Uncertainty Quantification of GAN-Based Image Inpainting of Material's Stress Fields using Monte Carlo Dropout

By

Sardar Rehan Yasin
(Registration No: 00000402160)

Department of Engineering

School of Interdisciplinary Engineering and Sciences

National University of Sciences and Technology (NUST)

Islamabad, Pakistan

(2024)

Uncertainty Quantification of GAN-Based Image Inpainting of Material's Stress Fields using Monte Carlo Dropout

By

Sardar Rehan Yasin

(Registration No: 00000402160)

A thesis submitted to the National University of Sciences and Technology, Islamabad,

in partial fulfillment of the requirements for the degree of

Master of Science in

Computational Science and Engineering

Supervisor: Dr. Absaar Ul Jabbar

School of Interdisciplinary Engineering and Sciences

National University of Sciences and Technology (NUST)

Islamabad, Pakistan

(2024)

Annex A to NUST Letter No. 0972/102/Exams/Thesis-Cert dated 23 Dec 16.

THESIS ACCEPTANCE CERTIFICATE

Certified that final copy of MS/MPhil thesis written by Mr/Ms Sardar Rehan Yasin Registration No. 00000402160 of SINES has been vetted by undersigned, found complete in all aspects as per NUST Statutes/Regulations, is free of plagiarism, errors, and mistakes and is accepted as partial fulfillment for award of MS/MPhil degree. It is further certified that necessary amendments as pointed out by GEC members of the scholar have also been incorporated in the said thesis.

Signature with stamp: _____
Assistant Professor
SINES - NUST, Sector H-12
Islamabad

Name of Supervisor: Dr. Ahsan Ul Jabbar

Date: 19 Nov 2024

Signature of HoD with stamp: _____
Dr. Mian Ilyas Ahmad
HOD Engineering
Professor
SINES - NUST, Sector H-12
Islamabad

Date: 28 Nov 2024

Countersign by

Signature (Dean/Principal): _____
Dr. SYED IRTIZA ALI SHAH
Principal & Dean
SINES - NUST, Sector H-12
Islamabad.

Date: 28 NOV 2024

AUTHOR'S DECLARATION

I Sardar Rehan Yasin hereby state that my MS thesis titled "Uncertainty Quantification of GAN-Based Image Inpainting of Material's Stress Fields using Monte Carlo Dropout" is my own work and has not been submitted previously by me for taking any degree from National University of Sciences and Technology, Islamabad or anywhere else in the country/ world.

At any time if my statement is found to be incorrect even after I graduate, the university has the right to withdraw my MS degree.

Student Signature:_____

Name:_____

Date: _____

PLAGIARISM UNDERTAKING

I solemnly declare that research work presented in the thesis titled "Uncertainty Quantification of GAN-Based Image Inpainting of Material's Stress Fields using Monte Carlo Dropout" is solely my research work with no significant contribution from any other person. Small contribution/ help wherever taken has been duly acknowledged and that complete thesis has been written by me.

I understand the zero tolerance policy of the HEC and National University of Sciences and Technology (NUST), Islamabad towards plagiarism. Therefore, I as an author of the above titled thesis declare that no portion of my thesis has been plagiarized and any material used as reference is properly referred/cited.

I undertake that if I am found guilty of any formal plagiarism in the above titled thesis even after award of MS degree, the University reserves the rights to withdraw/revoke my MS degree and that HEC and NUST, Islamabad has the right to publish my name on the HEC/University website on which names of students are placed who submitted plagiarized thesis.

Student Signature:_____

Name:_____

Date: _____

THIS RESEARCH IS DEDICATED TO ALL THOSE GREAT MINDS WHO HAVE LAID THE FOUNDATIONS OF SCIENCE AND ENGINEERING.

ACKNOWLEDGEMENTS

First and foremost, I would like to express my gratitude to Allah (S.W.T) for making me one of the seekers of knowledge. I am extremely greatful to my mother and father for being the torch bearers in my life. I could not have achieved any milestone without the constant support of my siblings.

I would also like to acknowledge my supervisor Dr. Absaar Ul Jabbar and my G.E.C. members for guiding me throughout my research work and improving my research skills. My sincere appreciation goes to my dear friends who stood by me during challenging times and raised my spirits. Additionally, I am thankful to the technical and administrative staff at SINES, NUST, for promptly addressing all my queries.

Last but not the least, I would like to thank me, for believing in me, for doing all this hard work, for having no days off, for never quitting, for being me at all times.

Contents

List of Tables

List of Figures

LIST OF SYMBOLS, ABBREVIATIONS AND ACRONYMS

GANs: Generative Adversarial Networks
NDT: Non-destructive Testing
AI: Artificial Intelligence
BDL: Bayesian Deep Learning
MC Dropout: Monte Carlo Dropout
MCMC: Markov Chain Monte Carlo
VI: Variational Inference
ELBO: Evidence Lower Bound
VAEs: Variational Autoencoders
DGPs: Deep Gaussian Processes
CNNs: Convolutional Neural Networks
RNNs: Recurrent Neural Networks
BNNs: Bayesian Neural Networks
UQ: Uncertainty Quantification
FEA: Finite Element Analysis
DFT: Density Functional Theory
cGANs: Conditional Generative Adversarial Networks
DCGANs: Deep Convolutional Generative Adversarial Networks
WGANs: Wasserstein Generative Adversarial Networks
SN-PatchGAN: Spectral Normalized Patch Generative Adversarial Network
MNPS: Multi-scale Neural Patch Synthesis
SPDNorm: Spatially Probabilistic Diversity Normalization
DE-GAN: Domain Embeded Generative Adversarial Network
WGAN-GP: Wasserstein Generative Adversarial Network with Gadient Penalty
lsGAN: Least Squared Generative Adversarial Network
MaskGIT: Masked Generative Image Transformer
MVTM: Masked Visual Token Modeling

CCDF: Come Closer Diffuse Faster
MRI: Magnetic Resonance Imaging
MCG: Manifold Constrained Gradient
CT: Computerized Tomography
DDPM: Denoising Diffusion Probabilistic Model
SSIM: Structural Similarity Index Measure
TES: Total En Bloc Spondylectomy
NNs: Neural Networks
RMSE: Root Mean Squared Error
U2AFN: Uncertainty Aware Adaptive Feedback Network
PSNR: Peak Signal to Noise Ratio
PEM: Principal Eigen value Moments
GPR: Gaussian Process Regression
BUQO: Bayesian Uncertainty Quantification
ReLU: Rectified Linear Unit
MAE: Mean Absolute Error

Abstract

This research focuses on uncertainty quantification in reconstructing incomplete mechanical stress field maps by utilizing Generative Adversarial Networks (GANs) integrated with dropout mechanisms to introduce noise and variability. Specifically, the study investigates the effects of selectively applying dropout as a source of controlled noise to a GAN-based model known as DeepFill. Two distinct configurations were examined in terms of dropout application: one where dropout was confined exclusively to the first hidden layer and another where it was uniformly distributed across all hidden layers. The dropout rates were varied at 0.2, 0.4, and 0.6 to assess their impact on model performance and reliability. Results demonstrate that applying dropout exclusively to the first hidden layer yields lower uncertainties, even for higher dropout rates. By restricting the introduction of noise to the initial hidden layer, the network benefits from regularization without significantly impacting the stability of deeper layers. This selective application ensures that noise is managed in a controlled manner, allowing the model to generalize better while maintaining consistent predictions. In contrast, applying dropout across all hidden layers amplifies noise throughout the network, leading to larger uncertainties. Moreover, it is observed that the uncertainty increases linearly with the increase in dropout rates. These findings are particularly significant for applications in stress field reconstructions, such as fracture mechanics and non-destructive testing (NDT).

Keywords: Uncertainty Quantification, Inpainting, Mechanical Stress Fields Reconstruction, GANs

Chapter 1

Introduction

"All models are wrong, some are useful", George E.P. Box. This quote questions the reliability of AI models. All AI based models are not hundred percent accurate due to certain reasons which include but are not limited to poor feature learning by the model, insufficient data used for training, out of distribution test cases, etc. Therefore, every prediction made by an AI model has a certain level of uncertainty associated. This uncertainty can manifest in the form of confidence intervals, probabilistic outputs, or variability. It highlights the range of possible outcomes and the likelihood of each one rather than providing a single, definitive answer. Understanding uncertainty is critical for engineers and researchers, as it directly influences decision-making processes, risk assessment, and the reliability of simulations.

1.1 Introduction to Uncertainty

Uncertainty is a fundamental concept that pervades various fields, from science and engineering to finance and healthcare. It refers to the degree of confidence or doubt about a particular outcome or measurement.

1.1.1 Types of Uncertainty

Uncertainty can be broadly categorized into two types: **aleatoric** and **epistemic** [3].

i) Aleatoric Uncertainty Aleatoric uncertainty refers to the inherent variability or randomness present in a system or process. It arises from the natural fluctuations in phenomena and is often characterized by unpredictable outcomes.

Sources: This type of uncertainty can originate from various factors, such as:

- Environmental Variability: Changes in external conditions, like temperature, humidity, or pressure, can introduce randomness into experimental results.
- Measurement Errors: Limitations in measuring instruments can lead to variations in data due to instrument precision or human error.
- Natural Variability: Inherent differences in populations, materials, or systems can lead to variability in outcomes.

Characteristics:

- Statistical Nature: Aleatoric uncertainty is typically described using statistical methods. It can be quantified using probability distributions, such as normal or binomial distributions, to capture the range of possible outcomes.
- Irreducibility: This type of uncertainty cannot be reduced or eliminated; it is an intrinsic part of the system. However, it can be characterized and understood through statistical analysis and data collection.

ii) Epistemic Uncertainty Epistemic uncertainty arises from a lack of knowledge or information about a system or process. It reflects the uncertainty due to incomplete understanding or inadequate models.

Sources: Epistemic uncertainty can stem from various factors, including:

- Model Uncertainty: Simplifications or assumptions made in models can lead to inaccuracies. For example, using an oversimplified model to represent a complex phenomenon can introduce significant epistemic uncertainty.
- Parameter Uncertainty: Uncertainties in the values of model parameters, which may be estimated from limited data, contribute to epistemic uncertainty. For example, if the input parameters in a predictive model are not known with high confidence, the predictions made by that model will reflect epistemic uncertainty.
- Data Limitations: Insufficient or low-quality data can lead to a lack of confidence in the predictions or insights drawn from analyses. If data is missing, outdated, or noisy, it can skew the understanding of the system.

Characteristics:

- Reducibility: Unlike aleatoric uncertainty, epistemic uncertainty can often be reduced through additional research, improved data collection, and enhanced modeling techniques. Gathering more data or refining models can help diminish this type of uncertainty.

- Subjective Nature: Epistemic uncertainty can also be influenced by the knowledge, experience, and judgment of individuals or groups involved in the analysis. This subjectivity can lead to variability in how different analysts perceive and quantify uncertainty.

In the context of machine learning, understanding the types of uncertainty—aleatoric and epistemic—is essential for developing robust models that can make reliable predictions. **Aleatoric uncertainty** pertains to the noise inherent in the data itself, reflecting variability that arises from factors such as measurement errors or natural fluctuations within the dataset. This type of uncertainty can be modeled using probabilistic frameworks, allowing the machine learning algorithms to predict a distribution of possible outcomes rather than a single deterministic value. On the other hand, **epistemic uncertainty** emerges from limitations in the model or insufficient knowledge about the data-generating process. This uncertainty is often a result of the model's architecture, parameter choices, or the volume and quality of the training data. Techniques such as Bayesian inference, Monte Carlo dropout, and ensemble methods are commonly employed to quantify epistemic uncertainty, enabling models to express their confidence in predictions. By effectively distinguishing between these two types of uncertainty, machine learning practitioners can better assess the reliability of model outputs and improve decision-making processes in applications ranging from healthcare to autonomous systems.

1.1.2 Methods to Calculate Uncertainty

Uncertainty can be estimated via various techniques. This section gives an overview of most common ones used.

i) Bayesian Deep Learning (BDL)
Bayesian Deep Learning integrates Bayesian principles into deep learning frameworks to quantify uncertainty in predictions. In BDL, the model parameters are treated as random variables with prior distributions, rather than fixed values. This approach allows the model to express uncertainty regarding its predictions, as it can output a distribution over possible outcomes instead of a single point estimate. By updating the prior distribution with observed data through Bayes' theorem, BDL captures both aleatoric and epistemic uncertainty, enabling better decision-making in applications where uncertainty is critical, such as medical diagnosis or autonomous systems.

ii) Monte Carlo Dropout (MC Dropout)
Monte Carlo Dropout is a technique for estimating uncertainty in neural networks by applying dropout during inference, not just during training. By randomly dropping

out a subset of neurons in the network, MC Dropout generates multiple stochastic forward passes for the same input, yielding a distribution of predictions. The mean of these predictions serves as the model's output, while the standard deviation provides an estimate of uncertainty. This approach is particularly useful for quantifying epistemic uncertainty and is computationally efficient, as it leverages existing dropout mechanisms already employed during training.

iii) Markov Chain Monte Carlo (MCMC)
Markov Chain Monte Carlo [4] is a family of algorithms used for sampling from probability distributions when direct sampling is challenging. MCMC methods create a Markov chain that converges to the target distribution, allowing for the estimation of its properties. The most common MCMC method is the Metropolis-Hastings algorithm, which generates samples through a random walk. MCMC is valuable in Bayesian statistics for posterior inference, enabling the computation of complex integrals and uncertainty quantification in models with high dimensional parameter spaces.

iv) Variational Inference (VI)
Variational Inference is a technique used to approximate complex posterior distributions in Bayesian models. Rather than sampling from the posterior directly, VI transforms the problem into an optimization task, where a simpler, parameterized distribution is fitted to approximate the true posterior. By minimizing the Kullback-Leibler divergence between the two distributions, VI enables efficient inference in large-scale models and provides a deterministic way to quantify uncertainty. VI is widely used in deep learning and graphical models due to its scalability and computational efficiency.

v) Bayesian Active Learning
Bayesian Active Learning combines Bayesian inference with an active learning framework to iteratively select the most informative samples for labeling. In this approach, the model estimates uncertainty in its predictions and identifies instances where the model is most uncertain. By querying the oracle (e.g., a human annotator) for labels on these uncertain samples, the model can improve its learning efficiency and reduce the number of labeled instances required. This technique is especially beneficial in scenarios with limited labeled data, such as medical imaging or natural language processing.

vi) Bayes by Backprop
Bayes by Backprop [5] is a method that combines the principles of Bayesian inference with backpropagation in neural networks. It involves learning a distribution over the weights of the network instead of point estimates, allowing for uncertainty quantification in predictions. During training, the weights are updated using a variational approach, optimizing the evidence lower bound (ELBO) to approximate the posterior

distribution. This technique enables the model to capture uncertainty and is suitable for applications where interpretability and robustness are essential.

vii) Variational Autoencoders (VAEs)
Variational Autoencoders are a type of generative model that combines neural networks with variational inference. VAEs learn a latent representation of input data by encoding it into a lower dimensional space, where the latent variables are treated as random variables with a prior distribution. The decoder reconstructs the input from the latent space, and the model is trained to maximize the ELBO, balancing reconstruction accuracy and the KL divergence between the learned and prior distributions. VAEs provide a powerful framework for uncertainty quantification, enabling the generation of new samples and capturing the variability inherent in the data.

viii) Deep Gaussian Processes (DGPs)
Deep Gaussian Processes [6] extend the concept of Gaussian Processes (GPs) to multiple layers, creating a hierarchical model that captures complex, nonlinear relationships in data. In a DGP, each layer is a Gaussian Process, and the outputs of one layer serve as the inputs to the next. This structure allows DGPs to model intricate functions with a rich set of latent variables, capturing uncertainty at various levels of abstraction.

The primary advantage of DGPs is their ability to learn complex distributions over functions, providing a flexible framework for regression and classification tasks. The uncertainty estimates derived from DGPs reflect both aleatoric (inherent noise in the observations) and epistemic (uncertainty due to model limitations) uncertainty. However, DGPs can be computationally intensive and challenging to train, often requiring approximate inference techniques for practical applications.

ix) Laplace Approximations
Laplace Approximations are a method for approximating complex posterior distributions in Bayesian inference. The basic idea is to find a Gaussian approximation to the posterior by identifying the mode (maximum a posteriori estimate) and approximating the curvature of the log posterior around that point using the Hessian matrix. This results in a Gaussian distribution that captures the uncertainty around the mode.

Laplace Approximations are computationally efficient and provide a quick way to derive uncertainty estimates for model parameters. They are particularly useful in high-dimensional spaces where direct sampling or other more complex inference methods (like MCMC) may be computationally prohibitive. However, the accuracy of Laplace Approximations relies on the assumption that the posterior is approximately Gaussian, which may not hold in cases with multimodal distributions or significant nonlinearity.

1.1.3 Uncertainty Quantification using Monte Carlo Dropout Method

Monte Carlo methods, particularly Monte Carlo dropout (MC dropout), are widely used for uncertainty quantification in machine learning due to several advantages that make them effective and appealing compared to other methods:

i) Simplicity and Ease of Implementation
MC dropout is straightforward to implement, requiring minimal modifications to existing neural network architectures. By simply enabling dropout during inference, it can generate multiple predictions for the same input, allowing for the estimation of uncertainty. This makes it accessible to a wide range of practitioners and researchers without requiring extensive changes to model training procedures.

The authors in [7] show that MC dropout can be interpreted as a variational inference method, representing a simple and efficient way to incorporate uncertainty estimation into existing neural networks without complex modifications.

ii) Robustness to Overfitting
By leveraging dropout, MC dropout inherently acts as a regularizer, which helps reduce overfitting. This property can enhance model generalization, particularly in cases with limited data. It captures the model's uncertainty while still promoting robust feature learning.

This seminal paper [8] on dropout illustrates how dropout layers help reduce overfitting in neural networks, thereby enhancing generalization and indirectly supporting the robustness of MC dropout for uncertainty quantification.

iii) Flexibility Across Models
MC dropout can be applied to various neural network architectures, including convolutional neural networks (CNNs) and recurrent neural networks (RNNs). This versatility allows it to be used across different domains and types of problems, making it a valuable tool in the machine learning toolbox.

The study [7] emphasizes the applicability of MC dropout across various neural network architectures, highlighting its flexibility for different types of machine learning tasks.

iv) Effective Uncertainty Estimation MC dropout effectively captures both epistemic (model uncertainty) and aleatoric (data noise) uncertainties. By producing a distribution of predictions, it provides not just point estimates but also a measure of

uncertainty for each prediction. This is critical in applications where understanding uncertainty is essential for decision-making, such as in medical diagnosis or autonomous driving.

This work [9] compares MC dropout with other Bayesian methods and shows that MC dropout effectively estimates both epistemic and aleatoric uncertainty in tasks involving image classification.

v) Empirical Success
Numerous studies have demonstrated the empirical success of MC dropout in various tasks, showing it often outperforms other uncertainty quantification methods like Bayesian Neural Networks (BNNs) [9] and Variational Inference (VI) in practice. Its ability to provide reliable uncertainty estimates has been validated in different settings, contributing to its popularity.

vi) Computational Efficiency
While other methods, such as Markov Chain Monte Carlo (MCMC) or Variational Inference, can be computationally intensive and time-consuming, MC dropout offers a more efficient alternative. It does not require complex sampling or optimization procedures, making it feasible for large datasets and complex models.

This research [10] illustrates the computational efficiency of MC dropout compared to more complex Bayesian methods such as MCMC, emphasizing the feasibility of MC dropout for large scale problems.

vii) Direct Interpretation
The outputs from MC dropout, including mean predictions and standard deviations, can be directly interpreted as measures of confidence in predictions. This interpretability is crucial for many applications where understanding the model's confidence is as important as the predictions themselves.

The authors in [11] present a comparative analysis of MC dropout and deep ensembles, noting that MC dropout provides interpretable uncertainty estimates (mean predictions and standard deviations) that are critical for decision-making processes.

1.1.4 Importance of Uncertainty Quantification

Uncertainty quantification (UQ) is a critical aspect of modern scientific and engineering applications, particularly in fields such as machine learning, materials science, and risk assessment. The significance of UQ arises from its ability to provide a comprehensive

understanding of the reliability and robustness of predictive models. This section discusses the importance of UQ in various domains, its implications for decision-making, and its role in advancing scientific research.

i) Enhancing Model Reliability
UQ plays a crucial role in enhancing the reliability of predictive models by assessing the impact of uncertainties in input parameters and model structures. In many real-world applications, data is often noisy, incomplete, or subject to variability. By quantifying uncertainty, researchers can identify the range of possible outcomes, enabling them to better understand model behavior under different conditions. This knowledge is vital for developing robust models that can withstand variations in data and parameters, ultimately leading to more accurate predictions.

ii) Guiding Decision-Making
Uncertainty quantification informs decision-making processes by providing a framework to evaluate risks associated with different choices. For example, in healthcare, understanding the uncertainty in a model predicting patient outcomes can assist clinicians in evaluating treatment options, leading to more personalized and effective care.

iii) Facilitating Risk Assessment
In engineering and safety-critical applications, UQ is essential for risk assessment. By quantifying uncertainty, engineers can evaluate the potential failures of systems and the consequences of those failures. This assessment is crucial in designing safer systems and infrastructure. For instance, in structural engineering, UQ helps determine the safety margins of buildings and bridges under uncertain loading conditions, ensuring that designs meet safety standards.

iv) Improving Model Interpretability
Uncertainty quantification enhances model interpretability by providing insights into the confidence of predictions. When models output uncertainty estimates alongside predictions, users can gauge how much trust to place in the results. This is particularly important in high stakes domains such as autonomous driving and medical diagnosis, where understanding the uncertainty associated with predictions can be critical for user trust and acceptance.

v) Supporting Model Improvement
UQ can guide the improvement of models by identifying areas where predictions are most uncertain. By analyzing uncertainty sources, researchers can focus on refining specific aspects of the model or gathering additional data to reduce uncertainty. This iterative process not only enhances model performance but also drives innovation by revealing gaps in knowledge that can be addressed through further research.

vi) Advancing Scientific Discovery
In scientific research, UQ is vital for advancing knowledge and understanding complex phenomena. By quantifying uncertainties in models and simulations, researchers can better evaluate the implications of their findings and draw more robust conclusions. For example, in climate modeling, UQ helps scientists understand the range of possible future scenarios and the associated uncertainties, informing policy decisions and societal responses to climate change.

vii) Enhancing Machine Learning Performance
In machine learning, UQ is particularly important for tasks such as classification, regression, and reinforcement learning. Quantifying uncertainty helps in developing models that can generalize better to unseen data, improving overall performance. Techniques such as Bayesian inference and Monte Carlo dropout provide a systematic way to incorporate uncertainty into neural networks, allowing for better decision-making and risk assessment in applications ranging from image recognition to natural language processing.

1.2 Importance of AI in Materials Science

Artificial Intelligence (AI) has become an indispensable tool in modern material science. Its ability to process and analyze large datasets quickly and accurately has transformed how scientists and engineers approach material design, analysis, and optimization. AI techniques, including machine learning and deep learning, enable the discovery of new materials, prediction of material properties, and optimization of manufacturing processes. For example, AI can help predict the properties of a new alloy based on its composition and microstructure, accelerating the development of new materials.

In the context of material science, AI is particularly valuable for solving complex inverse design problems. Inverse problems involve determining the underlying causes or parameters from observed effects or data. This approach enables scientists to infer material structures, compositions, and mechanisms that are not directly observable. Solving inverse problems in materials science often requires sophisticated mathematical and computational techniques, including optimization algorithms and machine learning methods. These techniques help translate experimental data into meaningful insights, driving innovations in material design, development, and application. By addressing inverse problems, researchers can develop new materials with tailored properties, optimize manufacturing processes, and advance the understanding of complex material systems. Traditional methods often struggle with the complexity and computational

intensity required for such tasks. However, AI, with its ability to learn from data and make predictions, provides a more efficient and effective approach.

1.2.1 Computational Solutions to Material Science Problems

Computational methods have long been a cornerstone of material science, enabling researchers to simulate and predict material behavior under various conditions. Techniques like finite element analysis (FEA), molecular dynamics simulations, and density functional theory (DFT) are widely used to model the properties and behavior of materials at different scales. However, these methods can be computationally expensive and may not always provide unique or stable solutions, especially in the case of inverse design problems.

The integration of AI into these computational methods offers a powerful synergy. AI can enhance traditional simulations by providing more accurate predictions, identifying patterns in complex data, and exploring vast design spaces more efficiently. For instance, AI can assist in the optimization of material properties by exploring combinations of microstructural features that lead to desired mechanical properties. In inverse design problems, AI can generate multiple potential solutions that fit the observed data, helping researchers identify the most probable configurations.

1.3 Advancements in Deep Learning and GANs

In recent years, deep learning has revolutionized several fields, including computer vision, natural language processing, and biomedical imaging. The ability of deep learning models to learn from vast amounts of data and make accurate predictions has spurred interest in applying these techniques to materials science. GANs, in particular, have demonstrated significant potential in tasks such as image synthesis, style transfer, and more recently, in the characterization and analysis of mechanical stress fields.

GANs operate by training two networks: the generator, which tries to produce realistic data, and the discriminator, which attempts to distinguish between real and generated data. This training process, known as adversarial training, is powerful because it pushes the generator to improve its output continually. A simple GAN architecture is shown in Figure 1.1. There are various types of GANs designed to address specific challenges:

1. Vanilla GANs: The original form of GANs, where the generator and discriminator are simple neural networks.

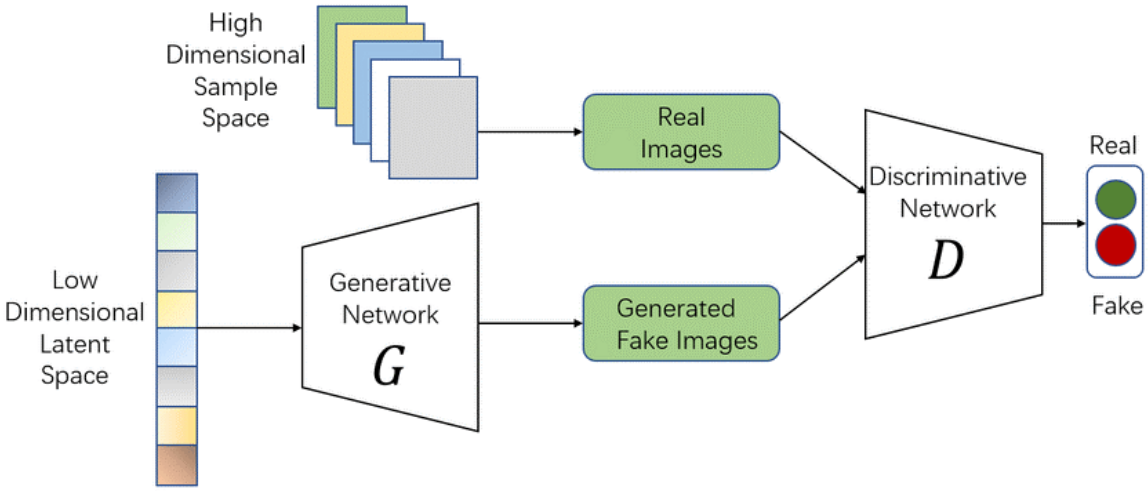

Figure 1.1: The architecture of Vanilla GAN [1]

2. Conditional GANs (cGANs): Extend GANs by conditioning the generation process on additional information, such as labels or specific features, allowing for more controlled data generation.

3. Deep Convolutional GANs (DCGANs): Incorporate convolutional layers, making them more effective for generating high-quality images.

4. Wasserstein GANs (WGANs): Modify the loss function to provide better gradients for training, addressing issues like mode collapse.

5. Cycle GANs: Enable image-to-image translation without paired training examples, useful for tasks like converting images from one domain to another (e.g., photos to paintings).

In addition to GANs, there lies a subclass of GANs which is called **Diverse GANs**. Diverse GANs are a class of Generative Adversarial Networks designed to produce a wide range of outputs that reflect the variability and diversity present in real-world data. Unlike standard GANs, which may generate high-quality but often similar outputs, diverse GANs focus on capturing the underlying distribution of the data, thereby enabling the generation of multiple plausible solutions for a given input. This diversity is particularly valuable in tasks such as image synthesis, style transfer, and data augmentation, where capturing the full spectrum of possible variations is crucial. Diverse GANs achieve this by incorporating techniques such as latent space manipulation, mode seeking regularization, and architectural modifications that encourage the net-

work to explore and represent multiple modes of the data distribution. By promoting variability in the generated outputs, diverse GANs enhance the model's ability to generalize and adapt to different contexts, making them powerful tools in applications that require a rich and varied output space.

1.3.1 Image Inpainting and the Role of GANs

Image inpainting is a technique used to restore or reconstruct missing or damaged parts of an image. This process is crucial in various applications, including photo editing, restoration of old or damaged photographs, removal of unwanted objects from images, and even medical imaging where certain parts of an image might be obscured. The goal of image inpainting is to fill in the missing regions in a way that is seamless and visually coherent with the surrounding parts of the image, making it appear as if the inpainted area was never missing or altered.

Traditional image inpainting methods rely on techniques such as interpolation, diffusion, and patch-based methods. These techniques use the surrounding pixel information to estimate the missing areas. However, these methods can struggle with maintaining high-level semantic coherence and often fail to generate realistic textures, especially when the missing areas are large or complex. This is where deep learning, particularly Generative Adversarial Networks (GANs), has made significant advancements.

GANs have revolutionized the field of image inpainting due to their ability to generate highly realistic images. The use of GANs in image inpainting offers several advantages:

1. High-Quality Outputs: GANs can generate high-resolution and photorealistic images. The adversarial training setup forces the generator to produce inpainted areas that are not only plausible but also blend seamlessly with the surrounding image context.

2. Learning from Data: Unlike traditional methods that rely on handcrafted features or simple pixel interpolation, GANs learn directly from data. This allows them to capture complex patterns and structures in the data, leading to more accurate and contextually appropriate inpainting.

3. Flexibility: GANs can be trained on diverse datasets, making them adaptable to a wide range of inpainting tasks. Whether it's filling in small details or reconstructing large missing regions, GANs can handle various levels of image complexity.

4. Semantic Understanding: One of the critical strengths of GANs in image in-

Figure 1.2: Diverse image inpainting results [2]

painting is their ability to understand the semantic context of an image. This means that they can generate content that is not only visually coherent but also contextually appropriate. For example, if a part of a face is missing, a GAN can infer the likely appearance of the missing features based on the rest of the face as shown in Figure 1.2.

In practice, various types of GAN architectures can be employed for image inpainting, including Deep Convolutional GANs (DCGANs) for better handling of image data, and Conditional GANs (cGANs) that allow the generation process to be guided by additional information, such as class labels or specific conditions. More advanced techniques, like the use of contextual attention mechanisms, have been developed to improve the quality of inpainting by focusing on relevant areas of the image and better blending the generated content.

Overall, GANs have set a new standard in image inpainting by providing a powerful tool for creating realistic and high-quality images, even in the presence of significant missing or corrupted areas. Their ability to learn from large datasets and produce contextually and visually coherent inpainting has made them indispensable in fields ranging from digital media and entertainment to scientific research and medical imaging.

1.4 Background and Motivation

The advancement of materials engineering plays a critical role in modern technological developments, where understanding and predicting the behavior of materials under various conditions is of utmost importance. Materials engineering focuses on studying how materials respond to external forces, their mechanical properties, and how they fail under stress. Accurate prediction of these behaviors can enable the design of safer and more efficient materials, impacting industries such as aerospace, automotive, and construction. However, predicting material behavior is complex due to various internal factors, such as microstructural variations, defects, and boundary conditions, which introduce significant uncertainty into the models.

1.4.1 Importance of Predicting Material Behavior

Predicting material behavior under different conditions is crucial for ensuring the reliability and safety of structures. Engineers use stress analysis to predict how materials will deform, fracture, or fail when subjected to forces like tension, compression, or shear. These predictions are essential in designing materials for critical applications, where failure could have catastrophic consequences. Traditionally, techniques such as finite element analysis (FEA) have been used to predict material behavior, but they rely on complete and precise data, which is often not available. This is where data-driven approaches, especially machine learning models, step in to fill the gap.

1.4.2 Inpainting of Material Stress Fields Using GANs

Generative Adversarial Networks (GANs) have emerged as powerful tools in materials engineering, particularly for inpainting tasks where missing or incomplete data needs to be reconstructed. Stress field inpainting, which involves predicting the missing stress distributions in a material, is a critical challenge in non-destructive testing and failure analysis. By learning from existing data, GANs can predict the missing parts of the stress fields, offering engineers an efficient method to reconstruct internal material behaviors based on partial observations. This method has shown promise in capturing both local and global stress patterns, making it a valuable tool for material scientists.

However, despite the power of GANs in inpainting tasks, the predictions are not always accurate or reliable, especially in regions with sparse data. In critical applications

where high precision is required, any uncertainty in the reconstructed stress fields could lead to incorrect predictions of material failure or deformation, thereby affecting safety and performance.

1.4.3 The Role of Uncertainty in GAN Predictions

Given the limitations of GANs and the high stakes involved in predicting material behavior, it becomes necessary to quantify the uncertainty in the GAN predictions. Uncertainty quantification (UQ) allows engineers to gauge the confidence of the model's predictions, identifying areas where the model is less certain. This is crucial because it helps guide engineers in making informed decisions based on the model's output. Monte Carlo dropout, a popular method for UQ, can be applied to GANs to generate multiple predictions for the same input, providing a measure of variability and uncertainty in the reconstructed stress fields.

Integrating uncertainty into GAN predictions not only improves the reliability of the inpainting process but also highlights the regions where predictions may require further scrutiny or additional data. This is especially important in materials engineering, where high-confidence predictions are required to ensure the safety and performance of engineered materials. By understanding the uncertainty, engineers can make better-informed decisions and optimize the design of materials for various applications.

In conclusion, the need for accurate prediction of material behavior and the challenges posed by incomplete data highlight the importance of uncertainty quantification in materials engineering. GANs, coupled with uncertainty estimation methods like Monte Carlo dropout, offer a promising pathway to improve the reliability of predictions in stress field inpainting, thus enhancing our ability to design and evaluate materials more effectively.

1.4.4 Research Objective

The primary objective of this research is to quantify the uncertainty in image inpainting of mechanical stress fields. Image inpainting in this context refers to the process of filling in missing parts of an image of a material's stress fields, based on the surrounding context. This study utilized a technique known as Monte Carlo dropout, a regularization method commonly used during the training phase of neural networks, into the inference phase of GANs. By applying varying dropout rates in different layers of the generator, the research aims to assess the impact on the standard deviation and error of the generated images. These metrics provide a measure of the confidence in

the GANs' predictions, thereby enhancing the models' robustness and reliability.

The research builds upon the methodology presented in Yang and Buehler's work [12], which focused on transferrable deep learning approaches for recovering missing physical field information. In this study, the introduction of dropout during inference is a novel modification aimed at exploring the uncertainty associated with the predictions made by GANs. The performance of these models is evaluated using standard deviation and error metrics, providing valuable insights into the degree of uncertainty and the reliability of the inpainting process.

1.5 Structure of the Thesis

This thesis is organized to provide a comprehensive exploration of the research topic, covering all aspects from theoretical background to practical applications. The chapters are structured as follows:

- **Chapter 1: Introduction** - This chapter provides an overview of the background and motivation for the research, outlines the research objectives, and presents the overall structure of the thesis.
- **Chapter 2: Literature Review** - This chapter reviews the relevant literature, focusing on the advancements in deep learning, particularly GANs, and their applications in material science. It also discusses previous work on uncertainty quantification in deep learning models.
- **Chapter 3: Methodology** - This chapter details the methodology used in the research, including the modifications made to the existing code and the experimental setup for analyzing the impact of dropout during inference.
- **Chapter 4: Results** - This chapter presents the results of the experiments, analyzing the impact of varying dropout rates on the inpainting performance. It includes detailed discussions on the standard deviation and error metrics and their implications for uncertainty quantification.
- **Chapter 5: Discussion and Conclusion** - This final chapter summarizes the key findings of the research, discusses their implications, and provides recommendations for future work. It also highlights the contributions of the research to the field of material science and deep learning.

Chapter 2

Literature Review

The integration of uncertainty quantification into machine learning, particularly in materials engineering, has gained significant traction in recent years. Traditional deterministic models, while effective in certain controlled environments, often fail to capture the variability and unpredictability inherent in real-world materials. In response, researchers have explored various probabilistic approaches to better account for uncertainty. Bayesian Deep Learning, Monte Carlo dropout, and Generative Adversarial Networks (GANs) have emerged as key methodologies, offering frameworks to not only predict material behavior but also access the confidence of these predictions. This shift has been instrumental in addressing challenges in fields like stress field reconstructions, where incomplete or noisy data is prevalent. By incorporating uncertainty estimation techniques, these models provide more reliable predictions, enabling engineers to design and analyze materials with greater confidence and precision.

2.1 Advancements in GANs for Inpainting Tasks

The literature on GANs and their applications in image inpainting and material science is vast and continually growing. The development of GANs has opened new possibilities for generating realistic data, solving complex inverse design problems, and enhancing data-driven approaches in various fields. The incorporation of uncertainty quantification techniques further strengthens the utility of GANs in scientific research, providing more reliable and interpretable results. This literature review provides a foundation for understanding GANs and uncertainty quantification research and highlights key areas of their application in materials science.

A research [13] explores the use of conditional adversarial networks (cGANs) for image-to-image translation tasks. By learning both the mapping from input to output images

and the loss function, cGANs eliminate the need for hand-engineered loss functions, thereby providing a unified framework for image synthesis. The core idea is to use these networks not only to learn the mapping from input to output images but also to simultaneously learn the loss function necessary to train this mapping. This approach enables the application of a common framework to diverse problems traditionally requiring different loss formulations. The paper demonstrates the effectiveness of cGANs in various applications such as synthesizing photos from label maps, reconstructing objects from edge maps, and colorizing images. The authors introduce a simple yet effective cGAN architecture, utilizing a “”U-Net” generator and a “PatchGAN” discriminator, which penalizes structure at the scale of image patches. This architecture has proven effective in capturing high-frequency details and producing sharp, realistic outputs. The paper emphasizes that cGANs can replace hand-engineered loss functions, providing a unified and automatic approach to a wide range of image translation tasks. Through extensive experiments, the authors show that cGANs can handle tasks like semantic labeling, map-to-photo translation, and sketch-to-photo synthesis, producing visually convincing results. The work highlights the community's positive reception and the versatility of the cGAN framework in creative and technical applications, as evidenced by various projects shared under the #pix2pix hashtag.

The principle of automating complex processes with neural networks is further expanded in [14] where the authors introduce a gated convolutional system that dynamically selects features for improved inpainting results. This paper presents a novel approach to image inpainting, focusing on completing images with free-form masks and user guidance. The authors introduce a system based on gated convolutions, which are designed to handle the limitations of traditional convolutions that assume all input pixels are valid. Unlike partial convolutions, which only use valid pixels, gated convolutions provide a dynamic, learnable feature selection mechanism for each channel and spatial location, improving inpainting quality by addressing issues like color discrepancy and edge artifacts. To address the challenges posed by free-form masks, the authors propose the SN-PatchGAN, a patch-based GAN loss that applies spectral normalization to the discriminator, making it stable and efficient in training. This technique ensures high-quality inpainting results by focusing on different patches and semantic details in the image. The system is capable of handling various tasks such as removing objects, modifying image layouts, and clearing watermarks, offering higher quality and flexibility compared to previous methods.

Traditional CNN-based inpainting methods often produce semantically plausible but blurry results due to their reliance on fully connected layers. Shift-Net [15] introduces a shift connection layer that borrows from exemplar-based methods to rearrange known image features into missing regions. The main innovation of Shift-Net is the inclusion of a shift-connection layer within a U-Net architecture. This layer enables the network

to leverage features from the known regions of an image to estimate the missing parts, thereby combining the strengths of both exemplar-based and CNN-based inpainting techniques. Shift-Net's architecture allows it to handle arbitrary shapes of missing regions and to generate results with fine details and sharp structures. The model is trained using a combination of guidance, reconstruction, and adversarial losses, optimizing it to produce realistic and coherent inpainting outputs. The Shift Net model significantly outperforms previous methods, such as Content-Aware Fill, context encoders, and multi-scale neural patch synthesis (MNPS), both in quality and efficiency. It is capable of inpainting tasks like object removal and region completion, demonstrating effectiveness on datasets such as Paris StreetView and Places365.

To further enhance the effective ness of inpainting tasks, a contextual attention layer in inpainting was introduced. The paper [16] presents an advanced image inpainting approach that addresses the challenge of filling large missing regions with contextually appropriate and visually realistic content. The proposed method integrates a novel contextual attention layer, which explicitly borrows relevant features from surrounding areas of the image, enhancing the coherence and reducing artifacts. The network utilizes a coarse-to-fine architecture, where an initial rough estimate of the missing content is refined in a second stage to produce a high-quality final result. Additionally, the use of global and local Wasserstein GANs (WGANs) ensures that the inpainted images are consistent both globally and locally. This method significantly improves efficiency, reducing training time and eliminating the need for post-processing steps like image blending. Tested on datasets such as CelebA, CelebA-HQ, DTD, ImageNet, and Places2, the approach outperformed existing methods by producing fewer distortions and more seamlessly integrated inpainted regions. The paper concludes that the new inpainting framework, with its contextual attention mechanism and WGANs, offers a robust and efficient solution for generating detailed and contextually accurate images, with broad applications in image editing and computational photography.

The paper [17] explores two contemporary deep learning approaches for semantic inpainting. The first approach, more commonly used, involves training an offline deep regression network on masked pixels with adversarial training refinement, allowing for single-pass inference. The second approach, less explored, involves training a generative model to map a latent prior distribution to a natural image manifold and using iterative optimization at inference to find the best-matching prior. The paper focuses on the latter, proposing a method to predict a matching prior directly from a masked image, thus converting the iterative process into a single feed forward inference, significantly speeding up the process by 800 times. The proposed model also includes structural priors to better preserve pose and size, enhancing image quality and consistency. The method scales to high resolutions and improves performance on datasets such as SVHN, Stanford Cars, CelebA, CelebA-HQ, and ImageNet. The paper also

addresses video inpainting, introducing a recurrent network for grouped latent prior learning, which maintains temporal consistency. The approach achieves state-of-the-art results, outperforming existing benchmarks. The study emphasizes inpainting as a search for the best latent prior rather than image refinement, offering a new perspective in the field.

The idea of diversity in image inpainting was further extended by PD-GAN [18] by generating multiple inpainting possibilities from a single input image. PD-GAN, or Probabilistic Diverse Generative Adversarial Network, is an advanced image inpainting method designed to generate multiple diverse and visually realistic completions for images with missing regions. This network builds upon a vanilla GAN framework, leveraging random noise inputs to produce varied outputs. A key feature of PD-GAN is the introduction of Spatially Probabilistic Diversity Normalization (SPDNorm), which modulates deep features during the image generation process. SPDNorm includes both hard and soft variants: hard SPDNorm ensures that pixels near the hole boundaries are more deterministic, thereby maintaining consistency with the surrounding context, while soft SPDNorm allows greater flexibility and diversity, particularly in the center of the hole. Additionally, PD-GAN employs a perceptual diversity loss, which enhances the diversity of the generated images by ensuring that different outputs, generated from the same initial conditions but different noise vectors, are perceptually distinct. Tested on datasets like CelebA-HQ, Places2, and Paris Street View, PD-GAN demonstrated superior performance compared to existing methods, producing multiple inpainting results that are not only diverse but also visually coherent. Ablation studies confirmed the critical role of SPDNorm and perceptual diversity loss in achieving these results, establishing PD-GAN as a significant advancement in image inpainting technology.

Unlike general inpainting methods that primarily focus on image resolution, DE-GAN [19] embeds domain-specific knowledge of faces, including face masks, parts, and landmarks, to guide the inpainting process. This approach helps produce more natural and harmonious facial features. DE-GAN uses a hierarchical variational auto-encoder (HVAE) to embed these three types of face domain knowledge into a latent variable space. The network is trained with two adversarial discriminators—a global discriminator and a patch discriminator—to ensure the generated images are both globally and locally consistent with real faces. The method is evaluated on two public face datasets, CelebA and CelebA-HQ, and demonstrates superior performance in generating high-quality, realistic inpainted faces compared to existing methods. DE-GAN is particularly effective in handling under-pose variations, a challenging scenario where the face is not in a standard frontal view. The results show that DE-GAN produces more consistent and natural facial structures and appearances than previous methods, making it a state-of-the-art solution in face inpainting. Additionally, the authors discuss potential applications of the domain knowledge extracted by DE-GAN, such as

face attribute editing and makeup transfer. However, the method has limitations, such as difficulties in handling images with multiple faces, which the authors plan to address in future work.

A paper [20] explores the use of Conditional Generative Adversarial Networks (cGAN) for image-to-image translation tasks, which involve transforming an image from one domain to another. The authors modified the cGAN's loss function by combining adversarial loss with various non-adversarial loss functions, such as L1, content, and structural losses. They evaluated the impact of different adversarial losses, including Wasserstein GAN (WGAN), WGAN with Gradient Penalty (WGAN-GP), and least Squared GAN (lsGAN), on the quality of the generated images. The experiments, conducted on the Facades dataset, demonstrated that combining WGAN's adversarial loss with L1 and content losses produced the best results, yielding high-quality images with both low and high-frequency details. The study also extended to applications in image inpainting and lesion segmentation, highlighting the practical utility of the proposed approach. The findings suggest that combining appropriate loss functions can significantly enhance the performance of cGANs in various image processing tasks.

2.2 Various Methods for Inverse Problems

Apart from conventional GANs based inpainting, new techniques have also emerged which are used for inverse problems and offer a wider range of possibilities.

Unlike traditional methods that generate images sequentially, MaskGIT [21] which uses a bidirectional transformer decoder, predicts all image tokens simultaneously, refining them iteratively. This approach significantly speeds up the image generation process and produces high-quality results. MaskGIT uses a technique called Masked Visual Token Modeling (MVTM) during training, where randomly masked tokens are predicted using context from all directions. The model achieves state-of-the-art results on the ImageNet dataset, outperforming previous transformer-based models in both quality and speed. Additionally, MaskGIT is versatile, capable of handling various image editing tasks like inpainting and extrapolation. The paper also discusses the importance of mask scheduling in the iterative decoding process and finds that a cosine schedule works best. Overall, MaskGIT represents a significant advancement in image synthesis and editing, providing a promising alternative to GANs.

In addition to GANs, a new class of generative models known as diffusion models have gained a significant appreciation for inpainting tasks. The paper [22] presents a new approach called Come-Closer-Diffuse-Faster (CCDF) to enhance the efficiency of

diffusion models, which have gained prominence for their outstanding performance in generative tasks but are limited by slow sampling speeds. Traditionally, these models generate images starting from Gaussian noise, a process that typically requires thousands of iterative steps. CCDF addresses this limitation by initiating the image generation process with a forward diffusion using better initializations, drastically cutting down the number of reverse diffusion steps needed. This innovative technique is grounded in the contraction property of stochastic difference equations, allowing for faster error reduction during the reverse diffusion phase. The paper delves into the background and methodology of diffusion models, highlighting their widespread applications in tasks such as super-resolution, image inpainting, and MRI reconstruction. Despite their versatility and strong performance, these models are hindered by the extensive time required for sampling, a drawback that CCDF aims to overcome. By starting with a more refined initial state, rather than pure Gaussian noise, CCDF leverages the inherent properties of these models to accelerate the convergence process significantly. Experiments conducted across various tasks demonstrate the efficacy of CCDF in achieving state-of-the-art results with fewer sampling steps. The study includes applications in super-resolution, where CCDF shows marked improvements in image quality and fidelity with fewer iterations. Similarly, in tasks like inpainting and MRI reconstruction, CCDF not only accelerates the process but also maintains, or even enhances, the quality of the output images compared to traditional methods.

The paper [23] highlights the strengths of diffusion models, particularly their ability to generate high-quality samples across different inverse problems without specific training. The authors demonstrate that by integrating the Manifold Constrained Gradient (MCG) with traditional diffusion steps, the process becomes more stable and accurate, offering better results in tasks like image inpainting, colorization, and sparse-view computed tomography. This paper discusses the use of diffusion models in solving various inverse problems in an unsupervised manner. Traditional methods, which apply a reverse diffusion step followed by a projection-based measurement consistency step, often result in suboptimal outputs due to deviation from the data manifold. The authors introduce a novel correction term called the MCG to address this issue. MCG helps to maintain the sample path closer to the data manifold, significantly enhancing the performance of these models. The authors provide a theoretical foundation for their method, showing that MCG can guide the diffusion process to stay on the data manifold, thus preventing errors. Extensive experiments show that the proposed method outperforms existing state-of-the-art methods both in terms of quality and efficiency. However, the paper notes some limitations, such as the inherent stochasticity and potential for slower sampling speeds, as well as concerns about possible misuse in creating biased or malicious content.

A novel approach to solving physics-based Bayesian inverse problems, which are preva-

lent in various scientific and engineering disciplines, was introduced in [24]. It presents a novel approach to solving physics-based Bayesian inverse problems, which are prevalent in various scientific and engineering disciplines. These problems are often ill-posed, making it challenging to obtain unique and stable solutions. The authors introduce a method leveraging deep generative models, specifically Generative Adversarial Networks (GANs), to enhance the efficiency and accuracy of Bayesian inference. The key innovation lies in using GANs to learn an approximate distribution from prior data, which serves as a prior in the Bayesian framework. This significantly reduces the dimensionality of the problem, mapping it to the low-dimensional latent space of the GAN, thereby facilitating efficient exploration of the posterior distribution. This approach is particularly effective in handling complex prior distributions, providing reliable uncertainty estimates, and incorporating prior information from historical data or previously acquired solutions. The method is applied to various inverse problems, including heat conduction, Radon transform in CT imaging, and elasticity imaging for inferring mechanical properties of tissues. The results demonstrate that the GAN-based method offers accurate reconstructions and uncertainty quantification, outperforming traditional methods, particularly in high-dimensional cases. The paper concludes that employing GANs as priors in Bayesian inference offers a promising solution for complex inverse problems, especially when rich prior information is available.

Traditional inpainting methods often struggle with various types of masks and can produce results that are not semantically meaningful. RePaint [25] addresses these issues by leveraging a pretrained, unconditional Denoising Diffusion Probabilistic Model (DDPM) as a generative prior, altering the reverse diffusion process to condition the image generation on available image data. This method allows RePaint to handle extreme and diverse mask types without specific mask training, producing high-quality and diverse outputs. The authors demonstrate that RePaint outperforms state-of-the-art Autoregressive and GAN-based methods in several inpainting scenarios, particularly in handling challenging masks like super-resolution and thick masks. The approach maintains high performance across different datasets, including CelebA-HQ and ImageNet, and shows robustness in generating semantically coherent images even under extreme conditions. Despite its strengths, RePaint's per-image optimization process is slower compared to other methods, posing challenges for real-time applications. Additionally, the method's reliance on a pretrained DDPM could inherit biases from the training data. Nonetheless, RePaint offers significant advancements in free-form inpainting and sets a new benchmark for future research in this field.

2.3 Inpainting in Materials Engineering

Extensive research has been conducted in materials engineering, in which inpainting has been applied. Inpainting has been effective in solving inverse problems in the field of materials engineering and other areas of research.

A research paper [26] delves into using image inpainting algorithms to remove irrelevant elements like letters, scale bars, and arrows, from microscopic images in material science documents. The authors employed four image inpainting methods: shift-net, global and local, contextual attention, and gated convolution. They evaluated these methods using the Structural Similarity Index Measure (SSIM) and ℓ_1/ℓ_2 errors to assess image quality. The study found that gated convolution performed the best in restoring the images. The authors also discussed the challenges of using statistical and threshold-based masks for preprocessing the images. The paper concludes that while the methods showed promising results for microscopic images, future work is needed to adapt these techniques for other types of images in material science documents.

Building on the previous study's foundation of image inpainting, a research [27] introduces a fully automated system designed for microscopic image restoration, emphasizing its application in material science. The paper presents a deep learning-based method for automatically detecting and inpainting damaged regions in microscopic images, specifically targeting material science applications. Microscopic images are crucial for analyzing material microstructure, but they often contain defects such as scratches or stains, which can hinder accurate analysis. The authors developed a fully automatic method that utilizes three main components: a damaged region segmentation network, an image inpainting network, and data augmentation using generative adversarial networks (GANs). The method was tested on Al–La alloy microscopic images and demonstrated superior performance in accurately recovering damaged regions compared to existing software. The approach uses a U-Net architecture for segmentation, followed by GAN-based inpainting, which fills in missing regions based on surrounding contextual information. Data augmentation was used to enhance the training dataset, improving the model's robustness. The study showed that this method could significantly improve the accuracy and efficiency of image processing in material science, and the approach could be extended to other fields requiring image restoration.

The previous methods of restoring damaged images naturally evolve into a study [28] focusing on high-fidelity microstructure reconstruction. The paper addresses three contexts of high fidelity microstructure reconstruction: interpolation (reconstructing missing or damaged regions), extrapolation (extending images for a larger view), and assembly (combining disjoint images into a complete set). The authors propose using image inpainting techniques to overcome the data scarcity problem and augment

datasets for machine learning applications. This approach helps generate statistically equivalent synthetic microstructures, crucial for simulations and uncertainty quantification in materials science. The proposed technique aims to advance the state of the art in microstructure reconstruction, complementing existing machine learning methods by providing high-fidelity data.

In painting techniques, originally developed for image restoration and artistic editing, have found critical applications in biomedical engineering, particularly in the analysis of complex medical images. A paper [29] explores a novel method for reconstructing a complete vertebral body using 3D printing, aimed at aiding Total En Bloc Spondylectomy (TES) for spinal tumor treatment. The process begins with the segmentation of spinal CT images using a neural network to create binary images. A Generative Adversarial Network (GAN) is then trained on these images to generate normal spinal CTs. For lesion detection, a deep convolutional neural network identifies the tumor areas in CT scans. The GAN model is subsequently used for inpainting these regions to restore a normal vertebral structure. This reconstructed model serves as a base for creating 3D printed implants. The study demonstrates that the proposed method effectively restores and reconstructs spinal images, potentially aiding in precise surgical planning and reducing the time and cost associated with creating accurate 3D models for medical use. However, the study acknowledges the limitations of relying solely on medical images for diagnosis and the need for large, annotated datasets for further research and improvement in CAD technologies.

Following the theme of reconstructing incomplete data, the base paper [12] presents an innovative deep learning framework to address inverse problems in material science. These inverse problems often involve deducing internal structures or properties of materials from indirect measurements, such as those taken at boundaries or interfaces. This is a common challenge in areas like failure analysis, nondestructive testing, and mechanical metamaterials design. The authors propose a two-pronged deep learning approach for both 2D and 3D scenarios. For 2D cases, they employ a conditional generative adversarial network (cGAN) to reconstruct incomplete mechanical field maps, such as stress and strain fields, from partially masked data. Following this, a convolutional neural network (CNN) translates these reconstructed fields into detailed descriptions of the composite microstructures. The 2D model showcases its robustness by handling a wide range of scenarios, including mixed components, different grid sizes, and complex microstructures such as those governed by the Cahn-Hilliard equation. The model also proves effective in handling non-linear material behaviors like plasticity. In the 3D context, the authors utilize a Transformer-based model to predict complete 3D mechanical fields from partial input data. This model demonstrates the capability to reconstruct internal material structures accurately even when only surface data is available. The use of a stack of 2D field frames to represent 3D data allows for the

prediction of entire fields from just a few initial frames. This method is particularly useful in practical applications like nondestructive testing, where accessing internal structures directly might not be feasible. A significant contribution of this work is its handling of ill-posed inverse problems, where multiple internal structures can produce similar boundary field data. The authors address this by generating multiple possible solutions, providing a probabilistic representation of potential internal structures. This is achieved by training multiple deep learning models on different subsets of the data, thus enabling the exploration of a range of feasible solutions. The frameworks developed in this study not only offer efficient and accurate solutions to inverse problems but also pave the way for potential applications in various fields beyond materials science, such as biomedical imaging, fluid dynamics, and more. The approaches do not rely on prior knowledge of governing physical laws, making them broadly applicable across different types of physical fields. The authors suggest that these methods could also be combined with optimization algorithms for inverse design problems, enabling the iterative search for optimal material configurations based on desired mechanical properties.

2.4 Uncertainty Quantification

Uncertainty quantification has become an integral part of modern machine learning models, especially in high-stakes fields like materials science and biomedical engineering. In these domains, understanding the confidence and reliability of predictions is crucial for decision making and risk assessment. Various techniques, such as Bayesian deep learning, Monte Carlo dropout, and Gaussian processes, have been developed to model and estimate uncertainties, enhancing the interpretability and trustworthiness of predictive models.

The journey begins with the foundational concept of using dropout in neural networks (NNs) to capture uncertainty. Traditionally employed as a regularization method to avoid overfitting, dropout can be repurposed as a Bayesian approximation technique [7]. This paper explores the use of dropout NNs to extract information typically discarded, addressing the challenge of representing uncertainty in deep learning. The study covers various network architectures and nonlinearities, demonstrating improvements in predictive log-likelihood and root mean square error (RMSE) compared to other methods. The framework is also applied in deep reinforcement learning, where uncertainty helps agents decide between exploiting known rewards and exploring new possibilities. Theoretical developments link dropout to Bayesian models, specifically deep Gaussian processes, by approximating the probabilistic nature of these processes. The paper provides a detailed explanation of how dropout, typically used to prevent overfitting, can be seen as a form of Bayesian model averaging. This approach improves

the handling of model uncertainty in NNs, crucial for applications requiring reliable uncertainty estimates, such as in critical infrastructure or reinforcement learning scenarios. Experiments demonstrate the practical benefits of this approach in regression and classification tasks, showing that dropout's uncertainty estimates can enhance performance in predicting outcomes and understanding model confidence. The paper concludes by discussing future research directions, including the exploration of other regularization techniques and their potential to further improve uncertainty modeling in deep learning.

A novel application of dropout within GANs was introduced in [30] to produce multiple diverse outputs from a single input. Typically, GANs use latent space exploration for this purpose, but Dropout offers an alternative that preserves input constraints, useful in tasks like A-to-B translation, style transfer, image inpainting, and text-to-image synthesis. Dropout is commonly used in GANs to prevent overfitting by deactivating units with a certain probability, which reduces co-adaptations and improves generalization. In this study, Dropout is applied not only during training but also during the generation phase to induce noise and create a variety of outputs. The paper examines different configurations, including applying Dropout to all hidden layers or just the first one, and varying the Dropout rates. Experiments using the MNIST dataset and a DCGAN architecture demonstrate that higher Dropout rates result in more diverse outputs but can also lead to a loss of detail. The study also shows that omitting scaling in the generation phase leads to noisy and incoherent images. When Dropout is applied only to the first hidden layers, subsequent layers can sometimes correct errors, improving image quality. The paper concludes that Dropout is a viable method for generating diverse outputs in GANs, especially when other methods like latent space exploration are not suitable. However, careful tuning of Dropout rates is necessary to balance output diversity with image coherence. Further research is suggested to explore the benefits of applying Dropout only to specific layers.

U2AFN [31] introduces a feedback mechanism and uncertainty maps that focus on high uncertainty areas, refining the inpainting quality iteratively. This paper presents the Uncertainty Aware Adaptive Feedback Network (U2AFN) for image inpainting, which addresses the challenge of filling large image holes that are difficult for conventional methods to handle. U2AFN utilizes an adaptive feedback mechanism and uncertainty estimation to iteratively refine the inpainting results. Key components include:

- **Adaptive Integration Feedback Block:** This block integrates low-level and high-level information, enhancing the model's ability to reconstruct missing regions progressively.
- **Uncertainty Map Estimation:** U2AFN predicts both inpainting results and

uncertainty maps, which indicate areas with high uncertainty. The model focuses on these high-uncertainty areas in subsequent iterations, gradually improving the inpainting quality.

- **Loss Functions:** The method employs several loss functions, including perceptual loss, style loss, total variation loss, and adversarial loss, to ensure the generated images are visually realistic and semantically coherent.

The paper demonstrates U2AFN's superior performance on multiple datasets, including CelebA-HQ, Places2, and Paris StreetView, compared to other state-of-the-art inpainting methods. U2AFN shows improvements in metrics like PSNR, SSIM, and mean L1 error, indicating better preservation of structural and semantic information in the inpainted images.

To quantify uncertainty in high-dimensional spaces the use of GANs as priors in Bayesian inference was explored in [32]. The authors demonstrate this approach in two main problem classes: inverse problems, where the forward mapping is known (e.g., image denoising and physics-driven inference), and data-driven problems, where the mapping must be inferred from data (e.g., image classification and inpainting). By leveraging GANs, the method efficiently reduces the dimensionality of the inference problem, allowing for more effective sampling from the posterior distribution. This approach proves useful in tasks like out-of-distribution detection and active learning, showcasing its potential across a range of applications.

Stepping into the materials science domain, a paper [33] dives into the challenge of reconstructing metallic microstructures while accounting for both aleatoric (data-related) and epistemic (model-related) uncertainties. This paper explores microstructure reconstruction techniques and the uncertainties they introduce, particularly focusing on metallic alloys. The study uses Principal Eigenvalue Moments (PEM) derived from Hu moments to quantify these uncertainties in 2D and 3D microstructures. The paper discusses the impact of these uncertainties on homogenized mechanical properties, using both analytical uncertainty quantification and Gaussian Process Regression (GPR). The study aims to improve understanding and analysis of microstructure uncertainties, particularly in additively manufactured materials.

An innovative, data-driven approach to Bayesian Uncertainty Quantification (BUQO) in image restoration, replacing traditional inpainting operators with convolutional neural networks (CNNs) was proposed [34]. The proposed method introduces a convolutional neural network for inpainting, replacing the conventional handcrafted inpainting operators. This allows for the creation of a plug-and-play BUQO algorithm that is adaptable to various structures without requiring specific tuning for each case. The method is tested on magnetic resonance imaging (MRI) simulations, demonstrating

improved qualitative outputs compared to the original BUQO method. The CNN inpainting operator is shown to provide more natural and explainable images, making the approach more robust and scalable for different applications, including medical imaging and potentially other fields like astronomical imaging. The authors also discuss the computational aspects and performance of the proposed method, highlighting its advantages over traditional approaches.

In conclusion, the literature surrounding uncertainty quantification in deep learning and image inpainting presents a broad spectrum of methodologies, each contributing unique insights to the field. From the foundational use of dropout as a Bayesian approximation to advanced GAN-based priors, the research consistently highlights the critical role of uncertainty in improving model reliability and performance. Across domains such as material science, generative models, and biomedical imaging, uncertainty-aware approaches are shown to enhance predictions, provide diverse outputs, and offer deeper insights into model confidence. Collectively, these works underline the importance of integrating uncertainty quantification into AI-driven solutions, not only to improve accuracy but also to ensure robustness in applications ranging from critical infrastructure to creative AI systems. Future research will likely continue refining these techniques, exploring new domains, and further bridging the gap between deterministic models and probabilistic interpretations.

Chapter 3

Methodology

The methodology of this study is designed to address the challenges of reconstructing incomplete mechanical field maps using deep learning techniques. The study employs a Generative Adversarial Network (GAN) framework, integrating various architectural innovations and regularization techniques to enhance the quality and reliability of the reconstructed images. The primary goal is to quantify uncertainty of the reconstructions using a robust model capable of handling missing data. This section details the dataset preparation, the architecture of the GAN, the implementation of dropout for regularization and uncertainty estimation, and the evaluation metrics used to assess the model's performance.

3.1 Dataset Overview

The dataset used in this research is derived from the study "Transferrable Deep Learning Approaches to Recover Missing Physical Field Information" by Zhenze Yang and Markus J. Buehler [12]. This dataset includes synthetic 2D mechanical field data generated through finite element analysis (FEA), designed to simulate various scenarios in materials science. The primary focus of this dataset is on 2D composite microstructures subjected to mechanical loading, specifically uniaxial tension tests. The dataset captures the resulting stress fields across different composite configurations as shown in Figure 3.1, providing a comprehensive basis for deep learning model training and validation.

3.1.1 Dataset Generation and Structure

This section elaborates how dataset is generated and explains its structure.

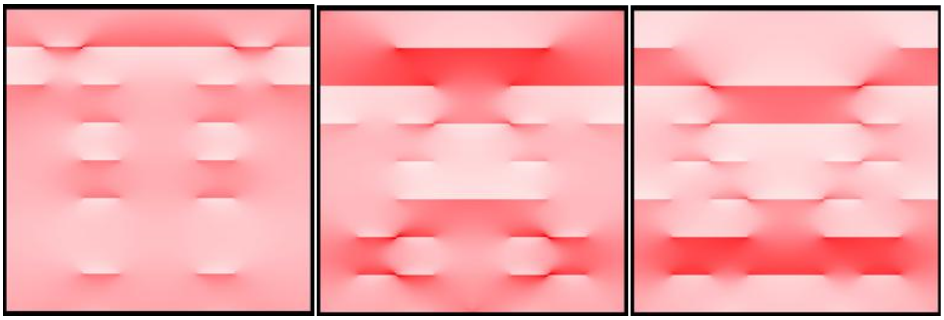

Figure 3.1: Mechanical stress fields dataset

- **Mechanical Fields and Components:** The dataset represents stress fields (σ_{11}), (where 1 denotes the x-direction) in 2D digital composites with linear elasticity under uniaxial tension. These stress fields are critical for understanding the mechanical response of the composites under loading conditions. The dataset includes both regular (e.g., square, triangle, circle) and irregular (e.g., brush stroke) masked regions to simulate incomplete or partially observed data scenarios.

- **Grid and Mask Variations:** The dataset features stress fields arranged on an 8x8 grid, providing a controlled environment for testing the robustness of the deep learning models. Masks of varying shapes and sizes [14] are applied to the field maps, creating scenarios where parts of the data are missing or obscured. This simulates real-world situations in nondestructive testing and other applications where complete data is often unavailable.

- **Training and Test Splits:** For the training and validation of the models, the dataset is split into training and test subsets. The training set is used to develop the model's predictive capabilities, while the test set evaluates the model's performance on unseen data. The training dataset comprises 800 examples, while the test dataset includes 200 examples, ensuring a rigorous assessment of model accuracy and generalizability.

- **Data Augmentation and Preprocessing:** To enhance the model's ability to generalize across different scenarios, data augmentation techniques such as rotation, scaling, and flipping are applied. Additionally, the dataset is normalized to ensure consistent input ranges for the models, improving convergence during training.

3.1.2 Importance of the Dataset

The dataset is crucial for developing and testing deep learning models that can accurately reconstruct mechanical fields from incomplete data. Its diverse range of stress

fields and masking scenarios provides a robust test bed for evaluating model performance in handling complex and ill-posed inverse problems. The use of synthetic data ensures that the models can be rigorously tested in controlled environments before being applied to real-world data.

This comprehensive dataset forms the foundation for the methodologies and results discussed in the following sections of this thesis. The next sections will detail the specific deep learning architectures and algorithms used to process this data and the outcomes of these methods in terms of accuracy and practical applicability.

3.2 GAN Architecture

A Conditional Generative Adversarial Network (cGAN) from [12] is employed to reconstruct incomplete mechanical field maps. This architecture is particularly suited for situations where only partial data is available, as it can learn to generate the missing portions of the data based on the observed portions. The cGAN model consists of two main components: the Generator and the Discriminator, which are trained together in a competitive setting.

3.2.1 Generator Architecture

The Generator in the cGAN architecture is responsible for generating the missing parts of the mechanical field maps. It takes as input the masked (incomplete) field map and aims to produce a complete map that includes both the observed and the missing data. The architecture of the Generator is as follows:

- **Input Layer:** The input consists of the masked field map, which includes both the observed and the masked regions.
- **Convolutional Layers:** The Generator uses a series of convolutional layers to capture local patterns and features within the input data. These layers are equipped with rectified linear unit (ReLU) activations, which introduce non-linearity to the model and help in learning complex patterns. Gated convolutions are employed to dynamically control the feature flow and effectively fill the missing regions. This mechanism allows the model to decide which features to use and which to ignore, enhancing the quality of the generated output.

- **Upsampling Layers:** The convolutional layers are followed by upsampling layers, which increase the spatial resolution of the feature maps. This is crucial for generating high-resolution output that matches the size of the input field map. Skip connections are used to incorporate information from earlier layers, preserving fine details and improving the quality of the generated output.
- **Output Layer:** The final layer of the Generator outputs the complete field map, including the previously missing regions. The output is typically passed through a Tanh activation function to ensure the generated values are within a specified range.

3.2.2 Discriminator Architecture

The Discriminator's role is to distinguish between real (ground truth) and fake (generated) field maps. It serves as a critic that helps the Generator improve its outputs by providing feedback on their realism. The architecture of the Discriminator is as follows:

- **Input Layer:** The Discriminator receives both the real and generated field maps as inputs. These maps are labeled accordingly, allowing the Discriminator to learn to differentiate between them.
- **Convolutional Layers:** Similar to the Generator, the Discriminator employs convolutional layers to extract features from the input maps. However, these layers are followed by Leaky ReLU activations, which prevent the model from dying gradients and improve learning stability.
- **Pooling and Flattening Layers:** Pooling layers reduce the spatial dimensions of the feature maps, summarizing the information and reducing the computational load. The feature maps are then flattened into a one-dimensional vector.
- **Fully Connected Layers:** The flattened vector is passed through fully connected layers, which output a single value representing the probability that the input map is real. This probability is typically passed through a sigmoid activation function.

3.2.3 Loss Function and Training

The cGAN model is trained using a combination of adversarial and reconstruction losses:

- **Adversarial Loss:** This loss encourages the Generator to produce outputs that are indistinguishable from real data. It is calculated based on the Discriminator's ability to correctly classify real and generated maps.
- **Reconstruction Loss:** This loss measures the difference between the generated and real field maps in the unmasked regions. It ensures that the generated maps are not only realistic but also accurate.

The training process involves alternating between optimizing the Generator and Discriminator, allowing the Generator to improve its outputs while the Discriminator becomes better at detecting fakes.

The cGAN architecture is a powerful tool for reconstructing incomplete mechanical field maps in the 2D case. Its ability to generate realistic and accurate data based on partial observations makes it highly suitable for addressing inverse problems in materials science.

3.3 Incorporating Monte Carlo Dropout in GAN Architecture

Dropout is a well-established regularization technique that mitigates overfitting in deep neural networks. It works by randomly deactivating a fraction of neurons during training, preventing the model from relying too heavily on specific paths through the network. This stochastic element encourages the network to learn redundant representations of data, leading to improved generalization when the model is exposed to unseen data. In the context of Generative Adversarial Networks (GANs), dropout can serve dual purposes: it helps stabilize the training process and provides a mechanism for uncertainty estimation. In this study, we integrated dropout into the Generator component of a Conditional GAN (cGAN) architecture, focusing on its effects on generating complete mechanical field maps from partially observed data.

3.3.1 Dropout Configurations

This research explores two configurations of applying dropout in hidden layers. First, the dropout is applied in first layer of coarse and fine generator with varying dropout rates. Second case involves application of dropout in all hidden layers which involves, down sampling block, up sampling block, coarse generator and fine generator, with varying dropout rates for each experiment. Figure 3.2 illustrates these two cases of

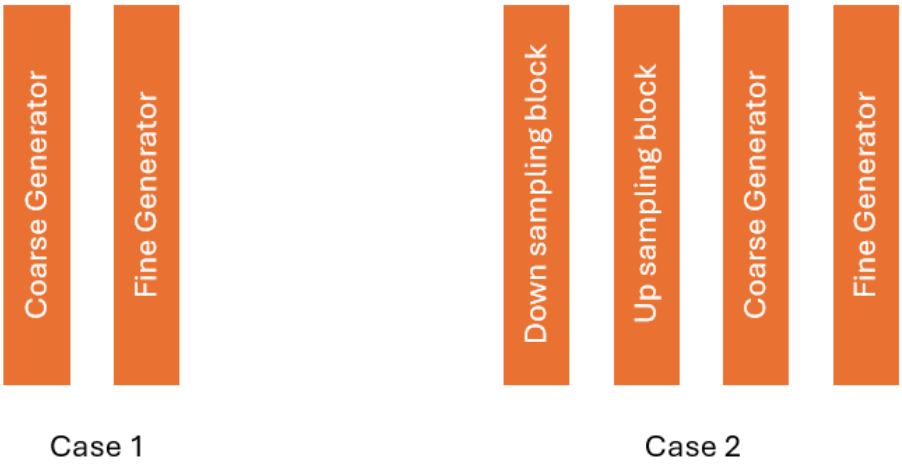

Figure 3.2: Dropout applied at the Generator for both the cases

dropout configuration.

In deep learning architectures, down sampling blocks reduce the spatial resolution of input data while capturing higher-level features, whereas up sampling blocks restore the resolution by increasing the spatial dimensions of feature maps. These processes are often used in encoder-decoder structures like U-Nets and GANs. In image generation tasks, a coarse generator first creates a rough, low-resolution version of the target output, capturing broad structures. The fine generator then refines this coarse output by adding detailed textures and improving visual quality. Together, these components enable efficient image inpainting, super-resolution, and other tasks requiring high-quality image reconstruction.

Also, the application of dropout in the discriminator of Generative Adversarial Networks (GANs) is not considered because the role of the discriminator is to provide strong, reliable feedback to the generator during training. The discriminator's task is to distinguish between real and generated data, and it needs to be as accurate as possible to guide the generator in improving its output. Introducing dropout in the discriminator would introduce randomness and weaken its ability to accurately classify images, leading to less effective feedback for the generator. Additionally, the primary purpose of dropout is to prevent overfitting, but in GANs, the discriminator is trained in conjunction with the generator, which reduces the risk of overfitting since the data distribution changes throughout the training process. Thus, maintaining a strong, stable discriminator without dropout is crucial for GAN performance.

3.3.1.1 Dropout in the First Hidden Layer

In the first configuration, dropout was applied exclusively to the first hidden layer of the coarse and fine generator. The rationale behind this choice was to assess how regularization at the initial stages of data transformation impacts the learning process. The dropout rates experimented with were 0.2, 0.4, and 0.6.

- **Dropout Rate (0.2):** This relatively low dropout rate was intended to provide moderate regularization, allowing the network to maintain a balance between learning useful features and avoiding overfitting.
- **Dropout Rate (0.4):** As a more aggressive regularization strategy, this dropout rate aimed to further reduce reliance on specific features, promoting the learning of more robust, generalized features.
- **Dropout Rate (0.6):** This high dropout rate was used to test the limits of the model's capacity to generalize under severe regularization. While high dropout rates can improve generalization, they can also hinder the network's ability to learn if too much information is discarded.

The application of dropout in the first hidden layer specifically targets the network's initial feature extraction phase. By forcing the model to learn diverse representations, we aimed to enhance its capacity to generalize from incomplete data.

3.3.1.2 Dropout across all Hidden Layers

In the second configuration, dropout was applied uniformly across all hidden layers of the Generator. The same set of dropout rates (0.2, 0.4, and 0.6) was used. This comprehensive application of dropout was intended to ensure that all stages of the data transformation process were regularized.

3.3.2 Motivation for Uniform Dropout Rates

Applying dropout across all layers prevents the network from overfitting at any stage, not just at the initial feature extraction phase. This comprehensive approach encourages the model to develop a more holistic understanding of the data, making it less susceptible to overfitting specific patterns.

3.4 Monte Carlo Dropout for Uncertainty Estimation

Beyond its regularization effects, dropout was also utilized as a tool for uncertainty estimation during testing, a technique known as Monte Carlo Dropout. In this approach, dropout is applied during the testing phase as well, and the model is run multiple times with different random drops, generating a distribution of outputs.

3.4.1 Uncertainty Quantification

By analyzing the variance in these outputs, we can quantify the model's uncertainty in its predictions. This is particularly valuable in scenarios involving inverse problems, where multiple internal structures may correspond to similar boundary observations. The ability to provide not just a single solution but a range of possible solutions, each with an associated confidence level, is crucial for practical applications.

3.5 Experimental Setup and Training Protocol

The cGAN architecture was trained using a dataset consisting of mechanical field maps, with both incomplete and complete data provided as training pairs. The training involved the following steps:

- **Data Preprocessing:** The input data was normalized and augmented with random noise to simulate partial observations. Masks representing the missing parts of the data were also generated.
- **Model Training:** The Generator and Discriminator were trained simultaneously in an adversarial framework. The GAN was trained using the Adam optimizer, with hyperparameters fine-tuned based on the validation set performance. The learning rate was set with a decay schedule to ensure smooth convergence. The model was trained until convergence, with early stopping employed to prevent overfitting.
- **Dropout Activation:** During training, the specified dropout rates were applied. This involved randomly deactivating neurons in the hidden layers, as described in the previous sections. The training process was monitored for overfitting, with early stopping implemented if necessary.

- **Testing and Validation:** The model's performance was evaluated using a separate validation set. Monte Carlo Dropout was employed at inference to generate multiple outputs for each test input, allowing for the assessment of both accuracy and uncertainty. Performance metrics such as mean absolute error, and standard deviation were calculate. Visual inspections of the reconstructed images and uncertainty maps were also performed to qualitatively assess the model's outputs.

Monte Carlo Dropout was employed during the testing phase to generate multiple predictions for the same input, allowing for the estimation of both the mean prediction and the associated uncertainty. Whereas a dropout rate determines what fraction of those neurons shall be dropped. This approach is particularly useful for quantifying uncertainty in inverse problems where solutions may not be unique. In this process, multiple predictions were generated for the same input image by enabling MC dropout during inference. When used during inference, this randomness creates slightly different outputs for the same input image and each time the model's performance is evaluated using several metrics, including standard deviation and error metrics. The evaluation focuses on the masked regions where reconstruction is necessary. The training procedure involves the Adam optimizer with a decay schedule, early stopping to prevent overfitting, and extensive validation to fine-tune hyperparameters.

For this research, we generated 50 images per input image, this means we are running the model 50 times on the same input image, with dropout enabled each time. These multiple predictions allow us to estimate uncertainty. Uncertainty is calculated as the standard deviation across the predictions for each pixel is given by

$$\text{Uncertainty at pixel } (i,j) = \sqrt{\frac{1}{n_{iter}} \sum_{k=1}^{n_{ensemble}} (x_{k,ij} - mean(x_{ij}))^2}, \tag{3.1}$$

where

$$n_{ensemble} = 50, \tag{3.2}$$

for our case and it is the number of times the model generates predictions for the same input with dropout enabled. k is the iteration number.

In addition, average uncertainty which gives a sense of the overall uncertainty level in the regions where inpainting was performed, is calculated is calculated as given by

$$\text{Average Uncertainty} = \frac{1}{N} \sum_{i-1}^{N} \text{uncertainty } (i), \tag{3.3}$$

where N is the number of pixels with non-zero uncertainty.

Mean Absolute Error (MAE) is calculated to quantify the average magnitude of prediction errors without considering their direction. It reflects how close the model's predictions are to the actual values, particular for each pixel in image inpainting. By averaging the absolute differences, MAE offers a reliable indication of the model's overall performance, highlighting the error size in reconstructed areas. A lower MAE indicates that the model is accurately filling in missing information.

MAE is calculated by first computing the absolute difference between predicted and ground truth pixel values within the masked (inpainted) regions. These absolute errors are averaged only over the masked pixels to provide a single MAE value specific to the reconstructed areas, capturing the average magnitude of prediction errors in these regions.

The MAE calculation can be expressed as:

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^{N} |y_i - \hat{y}_i| \tag{3.4}$$

where, y_i is the ground truth pixel value, $\hat{y}_i$ is the predicted pixel value, and N is the number of pixels in the masked region.

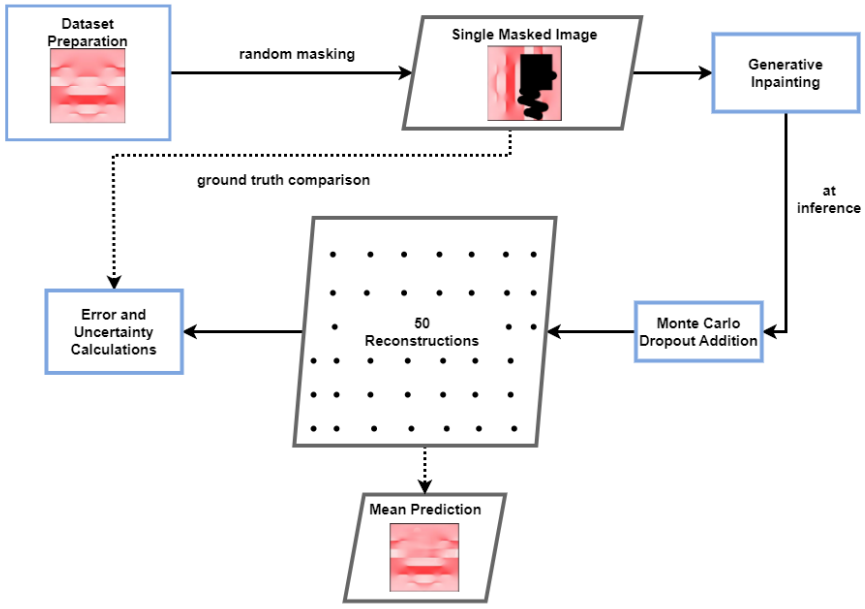

Figure 3.3: Methodology for uncertainty quantification of inpainted stress field images using GANs

In addition, a final single predicted/ mean image is also generated by computing the

mean of all 50 multiple predictions for a single image. This single predicted image balances out the noise and variability from individual predictions and it offers a measure of confidence based on how much the predictions vary.

Analyses of the results were conducted to assess the impact of different dropout rates on model performance, focusing on mean error and standard deviation. Visualizations, including reconstructed images, uncertainty surface plot, and error plots, were generated to qualitatively and quantitatively assess the model's outputs.

The methodology shown in Figure 3.3, which is employed in this study demonstrates the effectiveness of GANs, particularly with dropout for uncertainty estimation, in handling inverse problems in materials science. The use of dropout not only prevents overfitting but also provides valuable probabilistic information about the model's predictions. This capability is crucial in applications where multiple plausible solutions exist, and decision-making requires understanding the uncertainty associated with each solution.

Chapter 4

Results and Discussion

In this study, we thoroughly examined the effects of varying dropout configurations on the performance of a Generative Adversarial Network (GAN) used to reconstruct incomplete mechanical field maps. Our focus centered on two key configurations: applying dropout solely to the first hidden layer of the generator and applying dropout across all hidden layers. We evaluated the model's performance by analyzing the standard deviation of predictions, indicative of uncertainty, and the error metrics, across different dropout rates (0.2, 0.4, and 0.6).

4.1 Comparative Analysis: First Hidden Layer vs All Hidden Layers

4.1.1 First Hidden Layer

The results highlighted distinct behavioral patterns depending on where and how dropout was applied. When dropout was introduced exclusively to the first hidden layer of both the coarse and fine generators, the model exhibited stability in terms of uncertainty and error.

At a dropout rate of 0.2, the model achieved a low average uncertainty (0.016) and mean absolute error (11.8), indicating that the model was effectively regularized without sacrificing accuracy. Dropout in this layer prevented overfitting by deactivating a fraction of neurons, which led the network to generalize better while still maintaining its capacity to predict accurately. As the dropout rate was increased to 0.4, there was a slight rise in both uncertainty (0.028) and MAE (14.75), but these changes were modest. However, when the dropout rate was further increased to 0.6, the results became

more surprising. The average uncertainty slightly decreased to 0.012, along with MAE (13.38). This suggests that, in some cases, applying a higher level of dropout may help the model focus on learning more essential, invariant features by effectively ignoring noise or irrelevant features in the data. As a result, the network becomes more stable and less sensitive to small variations in input.

4.1.2 Effects of Dropout across all Hidden Layers

When dropout was applied uniformly across all hidden layers, the model's behavior was markedly different. At a dropout rate of 0.2, the average uncertainty increased significantly to 0.061, with the mean absolute error rising to 25.78, which is a surge in its value compared to the first-layer configuration. This suggests that applying dropout across all layers introduces greater variability in feature learning. While dropout in the first layer helps to regularize early feature extraction, applying it across all layers disrupts the flow of information through the network, leading to greater uncertainty and less reliable predictions. At a dropout rate of 0.4, the uncertainty and error spiked even further, with uncertainty reaching 0.185 and MAE adjusting to 44.31. This sharp increase in uncertainty suggests that the model struggled to maintain consistency as regularization intensified. Dropout across multiple layers increases the likelihood of essential neurons being deactivated, causing the network to lose important features and making it more difficult for the network to form stable representations. This leads to higher uncertainty and less precise outputs. At the highest dropout rate of 0.6, the effects were even more pronounced, with average uncertainty surging to 0.466 and the mean absolute error increasing sharply to 109.04. This degradation in performance highlights the dangers of over-regularization. When dropout is applied too aggressively across all layers, the network's ability to learn meaningful patterns is severely compromised, leading to poor generalization and significantly higher uncertainty. This behavior is a clear indication that dropout, when applied across all layers, disrupts the model's capacity to extract and retain important information, ultimately impairing its performance. Table 4.1 shows the results for both the cases.

To elaborate the results more, a graph in Figure 4.1 depicts the average uncertainty across multiple test data images for different dropout configurations and dropout rates (0.2, 0.4, 0.6). The analysis is divided into the two key configurations: applying dropout solely to the first hidden layer and applying dropout uniformly across all hidden layers.

Dropout in the First Layer: Dropout confined to the first hidden layer maintains relatively low uncertainty across all dropout rates. This configuration provides effective regularization without causing excessive instability in the model's predictions, making it a reliable option for tasks requiring confident reconstructions.

Table 4.1. Average uncertainty and average error for various dropout rates for both cases

Dropout rate	Average Uncertainty		Average Error	
	First layer	All layers	First layer	All layers
0.2	0.016	0.061	11.8	25.78
0.4	0.028	0.185	14.75	44.31
0.6	0.012	0.466	13.38	109.04

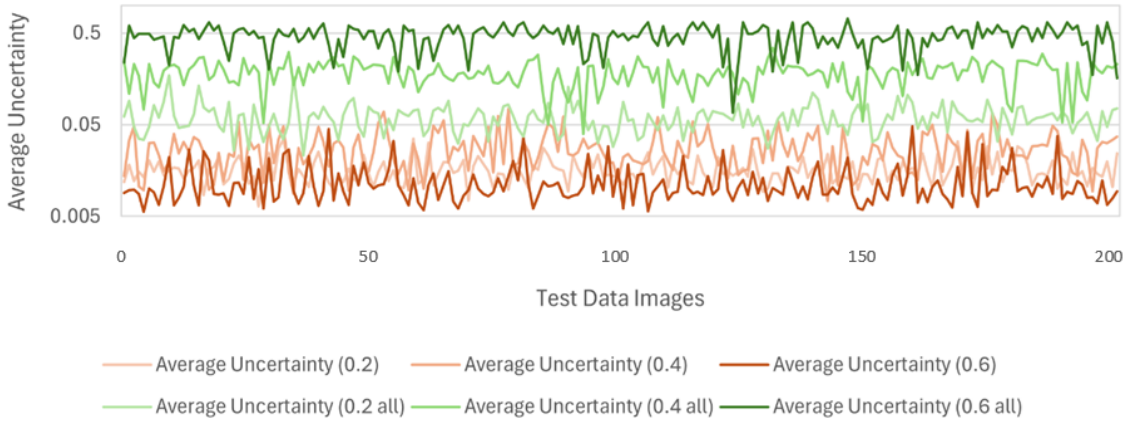

Figure 4.1: Average uncertainty graph for all the cases

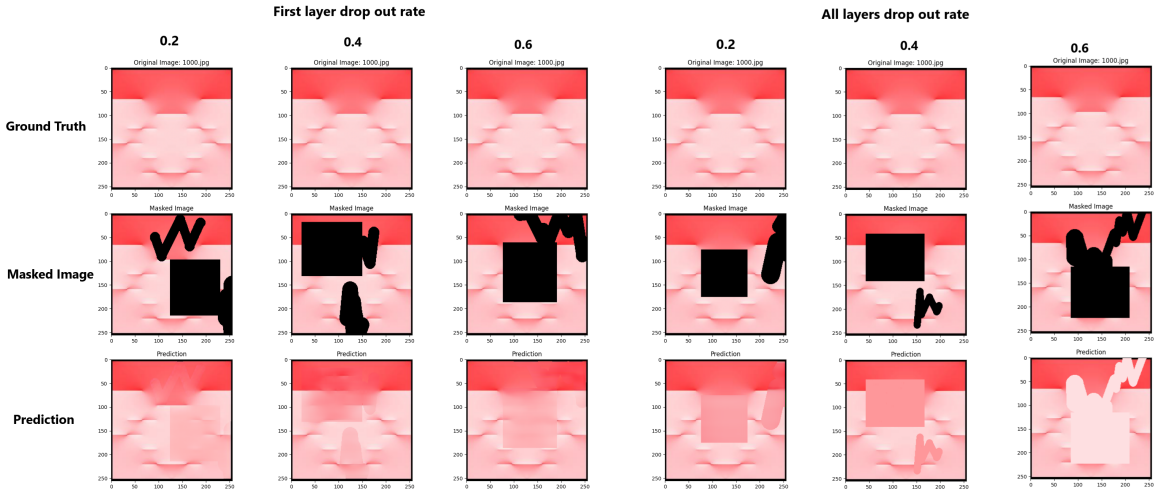

Figure 4.2: Increasing drop-out enabled results for first layer and all layers (L to R)

Dropout Across All Layers: In contrast, applying dropout across all layers results in higher average uncertainty and greater variance, particularly at higher dropout rates. While this may encourage the exploration of a broader range of features, it also introduces more variability, which could negatively impact the model's ability to make consistent predictions.

Figure 4.2 shows visually the effect of applying drop-out on inpainting in first layer with increasing drop out and then all layers. The first row in this image represents the original, unmasked image, which is then masked (second row). The masked image is then fed into GAN based model and after the reconstructions are generated, a mean prediction is generated which is a single image (third row). The first three columns in this image corresponds to the application of dropout in first layer with increasing dropout rates i.e. 0.2, 0.4, and 0.6. The last three columns correspond to the application of dropout in all the hidden layers with increasing dropout rates i.e 0.2, 0.4, and 0.6. The image clearly depicts that the stress field reconstructions deviate from the original image or the ground truth with increasing dropout rates and also with changing dropout configuration.

The dropout effect can be observed better visually through uncertainty quantification. To view how much uncertainty is present pixel wise, 3D surface plots were generated as shown in Figure 4.3. Here the brighter regions show higher uncertainty, and dark regions show low uncertainty. It can be easily observed that as the dropout rate shifts from first hidden layer to all hidden layers, uncertainty increase drastically with increasing dropout rates. This can help us in identifying easily which areas of the

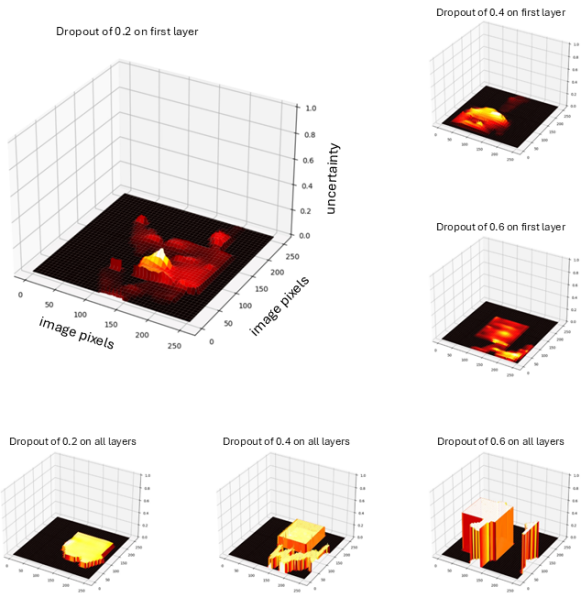

Figure 4.3: Pixel-wise surface plots of uncertainty for both the cases

reconstructed stress fields needs to be questioned and which areas can we rely upon our inpainting model.

To validate our results, we compared our findings on the impact of dropout configurations in Generative Adversarial Networks (GANs) with similar research [30]. Our study focused on how varying dropout rates (0.2, 0.4, 0.6) affected the reconstruction of mechanical field maps, specifically contrasting dropout applied to the first hidden layer versus all hidden layers. Our results indicated that increased dropout across all layers led to performance degradation, aligning with previous observations of higher dropout rates causing variability and image distortions. The referenced study also applied dropout during both training and generation, highlighting that while higher rates increased output diversity, they compromised image quality and coherence, which echoed our findings of increased uncertainty and error with the "all hidden layers" configuration. Both studies underscored the necessity of selective dropout application; in our case, dropout applied only to the first hidden layer preserved performance and mitigated overfitting, whereas the reference study found that early-layer dropout enhanced image quality through a "repair" mechanism in later layers. Overall, this comparison reinforces the validity of our results, suggesting that targeted dropout application, particularly in early layers, sustains model stability and performance while avoiding the pitfalls of over-regularization.

4.2 Implications for Model Confidence and Uncertainty Estimation

The integration of Monte Carlo dropout into the GAN architecture for stress fields reconstruction offers critical insights into the reliability of deep learning models in handling inverse problems. The evaluation of uncertainty serves as a confidence measure, allowing for more informed decision-making in areas where the model's predictions may be less certain. This aspect is particularly valuable in applications such as failure analysis or inefficiencies in predicting stress distributions. By assessing pixel-level uncertainty through multiple stochastic forward passes, this study reveals how different dropout configurations influence the stability and confidence of model predictions.

4.2.1 Significance of Findings

The findings of this study provide important advancements in the field of materials science, particularly in the context of using deep learning techniques for solving inverse problems related to stress field reconstruction. By integrating Monte Carlo (MC) dropout into a conditional GAN (cGAN) architecture, this research demonstrates that uncertainty quantification can effectively highlight the areas where model predictions may be unreliable. This has significant implications for applications in failure analysis, non-destructive testing, and other domains where the integrity of material predictions is critical. The ability to visualize uncertainty at a pixel-wise level introduces a new layer of confidence in AI-driven methods, allowing for informed decision making based on the degree of reliability in the reconstructed fields. Furthermore, the comparative analysis between first-layer dropout and dropout across all layers offers practical insights for configuring dropout to achieve a balance between regularization and predictive accuracy.

The results also highlight the critical trade-off between applying dropout for model generalization versus maintaining reliable predictions. Dropout applied across all layers increased the exploration of feature space but introduced higher uncertainty, which can be detrimental for applications requiring consistent predictions. Conversely, dropout applied only in the initial layers maintained stability and led to more dependable reconstructions. These insights are essential for optimizing neural networks in material science and other fields that depend on precise predictions under incomplete or noisy data conditions.

Chapter 5

Conclusion

This thesis explored the role of uncertainty quantification in the reconstruction of missing mechanical field maps using Generative Adversarial Networks (GANs) with a focus on dropout mechanisms. The primary objective of this research was to quantify and analyze the effects of varying dropout configurations on GAN performance in terms of uncertainty and error. Through the investigation, we aimed to improve the reliability and accuracy of predictive models in reconstructing incomplete mechanical fields, while providing a deeper understanding of uncertainty quantification, a crucial factor in many high-stakes applications within materials engineering, machine learning, and beyond.

5.1 Uncertainty Quantification and Its Importance

Uncertainty quantification (UQ) has become an essential tool in many scientific and engineering fields, particularly those where accurate predictions are necessary for decision-making. The ability to assess the reliability of a model's predictions, especially in scenarios with incomplete or noisy data, allows researchers and engineers to make more informed choices and manage risks effectively. Throughout this thesis, UQ served as the foundation for understanding the confidence in the outputs generated by GANs, a neural network architecture known for its strengths in tasks like image generation, inpainting, and field reconstruction.

In materials engineering, the reconstruction of mechanical stress fields is often hindered by incomplete data due to the challenges of obtaining full-field measurements in real-world conditions. Traditional deterministic models struggle to accurately predict these missing fields, especially when dealing with complex stress fields and non-linear

material behavior. By incorporating uncertainty quantification into GAN-based models, we have taken a critical step toward bridging this gap, providing more reliable predictions and enhancing the model's generalization ability across various datasets.

5.2 Dropout Mechanisms and GAN Performance

Dropout, traditionally used as a regularization technique to prevent overfitting in neural networks, was employed in this study as a means of quantifying uncertainty. The focus was placed on applying dropout selectively in different parts of the GAN architecture to observe how it affects the network's ability to reconstruct missing mechanical fields. Two primary configurations were examined: applying dropout solely to the first hidden layer and applying dropout uniformly across all hidden layers. The results of these configurations revealed key insights into the behavior of dropout and its implications for uncertainty quantification.

When dropout was applied exclusively to the first hidden layer of the coarse and fine generators, the model exhibited stability and robustness. The results showed that moderate dropout rates, particularly 0.4, struck a balance between reducing overfitting and preserving the model's capacity to predict accurately. At this rate, the network generated reconstructions with low uncertainty and error, highlighting the effectiveness of first-layer dropout in promoting generalization while avoiding the pitfalls of co-adaptation between neurons. By forcing the network to rely less on specific neurons, dropout at this layer encouraged the model to learn more distributed and resilient feature representations, which is critical for complex tasks like mechanical field inpainting.

Conversely, applying dropout across all hidden layers introduced significant variability and instability into the model's predictions. At a low dropout rate of 0.2, uncertainty and error were already higher compared to the first-layer configuration, and as the dropout rate increased, these metrics worsened. This configuration disrupted the flow of information throughout the network, leading to inconsistent feature learning and degraded performance. The results suggest that over regularization, particularly when applied uniformly across all layers, severely compromises the network's ability to retain meaningful patterns, resulting in poor generalization and unreliable predictions.

5.3 Comparative Analysis with Existing Research

The findings of this thesis were compared against similar studies in the field, particularly those exploring the use of dropout in GANs for image generation tasks. In these

studies, dropout was applied both during the training and generation phases to increase output diversity. While these studies primarily focused on generating multiple outputs from a single input, they also highlighted the importance of selective dropout application to maintain image coherence and quality. The parallels drawn between these works and our study reinforced the conclusion that moderate dropout in early layers offers a balance between diversity and performance, whereas excessive dropout across all layers leads to instability and loss of crucial information.

Additionally, the reference studies illustrated that applying dropout only to the first hidden layer allowed subsequent layers to "repair" errors induced by dropout, a phenomenon observed in our own results. This repair mechanism is essential for maintaining the structural and semantic integrity of the generated outputs, particularly in tasks requiring high levels of detail and accuracy, such as mechanical field reconstruction. By comparing our findings with those from other domains, we established that the principles of selective dropout application hold across various tasks, further validating our conclusions.

5.4 Significance of this Research

The significance of this research lies in several key areas:

- **Enhancement of Deep Learning based Inpainting Techniques:** This research advances the field of deep learning-based image inpainting by incorporating a probabilistic understanding of the model's outputs. The methodology developed here can be applied to enhance the robustness of deep fill techniques, particularly in scenarios where the completeness and accuracy of the reconstruction are crucial. For instance, in the reconstruction of medical images or nondestructive testing data, knowing the areas of high uncertainty can help focus further analysis or data collection efforts.
- **Broader Implications for Inverse Problems:** The approach and findings have broader implications beyond the specific task of mechanical field reconstruction. The ability to estimate uncertainty is a valuable addition to any inverse problem-solving framework, providing a way to account for the inherent ambiguities and variabilities in real-world data. This is particularly relevant in fields like biomedical imaging, geophysics, non-destructive testing and remote sensing, where the data often come with significant noise and incomplete information.

- **Guidance for Neural Network Regularization:** The detailed analysis of dropout's effects across different layers provides valuable guidance for neural network design and regularization. It suggests that the application of dropout should be carefully tailored based on the specific layers and the nature of the task. For example, in tasks where fine-grained details are crucial, reducing dropout rates or limiting dropout to specific layers might be more beneficial. Conversely, for tasks where generalization is more important, a higher dropout rate might be justified to prevent overfitting.
- **Implications for Real-World Applications:** The practical implications of these findings are far reaching. In fields such as autonomous driving, aerospace engineering, and financial modeling, where predictive accuracy and understanding of uncertainty are paramount, the techniques developed in this research can significantly enhance the robustness and reliability of deployed models. By providing a measure of uncertainty, these techniques enable more cautious and informed decision-making processes, potentially preventing costly errors and improving overall system safety and effectiveness.

5.5 Implications for Materials Engineering and Machine Learning

The application of GANs with uncertainty quantification in materials engineering holds significant potential for advancing the field. In this thesis, we demonstrated that incorporating dropout into GAN architectures allows for more reliable reconstruction of mechanical stress fields, a critical task for predicting material behavior under load. The ability to accurately predict missing fields, while also providing a measure of uncertainty, is invaluable for engineers designing materials for high-stakes applications, such as aerospace, automotive, and civil engineering.

The implications extend beyond materials science. In machine learning, uncertainty quantification is becoming increasingly important in domains such as autonomous driving, healthcare, and financial modeling, where making confident decisions based on model outputs is crucial. By showing how dropout can be effectively leveraged for uncertainty estimation in GANs, this research contributes to the broader field of machine learning by providing insights into how to make neural networks more interpretable and trustworthy. These advancements are particularly relevant for tasks that involve incomplete data, high-dimensional spaces, or situations where the consequences of incorrect predictions are severe.

5.6 Future Research Directions

While this thesis provided significant insights into the use of dropout for uncertainty quantification in GANs, several avenues for future research remain. While this study focused on dropout, other regularization techniques such as batch normalization, spectral normalization, and weight decay could also be explored. Comparative studies could reveal the relative advantages and trade-offs of these techniques in different neural network architectures and applications.

Additionally, the scope of this study was limited to reconstructing 2D mechanical fields, but extending these methods to 3D scenarios would be an important next step. Many real-world applications, particularly in materials science and biomedical imaging, involve 3D data, and developing GAN architectures capable of handling such data would significantly expand the applicability of this research.

Another promising direction is the incorporation of domain-specific knowledge into the GAN architecture. For instance, in materials science, understanding the physical laws governing material behavior could be integrated into the model to further improve prediction accuracy. This approach would combine data-driven learning with traditional physics-based methods, creating hybrid models that leverage the strengths of both approaches.

Future work could explore the integration of the GAN-based framework with optimization algorithms. This combined approach could enhance the framework's capability in solving inverse design problems, where the goal is to iteratively refine material configurations or structural designs to achieve desired properties. Such an integration could leverage the probabilistic outputs of the GAN to explore a broader solution space, potentially discovering novel and optimal designs.

The methodologies developed in this research could be applied to other domains beyond materials science. For instance, in biomedical imaging, the ability to reconstruct missing information with quantified uncertainty could improve diagnostic tools. Similarly, in fluid dynamics, accurate and uncertain reconstructions of flow fields could aid in more robust simulations and analyses.

Finally, exploring the potential of adaptive dropout rates that change based on the complexity of the data or the training stage could yield further improvements in performance. Such techniques would allow the network to dynamically adjust its regularization, providing a more tailored approach to uncertainty quantification and model optimization.

5.7 Concluding Remarks

In conclusion, this thesis has demonstrated the power of uncertainty quantification in enhancing the performance of GANs for reconstructing incomplete mechanical field maps. By investigating the effects of varying dropout configurations and rates, we provided valuable insights into how dropout can be used not only as a regularization technique but also as a tool for estimating uncertainty in complex tasks. The findings of this research contribute to both the fields of materials engineering and machine learning, offering a pathway toward more reliable and interpretable models.

The future of uncertainty quantification in GANs is bright, with numerous opportunities for further exploration and refinement. As the demand for more accurate and trustworthy predictive models grows, the insights gained from this research will play a crucial role in shaping the development of next-generation machine learning systems capable of addressing the challenges of incomplete data and uncertainty in high-stakes applications.

Bibliography

[1] L. Cai, Y. Chen, N. Cai, W. Cheng, and H. Wang, "Utilizing amari-alpha divergence to stabilize the training of generative adversarial networks," *Entropy*, vol. 22, p. 410, 04 2020.

[2] W. Cai and Z. Wei, "Piigan: generative adversarial networks for pluralistic image inpainting," *IEEE Access*, vol. 8, pp. 48451–48463, 2020.

[3] E. Hüllermeier and W. Waegeman, "Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods," *Machine learning*, vol. 110, no. 3, pp. 457–506, 2021.

[4] C. Sammut, *Markov Chain Monte Carlo*, pp. 639–642. Boston, MA: Springer US, 2010.

[5] M. Fortunato, C. Blundell, and O. Vinyals, "Revisiting bayes by backprop," 2018.

[6] A. Damianou and N. D. Lawrence, "Deep gaussian processes," in *Artificial intelligence and statistics*, pp. 207–215, PMLR, 2013.

[7] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *Proceedings of The 33rd International Conference on Machine Learning* (M. F. Balcan and K. Q. Weinberger, eds.), vol. 48 of *Proceedings of Machine Learning Research*, (New York, New York, USA), pp. 1050–1059, PMLR, 20–22 Jun 2016.

[8] S. Nitish, "Dropout: a simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, p. 1, 2014.

[9] Y. Gal and Z. Ghahramani, "Bayesian convolutional neural networks with bernoulli approximate variational inference," *arXiv preprint arXiv:1506.02158*, 2015.

[10] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," in *Advances in Neural Information Processing Systems* (F. Pereira, C. Burges, L. Bottou, and K. Weinberger, eds.), vol. 25, Curran Associates, Inc., 2012.

[11] S. Fort, H. Hu, and B. Lakshminarayanan, "Deep ensembles: A loss landscape perspective," *arXiv preprint arXiv:1912.02757*, 2019.

[12] Z. Yang and M. J. Buehler, "Fill in the blank: transferrable deep learning approaches to recover missing physical field information," *Advanced Materials*, vol. 35, no. 23, p. 2301449, 2023.

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

[14] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang, "Free-form image inpainting with gated convolution," in *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4471–4480, 2019.

[15] Z. Yan, X. Li, M. Li, W. Zuo, and S. Shan, "Shift-net: Image inpainting via deep feature rearrangement," in *Proceedings of the European conference on computer vision (ECCV)*, pp. 1–17, 2018.

[16] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang, "Generative image inpainting with contextual attention," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5505–5514, 2018.

[17] A. Lahiri, A. K. Jain, S. Agrawal, P. Mitra, and P. K. Biswas, "Prior guided gan based semantic inpainting," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 13696–13705, 2020.

[18] H. Liu, Z. Wan, W. Huang, Y. Song, X. Han, and J. Liao, "Pd-gan: Probabilistic diverse gan for image inpainting," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9371–9381, 2021.

[19] X. Zhang, X. Wang, C. Shi, Z. Yan, X. Li, B. Kong, S. Lyu, B. Zhu, J. Lv, Y. Yin, *et al.*, "De-gan: Domain embedded gan for high quality face image inpainting," *Pattern Recognition*, vol. 124, p. 108415, 2022.

[20] A. Abu-Srhan, M. A. Abushariah, and O. S. Al-Kadi, "The effect of loss function on conditional generative adversarial networks," *Journal of King Saud University-Computer and Information Sciences*, vol. 34, no. 9, pp. 6977–6988, 2022.

[21] H. Chang, H. Zhang, L. Jiang, C. Liu, and W. T. Freeman, "Maskgit: Masked generative image transformer," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11315–11325, 2022.

[22] H. Chung, B. Sim, and J. C. Ye, "Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12413–12422, 2022.

[23] H. Chung, B. Sim, D. Ryu, and J. C. Ye, "Improving diffusion models for inverse problems using manifold constraints," *Advances in Neural Information Processing Systems*, vol. 35, pp. 25683–25696, 2022.

[24] D. V. Patel, D. Ray, and A. A. Oberai, "Solution of physics-based bayesian inverse problems with deep generative priors," *Computer Methods in Applied Mechanics and Engineering*, vol. 400, p. 115428, 2022.

[25] A. Lugmayr, M. Danelljan, A. Romero, F. Yu, R. Timofte, and L. Van Gool, "Repaint: Inpainting using denoising diffusion probabilistic models," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11461–11471, 2022.

[26] T. Kim and B. C. Yeo, "Recovering microscopic images in material science documents by image inpainting," *Applied Sciences*, vol. 13, no. 6, p. 4071, 2023.

[27] B. Ma, B. Ma, M. Gao, Z. Wang, X. Ban, H. Huang, and W. Wu, "Deep learning-based automatic inpainting for material microscopic images," *Journal of Microscopy*, vol. 281, no. 3, pp. 177–189, 2021.

[28] A. Tran and H. Tran, "Data-driven high-fidelity 2d microstructure reconstruction via non-local patch-based image inpainting," *Acta Materialia*, vol. 178, pp. 207–218, 2019.

[29] Y. Miao, Y. Sun, S. Li, P. Zhang, Y. Yang, and Y. Hu, "Spinal neoplasm image inpainting with deep convolutional neutral networks," in *2019 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pp. 2619–2624, IEEE, 2019.

[30] S. Wieluch and F. Schwenker, "Dropout induced noise for co-creative gan systems," in *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pp. 0–0, 2019.

[31] X. Ma, X. Zhou, H. Huang, G. Jia, Y. Wang, X. Chen, and C. Chen, "Uncertainty-aware image inpainting with adaptive feedback network," *Expert Systems with Applications*, vol. 235, p. 121148, 2024.

[32] D. V. Patel and A. A. Oberai, "Gan-based priors for quantifying uncertainty," *arXiv preprint arXiv:2003.12597*, 2020.

[33] A. Senthilnathan, I. Javaheri, H. Zhao, V. Sundararaghavan, M. DeGraef, and P. Acar, "Uncertainty quantification of metallic microstructures using principal image moments," *Computational Materials Science*, vol. 215, p. 111775, 2022.

[34] M. Tang and A. Repetti, "A data-driven approach for bayesian uncertainty quantification in imaging," *arXiv preprint arXiv:2304.11200*, 2023.