

Text2Data: Low-Resource Data Generation with Textual Control

Shiyu Wang, Yihao Feng, Tian Lan, Ning Yu, Yu Bai, Ran Xu, Huan Wang*, Caiming Xiong, Silvio Savarese

Salesforce AI Research
181 Lytton Ave, Palo Alto, CA 94301

Abstract

Natural language serves as a common and straightforward signal for humans to interact seamlessly with machines. Recognizing the importance of this interface, the machine learning community is investing considerable effort in generating data that is semantically coherent with textual instructions. While strides have been made in text-to-data generation spanning image editing, audio synthesis, video creation, and beyond, low-resource areas characterized by expensive annotations or complex data structures, such as molecules, motion dynamics, and time series, often lack textual labels. This deficiency impedes supervised learning, thereby constraining the application of advanced generative models for text-to-data tasks. In response to these challenges in the low-resource scenario, we propose Text2Data, a novel approach that utilizes unlabeled data to understand the underlying data distribution through an unsupervised diffusion model. Subsequently, it undergoes controllable finetuning via a novel constraint optimization-based learning objective that ensures controllability and effectively counteracts catastrophic forgetting. Comprehensive experiments demonstrate that Text2Data is able to achieve enhanced performance regarding controllability across various modalities, including molecules, motions and time series, when compared to existing baselines.

Code —

<https://github.com/SalesforceAIResearch/text2data>

Extended version — <https://arxiv.org/abs/2402.10941>

Introduction

Autonomy and controllability stand as twin primary pillars of generative AI (Gozalo-Brizuela and Garrido-Merchan 2023; Wang et al. 2022). While the challenge of autonomy has been substantially addressed through the rapid advancements of generative models, controllability is now ascending as a fervently explored arena within the machine learning community. As natural languages are one of the most common and simplest control signal for human beings to interact with machines, the machine learning community has increasingly focused on generating data that aligns semantically with textual descriptions, given its wide-ranging applications such as image editing (Zhang et al. 2023; Kawar

et al. 2023; Shen et al. 2024c), audio synthesis (Liu et al. 2023; Huang et al. 2023), video generation (Li et al. 2018; Hu, Luo, and Chen 2022), and many more (Tevet et al. 2023; Sanghi et al. 2022; Shen and Tang 2024).

Recent breakthroughs in text-to-data generative models, particularly those using diffusion (Li et al. 2023; Yang et al. 2023; Kumari et al. 2023), have demonstrated remarkable proficiency by harnessing the rich semantic insights from vast datasets of data-text pairs. Despite the broad application of text-to-data generative models, not all modalities can meet the substantial data-text pair requirements for achieving optimal controllability during model training. This is often due to costly annotations or intricate data structures, a scenario we refer to as the low-resource situation. The lack of text labels in certain areas, such as molecules (Ramakrishnan et al. 2014; Irwin et al. 2012), motions (Guo et al. 2020; Mahmood et al. 2019), and time series (Du et al. 2020), primarily restricts supervised learning and hinders the use of advanced generative models for text-to-data generation tasks. The low-resource situation when training generative models unsurprisingly results in issues like undesirable generation quality, model overfitting, bias, and lack of diversity. However, the optimization for scarce text representations to improve the alignment between generated data and input texts in generative models is still under-explored.

To mitigate the issues in the low-resource scenario, strategies such as data augmentation (Hedderich et al. 2020; Meng et al. 2021), semi-supervised learning (Thomas et al. 2013; Cheuk, Herremans, and Su 2021), and transfer learning (Tits, El Haddad, and Dutoit 2020; Yi et al. 2018) are utilized. Yet, each comes across challenges. Data augmentation, for example, cannot always replicate genuine data fidelity to align accurately with initial text descriptions, and potentially leads to overfitting due to over-reliance on augmented samples. It also exacerbates the training complexity, intensifying the already high computational demand of diffusion models. For semi-supervised learning, text inherently carries nuances, ambiguities, and multiple meanings. Ensuring that the model infers the correct interpretation when leveraging unlabeled data is not straightforward. Lastly, while transfer learning offers a solution for limited datasets, it is prone to catastrophic forgetting (Iman, Arabia, and Rasheed 2023), where previous knowledge dimin-

*Corresponding author: huan.wang@salesforce.com
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

ishes as new text descriptions are introduced.

Alternative to existing solutions, we propose Text2Data, a diffusion-based framework achieving enhanced text-to-data controllability even under low-resource situation. Specially, Text2Data operates in two pivotal stages: **(1) Distribution mastery by leveraging unlabeled data.** Distinct from conventional semi-supervised learning methods, Text2Data does not aim to deduce labels for unlabeled data. Instead, this step uses unlabeled data to discern the overarching data distribution via an unsupervised diffusion model, eliminating the semantic ambiguity often associated with semi-supervised approaches. **(2) Controllable finetuning on text-labeled data.** The learned diffusion model is then finetuned by text-labeled data. Distinct from methods reliant on data augmentation, Text2Data abstains from inflating the training dataset. Instead, we introduce a novel constraint optimization-based learning objective, aiming to mitigate catastrophic forgetting by regularizing the model parameter space closely to its preliminary space before finetuning. Our contributions are summarized as follows:

- We introduce Text2Data, a novel framework designed for text-to-data generation in the low-resource scenario. This approach maintains the fine-grained data distribution by fully harnessing both labeled and unlabeled data.
- We design a novel learning objective based on constraint optimization to achieve controllability and overcome catastrophic forgetting during finetuning.
- We theoretically validate optimization constraint selection and generalization bounds for our learning objective.
- We compile real-world datasets across three modalities and conduct comprehensive experiments to show the effectiveness of Text2Data. The results demonstrate that Text2Data achieves superior performance baselines regarding both generation quality and controllability.

Related Works

Text-to-Data Diffusion-Based Generation

Diffusion models, notably divided into classifier-guided (Dhariwal and Nichol 2021) and classifier-free (Ho and Salimans 2022) categories, have significantly impacted data generation across various domains. (Hoogeboom et al. 2022; Yang et al. 2023; Ho et al. 2022; Shen et al. 2024a). The classifier-guided diffusion guides the model during inference phase by independently training a classifier and supervising the model with its gradient, which is inefficient when computing gradient at each time step and sometimes the generation quality is deficient as the guidance is not involved in the training. By contrast, classifier-free diffusion guidance blends score estimates from both a conditional diffusion model and an unconditional one with time step as a parameter, exemplified by E(3) Equivariant Diffusion Model (EDM) (Hoogeboom et al. 2022) and Motion Diffusion Model (MDM) (Tevet et al. 2023) for controllable molecule and motion generation, respectively. Furthermore, since natural languages are a prevalent medium for human to communicate with the world, the text-to-data generation

paradigm has gained traction, with diffusion models being instrumental in generating high-quality data aligned with textual inputs. The extensive applications encompass text-to-image generation (Ruiz et al. 2023; Zhang and Agrawala 2023; Shen et al. 2024b), text-to-speech generation (Huang et al. 2022; Kim, Kim, and Yoon 2022), text-to-shape generation (Li et al. 2023; Lin et al. 2023), and more, leveraging the abundant text descriptions for training potent generative models. Despite advancements in generating data from text across various modalities, many other modalities may not satisfy the stringent requirements for sufficient data-text pairs essential for attaining optimal controllability during the training of models.

Low-Resource Learning

In response to the challenges of low-resource training for controllable generative models, several strategies have been formulated. For instance, Yin et al. (2023) proposes Text-to-Text-to-Image Data Augmentation that employed both large-scale pretrained Text-to-Text and Text-to-Image generative models for data augmentation to generate photo-realistic labeled images in a controllable manner. Zang and Wan (2019) utilizes a semi-supervised approach to augment the training of both the encoder and decoder with both labeled and unlabeled data. Tu et al. (2019) proposes to learn a mapping between source and target linguistic symbols and employs transfer learning to transfer knowledge from a high-resource language to low-resource language. Nevertheless, all those strategies have their own limitations, such as high computational complexity for data augmentation, difficulty in maintaining the correct interpretation of text when leveraging unlabeled data during semi-supervised learning, and the potential catastrophic forgetting issues in transfer learning. Additionally, most of works on text-to-data generation lie in the modality of image, speeches, and texts, which have plenty of labeled datasets (Ito and Johnson 2017; Lin et al. 2014; Jiang et al. 2021; Wang, Wu, and Pino 2020) to train the models nowadays. Yet it is far under-explored for the low-source modalities such as molecules, motions and time series. Therefore, we propose Text2Data, a diffusion-based framework adept at harnessing limited text-labeled data to enhance the controllability of the model in text-to-data generation.

Problem Formulation

Suppose the dataset $\mathcal{D} = \{\mathbf{x}, \mathbf{c}\}$ contains N independent samples in total, where $\mathbf{x} = \{\mathbf{x}_i\}_{i=1}^N$ is the data samples such as molecules, motions, time series, etc. We assume that there is only a proportion of data in \mathbf{x} that has corresponding text description $\mathbf{c} = \{\mathbf{c}_i\}_{i=1}^{N_p}$ where $N_p \leq N$. We denote that data with text description is contained in \mathcal{D}_p and $\mathcal{D}_p \subset \mathcal{D}$. Using both text-labeled and unlabeled data in \mathcal{D} , we aim to learn a generative model, $p_\theta(\mathbf{x}|\mathbf{c})$ parameterized by θ that is able to generate data $\mathbf{x} \sim p_\theta(\mathbf{x}|\mathbf{c})$ corresponding to specific text description $\mathbf{c} = \mathbf{c}^*$.

Methods

Controllable data generation seeks to learn the conditional data distribution $p_\theta(\mathbf{x}|\mathbf{c})$ during training and subsequently




Figure 1: Overview of Text2Data. The model leverages unlabeled data (i.e., blue module) to discern the overall data distribution while the optimal set of model parameters Θ is obtained. Then the model is finetuned on labeled data (i.e., red module) by constraint optimization that gives the optimal set of parameters as $\Theta \cap \Theta'$, where Θ' is the optimal set of parameters if finetune the model without constraint.

draw samples from this assimilated distribution during the inference stage. Consequently, our primary objective during the training phase is to optimize the following:

$$\min_{\theta} \mathbb{E}_{\mathbf{x}, \mathbf{c} \sim p_{\mathcal{D}_p}(\mathbf{x}, \mathbf{c})} [-\log p_{\theta}(\mathbf{x}|\mathbf{c})]. \quad (1)$$

While the training of such generative models is contingent upon the supervision of text descriptions present in the dataset (i.e., \mathcal{D}_p), it is not always feasible to obtain an adequate number of data-text pairs to ensure optimal controllability (i.e., $|\mathcal{D}_p| < |\mathcal{D}|$), especially in modalities like molecular structures, motion patterns and time series. Such constraints can precipitate complications, including model overfitting when optimizing according to Eq. 1. Given the challenges, devising strategies to capitalize on the unlabeled data within \mathcal{D} —which is often more accessible and cost-effective—is pivotal to effectively learn $p_{\theta}(\mathbf{x}|\mathbf{c})$.

Notably, the marginal distributions learned from unlabeled data closely resemble those obtained from labeled data:

$$p_{\theta}(\mathbf{x}) \approx \int p_{\theta}(\mathbf{x}|\mathbf{c})p_{\mathcal{D}_p}(\mathbf{c})d\mathbf{c} = \mathbb{E}_{\mathbf{c} \sim p_{\mathcal{D}_p}(\mathbf{c})} [p_{\theta}(\mathbf{x}|\mathbf{c})], \quad (2)$$

where $p_{\mathcal{D}_p}(\mathbf{c})$ is the true underlying text generating distribution corresponding to $\mathbf{x} \in \mathcal{D}_p$. Hence, as per Figure 1, we are inspired to initially utilize the unlabeled data in \mathcal{D} to learn $p_{\theta}(\mathbf{x})$ and obtain the optimal set of model parameters $\hat{\theta} \in \Theta$, which serves as a robust approximation of $p_{\theta}(\mathbf{x}|\mathbf{c})$. Then, we finetune it using data-text pairs in \mathcal{D}_p to achieve desired model control. Crucially, we ensure that the parameters remain in close proximity to those established when learning $p_{\theta}(\mathbf{x})$ via constraint optimization to make Eq. 2 hold while mitigating catastrophic forgetting. Figure 1 shows our constraint to keep the finetuned parameters $\hat{\theta}'$ within $\Theta \cap \Theta'$, where Θ' represents the optimal parameters from unconstrained finetuning. Then we plug the generative diffusion implementation to our framework. Finally, we use the last section to offer theoretical insights into the selection

of the optimization constraint, underpinned by generalization bounds for our learning objective.

Learning Text-to-Data Generation under Low-Resource Situation

To capture general distribution of data, we leverage all data in \mathcal{D} with NULL tokens as conditions (i.e., $\mathbf{c} = \text{NULL}$) to facilitate subsequent finetuning. Specifically, we train the generative model $p_{\theta}(\mathbf{x}|\emptyset)$, where θ parameterizes the model and \emptyset represents the NULL token in practice. As the NULL token is independent to \mathbf{x} , we also have $p_{\theta}(\mathbf{x}|\emptyset) = p_{\theta}(\mathbf{x})$. Hence, we are equivalently optimizing the following:

$$\min_{\theta} \mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} [-\log p_{\theta}(\mathbf{x})], \quad (3)$$

where $p_{\mathcal{D}}(\mathbf{x})$ is the true underlying data generating distribution. Having discerned the general data distribution through Eq. 3, we proceed to finetune $p_{\theta}(\mathbf{x}|\mathbf{c})$ using text labels \mathbf{c} in \mathcal{D}_p to achieve the controllability of model. In alignment with Eq. 2, the anticipated finetuned parameter should approximate the parameter optimized in Eq. 3. We achieve this by finetuning $p_{\theta}(\mathbf{x}|\mathbf{c})$ using the labeled data in \mathcal{D}_p within the optimal set obtained in Eq. 3, leading to the subsequent learning objective for the finetuning phase:

$$\begin{aligned} \min_{\theta} \mathbb{E}_{\mathbf{x}, \mathbf{c} \sim p_{\mathcal{D}_p}(\mathbf{x}, \mathbf{c})} [-\log p_{\theta}(\mathbf{x}|\mathbf{c})] \\ \text{s.t. } \mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}_p}(\mathbf{x})} [-\log p_{\theta}(\mathbf{x})] \leq \xi, \\ \xi = \inf_{\theta \in \Theta} \mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} [-\log p_{\theta}(\mathbf{x})] \end{aligned} \quad (4)$$

where $p_{\mathcal{D}_p}(\mathbf{x}, \mathbf{c})$ is the true underlying data-text joint distribution. Θ denotes a localized parameter space where a minimum can be located. Specifically, we minimize $\mathbb{E}_{\mathbf{x}, \mathbf{c} \sim p_{\mathcal{D}_p}(\mathbf{x}, \mathbf{c})} [-\log p_{\theta}(\mathbf{x}|\mathbf{c})]$ using the labeled data in \mathcal{D}_p within the optimal set $\{\theta : \mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}_p}(\mathbf{x})} [-\log p_{\theta}(\mathbf{x})] \leq \xi\}$ to make the parameter not far from learned via Eq. 3, so that catastrophic forgetting is mitigated. The trade-off between the learning objective and its constraint aligns with the nature of a *lexicographic optimization* problem (Gong and Liu 2021) and therefore can be solved in that context.

Generative Objective on Empirical Samples

Eq. 4 is the population-level learning objective while empirically we follow the standard loss function (i.e., transformed evidence lower bound of Eq. 4) of the classifier-free diffusion guidance (Ho and Salimans 2022) by optimizing:

$$\min_{\theta} \mathcal{L}_2(\theta) \quad \text{s.t. } \mathcal{L}'_1(\theta) \leq \xi, \quad \xi = \inf_{\theta \in \Theta} \mathcal{L}_1(\theta), \quad (5)$$

where we have: $\mathcal{L}_1(\theta) = \mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x}), t} [|\epsilon_{\theta}(\mathbf{x}^{(t)}, t) - \epsilon|^2]$, $\mathcal{L}'_1(\theta) = \mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}_p}(\mathbf{x}), t} [|\epsilon_{\theta}(\mathbf{x}^{(t)}, t) - \epsilon|^2]$ and $\mathcal{L}_2(\theta) = \mathbb{E}_{\mathbf{x}, \mathbf{c} \sim p_{\mathcal{D}_p}(\mathbf{x}, \mathbf{c}), t} [|\epsilon_{\theta}(\mathbf{x}^{(t)}, \mathbf{c}, t) - \epsilon|^2]$. Also, t is sampled from uniform between 1 and T , T is the total diffusion steps, ϵ is the standard Gaussian random variable, and $\epsilon_{\theta}(\mathbf{x}^{(t)}, t)$ and $\epsilon_{\theta}(\mathbf{x}^{(t)}, \mathbf{c}, t)$ are functions we aim to fit at the t -th diffusion step. Note that $\epsilon_{\theta}(\mathbf{x}^{(t)}, t)$ and $\epsilon_{\theta}(\mathbf{x}^{(t)}, \mathbf{c}, t)$ share the same parameters but are just trained at different stages: distribution mastery on unlabeled data and controllable finetuning on labeled data, respectively. The framework of classifier-free diffusion model and the derivation of $\mathcal{L}_1(\theta)$, $\mathcal{L}'_1(\theta)$ and $\mathcal{L}_2(\theta)$ are introduced in Appendix.

As the true data generating distributions $p_{\mathcal{D}}(\mathbf{x})$, $p_{\mathcal{D}_p}(\mathbf{x})$ and $p_{\mathcal{D}_p}(\mathbf{x}, \mathbf{c})$ are unknown, we instead optimize the following empirical loss:

$$\min_{\theta} \hat{\mathcal{L}}_2(\theta) \quad \text{s.t.} \quad \hat{\mathcal{L}}_1(\theta) \leq \hat{\xi}, \quad \hat{\xi} = \inf_{\theta \in \hat{\Theta}} \hat{\mathcal{L}}_1(\theta), \quad (6)$$

where we have $\hat{\mathcal{L}}_1(\theta) = \mathbb{E}_{\mathbf{x} \sim \hat{p}_{\mathcal{D}}(\mathbf{x}), t} [\|\epsilon_{\theta}(\mathbf{x}^{(t)}, t) - \epsilon\|^2]$, $\hat{\mathcal{L}}'_1(\theta) = \mathbb{E}_{\mathbf{x} \sim \hat{p}_{\mathcal{D}_p}(\mathbf{x}), t} [\|\epsilon_{\theta}(\mathbf{x}^{(t)}, t) - \epsilon\|^2]$ and $\hat{\mathcal{L}}_2(\theta) = \mathbb{E}_{\mathbf{x}, \mathbf{c} \sim \hat{p}_{\mathcal{D}_p}(\mathbf{x}, \mathbf{c}), t} [\|\epsilon_{\theta}(\mathbf{x}^{(t)}, \mathbf{c}, t) - \epsilon\|^2]$. $\hat{p}_{\mathcal{D}}(\mathbf{x})$, $\hat{p}_{\mathcal{D}_p}(\mathbf{x})$ and $\hat{p}_{\mathcal{D}_p}(\mathbf{x}, \mathbf{c})$ are corresponding empirical data generating distributions. $\hat{\Theta}$ is the localized parameter space where a minimum can be located for $\hat{\mathcal{L}}_1(\theta)$. The lexicographic optimization of Eq. 6 is presented in Algorithm 1, which relies on learning the update direction on model parameter (i.e., $\nabla \hat{\mathcal{L}}_2(\theta) + \lambda \nabla \hat{\mathcal{L}}'_1(\theta)$ in Algorithm 1) to balance the trade-off between the minimization of $\hat{\mathcal{L}}_2(\theta)$ and $\hat{\mathcal{L}}'_1(\theta)$ using a dynamic gradient descent (Gong and Liu 2021):

$$\theta \leftarrow \theta - \omega \cdot (\nabla \hat{\mathcal{L}}_2(\theta) + \lambda \nabla \hat{\mathcal{L}}'_1(\theta)), \quad (7)$$

where ω is predefined positive step size, and λ is calculated based on whether the constraint is satisfied at the current gradient step:

$$\lambda = \max\left(\frac{\phi(\theta) - \nabla \hat{\mathcal{L}}_2(\theta)^T \nabla \hat{\mathcal{L}}'_1(\theta)}{\|\nabla \hat{\mathcal{L}}'_1(\theta)\|^2}, 0\right)$$

$$\phi(\theta) = \min(\alpha(\hat{\mathcal{L}}'_1(\theta) - \gamma \cdot \hat{\xi}), \beta \|\nabla \hat{\mathcal{L}}'_1(\theta)\|^2), \quad (8)$$

where α , β and γ are predefined positive hyperparameters. More details and intuition of lexicographic optimization are explained in Appendix. Furthermore, the lexicographic optimization-based constraint $\hat{\xi} = \inf_{\theta \in \hat{\Theta}} \hat{\mathcal{L}}_1(\theta)$ in Eq. 6 may be overly strict and could require relaxation to ease the training process. While we anticipate that the parameters derived from Eq. 6 should be close to those from Eq. 5, they do not necessarily have to be an exact subset of the parameters from Eq. 5.

Algorithm 1: Lexicographic optimization on Eq. 4

Input: $\hat{\xi} = \inf_{\theta \in \hat{\Theta}} \hat{\mathcal{L}}_1(\theta)$ by pretraining on \mathcal{D} ; Total diffusion steps T ;
Scheduled forward variance $\{\beta_t\}_{t=1}^T; \{\alpha_t = 1 - \beta_t\}_{t=1}^T$;
Predefined positive hyperparameters α, β, γ and ω ; probability of unconditional training p_{uncond}

while *Not converged* **do**
 Sample $\mathbf{x}, \mathbf{c} \sim \hat{p}_{\mathcal{D}_p}(\mathbf{x}, \mathbf{c}), \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), t \sim U(1, T)$
 Compute $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$
 Diffuse $\mathbf{x}^{(t)} = \sqrt{\bar{\alpha}_t} \mathbf{x} + \sqrt{1 - \bar{\alpha}_t} \epsilon$
 Replace \mathbf{c} with \emptyset with probability p_{uncond}
 Compute $\hat{\mathcal{L}}_2(\theta) = \|\epsilon_{\theta}(\mathbf{x}^{(t)}, \mathbf{c}, t) - \epsilon\|^2$,
 $\hat{\mathcal{L}}'_1(\theta) = \|\epsilon_{\theta}(\mathbf{x}^{(t)}, t) - \epsilon\|^2$
 Compute $\phi(\theta) = \min(\alpha(\hat{\mathcal{L}}'_1(\theta) - \gamma \cdot \hat{\xi}), \beta \|\nabla \hat{\mathcal{L}}'_1(\theta)\|^2)$
 Compute $\lambda = \max\left(\frac{\phi(\theta) - \nabla \hat{\mathcal{L}}_2(\theta)^T \nabla \hat{\mathcal{L}}'_1(\theta)}{\|\nabla \hat{\mathcal{L}}'_1(\theta)\|^2}, 0\right)$
 Update θ by $\theta - \omega \cdot (\nabla \hat{\mathcal{L}}_2(\theta) + \lambda \nabla \hat{\mathcal{L}}'_1(\theta))$

graphic optimization-based constraint $\hat{\xi} = \inf_{\theta \in \hat{\Theta}} \hat{\mathcal{L}}_1(\theta)$ in Eq. 6 may be overly strict and could require relaxation to ease the training process. While we anticipate that the parameters derived from Eq. 6 should be close to those from Eq. 5, they do not necessarily have to be an exact subset of the parameters from Eq. 5.

Generalization Bound of Learning Constraint

In this section, we deduce a confidence bound for the constraint to demonstrate that the optimal set of minimizing the empirical loss within the derived confidence bound encapsulates the true one, and guide the further relaxation on the

constraint if needed. First, we define sub-Gaussian random variable in Definition 0.1 followed by the deduced confidence bound in Theorem 0.2.

Definition 0.1 (Sub-Gaussian random variable). The random variable X with mean 0 is sub-Gaussian with variance σ^2 if $\forall s \in \mathbb{R}, \mathbb{E}_X[\exp(sX)] \leq \exp(\frac{\sigma^2 s^2}{2})$.

Based on the fact that $\mathbf{x}^{(t)}$ is diffused from the random variable \mathbf{x} and the standard Gaussian noise ϵ (i.e., Algorithm 1), therefore, $\epsilon_{\theta}(\mathbf{x}^{(t)}, t)$ and $\epsilon_{\theta}(\mathbf{x}^{(t)}, \mathbf{c}^{(t)}, t)$ are also random variables. Meanwhile, minimizing the loss in Eq. 5 pushes $\epsilon_{\theta}(\mathbf{x}^{(t)}, t)$ and $\epsilon_{\theta}(\mathbf{x}^{(t)}, \mathbf{c}^{(t)}, t)$ towards ϵ , which has the mean of 0. Then we introduce the following theorem 1 while assuming $\epsilon_{\theta}(\mathbf{x}^{(t)}, t)$ and $\epsilon_{\theta}(\mathbf{x}^{(t)}, \mathbf{c}^{(t)}, t)$ are sub-Gaussian random variables.

Theorem 0.2. For every θ and t , assume $\epsilon_{\theta}(\mathbf{x}^{(t)}, t)$ and $\epsilon_{\theta}(\mathbf{x}^{(t)}, \mathbf{c}_i, t)$ are sub-Gaussian random variables with mean 0 and variance σ^2 , and Θ is finite. Let $\Theta^* = \{\theta : \hat{\mathcal{L}}'_1(\theta) \leq \xi\}$, $\hat{\Theta}^* = \{\theta : \hat{\mathcal{L}}'_1(\theta) \leq \hat{\xi} + \epsilon\}$ where ϵ is the confidence bound with the probability of $1 - \delta$. Let θ^* be the solution to Eq. 5 and $\hat{\theta}^*$ be the solution to empirical Eq. 6, then we have the following:

1. $\Theta^* \subseteq \hat{\Theta}^*$: the set of θ by optimizing Eq. 6 within the confidence bound covers the true one.
2. $\mathcal{L}_2(\hat{\theta}^*) \leq \mathcal{L}_2(\theta^*) + 2\epsilon_{N_p}$: θ^* and $\hat{\theta}^*$ compete well on $\mathcal{L}_2(\theta)$.
3. $\mathcal{L}'_1(\hat{\theta}^*) \leq \xi + 2\epsilon_{N_p} + 2\epsilon_N$: $\hat{\theta}^*$ does not violate constraint of Eq. 5 too much on $\mathcal{L}'_1(\theta)$.

where $\epsilon = \epsilon_N + \epsilon_{N_p}$, $\epsilon_N = \sqrt{C\tilde{\sigma}^2} \cdot \sqrt{\frac{\log|\Theta| + \log\frac{2}{\delta}}{N}} \vee C\tilde{\sigma}^2 \cdot \frac{\log|\Theta| + \log\frac{2}{\delta}}{N}$, $\epsilon_{N_p} = \sqrt{C\tilde{\sigma}^2} \cdot \sqrt{\frac{\log|\Theta| + \log\frac{2}{\delta}}{N_p}} \vee C\tilde{\sigma}^2 \cdot \frac{\log|\Theta| + \log\frac{2}{\delta}}{N_p}$, $\tilde{\sigma}^2 = \sigma^2 + 1$ and $C = 8\sqrt{2}$.

Theorem 0.2 can be proved starting from the Bernstein's inequality and the union bound inequality on squared zero-mean sub-Gaussian random variable. More detailed proof is in Appendix. As a large amount of unlabeled data (i.e., N) is usually easy to obtain by either manual collection or simulation, ϵ_N is not large after taking logarithm on the number of model parameters (i.e., $|\Theta|$), even though it is usually much larger than N . Additionally, $\log|\Theta|$ is not significantly larger than N_p so that ϵ_{N_p} should not be large as well. For instance, in our experiments, around 45,000 samples and 14 million model parameters for motion generation result in rather small ϵ_N and ϵ_{N_p} . In practice, we use $\xi = \rho \cdot \inf_{\theta \in \Theta} \mathcal{L}_1(\theta)$ in Eq. 5 (i.e., $\hat{\xi} = \rho \cdot \inf_{\theta \in \hat{\Theta}} \hat{\mathcal{L}}_1(\theta)$ in Eq. 6) to relax the constraint, where ρ is an hyperparameter to keep the constraint within the confidence interval.

Experiments

Datasets

We employ datasets from three modalities that may suffer from low-resource scenario. Details are in Appendix.

Molecules. We extract 130,831 molecules from QM9 dataset (Ramakrishnan et al. 2014) with six molecular properties: polarizability (α), highest occupied molecular orbital energy (ϵ_{HOMO}), lowest unoccupied molecular orbital energy (ϵ_{LUMO}), the energy difference between HOMO and LUMO ($\Delta\epsilon$), dipole moment (μ) and heat capacity at 298.15K (C_v).

Motions. We employ HumanML3D that contains textually re-annotating motions captured from AMASS (Mahmood et al. 2019) and HumanAct12 (Guo et al. 2020). It contains 14,616 motions annotated by 44,970 textual descriptions.

Time Series. We assemble 24 stocks from Yahoo Finance from their IPO date to July 8, 2023. We tailor data to the length of 120 by slicing on the opening price for every 120 days, and scale by min-max normalization following Yoon, Jarrett, and Van der Schaar (2019). Totally, 210,964 time series are produced. We extract their features including frequency, skewness, mean, variance, linearity (i.e., R^2), and the number of peaks via "tsfresh" in Python.

Baseline Models

We compare Text2Data with a representative classifier-free diffusion model in each modality as the baselines.

E(3) Equivariant Diffusion Model (EDM) (Hoogetboom et al. 2022). To handle Molecule dataset, we employ EDM as the baseline. EDM utilizes an equivariant network to denoise diffusion processes by concurrently processing both continuous (i.e., atom coordinates) and categorical data (i.e., atom types). The controllability on molecular properties is realized by the classifier-free diffusion guidance conditioned on the embedding of text descriptions, which is encoded by a pretrained T5 (Raffel et al. 2020) encoder (i.e., t5-large).

Motion Diffusion Model (MDM) (Tevet et al. 2023). We use MDM, a classifier-free diffusion model, for text-to-human motion generation. The text descriptions are embedded by a pretrained T5 encoder to guide the motion generation, providing the mechanism for controllability.

Generation diffusion for time series (DiffTS). To generate time series, we design the classifier-free diffusion model conditioned on text embeddings encoded from pretrained T5 encoder. We employ the backbone of (Ho and Salimans 2022) by substituting image data to one-dimensional time series, and replacing the U-Net with one-dimensional convolutional neural network.

For implementation, we train baselines on specific proportions of labeled data. Text2Data is modified from a pretraining+finetuning strategy following Eq. 6. Additionally, we conduct ablation study where each baseline is still finetuned but without the constraint in Eq. 6 as a transfer learning-based baseline, or directly applied to augmented text-data pairs. Besides, we adapt from (You et al. 2024) to form a semi-supervised framework, by training a classifier to predict molecular properties and generate pseudo labels for unlabelled molecules. More details are in Appendix.

Evaluation Metrics

We follow different strategies to evaluate both the controllability and the generation quality of the propose approach compared to other baselines.

Controllability. We compare the controllability between Text2Data and the baselines by the similarity between generated data and the ground truth. To assess the generated molecules, we follow (Hoogetboom et al. 2022) and train a classifier for each property to extract specific properties from generated data. Then, we calculate the *Mean Absolute Error* (MAE) between the extracted property and the ground truth. To assess the controllability of motion generation, we compute *R precision* and *Multimodal distance* that measure the relevancy of the generated motions to the input prompts (Guo et al. 2022). To evaluate the controllability of time series generation, we extract properties via "tsfresh" and compute the MAE between properties of generated data and that of ground truth. Additionally, we also visualize generated data according to the specified properties.

Generation quality. Generation quality varies based on the modality. For molecular generation, we compute $-\log \text{likelihood}$ ($-\log p$) (Hoogetboom et al. 2022) and *validity* of generated molecules (Jin, Barzilay, and Jaakkola 2018), *molecular stability* and *atom stability* (Garcia Satorras et al. 2021) to evaluate their overall generation quality. For motion generation, we use *FID score* and *Diversity* following (Tevet et al. 2023). For time series generation, we follow (Yoon, Jarrett, and Van der Schaar 2019) and (Lim et al. 2023) by drawing t-SNE plots to visualize the overlap between generated data and the ground-truth. Better model tends to have a larger overlap, indicating more similar distribution.

Comparisons on Controllability

Figure 2 illustrates the MAE trend between properties of generated molecules and the intended one as the proportion of labeled training data rises. Text2Data achieves superior performance than EDM-finetune and EDM on all properties by a remarkable margin. The results also indicate that certain properties, such as ϵ_{LUMO} and C_v , are more readily controllable. For these properties, the performance of the three models converges as the amount of labeled training data becomes sufficiently large (Figure 2). We further increase the proportion of available labels in the dataset up to 100% and, as indicated in Appendix Table 3, the MAE keeps increasing when more labels are involved during training, and it gradually converges in the end. Appendix Table 4 also suggests that Text2Data surpasses the data augmentation-based method, which may suffer from the potentially poor alignment between text and data and the potential overfitting.

We depict the molecules generated as the text descriptor for polarizability shifts from "very low" to "very high" in Figure 3. Polarizability indicates the inclination of the molecule to form an electric dipole moment under an external electric field. As α values rise, we expect to see molecules with less symmetrical forms, as evidenced in Figure 3. This trend suggests the validity of generated molecules by Text2Data and its fine-grained controllability.

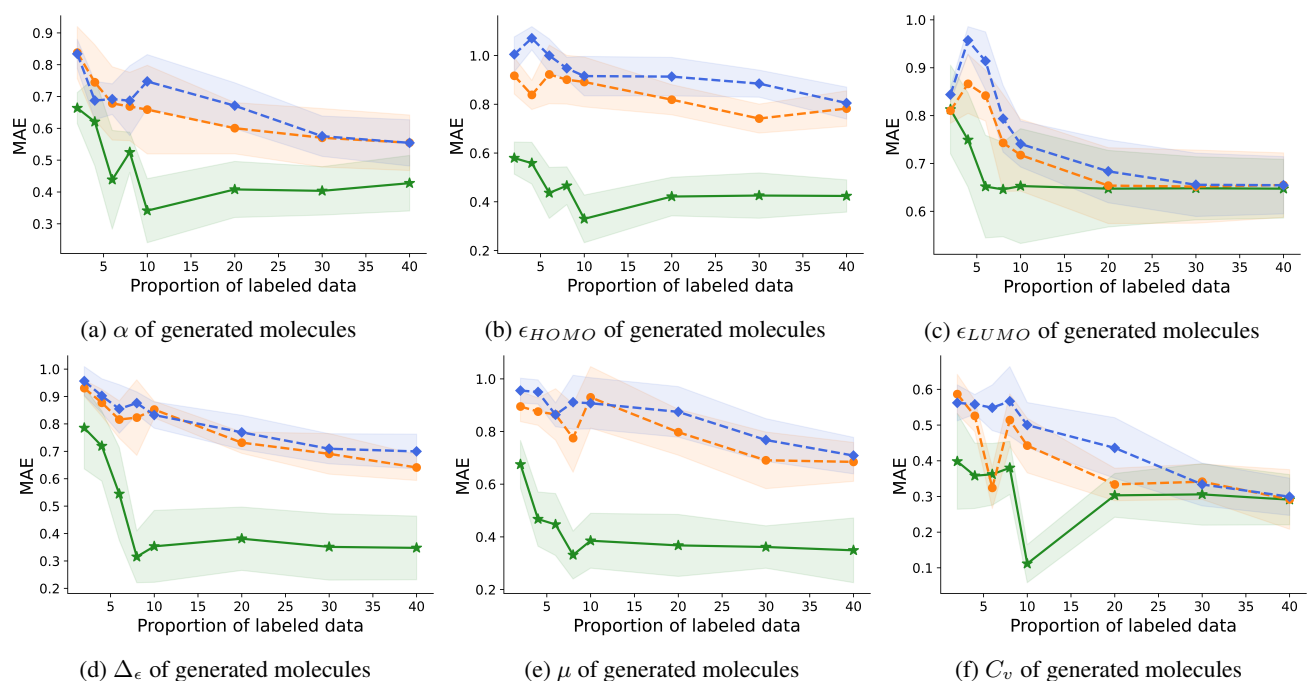


Figure 2: Evaluate controllability on Molecule dataset according to different proportions of paired data. Green solid line corresponds to Text2Data and two dashed lines are baseline comparisons, in which blue line is EDM and orange line is EDM-finetune. Properties of generated molecules are predicted by classifier ϕ_c . MAE is computed between properties of generated molecules and intended properties. Lower MAE indicates better performance.




Figure 3: Visualization of generated molecules when the polarizability increases from “very low” to “very high”.

Proportion (%)	R Precision \uparrow			Multimodal Dist. \downarrow		
	Text2Data	MDM-finetune	MDM	Text2Data	MDM-finetune	MDM
2	0.34 \pm 0.01	0.37\pm0.01	0.31 \pm 0.01	6.48 \pm 0.06	6.19\pm0.05	6.67 \pm 0.02
4	0.39 \pm 0.01	0.42\pm0.01	0.38 \pm 0.01	5.99 \pm 0.05	5.83\pm0.04	6.01 \pm 0.04
6	0.43 \pm 0.01	0.43\pm0.01	0.40 \pm 0.02	5.85 \pm 0.06	5.78\pm0.05	6.01 \pm 0.06
8	0.44\pm0.01	0.43 \pm 0.01	0.42 \pm 0.01	5.65\pm0.04	5.75 \pm 0.05	5.90 \pm 0.04
10	0.45\pm0.01	0.45 \pm 0.01	0.44 \pm 0.01	5.74\pm0.07	5.76 \pm 0.07	5.84 \pm 0.06
20	0.47\pm0.01	0.47 \pm 0.01	0.45 \pm 0.01	5.61\pm0.04	5.68 \pm 0.11	5.73 \pm 0.14
30	0.48\pm0.01	0.47 \pm 0.01	0.45 \pm 0.11	5.61\pm0.05	5.66 \pm 0.06	5.80 \pm 0.09
40	0.49\pm0.01	0.46 \pm 0.01	0.45 \pm 0.01	5.61\pm0.05	5.63 \pm 0.09	5.90 \pm 0.04

Table 1: Evaluate controllability on HumanML3D dataset by R Precision and Multimodal Distance according to different proportions of paired data.

As suggested in Table 1, Text2Data also outperforms MDM-finetune and MDM in the controllable generation of motions from text descriptions. While MDM-finetune is slightly better than Text2Data when the proportion of labeled training data is small-owing to milder catastrophic forgetting dur-

ing finetuning with a smaller sample size-Text2Data consistently surpasses both MDM-finetune and MDM as the volume of labeled training data increases. Specifically, in this situation, Text2Data surpasses MDM-finetune and MDM in R Precision with average margins of 2.31% and 5.57%, re-

Proportion (%)	Frequency ($\times 10^{-1}$)			Skewness			Mean ($\times 10^{-2}$)		
	Text2Data	DiffTS-finetune	DiffTS	Text2Data	DiffTS-finetune	DiffTS	Text2Data	DiffTS-finetune	DiffTS
2	2.59±0.20	2.62±0.20	2.60±0.17	1.68±0.20	2.34±0.30	1.84±0.22	0.63±0.39	0.63±0.39	0.79±0.41
4	2.55±0.18	2.59±0.19	2.59±0.19	1.63±0.14	1.77±0.29	2.80±0.28	0.60±0.40	0.61±0.39	0.73±0.40
6	2.52±0.18	2.57±0.19	2.57±0.19	1.00±0.16	1.85±0.22	1.81±0.21	0.56±0.38	0.58±0.38	0.71±0.36
8	2.54±0.18	2.56±0.19	2.57±0.18	1.10±0.18	1.56±0.18	1.78±0.09	0.57±0.38	0.63±0.40	0.62±0.36
10	2.54±0.19	2.57±0.18	2.55±0.20	0.87±0.13	1.20±0.17	1.12±0.11	0.55±0.37	0.57±0.36	0.62±0.39
20	2.54±0.18	2.55±0.18	2.58±0.21	1.05±0.15	1.06±0.12	1.26±0.14	0.55±0.40	0.57±0.37	0.65±0.38
30	2.53±0.18	2.56±0.17	2.56±0.18	1.03±0.12	1.16±0.25	1.71±0.23	0.51±0.33	0.53±0.33	0.59±0.34
40	2.53±0.18	2.55±0.18	2.55±0.19	1.03±0.11	1.15±0.24	1.19±0.18	0.51±0.33	0.57±0.32	0.57±0.36

Table 2: Evaluate controllability on time series by MAE on testing set, according to different proportions of paired data. Lower MAE indicates better performance.

spectively, and in Multimodal Distance with average margins of 0.93% and 3.30%, respectively. The results also indicate that an increase in labeled training data enhances the performance of controllability, which is expected as more supervision is involved.

Additionally, we evaluate controllability of Text2Data, along with its baseline comparisons, utilizing MAE to measure the congruence between the property of generated data and the intended one within the Time Series dataset. As indicated in Table 2, Text2Data consistently excels over two baselines, DiffTS-finetune and DiffTS, across all three properties assessed. Results of another three properties are presented in Appendix Table 5, which suggests the similar conclusion. Specifically, Text2Data and DiffTS-finetune both show a marked improvement over DiffTS in controlling frequency, variance, and skewness. They also exhibit a slight edge in controlling mean, number of peaks, and linearity. The enhanced performance of Text2Data correlates with its proficiency in alleviating the issue of catastrophic forgetting while maintaining a pursuit of controllability.

Comparisons on Generation Quality

Text2Data not only demonstrates superior performance in controllable text-to-data generation but also sustains competitive generation quality relative to baseline models.

When generating molecules from Text2Data and its baseline comparisons, as shown in Appendix Table 6, we compute $-\log p$ and validity to evaluate generation quality. The performance of Text2Data is consistently better. It surpasses EDM-finetune and EDM by average margins of 19.07% and 58.03%, respectively. It is 1.98% and 10.59% better than EDM-finetune and EDM on average, respectively, regarding validity of molecules. Besides, we evaluate Text2Data compared with EDM-finetune and EDM on molecular stability and atom stability (Appendix Table 7). Text2Data exceeds EDM-finetune and EDM by average margins of 2.34% and 17.31%, respectively, regarding molecular stability. It is also 0.29% and 1.21% better than EDM-finetune and EDM on average, respectively, regarding atom stability. The consistent improvements on all the three models result from our superior performance of Text2Data on properties (e.g., molecular stability) that are hard to control.

As indicated in Appendix Table 8, quantitative assessment of motion generation from text shows that Text2Data sur-

passes baseline methods in both quality and diversity. Particularly, Text2Data outperforms MDM-finetune and MDM by 2.73% and 36.05% on average, respectively, regarding FID. For diversity, Text2Data surpasses MDM-finetune and MDM by average margins of 0.81% and 3.71%, respectively. Enhanced performance is derived from Text2Data to fully leverage all samples in the dataset, while effectively mitigating catastrophic forgetting during finetuning.

We evaluate the overall quality of generating time series by making t-SNE plots of generated time series against ground truth. Substantial overlap between the generated time series and the ground truth suggests a closer distribution alignment, implying a better performance of Text2Data. As demonstrated in Appendix Figure 4, specifically, the red pattern represents the t-SNE of the ground-truth time series, whereas the blue pattern represents the t-SNE of the generated time series according to the same text description. Compared with DiffTS-finetune and DiffTS, Text2Data corresponds to the largest overlap between distributions of the generated and the ground-truth time series, suggesting its superior ability to precisely generate data according to the text description. The non-overlapping part may result from the diversity of generated time series or some other properties not labeled in the dataset so that not controlled by text description. The inferior performance of DiffTS stems from its training solely on labeled data, potentially leading to an incomplete understanding of the overall data distribution and a risk of overfitting due to limited size of labeled data. DiffTS may only partially capture the data distribution based on textual descriptions due to its susceptibility to catastrophic forgetting, which also heightens the risk of overfitting.

Conclusion

We propose Text2Data to improve quality and property control of text-to-data generation for various modalities in low-resource scenarios using diffusion models. Text2Data uses unlabeled data to capture the prevailing data distribution via unsupervised diffusion model. It is then finetuned on text-labeled data, with a novel constraint optimization-based learning objective to ensure controllability while reducing catastrophic forgetting. Experiments show consistently superior performance to recent baselines. While Text2Data is presented as a diffusion-based framework in this article, it can be seamlessly adapted to other generative models.

References

- Cheuk, K. W.; Herremans, D.; and Su, L. 2021. Reconvat: A semi-supervised automatic music transcription framework for low-resource real-world data. In *Proceedings of the 29th ACM International Conference on Multimedia*, 3918–3926.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- Du, S.; Li, T.; Yang, Y.; and Horng, S.-J. 2020. Multivariate time series forecasting via attention-based encoder–decoder framework. *Neurocomputing*, 388: 269–279.
- Garcia Satorras, V.; Hoogeboom, E.; Fuchs, F.; Posner, I.; and Welling, M. 2021. E (n) equivariant normalizing flows. *Advances in Neural Information Processing Systems*, 34: 4181–4192.
- Gong, C.; and Liu, X. 2021. Bi-objective trade-off with dynamic barrier gradient descent. *Advances in Neural Information Processing Systems*.
- Gozalo-Brizuela, R.; and Garrido-Merchan, E. C. 2023. ChatGPT is not all you need. A State of the Art Review of large Generative AI models. *arXiv preprint arXiv:2301.04655*.
- Guo, C.; Zou, S.; Zuo, X.; Wang, S.; Ji, W.; Li, X.; and Cheng, L. 2022. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5152–5161.
- Guo, C.; Zuo, X.; Wang, S.; Zou, S.; Sun, Q.; Deng, A.; Gong, M.; and Cheng, L. 2020. Action2motion: Conditioned generation of 3d human motions. In *Proceedings of the 28th ACM International Conference on Multimedia*, 2021–2029.
- Hedderich, M. A.; Lange, L.; Adel, H.; Strötgen, J.; and Klakow, D. 2020. A survey on recent approaches for natural language processing in low-resource scenarios. *arXiv preprint arXiv:2010.12309*.
- Ho, J.; Chan, W.; Saharia, C.; Whang, J.; Gao, R.; Gritsenko, A.; Kingma, D. P.; Poole, B.; Norouzi, M.; Fleet, D. J.; et al. 2022. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Hoogeboom, E.; Satorras, V. G.; Vignac, C.; and Welling, M. 2022. Equivariant diffusion for molecule generation in 3d. In *International conference on machine learning*, 8867–8887. PMLR.
- Hu, Y.; Luo, C.; and Chen, Z. 2022. Make it move: controllable image-to-video generation with text descriptions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18219–18228.
- Huang, R.; Huang, J.; Yang, D.; Ren, Y.; Liu, L.; Li, M.; Ye, Z.; Liu, J.; Yin, X.; and Zhao, Z. 2023. Make-an-audio: Text-to-audio generation with prompt-enhanced diffusion models. *arXiv preprint arXiv:2301.12661*.
- Huang, R.; Zhao, Z.; Liu, H.; Liu, J.; Cui, C.; and Ren, Y. 2022. Prodiff: Progressive fast diffusion model for high-quality text-to-speech. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2595–2605.
- Iman, M.; Arabnia, H. R.; and Rasheed, K. 2023. A review of deep transfer learning and recent advancements. *Technologies*, 11(2): 40.
- Irwin, J. J.; Sterling, T.; Mysinger, M. M.; Bolstad, E. S.; and Coleman, R. G. 2012. ZINC: a free tool to discover chemistry for biology. *Journal of chemical information and modeling*, 52(7): 1757–1768.
- Ito, K.; and Johnson, L. 2017. The LJ Speech Dataset. <https://keithito.com/LJ-Speech-Dataset/>.
- Jiang, Y.; Huang, Z.; Pan, X.; Loy, C. C.; and Liu, Z. 2021. Talk-to-Edit: Fine-Grained Facial Editing via Dialog. In *Proceedings of International Conference on Computer Vision (ICCV)*.
- Jin, W.; Barzilay, R.; and Jaakkola, T. 2018. Junction tree variational autoencoder for molecular graph generation. In *International conference on machine learning*, 2323–2332. PMLR.
- Kawar, B.; Zada, S.; Lang, O.; Tov, O.; Chang, H.; Dekel, T.; Mosseri, I.; and Irani, M. 2023. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6007–6017.
- Kim, H.; Kim, S.; and Yoon, S. 2022. Guided-tts: A diffusion model for text-to-speech via classifier guidance. In *International Conference on Machine Learning*, 11119–11133. PMLR.
- Kumari, N.; Zhang, B.; Zhang, R.; Shechtman, E.; and Zhu, J.-Y. 2023. Multi-concept customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1931–1941.
- Li, M.; Duan, Y.; Zhou, J.; and Lu, J. 2023. Diffusion-sdf: Text-to-shape via voxelized diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12642–12651.
- Li, Y.; Min, M.; Shen, D.; Carlson, D.; and Carin, L. 2018. Video generation from text. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Lim, H.; Kim, M.; Park, S.; and Park, N. 2023. Regular Time-series Generation using SGM. *arXiv preprint arXiv:2301.08518*.
- Lin, C.-H.; Gao, J.; Tang, L.; Takikawa, T.; Zeng, X.; Huang, X.; Kreis, K.; Fidler, S.; Liu, M.-Y.; and Lin, T.-Y. 2023. Magic3d: High-resolution text-to-3d content creation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 300–309.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, 740–755. Springer.

- Liu, H.; Chen, Z.; Yuan, Y.; Mei, X.; Liu, X.; Mandic, D.; Wang, W.; and Plumbley, M. D. 2023. AudioLDM: Text-to-Audio Generation with Latent Diffusion Models. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, 21450–21474. PMLR.
- Mahmood, N.; Ghorbani, N.; Troje, N. F.; Pons-Moll, G.; and Black, M. J. 2019. AMASS: Archive of motion capture as surface shapes. In *Proceedings of the International Conference on Computer vision (ICCV)*, 5442–5451.
- Meng, L.; Xu, J.; Tan, X.; Wang, J.; Qin, T.; and Xu, B. 2021. Mixspeech: Data augmentation for low-resource automatic speech recognition. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 7008–7012. IEEE.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research*, 21(140): 1–67.
- Ramakrishnan, R.; Dral, P. O.; Rupp, M.; and Von Lilienfeld, O. A. 2014. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1): 1–7.
- Ruiz, N.; Li, Y.; Jampani, V.; Pritch, Y.; Rubinstein, M.; and Aberman, K. 2023. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22500–22510.
- Sanghi, A.; Chu, H.; Lambourne, J. G.; Wang, Y.; Cheng, C.-Y.; Fumero, M.; and Malekshan, K. R. 2022. Clip-forge: Towards zero-shot text-to-shape generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18603–18613.
- Shen, F.; Jiang, X.; He, X.; Ye, H.; Wang, C.; Du, X.; Li, Z.; and Tang, J. 2024a. Imagdressing-v1: Customizable virtual dressing. *arXiv preprint arXiv:2407.12705*.
- Shen, F.; and Tang, J. 2024. IMAGPose: A Unified Conditional Framework for Pose-Guided Person Generation. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Shen, F.; Ye, H.; Liu, S.; Zhang, J.; Wang, C.; Han, X.; and Yang, W. 2024b. Boosting consistency in story visualization with rich-contextual conditional diffusion models. *arXiv preprint arXiv:2407.02482*.
- Shen, F.; Ye, H.; Zhang, J.; Wang, C.; Han, X.; and Wei, Y. 2024c. Advancing Pose-Guided Image Synthesis with Progressive Conditional Diffusion Models. In *The Twelfth International Conference on Learning Representations*.
- Tevet, G.; Raab, S.; Gordon, B.; Shafir, Y.; Cohen-or, D.; and Bermano, A. H. 2023. Human Motion Diffusion Model. In *The Eleventh International Conference on Learning Representations*.
- Thomas, S.; Seltzer, M. L.; Church, K.; and Hermansky, H. 2013. Deep neural network features and semi-supervised training for low resource speech recognition. In *2013 IEEE international conference on acoustics, speech and signal processing*, 6704–6708. IEEE.
- Tits, N.; El Haddad, K.; and Dutoit, T. 2020. Exploring transfer learning for low resource emotional tts. In *Intelligent Systems and Applications: Proceedings of the 2019 Intelligent Systems Conference (IntelliSys) Volume 1*, 52–60. Springer.
- Tu, T.; Chen, Y.-J.; Yeh, C.-c.; and Lee, H.-Y. 2019. End-to-end text-to-speech for low-resource languages by cross-lingual transfer learning. *arXiv preprint arXiv:1904.06508*.
- Wang, C.; Wu, A.; and Pino, J. 2020. CoVoST 2: A Massively Multilingual Speech-to-Text Translation Corpus. *arXiv:2007.10310*.
- Wang, S.; Du, Y.; Guo, X.; Pan, B.; Qin, Z.; and Zhao, L. 2022. Controllable data generation by deep learning: A review. *arXiv preprint arXiv:2207.09542*.
- Yang, D.; Yu, J.; Wang, H.; Wang, W.; Weng, C.; Zou, Y.; and Yu, D. 2023. Diffsound: Discrete diffusion model for text-to-sound generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Yi, J.; Tao, J.; Wen, Z.; and Bai, Y. 2018. Language-adversarial transfer learning for low-resource speech recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(3): 621–630.
- Yin, Y.; Kaddour, J.; Zhang, X.; Nie, Y.; Liu, Z.; Kong, L.; and Liu, Q. 2023. TTIDA: Controllable Generative Data Augmentation via Text-to-Text and Text-to-Image Models. *arXiv preprint arXiv:2304.08821*.
- Yoon, J.; Jarrett, D.; and Van der Schaar, M. 2019. Time-series generative adversarial networks. *Advances in neural information processing systems*, 32.
- You, Z.; Zhong, Y.; Bao, F.; Sun, J.; Li, C.; and Zhu, J. 2024. Diffusion models and semi-supervised learners benefit mutually with few labels. *Advances in Neural Information Processing Systems*, 36.
- Zang, H.; and Wan, X. 2019. A semi-supervised approach for low-resourced text generation. *arXiv preprint arXiv:1906.00584*.
- Zhang, L.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. *Proceedings of the International Conference on Computer Vision (ICCV)*.
- Zhang, Z.; Han, L.; Ghosh, A.; Metaxas, D. N.; and Ren, J. 2023. Sine: Single image editing with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6027–6037.