International Journal of Research Publication and Reviews, Vol (6), Issue (5), May (2025), Page – 5935-5943

International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

Exploring the Dynamics of Diffusion Models

Aditya kumar

Bachelor of Technology, Computer Science & Engineering, Arya College Of Engineering & I.T., Kukas, Jaipur-302028. aditya112k@gmail.com

Abstract—

Diffusion models have become a powerful and adaptable type of generative model in artificial intelligence, showing impressive capabilities in creating high-quality, varied, and realistic data. Unlike earlier generative models, such as Generative Adversarial Networks (GANs), diffusion models function through a process inspired by the physical concept of diffusion. The Diffusion Process In the forward process, a diffusion model systematically adds noise to an input data sample, slowly transforming it into pure noise. This approach can be mathematically described as a Markov chain, with each step involving the introduction of a small amount of noise to the previous state. The noise schedule, which dictates the volume of noise added at every step, is vital for the model's success. The essence of diffusion models is found in the reverse process, where the model learns to undo the noisy transformation, gradually removing noise until it recovers the original sample or creates a new, similar sample. This is accomplished by training a neural network to estimate the noise that was added during the forward process. Through iterative denoising, the model can produce remarkably realistic and diverse samples.

1.INTRODUCTION

The diffusion model is an important framework that helps explain how innovations, behaviors, and information spread within a population. This model clarifies the ways in which new ideas are embraced and highlights the different rates of acceptance among various individuals or groups. In today's interconnected world, where technology speeds up the sharing of information, it is essential to examine the details of diffusion for both academic study and practical use. The model suggests that the adoption process is affected by several factors, including the features of the innovation, the communication channels used, and the socio-economic environment of potential adopters. By investigating these aspects, researchers can understand the patterns and rates of diffusion, which can help in creating more effective strategies to promote the broad acceptance of beneficial advancements. This investigation not only enriches theoretical discussions but also has important consequences in fields such as marketing and public health.[1]

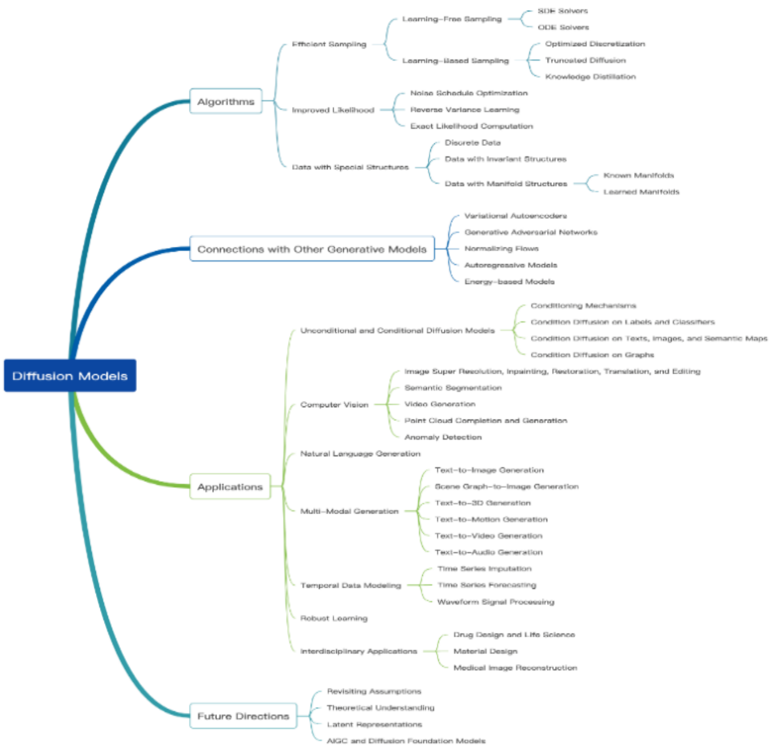

1.Definition and Significance of the Diffusion Model

The diffusion model serves as a statistical framework that describes how innovations, information, or behaviors spread through populations over time. Essentially, the model offers a structured approach to understanding the adoption process, with significant implications across various fields like marketing, sociology, and epidemiology. By utilizing mechanisms such as the replication of influence among individuals and the effect of network structures, the diffusion model helps researchers forecast trends and pinpoint key adopters within a population. Its significance notably extends to new technologies and products, allowing companies to develop strategies that enhance reach and adoption rates. This is especially relevant in areas like social media and e-commerce, where comprehending user behavior can greatly improve engagement.
strategies. Furthermore, the incorporation of advanced techniques, such as Laplacian mixture models, underscores the robustness of the models, enabling the analysis of complex data structures and improving dimensionality reduction to uncover underlying patterns.[2],[3],[4].

1.1Theoretical Framework of the Diffusion Model

The theoretical framework of the diffusion model provides a fundamental structure for comprehending how substances disseminate through different media. Key to this framework is the differentiation between local and nonlocal diffusion processes. Local diffusion follows classical principles, while nonlocal diffusion involves interactions across a wider spatial range. Recent innovations have introduced advanced coupling methods that effectively merge these two approaches, addressing interfacial inconsistencies and guaranteeing flux balance, as noted in recent research ((Du et al.)). Additionally, empirical data, exemplified by observations in thermotropic liquid crystals, indicates how temperature changes can profoundly affect the translational diffusion coefficient. This finding illustrates a complex relationship between orientational and translational order, challenging conventional views ((Bagchi et al.)). By creating a solid theoretical foundation that integrates these various elements, researchers can enhance the predictive capabilities and relevance of diffusion models in diverse fields.[5],[6].

1.2 Fundamental principles of diffusion models

- **Forward Diffusion Process:** In the forward process, noise is incrementally added to the input data (e.g., images, text, or other modalities), progressively transforming it into pure random noise. This process is modeled as a Markov chain, where each step depends only on the previous state.
- **Reverse Diffusion Process:** The reverse process reconstructs the original data by progressively denoising the random noise. This involves learning the probability distribution of the data at each noise level and performing iterative sampling in the reverse direction.
- **Noise Schedule:** The noise schedule determines the rate at which noise is added or removed during the forward and reverse processes.
- **Variational Inference:** This is statistical method that approximate the posterior distribution of latent variables, in this case, the original data given the noisy observations.
- **Score-Based Diffusion Models:** A variant of diffusion models that directly learns the gradient of the data distribution, often referred as the" scorefunction".Diffusion models have achieved phenomenal successin synthesizing High -images,text and all other data moalities with un-precedented control and diversity.[7]
- **Latent Variable Framework:** Diffusion models often utilize a latent variable approach, with variational inference methods employed to approximate the posterior distribution of the original data given noisy observations.
- **Iterative Refinement:** Generation in diffusion models is iterative, allowing fine-grained control over the output. This enables the production of high-quality, diverse results with excellent control over specific attributes.
- **Applications in Generative AI:** These principles have enabled diffusion models to excel in generating realistic images, videos, and audio. Their flexibility and robustness make them particularly effective for creative tasks and scientific research.[8]

1.3 A CATEGORIZATION OF DIFF US ION MODELS:

1.Based on the Nature of the Diffusion Process:

- **Denoising Diffusion Probabilistic Models (DDPMs):**
  - **Forward Diffusion (Noising Process):** This is a Markov process where Gaussian noise is progressively added to the data over a series of timesteps (T). Each timestep adds a small amount of noise, gradually transforming the original data into pure noise. This process is defined by a variance schedule (βt), which determines the amount of noise added at each timestep.
  - **Reverse Diffusion (Denoising Process):** This is the core of the model. It learns to reverse the noising process, starting from pure noise and iteratively removing noise to generate a data sample. This is done by training a neural network to predict the noise added at each timestep.
  - **Training with Variational Inference:** DDPMs are trained using variational inference, which involves maximizing the evidence lower bound (ELBO) on the log-likelihood of the data. This involves comparing the predicted noise to the actual noise added during the forward process.
  - **Key Characteristics:** Discrete timesteps, fixed noise schedule, training by noise prediction.
- **Score-Based Generative Models:**
  - **Score Function:** The score function of a probability distribution is the gradient of its log-density. It points in the direction of increasing probability.
  - **Score Matching:** This technique is used to train a neural network to estimate the score function of the data distribution. It involves minimizing the difference between the estimated score and the true score.
  - **Noise Conditional Score Networks (NCSNs):** These models learn the score function at different noise levels. This allows for more efficient sampling by starting with a high noise level and gradually reducing it.

- **Sampling with Langevin Dynamics:** Samples are generated by iteratively following the estimated score function, using a process called Langevin dynamics. This involves taking small steps in the direction of the score function, with added noise to avoid getting stuck in local optima.
- **Key Characteristics:** Focus on estimating the score function, sampling with Langevin dynamics, can be continuous-time.

- **Stochastic Differential Equations (SDEs):**
  - **Continuous-Time Formulation:** This approach formulates the diffusion process as a continuous-time SDE, which describes the evolution of a system over time under the influence of random noise.
  - **General Framework:** SDEs provide a more general framework that encompasses both DDPMs and score-based models as special cases.
  - **Flexibility:** This allows for more flexibility in defining the diffusion process and designing sampling methods.

2. Based on the Type of Data:

- **Image Generation:**
  - **High-Resolution Image Synthesis:** Diffusion models have achieved state-of-the-art results in generating high-resolution images with remarkable quality and diversity.
  - **Image Editing and Manipulation:** These models can also be used for image editing tasks, such as inpainting, super-resolution, and image manipulation.
- **Audio Synthesis:**
  - **Speech Synthesis:** Diffusion models can generate natural-sounding speech with high fidelity.
  - **Music Generation:** These models can create musical pieces with complex structures and harmonies.
- **Molecular Generation:**
  - **Drug Discovery:** Diffusion models can generate novel molecules with desired properties, which can be used for drug discovery and development.
  - **Materials Science:** These models can also be used to design new materials with specific characteristics.
- **Other Data Types:**
  - **Text Generation:** While less common than for images, diffusion models are being explored for text generation tasks.
  - **Point Clouds:** These models can generate 3D point clouds, which have applications in computer graphics and robotics.
  - **Graphs:** Diffusion models can also be used to generate graphs, which can be used to model relationships between objects.

3. Based on the Model Architecture:

- **U-Net:**
  - **Encoder-Decoder Structure:** The U-Net architecture has an encoder that downsamples the input and a decoder that upsamples it, with skip connections between corresponding layers.
  - **Capturing Local and Global Information:** This architecture is well-suited for diffusion models because it can capture both local and global information in the data, which is important for generating high-quality samples.
- **Transformers:**
  - **Attention Mechanism:** Transformers rely on the attention mechanism, which allows the model to focus on relevant parts of the input when making predictions.
  - **Sequence Modeling:** Transformers are particularly well-suited for sequence modeling tasks, such as text generation and audio synthesis.

4. Based on the Training Objective:

- **Variational Inference:**
  - **Evidence Lower Bound (ELBO):** This is a lower bound on the log-likelihood of the data, which is maximized during training.
  - **Tractable Objective:** The ELBO provides a tractable training objective that can be optimized using gradient-based methods.
- **Score Matching:**
  - **Minimizing Score Difference:** This technique involves minimizing the difference between the estimated score and the true score.
  - **Various Score Matching Techniques:** There are various score matching techniques, such as denoising score matching and sliced score matching.

5. Based on Conditional Generation:

- **Conditional Diffusion Models:**
  - **Controlling the Generation Process:** These models can generate data conditioned on certain inputs, such as class labels, text descriptions, or other modalities.

- **Applications:** This allows for more controlled generation of specific types of data, such as generating images of a specific object or generating text that matches a given style.[9],[10],[11]

2.DIFFUSION MODELS FOR DATA WITH SPECIAL STRUCTURES:

Diffusion models have demonstrated remarkable success in domains like images and audio. However, their application to other modalities often requires adjustments due to unique structural challenges inherent to various data types. For instance, issues arise when score functions are defined only on continuous domains or when data exist on low-dimensional manifolds. Adapting diffusion models to these scenarios is essential for effective functionality.

2.1 Discrete Data

Most diffusion models are geared towards continuous data domains, because Gaussian noise perturbation as used in DDPMs is not a natural fit for discrete data, and the score functions required by SGMs and Score SDEs are only defined on continuous data domains.to generate discrete data of high dimensions. Specifically, VQ-Diffusion replaces Gaussian noise with a random walk on the discrete data space, or a random masking operation. The resulting transition kernel for the forward process takes the form of $(x_t \mid x_{t-1}) = v^\top(x_t)Q_t v(x_{t-1})$ (39) where v(x) is a one-hot column vector, and Q_t is the transition kernel of a lazy random walk. D3PM accommodates discrete data in diffusion models by constructing the forward noising process with absorbing state kernels or discretized Gaussian kernels. Campbell et al. (2022) present the first continuous-time framework for discrete diffusion models. Leveraging Continuous Time Markov Chains, they are able to derive efficient samplers that outperform discrete counterparts, while providing a theoretical analysis on the error between the sample distribution and the true data distribution. Concrete Score Matching (CSM) proposes a generalization of the score function for discrete random variables. Concrete score is defined by the rate of change of the probabilities with respect to directional changes of the input, which can be seen as a finite-difference approximation to the continuous (Stein) score. The concrete score can be efficiently trained and applied to MCMC. Based on the theory of stochastic calculus, Liu et al. (2023) proposes a framework for diffusion models to generate data on constrained and structured domains, including discrete data as a special case. Using a fundamental theorem in stochastic calculus, the Doob's h-transform, one can constrain the data distribution on a specific area by including a special force term in the reverse diffusion process. They use a parameterization of the force term with an EM-based optimization algorithm. Furthermore, the loss function can be transformed to $L2$ loss using Girsanov theorem.

2.2 Data with Invariant Structures

Data with invariant structures in diffusion models refers to datasets that exhibit certain symmetries or invariances, such as translational, rotational, or scaling invariance. These properties are essential in applications like image generation, molecular modeling, and physics simulations, where maintaining such invariances ensures the generated data is consistent with real-world constraints.[12]

1. Invariant Structures Defined

Invariant structures are features or relationships within data that remain unchanged under specific transformations. For example:

- **Translational Invariance:** Shifting an image doesn't change its core structure.
- **Rotational Invariance:** Rotating an object retains its fundamental characteristics.
- **Scaling Invariance:** Resizing an object preserves its proportional relationships.
- **Permutation Invariance:** The ordering of elements in a set doesn't affect its overall properties (e.g., atoms in a molecule).

2. Why Invariant Structures Matter

a. Real-world Alignment

Many natural processes exhibit invariance. For instance:

- Molecular structures should be invariant to rotation and translation in 3D space.
- Physical simulations require invariance under specific symmetry transformations.

b. Model Efficiency

Incorporating invariances reduces the space of possible solutions the model has to learn, making training more efficient and the model more robust.

c. Consistency

Ensuring invariant structures prevents the generation of unrealistic outputs that violate physical or geometric rules.

3. Invariant Structures in Diffusion Models

a. Diffusion Models Overview

Diffusion models are generative models that learn to transform noise into data by iteratively denoising. They are particularly suited for generating high-quality, complex data distributions.

b. Challenges with Invariant Structures

In diffusion models, naively learning invariance can be challenging:

- Data augmentation (e.g., rotating or flipping images) helps but isn't always sufficient.
- Explicitly encoding invariance in the model architecture is often necessary.

c. Approaches to Handle Invariances

1. **Group Equivariant Networks**
   - Use networks like Convolutional Neural Networks (CNNs) designed to respect symmetries (e.g., rotation-equivariant CNNs).
2. **Position Encodings**
   - Use embeddings that respect invariances, such as spherical harmonics for 3D rotational invariance.
3. **Symmetry-aware Architectures**
   - Modify diffusion models to operate on structures like graphs or tensors, ensuring invariance under specific transformations.
4. **Score Matching with Invariance**
   - Adjust score-matching objectives to incorporate invariance properties directly.

4. Applications

a. Molecular and Protein Modeling

- Ensuring invariance to rotations and translations helps in predicting molecular conformations or drug interactions.

b. Physics Simulations

- Modeling invariant physical systems (e.g., fluid dynamics, particle interactions).

c. Image and 3D Object Generation

- Generating images or 3D objects that respect symmetries, like designing symmetric tools or artworks.

5. Mathematical Treatment

Invariant structures can be formalized using group theory:

- Let GGG be a symmetry group (e.g., rotations).
- A function f(x)f(x)f(x) is invariant under GGG if f(g·x)=f(x)f(g \cdot x) = f(x)f(g·x)=f(x) for all g∈Gg \in Gg∈G.

Diffusion models incorporate this by:

- Designing equivariant layers ϕ\phiϕ such that ϕ(g·x)=g·ϕ(x)\phi(g \cdot x) = g \cdot \phi(x)ϕ(g·x)=g·ϕ(x).
- Ensuring the denoising process respects these symmetries at each step.

6. Open Challenges

1. Balancing computational cost with invariance complexity.
2. Extending invariances to complex or less understood transformations.
3. Generalizing models across datasets with varying invariance properties.[13],[14]

3.CONNECTIONS WITH OTHER GENERATIVE MODELS:

Diffusion models, while powerful in their own right, don't exist in isolation. They have interesting connections and relationships with other generative models, both older and contemporary. Understanding these connections provides valuable insights into their strengths and weaknesses. Here's a detailed look:[15]

1. Variational Autoencoders (VAEs):

- Shared Idea: Latent Variables: Both VAEs and diffusion models rely on the concept of latent variables. VAEs learn a compressed representation (latent code) of the data, while diffusion models gradually transform data into a simple latent distribution (like Gaussian noise).
- Key Difference: How Latent Space is Used: VAEs directly map data to a latent space and learn to decode from it. Diffusion models, on the other hand, define a gradual diffusion process in the data space itself, effectively "exploring" the data manifold through a series of small steps.
- Connection: Training Objective: DDPMs are trained using a variational lower bound (ELBO), which is also a core component of VAE training. This connection highlights the shared probabilistic foundation.
- Latent Diffusion Models (LDMs): These models explicitly combine VAEs and diffusion models. An autoencoder (like a VAE) is used to learn a lower-dimensional latent space, and the diffusion process is then applied in this latent space. This can significantly improve efficiency, as the diffusion process operates in a smaller space.

2. Generative Adversarial Networks (GANs)[16]:

- Different Training Paradigms: GANs use a two-player game between a generator and a discriminator, while diffusion models are trained to reverse a noise process.
- Strengths and Weaknesses: GANs can generate very sharp and realistic samples but are notoriously difficult to train (prone to instability and mode collapse). Diffusion models are more stable to train and tend to produce higher quality samples, but they are generally slower at sampling.
- Connection: Score Matching: Some connections exist through score matching techniques, which can be used to train both diffusion models and certain types of GANs.
- Recent Trends: There's growing interest in combining aspects of GANs and diffusion models, trying to leverage the strengths of both.

3. Autoregressive Models:

- Different Generation Processes: Autoregressive models generate data sequentially, one element at a time, conditioned on the previously generated elements. Diffusion models generate data by iteratively denoising noise.
- Connection: Modeling Conditional Distributions: Both types of models need to model conditional distributions. In autoregressive models, it's the distribution of the next element given the previous ones. In diffusion models, it's the distribution of the data at the previous timestep given the data at the current timestep.
- Strengths and Weaknesses: Autoregressive models are well-suited for sequential data like text but can struggle with high-dimensional data like images. Diffusion models excel at high-dimensional data but are less naturally suited for sequential data (though there's ongoing research in this area).

4. Energy-Based Models (EBMs):

- Shared Goal: Modeling Probability Distributions: Both diffusion models and EBMs aim to learn the underlying probability distribution of the data.
- Different Approaches: EBMs define an energy function that assigns a scalar value to each data point, and the probability distribution is defined based on this energy function. Diffusion models, as we've discussed, use a diffusion process.
- Connection: Score Matching: Score matching, which is used to train score-based diffusion models, can also be used to train EBMs.

5. Normalizing Flows[17]:

- Different Transformation Strategies: Normalizing flows transform a simple probability distribution (like Gaussian) into a complex data distribution through a series of invertible transformations. Diffusion models gradually transform data into noise and then reverse this process.
- Connection: Change of Variables: Both techniques rely on the change of variables formula to compute probabilities.
  [18],[19]

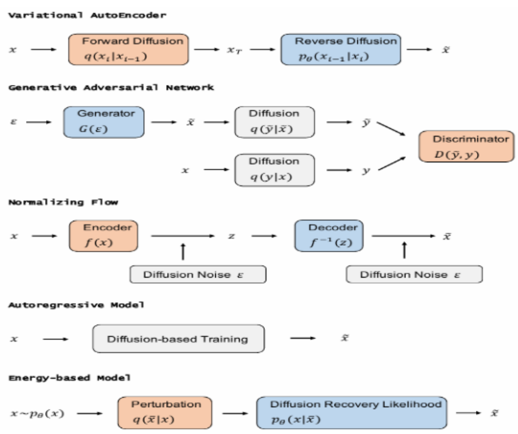

4. APPLICATIONS OF DIFFUSION MODELS:

Diffusion models have emerged as a powerful class of generative models, capable of producing high-quality, realistic data across various domains. Their iterative denoising process enables them to model complex distributions effectively. Below, we explore the key applications of diffusion models in detail:[21],[22],[23]

1. Image Generation and Editing

a. High-Resolution Image Synthesis

Diffusion models are extensively used for generating realistic images from noise:

- **Applications:** Creating artwork, digital illustrations, and photorealistic content.
- **Examples:** OpenAI's DALL-E, Google's Imagen, and Stability AI's Stable Diffusion.

b. Inpainting and Outpainting

- **Image Inpainting:** Filling in missing or corrupted parts of an image.
- **Outpainting:** Extending an image beyond its original boundaries while maintaining style and content consistency.
- **Use Cases:** Photo restoration, content creation, and visual effects in media.

c. Style Transfer and Manipulation

Diffusion models can generate images in different artistic styles or transform one image into another while preserving key features:

- **Use Cases:** Designing advertisements, game assets, or social media content.

2. Text-to-Image Generation

Diffusion models have been integrated with natural language processing to create images based on textual descriptions.

- **Applications:**
  - Visualizing concepts from descriptions.
  - Generating unique visual content for books, educational material, or marketing.
  - Accessibility tools for visually impaired users.
- **Example:** OpenAI's DALL-E combines GPT models with diffusion processes for text-to-image generation.

3. 3D Object and Scene Generation

Diffusion models are being adapted for 3D content creation.

- **Applications:**
  - Generating 3D models for games, movies, or simulations.
  - Virtual reality (VR) and augmented reality (AR) applications.
  - Digital twin creation for industrial use cases.
- **Advances:** Techniques like NeRF (Neural Radiance Fields) can be extended with diffusion models for photorealistic 3D renderings.

4. Audio and Speech Synthesis

Diffusion models have been applied to audio data, leveraging their ability to process temporal structures.

- **Applications:**
  - **Text-to-speech (TTS):** Generating realistic voices for virtual assistants and audiobooks.
  - **Music Generation:** Composing music in specific styles or genres.
  - **Noise Reduction:** Enhancing audio quality by removing background noise.
- **Example:** WaveGrad and DiffWave are diffusion-based models for speech synthesis.

5. Molecular and Protein Modeling[20],[21]

In scientific domains, diffusion models are used for simulating and designing molecules and proteins.

- **Applications:**
  - **Drug Discovery:** Predicting molecular structures and their binding affinities for drug development.
  - **Protein Folding:** Modeling protein structures from sequence data.
  - **Material Design:** Generating new materials with specific properties.
- **Examples:**
  - **Score-based Generative Models:** Used to predict 3D molecular structures.
  - **Applications in Chemistry:** Diffusion models can generate stable molecular conformations by respecting physical invariances (e.g., rotational symmetry).

6. Video Generation[22],[23],24],[25]

Extending diffusion models to temporal sequences enables realistic video generation.

- **Applications:**
  - **Video Synthesis:** Creating short videos from noise or textual prompts.
  - **Video Inpainting:** Restoring missing frames in videos.
  - **Animation Assistance:** Generating intermediate frames for smooth transitions in animation.
- **Challenges:** Capturing temporal consistency across frames while maintaining spatial quality.

7. Physics Simulations

Diffusion models are applied to simulate physical systems where modeling complex distributions is essential.

- **Applications:**
  - **Fluid Dynamics:** Simulating realistic water, air, or smoke flow.
  - **Particle Interactions:** Predicting the behavior of particles in systems like plasma physics.
  - **Weather Forecasting:** Enhancing predictions by modeling chaotic systems.
- **Example:** Modeling turbulence or diffusion of pollutants in the atmosphere.

8. Medical Imaging

Diffusion models contribute to advancements in healthcare through imaging applications.

- **Applications:**
  - **MRI Reconstruction:** Accelerating Magnetic Resonance Imaging (MRI) scans by reconstructing high-quality images from undersampled data.
  - **CT Scan Denoising:** Improving image clarity while reducing radiation doses.
  - **Synthetic Data Generation:** Generating synthetic medical data for training machine learning models without privacy concerns.

- **Examples:** Improving image segmentation, enhancing diagnostic accuracy, and reducing costs.

9. Generative Design

Diffusion models aid in the creation of optimized designs for engineering and architecture.

- **Applications:**
  - **Product Design:** Generating prototypes with specific functional and aesthetic requirements.
  - **Architecture:** Creating layouts and designs for buildings or urban planning.
  - **Fashion Design:** Developing new clothing patterns or accessories.

10. Robotics and Control

Diffusion models are being explored for decision-making and control in robotics.

- **Applications:**
  - **Trajectory Planning:** Predicting smooth and efficient paths for robots.
  - **Motion Imitation:** Generating realistic movements for humanoid robots or animated characters.
  - **Reinforcement Learning:** Modeling environment dynamics for training agents.

11. Anomaly Detection and Data Imputation

Diffusion models are effective in identifying unusual patterns in data and filling in missing information.

- **Applications:**
  - **Fraud Detection:** Identifying irregularities in financial transactions.
  - **Industrial Monitoring:** Detecting faults in manufacturing systems.
  - **Time-Series Analysis:** Imputing missing values in sequential data like stock prices or sensor readings.

12. Cryptography and Privacy

Diffusion models can contribute to secure and privacy-preserving technologies.

- **Applications:**
  - **Synthetic Data for Privacy:** Generating realistic yet synthetic data for machine learning, ensuring privacy.
  - **Encryption Schemes:** Exploring their use in designing new cryptographic techniques.

Challenges and Future Directions[26],[27]

1. **Scalability:** Extending diffusion models to handle larger datasets and higher resolutions efficiently.
2. **Domain Adaptation:** Generalizing models across different domains or types of data.
3. **Interpretability:** Improving understanding of the learned representations and generated outputs.
4. **Optimization:** Reducing computational requirements while maintaining performance.

CONCLUSION

Diffusion models are the new class of generative models that are redefining the field by producing top-of-the-line data that could be used for applications ranging from image synthesis to molecular modeling. This iterative denoising method assures stable training along with accurate representation of complex data distributions. Thus far, diffusion models are hindered by issues of excessive computational cost and also the slow speed of inference. However, the problem is being solved by using latent diffusion models. Because of their increasing footprint and potential, diffusion models are indeed going to provide a different dimension or horizon for future AI intervention, giving it a place in more and more diversified areas.

REFERENCES:

[1] Rogers, E. M. (2003). Diffusion of Innovations (5th ed.). arXiv,“arXiv.orge-Printarchive.”2021.[Online].Available: https://arxiv.org/

[2] Bass, F. M. (1969). A new product growth model for consumer durables. *Management Science, 15*(5), 215-227. Intel,“Intelhigh-levelsynthesiscompiler.”2021.[Online].Available:https://www.intel.com/content/www/us/en/software/programmable/quartus-prime/hls-compiler.html

[3] Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology, 83*(6), 1420-1443.Intel, “Intel FPGASDK for OpenCL software technology.”2020.[Online].Available:https://www.intel.com/content/www/us/en/software/programmable/sdk-for opencl/overview.html

[4] Coifman, R. R., & Lafon, S. (2006). Diffusion maps. *Applied and Computational Harmonic Analysis, 21*(1), 5-30.J. Ragan-Kelley, C. Barnes, A. Adams, S. Paris, F. Durand, and S. Amarasinghe, “Halide: A language and compiler for optimizing parallelism, locality, and recomputation in image processing pipelines,” SIGPLAN Not., vol. 48, no. 6, pp. 519–530, Jun. 2013.

[5] Samanta, S., & Bagchi, B. (2009). Translational diffusion of particles in nematic liquid crystals: A computer simulation study. *The Journal of Chemical Physics, 131*(6), 064303.Y.-H. Lai et al., “HeteroCL: A multi-paradigm programming infrastructure for software-defined reconfigurable computing,” in Proc. ACM/SIGDA Int. Symp. Field-Programmable Gate Arrays, 2019, pp. 242–251.

[6] Metzler, R., & Klafter, J. (2000). The random walk's guide to anomalous diffusion: A fractional dynamics approach. *Physics Reports, 339*(1), 1-77.SYCL overview. 2020. [Online]. Available: https://www.khronos. org/sycl/

[7] Ho, J., Jain, A., & Abbeel, P. (2020). Denoising Diffusion Probabilistic Models. *Advances in Neural Information Processing Systems (NeurIPS 2020)*.Falcon computing, “Merlin compiler.” 2021. [Online]. Available:

https://www. falconcomputing. com/ merlin-fpga-compiler/

[8] Song, Y., & Ermon, S. (2019). Generative Modeling by Estimating Gradients of the Data Distribution. *Advances in Neural Information Processing Systems (NeurIPS)*. Clash, “Clash: A modern, functional, hardware description language.” 2021. [Online]. Available: https://clash-lang.org/

[9] Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, *33*, 6840-6851.

[10] Hyvärinen, A. (2005). Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, *6*(Apr), 695-705.

[11] Hyvärinen, A. (2005). Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, *6*(Apr), 695-705.Y. Chi, J. Cong, P. Wei, and P. Zhou, “SODA: Stencil with optimized dataflow architecture,” in Proc. IEEE/ACM Int. Conf. Comput.-Aided Des., 2018, pp. 1–8.

[12] SchNet (Schütt et al., 2017) or SE(3)-Transformers (Fuchs et al., 2020).S. Grauer-Gray, L. Xu, R. Searles, S. Ayalasomayajula, and J. Cavazos, “Auto-tuning a high-level language targeted to GPU codes,” in Proc. Innovative Parallel Comput., 2012, pp. 1–10.

[13] Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, *33*, 6840-6851.

[14] Cohen, T. S., & Welling, M. (2016). Group equivariant convolutional networks. *International conference on machine learning*, 2990-2999.PYNQ-Z1: Python productivity for Zynq-7000 ARM/FPGA SoC.
2021. [Online]. Available: https://store.digilentinc.com/pynq-z1python-productivity-for-zynq-7000-arm-fpga-soc/

[15] Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, *33*, 6840-6851.Xilinx Alveo boards. 2021. [Online]. Available: https://www. xilinx.com/products/boards-and-kits/alveo.html

[16] Hyvärinen, A. (2005). Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, *6*(Apr), 695-705.

[17] Papamakarios, G., Coullon, T., & Murray, I. (2019). Review of normalizing flows. *arXiv preprint arXiv:1908.09257*.

[18] Rezende, D. J., & Mohamed, S. (2015). Variational inference with normalizing flows. *International conference on machine learning*, 1530-1538.

[19] Radford, A., Kim, J. W., Xu, C., Griffith, G., Salimans, T., Sutskever, I., & Amodei, D. (2021).

[20] Saharia, C., Chan, W., Saxena, S., Li, J., Whang, J., Denton, E. L., ... & Fleet, D. J. (2022).

[21] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022).

[22] Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, *33*, 6840-6851

[23] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022).

[24] Saharia, C., Chan, W., Saxena, S., Li, J., Whang, J., Denton, E. L., ... & Fleet, D. J. (2022).

[25] Gafni, O., Polyak, A., Ashual, O., Sheynin, S., Parikh, D., & Taigman, Y. (2022).

[26] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022).

[27] Lu, C., Zhou, Y., Li, F., Li, X., & Li, J. (2022).