

# To the Noise and Back: Diffusion for Shared Autonomy

Takuma Yoneda<sup>1</sup>, Luzhe Sun<sup>2</sup>, Ge Yang<sup>3,4</sup>, Bradly Stadie<sup>5</sup>, and Matthew R. Walter<sup>1</sup>

<sup>1</sup>Toyota Technological Institute at Chicago (TTIC)

<sup>2</sup>Department of Computer Science, University of Chicago

<sup>3</sup>Institute of Artificial Intelligence and Fundamental Interactions (IAIFI)

<sup>4</sup>Computer Science and Artificial Intelligence Laboratory (CSAIL), MIT

<sup>5</sup>Department of Statistics and Data Science, Northwestern University

**Abstract**—Shared autonomy is an operational concept in which a user and an autonomous agent collaboratively control a robotic system. It provides a number of advantages over the extremes of full-teleoperation and full-autonomy in many settings. Traditional approaches to shared autonomy rely on knowledge of the environment dynamics, a discrete space of user goals that is known a priori, or knowledge of the user’s policy—assumptions that are unrealistic in many domains. Recent works relax some of these assumptions by formulating shared autonomy with model-free deep reinforcement learning (RL). In particular, they no longer need knowledge of the goal space (e.g., that the goals are discrete or constrained) or environment dynamics. However, they need knowledge of a task-specific reward function to train the policy. Unfortunately, such reward specification can be a difficult and brittle process. On top of that, the formulations inherently rely on human-in-the-loop training, and that necessitates them to prepare a policy that mimics users’ behavior. In this paper, we present a new approach to shared autonomy that employs a modulation of the forward and reverse diffusion process of diffusion models. Our approach does not assume known environment dynamics or the space of user goals, and in contrast to previous work, it does not require any reward feedback, nor does it require access to the user’s policy during training. Instead, our framework learns a distribution over a space of desired behaviors. It then employs a diffusion model to translate the user’s actions to a sample from this distribution. Crucially, we show that it is possible to carry out this process in a manner that preserves the user’s control authority. We evaluate our framework on a series of challenging continuous control tasks, and analyze its ability to effectively correct user actions while maintaining their autonomy.

## I. INTRODUCTION

Contemporary robots primarily operate in one of two different ways—full teleoperation or full autonomy. Teleoperation is common in unstructured environments (e.g., underwater), where the proficiency with which robots are able to understand their surroundings is insufficient for fully autonomous robots to operate reliably. However, direct teleoperation requires users to interpret the robot’s environment observations while simultaneously controlling its low-level actions, a responsibility that is particularly challenging for highly dynamic tasks. This operational gap motivates a setting in which a human and an autonomous agent collaborate and share control of the robot.

Shared autonomy [2] is a framework in which a human user (also referred to as the *pilot*) performs a task with an assistance of an autonomous agent (also referred to as the




Fig. 1: Our framework utilizes a diffusion model to adapt a user’s action (red) to those from a demonstration distribution (green) in a manner (blue) that balances a user’s desire to maintain control authority with the benefits (e.g., safety) of conforming to the desired (demonstration) distribution. Without knowledge of the user’s specific goal (e.g., the landing location), the demonstration distribution reflects different goals that the demonstration trajectories previously reached.

*copilot*) [21, 41, 3, 13, 14, 22, 5]. The role of the agent is to complement the control authority of the user, whether to improve the robot’s performance on the current task or to encourage/ensure safe behavior. An important consideration when providing assistance via shared autonomy is the degree to which the agent balances the user’s preference for maintaining control authority (i.e., the *fidelity* of the assisted behavior relative to the user’s actions), and the potential benefits of endowing more control to the agent (i.e., the *conformity* of the assisted behavior to that of an autonomous agent).

A core difficulty of shared autonomy lies in the fact that the user’s goal (intent) is typically not known. Many approaches to shared autonomy assume that there is a fixed, discrete set of candidate goals and seek to infer the user’s specific goal at test-time based on observations, including the user’s control input [34, 28, 37, 24, 14]. Such assumptions may be reasonable in structured environments (e.g., in the context of a manipulation task when there is a small number of graspable objects sitting on a table). However, they can be limiting in unstructured environments that lack well-defined goals or that have a very large set of potential goals.

Bootstrapped by function approximation with neural-networks, recent deep reinforcement learning (RL) algorithms

seek to learn assistive policies without assumptions on the knowledge of the goal space, or the assumption that the environment dynamics are known. Reddy et al. [39] propose a deep RL approach to shared autonomy for domains with discrete actions, nominally relying on reward feedback from the user as an alternative to assuming that the goal space is known. In an effort to balance the user’s control authority with task performance, the assistant chooses the action most similar to that of the user while also satisfying a state-action value constraint. Schaff and Walter [43] treat the copilot as providing a residual that is added to the user’s actions to correct for unsafe behavior. They train their model to minimize the norm of the residual, subject to a goal-agnostic reward constraint that represents safe behavior.

These methods treat the pilot as a part of the environment, using an augmented state that includes the user’s action. Framing the problem in this way has a clear and significant advantage—it enables the direct utilization of the modern suite of tools for deep RL. However, these methods have two notable limitations. First, they nominally require human-in-the-loop interaction during training in order to generate user actions while learning the assistant’s policy. Since the sample complexity of deep RL makes this interaction intractable, these methods replace the human with a *surrogate policy*. If this surrogate is misspecified or invalid, this approach can lead to copilots that are incompatible with actual human pilots [43]. Second, these methods require access to task-specific reward during training, which may be difficult to obtain in practice.

In light of these limitations, we propose a model-free approach to shared autonomy that interpolates between the user’s action and an action sampled from a generative model that provides a distribution over desired behavior (Fig. 1).<sup>1</sup> Our approach has the distinct advantage that it does not require knowledge of or access to the user’s policy or any reward feedback during training, thus eliminating the need for reward engineering. Instead, our training process, which involves learning the generative model, only requires access to trajectories that are representative of desired behavior.

The generative model that underlies our approach is a diffusion model [49, 45, 46, 25], which has proven highly effective for complex generation tasks including image synthesis [12, 42]. Diffusion models consist of two key processes: the *forward process* and the *reverse process*. The forward process iteratively adds Gaussian noise to the input with an increasing noise scale, while the reverse process is trained to iteratively denoise a noisy input in order to arrive at the target distribution. As part of this denoising process, the model produces the gradient direction in which the likelihood of its input increases under the target distribution. Once the model is trained, generating a sample from the (unknown) target distribution involves running the reverse process on a sample drawn from a zero-mean isotropic Gaussian distribution.

As we will see in the following sections, a direct use of diffusion models for shared autonomy ends up in generating an

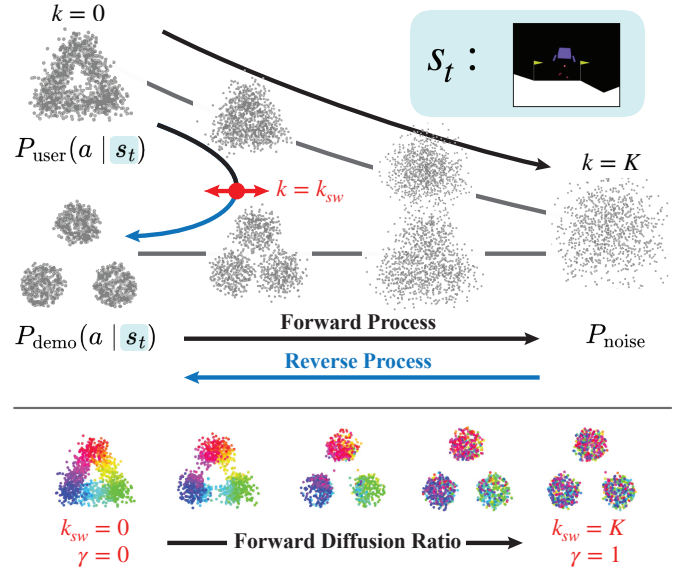


Fig. 2: (Top) A visualization of diffusion processes of action distributions at state  $s_t$ . The black arrow at the top shows the forward diffusion of a source user distribution  $P_{\text{user}}$  and the blue arrow below shows the reverse diffusion to the target demonstration distribution  $P_{\text{demo}}$ . We switch these two processes in the intermediate step  $k = k_{\text{sw}}$  to achieve partial forward and reverse diffusion shown in the black and blue arrow on the left. (Bottom) The result of forward and reverse diffusion for different switching times  $k_{\text{sw}}$ , where standard reverse diffusion process corresponds to  $k_{\text{sw}} = K$ .

action that ignores user’s intent (i.e., low *fidelity* to user intent), even though the action would be consistent with the desired behaviors (i.e., high *conformity* to the target behaviors). To address this, we propose a new algorithm that controls the effect of the forward and reverse process through a *forward diffusion ratio*  $\gamma$  that regulates the balance between the fidelity and the conformity of the generated actions. The forward diffusion ratio provides a formal bound on the extent to which the copilot’s action deviates from that of the user.

We evaluate our shared autonomy algorithm using a series of continuous control tasks. In each case, we demonstrate that our algorithm significantly improves the performance of a variety of different pilots, and we analyze the effects of a range of different forward diffusion ratios.

## II. METHOD

An integral part of our shared autonomy approach is the utilization of diffusion models as a generative model that serves to correct the actions of the user. As such, we begin this section with an in-depth review of diffusion models.

### A. Background on diffusion models

We first present the mathematical formulation of diffusion models, with a particular emphasis on the denoising diffusion probabilistic model (DDPM) [25], which we employ.

<sup>1</sup>For video and code, see <https://diffusion-for-shared-autonomy.github.io>.

A probabilistic diffusion model [45] is a type of generative model that is characterized by two processes (Fig. 2): a *forward diffusion process* and a *reverse diffusion process*. The forward diffusion process is an iterative first-order Markov chain that adds noise to a sample from a data distribution  $\mathbf{x}_0 \sim q(\mathbf{x}_0)$ . Assuming a predefined sequential noise schedule of  $K$  steps,  $\beta_1, \dots, \beta_K$  and  $\alpha_k := 1 - \beta_k$ , a single step in the forward process operates as

$$\mathbf{x}_k = \sqrt{\alpha_k} \mathbf{x}_{k-1} + \sqrt{1 - \alpha_k} \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (1)$$

Applying this recursively, multiple steps of forward process can be written in a closed form as

$$\mathbf{x}_k = \sqrt{\bar{\alpha}_k} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_k} \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (2)$$

where  $\bar{\alpha}_k := \prod_{s=1}^k \alpha_s$ . Equivalently,

$$q(\mathbf{x}_k | \mathbf{x}_{k-1}) = \mathcal{N}(\mathbf{x}_k; \sqrt{\alpha_k} \mathbf{x}_0, (1 - \alpha_k) \mathbf{I}) \quad (3a)$$

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

Meanwhile, the *reverse diffusion process* is a Markov chain that iteratively denoises a noisy input, starting from  $\mathbf{x}_K \sim N(\mathbf{0}, \mathbf{I})$ . The conditional distribution over the output at each step  $q(\mathbf{x}_{k-1} | \mathbf{x}_k)$  is Gaussian. However, computing these intermediate steps depends upon the entire data distribution and is generally not tractable in closed-form. Thus, we train a model that approximates it as  $p_\theta(\mathbf{x}_{k-1} | \mathbf{x}_k) = \mathcal{N}(\mathbf{x}_{k-1}; \boldsymbol{\mu}_\theta(\mathbf{x}_k, k), \sigma_k^2 \mathbf{I})$  to run the reverse process. The probability  $q(\mathbf{x}_{k-1} | \mathbf{x}_k)$  that we want to model becomes tractable when conditioned on  $\mathbf{x}_0$ . By Bayes' rule, we have

$$q(\mathbf{x}_{k-1} | \mathbf{x}_k, \mathbf{x}_0) = q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{x}_0) \frac{q(\mathbf{x}_{k-1} | \mathbf{x}_0)}{q(\mathbf{x}_k | \mathbf{x}_0)} \quad (4a)$$

$$= q(\mathbf{x}_k | \mathbf{x}_{k-1}) \frac{q(\mathbf{x}_{k-1} | \mathbf{x}_0)}{q(\mathbf{x}_k | \mathbf{x}_0)} \quad (4b)$$

$$= \mathcal{N}(\mathbf{x}_{k-1}; \tilde{\boldsymbol{\mu}}(\mathbf{x}_k, \mathbf{x}_0), \tilde{\beta}_k \mathbf{I}) \quad (4c)$$

Crucially, Equation 4b only contains known terms from the Gaussian distributions in Equations 3a and 3b. In the last expression, we have dropped the exact forms of mean and covariance for notational brevity.

Intuitively, we can train the model  $p_\theta(\mathbf{x}_{k-1} | \mathbf{x}_k)$  by minimizing its divergence from  $q(\mathbf{x}_{k-1} | \mathbf{x}_k, \mathbf{x}_0)$  over samples from the data distribution  $q(\mathbf{x}_0)$ . Utilizing the evidence lower bound of  $\mathbb{E}_q[-\log p_\theta(\mathbf{x}_0)]$ , we can formulate the loss as:

$$\mathcal{L} := \mathbb{E}_q \left[ \sum_{k=2}^{K-1} L_k - \log p_\theta(\mathbf{x}_0 | \mathbf{x}_1) + C \right], \quad (5)$$

$$L_k = D_{\text{KL}}(q(\mathbf{x}_{k-1} | \mathbf{x}_k, \mathbf{x}_0) \| p_\theta(\mathbf{x}_{k-1} | \mathbf{x}_k)),$$

where  $D_{\text{KL}}$  is the Kullback-Leibler (KL) divergence,  $C$  is a constant that corresponds to fixed parameters of the forward diffusion process [25]. The term  $L_k$  is the KL divergence between two Gaussian distributions, which can be computed in closed form as the squared error between two means

$$L_k = \mathbb{E}_{\mathbf{x}_k \sim q} \left[ \frac{1}{2\sigma_k^2} \|\tilde{\boldsymbol{\mu}}(\mathbf{x}_k, \mathbf{x}_0) - \boldsymbol{\mu}_\theta(\mathbf{x}_k, k)\|_2^2 \right] + C', \quad (6)$$

where  $C'$  is a constant independent of  $\theta$ . Now we note the fact that predicting  $\hat{\mathbf{x}}_{k-1} = \boldsymbol{\mu}_\theta(\mathbf{x}_k, k)$  given  $\mathbf{x}_k$  is equivalent to predicting the noise  $\boldsymbol{\epsilon}$  that was added to  $\mathbf{x}_{k-1}$  (see Equation 1). We can thus reformulate the problem as one of minimizing the error in the noise prediction. By parameterizing  $\boldsymbol{\mu}_\theta(\mathbf{x}_k, k)$  as a function of predicted noise  $\boldsymbol{\epsilon}_\theta(\mathbf{x}_k, k)$ , we can simplify Equation 6 to

$$\mathbb{E}_{q, \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [c_k(\beta_{1:K}, \sigma_k) \|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\theta(\mathbf{x}_k, k)\|_2^2], \quad (7)$$

where  $c_k(\beta_{1:K}, \sigma_k)$  is a constant computed from the beta schedule for each timestep  $k$ . Ho et al. [25] further simplify this equation to arrive at the final loss that, with a slight abuse of notation, becomes

$$\mathcal{L}_{\text{simple}} := \mathbb{E}_{k, \mathbf{x}_0, \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\theta(\mathbf{x}_k(\mathbf{x}_0, \boldsymbol{\epsilon}), k)\|_2^2] \quad (8a)$$

$$\mathbf{x}_k(\mathbf{x}_0, \boldsymbol{\epsilon}) := \sqrt{\bar{\alpha}_k} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_k} \boldsymbol{\epsilon}, \quad (8b)$$

where  $k$  is sampled uniformly as  $k \in [1, K]$ . After training, we can generate a sample  $\mathbf{x}_0$  by running the reverse diffusion process recursively from  $\mathbf{x}_K \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$

$$\mathbf{x}_{k-1} = \boldsymbol{\mu}_{\theta^*}(\mathbf{x}_k, k) + \sigma_k \mathbf{z}, \quad (9)$$

where  $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  for  $k > 1$ , else  $\mathbf{z} = \mathbf{0}$ .

## B. Diffusion model guidance for distribution transformation

Consider generally the problem of transforming one distribution  $P_{\text{src}}$  onto another distribution  $P_{\text{tgt}}$ . This problem is challenging, because finding such a correspondence between distributions has an inherent dependence on the intrinsic geometrical match between the two distributions. However, there is an enticing shortcut that presents itself in the form of diffusion processes. In particular, we note that the forward diffusion process should allow us to map  $P_{\text{src}}$  into a noise distribution  $P_{\text{noise}}$  (i.e., zero-mean isotropic Gaussian noise). Drawing samples from this noise distribution, we can execute the reverse diffusion process to guide the sample to being drawn from  $P_{\text{tgt}}$ , assuming of course that the reverse process was trained on data samples from  $P_{\text{tgt}}$ .

Although this diffusion process will ultimately take points  $\mathbf{x}_{\text{src}}$  from  $P_{\text{src}}$  and map to samples drawn from  $P_{\text{tgt}}$ , resulting in  $\hat{\mathbf{x}}_{\text{tgt}}$ , it will do so in a destructive manner. In particular, as the forward diffusion process is applied to  $\mathbf{x}_{\text{src}}$ , the distribution of these points becomes progressively more random (by design). In the end, its distribution reaches an isotropic Gaussian distribution, completely corrupting the information in the original sample. At this point, there is no relationship between the source points  $\mathbf{x}_{\text{src}}$  and the derived points  $\hat{\mathbf{x}}_{\text{tgt}}$ . We may as well have simply started from Gaussian noise and run reverse diffusion onto  $P_{\text{tgt}}$ .

We would like to learn a transformation  $\mathcal{F} : \mathbf{x}_{\text{src}} \rightarrow \hat{\mathbf{x}}_{\text{tgt}}$  that preserves information. Such a transformation between distributions can be achieved by running the forward diffusion process on  $P_{\text{src}}(x)$  *partway through*, up to  $k = k_{\text{sw}}$ , followed by the reverse diffusion process for the same number of steps ( $k_{\text{sw}}$ ), in order to transform these partially diffused points onto  $P_{\text{tgt}}$ . Crucially, this process trades off the amount of

information present in  $\mathbf{x}_{\text{src}}$  that is preserved as a result of forward diffusion, and the consistency (in terms of likelihood) of the sample generated via reverse diffusion with respect to the target distribution  $P_{\text{tgt}}$  (i.e., the distribution over desired behaviors).

Very recently, Meng et al. [32] proposed a similar partial diffusion process for the task of generating realistic images based upon a user-provided sketch. The authors derive a bound on the distance between  $\hat{\mathbf{x}}_{\text{tgt}}$  and  $\mathbf{x}_{\text{src}}$  as follows.<sup>2</sup> Assuming that  $\|\epsilon_{\theta}(\mathbf{x}, k)\| \leq \mathcal{K}$  for all  $\mathbf{x} \in X$  and  $k \in [0, K]$ , then for all  $\delta \in (0, 1)$  with probability at least  $(1 - \delta)$ ,

$$\|\mathbf{x}_{\text{src}} - \hat{\mathbf{x}}_{\text{tgt}}\|_2^2 \leq \sigma_{k_{\text{sw}}}^2 (\mathcal{K}\sigma_{k_{\text{sw}}}^2 + d + 2\sqrt{-d \cdot \log \delta} - 2 \log \delta), \quad (10)$$

where  $d$  is the dimensionality of  $\mathbf{x}$ . From this bound, we can see that the distance between  $\mathbf{x}_{\text{src}}$  and  $\hat{\mathbf{x}}_{\text{tgt}}$  increases with  $k_{\text{sw}}$ , since  $\sigma_{k_{\text{sw}}}$  increases with  $k_{\text{sw}}$  while other terms are not affected. This expression allows us to bound the difference between the action generated by the assistant  $\hat{\mathbf{x}}_{\text{tgt}}$  and that of the user  $\mathbf{x}_{\text{src}}$ .

Figure 2 visualizes this property in two dimensions. Here, temporarily ignoring state condition, samples from  $P_{\text{user}}$  (i.e., source distribution  $P_{\text{src}}$ ) form the shape of a triangle, and samples from  $P_{\text{demo}}$  (i.e., target distribution  $P_{\text{tgt}}$ ) follow a distribution with three modes, each centered on one of the vertices of a triangle. The black and blue arrow represents the idea of executing partial diffusion—running the forward (black) and reverse (blue) processes for  $k_{\text{sw}}$  steps. At the bottom of Figure 2, we visualize the result of partial forward and reverse diffusion for different values of  $k_{\text{sw}}$ . Here, we color points based on their original spatial location in  $P_{\text{src}}$  (e.g., green points in the lower-right, red at the top, and blue at the lower-left), and track their location over different degrees of forward and reverse diffusion. Based on the visualization of the resulting distributions, we see that the distance between the initial points  $\mathbf{x}_{\text{src}}$  and  $\hat{\mathbf{x}}_{\text{tgt}}$  increases with  $k_{\text{sw}}$ , consistent with Equation 10. More generally, we find that:

**When  $k_{\text{sw}}$  is small**

- The original information is well-preserved (small displacements; high *fidelity*)
- The obtained distribution is far from  $P_{\text{tgt}}$  (low *conformity*)

**When  $k_{\text{sw}}$  is large**

- The original information is corrupted (large displacements; low *fidelity*)
- The obtained distribution is close to  $P_{\text{tgt}}$  (high *conformity*)

To discuss the effect of  $k_{\text{sw}}$  independent of the number of diffusion steps  $K$ , we herein define the *Forward Diffusion Ratio* as  $\gamma := k_{\text{sw}}/K$ , and will refer to this throughout the paper.

A similar idea in diffusion models has been adopted in image generation, manipulation, and 3D geometry generation [32, 50, 38, 52].

<sup>2</sup>Their bound is for a variance exploding (VE) formulation of diffusion, whereas we employ DDPM that uses a variance preserving (VP) formulation. Despite the difference, they share the same mathematical intuition.

---

### Algorithm 1 Shared autonomy as partial diffusion

---

**Input:** Observation  $\mathbf{s}_t$  and pilot’s action  $\mathbf{a}_t^h$

**Output:** Shared action  $\mathbf{a}_t^s$

**Require:** A pretrained state-conditioned denoising model  $\mu_{\theta^*}(\mathbf{a}, k | \mathbf{s})$ , forward diffusion ratio  $\gamma$ , diffusion timestep  $K$

- 1: Compute the switching timestep  $k_{\text{sw}} \leftarrow \text{ToInteger}(\gamma K)$
  - 2: Sample a Gaussian noise  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
  - 3:  $\tilde{\mathbf{a}}_{t, k_{\text{sw}}}^h \leftarrow \sqrt{\bar{\alpha}_{k_{\text{sw}}}} \mathbf{a}_t^h + \sqrt{1 - \bar{\alpha}_{k_{\text{sw}}}} \epsilon$  ▷ Forward process
  - 4:  $\mathbf{a}_{t, k_{\text{sw}}}^s \leftarrow \tilde{\mathbf{a}}_{t, k_{\text{sw}}}^h$
  - 5: **for**  $k$  in  $k_{\text{sw}}, \dots, 2$  **do**
  - 6:     Sample a noise vector  $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
  - 7:      $\mathbf{a}_{t, k-1}^s \leftarrow \mu_{\theta^*}(\mathbf{a}_{t, k}^s, k | \mathbf{s}_t) + \sigma_k \mathbf{z}$  ▷ Reverse process
  - 8: **end for**
  - 9:  $\mathbf{a}_{t, 0}^s \leftarrow \mu_{\theta}(\mathbf{a}_{t, 1}^s, 1 | \mathbf{s}_t)$
  - 10: **return**  $\mathbf{a}_{t, 0}^s$
- 

### C. Distribution transformation for shared autonomy

In a shared autonomy task, a copilot is asked to produce a shared action  $\mathbf{a}_t^s$ , given the current state  $\mathbf{s}_t$  and the pilot’s action  $\mathbf{a}_t^h$ . The goal is for the copilot to intervene in a manner that preserves the pilot’s “intention” while correcting the action. However, formulating the pilot’s “intention” is challenging, especially when the space of goals is unknown or not well-defined.

Our algorithm assumes that the latent intent of the user has non-zero likelihood under the distribution over target behaviors that we learn to sample from using a diffusion model. In other words, if we have a finite set of demonstrations to train the diffusion model, the states or state-action pairs that a user would want to achieve need to exist in the demonstrations. With our formulation to shared autonomy, the copilot manages the trade-off between respecting pilot’s action (i.e., *fidelity* to the pilot) and executing an action that is likely under the learned behavior distribution (i.e., *conformity* to the target behaviors). This is illustrated in Fig. 8, and it bears similarity to the trade-off discussed in Section II-B based on Figure 2.

Now, given a state  $\mathbf{s}_t$ ,  $P_{\text{user}}$  in Figure 2 is an (unknown) distribution over pilot actions and  $P_{\text{demo}}$  is a distribution over target behaviors where each point represents a single action. We can interpret that each cluster in  $P_{\text{demo}}$  corresponds to a set of actions, each of which causes distinct transitions. Hence, each cluster can be thought of as a pilot’s “intention”, and the actions outside of the cluster as being undesirable. Considering the visualization of the different partially forward and reverse diffused distributions at the bottom of the figure, we see that as we increase  $\gamma$  from 0.0 to 1.0, most of the pilot’s actions begin to go inside of the set of ideal actions. This captures the notion that conformity to the target behavior increases with  $\gamma$ . On the other hand, once  $\gamma$  exceeds 0.5, the displacement of pilot actions (signified by the mixture of colors) begin to enlarge, which is equivalent to the copilot producing an action that ignores the pilot’s “intention”. This captures the notion that the fidelity to the pilot decreases as we increase  $\gamma$ .

Algorithm 1 summarizes the procedure of applying partial forward and reverse diffusion to a pilot action  $\mathbf{a}_t^h$  and generating a shared action  $\mathbf{a}_t^s$ . In Section III, we investigate the






Fig. 3: We evaluate our algorithm in the context of four shared autonomy environments including a (a) 2D Control task in which an agent navigates to one of two different goals, (b) Lunar Lander that tasks a drone with landing at a designated location, (c) a Lunar Reacher variant in which the objective is to reach a designated region in the environment, and (d) Block Pushing, in which the objective is to use a robot arm to push an object into one of two different goal regions.

fidelity-conformity trade-off by modulating  $\gamma$  empirically for various environments with different pilots.

### III. EXPERIMENTS

We evaluate our approach to shared autonomy by pairing our copilot with various pilots on a variety of simulated continuous control tasks (Fig. 3): 2D Control, Lunar Lander, Lunar Reacher and Block Pushing. Each of these tasks provide the opportunity for the pilot to execute one of several different behaviors that is not known to the copilot. We design each domain to include a randomly sampled target state the pilot intends to reach (herein referred to as the *goal*). For example, in Lunar Reacher, a goal location is sampled randomly above ground. Across all tasks, we reveal the goal only to the pilot, by including it as part of the state. The copilot never has access to the goal. This results in a scenario in which pilot’s intention (i.e., the goal) is unknown to the copilot.

We seek to answer the following questions: (1) How much can our copilot assist a pilot? (2) Does our copilot generalize to different pilots? and (3) What is the effect of the forward diffusion ratio  $\gamma$ , and how can we interpret it?

#### A. Continuous control domains

1) *2D Control*: We build a simple 2D continuous control domain based on the maze-2D environment in D4RL [18] built on the MuJoCo simulator [48], where the goal is located either at the lower-left or lower-right corners of the large open space (Fig. 3(a)). The agent is represented as a point mass, and the actions are 2D forces applied to itself. The state consists of the agent’s location and velocity, as well as the goal location. Episodes are terminated if a timeout of 300 steps is reached.

2) *Lunar Lander*: Lunar Lander (Fig. 3(b)) is a continuous control environment adopted from Open AI Gym [7] that involves landing a spaceship on a landing pad. The actions consist of continuous left, right and upward forces that emulate thrusters on the spaceship. The state contains the position, orientation, linear and rotational velocity of the spaceship, whether each leg touches the ground, and the landing pad location (provided only to the pilot). An episode ends when the spaceship lands on the landing pad and becomes idle, crashes, flies out of bounds, or it reaches a timeout of 1000 steps.




Fig. 4: A visualization of the resulting trajectories in the 2D Control environment for different settings for the forward diffusion ratio  $\gamma$ . The user’s objective is to reach the left-hand goal. Without assistance (a) the user successfully reaches the goal two times, while the eight others timeout. As we increase  $\gamma$ , we see that (b)–(d) the user reaches the desired goal a vast majority of the time. As  $\gamma$  gets closer to 1.0, (e) (f) the assisted policy conforms to the expert policy, which avoids timeouts, but without knowledge of the user’s goal distributes the trajectories evenly between the left and right goals.

3) *Lunar Reacher*: A variant of Lunar Lander adopted from previous work [43, 7], where the goal is not to land, but to reach a random target location above the ground (Fig. 3(c)). The setting otherwise matches that of Lunar Lander.

4) *Block Pushing*: A variant of the Simulated Pushing environment [17]. The environment (Fig. 3(d)) consists of a simulated six-DoF robot xArm6 in PyBullet [9] equipped with a small cylindrical end effector. The task is to push an object into one of two target zones in the robot’s workspace. The episode terminates when the target reaches one of the two locations or the number of steps exceeds a timeout of 100.

TABLE I: Success and crash/out-of-bounds (OOB) rates on Lunar Lander and Lunar Reacher for different pilots with ( $\gamma = 0.4$ ) and without assistance. Each entry corresponds to 10 episodes across 30 random seeds. Note that the Zero and Random pilots have no knowledge of the goal.

Pilot	Lunar Lander				Lunar Reacher			
	Success Rate $\uparrow$		Crash/OOB Rate $\downarrow$		Success Rate $\uparrow$		Crash/OOB Rate $\downarrow$	
	w/o Copilot	w/ Copilot	w/o Copilot	w/ Copilot	w/o Copilot	w/ Copilot	w/o Copilot	w/ Copilot
Noisy	20.67 $\pm$ 4.50	<b>68.00 <math>\pm</math> 5.35</b>	28.33 $\pm$ 2.62	<b>7.67 <math>\pm</math> 2.87</b>	14.33 $\pm$ 2.49	<b>45.33 <math>\pm</math> 3.30</b>	77.33 $\pm$ 3.09	<b>38.00 <math>\pm</math> 2.94</b>
Laggy	21.33 $\pm$ 2.05	<b>75.00 <math>\pm</math> 3.56</b>	76.67 $\pm$ 2.49	<b>9.67 <math>\pm</math> 3.86</b>	30.67 $\pm$ 5.56	<b>55.33 <math>\pm</math> 6.13</b>	69.33 $\pm$ 5.56	<b>31.33 <math>\pm</math> 2.87</b>
Zero	0.00 $\pm$ 0.00	27.00 $\pm$ 0.82	100.00 $\pm$ 0.00	19.00 $\pm$ 2.94	0.00 $\pm$ 0.00	19.67 $\pm$ 2.62	100.00 $\pm$ 0.00	58.33 $\pm$ 3.09
Random	0.00 $\pm$ 0.00	25.00 $\pm$ 4.32	100.00 $\pm$ 0.00	19.33 $\pm$ 5.25	4.33 $\pm$ 1.89	22.33 $\pm$ 2.87	95.33 $\pm$ 1.70	59.00 $\pm$ 0.82
Expert	77.67 $\pm$ 2.62	78.67 $\pm$ 2.87	12.33 $\pm$ 0.94	8.00 $\pm$ 1.63	49.33 $\pm$ 4.78	55.00 $\pm$ 2.16	44.00 $\pm$ 3.56	31.67 $\pm$ 2.87

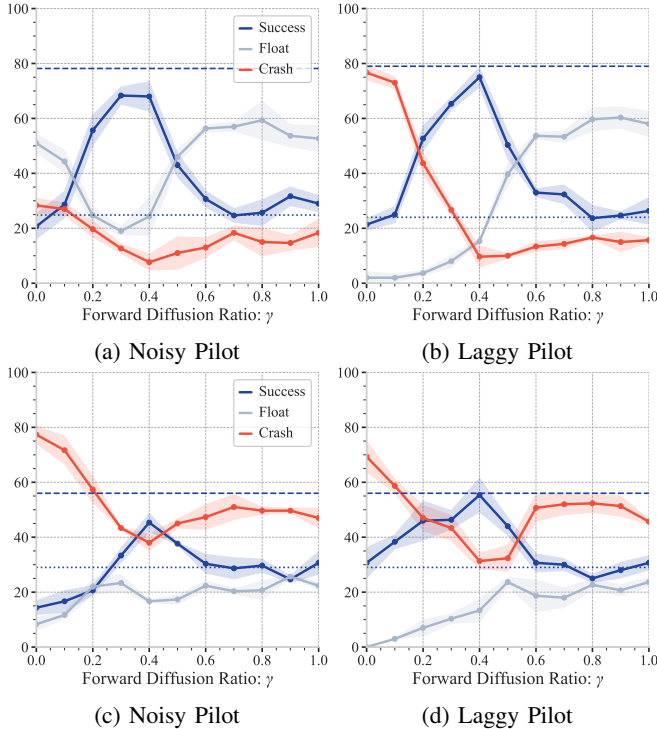


Fig. 5: Success (higher is better), floating, and crash (lower is better) rates for Lunar Lander (top) and Lunar Reacher (bottom) with (a) (c) noisy and (b) (d) laggy pilots. The dashed blue line denotes the success rate of an expert policy, while the dotted blue line denotes the success rate of our model with full-diffusion ( $\gamma = 1.0$ ).

The position and orientation of the block and end effector are randomly initialized at the start of each episode, while the target locations are fixed.

### B. Training

To understand the pilot’s intent, our copilot relies on a state-conditioned diffusion model  $p_{\theta}(a_t|s_t)$ . This diffusion model is trained as follows. First, we collect expert demonstrations, each containing a sequence of goal-embedded state-action pairs for each domain, where the goal is randomly sampled at each episode. We then use the resulting demonstrations to train a state-conditioned diffusion model using the DDPM loss




Fig. 6: The rates of block pushing task to terminate at correct goal, wrong goal or timeout with (c) noisy and (d) laggy pilot. The dashed blue line denotes the rate at which an expert policy reaches the correct goal, while the dotted blue line denotes the same rate for our model with full-diffusion ( $\gamma = 1.0$ ).

(Eqn. 8). The details of the model architecture are described in Appendix A. As we hide a goal from our copilot, we remove goal locations from each observation prior to training.

The following experiments use a single diffusion model trained separately for each task. We note that the copilot’s behavior changes only according to  $\gamma$  given a diffusion model.

To collect demonstrations, we first train an expert policy with soft actor-critic [23] for 3M timesteps in Lunar Lander and Lunar Reacher, and 1M timesteps in Block Pushing. We then roll out the policies in each environment to collect demonstrations of state-action pairs  $D_{\text{expert}}$  for various goals. Appendix A details the model architecture and the hyperparameter settings.

### C. Surrogate pilots for evaluation

Our method does not require access to a pilot (surrogate or otherwise) when training the copilot. However, surrogate pilots are useful when we want to perform a large number of evaluations in a reproducible manner. Thus, we prepare two surrogate pilots consistent with previous work [39, 43]: a *Laggy* pilot and a *Noisy* pilot, both of which are corrupted versions of a single expert. At each time step, the Laggy pilot repeats its previous action with probability  $p_{\text{laggy}}$ , and otherwise executes an action drawn from the expert’s policy.

TABLE II: Correct goal and timeout rates on the Block Pushing task for different pilots with ( $\gamma = 0.2$ ) and without assistance. Each entry corresponds to 10 episodes across 30 random seeds. Note that the Zero and Random pilots have no knowledge of the goal.

Pilot	Correct Goal Rate $\uparrow$		Timeout Rate $\downarrow$	
	w/o Copilot	w/ Copilot	w/o Copilot	w/ Copilot
Noisy	62.33 $\pm$ 2.49	<b>75.67 <math>\pm</math> 2.05</b>	37.33 $\pm$ 2.05	<b>20.00 <math>\pm</math> 3.27</b>
Laggy	42.00 $\pm$ 4.90	<b>74.33 <math>\pm</math> 2.87</b>	58.00 $\pm$ 4.90	<b>21.33 <math>\pm</math> 4.50</b>
Zero	0.00 $\pm$ 0.00	40.67 $\pm$ 2.62	100.00 $\pm$ 0.00	6.67 $\pm$ 1.70
Random	0.00 $\pm$ 0.00	16.33 $\pm$ 3.09	100.00 $\pm$ 0.00	68.00 $\pm$ 1.63
Expert	99.00 $\pm$ 0.82	94.67 $\pm$ 2.05	1.00 $\pm$ 0.82	5.33 $\pm$ 2.05

With probability  $p_{\text{noisy}}$ , the Noisy pilot samples an action from a uniform distribution over the action space, and otherwise executes an action sampled from the expert policy. We evaluate our shared autonomy algorithm for pilots across a broad range of parameters  $p_{\text{noisy}}$  and  $p_{\text{laggy}}$ . Due to space constraints, we only include results for a representative subset of the pilot parameters for Lunar Lander ( $p_{\text{laggy}} = 0.85$ ,  $p_{\text{noisy}} = 0.3$ ), Lunar Reacher ( $p_{\text{laggy}} = 0.85$ ,  $p_{\text{noisy}} = 0.6$ ), and Block Pushing ( $p_{\text{laggy}} = p_{\text{noisy}} = 0.6$ ), but include results for the full range of pilot parameters in Figures 9–14 of Appendix B.

#### D. Various pilots with our copilot

We evaluate our copilot with various surrogate pilots in all four environments. Table I shows success and crash/out-of-bound rates when various pilots are paired with our copilot in Lunar Lander and Lunar Reacher. We see that adopting our copilot significantly improves the success and crash rates for all but the expert policy (as expected). We note that each task involves reaching or landing on a randomly sampled target that is not known to the copilot. Consequently, one can not expect the copilot to have significant effect on the success rate of the Random or Zero pilot. Nevertheless, the results show that our copilot improves their success rates with performance similar to that of full diffusion ( $\gamma = 1.0$ ) and significantly decreases the rate at which the pilots crash.

For the Block Pushing task, the Laggy and Noisy policy are based on an expert whose intent is always to push the block to the green zone, which we refer to as the correct goal. The set of state-action pairs on which we trained the diffusion model also include trajectories that reach the red zone. As such, without knowledge of the pilot’s goal, the copilot may generate actions that push the block into the red zone (referred to as the incorrect goal), causing the episode to terminate. As we see in Table II, our copilot improves the rate at which all but an expert pilot reach the correct goal, while decreasing the rate at which episodes timeout without reaching any goal.

#### E. The effects of the forward diffusion ratio $\gamma$

The forward diffusion ratio  $\gamma$  determines how many steps of forward diffusion process to apply on pilot’s action  $a_t^h$ . This value changes the balance of interpolating between source and target distributions, which in turn controls the trade-off

between fidelity and conformity of an action. In this section, we demonstrate this trade-off with different  $\gamma$  values in various environments. We deployed our copilot with various surrogate pilots as discussed in Section III-C for each  $\gamma$  value.

Figure 4 shows the effect of  $\gamma$  for the 2D Control domain. Although the copilot is trained on demonstrations for both goals, the pilot’s goal is fixed to the left. As  $\gamma$  increases, we see that the number of trajectories that reach the correct goal increases. Qualitatively, we also see that the stochasticity of the trajectories lessens, and the paths towards the goal become smoother. However, when  $\gamma$  is too high, the copilot ignores the pilot’s actions. Without knowledge of the pilot’s goal, the copilot ends up distributing trajectories evenly between the left and right goals, consistent with the set of trajectories on which the model was trained.

Figures 5 summarizes the results on Lunar Lander and Lunar Reacher, respectively, by categorizing the trajectories into success, crash and float. Here, float denotes a trajectory where the spaceship does not crash nor go out of bounds, but nevertheless fails to complete the task within the time limit.

In all cases, we observe that the success rate follows a distinct pattern. First, the success rate monotonically increases with  $\gamma$  as the copilot improves task performance. After reaching a peak success rate at a critical value of  $\gamma$ , the success rate gradually decreases. The drop in the success rate reflects the copilots infidelity to the user’s intent. The pilot’s original actions provide a signal of which goal or landing pad is correct. When  $\gamma$  is too large, the copilot totally ignores the pilot’s actions and instead chooses to land in a manner following the copilot’s original training distribution.

In general, the crash rate monotonically decreases with  $\gamma$ , with the exception of Lunar Reacher. This demonstrates that the quality of the generated actions improves as  $\gamma$  increases. In Lunar Reacher, many of the demonstrations go straight to a target, which is often placed at the edge of the the environment. Consequently, if the agent overshoots and misses the target, it is very likely that it will go out of bounds, which is counted as a crash. This is likely why increasing  $\gamma$  does not necessarily result in lower crash rates for Lunar Reacher.

#### F. Real Human User Experiments

We conducted a set of human user experiments involving the Lunar Lander and Lunar Reacher continuous control domains. For the experiments, we recruited 17 participants (9 who identified as male, and 8 as female, with an average age of 25).<sup>3</sup> We asked each user to interact with one of two different co-pilots within each episode, one corresponding to direct teleoperation (i.e., no assistance) and the other being our diffusion-based shared autonomy assistant. The identity of the co-pilot was not disclosed to the participant. At the beginning of each experiment, we allowed the user to practice for 10 episodes with each co-pilot. In the subsequent testing phase, the user controlled the system for another 10 episodes

<sup>3</sup>None of the participants were co-authors or otherwise involved in this research.

with each of the co-pilots. We conducted this experiment for both Lunar Lander and Lunar Reacher and emphasize that the location of the goal (the landing pad for Lunar Lander and goal region for Lunar Reacher) was chosen randomly from the continuous space of goals for each episode.

TABLE III: Success, crash/out-of-bounds (OOB), and float rates on Lunar Lander and Lunar Reacher for human user experiments.

	Lunar Lander			Lunar Reacher		
	Success Rate $\uparrow$	Crash Rate $\downarrow$	Float Rate	Success Rate $\uparrow$	Crash Rate $\downarrow$	Float Rate
w/o Copilot	1.76	98.24	0.00	17.06	82.94	0.00
w/ Copilot	32.35	61.18	6.47	34.71	65.29	0.00

We consider both the quantitative and qualitative performance of our shared autonomy algorithm. Quantitatively, Table III compares the average performance of human pilots on Lunar Lander and Lunar Reacher in terms of success, crash/out-of-bounds, and float (i.e., neither crashing, going out-of-bounds, nor succeeding) with and without the assistance of our diffusion-based shared autonomy. We performed a Welch  $t$ -test for the success rate of each participant and find  $p$ -values of  $p = 3.363 \times 10^{-7}$  for Lunar Lander and  $p = 9.033 \times 10^{-3}$  for Lunar Reacher, based on which we can reject the null hypothesis that our diffusion-based co-pilot does not improve success rate.

Qualitatively, we asked participants to rate their agreement with five statements about whether they felt that each co-pilot behaved in a manner that was “helpful”, “consistent”, “responsive”, “collaborative”, and “trustworthy” using a five-point Likert scale. Figure 15 in Appendix C visualizes the survey results for all four combinations of tasks and co-pilots. The results reveal that human users rated our shared autonomy co-pilot higher than the no-assistance co-pilot (i.e., teleoperation) with regards to all five qualities for both environments.

### G. Real-Robot Experiments

We consider an real-world task (Fig. 7) in which a user commands a UR5 manipulator equipped with a pinch gripper to place a donut-shaped object (e.g., a disk with a hole in the center) held by the robot onto one of two posts in the robot’s workspace (e.g., analogous to Towers of Hanoi, but with two poles and one disk). Emulating remote operation, the user observes a real-time (i.e., little-to-no latency) video feed of the robot and its workspace from a third-person view. The user controls the three-dimensional velocity of the end-effector in Cartesian space using a game controller with two joysticks (velocity in the  $x$ - and  $y$ -directions using the right joystick, and the  $z$ -direction using the left joystick). The arm position is initialized to a random location.

As in the human user experiments discussed above, we presented the user with two unidentified co-pilots, one being our diffusion-based shared autonomy assistant and the other being direct teleoperation (i.e., no assistance). We allowed




Fig. 7: The real-robot experiments involved (left) placing a disk on one of two posts. Human operators (right) controlled the 3D motion of the end-effector while observing images of the robot’s workspace streamed in real-time.

TABLE IV: Success rate and episode length with and without copilot for the real-robot experiment with human users.

	success rate $\uparrow$	episode length $\downarrow$
w/o Copilot	0.89	209.23 $\pm$ 83.71
w/ Copilot	1.00	143.66 $\pm$ 68.42

each user to practice with each co-pilot for three episodes. In the subsequent testing phase, participants controlled the arm using both interfaces for five episodes each. We made sure to use the same set of (random) initial arm positions for both test phases. We set the time limit to be 500 steps (where each velocity command is considered a single step), and treat any episode that exceeds this limit as a failure. After the testing phase, we asked each participant to rate their agreement with the same quantitative questions in the same manner as the Lunar Lander and Lunar Reacher experiments using a five-point Likert scale.

We recruited 14 participants (10 who identified as male, 4 who identified as female, average age of 27.7)<sup>4</sup> for human user experiments with the real robot.

Table IV shows the success rate and episode length that users took to reach the goal. We found that providing users with the assistance of our diffusion-based shared autonomy algorithm both improved their success rate and decreased the time required to complete the task. The results of the qualitative evaluation reveal that users preferred commanding the manipulator with the assistance of our algorithm as detailed in Appendix C.

## IV. RELATED WORK

Shared autonomy has appeared in many problem domains, including remote telepresence [11, 21, 41], assistive robotic manipulation [29, 34, 44], and assistive navigation [4, 20]. In

<sup>4</sup>None of the participants were co-authors or otherwise involved in this research.



shared autonomy, one of the most persistent challenges has been correctly identifying the pilot’s intentions or goals.

Early work sidesteps this challenge by assuming a priori knowledge of the pilot’s goals [10, 30]. Recent work has managed to relax this assumption by treating the pilot’s goal as a latent random variable which can be inferred from environmental observations and pilot actions [1, 6, 14, 26, 28, 30, 33, 37, 53]. In spite of this forward progress, these methods still assume knowledge about some combination of the transition dynamics, the pilot’s goals, or pilot’s policy, making them difficult to deploy in many unstructured scenarios.

Reddy et al. [39], and subsequently [47, 43] introduce model-free deep reinforcement learning (RL) to the shared autonomy setting. Because these methods are model free, knowledge of environment dynamics is no longer required, allowing one to train a policy that is not limited to a specific model class. Several follow up works have adapted deep RL to a variety of shared autonomy problems [15, 40, 35, 8].

To the best of our knowledge, all previous work on shared autonomy with deep RL either explicitly or implicitly assumes a human-in-the-loop setting. In particular, the main training loop typically contains a step to query a pilot (i.e., human user) to obtain its action. This constraint often makes training inefficient or impractical. Chen et al. [8] addresses this issue by proposing a two-phase training scheme. In the first phase, an autonomous agent learns a task-conditioned policy that can be helpful for assisting a pilot. This is followed by the second phase, which incorporates sparse feedback from humans. Although such approach improves the efficiency of the training pipeline, it does not change the fact that these methods inherently rely on human feedback.

Diffusion models [45] have recently been applied to many problems including image generation, image editing, text-conditioned image generation and video generation. In the field of robotics, Janner et al. [27] trained a diffusion model over trajectories, demonstrating that diffusion is capable of generating a diverse set of trajectories reaching a goal location. Pearce et al. [36] applies diffusion models to imitation learning. Meanwhile, Wang et al. [51] show that the policy parameterized with diffusion can generate a multi-modal distributions over possible actions. As far as we are aware, ours is the first approach that uses diffusion models in a shared control scenario, where a copilot learns from expert demonstrations and provides guidance by denoising noisy pilot action.

The core idea of our approach is in the partial forward and reverse diffusion that enables us to generate an output that effectively interpolates between the original input and one from target distribution that we train our diffusion models on. Although we independently came up with this idea, a others have explored a similar approach for image editing [32]. Related, many recent papers have consider running reverse diffusion from some intermediate step rather than pure Gaussian noise [31, 50, 38, 52], as we do in this paper.

## V. CONCLUSION

In this paper, we presented a new approach to shared autonomy based on diffusion models. Our approach only requires access to demonstrations that are representative of desired behavior, and does not assume access to or knowledge of the user’s policy, reward feedback, or knowledge of the goal space or environment dynamics. Integral to our approach is its modulation of the forward and reverse diffusion processes in a manner that seeks to balance the user’s desire to maintain control authority with the benefits of generating actions that are consistent with the distribution over desired behaviors. We evaluated our copilot on various continuous control environments and demonstrated that our diffusion-based copilot generalizes across a variety of pilots, improving their performance, while preserving their intention. We further presented an analysis of the effects of different degrees of partial diffusion on task performance. One limitation of our approach is that there is no component that explicitly addresses the likely mismatch in state distributions between  $P_{\text{pilot}}(s)$  and  $P_{\text{target}}(s)$ . Intuitively, it is very likely that the target state visitation distribution (i.e., expert demonstrations) is different from that of the pilot. However, our empirical results suggest that this is not a critical limitation, possibly because executing corrected actions tends to encourage the agent to visit states that are close to those visited as part of the expert demonstrations. One means of addressing this is to design a goal-conditioned policy that can navigate itself to an in-distribution state at test time. We leave this for a future work.

## ACKNOWLEDGMENTS

Ge Yang is supported by the National Science Foundation Institute for Artificial Intelligence and Fundamental Interactions (IAIFI, <https://iaifi.org/>) under the Cooperative Agreement PHY-2019786. This work was supported in part by ARO grant W911NF-17-1-0188. We would like to thank the members of the TTIC robotics group for insightful discussions. We also appreciate the conversations we had with Haochen Wang. His in-depth knowledge of diffusion models was invaluable to us.

## REFERENCES

- [1] Daniel Aarno, Staffan Ekvall, and Danica Kragic. Adaptive virtual fixtures for machine-assisted teleoperation tasks. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2005.
- [2] David A Abbink, Tom Carlson, Mark Mulder, Joost C.F. de Winter, Farzad Aminravan, Tricia L. Gibo, and Erwin R Boer. A topology of shared control systems—finding common ground in diversity. *IEEE Transactions on Human-Machine Systems*, 48(5), 2018.
- [3] Peter Aigner and Brennan McCarragher. Human integration into robot control utilising potential fields. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1997.

- [4] Brenna D. Argall. Modular and adaptive wheelchair automation. In *Proceedings of the International Symposium on Experimental Robotics (ISER)*, 2016.
- [5] Gideon Billings, Matthew R. Walter, Oscar Pizarro, Matthew Johnson-Roberson, and Richard Camilli. Towards automated sample collection and return in extreme underwater environments. *arXiv preprint arXiv:2112.15127*, 2021.
- [6] Alexander Broad, Todd Murphey, and Brenna Argall. Highly parallelized data-driven MPC for minimal intervention shared control. In *Proceedings of Robotics: Science and Systems (RSS)*, 2019.
- [7] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym, 2016.
- [8] Sean Chen, Jensen Gao, Siddharth Reddy, Glen Berseth, Anca D. Dragan, and Sergey Levine. ASHA: Assistive teleoperation via human-in-the-loop reinforcement learning. *arXiv preprint arXiv:2202.02465*, 2022.
- [9] Erwin Coumans and Yunfei Bai. PyBullet, a Python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2021.
- [10] Jacob W. Crandall and Michael A. Goodrich. Characterizing efficiency of human robot interaction: A case study of shared-control teleoperation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2002.
- [11] Thomas Debus, Jeffrey Stoll, Robert D. Howe, and Pierre Dupont. Cooperative human and machine perception in teleoperated assembly. In *Proceedings of the International Symposium on Experimental Robotics (ISER)*, 2000.
- [12] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat GANs on image synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [13] Anca D. Dragan and Siddhartha S. Srinivasa. Formalizing assistive teleoperation. In *Proceedings of Robotics: Science and Systems (RSS)*, 2012.
- [14] Anca D. Dragan and Siddhartha S. Srinivasa. A policy-blending formalism for shared control. *International Journal of Robotics Research*, 32(7), 2013.
- [15] Yuqing Du, Stas Tiomkin, Emre Kiciman, Daniel Polani, Pieter Abbeel, and Anca Dragan. AvE: Assistance via empowerment. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [16] Charles Dugas, Yoshua Bengio, François Bélisle, Claude Nadeau, and René Garcia. Incorporating second-order functional knowledge for better option pricing. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2000.
- [17] Pete Florence, Corey Lynch, Andy Zeng, Oscar Ramirez, Ayzaan Wahid, Laura Downs, Adrian Wong, Johnny Lee, Igor Mordatch, and Jonathan Tompson. Implicit behavioral cloning. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2021.
- [18] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4RL: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*, 2020.
- [19] Cristian Garcia. Denoising diffusion probabilistic models, 2021. URL [https://github.com/acids-ircam/diffusion\\_models](https://github.com/acids-ircam/diffusion_models). Online.
- [20] Mahmoud Ghorbel, Joelle Pineau, Richard Gourdeau, Shervin Javdani, and Siddhartha Srinivasa. A decision-theoretic approach for the collaborative control of a smart wheelchair. *International Journal of Robotics Research*, 10:131–145, 2017.
- [21] Ray C. Goertz. Manipulators used for handling radioactive materials. *Human Factors in Technology*, 1963.
- [22] Matthew C Gombolay, Reymundo A Gutierrez, Giancarlo F Sturla, and Julie A Shah. Decision-making authority, team efficiency and human worker satisfaction in mixed human—robot teams. In *Proceedings of Robotics: Science and Systems (RSS)*, 2014.
- [23] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *CoRR*, abs/1801.01290, 2018.
- [24] Kris Hauser. Recognition, prediction, and planning for assisted teleoperation of freeform tasks. *Autonomous Robots*, 35(4), 2013.
- [25] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *CoRR*, abs/2006.11239, 2020.
- [26] Guy Hoffman and Cynthia Breazeal. Effects of anticipatory action on human-robot teamwork efficiency, fluency, and perception of team. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2007.
- [27] Michael Janner, Yilun Du, Joshua Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2022.
- [28] Shervin Javdani, Siddhartha S. Srinivasa, and J. Andrew Bagnell. Shared autonomy via hindsight optimization. In *Proceedings of Robotics: Science and Systems (RSS)*, 2015.
- [29] Hyun K. Kim, J. Biggs, W. Schloerb, M. Carmena, Mikhail A. Lebedev, Miguel A.L. Nicolelis, and Mandayam A. Srinivasan. Continuous shared control for stabilizing reaching and grasping with brain-machine interfaces. *IEEE Transactions on Biomedical Engineering*, 53(6), 2006.
- [30] Jonathan Kofman, Xianghai Wu, Timothy J Luu, and Siddharth Verma. Teleoperation of a robot manipulator using a vision-based human-robot interface. *IEEE Transactions on Industrial Electronics*, 52(5), 2005.
- [31] Zhaoyang Lyu, Xu Xudong, Ceyuan Yang, Dahua Lin, and Bo Dai. Accelerating diffusion models via early stop of the diffusion process. *arXiv preprint arXiv:2205.12524*, 2022.
- [32] Chenlin Meng, Yang Song, Jiaming Song, Jiajun Wu,

- Jun-Yan Zhu, and Stefano Ermon. SDEdit: Image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021.
- [33] Katharina Muelling, Abdeslam Boularias, Betty Mohler, Bernhard Schölkopf, and Jan Peters. Learning strategies in table tennis using inverse reinforcement learning. *Biological Cybernetics*, 108(5), 2014.
- [34] Katharina Muelling, Arun Venkatraman, Jean-Sebastien Valois, John E. Downey, Jeffrey Weiss, Shervin Javdani, Martial Hebert, Andrew B. Schwartz, Jennifer L. Collinger, and J. Andrew Bagnell. Autonomy infused teleoperation with application to brain computer interface controlled manipulation. *Autonomous Robots*, 41(6), August 2017.
- [35] Yoojin Oh, Marc Toussaint, and Jim Mainprice. Learning to arbitrate human and robot control using disagreement between sub-policies. *arXiv preprint arXiv:2108.10634*, 2021.
- [36] Tim Pearce, Tabish Rashid, Anssi Kanervisto, Dave Bignell, Mingfei Sun, Raluca Georgescu, Sergio Valcarcel Macua, Shan Zheng Tan, Ida Momennejad, Katja Hofmann, and Sam Devlin. Imitating human behaviour with diffusion models. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.
- [37] Claudia Pérez-D’Arpino and Julie A Shah. Fast target prediction of human reaching motion for cooperative human-robot manipulation tasks using time series classification. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2015.
- [38] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. DreamFusion: Text-to-3D using 2D diffusion. *arXiv preprint arXiv:2209.14988*, 2022.
- [39] Siddharth Reddy, Anca Dragan, and Sergey Levine. Shared autonomy via deep reinforcement learning. In *Proceedings of Robotics: Science and Systems (RSS)*, 2018.
- [40] Siddharth Reddy, Sergey Levine, and Anca D. Dragan. First contact: Unsupervised human-machine co-adaptation via mutual information maximization. *arXiv preprint arXiv:2205.12381*, 2022.
- [41] Louis B. Rosenberg. Virtual fixtures: Perceptual tools for telerobotic manipulation. In *Proceedings of the Virtual Reality Annual International Symposium (VRAIS)*, 1993.
- [42] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022.
- [43] Charles Schaff and Matthew R. Walter. Residual policy learning for shared autonomy. In *Proceedings of Robotics: Science and Systems (RSS)*, 2020.
- [44] Sebastian Schröer, Ingo Killmann, Barbara Frank, Martin Völker, Lukas Fiederer, Tonio Ball, and Wolfram Burgard. An autonomous robotic assistant for drinking. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2015.
- [45] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2015.
- [46] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [47] Weihao Tan, David Koleczek, Siddhant Pradhan, Nicholas Perello, Vivek Chettiar, Vishal Rohra, Aaslesha Rajaram, Soundararajan Srinivasan, H. M. Sajjad Hosain, and Yash Chandak. On optimizing interventions in shared autonomy. *arXiv preprint arXiv:2112.09169*, 2021.
- [48] Emanuel Todorov, Tom Erez, and Yuval Tassa. MuJoCo: A physics engine for model-based control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- [49] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7): 1661–1674, 2011.
- [50] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A. Yeh, and Greg Shakhnarovich. Score Jacobian chaining: Lifting pretrained 2D diffusion models for 3D generation. *arXiv preprint arXiv:2212.00774*, 2022.
- [51] Zhendong Wang, Jonathan J Hunt, and Mingyuan Zhou. Diffusion policies as an expressive policy class for offline reinforcement learning. *arXiv preprint arXiv:2208.06193*, 2022.
- [52] Wang Yinhuai, Yu Jiwen, and Zhang Jian. Zero shot image restoration using denoising diffusion null-space model. *arXiv preprint arXiv:2212.00490*, 2022.
- [53] Wentao Yu, Redwan Alqasemi, Rajiv Dubey, and Norali Pernalet. Telematipulation assistance based on motion intention recognition. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2005.

## APPENDIX A

### MODEL ARCHITECTURE AND HYPERPARAMETERS

We implement the state-conditioned denoising network  $\epsilon_{\theta}(\mathbf{a}_t, k \mid \mathbf{s}_t)$  that makes up the diffusion model as a 4-layer multi-layer perceptron (MLP) with latent dimension  $h_{\text{dim}} = 128$  and softplus [16] activation after every layer except for the last one. The network outputs a vector that has same dimension as in the input. Each layer is conditioned on a diffusion timestep  $k$  that looks up a corresponding embedding with dimension  $h_{\text{dim}} = 128$ , and is subsequently fused with the output of the linear layer by an element-wise product. To accommodate its conditioning on  $\mathbf{s}_t$ , our denoising network works on a concatenation of state-action pairs  $(\mathbf{s}_t, \mathbf{a}_t)$ , and during training, we add Gaussian noise only to the action component to produce a target output. The input  $\mathbf{x}$  and the target  $\mathbf{y}$  are then follow as

$$\mathbf{x} = (\mathbf{s}_t, \mathbf{a}_t + \epsilon), \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (11a)$$

$$\mathbf{y} = (\mathbf{0}, \epsilon). \quad (11b)$$

We generate the set of demonstrations used to train our diffusion model by collecting 100k transitions for each goal in the 2D Control environment using the scripted expert policy provided in the original implementation [18], 3M transitions for each goal in the Block Pushing environment using the pretrained expert, and 10M transitions for Lunar Lander and Lunar Reacher using the corresponding pretrained expert. When we collect expert demonstrations, we filter out episodes that do not succeed, effectively making the quality of the demonstrations better than the original expert. This explains why the assisted expert often performs better than the expert (e.g., Table I—Lunar Reacher).

We use a publicly available implementation of DDPM [19]. For the sake of faster action generation, we set the number of diffusion steps to  $K = 50$ , and use settings of  $\beta_{\min} = 10^{-4}$  and  $\beta_{\max} = 0.26$  with sigmoid scheduling.

An implementation of our method as source code together with video demonstrations are available at <https://diffusion-for-shared-autonomy.github.io/>.

## APPENDIX B

### EFFECT OF $\gamma$ FOR DIFFERENT SURROGATE PILOTS

Our framework modulates the amount of forward and reverse diffusion as a means of trading off a user’s desire to maintain control authority with the general benefits that come with emulating behavior that may be seen as more desirable. Figure 8 provides a visualization of how the forward diffusion ratio may influence the “fidelity” of the diffused state-conditioned action with respect to the user’s input and the extent to which it “conforms” to the desirable behavior.

As discussed in Section III-C, we investigate the effectiveness of our shared autonomy algorithm when used in conjunction with pilots that seek to emulate the deficiencies typical of human control. Following previous work [39, 43], we employ two surrogate policies for evaluation, namely a *Laggy* pilot and a *Noisy* pilot, both of which are corrupted versions of a single expert.




Fig. 8: Our algorithm uses the *forward diffusion ratio*  $\gamma$  to regulate the extent of forward and reverse diffusion. When  $\gamma$  is small, the assistant generates actions nearly identical to those of the user (i.e., high fidelity) but may be far from the desired behaviors (i.e., low conformity). Depending on the competency of the user, these actions may result in poor task performance (e.g., due to crashes). As with standard diffusion, when  $\gamma$  is large, the assistant will generate actions that have high likelihood under the distribution over desired behavior (i.e., high conformity), with little-to-no regard for their similarity to the user’s actions (i.e., low fidelity). While these actions will be safe, they will likely result in poor performance since the forward process has removed any information about user’s intention. By controlling  $\gamma$ , our algorithm identifies a regime that trades off between fidelity and conformity to generate actions that preserve the user’s intent and, in turn, the critical information to complete the task.

At each time step, with probability  $p_{\text{noisy}}$ , the Noisy pilot samples an action from a uniform distribution over the action space, and otherwise executes an action sampled from the expert policy. Similarly, the Laggy pilot repeats its previous action at each timestep with probability  $p_{\text{laggy}}$ , and otherwise carries out an action drawn from the expert policy. We evaluate our shared autonomy algorithm for pilots with a broad range of different parameters  $p_{\text{noisy}} \in \{0.2, 0.3, \dots, 0.8, 0.9\}$  and  $p_{\text{laggy}} \in \{0.2, 0.3, \dots, 0.8, 0.9\}$  for Lunar Lander, Lunar Reacher, and Block Pushing.

#### A. Lunar Lander and Lunar Reacher

Figure 9 plots the effect of different settings for the forward diffusion ratio  $\gamma$  in terms of success, crash, and float rates on Lunar Lander for different parameterizations of the Noisy pilot. At the smallest noise setting  $p_{\text{noisy}} = 0.2$ , the Noisy pilot is not significantly worse than the expert pilot without assistance ( $\gamma = 0.0$ ). The Noisy pilot successfully lands at the goal approximately 66% of the time compared to approximately 78% for the expert pilot, and crashes (or goes out of bounds) approximately 19% of the time compared to approximately 12% for the expert pilot. As the noise level increases, the performance of the noisy pilot significantly






Fig. 9: Lunar Lander Performance as a function of the forward diffusion ratio  $\gamma$  for a noisy pilot with noise values ranging from (a)  $p_{\text{noisy}} = 0.2$  to (h)  $p_{\text{noisy}} = 0.9$ . In all plots, the dashed blue line denotes the success rate of the expert pilot, while the dotted blue line is the success rate of our model when performing “full” diffusion (i.e.,  $\gamma = 1.0$ ) on an action sampled from a zero-mean isotropic Gaussian distribution, which we refer to as a “Random” pilot in the paper.

deteriorates without assistance—the pilot rarely lands and instead either crashes (or goes out-of-bounds) or floats around until the episode times out. However, for each parameterization of the Noisy pilot, we see that the assistance of the copilot gives rise to a noticeable increase in task success rate and decrease in crash rate as we increase the forward diffusion ratio to  $\gamma = 0.4$ . As the forward diffusion ratio increases further, the success rate decreases as expected given that the copilot does not have knowledge of the goal (i.e., the landing location). In this case, the copilot chooses actions that are consistent with the different goals that one or more previous experts reached. At the same time, we see only a slight or no increase in the crash rate as the forward diffusion ratio increases, particularly for pilots with higher noise settings.

Figure 10 presents the results for the same set of experiments on Lunar Lander with different parameterizations of the Laggy pilot. Consistent with previous work [39], Lunar Lander and, as we will see, Lunar Reacher are more tolerant of repeated actions than they are of noisy actions. Consequently, the performance of Laggy Pilots is only slightly worse than that of an expert pilot until  $p_{\text{laggy}}$  exceeds 0.8. As with the Noisy pilot, we see that providing a Laggy pilot with the assistance of a copilot increases the task success rate, while decreasing the crash rate up to a forward diffusion ratio of  $\gamma = 0.4$  for all parameterizations of the Laggy pilot. Increasing the forward diffusion ratio beyond  $\gamma = 0.4$  decreases the success rate since the copilot does not have access to the pilot’s goal, while the crash rate increases only slightly.

Figure 11 plots the performance of a Noisy pilot on Lunar Reacher for different noise settings. As with Lunar Lander,

the success rate of the unassisted Noisy pilot is similar to that of the expert pilot when  $p_{\text{noisy}}$  is low, yet for  $p_{\text{noisy}} \geq 0.6$  the success rate is below 20% and the pilot often goes out-of-bounds as indicated by the high crash rate. Again, we see that the assistance of a copilot increases the task success rate while decreasing the crash rate with a clear peak at  $\gamma = 0.4$  for all but the highest noise settings, for which performance only changes slightly as  $\gamma$  approaches 1.0. Again, this decrease in success rate is expected since without access to the region that the pilot is attempting to reach, the copilot chooses actions that are consistent with the set of expert demonstrations on which the diffusion model was trained.

Figure 12 plots the performance of the Laggy pilot on the Lunar Reacher task for different settings of  $p_{\text{laggy}}$ . Similar to Lunar Lander, we find that Lunar Reacher is tolerant of an otherwise (near-)expert policy that randomly repeats its previous actions. Due to the nature of the task, the success rate of the Laggy Pilot is only slightly worse than that of the expert until  $p_{\text{laggy}}$  exceeds 0.8. For all parameterizations of the Laggy pilot, we see that the assistance of a copilot with  $\gamma = 0.4$  decreases the crash rate, while simultaneously maintaining or, in the case of  $p_{\text{laggy}} \geq 0.8$ , increasing the success rate.

Table V expands upon the analysis presented in Table I and makes explicit the effect of different settings for the forward diffusion ratio between  $\gamma = 0.0$  (no assistance), our chosen setting of  $\gamma = 0.4$ , and to  $\gamma = 1.0$  (full assistance) when paired with different pilots. Consistent with the plots in Figures 9–12, setting the forward diffusion ratio to  $\gamma = 1.0$  results in similar performance across all pilots both in terms of success




Fig. 10: Lunar Lander performance as a function of the forward diffusion ratio  $\gamma$  for a laggy pilot with values ranging from (a)  $p_{laggy} = 0.2$  to (h)  $p_{laggy} = 0.9$ . In all plots, the dashed blue line denotes the success rate of the expert pilot, while the dotted blue line is the success rate of our model when performing “full” diffusion (i.e.,  $\gamma = 1.0$ ) on an action sampled from a zero-mean isotropic Gaussian distribution, which we refer to as a “Random” pilot in the paper.

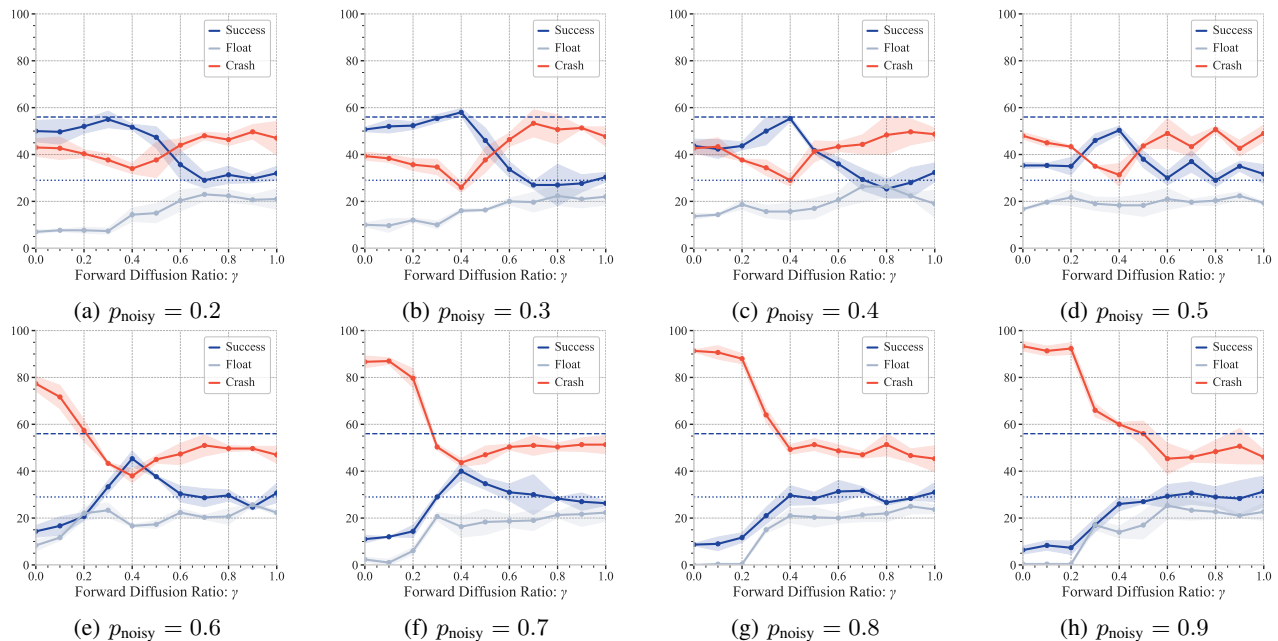


Fig. 11: Lunar Reacher performance as a function of the forward diffusion ratio  $\gamma$  for a Noisy pilot with noise values ranging from (a)  $p_{noisy} = 0.2$  to (h)  $p_{noisy} = 0.9$ . In all plots, the dashed blue line denotes the success rate of the expert pilot, while the dotted blue line is the success rate of our model when performing “full” diffusion (i.e.,  $\gamma = 1.0$ ) on an action sampled from a zero-mean isotropic Gaussian distribution, which we refer to as a “Random” pilot in the paper.

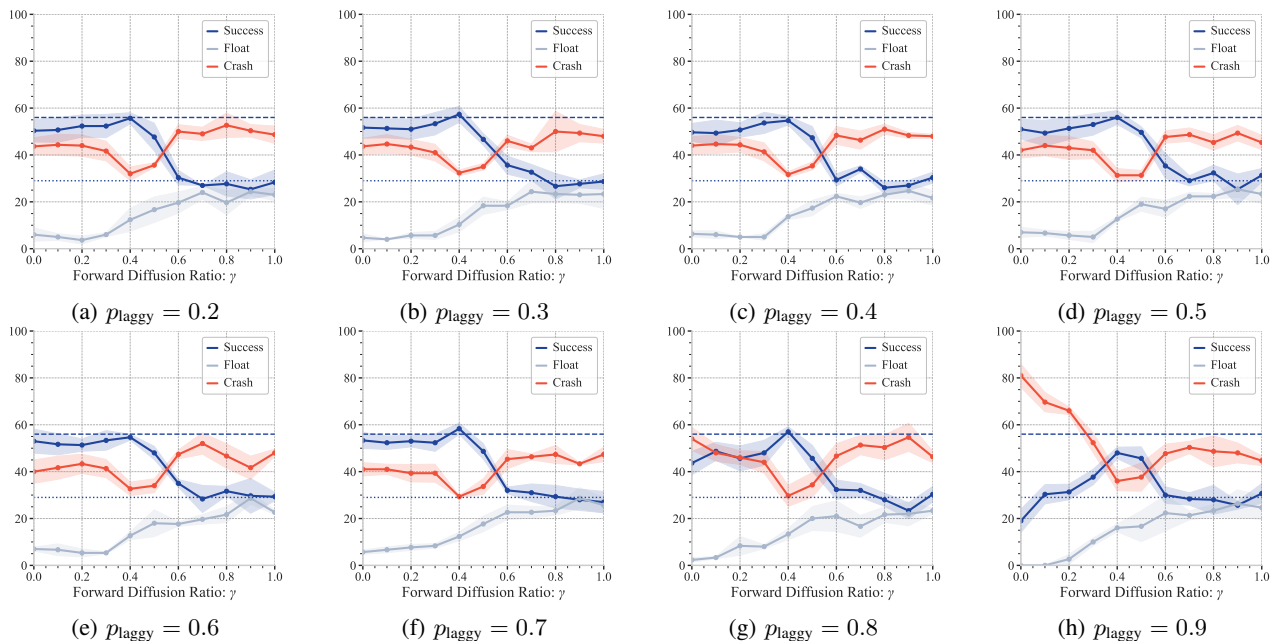


Fig. 12: Lunar Reacher performance as a function of the forward diffusion ratio  $\gamma$  for a Laggy pilot with values ranging from (a)  $p_{\text{laggy}} = 0.2$  to (h)  $p_{\text{laggy}} = 0.9$ . In all plots, the dashed blue line denotes the success rate of the expert pilot, while the dotted blue line is the success rate of our model when performing “full” diffusion (i.e.,  $\gamma = 1.0$ ) on an action sampled from a zero-mean isotropic Gaussian distribution, which we refer to as a “Random” pilot in the paper.

rate and crash rate, consistent with the copilot selecting actions that *conform* to the distribution over expert actions.

### B. Block Pushing

Figure 13 provides plots of the performance of the Noisy pilot on the Block Pushing task in terms of the correct goal, incorrect goal, and timeout rates for different settings of  $p_{\text{noisy}}$ . When the noise level is at or below 0.5, the Noisy pilot pushes the block to the correct goal more than 75% of the time without assistance (vs. nearly 100% for the expert) and never pushes the block to the incorrect goal. As the noise level increases from  $p_{\text{noisy}} = 0.6$ , the rate at which the unassisted Noisy pilot times out significantly increases, while the rate with which it reaches the correct goal precipitously drops for  $p_{\text{noisy}} \geq 0.8$ . With the assistance of a copilot, the rate at which the pilot pushes the block to the correct goal increases and the timeout rate decreases, both significantly, as the forward diffusion ratio increases to  $\gamma = 0.2$  for Noisy pilots with  $p_{\text{noisy}} \geq 0.6$ . As the forward diffusion ratio increases further and information about the correct goal is diminished, the assisted policy reaches one of the two candidate goals with roughly equal likelihood, while the timeout rate remains low, consistent with the set of expert demonstrations. For Noisy pilots with  $p_{\text{noisy}} \leq 0.5$ , assistance does not negatively affect

performance until the forward diffusion ratio exceeds  $\gamma = 0.2$ .

Figure 14 presents the performance of the Laggy pilot as a function of the forward diffusion ratio  $\gamma$  for different settings of  $p_{\text{laggy}}$ . The behavior is similar to that of the Noisy pilot (Fig. 13) except that increasing the level of copilot assistance to  $\gamma = 0.2$  improves the correct goal rate for all parameterizations of the Laggy pilot. Again, as the forward diffusion ratio approaches  $\gamma = 1.0$ , the copilot effectively chooses randomly between the two goals, while maintaining a low timeout rate.

## APPENDIX C

### REAL HUMAN USER EXPERIMENTS

In Lunar Lander, Lunar Reacher and real robot experiment, we asked the same set of quantitative questions using a five-point Likert scale. Figure 15 shows the results of Lunar Lander and Lunar Reacher. The users preferred our copilot over the placebo copilot in all of the metrics we asked in the question.

The results of real robot experiment is summarized in Figure 16. Our copilot got better scores in all metrics but for “responsive”. We believe that this is because unreasonable user control inputs (e.g., going away from any of the goal locations) are often strongly corrected by our copilot, which in turn gives the user an impression that their action is disrespected.

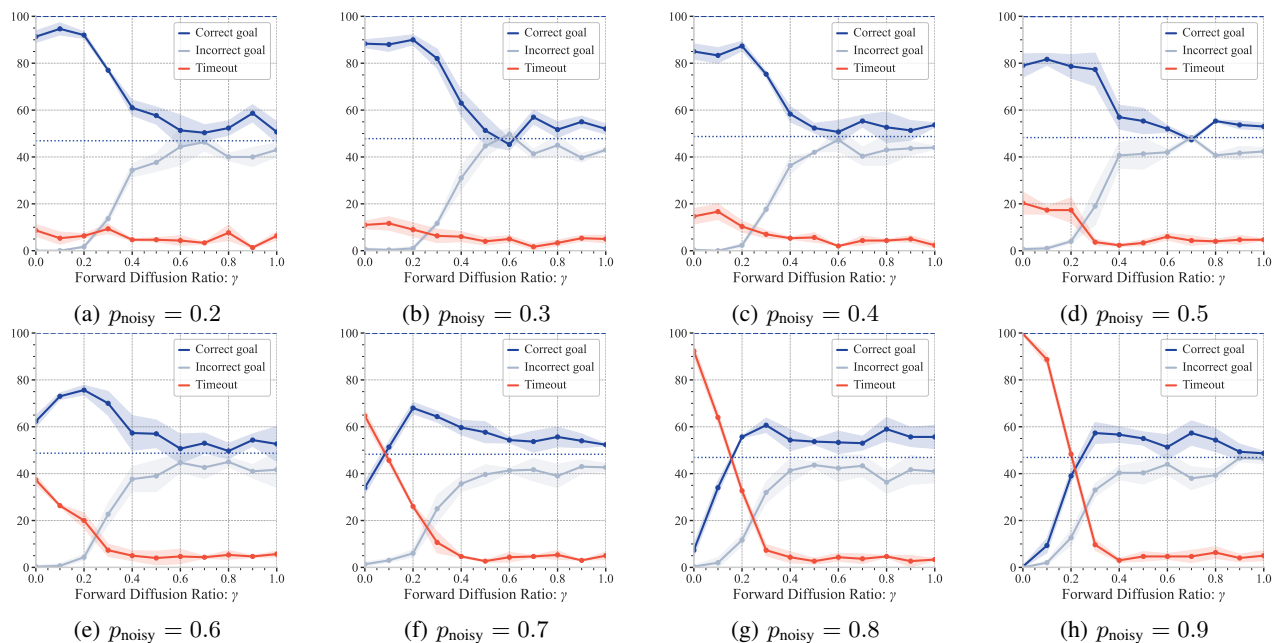


Fig. 13: Block Pushing performance in terms of the correct goal, wrong goal, and timeout rates as a function of the forward diffusion ratio  $\gamma$  for a Noisy pilot with noise values ranging from (a)  $p_{\text{noisy}} = 0.2$  (h)  $p_{\text{noisy}} = 0.9$ . In all plots, the dashed blue line denotes the success rate of the expert pilot, while the dotted blue line is the success rate of our model when performing “full” diffusion (i.e.,  $\gamma = 1.0$ ) on an action sampled from a zero-mean isotropic Gaussian distribution, which we refer to as a “Random” pilot in the paper.

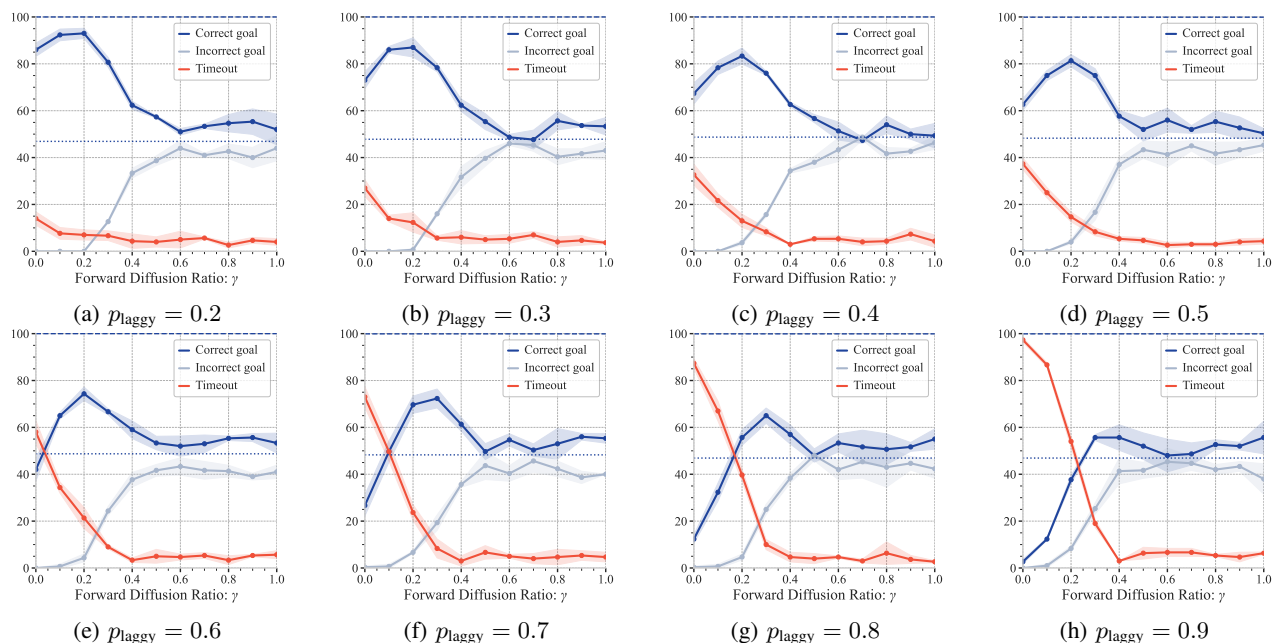


Fig. 14: Block Pushing performance in terms of the correct goal, wrong goal, and timeout rates as a function of the forward diffusion ratio  $\gamma$  for a Laggy pilot with values ranging from (a)  $p_{\text{laggy}} = 0.2$  to (h)  $p_{\text{laggy}} = 0.9$ . In all plots, the dashed blue line denotes the success rate of the expert pilot, while the dotted blue line is the success rate of our model when performing “full” diffusion (i.e.,  $\gamma = 1.0$ ) on an action sampled from a zero-mean isotropic Gaussian distribution, which we refer to as a “Random” pilot in the paper.



TABLE V: Success and crash/out-of-bounds (OOB) rates on Lunar Lander and Lunar Reacher for different pilots with and without assistance, where we show the results for our chosen value for  $\gamma = 0.4$  as well as the result of using full diffusion ( $\gamma = 1.0$ ), which has the effect of removing information about the goal provided by the pilot. The plots corresponding to the data in the table is figure 5. Each entry corresponds to 10 episodes across 30 random seeds. Note that the Zero and Random pilots have no knowledge of the goal.

Pilot	w/o Copilot ( $\gamma = 0.0$ )	Success Rate w/ Copilot (ours) ( $\gamma = 0.4$ )	w/ Copilot ( $\gamma = 1.0$ )	w/o Copilot ( $\gamma = 0.0$ )	Crash/OOB Rate w/ Copilot (ours) ( $\gamma = 0.4$ )	w/ Copilot ( $\gamma = 1.0$ )
<b>Lunar Lander</b>						
Noisy	20.67 $\pm$ 4.50	<b>68.00 <math>\pm</math> 5.35</b>	29.00 $\pm$ 2.94	28.33 $\pm$ 2.62	<b>7.67 <math>\pm</math> 2.87</b>	18.33 $\pm$ 4.92
Laggy	21.33 $\pm$ 2.05	<b>75.00 <math>\pm</math> 3.56</b>	26.33 $\pm$ 4.78	76.67 $\pm$ 2.49	<b>9.67 <math>\pm</math> 3.86</b>	15.67 $\pm$ 0.94
Zero	0.00 $\pm$ 0.00	27.00 $\pm$ 0.82	28.00 $\pm$ 2.16	100.00 $\pm$ 0.00	19.00 $\pm$ 2.94	13.33 $\pm$ 2.49
Random	0.00 $\pm$ 0.00	25.00 $\pm$ 4.32	27.67 $\pm$ 1.70	100.00 $\pm$ 0.00	19.33 $\pm$ 5.25	15.67 $\pm$ 1.25
Expert	77.67 $\pm$ 2.62	78.67 $\pm$ 2.87	29.33 $\pm$ 6.18	12.33 $\pm$ 0.94	8.00 $\pm$ 1.63	15.33 $\pm$ 1.70
<b>Lunar Reacher</b>						
Noisy	14.33 $\pm$ 2.49	<b>45.33 <math>\pm</math> 3.30</b>	30.67 $\pm$ 4.11	77.33 $\pm$ 3.09	<b>38.00 <math>\pm</math> 2.94</b>	47.00 $\pm$ 3.56
Laggy	30.67 $\pm$ 5.56	<b>55.33 <math>\pm</math> 6.13</b>	30.67 $\pm$ 2.62	69.33 $\pm$ 5.56	<b>31.33 <math>\pm</math> 2.87</b>	45.67 $\pm$ 0.94
Zero	0.00 $\pm$ 0.00	19.67 $\pm$ 2.62	31.00 $\pm$ 6.98	100.00 $\pm$ 0.00	58.33 $\pm$ 3.09	45.33 $\pm$ 5.56
Random	4.33 $\pm$ 1.89	22.33 $\pm$ 2.87	30.00 $\pm$ 4.08	95.33 $\pm$ 1.70	59.00 $\pm$ 0.82	45.00 $\pm$ 2.16
Expert	49.33 $\pm$ 4.78	55.00 $\pm$ 2.16	28.67 $\pm$ 4.50	44.00 $\pm$ 3.56	31.67 $\pm$ 2.87	50.00 $\pm$ 1.41

TABLE VI: Correct goal and timeout rates on the Block Pushing task for different pilots with and without assistance. where we show the results for our chosen value for  $\gamma = 0.2$  as well as the result of using full diffusion ( $\gamma = 1.0$ ), which has the effect of removing information about the goal provided by the pilot. The plot corresponding to the data in the table is figure 6. Each entry corresponds to 10 episodes across 30 random seeds. Note that the Zero and Random pilots have no knowledge of the goal.

Pilot	w/o Copilot ( $\gamma = 0.0$ )	Correct Goal Rate w/ Copilot (ours) ( $\gamma = 0.2$ )	w/ Copilot ( $\gamma = 1.0$ )	w/o Copilot ( $\gamma = 0.0$ )	Timeout Rate w/ Copilot (ours) ( $\gamma = 0.2$ )	w/ Copilot ( $\gamma = 1.0$ )
Noisy	62.33 $\pm$ 2.49	<b>75.67 <math>\pm</math> 2.05</b>	52.67 $\pm$ 7.41	37.33 $\pm$ 2.05	<b>20.00 <math>\pm</math> 3.27</b>	5.67 $\pm$ 1.25
Laggy	42.00 $\pm$ 4.90	<b>74.33 <math>\pm</math> 2.87</b>	53.33 $\pm$ 4.19	58.00 $\pm$ 4.90	<b>21.33 <math>\pm</math> 4.50</b>	5.67 $\pm$ 1.70
Zero	0.00 $\pm$ 0.00	40.67 $\pm$ 2.62	52.33 $\pm$ 0.47	100.00 $\pm$ 0.00	6.67 $\pm$ 1.70	3.67 $\pm$ 2.36
Random	0.00 $\pm$ 0.00	16.33 $\pm$ 3.09	52.00 $\pm$ 3.74	100.00 $\pm$ 0.00	68.00 $\pm$ 1.63	4.33 $\pm$ 2.05
Expert	99.00 $\pm$ 0.82	94.67 $\pm$ 2.05	51.00 $\pm$ 2.45	1.00 $\pm$ 0.82	5.33 $\pm$ 2.05	5.67 $\pm$ 2.49





Fig. 15: Results of the human user qualitative surveys for Lunar Lander and Lunar Reacher without and with the assistance of our shared autonomy algorithm.



(a) UR5 (w/o Co-pilot)

(b) UR5 (w/ Co-pilot)

Fig. 16: Results of the human user qualitative surveys for the UR5 manipulation experiments (left) without and (right) with the assistance of our method.