

Diffusion-Based Impedance Learning for Contact-Rich Manipulation Tasks

Noah Geiger, Tamim Asfour, Neville Hogan and Johannes Lachner

Abstract—Learning-based methods excel at robot motion generation but remain limited in contact-rich physical interaction. Impedance control provides stable and safe contact behavior but requires task-specific tuning of stiffness and damping parameters. We present Diffusion-Based Impedance Learning, a framework that bridges these paradigms by combining generative modeling with energy-consistent impedance control. A Transformer-based Diffusion Model, conditioned via cross-attention on measured external wrenches, reconstructs simulated Zero-Force Trajectories (sZFTs) that represent contact-consistent equilibrium behavior. A SLERP-based quaternion noise scheduler preserves geometric consistency for rotations on the unit sphere. The reconstructed sZFT is used by an energy-based estimator to adapt impedance online through directional stiffness and damping modulation. Trained on parkour and robot-assisted therapy demonstrations collected via Apple Vision Pro teleoperation, the model achieves sub-millimeter positional and sub-degree rotational accuracy using only tens of thousands of samples. Deployed in real-time torque control on a KUKA LBR iiwa, the approach enables smooth obstacle traversal and generalizes to unseen tasks, achieving 100% success in multi-geometry peg-in-hole insertion. The code for all experiments is publicly available on GitHub and videos of the experiments are available on the project website.

Index Terms—Imitation learning, Diffusion Models, Physical interaction, Impedance Control

I. INTRODUCTION

Robotic behavior arises from the interaction of two fundamentally different domains. On the one hand, motion planning and learning algorithms are based on data processing and operate on symbolic or numerical representations of virtual desired behavior, such as trajectories, poses, or task objectives. These representations are processed unidirectionally and are constrained only by temporal causality (“no output before input”) and boundedness. They are not subject to physical laws. In this paper, we refer to this as the *information domain* [1].

Physical interaction with the environment, on the other hand, is governed by bidirectional behavior. It is constrained by physical laws like energy conservation, entropy production, gravity, etc. In this paper, we refer to this as the *energy domain* [1]. Ignoring these physical constraints will yield robot control problems with stability (passivity) [2], [3] and safety [4].

Robots operating in contact-rich and unstructured environments must bridge these two domains: desired motions must be inferred and represented in the information domain, yet executed through control algorithms that take into account energy exchange with the environment. This requirement is critical in assembly, robot-assisted rehabilitation, and other contact-rich tasks where visual feedback may be limited and task success depends on regulating physical interaction rather than simply following a path.

The relation between motion planning (information domain) and physical interaction (energy domain) is well described by mechanical impedance [5]–[7]. Impedance Control has been widely adopted to guarantee stability [5] and safety [4] for robots that go in and out of contact. The separation of both domains can be formalized by equivalent networks, a generalization of the equivalent-circuit concept [8], to represent nonlinear dynamics and arbitrarily complex physical networks as a combination of primitives [9]–[11].

While Impedance Control provides a powerful framework for shaping robot–environment interaction [9], its effectiveness critically depends on the choice of stiffness and damping parameters [12]. Too much stiffness can lead to jamming, while too little can prevent task execution. Selecting these parameters is typically an iterative and time-consuming process, as current physics-based simulation tools are not yet sufficiently accurate or reliable to support fully offline tuning [13].

Recent advances in contact-rich manipulation can be divided into three main strands: 1) *model-based approaches*, 2) *statistical-impedance heuristics*, and 3) *learning-based approaches*.

Model-based approaches incorporate explicit contact dynamics including friction cones, complementarity constraints, and mixed-integer formulations [14]–[27]. These methods produce behavior consistent with physical contact but face challenges: high computational cost, parameter sensitivity, and limited robustness in unstructured settings. Most approaches require task-specific tuning of parameters. The selected control gains typically do not generalize across environments and tasks. Consequently, these approaches are not ideal for unstructured environments with unknown contact dynamics. These limitations motivate the development of learning-based frameworks capable of inferring and adapting stiffness online from observed interaction data.

Between explicit physics-based models and fully learning-based policies, a class of statistical-impedance heuristics have emerged within the learning-from-demonstration literature. These approaches typically model demonstrated motions using probabilistic representations such as Gaussian Mixture Mod-

N. Geiger and T. Asfour are with the Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology, Karlsruhe, Germany; e-mail: noah.geiger@student.kit.edu, asfour@kit.edu

N. Geiger, N. Hogan, and J. Lachner are with the Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, USA; N. Hogan, and J. Lachner are also with the Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, USA; e-mail: ngeiger@mit.edu, neville@mit.edu, jlachner@mit.edu

els, where the mean trajectory defines the nominal motion and the covariance is interpreted as a measure of task relevance or confidence [28], [29]. Impedance parameters are then modulated heuristically based on this statistical variability, with higher stiffness assigned to directions of low variance. These approaches have shown to be effective for structured tasks with consistent geometry. However, these methods are still limited in contact-rich settings where changing environmental constraints are not reflected in the statistical variability of the demonstrations.

Learning-based approaches, in contrast, primarily operate in the information domain. Reinforcement learning and sampling-based optimization methods [30]–[35] can achieve context-aware behavior, but typically lack an explicit representation of impedance, which prescribes a bidirectional relationship between motion and generalized force [36]. Instead, these methods rely on reward functions that penalize either force or position deviations. As a result, they tend to converge toward admittance-like behaviors that suppress interaction forces or motion deviations, rather than explicitly regulating impedance.

More recent generative approaches, including diffusion policies [37], flow-matching policies [38], [39], and Transformer-based policy models [40], [41], have demonstrated impressive performance in motion generation across tasks and embodiments [42]. However, these methods primarily focus on trajectory generation and are typically executed through PD controllers with fixed proportional and derivative gains to maintain stability during contact. Compliance in such systems is therefore not explicitly controlled but emerges indirectly through low control gains. While this strategy can prevent hardware damage, it provides limited control over physical interaction behavior in contact-rich or safety-critical scenarios where deliberate impedance modulation is essential.

This paper addresses this gap by formulating impedance adaptation as an equilibrium reconstruction problem rather than a trajectory learning problem. We present *Diffusion-Based Impedance Learning*, a framework that integrates generative modeling with energy-consistent impedance control. While the robot moves on a fixed nominal Zero-Force Trajectory (ZFT)¹, a conditional Transformer-based Diffusion Model reconstructs a simulated Zero-Force Trajectory (sZFT) from contact-perturbed motion and measured interaction wrenches. The reconstructed sZFT represents an inferred equilibrium that is consistent with the observed physical interaction and serves as the reference for impedance modulation. A directional adaptation scheme selectively reduces stiffness along non-task-relevant directions while preserving rigidity along directions essential for task execution.

Other work, in addition to that presented here, has also begun to combine generative models with compliant control for contact-rich manipulation. The Adaptive Compliance Policy (ACP) [43] shares several conceptual similarities with the present work, but differs in sensing modality, demonstration paradigm, controller architecture, and the role of the inferred equilibrium. ACP combines visual observations,

force sensing, and compliant kinesthetic teaching to directly predict compliance-related quantities for a low-level controller. In contrast, our approach operates solely on robot kinematics and measured interaction wrenches, reconstructing a contact-consistent equilibrium trajectory from teleoperated demonstrations and deriving impedance adaptation through an energy-based estimator. Conceptually, both ACP’s Virtual Target Pose and the proposed equilibrium reconstruction represent interaction-conditioned equilibria. However, the former is directly tracked as a force-generating target, whereas the latter serves as an intermediate representation for estimating impedance adaptation while the robot continues to follow its nominal motion reference.

In this way, our approach bridges data-driven learning and energy-consistent control: Diffusion Models explain physical interaction by reconstructing contact-consistent equilibria, while impedance control enforces stability and physical consistency during execution. Our method operates purely on robot kinematics and measured interaction wrenches, without relying on visual observations.

We validated the framework on a KUKA LBR iiwa in two contact-rich scenarios: parkour-style obstacle traversal and multi-geometry peg-in-hole insertion. Training data were collected through teleoperation using Apple Vision Pro (AVP) equipped with a markerless pose-tracking framework [44] integrated into the robot controller. Both tasks expose the limitations of fixed impedance and simple adaptation schemes, which require careful, task-specific parameter tuning and still tend to jam at obstacles or fail to achieve insertion. In contrast, Diffusion-Based Impedance Learning achieves consistent success without manual retuning.

This result is particularly noteworthy in the peg-in-hole experiments, as the Diffusion Model was not trained on this task. Despite the absence of peg-in-hole demonstrations during training, the system achieved a 100% success rate across all peg geometries. This generalization was obtained using only 1.6 hours of teleoperated training data (72,000 samples). The results summarized in Table I underscore the central contribution of this work: robust contact-rich manipulation in unstructured environments requires explicitly bridging learning in the information domain with control in the energy domain.

The code for the Transformer-based Diffusion Model, the robot controller, and the Apple Vision Pro teleoperation framework is available in our GitHub repository. Demonstration videos of all experiments can be found on the project website.

II. PRELIMINARIES

First, we outline the process in both domains: Transformer-based diffusion in the information domain and energy-based Impedance Control in the energy domain.

A. Diffusion Models

Diffusion Models learn to sample from a data distribution by approximating the reverse-time dynamics of a Gaussian diffusion process [45]. Let $\{z_i\}_{i=1}^N$ be independent and identically distributed samples from the data distribution p_{data} ; the goal is

¹The Zero-Force Trajectory (ZFT), introduced by Hogan [12], refers to the commanded equilibrium motion in the unconstrained case: the unique end-effector pose at which the interaction wrench vanishes.

TABLE I: Preview of experimental outcomes (full details in Section IV). Fixed impedance fails in contact-rich tasks, while Diffusion-Based Impedance Learning achieves robust success by reconstructing a contact-consistent sZFT.

Controller	Parkour	Cylindrical Peg	Square Peg	Star Peg
Fixed impedance	× (stopped at first obstacle)	100%	13%	0%
Diffusion-Based Impedance Learning	✓ (smooth traversal)	100%	100%	100%

to learn a parametric family of distributions $\{p_\theta\}_{\theta \in \Theta}$ such that $p_\theta \approx p_{\text{data}}$. A forward Markov chain maps a clean sample z_0 to (approximately) $\mathcal{N}(0, \mathbf{I})$ over T steps. The learned reverse-time chain iteratively denoises z_t to produce a sample that approximates p_{data} (Figure 1).

1) *Forward Diffusion*: The forward diffusion process defines a Markov chain with Gaussian transitions

$$q(z_t | z_{t-1}) = \mathcal{N}(\sqrt{\alpha_t} z_{t-1}, (1 - \alpha_t) \mathbf{I}), \quad (1)$$

where $\alpha_t := 1 - \beta_t$ and $\{\beta_t\}_{t=1}^T \subset (0, 1)$ is a predefined noise schedule. The quantity $(1 - \alpha_t) = \beta_t$ is the per-step diffusion variance, which controls the signal-to-noise ratio across time. After T steps with $\bar{\alpha}_T := \prod_{i=1}^T \alpha_i \approx 0$, the marginal distribution approaches a standard normal distribution, $z_T \approx \mathcal{N}(0, \mathbf{I})$. The marginal distribution at step t has the closed form

$$q(z_t | z_0) = \mathcal{N}(\sqrt{\bar{\alpha}_t} z_0, (1 - \bar{\alpha}_t) \mathbf{I}), \quad (2)$$

which admits the equivalent reparameterization

$$z_t = \sqrt{\bar{\alpha}_t} z_0 + \sqrt{1 - \bar{\alpha}_t} \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \mathbf{I}). \quad (3)$$

2) *Noise Schedule*: The per-step variances $\{\beta_t\}_{t=1}^T$ define the noise schedule, which controls how the signal-to-noise ratio $\text{SNR}_t = \bar{\alpha}_t / (1 - \bar{\alpha}_t)$ decays from ∞ to 0 over time. A well-chosen schedule balances the complexity of denoising, keeping early steps near-deterministic and later steps close to pure noise. [46], [47]

3) *Reverse Process and Learning Objective*: The true reverse $q(z_{t-1} | z_t)$ is intractable, as it marginalizes over z_0 . However, since the forward chain is linear-Gaussian, the posterior $q(z_{t-1} | z_t, z_0)$ is itself Gaussian with mean $\tilde{\mu}_t(z_t, z_0)$ and variance $\tilde{\beta}_t \mathbf{I}$ [48]. The model represents $p_\theta(z_{t-1} | z_t) = \mathcal{N}(\mu_\theta(z_t, t), \Sigma_\theta(z_t, t))$ and fit θ by minimizing the variational bound $\sum_t D_{\text{KL}}(q(z_{t-1} | z_t, z_0) \| p_\theta(z_{t-1} | z_t))$. With the ε -parameterization,

$$\mu_\theta(z_t, t) = \frac{1}{\sqrt{\alpha_t}} \left(z_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \varepsilon_\theta(z_t, t) \right),$$

and $z_t = \sqrt{\bar{\alpha}_t} z_0 + \sqrt{1 - \bar{\alpha}_t} \varepsilon$, the objective reduces (up to constants) to the noise-prediction loss [47]:

$$\mathcal{L}_{\text{simple}}(\theta) = \mathbb{E}_{t, z_0, \varepsilon} \left[\left\| \varepsilon - \varepsilon_\theta(\sqrt{\bar{\alpha}_t} z_0 + \sqrt{1 - \bar{\alpha}_t} \varepsilon, t) \right\|_2^2 \right].$$

4) *Training and Inference Algorithms*: Training follows a noise-prediction objective: at each iteration, a clean sample $z_0 \sim p_{\text{data}}$, a timestep $t \in \{1, \dots, T\}$ (uniform), and Gaussian noise $\varepsilon \sim \mathcal{N}(0, \mathbf{I})$ are sampled. The closed-form forward diffusion produces z_t , then the network ε_θ is updated to minimize mean-squared error between ε and its prediction $\varepsilon_\theta(z_t, t)$ [47], [49]. At inference, generation starts from $z_T \sim \mathcal{N}(0, \mathbf{I})$ and applies the learned reverse transitions $p_\theta(z_{t-1} | z_t)$ for

$t=T, \dots, 1$, yielding z_0 as a sample from the model [47]. For robotic trajectory generation, this iterative denoising maps pure noise into trajectories that are consistent with task goals (Figure 1).

5) *Conditional Diffusion*: To generate task-consistent behavior, the denoising process is conditioned on an additional context variable c (e.g., goals, scene state, or interaction wrench). The reverse model and noise predictor then become $p_\theta(z_{t-1} | z_t, c)$ and $\varepsilon_\theta(z_t, t, c)$.

Conditioning can be implemented in several ways, including (i) early fusion (e.g., feature concatenation), (ii) cross-attention over context tokens for Transformer denoisers, or (iii) lightweight fusion layers [50], [51].

In our approach, conditional diffusion is central. By conditioning on external wrenches, the model reconstructs sZFTs that are consistent with physical interaction. This ties the denoising process in the information domain to the regulation of interaction in the energy domain. More details about Diffusion Models can be found in [47], [49], [52], [53].

6) *Transformer-based Diffusion Model*: In this paper, we choose a Transformer-based denoising network since robotic trajectories are time series with context-dependent effects (e.g., contact events with delayed consequences). Self-attention provides a content-based global receptive field and fully parallelizable training, which is advantageous over Recurrent or Convolutional Neural Networks for sequence modeling [41], [54].

At each diffusion step, pose tokens are formed from translations and unit-quaternions. These tokens are embedded and enriched with learned positional and timestep embeddings. Interaction signals (forces and moments) are projected to context tokens and integrated through a multi-head cross-attention layer. Here, trajectory tokens serve as queries while context tokens serve as keys and values. The combined sequence is then processed by stacked self-attention and feed-forward layers, and a lightweight Multi-layer Perceptron (MLP) head predicts the 7D noise (3D translations and 4D unit-quaternions) for each token. This architecture captures temporal structure while conditioning on interaction signals, a requirement that is central for contact-rich applications [54]–[59]. More details about Transformer-based models can be found in [53], [54], [58], [59].

B. Energy-based Impedance Control

Impedance Control regulates the dynamic relationship between motion and interaction wrenches. Rather than commanding exact positions, the robot behaves like a virtual mechanical system characterized by stiffness, damping, and

inertia². In contrast to admittance control³, the impedance formulation enables the robot to maintain compliance while ensuring stability when interacting with passive environments [6], [36], [61], [62].

From an energy perspective, Impedance Control transforms the potential energy stored in a virtual spring to the work done by external wrenches [63]. Let the pose displacement be

$$\Delta \mathbf{x} = \begin{pmatrix} \mathbf{e}_t \\ \mathbf{e}_r \end{pmatrix} \in \mathfrak{se}(3), \quad (4a)$$

$$\mathbf{e}_t := \mathbf{p} - \mathbf{p}_0 \in \mathbb{R}^3, \quad (4b)$$

$$\mathbf{e}_r := \log(\mathbf{Q}_0 \mathbf{Q}^{-1}) \in \mathbb{R}^3 \simeq \mathfrak{so}(3). \quad (4c)$$

where the observed pose is $(\mathbf{p}, \mathbf{Q}) \in \mathbb{R}^3 \times \mathbb{S}^3$ and the equilibrium pose is $(\mathbf{p}_0, \mathbf{Q}_0) \in \mathbb{R}^3 \times \mathbb{S}^3$. Any unit quaternion admits the unit axis-angle representation

$$\mathbf{Q} = \left(\cos\left(\frac{\theta}{2}\right), \mathbf{u} \sin\left(\frac{\theta}{2}\right) \right), \quad \|\mathbf{u}\| = 1, \theta \in [0, \pi]. \quad (5)$$

In this formulation, the logarithm map is applied to the error quaternion $\mathbf{Q}_0 \mathbf{Q}^{-1}$ to return the axis-angle error vector $\mathbf{e}_r = \mathbf{u}_0 \theta_0$.

Let $\mathbf{K}_t, \mathbf{K}_r \in \mathbb{R}^{3 \times 3}$ be diagonal, positive semidefinite translational and rotational stiffness matrices and $\mathbf{K} = \text{diag}(\mathbf{K}_t, \mathbf{K}_r)$. Define the potential energy

$$\mathcal{U}(\Delta \mathbf{x}) = \frac{1}{2} \mathbf{e}_t^\top \mathbf{K}_t \mathbf{e}_t + \frac{1}{2} \mathbf{e}_r^\top \mathbf{K}_r \mathbf{e}_r \in \mathbb{R}_{\geq 0}. \quad (6)$$

The elastic wrench $\mathbf{F}_{\text{elastic}} \in \mathfrak{se}^*(3)$ can then be derived from the gradient

$$\mathbf{F}_{\text{elastic}} = \frac{\partial \mathcal{U}}{\partial \Delta \mathbf{x}} = \begin{pmatrix} \mathbf{f}_{\text{elastic}} \\ \mathbf{m}_{\text{elastic}} \end{pmatrix}, \quad (7)$$

with $\mathbf{f}_{\text{elastic}} \in \mathbb{R}^3$ and $\mathbf{m}_{\text{elastic}} \in \mathbb{R}^3$ being the elastic force and moment, respectively.

²And possibly higher-order dynamic effects.

³Admittance control requires additional methods to guarantee passivity [60].

1) *Damping design*: The damping design was formulated to ensure that energy dissipation scales consistently with both the commanded stiffness $\mathbf{K}_{\{t,r\}}$ and the task-space inertia for translations and rotations $\Lambda_{\{t,r\}} \in \mathbb{R}^{3 \times 3}$. To account for numerical issues near singular configurations, $\Lambda_{\{t,r\}}$ was computed using a damped least-squares method [2].

A user-selected positive time constant $d > 0$ defines the diagonal damping ratio matrix

$$\mathbf{D}_{\{t,r\}} = d \mathbf{I}_3. \quad (8)$$

The symmetric positive semidefinite square roots $\sqrt{\Lambda_{\{t,r\}}}$ and $\sqrt{\mathbf{K}_{\{t,r\}}}$ are obtained by eigendecomposition for $\Lambda_{\{t,r\}}$ and elementwise for $\mathbf{K}_{\{t,r\}}$. As presented in [64], an intermediate damping matrix $\mathbf{B}_{\{t,r\}}$ is constructed as

$$\mathbf{B}_{\{t,r\}} = \sqrt{\Lambda_{\{t,r\}}} \mathbf{D}_{\{t,r\}} \sqrt{\mathbf{K}_{\{t,r\}}} + \sqrt{\mathbf{K}_{\{t,r\}}} \mathbf{D}_{\{t,r\}} \sqrt{\Lambda_{\{t,r\}}}. \quad (9)$$

From $\mathbf{B}_{\{t,r\}}$, a scalar damping coefficient is computed to normalize the dissipation relative to the overall stiffness:

$$\lambda_{\{t,r\}} = \frac{2 \text{tr}(\mathbf{B}_{\{t,r\}})}{\text{tr}(\mathbf{K}_{\{t,r\}})} = \frac{2 \text{tr}(\mathbf{B}_{\{t,r\}})}{k_x + k_y + k_z}. \quad (10)$$

Finally, the damping matrices can be expressed as

$$\mathbf{B}_t = \lambda_t \mathbf{K}_t, \quad (11)$$

for translation and

$$\mathbf{B}_r = \lambda_r \mathbf{K}_r, \quad (12)$$

for rotation.

2) *Control Force*: The force $\mathbf{f}_{\text{ctrl}} \in \mathbb{R}^3$ is composed of $\mathbf{f}_{\text{elastic}}$ and damping through matrix $\mathbf{B}_t \in \mathbb{R}^{3 \times 3}$ that acts on the end-effector translational velocity $\dot{\mathbf{p}} \in \mathbb{R}^3$. It can be calculated as

$$\mathbf{f}_{\text{ctrl}} = \mathbf{f}_{\text{elastic}} - \mathbf{B}_t \dot{\mathbf{p}}. \quad (13)$$




Fig. 1: Diffusion overview. Bottom: dashed arrows depict the forward process $q(z_t | z_{t-1})$ which gradually corrupts a clean trajectory z_0 . Top: solid arrows depict the reverse process $p_\theta(z_{t-1} | z_t)$ which is learned to reconstruct structured robotic behavior. Figure adapted from [45].

3) *Control Moment*: The moment $\mathbf{m}_{\text{ctrl}} \in \mathbb{R}^3$ is composed of $\mathbf{m}_{\text{elastic}}$ and damping through matrix $\mathbf{B}_r \in \mathbb{R}^{3 \times 3}$ that acts on the angular velocity of the end-effector $\boldsymbol{\omega} \in \mathfrak{so}(3)$. It can be calculated as

$$\mathbf{m}_{\text{ctrl}} = \mathbf{m}_{\text{elastic}} - \mathbf{B}_r \boldsymbol{\omega}. \quad (14)$$

4) *Control Torque*: Finally, the transpose of the Jacobian $\mathbf{J}(\mathbf{q}) : \mathbb{R}^n \rightarrow \mathfrak{se}(3)$ provides the mapping between the 6-dimensional wrenches and the commanded joint torques $\boldsymbol{\tau}_{\text{ctrl}} \in \mathbb{R}^n$:

$$\boldsymbol{\tau}_{\text{ctrl}} = \mathbf{J}(\mathbf{q})^T \begin{pmatrix} \mathbf{f}_{\text{ctrl}} \\ \mathbf{m}_{\text{ctrl}} \end{pmatrix}. \quad (15)$$

In addition to the control wrench, the environment contributes external wrenches $\mathbf{F}_{\text{ext}} \in \mathfrak{se}^*(3)$.

The controller implemented in this paper can be found in our Github repository.

5) *Norton Equivalent Network*: The Impedance Control framework can be illustrated by the Norton Equivalent Network [6], [9], as represented in Figure 2. Here, two sources interact to manage physical interaction:

- The *motion source*, which represents a sequence of equilibrium poses $\{\mathbf{p}_0, \mathbf{Q}_0\}$ that together form the ZFT.
- The *impedance source*, which represents the baseline stiffness $\mathbf{K}$.

The motion source is part of the information domain. For contact-rich manipulation, it is a virtual nominal trajectory that lies beyond the physical constraints imposed by the environment. Impedance couples the motion source in the information domain to the energy domain (energy exchange through $\mathbf{F}$ and $\{\dot{\mathbf{p}}, \boldsymbol{\omega}\}$), a key property that we exploit later for energy-based stiffness estimation (Section III-E).

III. DIFFUSION-BASED IMPEDANCE LEARNING

A. Diffusion Model for Reconstruction of sZFT

From the perspective of diffusion, the displacement $\Delta \mathbf{x}$ can be interpreted as noise corrupting the equilibrium pose $(\mathbf{p}_0, \mathbf{Q}_0)$. The role of the Diffusion Model is to reconstruct the equilibrium pose by iteratively removing this noise. During training, the model learns to predict the difference between the observed pose $(\mathbf{p}, \mathbf{Q})$ and the ground truth equilibrium pose, guided by the measured external wrench $\mathbf{F}_{\text{ext}}$ (Figure 3). At inference, the model denoises $\Delta \mathbf{x}$ and outputs the reconstructed sZFT $(\hat{\mathbf{p}}_0, \hat{\mathbf{Q}}_0)$.

Throughout this paper, all quantities denoted by a hat ($\hat{\cdot}$) are predictions of the Diffusion Model and correspond to the sZFT.

It is important to distinguish between the nominal ZFT $(\mathbf{p}_0, \mathbf{Q}_0)$ and the reconstructed sZFT $(\hat{\mathbf{p}}_0, \hat{\mathbf{Q}}_0)$ in our formulation. The ZFT represents the programmed nominal motion reference and remains fixed during execution. In this work, ZFTs were generated programmatically by chaining waypoints using minimum-jerk interpolation and *Spherical Linear Interpolation* (SLERP) [65]. More generally, ZFTs can also be obtained through frame-based programming, imitation learning for free-space motion [66], [67], or teleoperation [68]. A video of the demonstrations is available on the project

website and the code for ZFT generation is publicly available on GitHub.

The sZFT, in contrast, is reconstructed online by the diffusion model from measured end-effector motion and interaction wrenches. It provides a denoised estimate of the equilibrium trajectory that explains the observed interaction. During deployment, the robot continues to track the nominal ZFT as the commanded motion reference, while the reconstructed sZFT is used exclusively for impedance adaptation. Specifically, the elastic energy in the virtual springs and external work exchanged with the environment (Equations (24) - (25)) are computed with respect to the reconstructed sZFT, enabling online estimation of stiffness and damping consistent with the observed contact behavior. In this way, the diffusion model does not directly generate impedance, but instead infers the interaction-consistent equilibrium underlying the impedance controller.

For example, when the robot encounters an unexpected obstacle, the planned ZFT drives the end-effector into contact, producing an external wrench. Conditioned on this wrench, the Diffusion Model reconstructs the sZFT and updates the controlled impedance. An illustration of the difference between ZFT and sZFT can be seen in Figure 4.

Having illustrated the effect conceptually, we now formalize the denoising process, which can be expressed in translational and rotational noise components. For translations, the full noise is defined as

$$\mathbf{p}_{\text{noise}} = \mathbf{p} - \mathbf{p}_0. \quad (16)$$

During training, a fraction of this noise is added to the clean equilibrium pose via the noise scheduler (Section III-B), and the network is trained to predict it. At inference, the predicted noise $\hat{\mathbf{p}}_{\text{noise}}$ is subtracted from the observation to reconstruct the equilibrium position iteratively,

$$\hat{\mathbf{p}}_0 = \mathbf{p} - \hat{\mathbf{p}}_{\text{noise}}. \quad (17)$$

For rotations, the full noise is represented by the relative quaternion

$$\mathbf{Q}_{\text{noise}} = \mathbf{Q} \mathbf{Q}_0^{-1}, \quad (18)$$

which encodes the rotational displacement between the observed quaternion $\mathbf{Q}$ and the equilibrium quaternion $\mathbf{Q}_0$. Partial training noise is injected to $\mathbf{Q}_0$ using SLERP [65] (Section III-B), and the network is trained to predict $\hat{\mathbf{Q}}_{\text{noise}}$. At inference, the equilibrium quaternion is reconstructed by

$$\hat{\mathbf{Q}}_0 = \mathbf{Q} \hat{\mathbf{Q}}_{\text{noise}}^{-1}. \quad (19)$$

The reconstructed sZFT provides the key link between the Diffusion Model and impedance adaptation. It serves as the input for the energy-based stiffness estimation described in Section III-E.

B. Noise Scheduling

For the denoising process, the maximum noise magnitude was computed separately for translations and rotations, and a variational noise schedule was applied across diffusion steps.




Fig. 2: Norton equivalent network to represent physical interaction as the combination of a virtual motion source and an impedance source. The displacement $\Delta \mathbf{x}$ is coupled to stiffness $\mathbf{K}$ to generate wrench.



Fig. 3: Diffusion-based reconstruction of the simulated Zero-Force Trajectory (sZFT). The observed pose $\{\mathbf{p}, \mathbf{Q}\}$ together with the external wrench $\mathbf{F}_{\text{ext}}$ is treated as noise relative to the equilibrium $\{\mathbf{p}_0, \mathbf{Q}_0\}$. The Diffusion Model iteratively denoises this input to reconstruct the equilibrium trajectory, which is then used in energy-based stiffness estimation.

1) *Translations*: For translations, the deviation between the observed end-effector position $\mathbf{p}$ and the ground-truth ZFT $\mathbf{p}_0$ was treated as noise. Additional zero-mean Gaussian noise $\epsilon_{\text{gauss}} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$ was injected during training to improve generalization and prevent overfitting [69]. The translational noise for the model was defined as

$$\text{noise}_t = \mathbf{p} - \mathbf{p}_0 + \epsilon_{\text{gauss}}, \quad (20)$$

with perturbed positions computed as

$$\mathbf{p}_{\text{noise}} = \mathbf{p}_0 + \beta_t \cdot \text{noise}_t, \quad (21)$$

where β_t is the noise scaling factor at diffusion step t .

2) *Rotations*: In our robot controller, orientations are represented with unit quaternions $\mathbf{Q}$. Injecting noise, as in the translational case, would violate the unit-norm constraint and distort the geometry of rotations. To address this, we introduce a SLERP-based [65] noise scheduler that perturbs quaternions directly on the unit sphere. This approach preserves valid orientations, provides perturbations consistent with geodesics, and allows rotational noise to be scheduled smoothly across diffusion steps. To our best knowledge, this is the first physically-consistent noise scheduler for rotations in robotics, formulated directly on a global representation of rotations [6].

The SLERP process is illustrated in Figure 5. The interpolation between two quaternions $\mathbf{Q}_a$ and $\mathbf{Q}_b$ is defined as

$$\begin{aligned} \text{SLERP}(\mathbf{Q}_a, \mathbf{Q}_b; \mu) &= \frac{\sin((1-\mu)\Omega)}{\sin(\Omega)} \mathbf{Q}_a \\ &+ \frac{\sin(\mu\Omega)}{\sin(\Omega)} \mathbf{Q}_b. \end{aligned} \quad (22)$$




Fig. 4: Illustration of the nominal Zero-Force Trajectory during deployment (ZFT, light blue) and the reconstructed simulated Zero-Force Trajectory (sZFT, skin color), generated by the Diffusion Model. The ground truth sZFT was created by teleoperation, which tracked the wrist frame $\{W\}$ and serves as an input to the robot controller to move the end-effector frame $\{EE\}$. In the illustration, the ZFT is just a straight line. However, due to the external wrench during contact, the Diffusion Model will reconstruct the sZFT to imitate the demonstrated behavior. Autonomously, the robot will adapt its compliant behavior through impedance to smoothly go over the obstacle.

where $\Omega = \arccos(\langle \mathbf{Q}_a, \mathbf{Q}_b \rangle)$ is the angle between $\mathbf{Q}_a$ and $\mathbf{Q}_b$, and $\mu \in [0, 1]$ is the interpolation factor.

C. Data Collection

All training data for Diffusion-Based Impedance Learning were collected through teleoperation using the VisionPro-Teleop framework [44] and Apple Vision Pro (AVP), which enabled markerless and spatially unconstrained teleoperation without external infrastructure. The AVP continuously tracked the operator's hand motion in six degrees of freedom, allowing teleoperation in free space, during contact interactions, and throughout transitions into and out of contact (Figure 6). Within the framework, the operator's hand pose was represented as a homogeneous transformation matrix based on a




Fig. 5: Rotational noise scheduling with Spherical Linear Interpolation (SLERP) [65]. Perturbed quaternions stay on the unit sphere, enabling geometry-consistent noise injection and removal.

27-joint kinematic hand model. In our implementation, the `handMiddleFingerKnuckle` (index 11) was used as the primary keypoint for translation and rotation tracking. During teleoperation, the tracked wrist pose served as the commanded ZFT and was subsequently treated as the ground-truth sZFT for model training. In parallel, external interaction wrenches at the robot end-effector were recorded using a force/torque sensor. Pose data were transferred via shared memory between the Python-based `VisionProTeleop` framework and the C++ robot controller. The implementation of the teleoperation interface is publicly available in our GitHub repository.

Unlike demonstrations with manipulandums and haptic interfaces, the AVP framework allowed spatially unrestricted teleoperation. In practice, the operator’s hand pose tracked by AVP could move freely in space, unconstrained by physical objects, whereas the robot end-effector was constrained by the actual contact surfaces. Therefore, the shift between the commanded ZFT and the robot pose generated an external wrench.

Two datasets were generated. The first consisted of a parkour-style scenario with multiple obstacles (Figure 6, right). A total of 17,457 samples were recorded. These data captured diverse interaction behaviors, including free-space motion, transitions into and out of contact, and sustained contact with unknown obstacles in an unstructured environment.

The second dataset focused on robotic-assisted upper-limb rehabilitation tasks (Figure 6, left). A total of 55,838 samples were recorded. Eight healthy adults (4 female, 4 male; age 20–64 years; height 162–192 cm; weight 64–91 kg) performed three types of movements derived from standard post-stroke rehabilitation protocols: (i) passive lifting, consisting of repeated vertical up-and-down motions of the arm; (ii) out-of-shoulder movements, involving arm motions that include shoulder rotation; and (iii) an activity-of-daily-living (ADL) task, in which subjects mimicked a feeding motion by moving the hand toward the mouth.

Each therapy type was performed under three interaction conditions designed to elicit different levels of subject participation: 1) passive condition: subjects were instructed to relax their arm and allow the robot to move it without resistance; 2) cooperative condition: subjects observed the motion demonstrated by the teleoperator wearing the AVP and actively attempted to minimize the interaction forces between their arm

and the robot hand; 3) mixed condition: subjects alternated between relaxing their arm and actively cooperating with the robot motion. Together, these conditions were chosen to capture varying levels of subject participation and assistance, providing diverse interaction behaviors for learning contact-consistent impedance adaptation.

The study was approved by MIT’s Committee on the Use of Humans as Experimental Subjects (COUHES), with informed consent obtained from all participants. This dataset was used to train the model presented in this paper; however, the deployment of upper-limb rehabilitation applications is part of a separate project and is not reported here.

The desired compliant robot behavior was chosen during the data collection phase and kept unchanged during the deployment phase (Section IV). For all trials, the baseline translational stiffness was set to $800 \frac{\text{N}}{\text{m}}$, and the baseline rotational stiffness was set to $150 \frac{\text{Nm}}{\text{rad}}$.

D. Model Selection to Reconstruct sZFT

The model selection was guided by two initial requirements: (i) sufficient accuracy in reconstructing both translational and rotational components and (ii) a small model capacity, suitable for real-time deployment. To evaluate this trade-off, we varied hidden dimensions, number of layers, attention heads, and loss-weighting strategies.

Performance was assessed using translational error (positional loss in mm) and rotational error (angular magnitude θ about unit-axis and unit-axis deviation α in degrees).

Figure 7 provides a graphical representation of the rotational losses, illustrating the decomposition into axis deviation α_{error} and magnitude error θ_{error} .

The initial requirements were 2 mm for the translational error and 2.8° (half of the small-angle approximation) for the angular error and unit-axis deviation. These values were determined heuristically based on previous contact-rich experiments [70].

Table II reports the results for the Transformer-based Diffusion Model on the parkour dataset. The configuration with 512 hidden dimensions, 4 attention heads, and 6 layers achieved the best balance between accuracy and efficiency, with positional error of 0.994 mm, θ error of 0.249° , and α error of 0.003° . Increasing the loss weight of the θ component (factor 5) yielded marginal improvements in rotational accuracy ($\theta = 0.240^\circ$, $\alpha = 0.003^\circ$) without degrading translational fidelity.

A notable result was that retraining this model on a combined dataset of parkour and robotic-assisted therapy motions—two completely different physical interaction tasks—further improved accuracy (positional error 0.883 mm, θ error 0.233° , α error 0.002°). Importantly, the two datasets were collected with different tools and inertial properties (lightweight tool for parkour and a 600 g prosthetic hand for physical therapy) yet the model compensated for these differences without explicit training. This robustness highlights the ability of Diffusion-Based Impedance Learning to generalize across different tasks and tool inertial properties on a single robot platform. The selected architecture was ultimately adopted as the final model because it satisfied the accuracy requirements while remaining small enough for real-time torque control.



Fig. 6: Examples of teleoperated demonstrations used for training. (Left) Physical therapy dataset with robotic-assisted upper-limb tasks. (Right) Parkour dataset with free-space and contact-rich motions across obstacles.

TABLE II: Transformer-based Diffusion Model results (parkour dataset). The selected model (highlighted) combines high accuracy with low computational complexity.

Hidden Dim	Heads	Layers	θ Loss [deg]	α Loss [deg]	Positional Loss [mm]
512	4	4	0.307	0.004	1.047
512	4	6	0.249	0.003	0.994
512	8	4	0.282	0.004	1.030
512	8	6	0.286	0.003	0.989
1024	4	4	0.228	0.003	1.033
1024	4	6	0.224	0.003	1.033
1024	8	4	0.227	0.003	1.016
1024	8	6	0.288	0.004	1.057




Fig. 7: Graphical representation of rotational losses, showing axis deviation α_{error} and magnitude error θ_{error} .

Another notable finding of this work is the small amount of training data required. The model already showed strong performance on the parkour dataset with only tens of thousands of samples (17,000 samples). For the combined data set, 72,000 samples at a controller rate of 0.005 s correspond to just 1.6 hours of demonstration. This is in contrast to image-based Diffusion Models, which require 100,000 samples even for learning single manipulation skills such as block pushing [68], [71].

E. Energy-based Stiffness Estimation

Impedance control can be interpreted as virtual springs that pull the robot end-effector toward an equilibrium pose. When the robot interacts with the environment, the deformation of these virtual springs is transformed into external wrenches, which perform mechanical work. A shift of the reconstructed

sZFT by the Diffusion Model changes the deformation, and thus the elastic energy stored in the virtual springs. By relating this change in elastic energy to the measured work, we can compute the stiffness equivalent that reproduces the observed interaction behavior.

For deployment, a user specifies baseline translational and rotational stiffness values that define the nominal interaction behavior. During execution, the Diffusion Model predicts the sZFT, which represents the equilibrium pose correction, consistent with the current interaction wrench. Rather than modifying the nominal ZFT, we adapt the impedance by reducing stiffness equivalent to the equilibrium reconstruction via sZFT.

Specifically, we compute nonnegative stiffness reduction terms $k_{t,i}^*$ and $k_{r,i}^*$ for each translational and rotational axis $i \in \{x, y, z\}$. These terms are subtracted from the predefined baseline stiffness values $K_{t,i,max}$ and $K_{r,i,max}$, yielding the effective stiffness used by the controller:

$$k_{t,i} = K_{t,i,max} - k_{t,i}^*, \quad (23a)$$

$$k_{r,i} = K_{r,i,max} - k_{r,i}^*. \quad (23b)$$

In our experiments, the baseline stiffness values were set deliberately high ($K_{t,i,max} = 800$ N/m, $K_{r,i,max} = 150$ Nm/rad). High baseline stiffness is desirable for tasks that require significant elastic energy storage, such as pressing or inserting a workpiece into another. We therefore selected baseline stiffness values that the robot can realistically render and that are sufficient for all tasks considered. The Diffusion Model

then autonomously reduced these stiffnesses when interaction occurred, adapting the impedance only where necessary.

1) *Translations*: For translations, the displacement was defined as $\hat{e}_t = \mathbf{p} - \hat{\mathbf{p}}_0$. The elastic energy $E_t \in \mathbb{R}$ of a virtual translational spring is

$$E_t^{\text{spring}} = \frac{1}{2} \mathbf{K}_t \|\hat{e}_t\|^2. \quad (24)$$

Here, $\|\cdot\|$ denotes the Euclidean norm. The work done by the external force $\mathbf{f}_{\text{ext}} \in \mathbb{R}^3$ along the virtual displacement $\hat{e}_t$ is given by the inner product $\langle \cdot, \cdot \rangle$:

$$E_t^{\text{work}} = \langle \mathbf{f}_{\text{ext}}, \hat{e}_t \rangle. \quad (25)$$

Equating the right hand sites of Equation (24) with Equation (25) yields the translational stiffness contribution

$$\mathbf{k}_t^* = \frac{2 \langle \mathbf{f}_{\text{ext}}, \hat{e}_t \rangle}{\|\hat{e}_t\|^2}. \quad (26)$$

To avoid oscillatory behavior, a damping term proportional to end-effector velocity was introduced:

$$\tilde{e}_t = \kappa_t \hat{e}_t - \gamma_t \dot{\mathbf{p}}, \quad (27)$$

where κ_t is a unit translational stiffness, $\dot{\mathbf{p}}$ is the translational end-effector velocity, and γ_t a damping coefficient.

2) *Rotations*: For rotations, the unit axis-angle representation $\hat{e}_r = \hat{\mathbf{u}}_0 \hat{\theta}_0$ was used (Equation (5)). The corresponding elastic energy $E_r \in \mathbb{R}$ of the virtual rotational spring is

$$E_r^{\text{spring}} = \frac{1}{2} \mathbf{K}_r \|\hat{e}_r\|^2. \quad (28)$$

The work done by the external moment $\mathbf{m}_{\text{ext}} \in \mathbb{R}^3$ about the rotational displacement $\hat{e}_r$ can be calculated by

$$E_r^{\text{work}} = \langle \mathbf{m}_{\text{ext}}, \hat{e}_r \rangle. \quad (29)$$

Equating the right hand sites of Equation (28) with Equation (29) results in the rotational stiffness contribution

$$\mathbf{k}_r^* = \frac{2 \langle \mathbf{m}_{\text{ext}}, \hat{e}_r \rangle}{\|\hat{e}_r\|^2}. \quad (30)$$

Dynamic effects were addressed by augmenting the error with a damping term:

$$\tilde{e}_r = \kappa_r \hat{e}_r - \gamma_r \boldsymbol{\omega}, \quad (31)$$

where κ_r is a unit rotational stiffness, $\boldsymbol{\omega}$ is the angular end-effector velocity, and γ_r a damping coefficient.

3) *Algorithm*: Algorithm 1 summarizes the procedure of the stiffness estimation. Stiffness was computed over a sliding window of 16 samples at 200 Hz (80 ms), which was sufficient given that the tasks were not highly dynamic. Thresholds of 1 N for force and 1 Nm for moment were applied, and values below these limits were ignored to suppress phases of low interaction.

Algorithm 1 Energy-Based Stiffness Estimation

```

1: Input:  $\hat{e}_t, \hat{e}_r, \dot{e}, \boldsymbol{\omega}, \mathbf{f}_{\text{ext}}, \mathbf{m}_{\text{ext}}, \kappa_t, \kappa_r, \gamma_t, \gamma_r,$ 
    $f_{\text{thres}}, m_{\text{thres}}, \mathbf{K}_{t,\text{max}}, \mathbf{K}_{r,\text{max}}, \varepsilon$ 
2: for  $i \in \{x, y, z\}$  do
3:    $\tilde{e}_{t,i} \leftarrow \kappa_t \hat{e}_{t,i} - \gamma_t \dot{e}_i$ 
4:   if  $|f_i| < f_{\text{thres}}$  or  $\tilde{e}_{t,i}^2 < \varepsilon$  then
5:      $k_{t,i}^* \leftarrow 0$ 
6:   else
7:      $k_{t,i}^* \leftarrow \frac{2 f_i \tilde{e}_{t,i}}{\tilde{e}_{t,i}^2 + \varepsilon}$ 
8:      $k_{t,i}^* \leftarrow \max(0, k_{t,i}^*)$ 
9:   end if
10:   $K_{t,i} \leftarrow \text{clip}(K_{t,\text{max},i} - k_{t,i}^*, 0, K_{t,\text{max},i})$ 
11:   $\tilde{e}_{r,i} \leftarrow \kappa_r \hat{e}_{r,i} - \gamma_r \omega_i$ 
12:  if  $|m_i| < m_{\text{thres}}$  or  $\tilde{e}_{r,i}^2 < \varepsilon$  then
13:     $k_{r,i}^* \leftarrow 0$ 
14:  else
15:     $k_{r,i}^* \leftarrow \frac{2 m_i \tilde{e}_{r,i}}{\tilde{e}_{r,i}^2 + \varepsilon}$ 
16:     $k_{r,i}^* \leftarrow \max(0, k_{r,i}^*)$ 
17:  end if
18:   $K_{r,i} \leftarrow \text{clip}(K_{r,\text{max},i} - k_{r,i}^*, 0, K_{r,\text{max},i})$ 
19: end for
20: Output:  $\mathbf{K}_t, \mathbf{K}_r$ 

```

4) *Directional Adaptation*: For deployment of our model in real-world experiments, the energy-based estimator was extended with a directional adaptation step. A uniform reduction of stiffness in response to external perturbations can prevent task execution, e.g., the end-effector may stop when encountering an obstacle. To avoid this problem, stiffness was modulated according to the directional relevance of each axis with respect to the generated sZFT of the Diffusion Model. The alignment factors below are defined relative to $\hat{\mathbf{p}}_0$ and $\hat{\mathbf{u}}_0 \hat{\theta}_0$ contained in $\hat{e}_t$ and $\hat{e}_r$, which are outputs of the Diffusion Model. Hence, the adaptation mechanism depended critically on the reconstructed sZFT: if the nominal free-space ZFT were used directly, the alignment factors would not reflect the equilibrium under contact, and directional adaptation would fail (Section IV).

For translations, the alignment factor

$$\psi_{t,i} = \frac{|\hat{e}_{t,i}|}{\|\hat{e}_t\|}, \quad i \in \{x, y, z\} \quad (32)$$

quantifies the contribution of each axis to the intended displacement. Analogously, for rotations, we define

$$\psi_{r,i} = \frac{|\hat{e}_{r,i}|}{\|\hat{e}_r\|}, \quad i \in \{x, y, z\}. \quad (33)$$

The relative translational and rotational factors $\rho_{t,i}$ and $\rho_{r,i}$ ensure that stiffness is decreased more strongly in directions that contribute little to the desired motion, while remaining close to the baseline stiffness along task-relevant axes:

$$\rho_{t,i} = 1 - \psi_{t,i}, \quad (34a)$$

$$\rho_{r,i} = 1 - \psi_{r,i}. \quad (34b)$$

The final adapted stiffness values $k_{t,i}$ and $k_{r,i}$ were implemented as:

$$k_{t,i} = K_{t,i,\max} - \rho_{t,i} k_{t,i}^*, \quad (35a)$$

$$k_{r,i} = K_{r,i,\max} - \rho_{r,i} k_{r,i}^*. \quad (35b)$$

As will be shown in the next section, this mechanism allowed the robot to move smoothly over obstacles in the parkour task and to achieve successful in peg-in-hole insertion.

An overview of the proposed model architecture and workflow can be seen in Figure 8.

IV. DEPLOYMENT OF DIFFUSION MODEL ON REAL ROBOT

The model was deployed in two robotic use-cases designed to evaluate its performance in contact-rich scenarios. In our experiments, we deliberately focused on Impedance Control rather than motion generation, since many real-world applications require active compliance, e.g., final assembly tasks where vision is occluded and only tactile sensing is available. The first was a parkour-style task requiring continuous interaction while traversing unknown obstacles. The second was a peg-in-hole insertion task with increasing geometric complexity.

The experiments were conducted on a KUKA LBR iiwa with seven DOFs, using KUKA's Fast Robot Interface (FRI) for torque control with 5 ms sample time. Built-in gravity and Coriolis/centrifugal compensation remained active throughout. Robot kinematics and dynamics were computed via the Exp[licit]-FRI interface⁴ [72]. A task-space impedance controller was implemented (Section II-B). The external wrenches were acquired with an ATI Gamma force/torque transducer, mounted on the robot flange. The code of the robot controller is publicly available on our public GitHub repository.

In both experiments, we compared a task-space impedance controller with constant stiffness values against our proposed diffusion-based directional stiffness adaptation. For all trials, the baseline translational stiffness was identical to that used during data collection and was set to $800 \frac{\text{N}}{\text{m}}$, while the baseline rotational stiffness was set to $150 \frac{\text{Nm}}{\text{rad}}$. The same damping design was applied throughout the experiments (Section II-B).

All workpieces were fabricated with a PRUSA i3 MK3 3D printer using PLA filament. The expected manufacturing error of the printer is on the order of 0.1 mm. All CAD models used in the experiments are available on our GitHub repository. Demonstration videos of both experiments are provided on the project website.

All plotted entities are expressed in the spatially-fixed robot base coordinate frame. The colors of the experimental plots are consistent with the colors of the coordinate axes in Figure 9.

A. Parkour

The first experiment consisted of a parkour-style task in which the robot traversed three obstacles while maintaining continuous contact with the table (Figure 9, left). The desired physical interaction behavior was demonstrated by a human

operator during data collection (Section III-C), who teleoperated the robot through the parkour task (Figure 6, right).

The nominal ZFT was generated by kinesthetic teaching of a number of key joint poses. A MATLAB script then computed the corresponding end-effector trajectory by smoothly interpolating between these configurations and manually shifting the trajectory along the negative z-direction to ensure constant contact with the table surface (see coordinate frame shown in Figure 9, left). It is important to note that the resulting nominal ZFT passed through the obstacles.

To protect the setup, stop conditions of $\|v\| = 0.24 \frac{\text{m}}{\text{s}}$ and $\|f_{\text{ext}}\| = 20 \text{ N}$ were defined. The constant-stiffness controller already failed at the first obstacle: the end-effector was stopped, energy accumulated in the virtual spring, and once this energy was released the end-effector accelerated into the velocity stop condition. Forces of up to 15 N were observed. In contrast, with Diffusion-Based Impedance Learning, the robot traversed all obstacles smoothly, i.e., without violating the velocity or force stop conditions. The stiffness adaptation mechanism prevented jamming in directions inconsistent with the ZFT while preserving rigidity along task-relevant axes.

Figure 10 shows the measured external wrenches together with the commanded translational and rotational stiffness values for the Diffusion-Based Impedance Learning. The bottom plots illustrate the relative translational and rotational factors (Equation (34)) and the work done by the external force along the virtual displacement (Equations (25) and (29)).

During the entire motion, the stiffness values changed dynamically in response to external wrenches. This continuous adaptation reflected the underlying Diffusion-Based Impedance Learning, which enabled the robot to remain compliant when unexpected contact occurred while maintaining sufficient stiffness to progress through the parkour.

Translational stiffness was adapted asymmetrically depending on the directional contribution of forces. Around 12 s, force peaks of 4 N along the y -axis and 15 N along the z -axes arose. Nevertheless, $K_{p,y}$ was reduced more strongly (reduced to near 400 N/m) than $K_{p,z}$ (remaining near 600 N/m) since the relative translational factor along the z -axis was high. A similar effect is visible near 30 s: while f_y reached an absolute of nearly 4 N with stiffness near 420 N/m, f_z was only 2 N yet stiffness dropped to about 200 N/m. These findings confirm that stiffness reduction was not simply proportional to force magnitude but was determined by the alignment with the sZFT, predicted by the Diffusion Model.

Rotational stiffness also adapted directionally, with reductions depending on which axis contributed less to the sZFT. Around 12 s, moments appeared simultaneously about the x - and y -axes. Despite higher magnitudes about the x -axis, the stiffness $K_{r,y}$ was reduced more strongly, showing that y -axis rotations were less important for the intended trajectory. This selective modulation allowed the robot to maintain stability while avoiding unnecessary resistance to contact.

B. Peg-In-Hole Insertion

The second experiment was a peg-in-hole insertion task with three different peg types: cylindrical, square, and star pegs. The

⁴<https://GitHub.com/explicit-robotics/Explicit-FRI>




Fig. 8: Overview of the proposed architecture and workflow. External wrenches and end-effector pose (position and unit-quaternion) are processed as time-series inputs to a conditional Transformer-based Diffusion Model. External wrenches are embedded as context tokens, while pose tokens are embedded with learned positional and diffusion-time encodings. The Diffusion Model reconstructs a simulated Zero-Force Trajectory (sZFT), representing a contact-consistent equilibrium. This reconstructed sZFT, together with the measured interaction signals, is passed to an energy-based stiffness estimator. Finally, directional stiffness adaptation is applied according to the directional relevance of each coordinate axis.

nominal clearance between each peg and its corresponding hole was 0.14 mm for the square peg and 0.20 mm for the circular and star pegs. All workpieces were fabricated with the same printer setup as for the parkour experiment. For all trials, task success was defined as completely inserting the peg.

For each peg type, the ZFT was generated from three manually demonstrated robot joint configurations, via kinesthetic teaching: a starting configuration, an initial contact (“touch”) configuration, and a final insertion configuration. These configurations were recorded by physically guiding the robot through the task. A smooth nominal ZFT was then generated using a MATLAB script by computing the corresponding end-effector poses and connecting them with a minimum-jerk interpolation. To induce a controlled pulling force during the final insertion phase, the equilibrium pose for final insertion was deliberately shifted along the negative z -direction. After completing the insertion, the robot moved straight upward along the z -axis (see coordinate frame in Figure 9, right). Importantly, during the touch pose the peg’s orientation was not aligned with the edges of the square and star pegs. As a result, the nominal ZFT did not encode any geometric alignment with the peg shape.

The three peg types represent increasing geometric complexity. The cylindrical peg is rotation-symmetric, which makes the final orientation about the end-effector z -axis irrelevant. In contrast, the square and star pegs require precise rotational alignment with an increasing number of edges, which significantly increases task complexity.

This trend is reflected in the constant-stiffness experiments.

The cylindrical peg succeeded in 30/30 trials. The square peg succeeded in 4/30 trials. The star peg failed in all 30 trials.

In practice, more complicated pegs like the square and star peg often require extensive manual tuning of the ZFT, of the impedance parameters, or both. This tuning is not only time-consuming but there is also no standard method that guarantees success. These results show that more complicated assembly processes require more advanced strategies.

In contrast, Diffusion-Based Impedance Learning achieved 30/30 successful insertions for all three peg types. This result was achieved by only demonstrating and interpolating three key poses per trajectory. Hence, no advanced method for trajectory generation was needed. Moreover, the outcome is particularly notable because the training data included only parkour and upper-limb rehabilitation data (Section III-C). No peg-in-hole demonstrations were used for training. Importantly, neither the Diffusion Model nor the impedance controller is provided with any information about the peg or hole geometry, pose, insertion axis, or clearance; all impedance adaptation was driven solely by measured interaction wrenches during the motion.

Across all three peg geometries, execution times remained comparable despite the different contact conditions and insertion complexities. Mean completion times for 30 trials were 10.8 ± 0.6 s for the circular peg, 10.2 ± 1.9 s for the square peg, and 12.3 ± 1.9 s for the star-shaped peg.

Figure 11 shows the stiffness adaptation for the star peg. Results for the other two pegs are provided in Appendix A. During all experiments, force and moment thresholds were set to 1 N and 1 Nm, respectively.




Fig. 9: Phases of the two experiments and the spatially fixed coordinate frames in which they are expressed. (Left) In the parkour experiments, three obstacles are encountered (numbered 1–3). (Right) During peg insertion, the robot moved from the starting pose to inclined contact with the hole (labeled “Ori”), followed by final insertion (with the start of insertion labeled “Ins”). The colors of the major axes of the coordinate frames are shown in the experimental plots (Figure 10 and 11).

As illustrated in Figure 11, stiffness values were reduced in response to external wrenches, with the amount of reduction governed by the relative translational and rotational factors. For most of the experiment, the relative translation factor along the z axis remained higher than along the y axis. For instance, at around 11 s, small forces of 0.5 N appeared along the y axis. Nevertheless, the stiffness $K_{p,y}$ was reduced to 600 N/m. During retraction of the peg at around 16 s, it can be seen that the relative translation factor along the z axis dropped and consequently the stiffness $K_{p,z}$ was reduced to 500 N/m. In contrast, the relative translation factor along the x axis increased and consequently the stiffness $K_{p,x}$ was increased to 700 N/m.

Due to the imperfect commanded ZFT, large moments about the x - and y -axes occurred during execution. The relative rotational factor about the z -axis remained low, leading to substantial reductions in $K_{r,z}$, which even reached 0 N/m between 8 s to 12 s. For the x -axis, the moments were initially small. When the moments increased significantly between 11 s to 16 s, the controller maintained high stiffness, since the relative rotational factor about the x -axis increased. A different behavior can be observed for the y -axis: during 11 s to 17 s, high moments reaching up to -10 N/m occurred. Since the relative rotational factor remained low, the stiffness dropped and reached 0 N/m.

C. Ablation Study

In addition to the two main controllers, we performed an ablation study to isolate the contributions of the reconstructed sZFT (Section III-A) and the directional impedance adaptation (Section III-E).

Importantly, the nominal ZFT was kept fixed throughout the ablation study. It was generated from simple kinesthetic demonstrations and was not optimized for the two tasks and pegs. Videos illustrating the kinesthetic teaching procedure for the parkour and peg-in-hole tasks are provided on the project website.

Specifically, we evaluated the following variants using the baseline impedance parameters ($800 \frac{\text{N}}{\text{m}}$ for translation and $150 \frac{\text{Nm}}{\text{rad}}$ for rotation):

- 1) Uniform stiffness adaptation without directional adaptation, using the reconstructed sZFT
- 2) Directional stiffness adaptation computed from the nominal ZFT

Uniform stiffness adaptation without directional modulation failed in both experimental scenarios. In the parkour task, the robot was blocked at the first obstacles, while none of the peg-in-hole insertions succeeded. Directional adaptation derived from the nominal ZFT also failed in contact-rich settings. In the parkour task, the robot exceeded safety limits and stopped. In the peg-in-hole experiments, this variant achieved 30/30 success only for the cylindrical peg, but failed for both the square and star pegs.

These results confirm that directional adaptation alone is insufficient if it is derived from the nominal ZFT. Instead, directional factors must be computed relative to the reconstructed sZFT for impedance adaptation to succeed.

To further evaluate how different uniform and non-uniform *fixed* impedance parameters influence task success in the peg-in-hole experiment, we additionally evaluated the following variants:

- a) Uniform high impedance: $800 \frac{\text{N}}{\text{m}}$ (translation), $150 \frac{\text{Nm}}{\text{rad}}$




Fig. 10: Measured external wrenches and commanded translational and rotational stiffness during the parkour experiment with stiffness adaptation. The relative translational and rotational factors show their respective contributions to stiffness modulation. The bottom row reports the performed translational and rotational work (Equations (25) and (29)). Shaded gray regions denote the three parkour obstacles. All quantities are expressed in the robot base frame; base-frame color coding and obstacle numbering are shown in Figure 9.

- (rotation)
- b) Uniform medium impedance: $400 \frac{\text{N}}{\text{m}}$ (translation), $50 \frac{\text{Nm}}{\text{rad}}$ (rotation)
- c) Uniform low impedance: $200 \frac{\text{N}}{\text{m}}$ (translation), $7 \frac{\text{Nm}}{\text{rad}}$ (rotation)
- d) Best-found, manually tuned non-uniform impedance: translation along z : $800 \frac{\text{N}}{\text{m}}$; translation along x and y : $\frac{\text{N}}{\text{m}}$; rotation about z : $5 \frac{\text{Nm}}{\text{rad}}$; rotation about x and y : $100 \frac{\text{Nm}}{\text{rad}}$

For a) (uniform high impedance), the cylindrical peg succeeded in 30/30 trials, the square peg in 4/30 trials, and the star peg failed in all trials. For b) (uniform medium impedance), yielded 30/30 success for the cylindrical peg, 13/30 success for the square peg, and 5/30 success for the star peg. For c) (uniform low impedance), it failed for all peg types, as the wrench required to orient the peg with the hole was insufficient.

For d) (manually tuned non-uniform impedance), the cylindrical peg succeeded in 30/30 trials, the square peg in 26/30 trials, and the star peg in 21/30 trials. These parameters were identified by an experienced impedance control programmer after testing 15 different parameter combinations. In contrast, inexperienced users would likely struggle to identify a single parameter set that performs robustly across all peg geometries.

Videos of all ablation experiments for the star peg example

are provided on our project website.

A summary of the ablation study for the peg-in-hole experiment can be found in Table III.

V. DISCUSSION

The experimental results confirm the central goal of this work: enabling robust contact-rich manipulation by bridging the information and the energy domains and therefore unifying state-of-the-art generative AI algorithms with model-based control. In particular, the success of peg-in-hole insertions with unseen tool and task geometries demonstrates that sZFT reconstruction generalizes beyond the training distribution.

Compared to prior learning-based approaches, our framework explicitly incorporates energy transformation, enabling stable impedance adaptation under contact. Compared to traditional model-based impedance controllers, it gains adaptability by reconstructing equilibrium trajectories directly from data. This ability to bridge domains accounts for the strong performance observed in unstructured environments.

Task success in contact-rich assembly is determined by the interplay between the nominal motion reference and the generated impedance, rather than by either component alone. For the peg-in-hole experiments of this paper, the initial contact (“touch”) phase is particularly important because it




Fig. 11: Measured external wrenches and commanded translational and rotational stiffness during the peg-in-hole star insertion experiment with stiffness adaptation. The relative translational and rotational factors show their respective contributions to stiffness modulation. The bottom row reports the performed translational and rotational work (Equations (25) and (29)). The shaded gray region denotes the insertion phase. All quantities are expressed in the robot base frame; base-frame color coding is shown in Figure 9.

TABLE III: Peg-in-hole success rates for the ablation study. All results are reported over 30 trials.

Controller Variant	Cylindrical Peg	Square Peg	Star Peg
<i>Time-varying Impedance</i>			
Diffusion Model (with directional adaptation)	100%	100%	100%
Diffusion Model w/o directional adaptation	0%	0%	0%
Directional adaptation w/o DM (nominal ZFT)	100%	0%	0%
<i>Fixed Impedance</i>			
Uniform high (baseline)	100%	13%	0%
Uniform medium	100%	43%	17%
Uniform low	0%	0%	0%
Best-found non-uniform	100%	87%	70%

establishes the conditions for the subsequent orientation and insertion phases. Nevertheless, the proposed framework exhibited considerable robustness to variations in the nominal trajectory. For both the cylindrical and square pegs, successful insertion was achieved for lateral offsets of up to approximately ± 3 mm in the x and y directions. Even for the more demanding star-shaped peg, which requires accurate rotational alignment, successful insertion remained possible for offsets of approximately ± 1 mm. This observation is consistent with the findings reported in [73].

A key element of the framework is directional stiffness adaptation. This mechanism is generic and can be tailored to applications in other domains. One example is robotic-assisted upper-limb therapy, where stiffness must remain high along

non-movement directions, while stiffness along the intended movement direction should adapt depending on whether the patient can execute the motion independently. In such settings, the same framework can be used by solely adapting the directional stiffness adaptation.

The force and moment thresholds in the stiffness estimator affect the responsiveness of impedance adaptation. In our experiments, values of 1 N and 1 Nm were selected to match the noise characteristics of the ATI Gamma force/torque sensor. These thresholds are not task-specific and can be adjusted for different sensors or applications.

At present, the adaptive impedance controller outputs diagonal, positive semidefinite stiffness matrices. This design simplifies analysis and implementation but omits cross-coupling

effects that have proven beneficial in traditional solutions such as remote-center compliance (RCC) devices [74], [75]. For more complex assembly tasks, integrating coupled stiffness parameters may further enhance robustness and alignment accuracy. Exploring this extension is an important direction for future work.

In this study, the baseline impedance behavior used during teleoperation was preserved during deployment, ensuring consistency between the demonstrations used for training and the interaction conditions encountered during execution. Future work should investigate robustness to changes in baseline impedance, including deployment with substantially different stiffness settings than those used during data collection. Such experiments would provide further insight into the ability of the learned model to generalize across different compliant behaviors and controller parameterizations.

The proposed approach guarantees passivity when stiffness is reduced: lowering stiffness dissipates stored potential energy and ensures that interaction with a passive environment remains output strictly passive. However, when stiffness is increased back toward the baseline, passivity is not generally guaranteed, as such changes can inject additional energy whenever displacements are nonzero. In these cases, the controller can be combined with previous work on passive impedance shaping [2] to enforce strict passivity.

In our work, we adopt a Transformer-based Diffusion Model due to its ability to capture long range temporal dependencies and flexibly condition on multimodal interaction signals. Future work should investigate a systematic comparison with alternative architectures such as convolutional or hybrid models under matched parameter and computational budgets.

Overall, our findings demonstrate that Diffusion-Based Impedance Learning is highly effective in the experimental setups investigated. While this provides strong evidence for the viability of the approach, further studies across a wider range of tasks and environments are needed to fully establish its generality. Nevertheless, the framework offers a promising design principle for stable, safe, and adaptable robot interaction in other domains including industrial assembly, physical Human-Robot Interaction, and robotic-assisted rehabilitation.

VI. CONCLUSION

This work introduced *Diffusion-Based Impedance Learning*, a framework that unifies generative trajectory models in the information domain with adaptive impedance regulation in the energy domain. By reconstructing simulated Zero-Force Trajectories (sZFTs) from contact-perturbed displacements and external wrench signals, the method enables real-time impedance adaptation for contact-rich manipulation.

The framework consistently outperformed fixed and heuristic strategies in experiments on a KUKA LBR iiwa. It enabled smooth traversal of obstacle-rich parkour environments where conventional impedance control failed, and achieved 30/30 success rates in cylindrical, square, and star peg-in-hole insertions without any task-specific demonstrations in the training data. The experiments in this work intentionally withhold geometric information from the controller, targeting scenarios in

which such knowledge is unavailable or unreliable due to unstructured environments. Diffusion-Based Impedance Learning should therefore be understood as a general, interaction-driven alternative rather than a replacement for highly engineered solutions in fully structured assembly lines.

Beyond these results, Diffusion-Based Impedance Learning establishes a general interface between learning and model-based control. It extends the role of generative models from visuomotor trajectory planning to physical interaction, demonstrating that contact-rich manipulation can be learned from relatively small datasets while maintaining real-time control performance.

Looking ahead, we aim to extend this framework from single-arm to dual-arm setups. While this work focuses on impedance learning and used a fixed ZFT, Diffusion-Based Impedance Learning can naturally be embedded in hierarchical architectures that couple high-level visuomotor motion policies with low-level impedance regulation.

REFERENCES

- [1] N. Hogan, "Physical interaction via dynamic primitives," in *Geometric and numerical foundations of movements*. Springer, 2017, pp. 269–299.
- [2] J. Lachner, F. Allmendinger, S. Stramigioli, and N. Hogan, "Shaping impedances to comply with constrained task dynamics," *IEEE Transactions on Robotics*, vol. 38, no. 5, pp. 2750–2767, 2022.
- [3] J. Lachner, M. C. Nah, and N. Hogan, "A physically consistent stiffness formulation for contact-rich manipulation," *The International Journal of Robotics Research*, vol. 0, no. 0, p. 02783649251393332, 0.
- [4] J. Lachner, F. Allmendinger, E. Hobert, N. Hogan, and S. Stramigioli, "Energy budgets for coordinate invariant robot control in physical human-robot interaction," *The International Journal of Robotics Research*, vol. 40, no. 8-9, pp. 968–985, Aug. 2021.
- [5] N. Hogan, "Stable execution of contact tasks using impedance control," in *Proceedings. 1987 IEEE International Conference on Robotics and Automation*, vol. 4. IEEE, 1987, pp. 1047–1054.
- [6] J. Lachner, "A geometric approach to robotic manipulation in physical human-robot interaction," Ph.D. dissertation, University of Twente, 2022.
- [7] S. Stramigioli, Ed., *Energy-Aware Robotics*, ser. Mathematical Control Theory I. Switzerland: Springer International Publishing, 2015, vol. 461.
- [8] H. v. Helmholtz, "II. über einige gesetze der vertheilung elektrischer ströme in körperlichen leitern mit anwendung auf die thierisch-elektrischen versuche," *Annalen der Physik und Chemie*, vol. 89, no. 6, pp. 211–233, 1853, english translation: "Some laws concerning the distribution of electrical currents in conductors with applications to experiments on animal electricity".
- [9] N. Hogan, "A general actuator model based on nonlinear equivalent networks," *IEEE/ASME Transactions on Mechatronics*, vol. 19, no. 6, pp. 1929–1939, 2014.
- [10] M. C. Nah, J. Lachner, and N. Hogan, "Robot control based on motor primitives: A comparison of two approaches," *The International Journal of Robotics Research*, vol. 43, no. 12, pp. 1959–1991, 2024.
- [11] M. C. Nah, J. Lachner, F. Tessari, and N. Hogan, "On the modularity of elementary dynamic actions," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 1398–1405.
- [12] N. Hogan, "Impedance control-an approach to manipulation. i-theory. ii-implementation. iii-applications," *ASME Journal of Dynamic Systems and Measurement Control B*, vol. 107, pp. 1–24, 1985.
- [13] J. Collins, D. Howard, and J. Leitner, "Quantifying the reality gap in robotic manipulation tasks," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 6706–6712.
- [14] M. Posa, C. Cantu, and R. Tedrake, "A direct method for trajectory optimization of rigid bodies through contact," *The International Journal of Robotics Research*, vol. 33, no. 1, pp. 69–81, 2014.
- [15] Z. Manchester and S. Kuindersma, "Variational contact-implicit trajectory optimization," in *Robotics Research: The 18th International Symposium ISRR*. Springer, 2020, pp. 985–1000.

- [16] T. Pang, H. T. Suh, L. Yang, and R. Tedrake, "Global planning for contact-rich manipulation via local smoothing of quasi-dynamic contact models," *IEEE Transactions on Robotics*, 2023.
- [17] T. Marcucci, R. Deits, M. Gabiccini, A. Bicchi, and R. Tedrake, "Approximate hybrid model predictive control for multi-contact push recovery in complex environments," in *2017 IEEE-RAS 17th international conference on humanoid robotics (Humanoids)*. IEEE, 2017, pp. 31–38.
- [18] F. R. Hogan and A. Rodriguez, "Reactive planar non-prehensile manipulation with hybrid model predictive control," *The International Journal of Robotics Research*, vol. 39, no. 7, pp. 755–773, 2020.
- [19] M. Anitescu and F. A. Potra, "Formulating dynamic multi-rigid-body contact problems with friction as solvable linear complementarity problems," *equivalentlinear Dynamics*, vol. 14, pp. 231–247, 1997.
- [20] A. Aydinoglu and M. Posa, "Real-time multi-contact model predictive control via admm," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 3414–3421.
- [21] H. J. T. Suh, T. Pang, and R. Tedrake, "Bundled gradients through contact via randomized smoothing," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4000–4007, 2022.
- [22] H. J. T. Suh, T. Pang, T. Zhao, and R. Tedrake, "Dexterous contact-rich manipulation via the contact trust region," 2025. [Online]. Available: <https://arxiv.org/abs/2505.02291>
- [23] Y. Li, G. Ganesh, N. Jarrassé, S. Haddadin, A. Albu-Schaeffer, and E. Burdet, "Force, impedance, and trajectory learning for contact tooling and haptic identification," *IEEE Transactions on Robotics*, vol. 34, no. 5, pp. 1170–1182, 2018.
- [24] Y. Li, L. Zheng, Y. Wang, E. Dong, and S. Zhang, "Impedance learning-based adaptive force tracking for robot on unknown terrains," *IEEE Transactions on Robotics*, vol. 41, pp. 1404–1420, 2025.
- [25] S. Shaw, B. Abbatematteo, and G. Konidaris, "Rmps for safe impedance control in contact-rich manipulation," 2022. [Online]. Available: <https://arxiv.org/abs/2109.12103>
- [26] A. S. Anand, J. T. Gravdahl, and F. J. Abu-Dakka, "Model-based variable impedance learning control for robotic manipulation," *Robotics and Autonomous Systems*, vol. 170, p. 104531, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0921889023001707>
- [27] Y. Michel, M. Saveriano, and D. Lee, "A novel safety-aware energy tank formulation based on control barrier functions," *IEEE Robotics and Automation Letters*, vol. 9, no. 6, pp. 5206–5213, 2024.
- [28] F. Abi-Farraj, T. Osa, N. P. J. Peters, G. Neumann, and P. R. Giordano, "A learning-based shared control architecture for interactive task execution," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 329–335.
- [29] M. Najafi, K. Adams, and M. Tavakoli, "Robotic learning from demonstration of therapist's time-varying assistance to a patient in trajectory-following tasks," in *International Conference on Rehabilitation Robotics (ICORR)*, 2017, pp. 888–894.
- [30] E. Theodorou, J. Buchli, and S. Schaal, "A generalized path integral control approach to reinforcement learning," *The Journal of Machine Learning Research*, vol. 11, pp. 3137–3181, 2010.
- [31] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal, "Stomp: Stochastic trajectory optimization for motion planning," in *2011 IEEE international conference on robotics and automation*. IEEE, 2011, pp. 4569–4574.
- [32] F. Stulp, E. A. Theodorou, and S. Schaal, "Reinforcement learning with sequences of motion primitives for robust manipulation," *IEEE Transactions on robotics*, vol. 28, no. 6, pp. 1360–1370, 2012.
- [33] G. Williams, A. Aldrich, and E. A. Theodorou, "Model predictive path integral control: From theory to parallel computation," *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 2, pp. 344–357, 2017.
- [34] Q. Yang, A. Dürr, E. A. Topp, J. A. Stork, and T. Stoyanov, "Variable impedance skill learning for contact-rich manipulation," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 8391–8398, 2022.
- [35] A. B. Tahmaz, R. Prakash, and J. Kober, "Impedance primitive-augmented hierarchical reinforcement learning for sequential tasks," 2025. [Online]. Available: <https://arxiv.org/abs/2508.19607>
- [36] N. Hogan, "Impedance control: An approach to manipulation — part i: Theory," *Journal of Dynamic Systems, Measurement, and Control*, vol. 107, no. 1, pp. 1–7, 1985.
- [37] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," *arXiv preprint arXiv:2303.04137*, 2023.
- [38] M. Braun, N. Jaquier, L. Rozo, and T. Asfour, "Riemannian flow matching policy for robot motion learning," 10 2024, pp. 5144–5151.
- [39] F. Zhang and M. Gienger, "Affordance-based robot manipulation with flow matching," 2025. [Online]. Available: <https://arxiv.org/abs/2409.01083>
- [40] M. Shridhar, L. Manuelli, and D. Fox, "Perceiver-actor: A multi-task transformer for robotic manipulation," in *Conference on Robot Learning (CoRL)*, 2022. [Online]. Available: <https://arxiv.org/abs/2209.05451>
- [41] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi, Q. Vuong, T. Kollar, B. Burchfiel, R. Tedrake, D. Sadigh, S. Levine, P. Liang, and C. Finn, "Openvla: An open-source vision-language-action model," *arXiv preprint arXiv:2406.09246*, 2024.
- [42] T. L. Team, J. Barreiros, A. Beaulieu, A. Bhat, R. Cory, E. Cousineau, H. Dai, C.-H. Fang, K. Hashimoto, M. Z. Irshad, M. Itkina, N. Kuppaswamy, K.-H. Lee, K. Liu, D. McConachie, I. McMahon, H. Nishimura, C. Phillips-Grafflin, C. Richter, P. Shah, K. Srinivasan, B. Wulfe, C. Xu, M. Zhang, A. Alspach, M. Angeles, K. Arora, V. C. Guizilini, A. Castro, D. Chen, T.-S. Chu, S. Creasey, S. Curtis, R. Denitto, E. Dixon, E. Dusel, M. Ferreira, A. Goncalves, G. Gould, D. Guoy, S. Gupta, X. Han, K. Hatch, B. Hathaway, A. Henry, H. Hochsztein, P. Horgan, S. Iwase, D. Jackson, S. Karamcheti, S. Keh, J. Masterjohn, J. Mercat, P. Miller, P. Mitiguy, T. Nguyen, J. Nimmer, Y. Noguchi, R. Ong, A. Onol, O. Pfannenstiel, R. Poyner, L. P. M. Rocha, G. Richardson, C. Rodriguez, D. Seale, M. Sherman, M. Smith-Jones, D. Tago, P. Tokmakov, M. Tran, B. V. Hoorick, I. Vasiljevic, S. Zakharov, M. Zolotas, R. Ambrus, K. Fetzter-Borelli, B. Burchfiel, H. Kress-Gazit, S. Feng, S. Ford, and R. Tedrake, "A careful examination of large behavior models for multitask dexterous manipulation," 2025. [Online]. Available: <https://arxiv.org/abs/2507.05331>
- [43] Y. Hou, Z. Liu, C. Chi, E. Cousineau, N. Kuppaswamy, S. Feng, B. Burchfiel, and S. Song, "Adaptive compliance policy: Learning approximate compliance for diffusion guided control," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 4829–4836.
- [44] Y. Park and P. Agrawal, "Using apple vision pro to train and control robots," 2024. [Online]. Available: <https://github.com/Improbable-AI/VisionProTeleop>
- [45] M. Janner, Y. Du, J. B. Tenenbaum, and S. Levine, "Planning with diffusion for flexible behavior synthesis," in *Proceedings of the 39th International Conference on Machine Learning (ICML)*, 2022.
- [46] A. Nichol and P. Dhariwal, "Improved denoising diffusion probabilistic models," *arXiv preprint arXiv:2102.09672*, 2021.
- [47] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *arXiv preprint arXiv:2006.11239*, 2020.
- [48] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, F. Bach and D. Blei, Eds., vol. 37. Lille, France: PMLR, 07–09 Jul 2015, pp. 2256–2265.
- [49] P. Holderrieth and E. Erives, "An introduction to flow matching and diffusion models," *arXiv preprint arXiv:2506.02070*, 2025.
- [50] X. Liu, K. Y. Ma, C. Gao, and M. Z. Shou, "Diffusion models in robotics: A survey," *TechRxiv*, April 2025.
- [51] U. A. Mishra and Y. Chen, "Reorientdiff: Diffusion model based reorientation for object manipulation," *arXiv preprint arXiv:2303.12700*, 2023.
- [52] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," *arXiv preprint arXiv:2010.02502*, 2020.
- [53] R. Wolf, Y. Shi, S. Liu, and R. Rayyes, "Diffusion models for robotic manipulation: A survey," *arXiv preprint arXiv:2504.08438*, 2025.
- [54] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, E. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017.
- [55] Y. Liang, H. Wen, Y. Nie, Y. Jiang, M. Jin, D. Song, S. Pan, and Q. Wen, "Foundation models for time series analysis: A tutorial and survey," in *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, ser. KDD '24. ACM, Aug. 2024, p. 6555–6565.
- [56] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations*, 2021.
- [57] A. Jaegle, F. Gimeno, A. Brock, O. Vinyals, A. Zisserman, and J. Carreira, "Perceiver: General perception with iterative attention," in *Proceedings of the 38th International Conference on Machine Learning*,

- ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 18–24 Jul 2021, pp. 4651–4664.
- [58] T. Wang, A. Roberts, D. Hesslow, T. L. Scao, H. W. Chung, I. Beltagy, J. Launay, and C. Raffel, “What language model architecture and pretraining objective work best for zero-shot generalization?” in *Proceedings of the 39th International Conference on Machine Learning (ICML)*, vol. 162. PMLR, 2022, pp. 23 671–23 694.
- [59] T. Lin, Y. Wang, X. Liu, and X. Qiu, “A survey of transformers,” *AI Open*, vol. 3, pp. 111–132, 2022.
- [60] A. Kramberger, E. Shahriari, A. Gams, B. Nemec, A. Ude, and S. Haddadin, “Passivity based iterative learning of admittance-coupled dynamic movement primitives for interaction with changing environments,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 6023–6028.
- [61] N. Hogan, “A general actuator model based on nonlinear equivalent networks,” in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 2647–2653.
- [62] —, “Impedance and interaction control,” in *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2015, pp. 1–7.
- [63] J. Lachner, F. Allmendinger, E. Hobert, N. Hogan, and S. Stramigioli, “Energy budgets for coordinate invariant robot control in physical human–robot interaction,” *The International Journal of Robotics Research*, vol. 40, no. 8-9, pp. 968–985, 2021.
- [64] A. Albu-Schaffer, C. Ott, U. Frese, and G. Hirzinger, “Cartesian impedance control of redundant robots: recent results with the DLR-light-weight-arms,” *IEEE International Conference on Robotics and Automation*, pp. 3704–3709, 2003.
- [65] K. Shoemake, “Animating rotation with quaternion curves,” *SIGGRAPH Comput. Graph.*, vol. 19, no. 3, p. 245–254, Jul. 1985.
- [66] S. Schaal, “Is imitation learning the route to humanoid robots?” *Trends in cognitive sciences*, vol. 3, no. 6, pp. 233–242, 1999.
- [67] A. J. Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal, “Dynamical movement primitives: learning attractor models for motor behaviors,” *Neural computation*, vol. 25, no. 2, pp. 328–373, 2013.
- [68] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, “Diffusion policy: Visuomotor policy learning via action diffusion,” *The International Journal of Robotics Research*, 2024.
- [69] C. M. Bishop, “Training with noise is equivalent to tikhonov regularization,” *Neural Computation*, vol. 7, no. 1, pp. 108–116, 01 1995.
- [70] J. Lachner, F. Tessari, A. M. W. Jr., M. C. Nah, and N. Hogan, “Divide et impera: Decoding impedance strategies for robotic peg-in-hole assembly,” *arXiv preprint arXiv:2410.01054*, 2025.
- [71] J. Carvalho, A. T. Le, M. Baierl, D. Koert, and J. Peters, “Motion planning diffusion: Learning and planning of robot motions with diffusion models,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 1916–1923.
- [72] J. Lachner, M. C. Nah, S. Stramigioli, and N. Hogan, “Exp [licit] an educational robot modeling software based on exponential maps,” in *2024 IEEE International Conference on Advanced Intelligent Mechatronics (AIM)*. IEEE, 2024, pp. 1359–1366.
- [73] J. Lachner, M. C. Nah, F. Tessari, and N. Hogan, “Elementary dynamic actions: key structures for contact-rich manipulation,” in *IROS 2023 Workshop on Leveraging Models for Contact-Rich Manipulation*, 2023.
- [74] D. E. Whitney, “Quasi-Static Assembly of Compliantly Supported Rigid Parts,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 104, no. 1, pp. 65–77, 03 1982. [Online]. Available: <https://doi.org/10.1115/1.3149634>
- [75] S. Hirai and K. Tanie, “Remote center compliance with adaptive compliance control,” in *Proceedings of the IEEE International Conference on Robotics and Automation*. IEEE, 1990, pp. 468–473.

APPENDIX A

PEG-IN-HOLE EXPERIMENTS

This appendix complements the main results by reporting stiffness adaptation for the circular and square pegs. The corresponding results for the star peg are already presented in the main text (Figure 11). The experimental setup, controllers, and parameters matched the main text: baseline stiffness $K_{t,i,\max} = 800 \text{ N/m}$ and $K_{r,i,\max} = 150 \text{ Nm/rad}$, damping as in Appendix II-B1, and force/moment thresholds of 1 N and 1 Nm. All quantities are expressed in the robot base frame.

For the circular peg, rotational symmetry about the end-effector z -axis makes the final yaw alignment irrelevant. Stiffness modulation is therefore dominated by translational factors, with only minor rotational adjustments. For the square peg, alignment with the edges increased sensitivity to off-axis contact. This yielded stronger directional reductions in both translation and rotation when axes contributed little to the sZFT, while stiffness remained high along task-relevant directions.



He began his professional experience as a Student Assistant at KIT, working on human pose estimation, and as a Working Student and Bachelor’s Candidate with Schaeffler, Karlsruhe, Germany, where he worked on machine learning for vehicle control systems. He then worked as a Data Scientist with Mercedes-Benz, Sindelfingen, Germany, focusing on scene understanding and end-to-end neural networks for autonomous driving, and as a Robotics Engineer with Bosch, Plymouth, MI, USA, where he implemented a mobile robot system for radar sensor validation. Since October 2025, he has been with the Junior Managers Program (AI/IT Track), Robert Bosch GmbH, Stuttgart, Germany, a program that prepares participants for leadership responsibility. His research interests include robotics, agentic AI, diffusion models, and physical AI. Mr. Geiger is Founder and Lead of the Studopolis Think Tank and was previously a Member of the Board of MIT VISTA.



Tamim Asfour (Fellow, IEEE) is Full Professor of Humanoid Robotics at Karlsruhe Institute of Technology (KIT), Germany. He directs the High Performance Humanoid Technologies Lab (H2T) at the Institute of Anthropomatics and Robotics. He received a diploma in Electrical Engineering from the University of Karlsruhe and a Ph.D. degree from KIT. His research focuses on the engineering humanoid robot systems that integrates artificial intelligence, informatics, and mechatronics to perform versatile tasks in the real world while learning from humans, experience and interaction with the environment. Tamim is the developer of the ARMAR humanoid robot family. He has been a visiting professor at Georgia Tech, Tokyo University of Agriculture and Technology, and the National University of Singapore. He is a IEEE Fellow, Editor-in-Chief of the Robotics and Automation Letters (RA-L), Founding Editor-in-Chief of the IEEE-RAS Humanoids Conference Editorial Board, the Spokesperson of the Robotics Institute Germany (RIG), and the scientific spokesperson of the KIT Center “Information · Systems · Technologies” (KCIST).




Fig. 12: Cylindrical peg insertion with stiffness adaptation. External wrenches, commanded translational and rotational stiffness, and relative motion factors are shown in the robot base frame. Base-frame color coding is given in Figure 9.



Neville Hogan (Member, IEEE) is Sun Jae Professor of Mechanical Engineering and Professor of Brain and Cognitive Sciences at the Massachusetts Institute of Technology. He earned a Diploma in Engineering (with distinction) from Dublin Institute of Technology and M.S., Mechanical Engineer and Ph.D. degrees from MIT. He joined MIT's faculty in 1979 and presently directs the Newman Laboratory for Biomechanics and Human Rehabilitation. He co-founded Interactive Motion Technologies, subsequently part of Bionik Laboratories. His research

includes robotics, motor neuroscience, and rehabilitation engineering, emphasizing the control of physical contact and dynamic interaction. Awards include: Honorary Doctorates from Delft University of Technology and Dublin Institute of Technology; the Silver Medal of the Royal Academy of Medicine in Ireland; the Saint Patrick's Day Medal for Academia from Science Foundation Ireland; the Henry M. Paynter Outstanding Investigator Award and the Rufus T. Oldenburger Medal from the American Society of Mechanical Engineers, Dynamic Systems and Control Division; and from the Institute of Electrical and Electronics Engineers, the Academic Career Achievement Award from the Engineering in Medicine and Biology Society and the Pioneer in Robotics Award from the Robotics and Automation Society.



Johannes Lachner (Member, IEEE) is an incoming Assistant Professor with the Elmore Family School of Electrical and Computer Engineering at Purdue University, West Lafayette, IN, USA. He is currently a Postdoctoral Researcher with the Department of Mechanical Engineering and the Department of Brain and Cognitive Sciences at the Massachusetts Institute of Technology (MIT). He received the M.Eng. double degree (with distinction) in Mechatronic Systems from the University of Ulster, Belfast, UK, and the University of Applied

Sciences Augsburg, Germany, and the Ph.D. degree (with highest distinction) in Robotics from the University of Twente, Enschede, The Netherlands. Prior to MIT, he was a Senior Research Engineer with the Corporate Research Department, KUKA Robotics GmbH, where he invented numerous patents in robotics, several of which have been transferred into commercial products. His research focuses on rehabilitation robotics, motor neuroscience, and AI for physical human-robot interaction, with an emphasis on geometric methods for contact-rich robot control and robot learning from human physical interaction. He received the MIT–Novo Nordisk AI Fellowship Award, the IROS Best Paper Award 2024, and two consecutive TU Munich Global Incentive Fund awards. He was a finalist for the European Georges Giralt Ph.D. Award for the best European Ph.D. thesis in robotics.




Fig. 13: Square peg insertion with stiffness adaptation. External wrenches, commanded translational and rotational stiffness, and relative motion factors are shown in the robot base frame. Base-frame color coding is given in Figure 9.