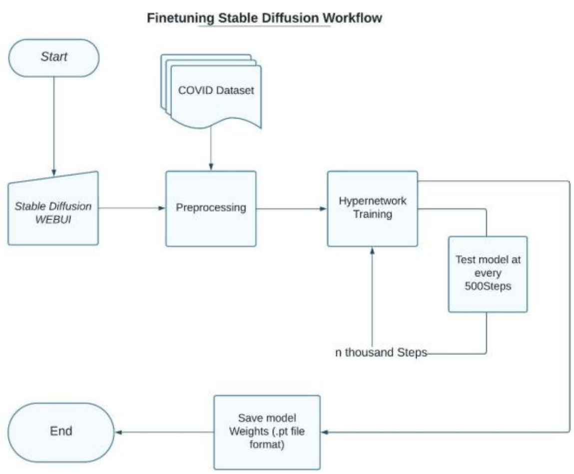

Fig. 2. Flowchart for Stable Diffusion Model of this Study

enhancing the performance and robustness of AI-based diagnostic systems, thereby enabling more accurate and reliable detection of COVID-19 cases.

III. Implementation

The implementation of the research involved utilizing the Stable diffusion repository available on GitHub as the foundation. This repository provided a well-established framework for generative modeling using Stable diffusion.

To train the model on a dataset of COVID-19 X-ray images, a curated dataset from Kaggle, consisting of images from COVID-19 patients, was obtained. The dataset was preprocessed to ensure uniform image sizes and incorporate appropriate data augmentation techniques, enhancing the model's robustness and generalization capabilities.

Subsequently, the Stable diffusion model was fine-tuned by adjusting its hyperparameters and optimizing its architecture to align with the specific characteristics of the dataset. To maximise the model's ability to create realistic synthetic X-ray pictures, an iterative fine-tuning approach was used.

After fine-tuning the model, it was used to produce synthetic pictures. The learnt weights and latent space representations were used to generate sets of 100 synthetic pictures at a time using a batch generation technique. This iterative generation approach allowed for the rapid production of a large number of different and realistic synthetic X-ray pictures. The technique flowchart used in this investigation is shown in Figure 2.

The quality and authenticity of the produced photographs were constantly monitored throughout the installation. To confirm that the synthetic pictures properly reflected the major radiological aspects of COVID-19, visual inspections, quantitative analyses, and comparisons with genuine patient X-rays were performed. The goal of this rigorous assessment method was to confirm the realism and authenticity of the synthetic X-ray pictures produced by the Stable diffusion model.

Powerful computational resources, including GPUs, were used to facilitate implementation and expedite the generating process. This enabled efficient and effective model inference and picture production, allowing for the creation of a large number of synthetic images. The Stable diffusion model was fine-tuned using a COVID-19 X-ray dataset, followed by batch-wise production of realistic synthetic pictures. The successful augmentation of the dataset revealed Stable diffusion's potential for creating useful synthetic data for enhancing COVID-19 diagnosis. The use of modern computer resources was critical in assuring the implementation's efficiency and scalability.

IV. Results

A cooperation with a medical practitioner with experience in diagnosing COVID-19 was created to test the correctness of the generated synthetic X-ray pictures. The generated images, produced using the Stable diffusion model, were presented to two professional doctors for their evaluation. With meticulous attention, the doctors thoroughly examined the images, conducting a comparative analysis with real patient X-rays. These doctors were given two parameters - Good and Bad for determining the quality of the X-ray images. Their professional expertise and insights proved invaluable in assessing the quality and fidelity of synthetic images. Examples of a good and bad augmented image are shown in Figure 3.

Based on the doctor's assessment, it was determined that approximately 68 percent of the generated images accurately represented the characteristic features of COVID-19 patient lungs in X-ray imaging. This result shown in Figure 4 demonstrates the potential of the Stable diffusion approach in generating synthetic X-ray images that closely resemble real cases. All the augmented images by stable diffusion model are collectively shown in Figure 5.

It is important to note that the evaluation conducted in this research involved subjective assessment by two medical professionals. While the doctor's expertise adds credibility to the evaluation, a larger and more diverse panel of medical experts would be necessary for a comprehensive and objective evaluation of the accuracy and reliability of the generated images.

Nevertheless, these initial findings provide a promising indication of the effectiveness of the Stable diffusion model in

Journal of Current Research in Engineering and Science
Bi-Annual Online Journal (ISSN : 2581 - 611X)

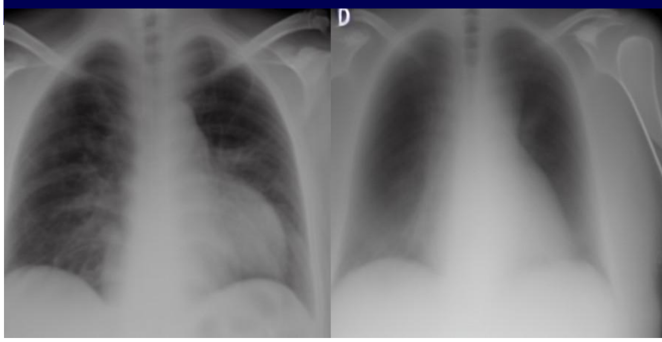

Fig. 3. Good augmented image(Left) and Badly augmented Image(Right)

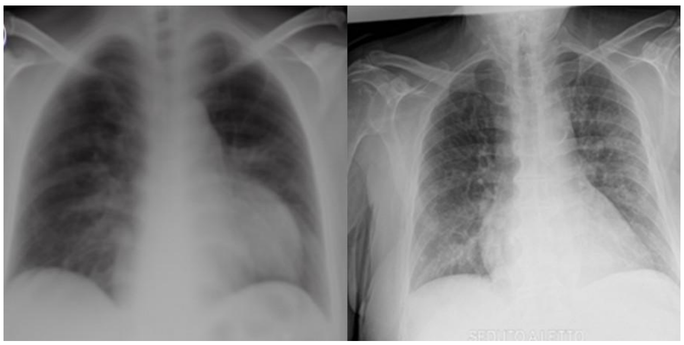

Fig. 4. Generated image(Left) and Original Image(Right)

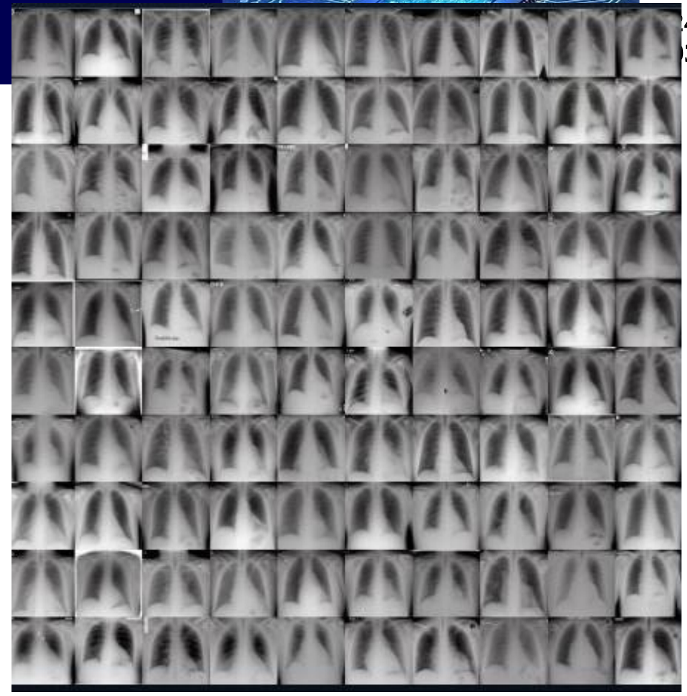

Fig. 5. Augmented images

generating synthetic X-ray images representative of COVID-19 cases. This result underscores the potential of generative modeling techniques to address data scarcity and augment the training datasets for improved COVID-19 diagnosis.

Further research and collaborations with a larger medical expert panel, along with quantitative evaluation metrics, are needed to validate the diagnostic accuracy and clinical utility of the synthetic X-ray images generated by the Stable diffusion approach. These efforts will contribute to enhancing the reliability and trustworthiness of AI-assisted diagnostic systems and their potential impact on patient care in the context of COVID-19.

V. Challenges

While the application of Stable diffusion for generating synthetic X-ray images in COVID-19 diagnosis shows promising results, there are several avenues for challenges that should be addressed. These include:

1. **Incorporating Clinical Metadata**: Future work can involve incorporating clinical metadata, such as patient demographics, comorbidities, and laboratory results, into the Stable diffusion framework. By incorporating this additional information, the generated synthetic images can better capture the complex relationships between imaging features and patient characteristics, leading to more accurate and personalized diagnostic models.

2. **Uncertainty Estimation**: Estimating and quantifying uncertainty in the generated synthetic images is an essential aspect to consider. Future research can focus on developing techniques to provide reliable uncertainty estimates in the Stable diffusion framework. This will enhance the interpretability and trustworthiness of the generated images, allowing clinicians to make more informed decisions based on the AI predictions.

3. **Evaluation on Diverse Data Sources:** The generalizability of the Stable diffusion model and the effectiveness of the generated synthetic images can be further evaluated on diverse data sources and external validation datasets. This will help assess the robustness and reliability of the approach across different healthcare systems, imaging equipment, and populations.

4. **Ethical Considerations and Bias Mitigation:** As with any AI-based technology, ethical considerations and bias mitigation are crucial aspects to address. Future work should focus on investigating potential biases in the generated synthetic images and developing strategies to ensure fairness, transparency, and equity in the diagnostic models trained on these images.

5. **Integration into Clinical Workflow:** Successfully integrating AI-generated synthetic images into the clinical workflow poses practical challenges. Future research can explore strategies for seamless integration, considering factors such as data privacy, regulatory compliance, and clinician acceptance. This will facilitate the translation of Stable diffusion-generated images into real-world clinical practice.

6. **Data Privacy and Security:** The privacy of patient data and the security of the created synthetic pictures are critical. Future research should focus on developing effective data anonymization and privacy preservation strategies to ensure patient confidentiality while using the benefits of generative modelling.

7. **Reliability of the evaluation:** A medical practitioner is used to verify results. It would have been much more efficient and reliable if a matrix was available

Journal of Current Research in Engineering and Science
Bi-Annual Online Journal (ISSN : 2581 - 611X)
Volume 7- Issue 1, January 2024
Paper : 93

as a credible benchmark to provide a more impartial evaluation of the correctness of the generated synthetic X-ray pictures. This matrix, which includes proven COVID-19 instances and X-ray pictures, would be gathered from a variety of sources, including reliable medical databases and expert comments.

The amount of authenticity and realism attained by comparing the synthetic pictures created by the Stable diffusion model can be assessed reference matrix. This method allows us to objectively examine the correctness of the produced pictures, minimising potential biases and subjective interpretations.

The solicited advice and experience of medical authorities that specialise in COVID-19 diagnosis was used. These specialists analysed the synthetic photos, offering useful insights and opinions. Their experience contributed to the review process by validating the veracity of the synthetic photographs.

A thorough assessment system that can strengthen the reliability and correctness of the generated synthetic X-ray pictures by can be guaranteed by combining objective comparison against thereference matrix with the experience of medical experts. This comprehensive assessment technique will increase the validity of our findings and conclusions on the Stable diffusion model's effectiveness and promise in resolving data shortage and enhancing COVID-19 diagnosis.

Addressing these future work areas and problems will progress the field of AI-based medical imaging, improve diagnostic model reliability, and make it easier to integrate Stable diffusion-generated synthetic pictures into clinical practise. By addressing these issues, we will be able to fully utilise the promise of generative modelling approaches in enhancing healthcare outcomes and patient care.

VI. Conclusion

The use of stable diffusion in the generation of synthetic X-ray pictures for COVID-19 diagnosis is a significant strategy for overcoming data shortages and improving the performance of AI-based diagnostic systems. We were able to develop realistic synthetic pictures that closely mimic genuine patient X-rays by fine-tuning the Stable diffusion model using a small dataset. These synthetic photos are a useful addition to the collection, allowing for more robust training and enhancing COVID-19 detection accuracy.

Stable diffusion's promise goes beyond COVID-19, providing chances for multi-modal data synthesis, modelling of uncommon illnesses, and data augmentation in a variety of respiratory conditions. However, obstacles remain, including as adding clinical metadata, calculating uncertainty, correcting biases, and assuring ethical concerns and data protection.

Despite these obstacles, the incorporation of Stable diffusion-generated synthetic pictures into clinical practise holds significant potential for future advancements in AI-assisted medical imaging and patient care. We may improve the capabilities of AI-based diagnostic systems by exploiting the benefits of stable diffusion, allowing for more accurate and efficient identification of respiratory disorders, eventually leading to improved patient outcomes.

VII. Future Work

The Stable diffusion technique, in addition to solving the lack of annotated X-ray images for COVID-19 diagnosis, provides up prospects for additional study and applications in the medical imaging sector. In this section, we outline potential future projects that may be investigated using the same approach:

1. **Multi-modal Data Synthesis:** Stable diffusion may be extended to provide synthetic data in a variety of modalities, including computed tomography (CT) scans and magnetic resonance imaging (MRI). This extension enables the synthesis of various medical pictures, allowing the construction of complete diagnostic models that can use several imaging modalities.
2. **Rare Disease Simulation:** Stable diffusion can be utilised to produce synthetic images of rare illnesses or circumstances with little data. By duplicating these rare events, the approach can aid in the development and training of machine learning models for accurate diagnosis and treatment planning.
3. **Data Augmentation for Other Respiratory Diseases:**Stable diffusion's use is not restricted to COVID-19. It may be used to create synthetic pictures for various respiratory disorders such pneumonia, TB, and lung cancer. The method, by supplementing current information, can contribute in the creation of more complete and reliable diagnosis models for a variety of respiratory disorders.
4. **Anomaly Detection and Outlier Identification:** Stable diffusion may be used to produce synthetic pictures of normal or healthy states. This may be used to discover anomalies or outliers in medical imaging, allowing for early detection of probable abnormalities or illnesses that may not have been fully represented in the original dataset.
5. **Transfer Learning and Domain Adaptation:**Stable diffusion may be used for domain adaptation tasks, where models trained on synthetic pictures created by the method can be fine-tuned on real-world data. When applied to diverse imaging datasets or healthcare situations, this transfer learning process can improve the generalizability and performance of AI models.

The Stable diffusion methodology can contribute to the development of robust AI-based diagnostic tools, data augmentation approaches, and anomaly detection methods in the field of medical imaging by investigating these connected study fields. Beyond COVID-19 diagnosis, the potential applications are numerous, addressing different healthcare difficulties and improving patient care across a wide spectrum of illnesses and ailments.

References

[1] Vegard Antun, Francesco Renna, Clarice Poon, Ben Adcock, and Anders C Hansen. On instabilities of deep learning in image reconstruction and the potential costs of ai. *Proceedings of the National Academy of Sciences*, 117(48):30088–30095, 2020.

Journal of Current Research in Engineering and Science
Bi-Annual Online Journal (ISSN : 2581 - 611X)
Volume 7- Issue 1, January 2024
Paper : 93

[2] Samuel G Armato III, Geoffrey McLennan, Luc Bidaut, Michael F McNitt-Gray, Charles R Meyer, Anthony P Reeves, Binsheng Zhao, Denise R Aberle, Claudia I Henschke, Eric A Hoffman, et al. The lung image database consortium (lidc) and image database resource initiative (idri): a completed reference database of lung nodules on ct scans. *Medical physics*, 38(2):915–931, 2011.

[3] Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin S Kirby, John B Freymann, Keyvan Farahani, and Christos Davatzikos. Advancing the cancer genome atlas glioma mri collections with expert segmentation labels and radiomic features. *Scientific data*, 4(1):1–13, 2017.

[4] Hyungjin Chung and Jong Chul Ye. Score-based diffusion models for accelerated mri, 2022.

[5] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021.

[6] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021.

[7] Amirhossein Kazerouni, Ehsan Khodapanah Aghdam, Moein Heidari, Reza Azad, Mohsen Fayyaz, Ilker Hacihaliloglu, and Dorit Merhof. Diffusion models in medical imaging: A comprehensive survey. *Medical Image Analysis*, 88:102846, 2023.

[8] Agnieszka Mikołajczyk and Michał Grochowski. Data augmentation for improving deep learning in image classification problem. In *2018 International Interdisciplinary PhD Workshop (IIPhDW)*, pages 117–122, 2018.

[9] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

[10] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models, 2022.