ESN-Based Reservoir Control for Single-Finger Dexterous Movement under Unknown Contact Conditions

Qingshuo Hu,Ziqi Wang,Fan Ye,Zhen Yin

***Abstract*—Contact-rich manipulation requires rapid joint-torque adjustment according to local contact feedback, but the mapping from fingertip force and joint states to multi-joint torques is nonlinear and temporally coupled. Existing policies use tactile or force feedback for reactive manipulation, yet such feedback is often coupled with high-level action representations rather than direct joint-torque generation, limiting timely adaptation to changing environment conditions. To isolate this contact-level problem, we use single-finger pressing as a controlled validation scenario and study torque generation conditioned on contact force and joint states. From a single 40s pressing sequence, we construct a time-series dataset for torque learning based on reservoir control. A reservoir composed of a 50-node Echo State Network (ESN) is trained to generate multi-joint torque commands. Under trajectory reference inputs, the trained model achieves offline torque-prediction R^2 scores above 0.997 for all three joint-torque outputs. When deployed as a closed-loop torque controller, it reduces tracking MAE and RMSE by 63.6% and 62.6%, respectively, compared with PID control. We further evaluate a reference-free ESN behavior model, which generates coordinated multi-joint motion from contact force and joint states alone and adjusts its motion pattern under changing environment conditions. These results indicate that ESN-based reservoir control can support small-data torque learning and adaptive motion generation at contact level.**

I. Introduction

Contact-rich manipulation requires dexterous robotic fingers to adjust joint torques rapidly according to local contact feedback during pressing, grasping, sliding, and pushing. Unlike free-space motion, contact interaction couples fingertip force, joint states, desired motion, and contact conditions. The resulting mapping from force and joint states to multi-joint torques is nonlinear and temporally coupled, making low-level contact execution difficult to model explicitly.

Recent studies have advanced contact-rich manipulation through model-based compliant control [2, 6], contact modeling [5], and tactile/force-guided policy learning [4, 1, 3]. Model-driven methods provide interpretable compliant responses, but often depend on contact assumptions, physical parameters, or online optimization. Learning-based visual-tactile policies improve reactive manipulation with tactile or force feedback, yet such feedback is commonly coupled with high-level action representations rather than direct joint-torque generation. As a result, timely torque-level adaptation from local contact force and joint states remains less explored.

In this work, we study short-term torque generation for single-finger contact interaction. Single-finger pressing is used as a controlled validation scenario to isolate the local coupling among fingertip force, joint states, and multi-joint torques. We formulate the problem as force–proprioception-to-torque learning with optional target motion states. An Echo State Network (ESN) is adopted because its recurrent reservoir provides short-term memory, while only the readout layer needs to be trained, enabling compact learning from limited contact data.

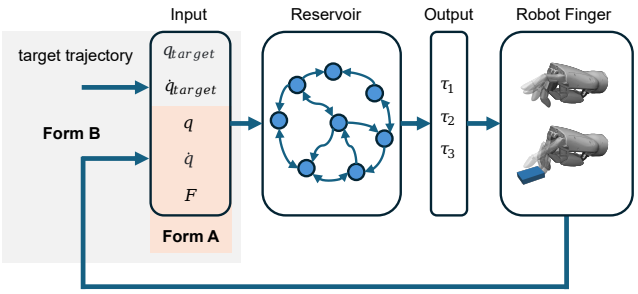

Fig. 1. Overview of the ESN-based reservoir framework for single-finger contact interaction. Form A uses $[q_t, \dot{q}_t, F_t]$ for reference-free behavior learning, while Form B uses $[q_t, \dot{q}_t, F_t, q_t^{\text{target}}, \dot{q}_t^{\text{target}}]$ for trajectory-conditioned torque control. Both forms output the multi-joint torque vector τ_t.

The main contributions of this work are:

- We demonstrate sample-efficient ESN-based torque learning for single-finger contact interaction from a short pressing sequence.
- We show that the learned ESN improves closed-loop trajectory tracking over PID control with a moderate increase in torque effort.
- We provide evidence that reference-free ESN dynamics can produce adaptive multi-joint behavior without explicit obstacle modeling or replanning.

II. Related Work

A. Model-based compliant control and contact modeling

Model-based compliant control provides a principled basis for robotic contact interaction. Virtual model control, model predictive control, and explicit contact modeling generate compliant responses through physical structures, constraints, or optimization objectives [6, 5]. These methods are interpretable when contact models are reliable, but usually require predefined assumptions, physical parameters, or online optimization. When local contact conditions change rapidly, explicitly modeling the short-term coupling among fingertip

force, joint states, and torques becomes difficult. This motivates lightweight torque-level learning for local finger contact.

B. Dexterous manipulation with tactile/force feedback

Tactile and force feedback are increasingly used to improve contact-rich manipulation. Reactive Diffusion Policy and ImplicitRDP use slow-fast visual-tactile or visual-force learning for reactive control under high-frequency contact feedback [4, 1]. Other work uses tactile feedback for dexterous non-prehensile manipulation or grasp-force regulation [3, 2]. These studies show the value of contact feedback for robustness and adaptability. However, tactile or force information is often used to guide action correction, end-effector motion, grasp-force planning, or object-level manipulation. In contrast, this work directly learns the short-term mapping from fingertip force and joint states to multi-joint torque commands.

III. Methodology

A. Problem Formulation

This work studies ESN-based torque learning and control for a single robotic finger during contact interaction. For a finger with n joints, the joint position, joint velocity, fingertip contact force, and joint torque at time step t are denoted by q_t, $\dot{q}_t$, F_t, and τ_t, respectively, where $q_t, \dot{q}_t, \tau_t \in \mathbb{R}^n$ and $F_t \in \mathbb{R}^m$.

For trajectory-conditioned control, the objective is to learn a torque policy that maps the current contact force and joint-state input, together with the target motion state to a multi-joint torque command:

$$u_t^{ctrl} = [q_t, \dot{q}_t, F_t, q_t^{\text{target}}, \dot{q}_t^{\text{target}}], \quad \tau_t^{ESN} = \pi_\theta^{ctrl}(u_t^{ctrl}). \tag{1}$$

where q_t^{target} and $\dot{q}_t^{\text{target}}$ denote the target joint position and velocity. The ESN directly outputs joint torques rather than torque corrections.

We also define a reduced contact-state input for reference-free behavior learning:

$$u_t^{beh} = [q_t, \dot{q}_t, F_t], \quad \tau_t^{ESN} = \pi_\theta^{beh}(u_t^{beh}). \tag{2}$$

This form evaluates whether the ESN can generate contact-induced torque behavior from contact force and joint states alone, without explicit target-trajectory conditioning.

B. ESN Model

Given input u_t, the ESN reservoir state is updated by

$$x_t = (1-\alpha)x_{t-1} + \alpha \tanh(W_{in}u_t + W_{res}x_{t-1} + b), \tag{3}$$

where x_t is the reservoir state, W_{in} is the input matrix, W_{res} is the recurrent reservoir matrix, α is the leaking rate, and b is the bias. During training, W_{in} and W_{res} are fixed, and only the readout layer is optimized.

The ESN outputs the multi-joint torque through a linear readout:

$$\hat{\tau}_t = W_{out}[x_t, u_t, 1]. \tag{4}$$

The readout weights are trained by ridge regression:

$$\min_{W_{out}} \sum_t \|\tau_t - W_{out}[x_t, u_t, 1]\|^2 + \lambda \|W_{out}\|^2, \tag{5}$$

where τ_t is the target joint torque and λ is the regularization coefficient.

C. Training and Control Procedure

The framework includes offline training and closed-loop deployment. During offline training, the ESN uses u_t^{ctrl} as input and the corresponding joint torque τ_t as the supervised output. This stage evaluates whether the ESN can learn target-conditioned torque generation from single-finger contact data.

During closed-loop deployment, the current joint state, contact force, and target trajectory state are fed into the trained ESN, and the torque command in Eq. (1) is directly applied to the finger. The tracking performance is compared with PID control under the same prescribed trajectories.

Finally, we evaluate a reference-free behavior model in which the ESN receives only the reduced contact-state input in Eq. (2). This setting tests whether reservoir dynamics can reproduce adaptive multi-joint torque behavior from the interaction state alone.

IV. Experiments

A. ESN Controller Training and Offline Validation

To train the ESN torque controller, we construct a contact-force–joint-state–torque time-series dataset from a single 40s pressing sequence. During pressing, q_t, $\dot{q}_t$, and F_t are recorded, the corresponding multi-joint torque τ_t is computed in simulation, and each time step is paired with q_t^{target} and $\dot{q}_t^{\text{target}}$ for trajectory-conditioned learning.

For controller training, the ESN uses the trajectory reference input vector defined in Eq. (1). In the offline training dataset, the target motion state is obtained from the next time step of the recorded trajectory:

$$q_t^{\text{target}} = q_{t+1}, \quad \dot{q}_t^{\text{target}} = \dot{q}_{t+1}. \tag{6}$$

This formulation trains the ESN to generate torque toward the next desired state while incorporating fingertip contact feedback.

The ESN controller is obtained by training only the readout weights while keeping the reservoir fixed. We set $N_r = 50$, evaluate 100 random reservoirs, and select the final controller according to validation torque-prediction error. The selected reservoir uses seed 30, input dimension 13, output dimension 3, leak rate $\alpha = 0.3$, and spectral radius $\rho = 1.5$. It achieves an offline torque-prediction R^2 score higher than 0.997 for all three joint-torque outputs in the test set and is exported as the runtime torque controller.

The ESN outputs the estimated torque command:

$$\hat{\tau}_t = \pi_\theta(u_t^{\text{ctrl}}). \tag{7}$$

The offline validation error is quantified by the mean squared error between the target torque and the ESN-generated torque:

$$E_\tau = \frac{1}{T}\sum_{t=1}^{T} \|\tau_t - \hat{\tau}_t\|^2. \tag{8}$$

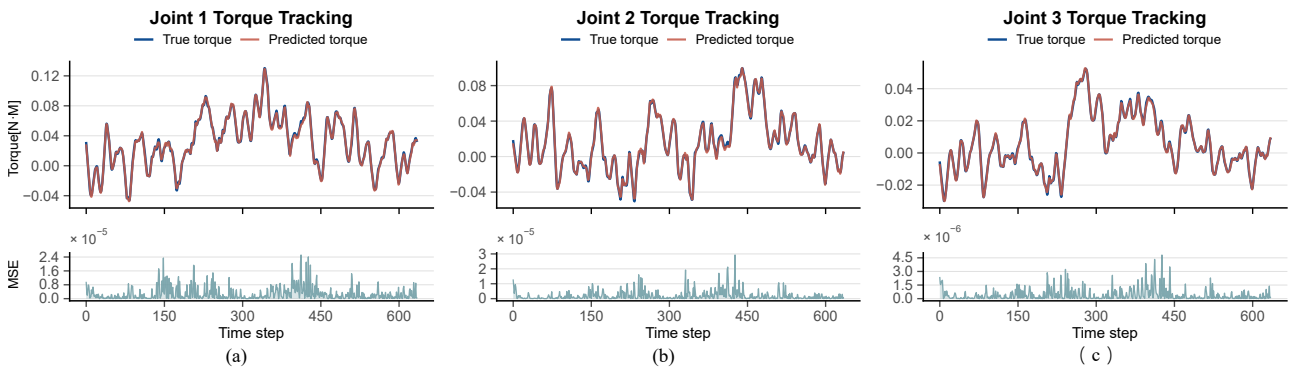

Fig. 2. Offline validation of the trained ESN torque controller. Ground-truth and ESN-generated torques are compared across joints, with the corresponding MSE distributions shown over time.

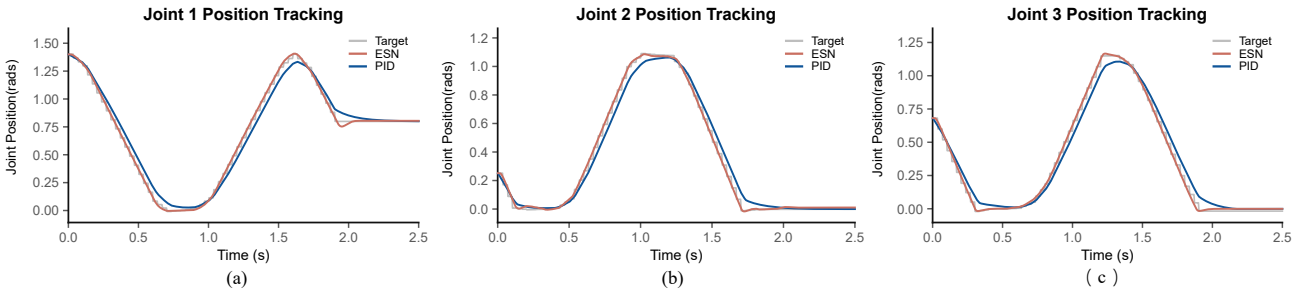

Fig. 3. Joint position tracking under prescribed target trajectories. Target, ESN, and PID trajectories are compared for Joints 1–3.

Fig. 2 shows that the ESN-generated torques closely follow the target torques across the evaluated joints, indicating that the reservoir model learns the trajectory-conditioned torque-generation mapping. Larger errors occur mainly near rapid torque transitions and local peaks, where contact-induced dynamics change more sharply.

B. Closed-Loop Control under Prescribed Trajectories

After offline validation, the trained ESN torque controller is deployed in a closed-loop trajectory-tracking task, with PID control used as the baseline.

At each time step, the ESN receives the trajectory reference input in Eq. (1); the torque command in Eq. (7) is directly applied to the simulated finger.

For comparison, a PID controller is implemented under the same target trajectories and simulation conditions. The gains are tuned on the pressing trajectory by first increasing K_p to improve tracking responsiveness and then adding K_d to suppress contact-induced oscillation. The integral term is omitted because the task focuses on dynamic tracking under contact transients rather than long-term static error compensation. The final gains are $K_p = 20.0$, $K_d = 2.0$, and $K_i = 0$.

The PID torque command is given by:

$$\tau_t^{\mathrm{PID}} = K_p e_t + K_d \dot{e}_t + K_i \sum_{k=1}^{t} e_k \Delta t, \tag{9}$$

where $e_t = q_t^{\mathrm{target}} - q_t$ and $\dot{e}_t = \dot{q}_t^{\mathrm{target}} - \dot{q}_t$ denote the joint position and velocity tracking errors, respectively.

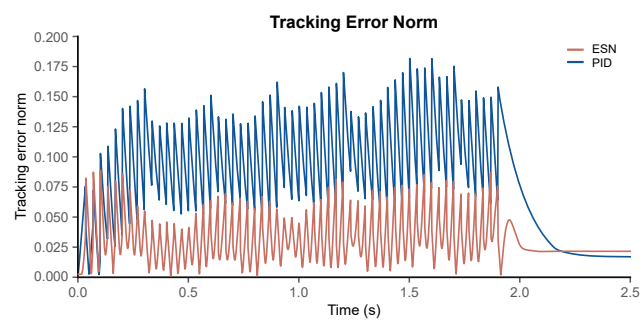

Fig. 4. Tracking-error norm comparison between ESN and PID control over time. The ESN maintains a lower tracking-error norm for most of the trajectory.

The two controllers are compared using tracking MAE, RMSE, tracking-error norm, and torque RMS.

Fig. 3 shows that the ESN follows the prescribed target trajectories more closely than PID for all three joints. During both rising and falling phases, the ESN shows smaller deviations, whereas PID exhibits more visible tracking lag and steady offset.

Fig. 4 shows that the ESN maintains a lower tracking-error norm than PID for most of the trajectory, indicating more accurate closed-loop tracking during contact.

Fig. 5 further summarizes the quantitative performance. The ESN achieves a tracking MAE of 0.0307 and RMSE of 0.0352,

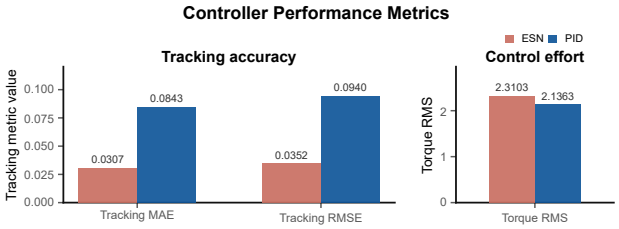

Fig. 5. Quantitative comparison of ESN and PID control in terms of tracking MAE, tracking RMSE, and torque RMS. The ESN improves tracking accuracy with a moderate increase in torque effort.

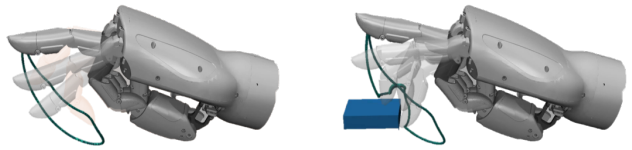

(a) (b)

Fig. 6. Reference-free ESN self-adaptive behavior learning. (a) Coordinated multi-joint motion generated by the learned ESN under the original condition. (b) When an additional obstacle changes the environment condition, the ESN produces a different motion pattern without explicit target-trajectory conditioning, obstacle modeling, or replanning.

compared with 0.0843 and 0.0940 for PID, corresponding to reductions of 63.6% and 62.6%, respectively. The torque RMS increases moderately from 2.1363 to 2.3103.

C. Reference-Free ESN Self-Adaptive Behavior Learning

We further evaluate whether the ESN can generate self-adaptive single-finger behavior without explicit target trajectories. Unlike the trajectory-conditioned controller, this model receives only the reduced contact-state input in Eq. (2). Therefore, no target joint position, target joint velocity, obstacle model, or replanning module is provided during deployment. The generated torque is determined only by the current contact force, joint states, and the internal reservoir dynamics.

As shown in Fig. 6(a), under the original condition, the reference-free ESN generates a coordinated multi-joint motion pattern from contact force and joint states alone. When an additional obstacle is introduced, the environment condition changes and produces different contact-force-and-joint-state feedback. As a result, the reservoir state evolves differently and generates a modified torque sequence, leading to an adjusted finger trajectory, as shown in Fig. 6(b). This behavior indicates that the ESN does not simply replay a fixed motion pattern. Instead, it modulates its torque output according to the current interaction state, suggesting a self-adaptive contact response without explicit obstacle models or trajectory replanning.

V. Conclusion

This work presents an ESN-based reservoir framework for short-term contact-level torque generation. Rather than learning a whole-hand task-level policy, the proposed method focuses on the local mapping from fingertip contact force and joint states to multi-joint torque commands. Single-finger pressing is used as a controlled validation scenario to isolate this contact-level problem. Two input forms are considered: a reference-free form $(q, \dot{q}, F)$ for behavior learning and a trajectory-conditioned form $(q, \dot{q}, F, q^{\text{target}}, \dot{q}^{\text{target}})$ for controller training and closed-loop tracking.

The results show that a compact ESN can learn accurate torque generation from limited contact data and can be deployed as a closed-loop torque controller. Compared with PID control, the ESN reduces tracking MAE and RMSE by 63.6% and 62.6%, respectively, with a moderate increase in torque RMS. The reference-free experiment further shows that the ESN can generate coordinated multi-joint motion from contact force and joint states alone and adjust its motion pattern under changing environment conditions. These findings indicate that ESN-based reservoir control provides a lightweight small-data approach for contact-level torque learning and adaptive motion generation. Future work will extend the framework to full-hand multi-finger coordination and hardware validation under richer multi-contact conditions.

Acknowledgment

This work was supported by Linkerbot Beijing Technology Co., Ltd. through the Linker Hand L20 platform.

References

[1] Wendi Chen, Han Xue, Yi Wang, Fangyuan Zhou, Jun Lv, Yang Jin, Shirun Tang, Chuan Wen, and Cewu Lu. Implicitrdp: An end-to-end visual-force diffusion policy with structural slow-fast learning. *arXiv preprint arXiv:2512.10946*, 2025.

[2] Cheng-Yu Kuo, Hirofumi Shin, and Takamitsu Matsubara. Tracing energy flow: Learning tactile-based grasping force control to reduce slippage in dynamic object interaction. *IEEE Robotics and Automation Letters*, 11(2):1386–1393, 2026. doi: 10.1109/LRA.2025.3640979.

[3] Pei Lin, Yuzhe Huang, Wanlin Li, Chenxi Xiao, and Ziyuan Jiao. Dexmove: Learning tactile-guided non-prehensile manipulation with dexterous hands. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=dT3ZciXvNX.

[4] Han Xue, Jieji Ren, Wendi Chen, Gu Zhang, Yuan Fang, Guoying Gu, Huazhe Xu, and Cewu Lu. Reactive diffusion policy: Slow-fast visual-tactile policy learning for contact-rich manipulation. In *Proceedings of Robotics: Science and Systems (RSS)*, 2025.

[5] Wen Yang and Wanxin Jin. Contactsdf: Signed distance functions as multi-contact models for dexterous manipulation. *IEEE Robotics and Automation Letters*, 10(5): 4212–4219, 2025. doi: 10.1109/LRA.2025.3548503.

[6] Yi Zhang, Daniel Larby, Fumiya Iida, and Fulvio Forni. Virtual model control for compliant reaching under uncertainties. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 795–801, 2024. doi: 10.1109/IROS58592.2024.10801592.