



Diff-E: Diffusion-based Learning for Decoding Imagined Speech EEG

Soowon Kim¹, Young-Eun Lee², Seo-Hyun Lee², Seong-Whan Lee¹

¹Department of Artificial Intelligence, Korea University, Republic of Korea

²Department of Brain and Cognitive Engineering, Korea University, Republic of Korea

soowon-kim@korea.ac.kr, ye_lee@korea.ac.kr, seohyunlee@korea.ac.kr, sw.lee@korea.ac.kr

Abstract

Decoding EEG signals for imagined speech is a challenging task due to the high-dimensional nature of the data and low signal-to-noise ratio. In recent years, denoising diffusion probabilistic models (DDPMs) have emerged as promising approaches for representation learning in various domains. Our study proposes a novel method for decoding EEG signals for imagined speech using DDPMs and a conditional autoencoder named Diff-E. Results indicate that Diff-E significantly improves the accuracy of decoding EEG signals for imagined speech compared to traditional machine learning techniques and baseline models. Our findings suggest that DDPMs can be an effective tool for EEG signal decoding, with potential implications for the development of brain-computer interfaces that enable communication through imagined speech.

Index Terms: silent communication, speech recognition, electroencephalography, imagined speech, brain-computer interface

1. Introduction

Speech is a crucial mode of communication for humans, allowing us to convey complex ideas and thoughts to others through sound patterns [1, 2, 3]. It is a fundamental part of our social and cultural interactions, enabling us to build relationships, share information, and express emotions. Speech is a highly complex process that involves the coordinated activity of various brain regions and muscles, including the larynx, tongue, lips, and jaw. Unfortunately, several individuals cannot speak due to physical limitations such as locked-in syndrome, a condition where they are fully conscious but unable to move or communicate verbally [4, 5]. The inability to speak can significantly impact an individual's quality of life, as it limits their ability to interact with others and express themselves effectively. Therefore, developing effective methods for restoring communication in individuals with locked-in syndrome and other speech impairments is a critical area of research. In this context, our research explores decoding brain signals as a potential method for enabling silent communication [6, 7, 8].

Electroencephalography (EEG) is a non-invasive technique for measuring the brain's electrical activity. EEG signals are acquired by placing electrodes on the scalp, which detect the electrical activity generated by the brain's neurons. EEG signals have been widely used in neuroscience research to study brain function, communication, and medical applications such as diagnosing neurological disorders. In recent years, EEG decoding has gained increasing attention to translate EEG signals into meaningful information, such as speech or attention detection [9, 10, 11, 12]. EEG decoding for speech processing is a particularly challenging task, as it involves decoding complex and highly variable neural patterns associated with speech




Figure 1: The figure illustrates the forward diffusion process applied to a section of the recorded EEG signal captured from Broca's area, FT7, of Subject 1.

production and perception [13]. Moreover, EEG signals are often contaminated with noise and artifacts, which can affect the accuracy of the decoding process [14]. Therefore, developing accurate and robust methods for decoding EEG signals associated with imagined speech is an active area of research, with potential applications in speech rehabilitation, human-computer interaction, and other related fields.

Denoising diffusion probabilistic models (DDPMs) have become a powerful tool for identifying sophisticated patterns within complex data, especially when dealing with high-dimensional data [15, 16, 17]. For DDPMs, the forward diffusion process introduces noise to the original data, transforming it into a corrupted version. This process is modeled by adding Gaussian noise at each time step, effectively simulating a random walk in the latent space. The aim of the denoising diffusion model is then to reverse this process, i.e., to denoise the corrupted data and recover the original signal. In recent years, DDPMs have gained attention as a promising approach for time-series data processing and analysis as well, including audio and video [18, 19]. In the context of EEG decoding, DDPMs may offer a promising approach for capturing the complex relationships between neural activity and language processing during imagined speech tasks. This can enable accurate decoding of EEG signals associated with imagined speech, which has potential applications in various fields such as speech rehabilitation and neuroscience.

DDPMs have also been employed in representation learning. Preechakul et al. [20] demonstrated that diffusion autoencoders can yield a more interpretable and decodable latent data representation than conventional autoencoders. Zhang et al. [21] displayed the effectiveness of pre-trained diffusion models using high-dimensional data representations. Their approach involves a trainable model that predicts a mean shift based on the encoded representation. It is trained to bridge as much of the gap as possible, a gap that emerges during the forward process of DDPMs. Consequently, the encoder is compelled to extract

as much information as it can from images to assist with this bridging process. In this study, building upon prior research, we present a novel approach to decode EEG signals using DDPMs and a Conditional Autoencoder (CAE). The CAE aids in learning meaningful features potentially lost during the forward process in DDPMs. Furthermore, we've integrated a jointly-trained classifier to improve decoding performance. To the best of our knowledge, this proposed approach represents the first attempt to utilize diffusion models to decode EEG signals associated with imagined speech.

2. Related Works

Decoding EEG signals using deep learning approaches is a challenging problem due to various factors, including the scarcity of data, poor signal-to-noise ratio, and high inter- and intra-individual variability [22]. Despite these challenges, several studies have explored different EEG decoding techniques for various applications, including speech decoding.

Schirrneister et al. [23] developed DeepConvNets for end-to-end learning of EEG signals in human subjects. They utilized machine learning techniques such as batch normalization and exponential linear units and a cropped training strategy to improve decoding performance, achieving accuracy levels similar to filter bank common spatial patterns algorithms. Another study by Lawhern et al. [24] proposed EEGNet, a compact CNN architecture specifically designed for EEG classification tasks. EEGNet uses depthwise and separable convolutions, which can extract features tailored to the unique characteristics of EEG control signals. The model has been tested on four BCI paradigms and shown to accurately classify EEG signals while being compact in terms of the number of parameters in the model.

Lee et al. [25] investigated the underlying characteristics affecting the decoding performance of two emerging BCI paradigms: imagined speech and visual imagery. The authors analyzed EEG signals filtered in six frequency ranges for 13 class classifications, including a rest class. They identified cortical regions relevant to both paradigms using semantic levels of word properties according to their concepts. The study showed high accuracy rates across all classes and multi-class scalability in both paradigms, providing crucial information for improving the efficiency of BCI applications.

Diffusion-based methods for time-series data have been extensively researched recently. One such method proposed in [26] is based on a structured state space model that incorporates a diffusion process to capture the underlying dynamics of the data. This method is able to impute and forecast missing values in time series data by utilizing the estimated diffusion process. The authors demonstrate the effectiveness of their approach on various real-world datasets and show that it outperforms existing methods for time-series imputation and forecasting. Jeong et al. [27] propose a method for improving the quality of synthetic speech generated by TTS systems. The proposed method utilizes DDPMs to remove noise from the intermediate representations of speech generated by TTS systems. By reducing noise in the intermediate representations, the proposed method is able to produce clearer and more natural-sounding speech. The authors evaluate the effectiveness of their approach on several real-world datasets and demonstrate that it outperforms existing methods for improving the quality of synthetic speech.

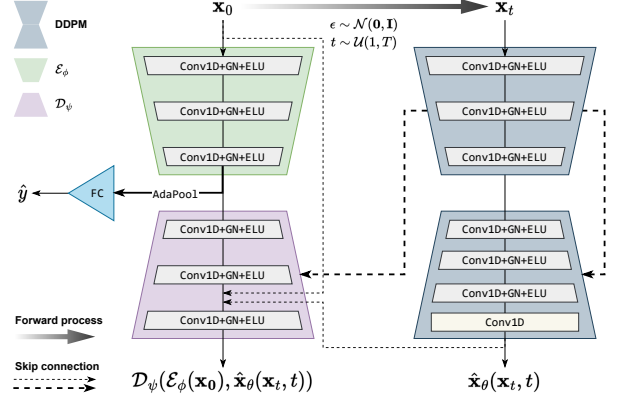


Figure 2: A flowchart of Diff-E for EEG signal decoding. The nodes labeled as Conv1D, GN, and Linear represent one-dimensional convolutional layer, group normalization, and linear layer, respectively. The node labeled as ELU represents a nonlinear activation function. Adaptive average pooling layer, AdaPool is used to make input for the classifier, FC.

3. Materials and Methods

3.1. Denoising Diffusion Models

DDPMs are a type of machine learning model that can learn complex probability distributions over data. The "forward process" in DDPMs is determined by a fixed Markov chain that progressively introduces Gaussian noise to the data. The forward process in DDPMs starts with a probability distribution called $q(\mathbf{x}_0)$, which represents the uncorrupted original data. This distribution is then transformed iteratively using a sequence of Markov diffusion kernels called $q(\mathbf{x}_t|\mathbf{x}_{t-1})$, which are Gaussian with a fixed variance schedule $\{\beta_t\}_{t=1}^T$. The process can be expressed by:

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I}) \quad (1)$$

$$q(\mathbf{x}_{1:T}|\mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t|\mathbf{x}_{t-1}) \quad (2)$$

Diffusion process with Gaussian noise can corrupt the original data with arbitrary stage, t with the notation $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$:

$$q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t}\mathbf{x}_0, (1 - \bar{\alpha}_t)\mathbf{I}) \quad (3)$$

Ho et al. [15] proposed a method to train a model that takes a noisy sample \mathbf{x}_t and predicts its noise by training a network $\epsilon\theta(\mathbf{x}_t, t)$. In contrast, our study trains Diff-E to predict the uncorrupted original signal, \mathbf{x}_0 , instead of predicting the injected noise:

$$\mathcal{L}_{DDPM}(\theta) = \|\mathbf{x}_0 - \hat{\mathbf{x}}_\theta(\mathbf{x}_t, t)\| \quad (4)$$

Here, θ are the DDPM parameters, and $\hat{\mathbf{x}}_\theta(\mathbf{x}_t, t)$ is the prediction of the DDPM. We randomly sample timestep t from a uniform distribution, $\mathcal{U}(1, T)$. The objective function encourages the model to denoise the noisy input and produce output that is close to the original input signal. We have adopted a time-conditional UNet architecture [28] as described in Fig. 2, similar to the model used in the original paper [15], with modifications to adapt it for EEG data.

3.2. Conditional Autoencoder

Similar to the work [21], the CAE aims to derive meaningful representations of EEG signals by leveraging the errors that occur during the reconstruction process by the DDPM. The forward pass of the DDPM causes information loss, which the CAE attempts to compensate for by identifying and rectifying these errors. This enables the CAE to generate more accurate representations of the original EEG signals. Specifically, we employ the objective function for the CAE as follows:

$$\mathcal{L}_{\text{CAE}}(\psi, \phi) = \|\mathcal{L}_{\text{DDPM}}(\theta) - \mathcal{D}_{\psi}(\mathcal{E}_{\phi}(\mathbf{x}_0), \hat{\mathbf{x}}_{\theta}(\mathbf{x}_t, t))\|, \quad (5)$$

The CAE is composed of an encoder network, denoted by \mathcal{E}_{ϕ} , and a decoder network, denoted by \mathcal{D}_{ψ} , which is similar in structure to the DDPM. However, \mathcal{D}_{ψ} receives skip connections not from the outputs of \mathcal{E}_{ϕ} , but from the DDPM layers. This enables \mathcal{D}_{ψ} to be implicitly conditioned on the corruption stage of the forward process of the DDPM, as shown in Fig. 2 as *dashed arrows*. Furthermore, to improve the reconstruction of $\mathcal{L}_{\text{DDPM}}$, the original signal, \mathbf{x}_0 , and the output of the DDPM, $\hat{\mathbf{x}}_{\theta}(\mathbf{x}_t, t)$, are also skip-connected as inputs to the last layer of \mathcal{D}_{ψ} as depicted in Fig. 2 as *thin dashed arrows*.

3.3. Classifier

After \mathcal{E}_{ϕ} processes the data, the output is transformed into a compact representation, \mathbf{z} , by collapsing the time dimension into a single dimension. This is achieved by applying an adaptive average pooling layer, which effectively creates a one-dimensional latent vector. The compressed representation is then fed into the linear classifier, \mathcal{C}_{ρ} , which is jointly trained with the CAE to further separate the representations of each class and perform the classification task. The dimension of \mathbf{z} is constant at 256 throughout the experiment. To include the classification loss in the objective function of the CAE, we modified it as follows to become the overall Diff-E objective:

$$\mathcal{L}_{\text{Diff-E}}(\psi, \phi, \rho) = \|\mathcal{L}_{\text{DDPM}}(\theta) - \mathcal{D}_{\psi}(\mathcal{E}_{\phi}(\mathbf{x}_0), \hat{\mathbf{x}}_{\theta}(\mathbf{x}_t, t))\| + \alpha \|\hat{y} - y\|_2 \quad (6)$$

where $\hat{y} = \mathcal{C}_{\rho}(\mathbf{z})$. Here, ρ is adjustable parameters for \mathcal{C}_{ρ} , and y is the true label of the input signal. The hyperparameter α governs the relative importance of the reconstruction loss and the classification loss. 0.1 is chosen for the value of α throughout the experiment. During the inference, only \mathcal{E}_{ϕ} and \mathcal{C}_{ρ} are used to classify the signals. Specifically, the predicted label is obtained as $\hat{y} = \mathcal{E}_{\phi}(\mathcal{C}_{\rho}(\mathbf{x}_0))$.

In order to evaluate the effectiveness of Diff-E, it is compared to other existing methods that have demonstrated their applicability to decoding EEG signals in various paradigms, such as motor imagery and event-related potentials. This comparison is conducted to assess the performance of Diff-E and to determine its suitability for imagined speech EEG signal decoding applications.

3.4. Model Implementation Details

The DDPM and CAE architectures in our study were designed with each layer comprising a one-dimensional convolution (CONV1D), group normalization (GN), and activation function (ELU). In \mathcal{E}_{ϕ} , adaptive average pooling (AdaPool) is applied to transform the output into a compressed representation, \mathbf{z} . The total number of trainable parameters for both DDPM and

CAE was approximately 300,000, while the linear classifier had 400,000 trainable parameters.

We trained the models using the RMSProp optimizer with a fixed random seed of 42 and a cyclic learning rate scheduler developed by Smith [29]. The base learning rate was set to $9 * 10^{-5}$, and the maximum learning rate was set to $1.5 * 10^{-3}$. The training process was performed over a total of 500 epochs. During training, we used L1 loss for both the DDPM and CAE components of the model. We also used the mean squared error loss against one-hot encoded labels for the linear classifier classification task. To evaluate the performance of the model, we set aside 20% of the data for testing with a fixed random seed. Our implementation is available at <https://github.com/yorgoon/DiffE/>.

3.5. Dataset

3.5.1. Data Description

The data utilized in this study was obtained from a prior study conducted by Lee et al. [25]. The study involved recruiting twenty-two healthy participants, of whom fifteen were male, with a mean age of 24.68 ± 2.15 . None of the participants had a history of neurological disease or language disorder, and they did not have any hearing or visual impairments. Additionally, they refrained from taking drugs for twelve hours before the session. All participants had received high-quality English education for more than 15 years. In the imagined speech task, twenty-two subjects were instructed to imagine saying twelve different words or sentences, such as “ambulance,” “clock,” “hello,” “help me,” “light,” “pain,” “stop,” “thank you,” “toilet,” “TV,” “water,” and “yes,” as well as a resting state, resulting in a total of thirteen classes that a communication board commonly used in hospitals for patients with paralysis or aphasia. These words are considered to be essential for patient communication. The researchers used a 64-channel EEG cap with active Ag/AgCl electrodes that followed the international 10-10 system to record the EEG signals. The FCz and FPz channels were set as the reference and ground electrodes, respectively. Brain Vision/Recorder software (BrainProduct GmbH, Germany) was used to collect the EEG signals, which were then operated using MATLAB 2018a software. The researchers ensured that the impedance of all electrodes was maintained below $10k\Omega$. The researchers randomly presented twenty-two blocks of twelve words and a rest class. Each of the twenty-two participants contributed one thousand three hundred samples, consisting of one hundred samples per category. The study was approved by the Korea University Institutional Review Board [KUIRB-2019-0143-01] and was conducted in accordance with the Declaration of Helsinki.

3.5.2. Preprocessing

In this study, we employed several preprocessing steps to ensure the quality and accuracy of the EEG data. The first step was to apply band-pass filtering to the signals between 0.5 and 125Hz, with additional notch filtering at 60 and 120Hz to eliminate power line interference. The data was then re-referenced using a common average reference method to reduce the impact of any noise present in the data. To eliminate ocular and muscular artifacts that can be caused by movement or sounds, we used automatic electrooculography and electromyography removal methods. After removing the artifacts, the EEG signals were selected in the high-gamma frequency band for training the model and data analysis. The dataset was then epoched into

Table 1: Accuracy and AUC scores for imagined speech classification

Subject	Accuracy (%)	AUC (%)
DeepConvNet	21.00 ± 4.27	69.55 ± 3.51
EEGNet	32.78 ± 4.98	78.07 ± 4.27
Lee et al.[25]	46.27 ± 6.79	79.97 ± 4.83
Diff-E	60.63 ± 8.03	90.39 ± 3.48

Table 2: Ablation study

Components	Accuracy (%)	AUC (%)
Diff-E	60.63 ± 8.03	90.39 ± 3.48
w/o DDPM	46.59 ± 8.35	86.54 ± 5.36
w/o DDPM & \mathcal{D}_ψ	47.01 ± 7.80	46.27 ± 6.79

2-second segments, with baseline correction applied 500ms before the task. All of the preprocessing steps were carried out using MATLAB-based tools, including the OpenBMI Toolbox [30, 31], BCI Toolbox [32], and EEGLAB [33].

4. Results and Discussion

4.1. Performance Comparison

In this study, we compared the performance of our proposed Diff-E method with three other methods, DeepConvNet [23], EEGNet [24], and Lee et al. [25], for EEG signal decoding. The results are presented in Table 1. Diff-E outperformed all other methods, including DeepConvNet, EEGNet, and Lee et al., in both accuracy and area under the curve (AUC). The average accuracy of Diff-E was 60.63%, with a standard deviation of 8.03%, and the average AUC was 90.39%, with a standard deviation of 3.48%. Notably, Diff-E achieved significantly higher accuracy and AUC than the other methods. The average accuracy of DeepConvNet, EEGNet, and Lee et al. were 21.00%, 32.78%, and 46.27%, respectively. The average AUC of DeepConvNet, EEGNet, and Lee et al. were 69.55%, 78.07%, and 79.97%, respectively. These findings highlight the superior performance of our proposed Diff-E method for EEG signal decoding.

Our study’s results were unexpected since commonly used, and well-known methods such as EEGNet and DeepConvNet for EEG decoding underperformed compared to a traditional machine-learning method, such as a common spatial pattern-support vector machine method used by Lee et al. [25]. These methods are frequently used for decoding EEG signals from various paradigms, including motor imagery and event-related potentials. Our findings emphasize the importance of carefully selecting the appropriate model architecture based on the task at hand, the EEG paradigms used, and other relevant factors.

4.2. Ablation Study

To evaluate the impact of each component of our proposed method, we performed an ablation study. The study involved removing the diffusion process and training the CAE’s encoder and classifier separately to assess the effects of learning from the DDPM’s errors and whether the autoencoder-style approach is advantageous for learning meaningful representations. We aimed to determine the significance of each component of Diff-E. During the training process, the CAE was trained only to reconstruct the original signals, and a linear classifier was used without the use of DDPM. Additionally, we conducted another

ablation study solely on the $(\mathcal{E}_\phi, \mathcal{C}_\rho)$ component for the classification task, which is essentially a CNN-based classifier.

As shown in Table 2, the results indicate that Diff-E outperformed the other models, and removing the DDPM or solely using $(\mathcal{E}_\phi, \mathcal{C}_\rho)$ led to a significant decrease in accuracy. These results indicate that without DDPM, CAE alone does not offer any advantage over traditional machine learning techniques.

4.3. Future Works

To enhance the effectiveness and robustness of our proposed method for EEG signal decoding, we intend to address several limitations in our future work. Our method was tested on a non-public dataset, as it had larger and more diverse classes than other public datasets, which usually have only a few classes [34]. Although our EEG dataset was relatively large, it may still be inadequate for large deep-learning models to optimize effectively. To improve the applicability of our method, we plan to collect larger datasets with more subjects and apply our method to various EEG paradigms, including motor imagery and event-related potentials. This will offer a more comprehensive evaluation of our method’s effectiveness and improve its robustness across different paradigms. Additionally, we plan to investigate methods for end-to-end learning directly from raw EEG signals, potentially reducing the need for extensive pre-processing and enhancing the efficiency of our method.

5. Conclusion

This work contributes to the growing body of research on deep-learning-based approaches to EEG signal decoding, specifically for the task of imagined speech. The use of DDPMs, a relatively new approach in this field, provides an innovative solution that has demonstrated superior performance compared to traditional methods such as DeepConvNet and EEGNet. The results of our study suggest that DDPMs have the potential to become a valuable tool in EEG signal processing and can be further explored for other applications in this domain. Our study also highlights the importance of selecting appropriate model architectures for different EEG paradigms. This consideration is crucial to achieving high accuracy and performance in EEG signal decoding, as demonstrated by the inferior performance of DeepConvNet and EEGNet in our study. The selection of an appropriate model architecture should take into account various factors, such as the characteristics of the EEG data and the task at hand. Overall, our study provides a promising approach to accurately decoding EEG signals related to imagined speech, which can have significant implications in the development of brain-computer interfaces for communication and assistive technology. This work also contributes to advancing the understanding and utilization of deep-learning-based approaches in EEG signal processing and decoding.

6. Acknowledgements

This work was supported by Institute for Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2021-0-02068, Artificial Intelligence Innovation Hub; No. 2019-0-00079, Artificial Intelligence Graduate School Program (Korea University); No. 2017-0-00451, Development of BCI based Brain and Cognitive Computing Technology for Recognizing User’s Intentions using Deep Learning).

7. References

- [1] P. Lieberman, *Uniquely human: The evolution of speech, thought, and selfless behavior*. Harvard University Press, 1991.
- [2] J. D. Peters, *Speaking into the air: A history of the idea of communication*. University of Chicago Press, 2012.
- [3] P. Lieberman, “The evolution of human speech: Its anatomical and neural bases,” *Current Anthropology*, vol. 48, no. 1, pp. 39–66, 2007.
- [4] S. Laureys, F. Pellas, P. Van Eeckhout, S. Ghorbel, C. Schnakers, F. Perrin, J. Berre, M.-E. Faymonville, K.-H. Pantke, F. Damas *et al.*, “The locked-in syndrome: what is it like to be conscious but paralyzed and voiceless?” *Progress in Brain Research*, vol. 150, pp. 495–611, 2005.
- [5] J. F. Houde and S. S. Nagarajan, “Speech production as state feedback control,” *Frontiers in Human Neuroscience*, vol. 5, p. 82, 2011.
- [6] U. Chaudhary, N. Birbaumer, and A. Ramos-Murguialday, “Brain–computer interfaces for communication and rehabilitation,” *Nature Reviews Neurology*, vol. 12, no. 9, pp. 513–525, 2016.
- [7] N. Birbaumer, A. R. Murguialday, and L. Cohen, “Brain–computer interface in paralysis,” *Current Opinion in Neurology*, vol. 21, no. 6, pp. 634–638, 2008.
- [8] U. Chaudhary, B. Xia, S. Silvoni, L. G. Cohen, and N. Birbaumer, “Brain–computer interface–based communication in the completely locked-in state,” *PLoS Biology*, vol. 15, no. 1, p. e1002593, 2017.
- [9] S. Cai, E. Su, L. Xie, and H. Li, “EEG-based auditory attention detection via frequency and channel neural attention,” *IEEE Transactions on Human-Machine Systems*, vol. 52, no. 2, pp. 256–266, 2021.
- [10] E. Su, S. Cai, L. Xie, H. Li, and T. Schultz, “STAnet: A spatiotemporal attention network for decoding auditory spatial attention from EEG,” *IEEE Transactions on Biomedical Engineering*, vol. 69, no. 7, pp. 2233–2242, 2022.
- [11] Y.-E. Lee, S.-H. Lee, S.-H. Kim, and S.-W. Lee, “Towards voice reconstruction from EEG during imagined speech,” *arXiv preprint arXiv:2301.07173*, 2023.
- [12] Y.-E. Lee and S.-H. Lee, “EEG-transformer: Self-attention from transformer architecture for decoding EEG of imagined speech,” *2022 10th International Winter Conference on Brain-Computer Interface*, pp. 1–4, 2022.
- [13] P. Li, S. Cai, E. Su, and L. Xie, “A biologically inspired attention network for EEG-based auditory attention detection,” *IEEE Signal Processing Letters*, vol. 29, pp. 284–288, 2021.
- [14] Y.-E. Lee, N.-S. Kwak, and S.-W. Lee, “A real-time movement artifact removal method for ambulatory brain–computer interfaces,” *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 28, no. 12, pp. 2660–2670, 2020.
- [15] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 6840–6851, 2020.
- [16] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, “Score-based generative modeling through stochastic differential equations,” *arXiv preprint arXiv:2011.13456*, 2020.
- [17] P. Dhariwal and A. Nichol, “Diffusion models beat gans on image synthesis,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 8780–8794, 2021.
- [18] Y. Tashiro, J. Song, Y. Song, and S. Ermon, “CSDI: Conditional score-based diffusion models for probabilistic time series imputation,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 24 804–24 816, 2021.
- [19] K. Rasul, C. Seward, I. Schuster, and R. Vollgraf, “Autoregressive denoising diffusion models for multivariate probabilistic time series forecasting,” *International Conference on Machine Learning*, pp. 8857–8868, 2021.
- [20] K. Preechakul, N. Chatthee, S. Wizadwongsa, and S. Suwanajakorn, “Diffusion autoencoders: Toward a meaningful and decodable representation,” *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10 619–10 629, 2022.
- [21] Z. Zhang, Z. Zhao, and Z. Lin, “Unsupervised representation learning from pre-trained diffusion probabilistic models,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 22 117–22 130, 2022.
- [22] Z. Tayeb, J. Fedjaev, N. Ghaboosi, C. Richter, L. Everding, X. Qu, Y. Wu, G. Cheng, and J. Conradt, “Validating deep neural networks for online decoding of motor imagery movements from EEG signals,” *Sensors*, vol. 19, no. 1, p. 210, 2019.
- [23] R. T. Schirrmester, J. T. Springenberg, L. D. J. Fiederer, M. Glasstetter, K. Eggenberger, M. Tangermann, F. Hutter, W. Burgard, and T. Ball, “Deep learning with convolutional neural networks for EEG decoding and visualization,” *Human Brain Mapping*, vol. 38, no. 11, pp. 5391–5420, 2017.
- [24] V. J. Lawhern, A. J. Solon, N. R. Waytowich, S. M. Gordon, C. P. Hung, and B. J. Lance, “EEGNet: a compact convolutional neural network for EEG-based brain–computer interfaces,” *Journal of Neural Engineering*, vol. 15, no. 5, p. 056013, 2018.
- [25] S.-H. Lee, M. Lee, and S.-W. Lee, “Neural decoding of imagined speech and visual imagery as intuitive paradigms for BCI communication,” *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 28, no. 12, pp. 2647–2659, 2020.
- [26] J. M. L. Alcaraz and N. Strodthoff, “Diffusion-based time series imputation and forecasting with structured state space models,” *arXiv preprint arXiv:2208.09399*, 2022.
- [27] M. Jeong, H. Kim, S. J. Cheon, B. J. Choi, and N. S. Kim, “Diff-tts: A denoising diffusion model for text-to-speech,” *arXiv preprint arXiv:2104.01409*, 2021.
- [28] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Medical Image Computing and Computer-Assisted Intervention 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*. Springer, 2015, pp. 234–241.
- [29] L. N. Smith, “Cyclical learning rates for training neural networks,” *2017 IEEE Winter Conference on Applications of Computer Vision*, pp. 464–472, 2017.
- [30] M.-H. Lee, O.-Y. Kwon, Y.-J. Kim, H.-K. Kim, Y.-E. Lee, J. Williamson, S. Fazli, and S.-W. Lee, “EEG dataset and OpenBMI toolbox for three BCI paradigms: an investigation into BCI illiteracy,” *GigaScience*, vol. 8, no. 5, p. giz002, 2019.
- [31] D.-Y. Lee, M. Lee, and S.-W. Lee, “Decoding imagined speech based on deep metric learning for intuitive BCI communication,” *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 29, pp. 1363–1374, 2021.
- [32] R. Krepki, B. Blankertz, G. Curio, and K.-R. Müller, “The berlin brain–computer interface (BBCI)–towards a new communication channel for online control in gaming applications,” *Multimedia Tools and Applications*, vol. 33, no. 1, pp. 73–90, 2007.
- [33] A. Delorme and S. Makeig, “EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis,” *Journal of Neuroscience Methods*, vol. 134, no. 1, pp. 9–21, 2004.
- [34] G. A. P. Coretto, I. E. Gareis, and H. L. Rufiner, “Open access database of eeg signals recorded during imagined speech,” *12th International Symposium on Medical Information Processing and Analysis*, vol. 10160, p. 1016002, 2017.