

Constrained Generative Modeling with Manually Bridged Diffusion Models

Saeid Naderiparizi^{*1,3}, Xiaoxuan Liang^{*1,3}, Berend Zwartsenberg³, Frank Wood^{1,2,3}

¹Department of Computer Science, University of British Columbia, Vancouver, Canada

²Alberta Machine Intelligence Institute (Amii), Edmonton, Canada

³InvertedAI, Vancouver, Canada

saeidnp@cs.ubc.ca, liang51@cs.ubc.ca, berend.zwartsenberg@inverted.ai, fwood@cs.ubc.ca

Abstract

In this paper we describe a novel framework for diffusion-based generative modeling on constrained spaces. In particular, we introduce manual bridges, a framework that expands the kinds of constraints that can be practically used to form so-called diffusion bridges. We develop a mechanism for combining multiple such constraints so that the resulting multiply-constrained model remains a manual bridge that respects all constraints. We also develop a mechanism for training a diffusion model that respects such multiple constraints while also adapting it to match a data distribution. We develop and extend theory demonstrating the mathematical validity of our mechanisms. Additionally, we demonstrate our mechanism in constrained generative modeling tasks, highlighting a particular high-value application in modeling trajectory initializations for path planning and control in autonomous vehicles.

Code and extended version —

github.com/plai-group/manually-bridged-models

1 Introduction

For generative models to become practically useful in embodied artificial intelligence domains, guaranteed constraint satisfaction is effectively required. Examples abound, such as path planning and control in autonomous vehicles and advanced driver-assistance (AV/ADAS) systems (Janner et al. 2022; Zhong et al. 2023), kinematic, dynamics, power, and other constraints in robotics (Schulman et al. 2014), safety critical plant operation (Knight 2002), etc. Strictly eliminating non-factual or offensive hallucinations from large language models (LLM) falls in this problem category too (Azamfirei, Kudchadkar, and Fackler 2023; Huang et al. 2023).

Our experimental focus in this paper will be on a particular subproblem in AV/ADAS planning and behavioral simulation, so starting here, we will motivate our work using language from this domain; however, note that the mechanisms and theory we develop are general.

Consider the problem of realistically distributing agents in top-down, two-dimensional space; cars, pedestrians, etc.

^{*}These authors contributed equally.

There are several characteristics of such a distribution that are sufficiently close to being constraints that they may as well be. Cars and other vehicles are both constrained to be “on road” and also to be not overlapping (colliding). Existing generative models fit to even very large datasets of such data struggle in the sense that samples from them exhibit “in-fractions” (constraint violations) at excessive rates (Zwartsenberg et al. 2023; Jiang et al. 2024; Niedoba et al. 2024); even when the data on which they are trained is cleaned to contain only non-infracting examples.

Why does this occur? In the non-parametric limit this problem would not exist. However, in any finite-data and finite-capacity model the particular generalization strategy the model employs remains a degree of freedom. We are motivated to seek modeling approaches that allow us to direct and exclude generalizations that place mass outside whatever problem specific constraint set there is.

Various approaches to this have been developed (Chang et al. 2024; Huang et al. 2024); this paper explores diffusion-bridge-like mechanisms, inspired by diffusion bridge theory, that empirically demonstrate superior constraint-satisfaction at little expense to generalization otherwise. We call these mechanisms “manual bridges” to distinguish them from the more formally mathematically delimited bridge functions employed in the greater diffusion bridge literature (Schauer, van der Meulen, and van Zanten 2017; De Bortoli et al. 2021; Wang et al. 2021; Heng et al. 2021; Chen, Liu, and Theodorou 2022; Zhou et al. 2024; Shi et al. 2024).

Our central contribution is an architecture that allows both for imposition of so called “manual bridges” to impose constraints in diffusion-based generative models and stable training of models that are constrained in this way, resulting in “manually bridged models,” a novel family of generative models that are capable of fitting complex distributions well while also respecting and representing the sharp boundaries imposed by constraints.

2 Related Work

A framework for constraints on the sample space of generative models was first theoretically described in (Hanneke et al. 2018). In their framework, constraints are represented through a black-box oracle function labeling samples as valid or invalid. This problem has been practically explored for generative adversarial networks (GANs) (Kong

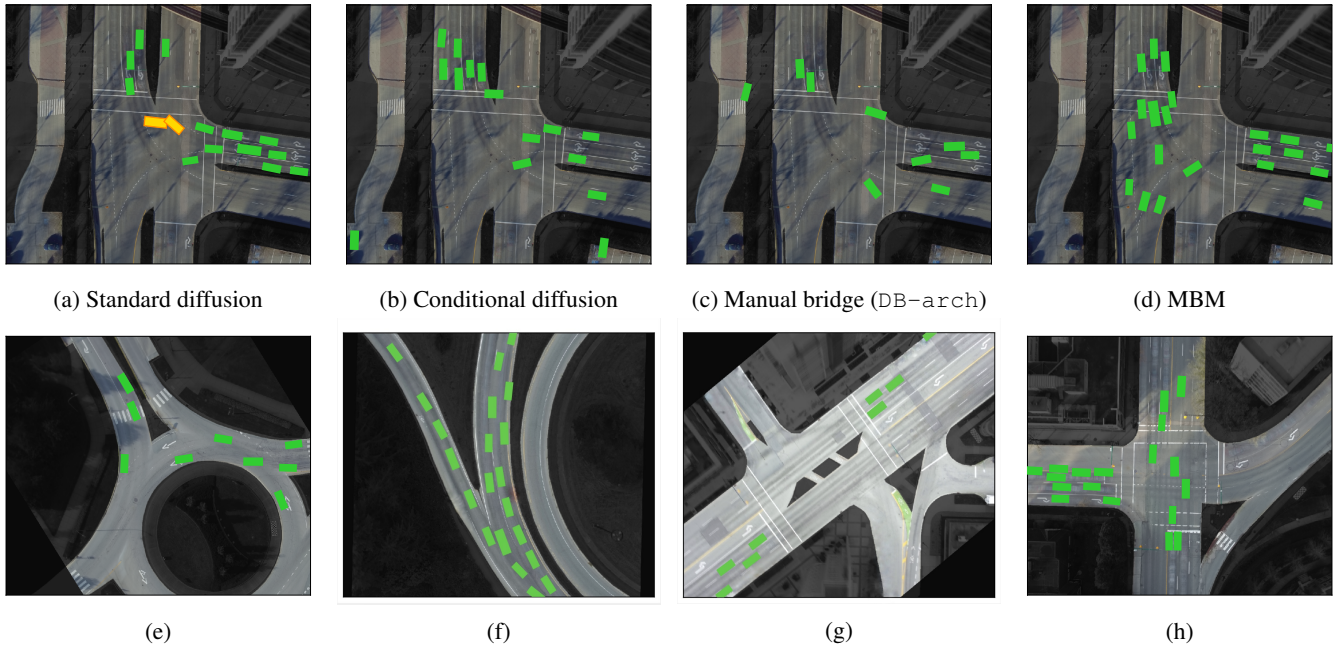


Figure 1: MBM applied to traffic scene generation. The goal is to place vehicles on a given bird’s-eye view image of a map; indicated here as the “light” region of an underlying aerial image. The model output is the set of “cars” (infraction-free cars are green; cars involved in infractions are yellow). The top row shows samples from different models given the same map. The standard diffusion sample (a) contains a collision infraction. The rest of the top row shows different architectural mechanisms to avoid infractions. Both conditional diffusion (b) and (c) are not realistic: they both distort the distribution, albeit in different ways, this effect being more apparent in (c). A sample from MBM in (d) shows no infractions while remaining realistic. The second row shows samples from MBM on additional maps.

and Chaudhuri 2023) and diffusion models (Naderiparizi et al. 2024). More recently, Christopher, Baek, and Fioretto (2024) proposed a method for generating constraint satisfying samples from pre-trained diffusion models. It requires a projection operator the constraint set, which is generally intractable for complex constraints. Moreover, they employ a Langevin dynamics-based (Welling and Teh 2011) sampler which is slow to converge. The main idea in all of these methods was to improve generative models by incorporating information from the constraints. In this paper, however, the goal is to construct a model family that does not generate invalid samples by design.

Recognizing the expressivity of diffusion models, various approaches to incorporating pre-defined constraints into them have emerged in the literature. Lou and Ermon (2023) proposed reflected diffusion models that enforce the whole diffusion sampling trajectory to remain bounded in a convex set. Fishman et al. (2023, 2024) extended reflected diffusion models to support more general constraints. However, their approaches are only evaluated on low-dimensional problems with simple constraints. Moreover, reflected diffusion makes the forward process, and consequently training, expensive. Fishman et al. (2023) also proposed a barrier function based approach for constrained diffusion models. Liu et al. (2024) used barrier functions to transform constrained domains into unconstrained dual ones. Both these methods only support

convex constraints.

Another closely related body of work is diffusion bridges, stochastic processes that are guaranteed to end in a given constraint set. Wu et al. (2022) developed a set of mathematically sufficient conditions for designing diffusion bridges to a given constraint set. The follow-up work of Liu et al. (2023) used bridges to formulate diffusion models on discrete sets. They also provided closed-form bridges for a restricted set of constraints such as product of intervals. These closed form bridges quickly become intractable as the constraint set gets more complex.

3 Background

3.1 Diffusion Models

Diffusion models (Sohl-Dickstein et al. 2015; Song and Ermon 2019; Song et al. 2021) are a class of generative models that learn to invert a stochastic process, known as the “forward process,” that gradually adds noise to samples from a data distribution $q_0(\mathbf{x}_0)$. The forward process is formulated as an SDE:

$$d\mathbf{x}_t = f(\mathbf{x}_t; t)dt + g(t)d\mathbf{w}, \quad \mathbf{x}_0 \sim q_0, \quad (1)$$

where f and g are drift and diffusion functions and \mathbf{w} is the standard Wiener process. The forward process is designed such that the SDE’s solution at time T is $q_T(\mathbf{x}_T) \approx \pi(\mathbf{x}_T)$ for some known π typically equal to $\mathcal{N}(\mathbf{0}, \mathbf{I})$.

Anderson (1982) showed that the path measure on the continuous trajectories following the forward process in Eq. (1) is identical to the one governed by the following time-reversed SDE:

$$d\mathbf{x}_t = [f(\mathbf{x}_t; t) - g^2(t)\nabla_x \log q_t(\mathbf{x}_t)] dt + g(t) d\bar{\mathbf{w}}, \quad (2)$$

where $\mathbf{x}_T \sim \pi$ and $\bar{\mathbf{w}}$ is the Wiener process when time flows backwards. Diffusion models learn the score function $s_\theta(\mathbf{x}_t; t) \approx \nabla_x \log q_t(\mathbf{x}_t)$ and approximate the reverse process in Eq. (2) by:

$$d\mathbf{x}_t = [f(\mathbf{x}_t; t) - g^2(t)s_\theta(\mathbf{x}_t; t)] dt + g(t) d\bar{\mathbf{w}}, \quad (3)$$

where $\mathbf{x}_T \sim \pi(\mathbf{x}_T)$. One can learn this score function by minimizing (Vincent 2011)

$$\mathbb{E}_{t, \mathbf{x}_0, \mathbf{x}_t} \left[\lambda(t) \|s_\theta(\mathbf{x}_t; t) - \nabla_x \log q(\mathbf{x}_t | \mathbf{x}_0)\|^2 \right], \quad (4)$$

where $\mathbf{x}_0 \sim q_0$ and $\mathbf{x}_t \sim q_t(\cdot | \mathbf{x}_0)$ in the expectation and $\lambda : [0, T] \rightarrow \mathbb{R}^+$ is a weighting function. Once trained, one can generate data from the model by sampling $\mathbf{x}_T \sim \pi$ and simulating the approximated reverse process in Eq. (3).

In the remainder of this paper we use q and p_θ respectively to denote the probability density function of the forward and reverse process. P_θ and Q denote the probability mass functions associated with p and q . This applies to the marginals, conditionals, and posteriors as well. Furthermore, to reduce notational clutter throughout the rest of the paper, we omit the explicit mention of θ and t when their meaning is evident from the context.

3.2 Constrained Generative Modeling

Constrained generative modeling tackles the problem of learning and generating from a distribution within a bounded domain $\Omega \subset \mathbb{R}^d$. This constrained domain Ω is either described by explicit constraints, e.g., a set of linear inequalities (Lou and Ermon 2023), or implicitly via binary functions taking $\mathbf{x} \in \mathbb{R}^d$ as inputs and indicating whether the constraints are satisfied (Naderiparizi et al. 2024). The training data in such problem is guaranteed to satisfy the given constraints. Formally, the dataset $\mathcal{D} = \{\mathbf{x}_0^i\}_{i=1}^N$ follows a data distribution q_0 such that $Q_0(\Omega) = 1$. The goal of constrained generative modeling is to approximate q_0 while being bounded to Ω . A maximum likelihood estimation objective for this problem is formulated as

$$\arg \min_{\theta} \text{KL}(q_0 || p_\theta) \quad \text{s.t. } P_\theta(\mathbf{x} \in \Omega) = 1. \quad (5)$$

3.3 Diffusion Bridges

Diffusion bridges for constrained generative modeling was introduced by Wu et al. (2022); Liu et al. (2023). It is a generalized framework of Brownian bridge processes and requires the constraint boundaries to be explicitly stated. For a constraint set $\Omega \subset \mathbb{R}^d$, a function $\mathcal{B}^\Omega(\mathbf{x}_t, t)$ defined on $\mathbb{R}^d \times \mathbb{R}^+$ is an Ω -bridge for the reverse process in Eq. (3) if the solutions of

$$d\mathbf{x}_t = [\nu_\theta(\mathbf{x}_t; t) + g^2(t)\mathcal{B}^\Omega(\mathbf{x}_t; t)]d\bar{t} + g(t)d\bar{\mathbf{w}}, \quad (6)$$

at final time $t = 0$ are guaranteed to be in Ω . Here $\mathbf{x}_T \sim \pi$ and $\nu_\theta(\mathbf{x}_t; t) := -f(\mathbf{x}_t; t) + g(t)^2 s_\theta(\mathbf{x}_t; t)$. Intuitively, the

extra injected term $\mathcal{B}^\Omega(\mathbf{x}_t; t)$ is quantified through external constraint functions, which drifts the particle to move towards the boundary and stay inside Ω . Wu et al. (2022) provides a set of sufficient conditions for Eq. (6) to constitute a valid diffusion bridge. However, these requirements on \mathcal{B}^Ω make them practically applicable only to a very limited set of problems. For completeness, we provide these requirements in the appendix¹. Liu et al. (2023) proposes a particular diffusion bridge and shows it satisfies the necessary conditions. It, however, requires closed form access an expectation that is only tractable for very simple constraints such as product of intervals. We discuss this more in Section 4.1. For completeness, we show the exact form of this bridge in the appendix. Furthermore, we show it corresponds to an optimal diffusion model trained on $\mathcal{U}(\Omega)$, a uniform distribution on Ω .

4 Methodology

In this section, we explain ‘‘Manually Bridged Models’’ (MBM), our approach to constrained generative modeling with manual bridges. MBM incorporates the complex constraint information into the model leading to a formulation similar to diffusion bridges. We show that such bridged models parameterize a family of sequence of distributions that only place mass on Ω at diffusion time $t = 0$. The models are then trained using the same objective as standard diffusion models.

4.1 Manual Bridges

Here, we first formally define the notion of manual bridges. Next, we explain how they are incorporated in diffusion models. Finally, we show how to combine multiple bridges to get a model that satisfies a set of given constraints.

Definition 1 (Ω -distance function). *Let $\ell^\Omega : \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}^{\geq 0}$ be a continuous and almost everywhere differentiable function with finite gradients w.r.t. x . We call ℓ^Ω an Ω -distance function when $\ell^\Omega(\mathbf{x}; 0) = 0$ if and only if $\mathbf{x} \in \Omega$.*

Definition 2 (Manually bridged model). *Given a diffusion model $s_\theta(\mathbf{x}_t, t)$, an Ω -distance function $\ell^\Omega(\mathbf{x}; t)$, and a C^1 -function $\gamma : [0, T] \rightarrow \mathbb{R}^+$ such that $\gamma(T) \approx 0$ and $\lim_{t \downarrow 0} \gamma(t) = \infty$, a **manual bridge** is defined as $\mathbf{b}^\Omega(\mathbf{x}; t) := -\gamma(t)\nabla_x \ell^\Omega(\mathbf{x}; t)$. A **manually bridged model** is defined as*

$$s_\theta^\Omega(\mathbf{x}; t, \gamma, \ell) := s_\theta(\mathbf{x}; t) + \mathbf{b}^\Omega(\mathbf{x}; t). \quad (7)$$

Intuitively, the added manual bridge term guides the distribution towards Ω . Manually bridged models correspond to score functions of distributions of the form $p^\Omega(\mathbf{x}; t) \propto p(\mathbf{x}; t) \exp(-\gamma(t)\ell^\Omega(\mathbf{x}; t))$. Since γ smoothly changes from zero at $t = T$ to infinity at $t = 0$, $p^\Omega(\mathbf{x}; t)$ smoothly interpolates between $p^\Omega(\mathbf{x}; T) = p(\mathbf{x}; T)$ at $t = T$ and $p^\Omega(\mathbf{x}; 0) \propto p(\mathbf{x}; 0)\mathbb{1}_\Omega(\mathbf{x})$.

Proposition 1. *Let $s_\theta(\mathbf{x}, t)$ be a score function corresponding to a density $p_\theta(\mathbf{x}, t)$. If $s_\theta(\mathbf{x}, t)$ is continuous in t and $p_\theta(\mathbf{x}, 0)$ is finite for $\mathbf{x} \notin \Omega$, then the manually bridged model in Definition 2 results in a sequence of distributions that only place mass on Ω at time $t = 0$.*

Proof of this proposition is provided in the appendix.

¹Please refer to the arXiv version of the paper for the appendix.

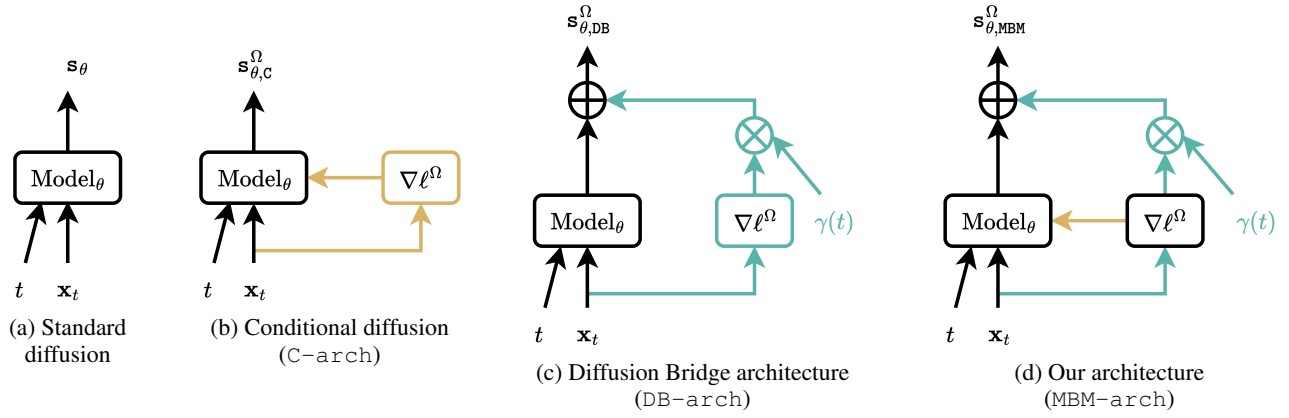


Figure 2: Score function architectural variants considered. The latter three use the same “manual bridge,” with the last notably including an additional path for the bridge function gradient not previously considered in the literature. Each diagram shows a single denoising step. In these diagrams the input t to the model is omitted for conciseness.

Proposition 2 (Combining Manual Bridges). *Let $\mathbf{b}^{\Omega_1}(\mathbf{x}; t) = -\gamma_1(t)\nabla_{\mathbf{x}}\ell^{\Omega_1}(\mathbf{x}; t)$ and $\mathbf{b}^{\Omega_2}(\mathbf{x}; t) = -\gamma_2(t)\nabla_{\mathbf{x}}\ell^{\Omega_2}(\mathbf{x}; t)$ be two manual bridges as defined in Definition 2. If $\bar{\Omega} := \Omega_1 \cap \Omega_2 \neq \emptyset$, the combined bridge $\mathbf{b}^{\bar{\Omega}}(\mathbf{x}; t) = \mathbf{b}^{\Omega_1}(\mathbf{x}; t) + \mathbf{b}^{\Omega_2}(\mathbf{x}; t)$ is also a manual bridge to $\bar{\Omega} \cap \Omega_2$. Therefore, the space of manual bridges is closed under addition.*

Proof. Without loss of generality, assume $\lim_{t \downarrow 0} \frac{\gamma_2(t)}{\gamma_1(t)} \neq 0$. Let $\ell^{\bar{\Omega}}(\mathbf{x}; t) := \ell^{\Omega_1}(\mathbf{x}; t) + \frac{\gamma_2(t)}{\gamma_1(t)}\ell^{\Omega_2}(\mathbf{x}; t)$. Since all functions are continuous in t , both ℓ^{Ω_1} and ℓ^{Ω_2} are distance functions and $\lim_{t \downarrow 0} \frac{\gamma_2(t)}{\gamma_1(t)} \neq 0$, $\ell^{\bar{\Omega}}(\mathbf{x}; 0)$ is zero if and only if both $\ell^{\Omega_1}(\mathbf{x}; 0)$ and $\ell^{\Omega_2}(\mathbf{x}; 0)$ are zero. Therefore, $\ell^{\bar{\Omega}}$ is a distance function to $\bar{\Omega}$. Further, by definition $\lim_{t \downarrow 0} \gamma_1(t) = \infty$. Therefore, $-\gamma_1(t)\nabla_{\mathbf{x}}\ell^{\bar{\Omega}}(\mathbf{x}; t) = \mathbf{b}^{\Omega_1}(\mathbf{x}; t) + \mathbf{b}^{\Omega_2}(\mathbf{x}; t) = \mathbf{b}^{\bar{\Omega}}(\mathbf{x}; t)$ is a manual bridge to $\bar{\Omega}$. \square

Generalizing Proposition 1, one can combine multiple bridges by $\mathbf{b}^{\cap_{i=1}^N \{\Omega_i\}} = \sum_{i=1}^N \mathbf{b}_i^{\Omega_i}(\mathbf{x}; t)$.

The objective function of MBM is the denoising loss with this particular parameterization of the score estimator model

$$\mathbb{E}_{t, \mathbf{x}_0, \mathbf{x}_t} \left[\lambda(t) \|\mathfrak{s}_{\theta}(\mathbf{x}_t; t, \gamma, \ell) - \nabla_{\mathbf{x}} \log q_t(\mathbf{x}_t | \mathbf{x}_0)\|^2 \right] \quad (8)$$

where $\mathbf{x}_0 \sim q_0^{\Omega}$, $\mathbf{x}_t \sim q_t(\cdot | \mathbf{x}_0)$ in the expectation.

Connection to diffusion bridges Note that our manually bridged models give rise to reverse SDEs similar to that of diffusion bridges in Eq. (6). To see this, we can plug the manually bridged score function in Eq. (3):

$$d\mathbf{x}_t = [-\nu(\mathbf{x}_t; t) - g^2(t)\mathbf{b}^{\Omega}(\mathbf{x}_t; t)] dt + g(t) d\bar{\mathbf{w}}, \quad (9)$$

where $\nu(\mathbf{x}_t; t) \equiv -f(\mathbf{x}_t; t) + g^2(t)\mathfrak{s}_{\theta}(\mathbf{x}_t; t)$. This is equivalent to Eq. (6) once $\mathcal{B}^{\Omega}(\mathbf{x}_t; t) \equiv \mathbf{b}^{\Omega}(\mathbf{x}_t; t)$. However, in order to guarantee that solutions of this SDE lie in Ω , the bridge function $\mathbf{b}^{\Omega}(\mathbf{x}, t) = -\gamma(t)\nabla_{\mathbf{x}}\ell^{\Omega}(\mathbf{x}; t)$ must satisfy requirements such as an expected Polyak-Lojasiewicz condition: $\mathbb{E}_{t, \mathbf{x}_t \sim p_t^{\Omega}} [\ell^{\Omega}(\mathbf{x}_t; t)] \leq \mathbb{E}_{t, \mathbf{x}_t \sim p_t^{\Omega}} \left[\|\nabla_{\mathbf{x}}\ell^{\Omega}(\mathbf{x}_t; t)\|^2 \right]$.

This greatly restricts the set of allowed Ω -distance functions. Furthermore, it restricts generality of combination of bridges. Our manually bridged models therefore are not guaranteed to have SDE solutions in Ω . However, as shown in the previous section, they still represent score functions of distributions that converge to one constrained to Ω . Further, since we train the model to approximate the reverse process in Eq. (2), an SDE with a trained model is likely to converge to this terminal distribution as well. As we show later in the Experiment section, manual bridges empirically perform similarly to diffusion bridges.

4.2 Architecture

In practice, the typical diffusion bridge method of incorporating manual bridges as in Eq. (7) leads to poor training. This is because the added bridge term increases variance of the loss. As illustrated in Fig. 2, we consider three different types of incorporating the bridge information into the model.

1. **Conditional diffusion (C-arch)**: modifies the score network to take (a re-weighted version of) the bridge as an additional input. This effectively becomes a conditional diffusion model, and those have been proven to be highly effective in, e.g., large-scale text-conditional image generation tasks. Note that this mechanism merely incorporates the information from the distance function to the score network. As such, it is not a bridged model and does not provide guarantees regarding the constraints. However, its training is stable.

$$\mathfrak{s}_{\theta, \text{C}}^{\Omega}(\mathbf{x}_t; t, \gamma, \ell) := \mathfrak{s}_{\theta}(\mathbf{x}_t; t, \nabla_{\mathbf{x}}\ell^{\Omega}(\mathbf{x}_t)) \quad (10)$$

2. **Diffusion Bridges (DB-arch)**: offsets the score function using the bridge function. This is the mechanism used in diffusion bridge (Wu et al. 2022; Liu et al. 2023). This mechanism is effective at constraining the distribution but results in unstable training.

$$\mathfrak{s}_{\theta, \text{DB}}^{\Omega}(\mathbf{x}_t; t, \gamma, \ell) := \mathfrak{s}_{\theta}(\mathbf{x}_t; t) + \mathbf{b}^{\Omega}(\mathbf{x}_t; t) \quad (11)$$

3. **Manually Bridged Models (MBM-arch; ours)**: combines the above two mechanisms. We have found that additionally providing (a re-weighted version of) the bridge

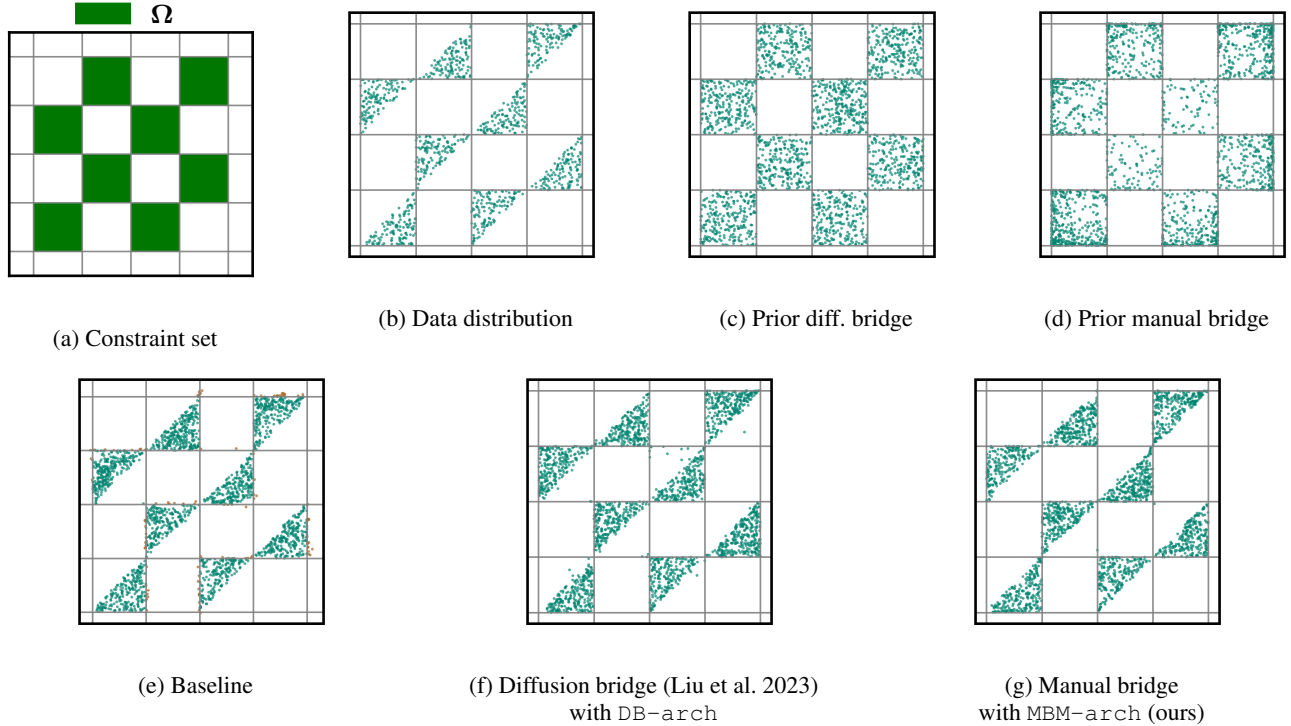


Figure 3: Visualization of the checkerboard constraint experiment results. The problem is constrained on a checkerboard pattern and the data has a uniform distribution over triangles within the checkerboard, shown in (b). Invalid samples are shown in brown. (c, d) show the diffusion bridge and manually bridged models without a trained diffusion model. This is effectively the prior distributions in these models. As shown on the second row, bridged models do not produce invalid samples. Further, the manually bridged model gives comparable samples to the diffusion bridges (Liu et al. 2023). Finally, incorporating the bridges in our proposed way (labelled with MBM) even improves the diffusion bridge models.

to the score network stabilizes training resulting in both a good model fit and good constraint satisfaction.

$$s_{\theta, \text{MBM}}^{\Omega}(\mathbf{x}_t; t, \gamma, \ell) := s_{\theta}(\mathbf{x}_t; t, \nabla_x \ell^{\Omega}(\mathbf{x}_t)) + \mathbf{b}^{\Omega}(\mathbf{x}_t; t) \quad (12)$$

While any of the three mechanisms above can be used to incorporate either manual or diffusion bridges, our proposed method, referred to as MBM, involves *manual bridges* applied through the `MBM-arch` mechanism.

5 Experiments

We demonstrate MBM on a simple 2D synthetic dataset and a traffic scenario generation experiment with collision and offroad avoidance. Additionally, we include an image watermarking experiment in the appendix. We release the source code implementing MBM together with the 2D synthetic and image watermarking experiments.

5.1 Checkerboard Constraint Experiment

Setup We first experiment with a simple 2D dataset with a checkerboard constraint (Fig. 3a). The data distribution q_0 is a mixture of uniform triangles enclosed in the allowed checkerboard area. Our dataset consists of 1,000 samples from this data distribution (Fig. 3b). We report results for

a baseline diffusion model in Fig. 3e. This is the standard diffusion as explained in Background section. Samples from this baseline model include quite a few constraint violations.

Diffusion bridge baseline As stated before, diffusion bridges (Liu et al. 2023) are only tractable on very simple constraints. The checkerboard pattern of this experiment is one example. As stated before and shown in the appendix, the diffusion bridge for this problem is equivalent an exact diffusion model for uniformly distributed data in the constraint set Ω . We empirically verify this in Fig. 3c, where we construct an SDE with only the bridge. This is equivalent to Eq. (6) without the ν term. As such, the results in Fig. 3c match $\mathcal{U}(\Omega)$. Consequently, we modify the typical checkerboard data distribution to be non-uniform in the constraint set. The “squares” in Fig. 3a are constraints. The data just happens to lie in half of the constraint set. This modification is necessary to highlight model training within the diffusion bridge framework – the bridge to the uniform checkerboard is literally the generative model for the typical checkerboard dataset.

Following Liu et al. (2023), we implement a model with `DB-arch` (see Fig. 2c). Once this bridged model is trained, it will match the data distribution while staying within the constraints (Fig. 3f).




Figure 4: Checkerboard constraint experiment results. Solid, dashed and dotted lines respectively correspond to manual bridge, diffusion bridge and baseline models. Different colors represent different mechanisms of incorporating bridges as shown in Fig. 2. The shaded area shows standard deviation of the metrics over three models trained with different random seeds.

Manual bridges For manual bridges we use $\ell^\Omega(\mathbf{x}; t) := \min_{\mathbf{y} \in \Omega} \|\mathbf{x} - \mathbf{y}\|_2^2$ as the distance function and $\gamma(t) := \frac{1}{\sigma^2(t)}$ where $\sigma(t)$ is the total amount of noise added at time t of the diffusion process. We demonstrate in Fig. 3d that our manually bridged model guides the generation differently from diffusion models. Similar to Fig. 3c, this is resulted from an SDE with only the bridge. Fig. 3g shows that a Manually Bridged Model implemented in our proposed mechanism MBM-arch (Fig. 2d) produces comparable results to a diffusion bridge model.

Quantitative results We report our quantitative results in Fig. 4. We report evidence lower bound (ELBO) and infraction rate to respectively measure distribution match and constraint satisfaction rate. We compare standard diffusion model with both types of bridged models (diffusion and manual bridges) as well as different mechanisms for incorporating bridges.

Although passing the bridge information to the model to condition on (as done in conditional diffusion models, labeled C-arch in the plot) reduces infraction rate, it fails to achieve zero infraction rate. With bridges incorporated, both diffusion bridges or manual bridges brings the infraction rate down to (almost) zero with both DB-arch and MBM-arch mechanism. However, the ELBO plot shows that the DB-arch models are much slower to converge. It takes the DB-arch models around 250k iterations to achieve their maximum validation ELBO while the other models achieve the maximum at around 20k-30k iterations and start over-fitting after. It demonstrates that our proposed MBM-arch mechanism or incorporating bridges is even effective for implementing diffusion bridges of Liu et al. (2023). This plot also verifies that while diffusion bridges have a more solid grounding with guaranteed convergence to an Ω -constrained distribution, our proposed manual bridges can successfully achieve similar results. This makes manual bridges a much simpler and more widely applicable alternatives to diffusion bridges.

5.2 Traffic Scene Generation

We continue to the experiment of generating realistic traffic scenes for placing various number of arbitrary-sized vehicles on different maps. Since a major part of traffic simulation tasks focuses on learning driving behaviors and predicting vehicle trajectories, the task of generating a realistic initial traffic scene before the traffic simulations is equally significant. A common practice to determine whether a generated traffic scene sample is valid, is to check if any vehicle is outside the drivable region (called “offroad”), or if vehicles overlap one another (called “collision”). To deal with the invalid samples being generated, the previous methods for this problem usually simply discard them (Tan et al. 2021; Zwartsenberg et al. 2023). This requires repeated sampling from the model and is computationally inefficient. A more recent method updates the model post-hoc by guiding it via a separately trained model to lower the infraction rate (Naderiparizi et al. 2024). While it lowers the infraction rate, it is still far from the ideal of zero infraction and requires extra computing resources. MBM is constructed to achieve (almost) zero infraction rate, hence eliminating the requirement of repeated sampling or post-hoc modifications.

The problem in this experiment is to generate up to 25 vehicles on a bird’s-eye view image of a road from one of 70 locations in the dataset. To condition on the road, we extract its image features using a convolutional neural network (CNN) encoder (LeCun et al. 1998) and pass them to the score network. The CNN is trained jointly with the score network. Each vehicle is represented by its center position, length, width, heading direction and velocity, which makes 7 dimensions per vehicle. Note that vehicles are considered jointly, and their interactions are modeled in full, so that the overall dimensionality of the problem is $N \times 7$, with N the total number of vehicles. As alluded to earlier, there are two constraints in this experiment: collision and offroad. For the collision constraint we compute the area of intersection between two vehicles as the Ω_c -distance function. For offroad, we consider a car that has all four wheels off the drivable area as offroad. Therefore, for the Ω_o -distance, we compute the shortest squared distance to the constraint set

Method	Collision (%) ↓	Offroad (%) ↓	Infraction (%) ↓	r-ELBO ↑
Standard diffusion (Karras et al. 2022)	17.40 ± 0.01	7.33 ± 0.18	23.00 ± 0.14	-0.94 ± 0.01
Guided diffusion	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	-1.53 ± 0.11
Manual bridge with C-arch	19.47 ± 0.20	6.87 ± 0.03	24.52 ± 0.23	-0.94 ± 0.00
Manual bridge with DB-arch	0.24 ± 0.03	0.00 ± 0.00	0.25 ± 0.03	-1.20 ± 0.01
Manual bridge with MBM-arch (ours)	0.10 ± 0.00	0.00 ± 0.00	0.10 ± 0.00	-0.95 ± 0.00

Table 1: Results for traffic scene generation. We compare our model MBM-arch against its alternative architectures C-arch and DB-arch, a standard diffusion model, and a guided variant using the bridge term as a guidance signal.

of each of the four corners of the vehicle and take the minimum. The γ functions for collision and offroad are respectively $\gamma_c(t) = \frac{1}{10\sigma^2(t)}$ and $\gamma_o(t) = \frac{1}{100\sigma^2(t)}$. The complex constraints of this problem renders diffusion bridges (Liu et al. 2023) and most of the existing diffusion-based constrained generative modeling approaches (Lou and Ermon 2023; Fishman et al. 2023; Liu et al. 2024; Fishman et al. 2024) inapplicable.

We implement a diffusion model based on EDM (Karras et al. 2022). Our model has a transformer-based architecture (Vaswani et al. 2017) composed of self-attention and cross-attention layers modeling the interactions between vehicles and between vehicles and the road map. For evaluation, we report collision, offroad, and overall infraction rate to measure constraint satisfaction and the validation loss which corresponds to a reweighted ELBO (r-ELBO) to evaluate the distribution match. The standard diffusion serves as the baseline. We also include a guided diffusion variant that guides this pre-trained diffusion with our manual bridges without further training. We compare these with different mechanisms for incorporating manual bridges as described in Section 4.2. Table 1 reports our results for this experiment. Adding the manual bridge as an offset to the model, DB-arch, strongly reduces infractions but it deteriorates model quality significantly as evidenced by the degraded r-ELBO. While guided diffusion achieves zero infraction, its significantly worse r-ELBO suggests a strong distribution shift. MBM-arch ultimately produces a model with close to zero infractions, and importantly recovers performance in resulting r-ELBO. In summary, MBM with our proposed MBM-arch mechanism is able to effectively nearly solve the problem of infraction free traffic scene generation.

6 Discussion

We introduced manually bridged models, a family of score functions corresponding to distributions that anneal between a known base distribution $\pi(x_T)$ to an Ω -constrained distribution at time $t = 0$. Manual bridges can, for instance, be used to impose prior knowledge in the form of safety constraints. We showed

- one can construct such manual bridges having access to a distance function to the constraint set,
- such bridges can be combined to get a multiply-constrained model,
- learning diffusion processes on top of manual bridges can fit observational data while maintaining safety con-

straints and overcoming potentially unwelcome biases imposed by suboptimal manual bridges,

- architecture matters when imposing bridges, particularly for effective learning. We propose a combination of (i) including the manual bridge in the score function in form of addition and (ii) conditioning the score network on the bridge. We empirically demonstrate that this combination is crucial in training of bridged models.

Limitations While MBM is much more flexible and widely applicable to constrained generative modeling problems, it still relies on differentiable distance functions, which may not be straightforward to implement for certain problems. Further, it is quite likely that there is a greater family of manual bridges that can be described and composed than we have established. This includes both the bridge primitives and the methods of combining bridge functions.

The proper functioning of manual bridges still requires a fair amount of hyperparameter fiddling, particularly integration schedule, minimum noise level, and bridge scaling. This is related to the gap between numerical and exact integration, however, it bears mentioning here as the asymptotic scaling of bridges and their gradients is extreme; sufficiently so to be fiddly. The specific inspiration for the architecture innovation we introduced came from this low noise asymptotic scaling difficulty. Even with our architectural innovation and manual bridge conditions, applying this method to new domains is not automatic. In this sense it is not terribly different to prior specification in Bayesian models, except here the error mode looks more like getting garbage samples that are “safe.”

Future Work There is a great deal of theory work to be done, particularly with respect to how manual bridges fit into the diffusion bridge formalism. We would ideally like to identify the mathematical conditions on manual bridges and their combinations that would give rise to formal guarantees of constraint satisfaction. Thus far, it appears that we need a generalization of Ito’s lemma that relaxes the primary twice differentiable condition to twice differentiable almost everywhere or even looser.

While solving the constrained, conditional agent placement distribution problem is of value; it is obvious that our approach should and could be extended to trajectory space as well. Static actor placements are, after all, merely short trajectories. Our initial attempts at this were promising, but, work remains to be done.

Acknowledgments

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), the Canada CIFAR AI Chairs Program, Inverted AI, MITACS, the Department of Energy through Lawrence Berkeley National Laboratory, and Google. This research was enabled in part by technical support and computational resources provided by the Digital Research Alliance of Canada Compute Canada (alliancecan.ca), the Advanced Research Computing at the University of British Columbia (arc.ubc.ca), and Amazon.

References

- Anderson, B. D. 1982. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3): 313–326.
- Azamfirei, R.; Kudchadkar, S. R.; and Fackler, J. 2023. Large language models and the perils of their hallucinations. *Critical Care*, 27(1): 120.
- Chang, W.-J.; Pittaluga, F.; Tomizuka, M.; Zhan, W.; and Chandraker, M. 2024. SAFE-SIM: Safety-Critical Closed-Loop Traffic Simulation with Diffusion-Controllable Adversaries. In *European Conference on Computer Vision*, 242–258. Springer.
- Chen, T.; Liu, G.-H.; and Theodorou, E. 2022. Likelihood Training of Schrödinger Bridge using Forward-Backward SDEs Theory. In *International Conference on Learning Representations*.
- Christopher, J. K.; Baek, S.; and Fioretto, F. 2024. Constrained Synthesis with Projected Diffusion Models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- De Bortoli, V.; Thornton, J.; Heng, J.; and Doucet, A. 2021. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34: 17695–17709.
- Fishman, N.; Klarner, L.; De Bortoli, V.; Mathieu, E.; and Hutchinson, M. J. 2023. Diffusion Models for Constrained Domains. *Transactions on Machine Learning Research*.
- Fishman, N.; Klarner, L.; Mathieu, E.; Hutchinson, M.; and De Bortoli, V. 2024. Metropolis Sampling for Constrained Diffusion Models. *Advances in Neural Information Processing Systems*, 36.
- Hanneke, S.; Kalai, A. T.; Kamath, G.; and Tzamos, C. 2018. Actively avoiding nonsense in generative models. In *Conference On Learning Theory*, 209–227. PMLR.
- Heng, J.; De Bortoli, V.; Doucet, A.; and Thornton, J. 2021. Simulating diffusion bridges with score matching. *arXiv preprint arXiv:2111.07243*.
- Huang, L.; Yu, W.; Ma, W.; Zhong, W.; Feng, Z.; Wang, H.; Chen, Q.; Peng, W.; Feng, X.; Qin, B.; et al. 2023. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *arXiv preprint arXiv:2311.05232*.
- Huang, Z.; Zhang, Z.; Vaidya, A.; Chen, Y.; Lv, C.; and Fisac, J. F. 2024. Versatile Scene-Consistent Traffic Scenario Generation as Optimization with Diffusion. *arXiv preprint arXiv:2404.02524*.
- Janner, M.; Du, Y.; Tenenbaum, J.; and Levine, S. 2022. Planning with Diffusion for Flexible Behavior Synthesis. In *International Conference on Machine Learning*, 9902–9915. PMLR.
- Jiang, C. M.; Bai, Y.; Cornman, A.; Davis, C.; Huang, X.; Jeon, H.; Kulshrestha, S.; Lambert, J. W.; Li, S.; Zhou, X.; Fuertes, C.; Yuan, C.; Tan, M.; Zhou, Y.; and Anguelov, D. 2024. SceneDiffuser: Efficient and Controllable Driving Simulation Initialization and Rollout. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Karras, T.; Aittala, M.; Aila, T.; and Laine, S. 2022. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems*, 35: 26565–26577.
- Knight, J. C. 2002. Safety critical systems: challenges and directions. In *Proceedings of the 24th international conference on software engineering*, 547–550.
- Kong, Z.; and Chaudhuri, K. 2023. Data redaction from pre-trained gans. In *2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, 638–677. IEEE.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.
- Liu, G.-H.; Chen, T.; Theodorou, E.; and Tao, M. 2024. Mirror diffusion models for constrained and watermarked generation. *Advances in Neural Information Processing Systems*, 36.
- Liu, X.; Wu, L.; Ye, M.; and qiang liu. 2023. Learning Diffusion Bridges on Constrained Domains. In *The Eleventh International Conference on Learning Representations*.
- Lou, A.; and Ermon, S. 2023. Reflected diffusion models. In *International Conference on Machine Learning*, 22675–22701. PMLR.
- Naderiparizi, S.; Liang, X.; Cohan, S.; Zwartsenberg, B.; and Wood, F. 2024. Don’t be so Negative! Score-based Generative Modeling with Oracle-assisted Guidance. In *Forty-first International Conference on Machine Learning*.
- Niedoba, M.; Lavington, J.; Liu, Y.; Lioutas, V.; Sefas, J.; Liang, X.; Green, D.; Dabiri, S.; Zwartsenberg, B.; Scibior, A.; et al. 2024. A Diffusion-Model of Joint Interactive Navigation. *Advances in Neural Information Processing Systems*, 36.
- Schauer, M.; van der Meulen, F.; and van Zanten, H. 2017. Guided proposals for simulating multi-dimensional diffusion bridges. *Bernoulli*, 23(4A): 2917 – 2950.
- Schulman, J.; Duan, Y.; Ho, J.; Lee, A.; Awwal, I.; Bradlow, H.; Pan, J.; Patil, S.; Goldberg, K.; and Abbeel, P. 2014. Motion planning with sequential convex optimization and convex collision checking. *The International Journal of Robotics Research*, 33(9): 1251–1270.

- Shi, Y.; De Bortoli, V.; Campbell, A.; and Doucet, A. 2024. Diffusion Schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 36.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, 2256–2265. PMLR.
- Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*.
- Tan, S.; Wong, K.; Wang, S.; Manivasagam, S.; Ren, M.; and Urtasun, R. 2021. Scenegen: Learning to generate realistic traffic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 892–901.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Vincent, P. 2011. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7): 1661–1674.
- Wang, G.; Jiao, Y.; Xu, Q.; Wang, Y.; and Yang, C. 2021. Deep generative learning via schrödinger bridge. In *International conference on machine learning*, 10794–10804. PMLR.
- Welling, M.; and Teh, Y. W. 2011. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, 681–688. Citeseer.
- Wu, L.; Gong, C.; Liu, X.; Ye, M.; and Liu, Q. 2022. Diffusion-based molecule generation with informative prior bridges. *Advances in Neural Information Processing Systems*, 35: 36533–36545.
- Zhong, Z.; Rempe, D.; Xu, D.; Chen, Y.; Veer, S.; Che, T.; Ray, B.; and Pavone, M. 2023. Guided conditional diffusion for controllable traffic simulation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 3560–3566. IEEE.
- Zhou, L.; Lou, A.; Khanna, S.; and Ermon, S. 2024. Denoising Diffusion Bridge Models. In *The Twelfth International Conference on Learning Representations*.
- Zwartsenberg, B.; Ścibior, A.; Niedoba, M.; Lioutas, V.; Sefas, J.; Liu, Y.; Dabiri, S.; Lavington, J. W.; Campbell, T.; and Wood, F. 2023. Conditional Permutation Invariant Flows. *Transactions on Machine Learning Research*.