TOUCH HAS MEMORY: LEARNING SHARED RESPONSE MODELS FOR FEW-SHOT CALIBRATION OF SOFT TACTILE SENSORS

Anonymous authors
Paper under double-blind review

ABSTRACT

Maintaining accurate force estimation continuously after calibration is challenging for soft tactile sensors, whose responses vary across sensors and depend on loading history. For large-scale manufacturing, the challenge is to capture these history-dependent responses with expressive models while keeping local calibration effort low. We present a few-shot calibration framework that separates shared response modeling from local output calibration. We use four response models: static, path-dependent, temporal relaxation, and hysteresis. Each predicts normal contact force through a compact neural readout. We train the models independently on source recordings, using initial segments for calibration and later segments to optimize force prediction. At each new sensing location, we freeze the models and use a few initial force references to calibrate their output combination. After calibration, all model and calibration parameters remain fixed, while local memory states update causally from the incoming electrical data stream, enabling force estimation without further force references. We evaluate the framework at fourteen sensing locations across three tactile sensor arrays excluded from training, each containing 48 sensing elements. With three fitting references selected from dense initial force recordings, our framework achieves an overall RMSE of 0.375 N under full-window calibration, outperforming the evaluated external baselines. These results suggest that shared response models can support accurate force estimation under changing loading conditions while limiting local learning to output calibration.

1 INTRODUCTION

Calibration of soft tactile sensors should support accurate force estimation throughout operation. This is challenging because the relationship between digital readings and applied force varies across sensors and depends on how each sensor has been loaded. For large-scale manufacturing of tactile sensor arrays, calibration should require only a few force references at each new sensing location.

Similar sensor readings can correspond to different forces depending on the preceding loading and unloading stages (Kim et al., 2020; Li et al., 2026). Calibration therefore needs to address both differences across sensing locations and the effects of loading history. Capturing these history-dependent responses requires expressive models, but learning them at each new sensing location from only a few force references poses an additional challenge.

Prior work has explored both history-dependent modeling and calibration transfer for soft sensors. Kim et al. (2020) combine a time-delay neural network with domain adaptation based on optimal transport to calibrate soft sensors across devices and during long-term use. GenForce transfers force prediction across tactile sensors by learning mappings between marker representations from location-paired data (Chen et al., 2026). Li et al. (2026) develop a calibration framework for soft strain sensors that addresses nonlinearity, hysteresis, cycling attenuation, and batch inconsistency. Methods such as CAVIA and CoDA adapt shared models through low-dimensional context variables (Zintgraf et al., 2019; Kirchmeyer et al., 2022). We thus ask whether response models learned from source tactile arrays can be reused at new sensing locations through few-shot output calibration. After calibration, the shared models and fitted calibration parameters remain fixed, while force estimation uses only current and past sensor readings.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

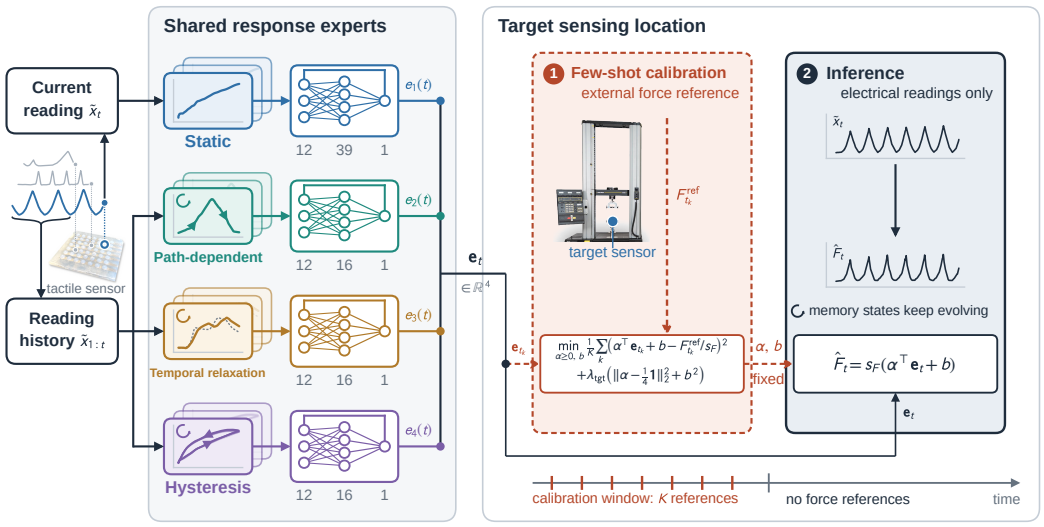


Figure 1: Overview of few-shot calibration with shared response experts. Left: the static expert uses current digital readings, while other experts encode reading history in local memory states. Each expert predicts normalized force $e_j(t)$ through an independently trained neural readout. Right: (1) few-shot calibration uses K initial force references to fit four nonnegative mixture coefficients $\boldsymbol{\alpha}$ and an offset b. (2) The calibrated mixture converts digital readings into force estimates. Parameters stay fixed while memory states evolve. $\tilde{x}_t$ denotes the digital reading, and $\widehat{F}_t$ the predicted force.

We present a framework for few-shot calibration that separates shared response modeling from local output calibration (Figure 1). The shared response models, referred to as experts, represent static, path-dependent, temporal relaxation, and hysteretic responses. Each model uses a compact neural network to map its response representation to force. We train each model independently, using initial segments for calibration and later segments to optimize force prediction. At each target location, the expert library remains fixed, while a regularized output combination is calibrated using a small set of initial force references. After calibration, all parameters remain fixed, while local memory states update causally from incoming sensor readings.

We evaluate the framework on three tactile sensor arrays excluded from training. Each array contains 48 sensing elements (taxels), and our evaluation covers fourteen tested sensing locations across the three arrays. With only three fitting references, our framework achieves an overall RMSE below 0.4 N under both full- and short-window calibration, outperforming the evaluated external baselines. The paper makes three contributions:

1. **Shared response modeling.** We learn shared response models and calibrate only their output combination at each new sensing location.
2. **Learning for prediction after calibration.** We train each model independently using force prediction errors after few-shot calibration.
3. **Evaluation under changing loading conditions.** We demonstrate accurate force estimation under changing loads without further force references after calibration.

2 Related Work

Physical response modeling. Physical response models describe nonlinear responses and loading history through hysteresis models and learned dynamics. Examples include Duhem-based compensation for CNT/PDMS force sensors (Kim & Kim, 2017), rate-dependent compensation for piezoresistive strain sensors (Oliveri et al., 2019), and LSTM models of piezoelectric actuators (Liu et al., 2019). Differentiable Preisach models integrate structured memory with learning for accelerator magnets (Roussel et al., 2022). More broadly, SINDy learns sparse dynamical models (Brunton et al., 2016), while APHYNITY learns constrained corrections to incomplete physical models (Yin et al., 2021). We learn response models from source recordings and reuse them across sensing locations.

Cross-sensor transfer and calibration. Cross-sensor calibration draws on signal alignment, physical modeling, and shared representations. Kim et al. (2020) combine a time-delay neural network with optimal transport for cross-sensor calibration and long-term use, using unlabeled target signals. Li et al. (2026) address hysteresis, nonlinearity, cyclic attenuation, and batch variation through hierarchical physics-guided calibration of soft strain sensors. GenForce transfers force prediction across tactile sensors through shared marker representations and material compensation (Chen et al., 2026). Our work studies how shared response models can be reused through calibration with a few initial force references. After calibration, we estimate force from sensor readings alone.

Few-shot adaptation of shared models. Few-shot methods often adapt a small set of parameters within a shared model. CAVIA updates task-specific context variables while keeping network weights fixed (Zintgraf et al., 2019), and CoDA adapts shared dynamics to new physical systems through low-dimensional context (Kirchmeyer et al., 2022). Chakrabarty et al. (2025) use meta-learning to adapt physically constrained neural state-space models. Bertinetto et al. (2019) differentiate through closed-form ridge regression to learn representations for few-shot adaptation. We differentiate through calibration to improve later force prediction. At new sensing locations, we calibrate only the output combination of the fixed response models.

Predictor combination. Mixtures of experts combine specialist predictions through an input-dependent gating network (Jacobs et al., 1991), while stacked regression estimates nonnegative combination weights from cross-validation predictions (Breiman, 1996). Calandra et al. (2015) use a mixture-of-contacts model to learn contact contributions to inverse dynamics for torque control. Our framework uses predictor combination to calibrate shared models of static and history-dependent sensor responses. Initial force references determine the combination coefficients, which remain fixed during subsequent force estimation while the models' memory states continue to update.

3 Problem Setup

Sensor arrays and target sensing locations. We study few-shot calibration of soft tactile sensors for continued force estimation. Let $i = (d, \ell)$ denote sensing location ℓ on physical tactile array d. At sample index t, $x_{i,t} \in \mathbb{R}$ is the digital reading, $\tau_{i,t}$ is its timestamp in seconds, and $F_{i,t} \in \mathbb{R}$ is the reference normal force in newtons. The reading history with timestamps is $H_{i,t} = ((x_{i,u}, \tau_{i,u}))_{u=1}^{t}$. Shared model parameters θ are learned from source recordings, using synchronized reference forces as ground truth. We calibrate each target independently.

Few-shot initial calibration. Let $\mathcal{W}_i$ denote the sample indices in an initial calibration window of duration Δ_i^{cal} seconds, ending at index c_i. Calibration uses $K_i > 0$ force references at distinct indices:

$$\mathcal{C}_i = \{(t_j, F_{i,t_j})\}_{j=1}^{K_i}, \qquad t_j \in \mathcal{W}_i, \tag{1}$$

The calibration procedure $\mathcal{A}_\theta$ uses H_{i,c_i} and $\mathcal{C}_i$ to fit local output calibration parameters ϕ_i, while keeping θ fixed:

$$\phi_i = \mathcal{A}_\theta(H_{i,c_i}, \mathcal{C}_i). \tag{2}$$

In this study, few-shot refers to parameter fitting with K_i force references. We collect ground-truth force measurements during the initial segment to detect contact and select references.

Force estimation after calibration. After calibration, force estimation uses only the observed reading history:

$$\widehat{F}_{i,t} = f_{\theta,\phi_i}(H_{i,t}), \qquad t > c_i. \tag{3}$$

Both θ and ϕ_i remain fixed, while local memory states update causally from incoming readings. Section 4 describes the estimator and calibration procedure; Section 5 details the evaluation and metrics.

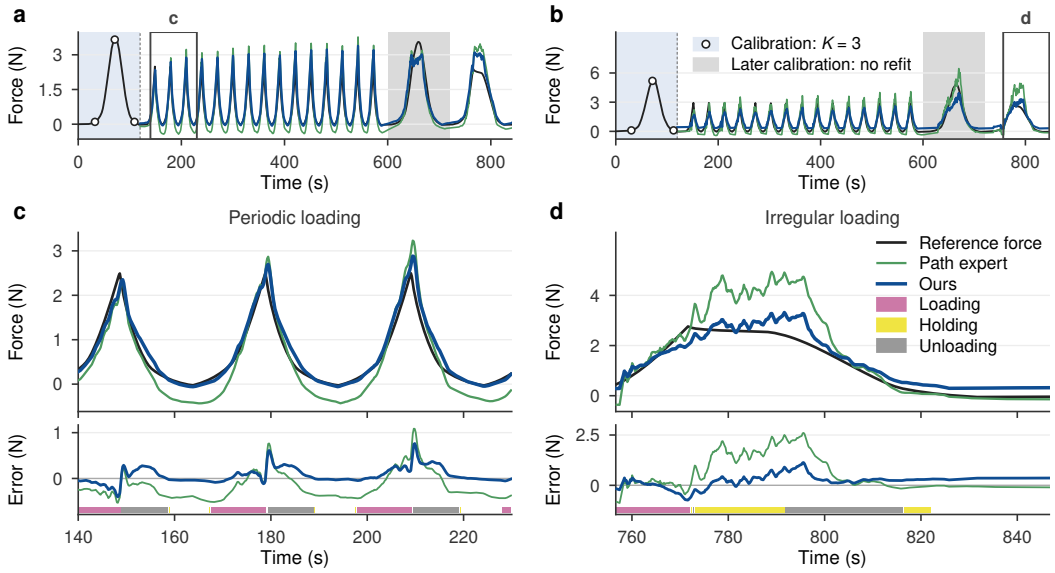


Figure 2: Force estimation after initial calibration. **a, b,** Full recordings from two sensing locations, using full-window calibration ($K = 3$). Grey bands mark later calibration actions without refitting. **c, d,** Enlarged views of the boxed periodic and final irregular loading segments in **a** and **b**, respectively. Each excerpt compares predictions from the mixture and path expert with the reference force. Lower plots show prediction errors, with coloured strips indicating loading, holding, and unloading states. Five-sample centered averages are used only to display predictions.

4 METHOD

We build four shared response models, which we refer to as experts. The neural networks are trained on source data, while calibration at each target sensing location is restricted to their output combination (Figure 1).

4.1 SHARED RESPONSE EXPERT LIBRARY

Each expert uses a fixed response representation and a trainable neural readout. The static expert uses the current digital reading, while the other experts capture path dependence, temporal relaxation, and hysteresis from reading history. We normalize digital readings using a scaler fitted to the training data, yielding $\tilde{x}_{i,t}$. Let $v_{i,t} = \tanh \tilde{x}_{i,t}$, $\Delta\tilde{x}_{i,t} = \tilde{x}_{i,t} - \tilde{x}_{i,t-1}$, and $\Delta\tau_{i,t} = \tau_{i,t} - \tau_{i,t-1}$. For expert j, $\mathbf{z}_{i,t}^{(j)} \in \mathbb{R}^8$ denotes its representation-specific feature vector, with components indexed by k.

Static representation. This expert uses $z_{i,t,k}^{(1)} = \tanh[1.5(\tilde{x}_{i,t} - \xi_k)]$, where ξ_k are fixed centers.

Path-dependent representation. The path expert encodes the direction of sensor reading changes in local memory states, with β_k controlling memory decay as the reading changes:

$$\begin{aligned} \rho_{i,t,k} &= \exp(-\beta_k |\Delta\tilde{x}_{i,t}|), \\ m_{i,t,k} &= \rho_{i,t,k} m_{i,t-1,k} + (1 - \rho_{i,t,k}) \operatorname{sgn}(\Delta\tilde{x}_{i,t}), \qquad z_{i,t,k}^{(2)} = m_{i,t,k}. \end{aligned} \tag{4}$$

Temporal relaxation representation. The temporal relaxation expert tracks sensor readings using memory states $q_{i,t,k}$ with time constants T_k, and uses $v_{i,t} - q_{i,t,k}$ to capture deviations from recent history:

$$\begin{aligned} \gamma_{i,t,k} &= \exp(-\Delta\tau_{i,t}/T_k), \\ q_{i,t,k} &= \gamma_{i,t,k} q_{i,t-1,k} + (1 - \gamma_{i,t,k}) v_{i,t}, \qquad z_{i,t,k}^{(3)} = v_{i,t} - q_{i,t,k}. \end{aligned} \tag{5}$$

Hysteresis representation. The hysteresis expert encodes loading and unloading history using eight play operators with different thresholds (Chen et al., 2013):

$$p_{i,t,k} = \max\{\tilde{x}_{i,t} - r_k, \min\{\tilde{x}_{i,t} + r_k, p_{i,t-1,k}\}\}, \qquad z^{(4)}_{i,t,k} = \frac{\tilde{x}_{i,t} - p_{i,t,k}}{r_k}. \tag{6}$$

The normalized differences between the input and these memory states form the hysteresis features. All representation parameters, including the positive scales β_k, T_k, r_k, remain fixed during training and calibration (values in Appendix B.1).

Independent neural readouts. For expert j, we concatenate $[1, \tilde{x}_{i,t}, v_{i,t}, \tilde{x}_{i,t}v_{i,t}]^\top$ with $\mathbf{z}^{(j)}_{i,t}$ to form $\mathbf{r}^{(j)}_{i,t} \in \mathbb{R}^{12}$. Its independently trained readout predicts normalized force:

$$e^{(j)}_{i,t} = g_{\theta_j}(\mathbf{r}^{(j)}_{i,t}) = \mathbf{c}_j^\top \mathbf{r}^{(j)}_{i,t} + \mathbf{w}_j^\top \tanh(W_j \mathbf{r}^{(j)}_{i,t} + \mathbf{d}_j) + d_{0,j}. \tag{7}$$

Here, $\theta = \{\theta_j\}_{j=1}^4$ comprises 1,270 shared readout parameters. We keep the expert library compact for inference across tactile arrays. A fixed force scale $s_F > 0$, computed from training data, converts normalized predictions to newtons.

4.2 Learning shared experts

We train each expert to predict force after few-shot calibration. Each episode uses an initial calibration set $\mathcal{C}$ and a later prediction segment $\mathcal{Q}$ from the same source recording. We compute memory states from the full preceding reading history. Let $y_t = F_t/s_F$ denote the normalized reference force. We fit a temporary gain a_j and normalized offset d_j for each expert j on $\mathcal{C}$:

$$(\widehat{a}_j, \widehat{d}_j) = \arg\min_{a,d} \sum_{(t,F_t)\in\mathcal{C}} \omega_t (a e^{(j)}_t + d - y_t)^2 + \lambda_{\text{src}}\big((a-1)^2 + d^2\big), \tag{8}$$

where $\omega_t \geq 0$, $\sum_{(t,F_t)\in\mathcal{C}} \omega_t = 1$, and $\lambda_{\text{src}} > 0$. We obtain $(\widehat{a}_j, \widehat{d}_j)$ using a differentiable linear solver (Bertinetto et al., 2019). Each expert independently minimizes the following loss on $\mathcal{Q}$, with normalized sample-duration weights $\nu_t \geq 0$ satisfying $\sum_{t\in\mathcal{Q}} \nu_t = 1$:

$$\mathcal{L}_j(\theta_j) = \mathbb{E}_{\text{episode}}\left[\sum_{t\in\mathcal{Q}} \nu_t \left\{0.75(\widehat{a}_j e^{(j)}_t + \widehat{d}_j - y_t)^2 + 0.25(e^{(j)}_t - y_t)^2\right\}\right]. \tag{9}$$

The calibrated term promotes accurate prediction after adaptation, while the uncalibrated term anchors the expert outputs to the normalized force targets. We update only θ_j, propagating gradients through the readout and calibration solve. The temporary gain and offset are discarded after each episode.

4.3 Few-shot calibration and causal prediction

Local output calibration. At target location i, we freeze θ and process H_{i,c_i} to obtain the expert predictions $\mathbf{e}_{i,t} = [e^{(1)}_{i,t}, \ldots, e^{(4)}_{i,t}]^\top$ at the reference times. Using the K_i force references in $\mathcal{C}_i$, we fit $\phi_i = (\boldsymbol{\alpha}_i, b_i)$, with four nonnegative coefficients and an unconstrained offset:

$$(\boldsymbol{\alpha}_i, b_i) = \arg\min_{\boldsymbol{\alpha}\geq 0,\, b} \frac{1}{K_i} \sum_{(t,F_{i,t})\in\mathcal{C}_i} \left(\boldsymbol{\alpha}^\top \mathbf{e}_{i,t} + b - \frac{F_{i,t}}{s_F}\right)^2 + \lambda_{\text{tgt}} \left(\left\|\boldsymbol{\alpha} - \tfrac{1}{4}\mathbf{1}\right\|_2^2 + b^2\right). \tag{10}$$

The coefficient sum can vary to adjust the overall gain. The regularizer favors coefficients of $1/4$ and a zero offset to stabilize calibration with few references. With $\lambda_{\text{tgt}} > 0$, the objective remains strictly convex even with only three references, yielding a unique calibration solution.

Causal force estimation. For $t > c_i$, the calibrated mixture predicts force as

$$\widehat{F}_{i,t} = s_F\big(\boldsymbol{\alpha}_i^\top \mathbf{e}_{i,t} + b_i\big). \tag{11}$$

For each incoming sensor reading, we update the memory states and compute force predictions using all four neural networks. States carry over from calibration, while θ and ϕ_i remain fixed.

5 EXPERIMENTS

We evaluate whether shared response models support accurate force estimation at sensing locations on new arrays after few-shot output calibration. We then examine how the number of fitting references, loading conditions, expert combination, and source training affect prediction.

5.1 EXPERIMENTAL SETUP

Data and model selection. Our source dataset contains 9 hours of recordings. We train on data from four source tactile arrays and construct training episodes using sliding windows. Evaluation covers 19 long recordings at 14 tested sensing locations on three additional tactile arrays. Each target array contains 48 sensing elements (taxels). Reference forces are measured at 10 Hz using an Instron 5969 equipped with a 50 N load cell. Digital readings are then synchronized offline with force references, causally filtered, and sampled at 5 Hz. Data from the evaluation arrays are not used when training the shared response models. Baseline hyperparameters are selected using validation on source tactile arrays and held fixed during target evaluation (Appendix A.2).

Baselines. We compare a Prandtl–Ishlinskii (PI) model pretrained on source data, a raw-signal GRU, concatenated-feature MLP and GRU models, TDNN+OT (Kim et al., 2020), D_Former (Li et al., 2026), and CAVIA (Zintgraf et al., 2019). For the configurations in Table 1, target calibration fits the PI coefficients, an output gain and offset for the GRUs, the MLP output layer, or the CAVIA context vector. TDNN+OT adapts using unlabeled target readings, while D_Former is evaluated without target parameter updates.

Calibration and evaluation. Each recording is calibrated independently using $K = 3$ fitting references in the main comparison and $K = 1, \ldots, 32$ in the budget study. We compare the full initial calibration window with a short window lasting up to 30 s after contact detection, ending no later than the full window. Default force-based selection uses a 0.1 N contact threshold and selects evenly spaced ranks among time-ordered samples above it. As an ablation study, time-based selection uses three references at preset times within the full window, without force-based contact detection, screening, or replacement. K counts only the force references used for fitting. All methods share the same preprocessing and window definitions, and methods using K references receive identical reference indices. All settings are evaluated on the same samples after the full window, with no further fitting. We report time-weighted RMSE (N), averaging seeds within recordings, repeat recordings within locations, and locations within arrays, then averaging target arrays equally. Appendix A.1 gives the window and preprocessing details.

Table 1: RMSE (N), averaged equally across target arrays. All settings are evaluated on the same samples after the full calibration window. Time-based selection also uses this window for calibration. Methods using target force labels share $K = 3$ references within each setting; † denotes methods using no target force labels for adaptation. Reduction reports the relative RMSE reduction of Ours versus each baseline. Underlining marks the best external baseline.

	Full window		Short window (≤30 s)		Time-based selection
Method	RMSE	Reduction	RMSE	Reduction	RMSE
Physics model					
Classical PI	0.820	54%	0.631	37%	0.767
Learned features					
Raw + GRU	1.058	65%	0.987	60%	0.895
Concat features + MLP	0.769	51%	0.915	57%	0.653
Concat features + GRU	0.748	50%	1.395	72%	<u>0.590</u>
Adaptation					
TDNN + OT† (Kim et al., 2020)	0.680	45%	0.850	53%	0.680
D_Former† (Li et al., 2026)	0.963	61%	0.963	59%	0.963
CAVIA (Zintgraf et al., 2019)	<u>0.652</u>	42%	<u>0.607</u>	35%	0.597
Ours	**0.375**	—	**0.397**	—	**0.415**

5.2 Comparison with baseline methods

With force-based selection and $K = 3$, our method achieves the lowest RMSE in Table 1 under both calibration windows: 0.375 N with the full window and 0.397 N with the short window. These errors are respectively 42.5% and 34.6% lower than those of CAVIA, the best-performing external baseline in Table 1 under both settings.

Under full-window calibration with force-based selection, our method achieves lower mean RMSE than each external baseline in Table 1 on all three tactile arrays, with RMSE for individual arrays ranging from 0.303 to 0.500 N. Table 3 reports additional adaptation variants. For the raw-signal GRU, last-layer adaptation yields a full-window RMSE of 0.614 N, below CAVIA's 0.652 N, but a short-window RMSE of 1.836 N. The mixture has lower mean RMSE than all of these external configurations under both windows.

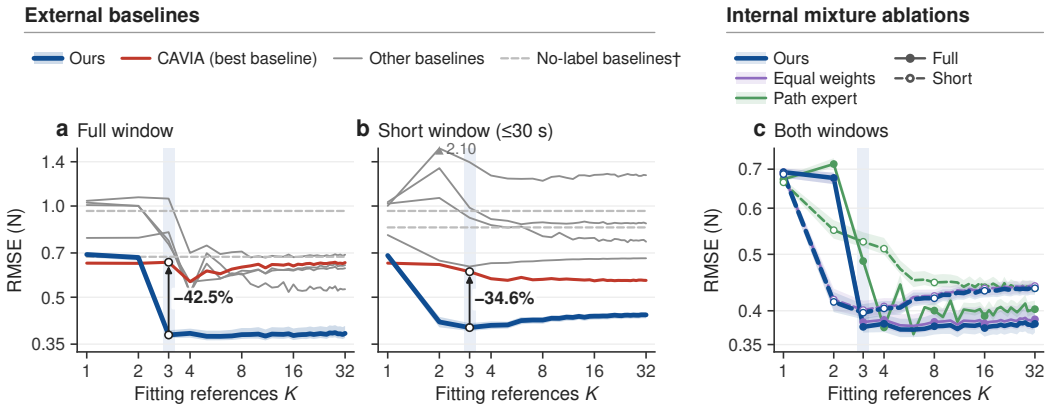


Figure 3: Force-estimation performance versus the number of fitting references K. **a, b,** Our method compared with external baselines under full- and short-window calibration, respectively. **c,** Comparison of the path expert, an equal-weight mixture, and our mixture with locally fitted weights. Solid and dashed lines indicate full- and short-window calibration, respectively. Errors are RMSE (N), and both axes are logarithmic. Shaded columns mark $K = 3$ (Table 1); arrows in **a, b** show relative RMSE reductions versus CAVIA. † indicates methods using no target force labels for adaptation.

5.3 Few-shot calibration

With full-window calibration, RMSE is 0.375 N using three fitting references and 0.374 N using sixteen, with little change over $K = 3, \ldots, 32$ (Figure 3a). Under short-window calibration, more references do not consistently lower error: RMSE is 0.397 N at $K = 3$ and 0.433 N at $K = 16$ (Figure 3b). Reference sets are selected separately for each K. Our method outperforms all evaluated external baselines for $K = 3, \ldots, 32$ with the full window and $K = 2, \ldots, 32$ with the short window; CAVIA performs best at the smaller budgets.

The mixture outperforms both the path expert and an equal-weight mixture at $K = 3$ under both windows (Figure 3c). Its advantage depends on the reference budget: the path expert achieves slightly lower RMSE at $K = 4$ under full-window calibration.

With time-based selection, three references yield an RMSE of 0.415 N, compared with 0.375 N using force-based selection (Table 1). Equal weights and the path expert yield 0.420 and 0.441 N, respectively. Offline synchronization and force interpolation remain unchanged (Appendix A.1).

5.4 Force estimation under changing loading conditions

After initial calibration, model parameters remain fixed while memory states update with incoming readings. We evaluate force estimation on two subsets: Rapid, comprising loaded samples ($F >$ 0.1 N) with an absolute force change exceeding 0.5 N over one second; and Irregular, the recorded segment of simulated free-form pressing. These subsets can overlap.

Table 2: Expert and mixture ablations at $K = 3$ (RMSE in N). Single experts and equal weights fit gain and offset; expert removal refits the remaining coefficients and offset. *Full window. Bold marks column minima; blue/red shading indicates lower/higher RMSE than Ours.

Method	Full window	Short window (≤30 s)	Rapid*	Irregular*
Individual experts				
Static expert	0.8100	0.7415	1.1269	1.2408
Path expert	0.4868	0.5254	0.9808	0.6880
Temporal relaxation expert	0.7707	0.6610	1.7008	0.7662
Hysteresis expert	0.5554	0.8163	0.8190	0.8829
Expert combinations				
Equal weights	0.3826	0.4012	0.6002	0.3681
MLP16 static readout	0.3766	0.3984	0.5769	0.3700
Ours	0.3754	**0.3970**	0.5864	0.3733
Expert removal				
w/o static	0.3724	0.4117	0.5623	0.3275
w/o path	0.3870	0.4301	0.5628	0.3259
w/o temporal relaxation	**0.3678**	0.4060	**0.5307**	0.3796
w/o hysteresis	0.4192	0.4614	0.6190	**0.3162**

With full-window calibration, the mixture reduces RMSE relative to the path expert from 0.981 to 0.586 N on Rapid and from 0.688 to 0.373 N on Irregular (Table 2). The best-performing expert combination varies with loading conditions: on Irregular, both equal weights and the mixture without the hysteresis expert achieve lower RMSE than our full mixture.

Figure 2 illustrates force estimation after calibration at two sensing locations. Compared with the path expert, the mixture reduces negative bias between periodic presses (Figure 2c) and overestimation during the hold in the irregular-loading example (Figure 2d), without further parameter fitting.

5.5 Ablation analysis

We examine expert combination, local output calibration, and training at $K = 3$. Appendix C provides the experimental details.

Expert combination. The mixture achieves lower overall RMSE than any individual expert under both calibration windows (Table 2). Removing any expert increases short-window RMSE, whereas removing the static or temporal relaxation expert slightly reduces full-window RMSE.

Local output calibration. With full-window calibration, an equal-weight mixture with fitted gain and offset reduces RMSE from the path expert's 0.487 to 0.383 N. Locally fitting relative weights further reduces it to 0.375 N (Figure 4a). Most of this improvement therefore comes from combining expert predictions with simple output calibration. The benefit of local weight fitting depends on the calibration window: source-learned proportions achieve lower full-window RMSE (0.367 N),

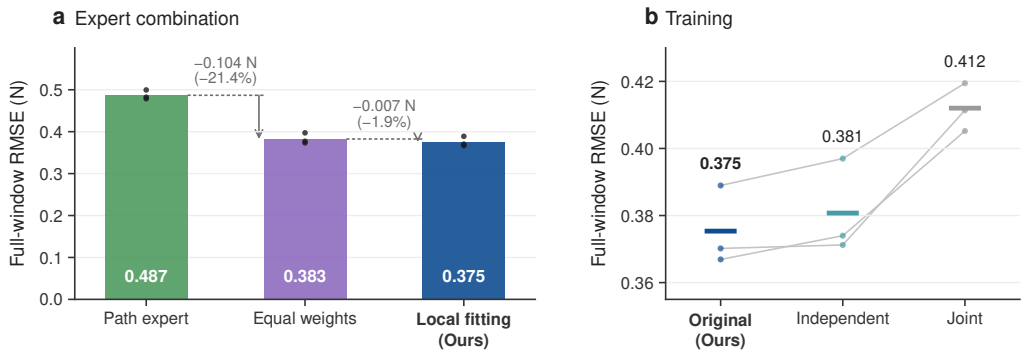


Figure 4: **Expert combination and continuation training** at $K = 3$ under full-window calibration (RMSE in N). **a,** Path expert, equal-weight mixture, and our mixture with locally fitted weights. **b,** Original model and independent or joint continuation from the same checkpoints.

while local fitting achieves the lowest short-window RMSE among the evaluated combination rules (Table 5).

Continuation training. Starting from the same final checkpoints, independent continuation yields lower overall target RMSE than joint continuation under both calibration windows. Neither improves on the original model (Figure 4b; Table 6).

6 Discussion and Conclusion

We present a few-shot calibration framework that separates shared response modeling from local output calibration. Response experts trained independently on source tactile arrays are reused at new sensing locations, where a few initial force references calibrate their output combination. After calibration, all parameters remain fixed while local memory states update from incoming readings, enabling force estimation under changing loads without further force references.

Across fourteen sensing locations on three tactile arrays excluded from training, the framework achieves an overall RMSE of 0.375 N under full-window calibration using three fitting references selected from dense initial force recordings. Under this setting, combining expert predictions with simple gain-and-offset calibration accounts for most of the improvement over the path expert. The additional benefit of local weight fitting depends on the calibration window.

Future work should evaluate transfer across more tactile arrays and different tactile sensor types, long-term operation, and calibration reuse when reading histories are interrupted or incomplete.

AI use statement

Generative AI tools were used to polish the language of the manuscript after the authors had drafted it and to assist with literature searches. Numerical results are verified using recorded data and the experimental code; no experimental data were generated by AI. The authors review AI-assisted content and take responsibility for all claims, citations, equations, code, and research artifacts.

Reproducibility statement

Sections 3–5 and Appendices A–C describe the task, model, calibration procedure, and evaluation protocol. After the review process, we will publicly release the raw measurements, source and target splits, training configurations, model checkpoints, calibration reference indices, and code for preprocessing, synchronization, training, and evaluation.

References

Luca Bertinetto, João F. Henriques, Philip H. S. Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=HyxnZh0ct7.

Leo Breiman. Stacked regressions. *Machine Learning*, 24(1):49–64, 1996. doi: 10.1023/A:1018046112532. URL https://statistics.berkeley.edu/sites/default/files/tech-reports/367.pdf.

Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016. doi: 10.1073/pnas.1517384113. URL https://arxiv.org/abs/1509.03580.

Roberto Calandra, Serena Ivaldi, Marc Peter Deisenroth, and Jan Peters. Learning torque control in presence of contacts using tactile sensing from robot skin. In *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, pp. 690–695. IEEE, 2015. doi: 10.1109/HUMANOIDS.2015.7363429. URL https://www.dfki.de/web/forschung/projekte-publikationen/publikation/13804.

Ankush Chakrabarty, Gordon Wichern, Vedang M. Deshpande, Abraham P. Vinod, Karl Berntorp, and Christopher R. Laughman. Meta-learning for physically-constrained neural system identification. *Neurocomputing*, 651:130945, 2025. doi: 10.1016/j.neucom.2025.130945. URL https://www.merl.com/publications/TR2025-159.

Yuansheng Chen, Jinhao Qiu, and Hongjun Sun. A hybrid model of Prandtl–Ishlinskii operator and neural network for hysteresis compensation in piezoelectric actuators. *International Journal of Applied Electromagnetics and Mechanics*, 41(3):335–347, 2013. doi: 10.3233/JAE-131648.

Zhuo Chen, Ni Ou, Xuyang Zhang, Zhiyuan Wu, Yongqiang Zhao, Yupeng Wang, Emmanouil Spyrakos Papastavridis, Nathan Lepora, Lorenzo Jamone, Jiankang Deng, and Shan Luo. Training tactile sensors to learn force sensing from each other. *Nature Communications*, 17 (1):2101, 2026. doi: 10.1038/s41467-026-68753-1.

Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. Adaptive mixtures of local experts. *Neural Computation*, 3(1):79–87, 1991. doi: 10.1162/neco.1991.3.1.79. URL https://www.cs.toronto.edu/~hinton/absps/jjnh91.pdf.

DongWook Kim, Junghan Kwon, Byungjun Jeon, and Yong-Lae Park. Adaptive calibration of soft sensors using optimal transportation transfer learning for mass production and long-term usage. *Advanced Intelligent Systems*, 2(6):1900178, 2020. doi: 10.1002/aisy.201900178.

Ji-Sik Kim and Gi-Woo Kim. Hysteresis compensation of piezoresistive carbon nanotube/polydimethylsiloxane composite-based force sensors. *Sensors*, 17(2):229, 2017. doi: 10.3390/s17020229.

Matthieu Kirchmeyer, Yuan Yin, Jeremie Dona, Nicolas Baskiotis, Alain Rakotomamonjy, and Patrick Gallinari. Generalizing to new physical systems via context-informed dynamics model. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 11283–11301. PMLR, 2022. URL https://proceedings.mlr.press/v162/kirchmeyer22a.html.

Jiali Li, Haitao Yang, Lanjing Wang, Nungsiong Lai, Chong Sun, Fuhui Zhou, Qiulei Liu, Wei Wang, Jinghan Li, Chengzhi Hu, Xiaodong Chen, and Xiaonan Wang. Computationally intelligent calibration framework for durable soft strain sensors. *Nature Communications*, 17(1):6419, 2026. doi: 10.1038/s41467-026-72113-4.

Yanfang Liu, Rui Zhou, and Mingying Huo. Long short term memory network is capable of capturing complex hysteretic dynamics in piezoelectric actuators. *Electronics Letters*, 55(2):80–82, 2019. doi: 10.1049/el.2018.7490.

Alberto Oliveri, Martina Maselli, Matteo Lodi, Marco Storace, and Matteo Cianchetti. Model-based compensation of rate-dependent hysteresis in a piezoresistive strain sensor. *IEEE Transactions on Industrial Electronics*, 66(10):8205–8213, 2019. doi: 10.1109/TIE.2018.2884204.

R. Roussel, A. Edelen, D. Ratner, K. Dubey, J. P. Gonzalez-Aguilera, Y. K. Kim, and N. Kuklev. Differentiable Preisach modeling for characterization and optimization of particle accelerator systems with hysteresis. *Physical Review Letters*, 128(20):204801, 2022. doi: 10.1103/PhysRevLett.128.204801.

Yuan Yin, Vincent Le Guen, Jérémie Dona, Emmanuel de Bézenac, Ibrahim Ayed, Nicolas Thome, and Patrick Gallinari. Augmenting physical models with deep networks for complex dynamics forecasting. In *International Conference on Learning Representations*, 2021. URL https://arxiv.org/abs/2010.04456.

Luisa Zintgraf, Kyriacos Shiarli, Vitaly Kurin, Katja Hofmann, and Shimon Whiteson. Fast context adaptation via meta-learning. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 7693–7702. PMLR, 2019. URL https://proceedings.mlr.press/v97/zintgraf19a.html.