

Uncertainty-Aware Finger Force Estimation for Reliable Box Manipulation

Anonymous Authors

Abstract—Reliable box manipulation requires successful placement without damage, which depends on maintaining sufficient contact force without excessive loading. We hypothesize that per-finger force feedback improves manipulation reliability relative to open-loop finger commands. To this end, we present **Finger External-Force Estimation from Proprioceptive Signals (FEEL)**, an uncertainty-aware residual method that estimates per-finger external force magnitude from proprioceptive histories. Our method combines a calibrated virtual-spring estimate with a residual distribution predicted by a Mixture Density Network (MDN) and requires no dedicated contact sensors during deployment. The estimated force is used for continuous per-finger adaptation, while predictive uncertainty is evaluated as an empirical signal of ambiguity in proprioceptive force estimation. Across ten fingers and contact forces spanning 0–20 N, our method achieves a mean RMSE of 0.251 N, reducing the error by 92.8% relative to the virtual spring and by 9.4% relative to an end-to-end MDN. In box pick-and-place, force adaptation maintains task success at 10/10 and increases both damage-free outcomes and reliability from 7/10 to 10/10. A video demonstrating the proposed method is available at <https://anonymous-webpage.pages.dev/>.

I. INTRODUCTION

Box pick-and-place is a common operation in logistics and manufacturing [1], [2]. Automated systems often repeat this operation across many packages. Boxes with similar dimensions can differ in mass, contents, and structural strength. These properties may not be known before manipulation. The same grasping motion can therefore produce different outcomes across packages. This variability makes consistent box pick-and-place success difficult to achieve.

The contact force applied during manipulation directly affects these outcomes. Insufficient force can cause grasp or placement failure. Excessive force can deform the box or damage its contents. Task completion alone therefore does not characterize handling quality. Handling damage can cause product loss, rework, and additional operating costs [3]. A trial is reliable only when the task succeeds without box damage.

Force feedback can help maintain a stable grasp while limiting excessive interaction force [4], [5]. Dedicated force or tactile sensors mounted on a robot hand provide direct measurements at the contact surface [6]. However, installing these sensors on multiple fingers increases hardware cost and requires calibration at each sensing location [7], [8]. Sensor coverage also depends on placement, which can leave some contact forces undetected [9]. These requirements motivate force estimation from robot proprioception [7], [10]–[12], avoiding dedicated contact sensors during deployment. However, unobserved contact locations can make the relationship between joint responses and interaction forces difficult to

capture with a fixed analytical model. This motivates learning a correction to the nominal force model.

We hypothesize that per-finger force feedback improves manipulation reliability relative to open-loop finger commands. To this end, we present **Finger External-Force Estimation from Proprioceptive Signals (FEEL)**, an uncertainty-aware method that estimates per-finger external force magnitude by correcting a calibrated virtual-spring estimate with a residual distribution learned from proprioceptive histories. The corrected force estimate provides continuous per-finger feedback without dedicated contact sensors during deployment. The residual distribution also provides predictive uncertainty, which we examine as an empirical signal of ambiguity when force variations are weakly reflected in proprioceptive histories. We evaluate force-estimation accuracy, the contribution of force adaptation to reliable box pick-and-place, and whether predictive uncertainty reflects this ambiguity.

The main contributions of this paper are twofold: 1) an uncertainty-aware residual method that estimates per-finger external force magnitude from proprioception for reliable box manipulation; and 2) a real-robot evaluation of its force-estimation accuracy, predictive uncertainty, and use for per-finger force adaptation in box pick-and-place.

II. RELATED WORK

Damage during robotic handling is a practical concern in warehouse logistics. Package damage can result in product loss and higher operating costs [3]. Prior work addresses manipulation damage before and during task execution. Damage-prediction methods identify risky actions before execution [3]. Force-sensitive grasping and gentle manipulation regulate physical interaction during execution to limit object damage [4], [5]. Package-handling studies also examine box deformation and loss of grasp during physical transfer [1], [2]. Force feedback is therefore relevant to completing the task while limiting box damage.

Force feedback can be obtained through direct sensing or proprioceptive estimation. Dedicated force and tactile sensors provide direct measurements at the contact surface [6]. However, equipping multiple fingers with dedicated sensors increases hardware cost and requires calibration at each sensing location [7], [8]. Model-based methods estimate interaction without external contact sensors using robot dynamics, actuator measurements, or observed motion [10], [11], [13]. A common model represents compliant interaction as a virtual spring that relates elastic displacement to force [14]. The spring stiffness can be identified from measured interaction

data, including through data-driven regression [15]. A calibrated virtual spring provides the nominal force magnitude estimate in our residual formulation.

Data-driven force estimation can use direct regression or residual learning. Direct methods estimate force or torque from synchronized robot measurements and reference labels [7], [12], [16]. Hybrid methods retain model-based predictions by learning their residual errors [17], [18]. Residual learning retains the nominal estimate while accounting for effects that are difficult to model. Our formulation adds a proprioceptive residual to each finger's calibrated virtual-spring estimate.

The residual mapping can remain ambiguous when proprioceptive measurements do not fully describe the physical interaction. A Mixture Density Network (MDN) represents the conditional residual distribution as a Gaussian mixture [19]. The conditional distribution admits a closed-form decomposition of predictive variance into within-component and between-component terms [20]. We refer to these terms collectively as predictive uncertainty and analyze whether they reflect ambiguity in proprioceptive force estimation.

III. PROBLEM FORMULATION

At deployment, the robot observes a history of proprioceptive signals. Our proposed method uses this history to estimate per-finger external force magnitudes without dedicated force sensors. The force estimates support continuous finger regulation, while predictive uncertainty is analyzed as an empirical descriptor of ambiguity in the residual prediction.

A. Robot Configuration and Modeling Assumptions

We consider a bimanual robot equipped with two position-controlled, five-fingered hands. Each finger has four actuated joints. The four non-thumb fingers have one abduction–adduction joint and three flexion joints, while the thumb uses a separate four-joint chain for opposition and flexion. For finger $i \in \{1, \dots, N\}$ at time t , the measured joint position is $\mathbf{q}_{i,t} \in \mathbb{R}^4$, and the commanded joint position is $\mathbf{q}_{i,t}^{\text{cmd}} \in \mathbb{R}^4$. Measured and commanded joint positions are updated at 100 Hz, matching the hand-control loop. The main experiments use the proportional–derivative (PD) gain selected from three settings in Section V-E.

We adopt a virtual-spring model to obtain an initial force magnitude estimate for each finger. Computing the spring displacement requires a defined point on the finger. Exact contact location is difficult to determine from proprioceptive measurements alone. We therefore assign one nominal contact point to the center of each fingertip surface. For finger i , its current position $\mathbf{p}_i(\mathbf{q}_{i,t})$ is obtained from the measured joint position through forward kinematics. We approximate the finger–object interaction as quasi-static when constructing the virtual-spring estimate. The scope of this work is to estimate the corresponding external-force magnitude for each finger.

B. Residual Force Magnitude Formulation

Under these assumptions, the displacement vector from the measured contact-point position to its commanded position

provides a kinematic proxy for the loading direction. Its norm defines the spring displacement. For finger i , let $\delta_{i,t} = \|\mathbf{p}_i(\mathbf{q}_{i,t}^{\text{cmd}}) - \mathbf{p}_i(\mathbf{q}_{i,t})\|_2$ denote the spring displacement. The virtual-spring force magnitude estimate is

$$f_{i,t}^{\text{spr.}} = k_i \delta_{i,t}, \quad (1)$$

where $k_i \geq 0$ is the calibrated virtual stiffness for finger i . The superscript *spr.* denotes the virtual-spring estimate. Section IV-B describes the stiffness calibration using a separate dataset. The estimate represents an external-force magnitude in newtons under the assumed contact geometry.

The goal is to learn the residual between the reference force magnitude and the virtual-spring estimate from proprioceptive history. Let $f_{i,t}^{\text{ref.}}$ denote the reference force magnitude used during training and evaluation. The residual target is $r_{i,t} = f_{i,t}^{\text{ref.}} - f_{i,t}^{\text{spr.}}$. We model the relation between the proprioceptive history $\mathbf{x}_{i,t}$ and the force magnitude residual $r_{i,t}$ as the conditional distribution $p_\theta(r_{i,t} | \mathbf{x}_{i,t})$. At deployment, $r_\theta(\mathbf{x}_{i,t})$ denotes the residual estimate obtained from this distribution. The estimated external-force magnitude is

$$\hat{f}_{i,t} = f_{i,t}^{\text{spr.}} + r_\theta(\mathbf{x}_{i,t}). \quad (2)$$

In addition to the force estimate, we seek a predictive uncertainty signal that characterizes ambiguity in the residual prediction. The force estimate is used for continuous per-finger regulation, while predictive uncertainty is analyzed as an empirical signal of force variations that are weakly reflected in the proprioceptive history. Residual inference and predictive-uncertainty computation are described in Section IV-C.

C. Reliable Box Manipulation

We consider box pick-and-place as a representative handling operation in logistics. The robot lifts a box placed in front of it and transfers the box to a designated shelf. A trial is successful when the planned transfer and placement are completed. A completed placement may still damage the box, whereas damage avoidance alone is insufficient if placement fails. We therefore evaluate both outcomes.

We evaluate manipulation performance using trial-level reliability. Let S_n and B_n denote task success and a damage-free box outcome in trial n , respectively. The trial-level reliability indicator is defined as

$$\rho_n = \mathbb{I}[S_n \wedge B_n]. \quad (3)$$

Thus, $\rho_n = 1$ only when placement succeeds without box damage. The task evaluation reports the proportion of reliable trials together with the separate success and damage rates. Section V-D specifies the seam-damage proxy.

IV. PROPOSED METHOD

This section presents the proposed external-force magnitude estimation module and its use for box manipulation. Figure 1 summarizes the pipeline from real force-data collection to residual force estimation, force-based finger adaptation,




Fig. 1: Overview of FEEL. (A) A wearable F/T sensor collects reference force magnitudes synchronized with proprioceptive histories during finger motion. (B) The proprioceptive history $\mathbf{x}_{i,t}$ serves as input to FEEL. The calibrated virtual spring provides the model-based force estimate $f_{i,t}^{\text{SPR}}$, while the MDN predicts the residual $r_{\theta}(\mathbf{x}_{i,t})$ and its predictive uncertainty. (C) The two force terms are added to obtain the per-finger estimate, $\hat{f}_{i,t} = f_{i,t}^{\text{SPR}} + r_{\theta}(\mathbf{x}_{i,t})$. During deployment, this estimate provides feedback for force-based finger adaptation without dedicated contact sensors.

and predictive uncertainty analysis. Sections IV-A and IV-B describe data collection and virtual-spring calibration. Section IV-C presents residual force magnitude estimation and predictive uncertainty. Section IV-D describes the use of the estimated force for per-finger adaptation.

A. Real Force-Data Collection

We train all reported force magnitude estimators on real-robot measurements based on the simulation–real comparison reported in Section V-B. These measurements are collected using a custom wearable device built around an AFT20-D15 force/torque (F/T) sensor (AIDIN Robotics). Figure 2 shows the sensor attached to the operator’s fingertip using a custom-designed fingertip mount. An ESP32 microcontroller and an SN65HVD230 CAN transceiver are mounted on the wrist. The microcontroller communicates with the F/T sensor over CAN through the transceiver. This arrangement allows the sensing surface to be positioned at selected contact locations on the robot finger. The three-axis force vector $\mathbf{f}_{i,t}^{\text{f.t.}} \in \mathbb{R}^3$ and the torque channels are transmitted to the host computer through a serial connection and recorded at 90 Hz. The measured force magnitude $f_{i,t}^{\text{ref.}} = \|\mathbf{f}_{i,t}^{\text{f.t.}}\|_2$ serves as the reference force measurement during training and evaluation. The estimator predicts this force magnitude. The F/T force direction and torque are not used for supervision.

During each collection sequence, the robot replays the same predefined motion for the selected finger. The operator applies the instrumented surface at multiple locations on the finger while varying the contact force. Contact location is varied during collection but is not provided to the estimator. The collected force magnitude spans 0–20 N, and samples before, during, and after contact are retained. The estimator dataset contains approximately 25,000 synchronized samples per finger. The F/T sensor supplies reference force measurements only for dataset construction and offline force magnitude estimator evaluation and is not used during box manipulation.

Measured and commanded joint positions are recorded at 100 Hz, while the F/T channels are recorded at 90 Hz. The



Fig. 2: Real force-data collection using a wearable F/T sensor. Contact interactions at multiple finger locations are synchronized with proprioceptive measurements.

clocks of the two computers are synchronized before data collection, and the two streams are aligned by timestamp on the host computer. For each F/T measurement, the most recent $H = 10$ robot-state samples up to the aligned timestamp form a 0.1-s proprioceptive history. The corresponding reference force magnitude provides $f_{i,t}^{\text{ref.}}$ and is paired with $f_{i,t}^{\text{SPR}}$ to compute the residual label.

B. Virtual-Spring Calibration

We identify the stiffness in Eq. (1) from a separate calibration dataset using least-squares fitting [14]. This dataset contains approximately 5,000 synchronized samples per finger and does not overlap with the estimator dataset. For each finger, the nonnegative stiffness k_i is identified over its calibration samples:

$$k_i^* = \arg \min_{k_i \geq 0} \sum_j (f_{i,j}^{\text{ref.}} - k_i \delta_{i,j})^2. \quad (4)$$

The identified stiffness is fixed before residual-model training. The nominal fingertip point defined in Section III-A is used to compute the spring displacement δ . Together with the identified stiffness, it defines the virtual-spring estimate used to construct the residual labels.

C. Residual Force Magnitude Estimation

The virtual spring provides a physically grounded first-order estimate by relating contact-induced tracking displacement to external load [14]. Its fixed stiffness and nominal contact point, however, cannot fully capture variations in contact geometry and finger compliance. We therefore retain the virtual-spring estimate and learn only its residual error [17]. This formulation anchors the prediction to the nominal physical model while using data-driven learning to account for unmodeled effects.

External contact can prevent a position-controlled finger from following its commanded motion. The resulting deviation provides a signal related to the external load, while the measured joint position provides the configuration needed to interpret this signal [7], [10]. For finger i , we define the joint tracking error as $\mathbf{e}_{i,t} = \mathbf{q}_{i,t}^{\text{cmd}} - \mathbf{q}_{i,t}$. A short history captures how the joint position and tracking error evolve during contact:

$$\mathbf{x}_{i,t} = [\mathbf{q}_{i,t-H+1:t}, \mathbf{e}_{i,t-H+1:t}]. \quad (5)$$

We use $H = 10$, corresponding to a 0.1-s history at 100 Hz, without external force or tactile measurements at deployment. Variations in unobserved contact conditions can produce different force residuals from similar proprioceptive histories. To represent this conditional variability, we model $p_\theta(r | \mathbf{x})$ using an MDN rather than assuming a single deterministic residual [19], [20].

For K Gaussian components, the residual distribution is

$$p_\theta(r | \mathbf{x}) = \sum_{k=1}^K \pi_k(\mathbf{x}) \mathcal{N}(r; \mu_k(\mathbf{x}), \sigma_k^2(\mathbf{x})), \quad (6)$$

where π_k , μ_k , and σ_k denote the weight, mean, and standard deviation of component k , respectively. Each component represents a possible range of the residual force magnitude, and its weight describes the relative likelihood of that range. The network parameters are learned by minimizing the negative log likelihood of the residual labels:

$$\mathcal{L}_{\text{res}}(\theta) = - \sum_{(\mathbf{x}, r) \in \mathcal{D}_{\text{train}}} \log p_\theta(r | \mathbf{x}). \quad (7)$$

This objective encourages the predicted distribution to assign high likelihood to the residual observed for each proprioceptive history. At inference, the mean of the highest-weight component provides the residual estimate:

$$r_\theta(\mathbf{x}) = \mu_{\arg \max_k \pi_k(\mathbf{x})}(\mathbf{x}). \quad (8)$$

The force magnitude estimate follows Eq. (2) and is clipped at zero to maintain a nonnegative value.

An additional property of the residual MDN is that its predictive variance has a closed-form decomposition by the law of total variance [20]. Let $\bar{r}(\mathbf{x}) = \sum_k \pi_k(\mathbf{x}) \mu_k(\mathbf{x})$ denote the mixture mean. The predictive variance of the residual decomposes as

$$\begin{aligned} \mathbb{V}_\theta(r | \mathbf{x}) &= \mathbb{E}_{k \sim \pi}[\mathbb{V}(r | \mathbf{x}, k)] + \mathbb{V}_{k \sim \pi}(\mathbb{E}[r | \mathbf{x}, k]) \\ &= \underbrace{\sum_{k=1}^K \pi_k(\mathbf{x}) \sigma_k^2(\mathbf{x})}_{u_\theta^{\text{alea.}}(\mathbf{x})} + \underbrace{\sum_{k=1}^K \pi_k(\mathbf{x}) (\mu_k(\mathbf{x}) - \bar{r}(\mathbf{x}))^2}_{u_\theta^{\text{epis.}}(\mathbf{x})}. \end{aligned} \quad (9)$$

Following [20], the first term, $u_\theta^{\text{alea.}}$, is the within-component variance and is interpreted as aleatoric uncertainty. The second term, $u_\theta^{\text{epis.}}$, is the between-component variance and is interpreted as epistemic uncertainty. Together, these terms constitute the predictive variance of the residual MDN and are obtained from the same forward pass without sampling. We refer to them collectively as predictive uncertainty in the remainder of the paper and evaluate them as empirical signals rather than calibrated measures of model uncertainty. Section V-F examines whether predictive uncertainty reflects ambiguity in proprioceptive force estimation.

D. Force-Based Finger Adaptation

We assume that the object contacts the inner surfaces of the fingers during grasping. The nominal finger motion closes each finger from the measured initial configuration toward a fixed target hand configuration [22], [23]. We represent this movement by a normalized motion progress from zero to one. Our objective is to adapt each finger command to keep its estimated force magnitude near a prescribed level for stable grasping without excessive contact force.

Let $\Delta \mathbf{q}_{i,t}$ denote the joint-position increment toward the target configuration, and let $f_i^{\text{th.}}$ denote the force magnitude threshold. A ± 1 -N deadband around this threshold reduces command changes caused by small estimation fluctuations. The coefficient $\alpha_{i,t}$ takes values of 1, 0, and -1 to close, hold, and open the finger, respectively. Here, $\mathcal{Q}_i$ denotes the admissible joint-position range, and $\Pi_{\mathcal{Q}_i}$ projects the command into this range. The command is updated as

$$\mathbf{q}_{i,t+1}^{\text{cmd}} = \Pi_{\mathcal{Q}_i}(\mathbf{q}_{i,t}^{\text{cmd}} + \alpha_{i,t} \Delta \mathbf{q}_{i,t}). \quad (10)$$

Each finger closes below the deadband, holds within it, and moves in the opening direction above it. This force-based regulation adapts each finger command independently according to its estimated force magnitude. Fingers that reach contact earlier can hold or open while the others continue closing.

V. EXPERIMENTS

This section evaluates FEEL from force estimation to box manipulation. Section V-A describes the common experimental protocol. Section V-B examines the simulation–real discrepancy, and Section V-C evaluates residual force estimation. Section V-D isolates the effect of force adaptation, Section V-E examines controller-gain sensitivity with force adaptation, and Section V-F evaluates predictive uncertainty when external loading is weakly reflected in proprioception. Section V-G discusses the scope and limitations.




Fig. 3: Comparison of force magnitude traces from matched left-index-finger tapping motions in MuJoCo and on the real robot.

A. Experimental Setup

Experiments use the bimanual robot and five-fingered hands described in Section III-A. Force data are collected with the wearable F/T tool in Section IV-A. The main estimator dataset contains approximately 25,000 synchronized samples per finger over 0–20 N at the selected proportional-derivative (PD) gain setting $(K_p, K_d) = (200, 400)$. A separate virtual-spring calibration set contains approximately 5,000 synchronized samples per finger. Robot states are sampled at 100 Hz, and the F/T measurements are sampled at 90 Hz. The two streams are aligned by timestamp on the host computer. The three-axis force magnitude provides the reference value.

We train one estimator per finger from a 0.1-s proprioceptive history of $H = 10$ robot observations. Each observation contains four measured joint positions and four tracking errors, giving an 80-dimensional input. All learned estimators use two 64-unit fully connected hidden layers with tanh activations. Each MDN has $K = 10$ Gaussian components and 11,294 trainable parameters. Training uses AdamW for 1,000 epochs with a batch size of 512.

Evaluation uses row-level five-fold cross-validation stratified by force magnitude bins. Each fold contains 20,000 training samples and 5,000 test samples per finger, with identical splits across the learned methods. Input standardization is computed only from the training samples in each fold. We evaluate the model on the corresponding test samples and report root-mean-square error (RMSE) in newtons as the mean $\pm$ sample standard deviation across the five folds. For the manipulation experiments, the residual MDN is retrained on the complete estimator dataset using the same architecture and hyperparameter settings. Force magnitude inference runs at 100 Hz, matching the hand-control loop.

For the predictive-uncertainty evaluation, we additionally collect 3,000 synchronized samples per finger under a loading condition designed to produce substantial force variation with limited changes in proprioceptive signals. A vertical load is applied at the distal end of each finger while the commanded finger configuration is held constant. The same

F/T measurement setup provides the corresponding reference force magnitudes.

The manipulation experiments use a cardboard box measuring $395 \times 275 \times 240$ mm (width $\times$ depth $\times$ height), with a mass of 2.0 kg, and a designated shelf as the placement target. A trial is considered successful when the box remains on the shelf at the end of the predefined pick-and-place motion. Tape separation from either side of the designated seam is recorded as box damage. Each box is prepared with 4.5 cm of low-adhesion packaging tape adhered to either side of the seam, and a new box is used for every trial. Reliability follows Eq. (3) and counts trials that complete the task without recorded damage. Within each comparison, the initial box pose, box specification, and controller parameters other than the evaluated component remain fixed. External force and tactile sensors are not used during manipulation. We use a 6-N threshold for all fingers, selected in preliminary grasp trials and fixed for task evaluation.

B. Simulation–Real Force Magnitude Comparison

This experiment compares simulated and measured force responses to the same finger motion. The left index finger executes the same three-tap command sequence in MuJoCo and on the physical robot [21]. The physical experiment uses an F/T-instrumented flat surface, while the simulation uses its nominal contact geometry. The commanded trajectory, initial configuration, PD gains, and effort limit are matched. Figure 3 compares the resulting force magnitude traces.

The simulated and measured traces differ in profile and yield an RMSE of 5.287 N. Simulator contact forces depend on model assumptions and numerical approximations that can depart from physical behavior [24]. Unmodeled hardware compliance and actuator response may also contribute to this discrepancy. Although limited to one finger and one three-tap motion, the comparison motivates training the physical-robot estimators with real F/T measurements.

C. Force Magnitude Estimation

This experiment evaluates whether learning the residual of a virtual-spring estimate improves per-finger force magnitude estimation. We compare three estimators on the same cross-validation folds of the real contact data. The virtual spring [14] uses the calibrated stiffness and the displacement at the assumed fingertip contact point. The end-to-end MDN predicts the reference force magnitude directly from the proprioceptive history, following prior MDN-based force prediction [25]. The proposed residual MDN instead predicts the error of the virtual-spring estimate. The two MDNs use the same input history and network architecture, allowing the comparison to focus on the residual formulation.

Table I shows that the residual MDN achieves the lowest mean RMSE for both hands. Relative to the virtual spring, it reduces RMSE by 92.7% for the left hand and 93.0% for the right hand. The corresponding reductions over the end-to-end MDN are 7.2% and 11.7%. The virtual spring relies on one assumed fingertip contact point, although the dataset includes contacts at multiple finger locations. Learning its

TABLE I: Comparison of per-finger external force magnitude estimation accuracy using row-level five-fold cross-validation stratified by force magnitude bins.

Hand	Finger	Virtual spring [14]	End-to-end MDN [25]	Residual MDN (Ours)
		RMSE (N) ↓		
Left	Thumb	5.327 ± 0.012	0.790 ± 0.029	0.723 ± 0.024
	Index	3.203 ± 0.042	0.192 ± 0.030	0.178 ± 0.044
	Middle	2.545 ± 0.026	0.093 ± 0.002	0.096 ± 0.004
	Ring	3.467 ± 0.030	0.140 ± 0.059	0.119 ± 0.021
	Little	2.135 ± 0.033	0.103 ± 0.004	0.106 ± 0.006
	<i>Mean</i>	3.335 ± 0.011	0.264 ± 0.016	0.245 ± 0.008
Right	Thumb	5.361 ± 0.034	0.718 ± 0.052	0.669 ± 0.056
	Index	3.092 ± 0.033	0.207 ± 0.014	0.161 ± 0.011
	Middle	3.472 ± 0.047	0.132 ± 0.019	0.147 ± 0.027
	Ring	3.075 ± 0.035	0.184 ± 0.087	0.123 ± 0.008
	Little	3.254 ± 0.040	0.214 ± 0.018	0.186 ± 0.025
	<i>Mean</i>	3.651 ± 0.014	0.291 ± 0.016	0.257 ± 0.016
Overall	<i>Mean</i>	3.493 ± 0.007	0.277 ± 0.014	0.251 ± 0.010

TABLE II: Ablation of residual output models for external force magnitude estimation.

Hand	Residual MLP	Residual MDN (Ours)
	RMSE (N) ↓	
Left	0.261 ± 0.004	0.245 ± 0.008
Right	0.280 ± 0.004	0.257 ± 0.016
Overall	0.270 ± 0.004	0.251 ± 0.010

residual partly corrects this discrepancy and yields the lowest mean error among the evaluated methods.

We further evaluate whether representing the residual as a conditional distribution improves upon predicting a single residual value. The residual MLP differs only in replacing the MDN output with a scalar prediction. Table II reports each hand’s five-finger average and the overall ten-finger average. The residual MDN reduces the overall RMSE from 0.270 N to 0.251 N. This reduction may reflect the conditional mixture’s ability to represent alternative residuals when different contact locations produce similar proprioceptive histories.

D. Reliable Box Pick-and-Place

This experiment evaluates whether per-finger force adaptation improves reliability in box pick-and-place. Both conditions use the same arm motion, target hand configuration, and placement target. The baseline uses fixed finger commands, whereas the comparison uses residual MDN estimates for the adaptation in Section IV-D. The comparison therefore isolates the effect of force adaptation. Each condition is evaluated over 10 trials using the success, damage, and reliability criteria defined in Section V-A.

Table III shows that both conditions complete all 10 trials, while damage-free outcomes increase from 7/10 without force adaptation to 10/10 with adaptation. Reliability consequently increases from 7/10 to 10/10. The main observed effect is a reduction in box damage rather than an increase

TABLE III: Comparison of force adaptation under one predefined box pick-and-place motion over 10 trials per condition.

Method	Success ↑	Damage-free ↑	Reliability ↑
No force adaptation	10/10	7/10	7/10
Force adaptation (Ours)	10/10	10/10	10/10

in task completion. With fixed finger commands, fingers continue closing after contact, which may increase local loading on the box. Force adaptation instead allows each finger to hold or open after reaching the force threshold while the remaining fingers continue closing. Under the evaluated box and motion, this behavior improves reliability by reducing damage while maintaining task success.

E. Controller-Gain Sensitivity

The finger PD gain influences both the proprioceptive signal used for force estimation and the physical response of the force-adaptive controller. With $(K_p, K_d) = (100, 200)$, external contact is expected to produce larger tracking deviations, which may make force magnitude easier to estimate, but weaker command tracking may limit grasp execution. With $(K_p, K_d) = (300, 600)$, stronger command tracking is expected to reduce those deviations and can increase local loading during contact even as the finger commands are adapted. We therefore hypothesize a trade-off: the setting that minimizes force-estimation error does not necessarily maximize manipulation reliability under force adaptation.

We evaluate three PD-gain pairs. The tested (K_p, K_d) values are $(100, 200)$, $(200, 400)$, and $(300, 600)$. For each setting, the estimator is trained and evaluated on 25,000 force samples per finger collected at the corresponding gain. During manipulation, each gain setting uses its corresponding residual MDN for the force-based finger adaptation in Section IV-D. All settings use the same force threshold, adaptation rule, box, and predefined motion for 10 trials each.




Fig. 4: Snapshots of real-robot box pick-and-place with and without force-based finger adaptation. The enlarged views show box deformation and damage at the designated seam without adaptation, whereas the seam remains intact with adaptation.

TABLE IV: Controller-gain sensitivity of force magnitude estimation and force-adaptive box manipulation.

(K_p, K_d)	RMSE (N) ↓	Box manipulation (/10)		
		Success ↑	Damage-free ↑	Reliability ↑
(100, 200)	0.233 ± 0.005	4/10	10/10	4/10
(200, 400)	0.251 ± 0.010	10/10	10/10	10/10
(300, 600)	0.255 ± 0.008	9/10	8/10	7/10

Table IV summarizes the estimation and manipulation results across the three gain settings. The (100, 200) setting achieves the lowest RMSE but produces the lowest task-success rate. This result suggests that larger contact-induced tracking errors provide a clearer signal for force estimation, while weaker position tracking limits task execution. The (300, 600) setting achieves task success in 9/10 trials and damage-free outcomes in 8/10 trials. Interestingly, we observed that the robot could not form a stable grasp because the high gains made the fingers too stiff. The (200, 400) setting balances these effects and achieves the highest reliability of 10/10. This comparison shows that the controller gain should be selected jointly for proprioceptive force estimation and force-adaptive manipulation.

F. Predictive Uncertainty Under Proprioceptive Ambiguity

Proprioception enables force feedback without dedicated contact sensors, but external loading is reflected only through changes in joint position and tracking error. Consider an interaction in which the finger structure bears a vertical load while the commanded finger configuration remains fixed. The applied force can vary substantially with little change in the proprioceptive history. Our proposed method (FEEL) uses an MDN to represent this ambiguity as distinct force-residual modes conditioned on a given history. This distribution can express alternative residual outcomes, but it cannot resolve

TABLE V: Comparison of predictive uncertainty for Force through joints and Force through structure.

Interaction condition	Aleatoric uncertainty	Epistemic uncertainty
Force through joints	0.370 ± 0.001	0.102 ± 0.005
Force through structure	0.364 ± 0.005	0.172 ± 0.088

interaction information absent from the proprioceptive input. Recognizing this unresolved ambiguity is important when interpreting the resulting force estimate. We therefore hypothesize that mixture-component disagreement can provide an empirical signal of such cases.

To test this hypothesis, we compare Force through joints and Force through structure. Force through joints uses held-out test samples collected with the same predefined finger-motion protocol as the training data. Force through structure uses 3,000 separately collected samples per finger across all ten fingers. These samples are obtained by applying a vertical load at the distal end of each finger while its commanded configuration remains fixed. The samples have a mean reference force magnitude of 10.909 ± 6.030 N across both hands and are not used for training or input standardization. Both conditions are evaluated using the same per-finger checkpoints and fold-specific standardization statistics. Table V reports the mean $\pm$ sample standard deviation across five folds after averaging each uncertainty term over samples and fingers.

Mean aleatoric uncertainty remains similar across the two conditions, whereas mean epistemic uncertainty is higher for Force through structure. This contrast indicates greater separation among the force-residual modes rather than greater variance within each mode. Component disagreement can therefore be interpreted as a practical proxy for interactions in which force is not uniquely determined by proprioception. It signals unresolved ambiguity rather than recovering the

missing force information. Future work may address this limitation by combining proprioception with complementary contact-sensitive measurements to capture force variations that are weakly reflected in joint responses.

G. Limitations and Future Work

The virtual spring assumes a single nominal fingertip point, and the residual model predicts one force magnitude per finger. This formulation cannot localize contact, recover force direction, or resolve simultaneous contacts, and its accuracy decreases when different force conditions produce similar proprioceptive histories. Training requires per-finger stiffness calibration and synchronized F/T supervision, although deployment does not require contact sensors. The estimator dataset was collected as continuous interaction streams rather than independent trials, and we therefore use row-level stratified cross-validation. As a result, temporally adjacent history windows may appear in different folds, while generalization across independently collected sequences remains to be evaluated. Manipulation was evaluated over 10 trials per condition using one box specification and one motion, while the simulation–real comparison covers one finger and one three-tap sequence. Future work will investigate richer contact representations, fuse complementary contact-sensitive measurements with proprioception, use calibrated simulation to reduce real-data collection, and evaluate additional objects and tasks.

VI. CONCLUSION

We presented FEEL, an uncertainty-aware residual method for estimating per-finger external force magnitude from proprioception. FEEL improved force-estimation accuracy over the evaluated baselines and enabled force adaptation that reduced box damage while maintaining task success. Mixture-component disagreement provided an empirical signal when force variations were weakly reflected in proprioception. These results establish FEEL as a practical approach to reliable box manipulation without dedicated contact sensors during deployment.

REFERENCES

- [1] J. Tanaka and A. Ogawa, “Cardboard box depalletizing robot using two-surface suction and elastic joint mechanisms: Mechanism proposal and verification,” *Journal of Robotics and Mechatronics (JRM)*, 2019.
- [2] S. Li, A. Keipour, K. Jamieson, N. Hudson, C. Swan, and K. Bekris, “Demonstrating large-scale package manipulation via learned metrics of pick success,” in *Proc. of Robotics: Science and Systems (RSS)*, 2023.
- [3] E. Aduh, F. Wang, D. Randle, K. Wang, P. Shah, C. Mitash, and M. Nambi, “Avoiding object damage in robotic manipulation,” in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024.
- [4] V. Gabler, G. Huber, and D. Wollherr, “A force-sensitive grasping controller using tactile gripper fingers and an industrial position-controlled robot,” in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [5] C. J. Ford, H. Li, J. Lloyd, M. G. Catalano, M. Bianchi, E. Psomopoulou, and N. F. Lepora, “Tactile-driven gentle grasping for human–robot collaborative tasks,” in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2023.
- [6] M. Wang, W. Li, H. Liang, B. Li, K. Althoefer, Y. Su, and H. Liu, “Large-scale deployment of vision-based tactile sensors on multi-fingered grippers,” in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024.
- [7] S. Oh, J. J. Liu, T. Tao, P. Han, K. Shaw, S. Funabashi, R. Salakhutdinov, and D. Pathak, “FACTR 2: Learning external force sensing for commodity robot arms improves policy learning,” *arXiv preprint arXiv:2606.12406*, 2026.
- [8] H. Yu, J. Liu, and D. Ma, “Multi-tactile sensor calibration via motion constraints with tactile measurements,” *The International Journal of Robotics Research (IJRR)*, 2025.
- [9] S. Yun, S. Kang, M. Kim, and S.-S. Yoon, “Input preshaping control of the safe arm with MR-based passive compliant joints,” in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2004, pp. 2788–2793.
- [10] C. B. Yigit, E. Bayraktar, O. Kaya, and P. Boyraz, “External force/torque estimation with only position sensors for antagonistic VSAs,” *IEEE Transactions on Robotics*, vol. 37, no. 2, pp. 675–682, Apr. 2021. *IEEE Transactions on Robotics (T-RO)*, 2021.
- [11] H. T. L. Doan, H. Arita, and K. Tahara, “External sensor-less fingertip force/position estimation framework for a linkage-based under-actuated hand with self-locking mechanism,” in *Proc. of the IEEE/SICE International Symposium on System Integration (SII)*, 2024.
- [12] J. Jung, S. You, D. Kim, and J. Park, “Variable stiffness control via external torque estimation using LSTM,” in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [13] A. Colomé, D. Pardo, G. Alenyà, and C. Torras, “External force estimation during compliant robot manipulation,” in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2013.
- [14] D. Erickson, M. Weber, and I. Sharf, “Contact stiffness and damping estimation for robotic systems,” *The International Journal of Robotics Research (IJRR)*, 2003.
- [15] E. Uhlmann, M. Polte, and J. Blumberg, “Estimation of external force-torque vector based on double encoders of industrial robots using a hybrid Gaussian process regression and joint stiffness model,” *Journal of Machine Engineering*, 2023.
- [16] J. Ballesteros, F. Pastor, J. M. Gómez-de-Gabriel, J. M. Gandarias, A. J. García-Cerezo, and C. Urdiales, “Proprioceptive estimation of forces using underactuated fingers for robot-initiated pHRI,” *Sensors*, 2020.
- [17] D. Nguyen-Tuong and J. Peters, “Using model knowledge for learning inverse dynamics,” in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2010.
- [18] D. Lim, M.-J. Kim, J. Cha, and J. Park, “MOB-Net: Limb-modularized uncertainty torque learning of humanoids for sensorless external torque estimation,” *The International Journal of Robotics Research (IJRR)*, 2025.
- [19] C. M. Bishop, “Mixture density networks,” *Neural Computing Research Group, Aston University, Technical Report NCRG/94/004*, 1994.
- [20] S. Choi, K. Lee, S. Lim, and S. Oh, “Uncertainty-aware learning from demonstration using mixture density networks with sampling-free variance modeling,” in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [21] E. Todorov, T. Erez, and Y. Tassa, “MuJoCo: A physics engine for model-based control,” in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- [22] J. Park, T. Yoon, S. Kim, J. G. Kim, W. Ye, J. Park, Y. Chai, G. Cho, G. Cho, D. Kim, K. Lee, Y.-J. Kim, and S. Choi, “Learning dexterous grasping from sparse taxonomy guidance,” in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2026.
- [23] Y. Lin, Y.-L. Wei, H. Liao, M. Lin, C. Xing, H. Li, D. Zhang, M. Cukosky, and W.-S. Zheng, “TypeTele: Releasing dexterity in teleoperation by dexterous manipulation types,” in *Proc. of the 9th Conference on Robot Learning (CoRL)*, PMLR, vol. 305, pp. 4975–4993, 2025.
- [24] Q. Le Lidec, W. Jallet, L. Montaut, I. Laptev, C. Schmid, and J. Carpentier, “Contact models in robotics: A comparative analysis,” *IEEE Transactions on Robotics (T-RO)*, 2024.
- [25] J. Gao, Y. Zhou, and T. Asfour, “Learning compliance adaptation in contact-rich manipulation,” *arXiv preprint arXiv:2005.00227*, 2020.