

COMMUN: SEMANTIC IMAGE COMMUNICATIONS AS AN INVERSE PROBLEM WITH INN-GUIDED DIFFUSION MODELS

Jiakang Chen[†], Di You[†], Deniz Gündüz, Pier Luigi Dragotti

Department of Electrical and Electronic Engineering, Imperial College London, UK

ABSTRACT

Joint source-channel coding schemes based on deep neural networks (DeepJSCC) have recently achieved remarkable performance for wireless image transmission. However, these methods usually focus only on the distortion of the reconstructed signal at the receiver side with respect to the source at the transmitter side, rather than the perceptual quality of the reconstruction which carries more semantic information. As a result, severe perceptual distortion can be introduced under extreme conditions such as low bandwidth and low signal-to-noise ratio. In this work, we propose CommIN, which views the recovery of high-quality source images from degraded reconstructions as an inverse problem. To address this, CommIN combines Invertible Neural Networks (INN) with diffusion models, aiming for superior perceptual quality. Through experiments, we show that our CommIN significantly improves the perceptual quality compared to DeepJSCC under extreme conditions and outperforms other inverse problem approaches used in DeepJSCC.

Index Terms— Semantic communications, joint source-channel coding, inverse problems, invertible neural networks, diffusion models

1. INTRODUCTION

Shannon’s separation theorem [1] forms the cornerstone of modern communication systems. In a typical point-to-point transmission, the transmitter that performs the encoding process usually consists of two steps, the source encoder that removes redundant information from the source to achieve compression, and the channel encoder that adds redundancy to the compressed information to correct for errors caused by noisy communication channels. However, the theoretical optimality of designing the source encoder and channel encoder separately can only be guaranteed if there exists an infinite code length, which is not achievable in practice [2], in particular under extreme bandwidth and delay constraints. In such cases, a Joint Source-Channel Coding (JSCC) approach can potentially lead to better performance by integrating source compression and channel encoding into a single step.




Fig. 1: Visual comparison examples of CelebA-HQ images under complex AWGN channel at SNR = -5dB and bandwidth compression ratio $\rho = 0.0013$. We observe that in this very challenging communication setting, the signal reconstructed by DeepJSCC is highly distorted. While InverseJSCC improves the perceptual quality of the reconstructed images by utilizing the prior knowledge of a pre-trained StyleGAN-2 generator, our approach produces even higher perceptual quality reconstructions.

Recently JSCC approaches based on neural networks called DeepJSCC, have demonstrated excellent performance through powerful learning capabilities [3][4][5][6]. However, traditional JSCC methods focus on pixel-level distortion or structural similarity rather than perceptual distortion, which can be measured by learned perceptual image patch similarity (LPIPS) metric [7]. In [8], the authors discuss the use of LPIPS directly as a loss function during training.

Generative models based on neural networks have demonstrated a strong ability to generate high perceptual quality data. Examples include variational autoencoder (VAE), generative adversarial network (GAN) and the cutting-edge diffusion models. There are several works that attempt to leverage the power of these generative models into JSCC. For example, [9] implements a JSCC scheme over binary channels via VAE. In [10], adversarial training is employed to achieve secure JSCC in the presence of an eavesdropper. In [11], GAN-style loss is employed to train the legitimate decoder. A diffusion denoising model is employed at the decoder in [12] as part of a hybrid transmission framework. The papers [13] and [14] perform noise reduction by placing the diffusion model on the noisy channel symbols of JSCC to improve the reconstruction

[†]Equal contribution




Fig. 2: Our key insight is to simulate the degradation process of the communication pipeline (top) with INN (bottom), which can transform the input original signal $\mathbf{x}$ into the simulated degraded observation $\hat{\mathbf{x}}$ and estimate the lost details $\mathbf{d}$ by the degradation.

quality. However, the above approaches do not directly apply the state-of-the-art generative model structure to JSCC. Moreover, they have not fully utilised the strong prior of pretrained GANs such as StyleGAN-2, which has potential to improve the perceptual quality of the reconstruction significantly.

In [15], a pretrained generative model is employed to solve a classical inverse problem in compressed sensing. Outstanding performance in single-image super-resolution is achieved in [16] through the utilization of generative models, highlighting its efficacy in addressing practical inverse problems. The similarity between communication systems and inverse problems suggests the potential of applying a powerful generative model to JSCC, as represented by works such as [17][10][18]. Among them, [18] proposes InverseJSCC, which considers recovering high quality source images with better perceptual quality from degraded reconstructions as an inverse problem. The inverse problem is solved by a pre-trained StyleGAN-2 generator and combined with the ILO [19] method, which uses the state-of-the-art GAN prior to achieve a never-before-attained reconstructed perceptual quality under very extreme channel conditions.

In this paper, inspired by [18], we propose CommIN, which treats the JSCC problem as an inverse problem and, by leveraging results from our recent work [20], we reconstruct the source using Invertible Neural Networks (INN) and generative diffusion models. The key insight here is that the INN splits the signal (typically an image) into a coarse version $\mathbf{c}$ and details $\mathbf{d}$. In [20], we train the forward part of the INN to mimic the degradation process in an inverse problem so that $\mathbf{c}$ is very close to the observed degraded image $\mathbf{y}$. In this paper we treat the communication process composed of the transmission, corruption with channel noise and reconstruction as the degradation process and mimic it using an INN (see Fig. 2). We treat the reconstructed image $\hat{\mathbf{x}}$ after transmission over a noisy channel using DeepJSCC as the degraded image in an inverse problem and apply variations of the approach in [20] using diffusion models to further improve the reconstruction.

Our overall proposed approach achieves state-of-the-art performance and has the advantage of being computationally efficient. Our primary achievements include: CommIN stands as the first work employing INNs to simulate the degradation process in DeepJSCC. It is also the first initiative to leverage the state-of-the-art generative prior of diffusion models to enhance the perceptual quality of DeepJSCC's reconstructions, especially in extreme conditions.

The paper is organized as follows: In Section 2, we briefly introduce the communication scenario and DeepJSCC. In Section 3, we then explain how to treat DeepJSCC as an inverse problem, and introduce the INN-Guided Diffusion model to solve this inverse problem. In Section 4, we present numerical results showing the superior performance of CommIN. Finally, Section 5 concludes the paper.

2. SYSTEM MODEL

We consider the point-to-point wireless transmission of images over a noisy channel (see Fig. 2). As in [3], we model the transmitter and receiver as two neural networks. Specifically, the transmitter is the encoder and the receiver is the decoder of an auto-encoder system. The transmitter maps the source signal, a real vector $\mathbf{x} \in \mathbb{R}^m$ to the complex channel input vector $\mathbf{z} \in \mathbb{C}^k$ using a parameterised encoding function $\mathbf{z} = f(\mathbf{x}; \theta)$. Here m represents the source bandwidth and k the channel bandwidth. We define the bandwidth compression ratio (BCR) as $\rho = k/m$. This ratio reflects the compression applied to the signal by the communication system. In the image transmission problem, $m = H \times W \times C$, where H , W and C are the height, width and colour channel of the input image. The channel input vector needs to satisfy the following average power constraint:

$$\frac{1}{k} \mathbb{E}_{\mathbf{z}} [\|\mathbf{z}\|_2^2] \leq \bar{P}. \quad (1)$$

This is achieved by normalising the signal $\tilde{\mathbf{z}}$, output of the last layer of the encoder, as follows:

$$\mathbf{z} = \sqrt{k\bar{P}} \frac{\tilde{\mathbf{z}}}{\sqrt{\tilde{\mathbf{z}}^H \tilde{\mathbf{z}}}}, \quad (2)$$

where H is the Hermitian transpose. Then the encoded signal $\mathbf{z}$ is sent to a noisy channel, which introduces random corruptions to the transmitted symbols. We consider a complex AWGN channel so that the received signal can be expressed as follows:

$$\hat{\mathbf{z}} = \eta(\mathbf{z}, \sigma^2) = \mathbf{z} + \mathbf{n}_C, \quad (3)$$

where $\mathbf{n}_C$ is sampled in an independent identically distributed (i.i.d.) way from a complex Gaussian distribution with variance σ^2 : $\mathbf{n}_C \sim \mathcal{CN}(0, \sigma^2 I_{k \times k})$. The receiver uses a parameterised decoding function $\hat{\mathbf{x}} = g(\hat{\mathbf{z}}; \phi)$ to get the reconstruction $\hat{\mathbf{x}}$ from $\hat{\mathbf{z}}$. To represent the condition of the noisy channel, we




Fig. 3: Details of the framework of INN in our approach.

define channel SNR as follows:

$$\text{Channel SNR} = 10 \log_{10} \frac{\bar{P}}{\sigma^2} \text{ dB}. \quad (4)$$

The encoder and decoder are jointly optimised by minimising:

$$\arg \min_{\theta, \phi} \mathbb{E}_{\mathbf{x} \sim p_{\mathbf{x}}} \mathbb{E}_{\hat{\mathbf{x}} \sim p_{\hat{\mathbf{x}}|\mathbf{x}}} [d(\mathbf{x}, \hat{\mathbf{x}})], \quad (5)$$

where d can be any distortion measure.

We use the same DeepJSCC architecture as in [18] but reduce the number of convolution channels from 512 to 256. The resulting DeepJSCC is optimised in an end-to-end fashion according to the Mean Squared Error (MSE) loss function:

$$d(\mathbf{x}, \hat{\mathbf{x}}) = \frac{1}{N} \sum_{i=1}^N \|\mathbf{x}_i - \hat{\mathbf{x}}_i\|_2^2. \quad (6)$$

3. COMMUN APPROACH

3.1. DeepJSCC as an Inverse Problem

Given the DeepJSCC encoder/decoder pair, we now aim to further improve the perceptual quality of the reconstructed image by treating it as the result of the degradation process in an inverse problem. Specifically, we interpret the received signal $\hat{\mathbf{x}}$ as the degraded measurements $\mathbf{y} = A(\mathbf{x})$ in a standard inverse problem, where A is the forward operator and corresponds, in this case, to the complete communication chain of DeepJSCC including the encoder, noisy channel and decoder (see Fig.2). More precisely, we can express the forward operator as

$$\mathbf{y} = A(\mathbf{x}, \eta) = g_{\phi}(\eta(f_{\theta}(\mathbf{x}), \sigma^2)) = \hat{\mathbf{x}}, \quad (7)$$

which is a non-linear process with AWGN η .

With the forward model in place and given measurements $\mathbf{y}$, one typically seeks to restore the original signal $\mathbf{x}$ by simultaneously imposing a fidelity constraint and a proper prior on the expected $\mathbf{x}$. The fidelity constraint is typically the MSE between the output of A and the measurements $\mathbf{y}$:

$$\text{MSE}(A(\mathbf{x}), \mathbf{y}) = \|A(\mathbf{x}) - \mathbf{y}\|_2^2. \quad (8)$$

The prior can come in many forms and here we aim to use diffusion models as prior.

Under this setting, we are faced with two main challenges. First the degradation process is highly non-linear and cannot

be expressed in closed-form. Secondly, imposing a diffusion model as prior in an inverse problem is often computationally demanding. We address these challenges by resorting to INNs. As shown in Fig. 2, our basic idea is to simulate the degradation process by training an INN, which can decompose the original signal $\mathbf{x}$ into two parts: the estimated degraded observation $\hat{\mathbf{x}}' (\approx \hat{\mathbf{x}})$ and lost details $\mathbf{d}$ during the degradation. Due to the invertibility of INN, the inverse transform of INN has the property of perfectly recovering the input original image $\mathbf{x}$ from $\hat{\mathbf{x}}'$ and $\mathbf{d}$ by construction. During the inference stage, the pre-trained INN is utilized to guide the diffusion process, ensuring that it simultaneously integrates the observed measurements $\hat{\mathbf{x}}$ and preserves the intricate details produced by the diffusion process.

In the following subsections, we first introduce the design of the INN in our algorithm, and then we describe how it works with the diffusion process to reconstruct the source signal with high perceptual quality.

3.2. Exploring the Degradation Process with INN

The framework of our INN is shown in Fig. 3. The original image is split into two parts by a splitting operator. Then the Prediction Network (PNet) conditioned on the coarse part aims to predict the detail part, while the Update Network (UNet) conditioned on the detail part is used to adjust the coarse part to make it smoother. The PNet and UNet are applied alternatively to generate the coarse and detail parts, $\mathbf{c}$ and $\mathbf{d}$, respectively. The PNet and UNet can be any complex non-linear functions and their properties will not affect the invertibility of INN.


To model the degradation process, we impose that $\mathbf{c}$ resembles $\mathbf{y}$. Given a training set $\{\mathbf{x}^i, \mathbf{y}^i\}_{i=1}^N$, which contains N high-quality images and their low-quality counterparts recovered by DeepJSCC, we optimize our INN with the following loss function:

$$L(\Theta) = \frac{1}{N} \sum_{i=1}^N \|\mathbf{c}_i - \mathbf{y}_i\|_2^2, \quad (9)$$

where Θ denotes the learnable parameters of our INN. Once we constrain one part of the output of $\text{INN}_{\text{Forward}}(\mathbf{x})$ to be close to $\mathbf{y}$, due to invertibility, the other part of the output $\mathbf{d}$ will inevitably represent the detailed information lost during the degradation process.

3.3. INN-Guided Probabilistic Diffusion Algorithm

We follow the diffusion model defined in denoising diffusion probabilistic models (DDPM [21]). As for other generative models, DDPM learns a distribution of images given a training set. Its inference process works by sampling a random noise vector $\mathbf{x}_T$ and gradually denoising it until it reaches a high-quality output image $\mathbf{x}_0$. Specifically, the reverse diffusion process iteratively samples $\mathbf{x}_{t-1}$ from $p(\mathbf{x}_{t-1}|\mathbf{x}_t)$ to yield clean images $\mathbf{x}_0 \sim q(\mathbf{x})$ from initial random noise $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. The following equations illustrate this process:



(a) PSNR versus SNR (higher better) (b) LPIPS versus SNR (lower better)

Fig. 4: Performance comparison in terms of PSNR and LPIPS.

$$\mathbf{x}_{0,t} = \frac{1}{\sqrt{\alpha_t}}(\mathbf{x}_t - \sqrt{1 - \alpha_t}\epsilon_\theta(\mathbf{x}_t, t)), \quad (10)$$

and

$$\mathbf{x}_{t-1} = \frac{\sqrt{\alpha_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t}\mathbf{x}_t + \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1 - \bar{\alpha}_t}\mathbf{x}_{0,t} + \sigma_t\mathbf{z}, \quad (11)$$

where $\epsilon_\theta(\mathbf{x}_t, t)$ is the pre-trained denoising model and $\mathbf{x}_{0,t}$ is the estimation of $\mathbf{x}_0$ from the noisy image $\mathbf{x}_t$.

To solve inverse problems with the above generative model, we need to refine each unconditional transition using $\mathbf{y}$ to ensure data consistency. In our proposed algorithm, we impose our data-consistency step by optimizing the clean image $\mathbf{x}_{0,t}$ instead of the noisy image $\mathbf{x}_t$. As shown in Algorithm 1, we impose an additional data consistency step (in blue) after each original unconditional sampling update. At this additional step, we decompose the intermediate result $\mathbf{x}_{0,t}$ with $INN_{Forward}$ into coarse $\mathbf{c}_t$ and detail part $\mathbf{d}_t$ and replace the coarse part $\mathbf{c}_t$ with the output of the DeepJSCC decoder, $\mathbf{y}$. The INN-optimized $\hat{\mathbf{x}}_{0,t}$ is generated by the inverse process $INN_{Inverse}$. Thus, the INN-optimized $\hat{\mathbf{x}}_{0,t}$ is composed of the coarse information $\mathbf{y}$ and the details generated by the diffusion process. To incorporate the INN-optimized $\hat{\mathbf{x}}_{0,t}$ into the DDPM algorithm, we update $\mathbf{x}_t$ with the guidance of the gradient of $\|\hat{\mathbf{x}}_{0,t} - \mathbf{x}_{0,t}\|_2^2$. With the help of INN, our algorithm effectively estimates the details lost in the degradation of the communication process and avoids being limited by the requirement of knowing the closed-form expression of the degradation model.

Algorithm 1 DDPM Sampling with INN

```

1:  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
2: for  $t = T, \dots, 1$  do
3:    $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  if  $t > 1$ , else  $\mathbf{z} = \mathbf{0}$ 
4:    $\mathbf{x}_{0,t} = \frac{1}{\sqrt{\alpha_t}}(\mathbf{x}_t - \sqrt{1 - \alpha_t}\epsilon_\theta(\mathbf{x}_t, t))$ 
5:    $\tilde{\mathbf{x}}_{t-1} = \frac{\sqrt{\alpha_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t}\mathbf{x}_t + \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1 - \bar{\alpha}_t}\mathbf{x}_{0,t} + \sigma_t\mathbf{z}$ 
6:    $\mathbf{c}_t, \mathbf{d}_t = INN_{Forward}(\mathbf{x}_{0,t})$ 
7:    $\hat{\mathbf{x}}_{0,t} = INN_{Inverse}(\mathbf{y}, \mathbf{d}_t)$ 
8:    $\mathbf{x}_{t-1} = \tilde{\mathbf{x}}_{t-1} - \zeta \nabla_{\mathbf{x}_t} \|\hat{\mathbf{x}}_{0,t} - \mathbf{x}_{0,t}\|_2^2$ 
9: end for
10: return  $\mathbf{x}_0$ 

```



Fig. 5: Reconstructed CelebA-HQ images by DeepJSCC, InverseJSCC and our approach for $\rho=0.0013$ and $\text{SNR}=1\text{dB}$.

4. EXPERIMENT RESULTS

To evaluate the performance of our proposed approach, we compare CommIN with state-of-the-art JSCC methods: DeepJSCC [6] and InverseJSCC [18]. We use CelebA-HQ train datasets for training. For fair comparisons, we use the same pre-trained encoder and decoder for all three approaches. All our experiments are implemented in PyTorch with Adam optimiser. In training our DeepJSCC, our channel SNR is uniformly sampled from $[-5, 5]$ dB. We train separate INNs for each degradation setting in the inference stage. We assume perfect channel knowledge at the encoder and the decoder. For validation, we choose CelebA-HQ 1K test datasets with image sizes 256×256 , and we experiment on channel SNRs of $\{-5, -3, -1, 1, 3, 5\}$ dB and BCR $\rho = 0.0013$, which correspond to highly challenging communication scenarios. Empirically, we set $T = 1000$ and set step size ζ to 0.3, 0.4, and 0.5 for $\text{SNR} = \{-5, -3\}$ dB, $\{-1, 1\}$ dB, and $\{3, 5\}$ dB respectively. We analyse the performance of the proposed scheme using both the widely used pixel-wise metric PSNR and the perceptual metric LPIPS [7].

Fig. 4 shows our results in terms of PSNR and LPIPS with respect to channel SNR on the CelebA-HQ dataset. For all the settings, our approach achieves the lowest LPIPS, while maintaining the similar PSNR with DeepJSCC. As shown in Fig. 5, our algorithm produces high-quality reconstruction results and preserves more details than other methods.

5. CONCLUSION

In this paper, we have introduced CommIN, a new JSCC scheme for wireless image transmission which treats the reconstruction of received corrupted images as an inverse problem. Key features of the approach include the use of INN to model the degradation introduced by the channel and DeepJSCC. Moreover, INNs facilitate the integration of diffusion models in the reconstruction process. CommIN achieves image reconstructions with perceptual quality beyond DeepJSCC in extreme environments with extremely low BCR with SNR. It also achieves competitive results compared with the existing inverse problem approach to JSCC in [18].

References

- [1] C. E. Shannon, “A mathematical theory of communication,” *The Bell system technical journal*, vol. 27, no. 3, pp. 379–423, 1948.
- [2] V. Kostina and S. Verdú, “Lossy joint source-channel coding in the finite blocklength regime,” *IEEE Transactions on Information Theory*, vol. 59, no. 5, pp. 2545–2575, 2013.
- [3] E. Boursoulatz, D. B. Kurka, and D. Gündüz, “Deep joint source-channel coding for wireless image transmission,” *IEEE Transactions on Cognitive Communications and Networking*, vol. 5, no. 3, pp. 567–579, 2019.
- [4] D. B. Kurka and D. Gündüz, “Deep joint source-channel coding of images with feedback,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5235–5239, IEEE, 2020.
- [5] T.-Y. Tung, D. B. Kurka, M. Jankowski, and D. Gündüz, “Deepjscc-q: Constellation constrained deep joint source-channel coding,” *IEEE Journal on Selected Areas in Information Theory*, 2022.
- [6] T.-Y. Tung and D. Gündüz, “Deepwive: Deep-learning-aided wireless video transmission,” *IEEE Journal on Selected Areas in Communications*, vol. 40, no. 9, pp. 2570–2583, 2022.
- [7] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, “The unreasonable effectiveness of deep features as a perceptual metric,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 586–595, 2018.
- [8] J. Xu, T.-Y. Tung, B. Ai, W. Chen, Y. Sun, and D. Gündüz, “Deep joint source-channel coding for semantic communications,” *to appear in IEEE Communications Magazine*, 2023.
- [9] K. Choi, K. Tatwawadi, A. Grover, T. Weissman, and S. Ermon, “Neural joint source-channel coding,” in *International Conference on Machine Learning*, pp. 1182–1192, 2019.
- [10] T. Marchioro, N. Laurenti, and D. Gündüz, “Adversarial networks for secure wireless communications,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8748–8752, IEEE, 2020.
- [11] M. Yang, C. Bian, and H.-S. Kim, “OFDM-guided deep joint source channel coding for wireless multipath fading channels,” *IEEE Transactions on Cognitive Communications and Networking*, vol. 8, no. 2, pp. 584–599, 2022.
- [12] X. Niu, X. Wang, D. Gündüz, B. Bai, W. Chen, and G. Zhou, “A hybrid wireless image transmission scheme with diffusion,” *arXiv preprint arXiv:2308.08244*, 2023.
- [13] T. Wu, Z. Chen, D. He, L. Qian, Y. Xu, M. Tao, and W. Zhang, “Channel denoising diffusion models for wireless communications,” *arXiv preprint arXiv:2305.09161*, 2023.
- [14] B. Xu, R. Meng, Y. Chen, X. Xu, C. Dong, and H. Sun, “Latent semantic diffusion-based channel adaptive denoising semcom for future 6G systems,” *arXiv preprint arXiv:2304.09420*, 2023.
- [15] A. Bora, A. Jalal, E. Price, and A. G. Dimakis, “Compressed sensing using generative models,” in *International conference on machine learning*, pp. 537–546, 2017.
- [16] S. Menon, A. Damian, S. Hu, N. Ravi, and C. Rudin, “Pulse: Self-supervised photo upsampling via latent space exploration of generative models,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2437–2445, 2020.
- [17] E. Erdemir, P. L. Dragotti, and D. Gündüz, “Privacy-aware communication over a wiretap channel with generative networks,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2989–2993, IEEE, 2022.
- [18] E. Erdemir, T.-Y. Tung, P. L. Dragotti, and D. Gündüz, “Generative joint source-channel coding for semantic image transmission,” *IEEE Journal on Selected Areas in Communications*, vol. 41, no. 8, pp. 2645–2657, 2023.
- [19] G. Daras, J. Dean, A. Jalal, and A. G. Dimakis, “Intermediate layer optimization for inverse problems using deep generative models,” in *International Conference on Machine Learning*, 2021.
- [20] D. You, A. Floros, and P. L. Dragotti, “Indigo: An inn-guided probabilistic diffusion algorithm for inverse problems,” *arXiv preprint arXiv:2306.02949*, 2023.
- [21] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 6840–6851, 2020.