

Tactile-based Active Inference for Force-Controlled Peg-in-Hole Insertions

Tatsuya Kamijo¹, Ixchel G. Ramirez-Alpizar^{*2}, Enrique Coronado², Gentiane Venture^{1,3}

Abstract— Reinforcement Learning (RL) has shown great promise for efficiently learning force control policies in peg-in-hole tasks. However, robots often face difficulties due to visual occlusions by the gripper and uncertainties in the initial grasping pose of the peg. In this paper, we propose a robust tactile insertion policy that can align the tilted peg with the hole using active inference, without the need for extensive training on large datasets. Our approach employs a dual-policy architecture: one policy focuses on insertion, integrating force control and RL to guide the object into the hole, while the other policy performs active inference based on tactile feedback to align the tilted peg with the hole. In real-world experiments, our dual-policy architecture achieved 90% success rate into a hole with a clearance of less than 0.1 mm, significantly outperforming previous methods that lack tactile sensory feedback (5%). To assess the generalizability of our alignment policy, we conducted experiments with five different pegs, demonstrating its effective adaptation to multiple objects.

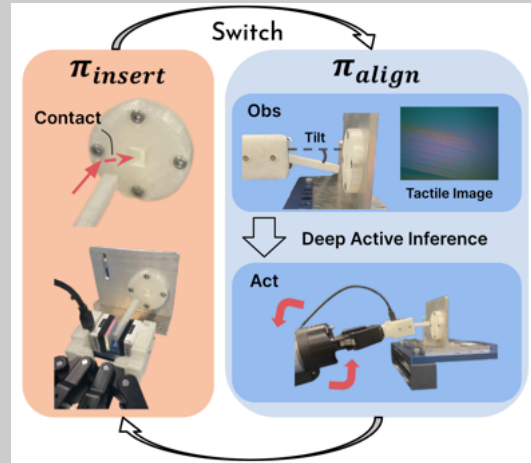


Fig. 1: Proposed dual-policy system for robust insertion. It leverages deep active inference for tactile pose alignment.

I. INTRODUCTION

Expanding industrial robot use requires safe, precise contact-rich manipulation like peg insertion, which is difficult due to grasp uncertainty, visual occlusion, and complex object–environment interactions. Early work on peg-in-hole tasks leverage mathematical or geometrical model of the environment [1], [2]. In recent years, learning-based approaches have shown great success in various manipulation tasks by learning complex behaviors, thus avoiding the need to model physical interactions between the manipulated object and the environment [3]. Reinforcement Learning (RL) methods have especially been shown to be effective in object insertion tasks. While RL policies can successfully learn peg-in-hole insertion tasks, most of the work either consider the peg to be part of the robot’s end-effector or fix it to the gripper, assuming no in-hand slippage occurs during the insertion process [4]. This paper addresses this assumption by using feedback signal from the vision-based tactile sensors.

We propose a safe tactile insertion policy that leverages vision-based tactile sensors for robot pose alignment. Our method combines the strengths of force control and active inference, a promising neuroscience theory that unifies perception and action under a single objective of minimizing

free energy. In our system, the agent predicts tactile sensation from its own internal state, namely the tilt of the peg. This internal state (i.e. inferred tilt) is updated in a way that minimizes discrepancies between predicted and current tactile image. The key insight is that by formulating tactile perception under free energy minimization problem, the active inference agent can adapt to diverse objects without extensive training on vast datasets. Our approach consists of two policies as depicted in Fig. 1:

- **Insertion policy** that executes force-controlled policy learned by RL in simulation to insert the grasped peg.
- **Alignment policy** that employs active inference based on tactile feedback to change the robot pose.

The insertion policy handles the main process of inserting the peg. When the peg slips while in contact with the environment, the alignment policy takes over. It executes active inference to adjust the robot’s end effector pose in a way that aligns the peg with the hole.

This paper presents three main contributions. First, we propose a novel deep active inference approach to tactile pose alignment. Building on the free energy principle, this approach adapts to various objects without a pre-trained model built on a large dataset or prior knowledge of their shape. Second, we introduce a self-data augmentation method to realize deep active inference without collecting extensive data in the real world. Third, we integrate a force-controlled insertion policy with an active inference–based tactile alignment policy, validated on a UR5e robot with a Gelsight Mini sensor [5].

*corresponding author

¹Department of Mechanical Engineering, Faculty of Engineering, The University of Tokyo, Japan.

²Task Intelligent Research Team, Intelligent Systems Research Institute, National Institute of Advanced Industrial Science and Technology (AIST), Tokyo, Japan. E-mail: ixchel-ramirezalpizar@aist.go.jp

³CNRS-AIST JRL (Joint Robotics Laboratory), National Institute of Advanced Industrial Science and Technology (AIST), Tsukuba, Japan.

II. BACKGROUND AND RELATED WORK

A. Learning Contact-Rich Manipulation Tasks

Learning force control by RL has shown significant promise for handling complex, contact-rich manipulation tasks [3]. Beltran-Hernandez et al. [6] proposes a learning-based force control framework that combines RL techniques with traditional force control methods. The authors use transfer-learning techniques (Sim2Real) and domain randomization to close the reality gap. Although these methods have advanced work in object-insertion tasks, they typically assume that the peg is rigidly fixed to the robot's end-effector to prevent in-hand slippage during the insertion process. However, visual occlusions by the gripper often make it difficult for a robot to receive feedback about the physical state around the end-effector.

Recent advances in high-resolution vision-based tactile sensors [5] have enabled robots to obtain rich tactile information directly from their grippers. Approaches leveraging tactile information for manipulation tasks include tactile mapping [7], pose estimation [8] and supervised learning [9]. Learning a tactile insertion policy through RL is one way to utilize this tactile information. Dong et al. [10] study different design choices for tactile-based feedback insertion policies and propose an RL-based insertion policy that incorporates curriculum training and tactile flow representation. They also argue that raw tactile RGB images contain detailed object-specific features, which can cause a learning-based agent to easily overfit to the training objects. While tactile flow representation is effective in dynamical situation where markers on the surface change position in response to actions, it is not the best option when you need to obtain the static state of the peg such as tilt. Contrary to this, our work demonstrates that static RGB images can also generalize across various objects when you use active inference in combination with contact area estimation.

B. Free Energy Principle and Active Inference

The free energy principle is a theoretical framework from neuroscience that proposes that living systems seek to minimize a statistical quantity known as free energy [11]. Building upon this theory, body perception can be modeled as inferring the hidden state z from the sensory observation \mathbf{o} . The objective in perception is to find a hidden state z that is consistent with Bayes' rule given an observation \mathbf{o} :

$$p(z|\mathbf{o}) = \frac{p(\mathbf{o}|z)p(z)}{p(\mathbf{o})}. \quad (1)$$

Here, $p(z|\mathbf{o})$ is the posterior probability of the internal state z given a sensory observation \mathbf{o} . To avoid calculating the marginal likelihood $p(\mathbf{o}) = \int_z p(\mathbf{o}|z)p(z)dz$, a reference distribution, also called the recognition density, $q(z)$ is introduced. By minimizing the Kullback-Leibler divergence D_{KL} between the true posterior and the recognition density,

the hidden state z can be optimized.

$$D_{\text{KL}}(q(z) \parallel p(z|\mathbf{o})) = \int q(z) \ln \frac{q(z)}{p(\mathbf{o}, z)} dz = F + \ln p(\mathbf{o}) \quad (2)$$

Rearranging (2), free energy can be written as

$$F = D_{\text{KL}}(q(z) \parallel p(z|\mathbf{o})) - \ln p(\mathbf{o}). \quad (3)$$

The first term is associated with perception, where the agent performs Bayesian inference to better model its environment. The second term $-\ln p(\mathbf{o})$ represents the sensory surprise and is related to the action component, where the agent acts to minimize this sensory surprise and thereby achieve desired sensory outcomes.

Active inference is a theory that outlines the specific mechanism by which a system acts on its world to change sensory inputs, thereby minimizing free energy. From Equation (3), minimizing D_{KL} is equivalent to minimizing free energy F when the observation \mathbf{o} is fixed. This is known as perceptual inference. Perceptual inference can tightly constrain the level of surprise by approximating the world but cannot lessen the surprise within the observations themselves $-\ln p(\mathbf{o})$. In active inference, the agent can decrease this sensory surprise by acting upon the environment to change sensory observations \mathbf{o} , thereby minimizing free energy. Thus, both perception and action can be done under the single free energy minimization:

$$z = \arg \min_z F(z, \mathbf{o}), \quad (4)$$

$$\mathbf{a} = \arg \min_{\mathbf{a}} F(z, \mathbf{o}(\mathbf{a})). \quad (5)$$

In applying active inference to robotics [12], generative functions are critical. Generative functions are learned mappings that approximate the hidden state dynamics and how sensory observations are generated from this hidden state, essentially serving as the model's understanding of the underlying causal structure of its environment. The work in [13] presents a deep learning approach to learn the generative function, proposing PixelAI that deals with high-dimensional RGB input for body perception and action. This type of learning-based method called Deep Active Inference (Deep AIF) in general has extended the active inference framework to high-dimensional inputs by learning generative functions [14], [15]. As for the choice of hidden state, previous work that apply active inference to robotics often assume the robots' joint angles as internal state [16], [17]. Our work sets apart from these approaches in that we treat external objects within the framework of active inference. We adopted the tilt of the grasped object as the internal state and learned the generative model that predicts tactile sensations from this tilt.

III. PROBLEM STATEMENT

We tackle the peg-in-hole task with a 6 DoF UR5e robot arm equipped with Gelsight sensors in its fingers. The objective is to safely insert various pegs into their corresponding holes given the initial and the goal end-effector pose. During the whole insertion process, the robot

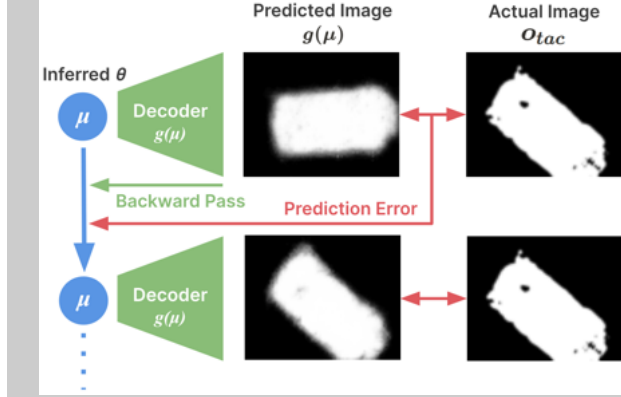


Fig. 2: The internal state is updated in real-time to minimize free energy, thereby reducing discrepancies between the predicted and the actual tactile images of the contact area estimation.

has access to its joint angles, force-torque readings and tactile readings from one of the tactile sensors. Besides the information from the vision-based tactile sensors, no other visual information was given to the robot. In the experiments we made the following assumptions:

- 1) The width of the peg to be grasped is small enough such that a “tilt” can be perceived in the tactile image.
- 2) The inferred tilt obtained from the iteration of free energy minimization is accurate enough to be used as an observation of the robot.
- 3) The target pose of the end effector is known.

IV. METHODOLOGY

A. Overview

To achieve robust and adaptive insertion, we propose a dual-policy structure, as depicted in Fig. 1. The system initially employs a learned force control policy to insert the grasped peg into the hole while avoiding excessive force [6]. When the grasped peg slips or tilts within the gripper after contacting the environment, the system activates active inference. In this phase, the agent infers the tilt of the peg and uses this inference to decide what pose to take.

B. Insertion Policy

The force-controlled insertion policy uses a PID-based parallel position-force control scheme with a selection matrix to determine the degree of control exerted on position and force along each direction [6]. The control command given to the robot \mathbf{x}_c is composed of a PD action on position, a PI action on force, a selection matrix S , and a position action from the RL policy \mathbf{a}_x :

$$\mathbf{x}_c = S(K_p^x \mathbf{x}_e + K_d^x \dot{\mathbf{x}}_e) + \mathbf{a}_x + (I - S)(K_p^f F_e + K_i^f \int F_e dt), \quad (6)$$

where F_e and \mathbf{x}_e represent the target force and position error, respectively. K_p^x , K_d^x , K_p^f , and K_i^f are the controller gain parameters. In this work, we exclude rotational action from the control command, as the alignment policy handles that aspect. All the PID and selection matrix parameters are learned by an RL policy in simulation and transferred to the real world (Sim2Real). For more details, see [6].

C. Alignment Policy

Fig. 2 outlines the agent’s perception of tactile observations based on its internal state.

1) **Active Inference Model:** We model perception as the inference of an unobservable internal state μ , which represents the tilt of the peg relative to the end effector. This internal state is modeled as a one-dimensional scalar. The robot’s observations, denoted as \mathbf{o} , comprise a preprocessed tactile image \mathbf{o}_{tac} and the relative angle θ between the peg and the hole. Since μ is belief of the tilt angle between the peg and the end effector, the angle θ may differ from μ in situations where the peg undergoes multiple tilts. Fig. 3 illustrates the definition of μ and θ . The generative model g maps this internal state to predicted tactile sensations as follows:

$$\mathbf{o}_{tac} = g(\mu) + \epsilon_o, \quad (7)$$

where ϵ_o is zero-mean Gaussian noise with variance Σ_o . We assume a Gaussian prior for μ with mean zero and variance σ_μ^2 . Using these assumptions, the free energy F can be expressed as:

$$F = -\ln p(\mathbf{o}_{tac}|\mu) - \ln p(\mu) - \ln p(\theta) + \text{Const}. \quad (8)$$

Here, $p(\theta)$ is also assumed Gaussian with zero mean and variance σ_θ^2 . Under the Laplace approximation, we assume a Gaussian form for the recognition density $q(\mu)$, defined by its mean μ and variance σ^2 .

Given that the agent’s action of aligning its pose does not affect the internal state μ , the optimal state that minimizes free energy can be found through gradient descent. The update rule is:

$$\begin{aligned} \frac{d}{dt}\mu &= -\frac{\partial F(\mu)}{\partial \mu} \\ &= \frac{\partial g(\mu)^T}{\partial \mu} \Sigma_{tac}^{-1}(\mathbf{o}_{tac} - g(\mu)) - \sigma_\mu^{-2}\mu. \end{aligned} \quad (9)$$

In this equation, the term $\mathbf{o}_{tac} - g(\mu)$ represents the prediction error, while Σ_{tac}^{-1} serves as a precision parameter. The expression $\partial g(\mu)^T / \partial \mu$ maps the tactile sensory space to the internal state space, which can be obtained through a backward pass of the network when the generative model is approximated by a deep neural network [16]. The second term acts as a regularization term that mitigates excessive updates. Upon updating the internal state to $\mu' = \mu + \Delta_t \dot{\mu}$, the agent performs inverse kinematics to correct the tilt by rotating the end effector around the Tool Center Point (TCP), which is defined at the center of the two fingers.

Updating μ according to Equation (9) constitutes the perceptual component of our model, as it aims to minimize the first and second terms of the free energy in Equation (8). Conversely, the action performed by the agent targets the minimization of the third term in the free energy equation, reflecting its goal to align the peg with the hole.

2) **Generative Model:** For the generative model g (i.e. generative function) to predict tactile sensation from the inferred tilt, we approximate it using a neural network architecture closely following the design proposed in [16]. The




Fig. 3: **Left:** While μ does not change after the action, θ becomes close to zero. **Right:** Depending on the surface of the peg, the estimated contact area can become noisy (top right figure), which poses challenges for accurate tilt estimation.




Fig. 4: Overview of the experimental setup. The bottom right shows the pegs used, arranged from left to right as shaft, pulley, cylinder, cuboid, and elliptical cylinder.

Algorithm 1: π_{align} by Deep Active Inference

Instant Decoder Training

$\sigma_{init} \leftarrow \text{ContactArea}(\text{tactile image in a straight pose})$
 Dataset $\leftarrow \text{DataAugmentation}(\sigma_{init})$
 Decoder.train(Dataset)

Active Inference Loop

```

 $\mu \leftarrow 0$ 
while True do
     $\sigma_{tac} \leftarrow \text{ContactArea}(\text{tactile image})$ 
     $\mathbf{g}(\mu) \leftarrow \text{Decoder.forward}(\mu)$ 
     $\partial \mathbf{g} \leftarrow \text{Decoder.backward}(\mu)$ 
     $\mathbf{e}_{tac} = \sigma_{tac} - \mathbf{g}(\mu)$  ▷ Prediction error
     $\mu = \text{Update}(\mu, \partial \mathbf{g}, \mathbf{e}_{tac})$  ▷ Perception
    if  $\mu > \text{threshold}$  then ▷ Action
        InverseKinematics( $\mu$ )
    end
end
    
```

54 pt
0.75 in
19.1 mm

network consists of two fully connected layers followed by alternating transposed convolution and standard convolution layers, with a final 1D-Dropout layer to mitigate overfitting. For a comprehensive understanding, readers are referred to [16]. Our modification to the architecture involves reducing the input dimension from 4 to 1 in order to decode the tactile image from a single numerical value, which is the tilt.

3) **Instant Decoder Training:** Our framework trains a decoder between the initial grasp and the insertion stage in just a few seconds, removing the need for a pre-trained model. Two key components enable this: contact area estimation and self-data augmentation. We adopt the neural network from [18] to estimate the peg’s contact area per frame, as only shape information is essential for tilt inference. This serves as the first component.

The second component is self-data augmentation. A notable limitation of using vision-based tactile sensors is their susceptibility to damage. To avoid damaging the sensor surface by collecting large datasets in a real-world setting, we developed a self-data augmentation method that creates a dataset from a single tactile image of the peg in a straight pose. Assuming that we have access to this tactile image of a peg in a straight pose (aligned with the hole).

To create a dataset for decoder training, we use computer vision techniques to rotate this estimated contact area of

the peg in a straight pose σ_{init} by a specified degree to generate new tactile data. As discussed in [10], raw tactile RGB images can be complex and contain task-irrelevant features (see Fig. 3). Our method overcomes this issue by performing active inference based on prediction error minimization in conjunction with contact area estimation. Extracting the contact area of the grasped peg enables the agent to quickly learn the decoder every time it grasps a new object.

The core advantage of this rapid training lies in its efficiency: our agent updates its internal state, denoted as μ , by minimizing prediction error. This means the decoder does not need to capture detailed shapes of the contact area for effective performance, as illustrated in Fig. 2. Consequently, this allows our system to swiftly adapt to new pegs without the need for a large, diverse pre-trained model.

Algorithm 1 summarizes the flow of the proposed alignment policy. During the force-controlled insertion process, a perceptual inference loop runs in real-time, updating μ following Equation (9) at regular intervals. The agent takes an action only when μ exceeds a specified threshold, indicating that the peg is tilted. The action is designed to reduce the relative angle between the peg and the hole, θ becomes zero by rotating the end effector around TCP, as calculated through inverse kinematics.

V. EXPERIMENTS AND RESULTS

A. Experimental Setup

We evaluate the proposed method on several peg-in-hole insertion tasks. We use a UR5e robot arm with a control frequency of 500 Hz. The robotic arm has a built-in force/torque sensor at its wrist and a Robotiq 2F-85 Adaptive Gripper. Two GelSight Mini sensors are mounted on the gripper, where only one of them is used during the experiments. We use the assembly task (motor, pulley, and shaft) from the Assembly Challenge of the World Robot Summit 2020 edition [19]. Additionally, we use other 3D-printed pegs such as a cuboid, a cylinder, and an elliptical

37 pt
0.514 in
13.1 mm

cylinder. We observed that working with the shaft resulted in significant slippage along the x -axis. Since translations along the x -axis are not the focus of our study, we opted for a 3D-printed cuboid peg instead of the shaft. This cuboid has an industrial-level clearance of 0.08 mm (measured with a caliper), compared to the shaft’s 0.05 mm. The motor pulley has a clearance of 0.3 mm. The average width of the pegs is 8mm. We used Gazebo simulator [20] version 9 for the simulation environment to train the insertion policy. To control the robot, the Robot Operating System (ROS) [21] with the Universal Robot ROS Driver is used.

B. Experimental Procedure

To validate our approach, we conduct two independent experiments. The first one is the perceptual inference experiment, where we focus on evaluating our alignment policy’s capability of inferring the tilt of multiple objects. The second one is the active inference experiment, where the whole dual-policy system is evaluated.

1) *Perceptual Inference Procedure*: For this experiment, we aim to validate our alignment policy’s ability to infer the tilts of multiple objects accurately. Assuming the availability of a tactile image from a straightly grasped peg, the agent’s goal is to infer peg tilts based on tactile images captured by the GelSight sensor. The initial tactile image is processed through a pre-trained neural network model to estimate the contact area [18] for all tactile images in the experiments. A training dataset (\mathcal{D}_{train}) of 500 tactile images, along with their corresponding tilts, is then created by rotating the initial contact area observation by a randomly sampled value from $[-20, 20]$ degrees. Test sets (\mathcal{D}_{test}) containing 100 samples are created using the tactile image in a straight pose at a different time step to avoid using the same exact shape among the train and the test data. The decoder model was trained over five epochs, taking approximately one second on average using an Intel Core i9-9900K CPU and an Nvidia RTX-2080 Ti GPU.

Upon completing the instant decoder training, perceptual inference is performed for each of the 100 test data \mathcal{D}_{test} . The initial value of μ is set to 0.0, with a maximum of 500 inference iterations. The update for μ is based on Equation (9) and involves two parameters: Σ_{tac}^{-1} and σ_{μ}^2 . We empirically set $\Sigma_{tac}^{-1} = 2 \times 10^4$ and $\sigma_{\mu}^2 = 1 \times 10^{-2}$.

2) *Dual-policy Procedure*: Before deploying the dual-policy architecture, the force-controlled insertion policy was trained in Gazebo simulation environment. During the training, the agent’s task was to insert a cuboid peg fixed to the end effector into a hole. The goal position, object-surface stiffness, goal pose uncertainty, and desired insertion force were randomized for each episode. The agent was trained for 100000 time steps, which takes about an hour to complete. The active inference experiment follows the same procedure as perceptual inference for decoder training, except that we do not create test data \mathcal{D}_{test} . The robot initially trains the decoder using a tactile image of the peg in a straight pose. The agent creates a dataset of 1000 tactile images whose tilts were drawn randomly from $[-10, 10]$. We narrowed the

TABLE I: Mean absolute error of estimated tilts [deg]

Object	Supervised	Proposed method
Shaft	4.5	1.0
Pulley	0.3	1.3
Cylinder	2.7	1.5
Cuboid	4.4	1.5
Elliptical Cylinder	1.0	1.2

range of degrees since we noticed that in real situations the tilts were mostly within this range. The robot then moves to an initial position, which is about 1 cm far in x axes from the hole, and $[-1, 1]$ cm in y and z axes. This initial position is changed after 10 insertions, a total of 40 insertions were done. At this initial position, the robot grasps the peg with a random tilt, assuming that a tilt has occurred due to the contact with the environment, or initial grasp pose uncertainty. While grasping the peg, the perceptual inference runs in real-time with 0.5 Hz, starting from $\mu = 0$ to update until it reaches the maximum of 1000 iterations. The updated μ is used as the current belief of the internal state. The agent first performs the active inference-based policy to align the tilted peg with the hole and transitions to the insertion policy. When θ becomes greater than a threshold, the agent switches to the alignment policy to make the peg straight. This threshold is empirically set to 0.7.

C. Comparison

To validate the adaptability and accuracy of our alignment policy on multiple objects, we also tested a baseline method that employs supervised learning with Convolutional Neural Networks (CNNs). The baseline model predicts the peg’s tilt from the tactile image of the estimated contact area, trained on the same dataset and epochs as the proposed method.

For the dual-policy experiments, we employed our insertion policy without tactile feedback [6] as a baseline for comparison. To make a fair comparison (using the same friction), we conducted the baseline experiments with the tactile sensor attached to the finger.

D. Results

1) *Perceptual Inference Result*: The results of the perceptual inference experiments are summarized in Table I. The maximum mean absolute error for estimated tilts using the supervised learning method is 4.5 degrees, while maximum mean absolute error for our proposed method is 1.5 degrees. As evidenced by the results for the shaft (4.5 deg) and the cylinder (2.7 deg), the supervised learning method struggles to extract meaningful features from small datasets, especially when the peg is round and has an indistinct border in the tactile image of the contact area. In contrast, our proposed method is robust against such noise, as the tilt inference is based on maximizing the pixel overlap between the predicted (i.e., inferred) contact area $g(\mu)$ and the contact area derived from the actual tactile sensation. This result indicates that our agent’s inference is generally applicable to multiple pegs with different geometry. It’s worth noting that the baseline supervised learning method was only feasible because the

TABLE II: Success rate of insertions

Object	Without tactile feedback	Proposed method
Cuboid	1/20 (5%)	36/40 (90%)
Pulley	14/20 (70%)	37/40 (93%)

internal state was defined as a single scalar. In contrast, our proposed method can be applied to internal states with higher dimensions simply by altering the decoder input dimension.

2) *Dual-policy Result:* The success rates of insertions for different objects are summarized in Table. II. For the cuboid peg with a clearance of 0.08 mm, the baseline method without tactile feedback achieved a success rate of just 5%, corresponding to 1 out of 20 trials. Given the stringent clearance, the baseline method generally failed to insert the peg with an initial random grasp pose, except in one trial where it managed to grasp the peg in almost a straight pose by chance. In contrast, our proposed method significantly improved the success rate to 90% with the aid of the tactile alignment policy, succeeding in 36 out of 40 trials. Similarly, for the pulley, the baseline method achieved a relatively better but still suboptimal success rate of 70%, or 14 out of 20 trials. Our proposed method again outperformed the baseline, with a success rate of 93% based on 37 successful insertions out of 40 trials. The relatively higher success rate of motor pulley insertions by baseline method compared with the cuboid insertions comes not only from the higher clearance (0.3 mm) but also from the shape and the material of pulley. As illustrated in bottom right of Fig. 4, the pulley has a hole in the center, whose position itself does not change due to slippage, or rotation. We also noticed that the motor pulley insertion has less friction than the 3D printed cuboid-hole, guiding the hole of the pulley to the motor. However, while this small friction enabled the baseline method to insert with 70% success rate, it greatly damaged the gel of the tactile sensors. Due to this significant damage on the tactile sensor observed during the experiments, we decided to run only 20 insertions for the baseline method.

VI. CONCLUSION

This paper presents a novel tactile-based active inference method for force-controlled peg-in-hole insertion. Our method achieved a 90% success rate in physical robot experiments, with a clearance of less than 0.1 mm by utilizing active inference on tactile data for peg alignment. Furthermore, we demonstrated that our approach generalizes to five different objects without the need for a pre-trained model or the collection of real-world data. This addresses a significant limitation associated with the fragility of vision-based tactile sensors. Future work will integrate vision to improve robustness against slippage in insertion and rotation. Additionally, we plan to extend the active inference model's internal states to tackle more complex tactile tasks.

REFERENCES

[1] H. Bruyninckx, S. Dutre, and J. De Schutter. Peg-on-hole: a model based solution to peg and hole alignment. In *Proceedings of 1995 IEEE International Conference on Robotics and Automation*, volume 2, pages 1919–1924 vol.2, 1995.

[2] M.E. Caine, T. Lozano-Perez, and W.P. Seering. Assembly strategies for chamferless parts. In *Proceedings, 1989 International Conference on Robotics and Automation*, pages 472–477 vol.1, 1989.

[3] Íñigo Elguea-Aguinaco, Antonio Serrano-Muñoz, Dimitrios Chrysostomou, Ibai Inziarte-Hidalgo, Simon Bøgh, and Nestor Arana-Arexolaleiba. A review on reinforcement learning for contact-rich robotic manipulation tasks. *Robotics and Computer-Integrated Manufacturing*, 81:102517, 2023.

[4] Cristian Camilo Beltran-Hernandez, Damien Petit, Ixchel Georgina Ramirez-Alpizar, Takayuki Nishi, Shinichi Kikuchi, Takamitsu Matsubara, and Kensuke Harada. Learning force control for contact-rich manipulation tasks with rigid position-controlled robots. *IEEE Robotics and Automation Letters*, 5(4):5709–5716, oct 2020.

[5] Wenzhen Yuan, Siyuan Dong, and Edward H. Adelson. Gelsight: High-resolution robot tactile sensors for estimating geometry and force. *Sensors (Basel, Switzerland)*, 2017.

[6] Cristian C. Beltran-Hernandez, Damien Petit, Ixchel G. Ramirez-Alpizar, and Kensuke Harada. Variable compliance control for robotic peg-in-hole assembly: A deep-reinforcement-learning approach. *Applied Sciences*, 10(19), 2020.

[7] Maria Bauzá, Oleguer Canal, and Alberto Rodriguez. Tactile mapping and localization from high-resolution tactile imprints. *2019 International Conference on Robotics and Automation (ICRA)*, pages 3811–3817, 2019.

[8] Maria Bauzá, Eric Valls, Bryan Lim, Theo Sechopoulos, and Alberto Rodriguez. Tactile object pose estimation from the first touch with geometric contact rendering. In *Conference on Robot Learning*, 2020.

[9] Letian Fu, Huang Huang, Lars Berscheid, Hui Li, Ken Goldberg, and Sachin Chitta. Safe self-supervised learning in real of visuo-tactile feedback policies for industrial insertion. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10380–10386, 2023.

[10] Siyuan Dong, Devesh Jha, Diego Romeres, Sangwoon Kim, Daniel Nikovski, and Alberto Rodriguez. Tactile-rl for insertion: Generalization to objects of unknown geometry. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021.

[11] Karl Friston. Friston, k.j.: The free-energy principle: a unified brain theory? *nat. rev. neurosci.* 11, 127–138. *Nature reviews. Neuroscience*, 11:127–38, 02 2010.

[12] Pablo Lanillos, Cristian Meo, Corrado Pezzato, Ajith Anil Meera, Mohamed Baioumy, Wataru Ohata, Alexander Tschantz, Beren Millidge, Martijn Wisse, Christopher L. Buckley, and Jun Tani. Active inference in robotics and artificial agents: Survey and challenges, 2021.

[13] Cansu Sancaktar, Marcel A. J. van Gerven, and Pablo Lanillos. End-to-end pixel-based deep active inference for body perception and action. In *2020 Joint IEEE 10th International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, 2020.

[14] Kai Ueltzhöffer. Deep active inference. *Biological Cybernetics*, 112(6):547–573, oct 2018.

[15] Beren Millidge. Deep active inference as variational policy gradients. *Journal of Mathematical Psychology*, 96:102348, 2020.

[16] Cansu Sancaktar, Marcel A. J. van Gerven, and Pablo Lanillos. End-to-end pixel-based deep active inference for body perception and action. In *2020 Joint IEEE 10th International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, 2020.

[17] Corrado Pezzato, Riccardo Ferrari, and Carlos Hernández Corbato. A novel adaptive controller for robot manipulators based on active inference. *IEEE Robotics and Automation Letters*, PP:1–1, 02 2020.

[18] Niklas Funk, Paul-Otto Müller, Boris Belousov, Anton Savchenko, Rolf Findeisen, and Jan Peters. Canfnet: High-resolution pixelwise contact area and normal force estimation for visuotactile sensors using neural networks. In *2023 IEEE ICRA full day workshop: Embracing Contacts*, 2023. <https://sites.google.com/view/canfnet>.

[19] World Robot Summit. Assembly Challenge. <https://wrs.nedo.go.jp/en/wrs2020/challenge/industrial/assembly.html>, 2020. Accessed: 2025-11-7.

[20] N. Koenig and A. Howard. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566)*, volume 3, pages 2149–2154 vol.3, 2004.

[21] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Ng. Ros: an open-source robot operating system. In *ICRA Workshop on Open Source Software*, volume 3, 01 2009.