Appendix

Anonymous Author(s)
Affiliation
Address
email

1 Robotic Platform

The rotating pressures are differentially pressurized, with both R_1 and R_2 being actuated to 35 psi or 241.32 kPa in the home position. A single variable R can be used to control both R_1 and R_2. Let $R = R_1 - R_2$, then if $R > 0$, $R_1 = 35$ and $R_2 = 35 - R$, otherwise $R_1 = 35 - R$ and $R_2 = 35$. Therefore, even though there are 6 actuations, the controller developed only needs 5 change in actuation.

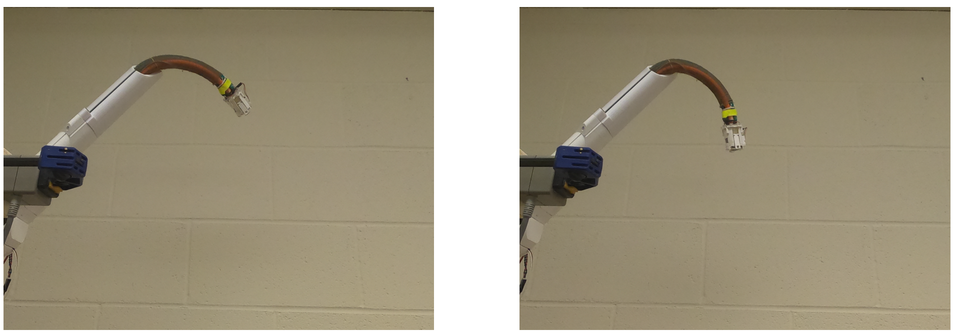

Figure 1: An example of rotating pressures differentially (left) vs non-differentially pressurized (right) where, $\theta 1 = -45, \theta_2 = 90, L = 10, B = 7, R = 0$. Differential pressurization increases the manipulator's stiffness, increases its load bearing capacity, and also dampens oscillations during movement.

2 Forward Model Data Collection

The data collection method for the full workspace is given in Algorithm 1.

Algorithm 1 Data Collection Method for Full Workspace

1: Set bounds on rigid actuations, θ_1, θ_2, pressures, B, R_1, R_2, and extrusion length, L
2: Set bounds on actuation changes $\Delta\theta_1, \Delta\theta_2, \Delta L, \Delta B, \Delta R$
3: Extrude to the lower bound, L_0
4: Randomly initialize pressures to within the bounds set in (1)
5: **while** time horizon is not reached **do**
6:     Sample uniformly within the bounds defined in (2)
7:     Increment the system by the control input sampled
8:     Collect end effector position and orientation
9:     Wait 0.25 seconds (4 Hz collection rate)
10: **end while**

A sample collected using Algorithm 1 is shown below, where, $s(t)$, $s(t-1)$ denote the poses, $u(t)$, $u(t-1)$ denote the actuations at current and previous time steps t and $t-1$ respectively. $\Delta u(t)$ denotes the change in actuation at time step t, $\Delta u(t-1))$ denotes the change in actuation at time step $t-1$, and the pose at next time step is denoted by $s(t+1)$.

$$\Big((s(t), s(t-1), u(t), u(t-1), \Delta u(t), \Delta u(t-1)), s(t+1)\Big)$$

Submitted to the 7th Conference on Robot Learning (CoRL 2023). Do not distribute.

3 Forward Model Training Graphs

Forward model training using different models is summarized in Figure 2. Validation errors are all very similar except for the MLP which is not surprising because recurrent networks capture temporal dynamics more easily. Interestingly, the LMU outperforms the LSTM on the robotic platform even with high validation errors.

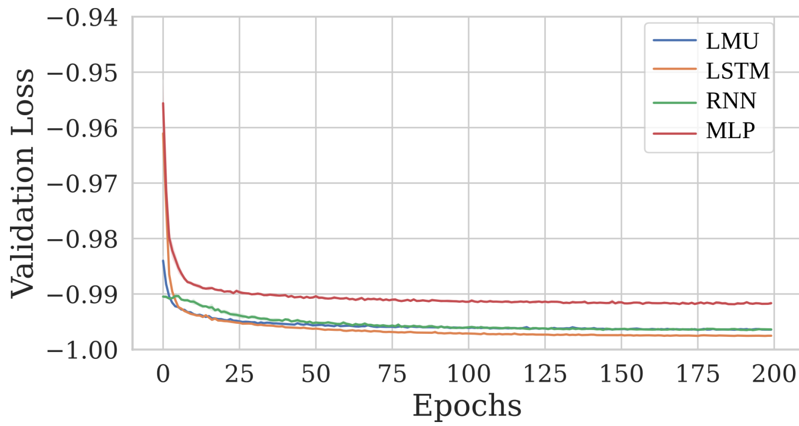

Figure 2: Forward Model Training Validation Errors for Various Architectures. Even though the LMU does not have the lowest validation error, it performs better on the withheld data as seen in Table 1.

4 Control Policy Network

Control policy training errors are plotted for each model tested (Figure 3). Surprisingly, the LMU-based policy outperforms the other models on the physical platform despite having higher validation errors compared to MLP and LSTM.

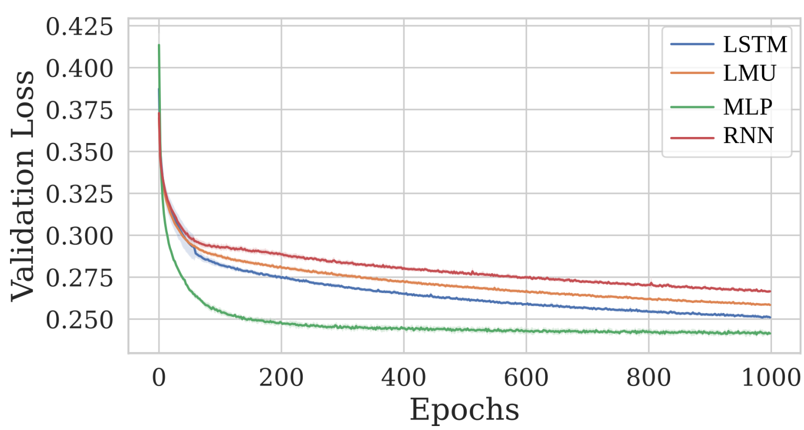

Figure 3: Policy training validation errors for various architectures. The resulting policies are evaluated on the physical system. Although validation loss does not converge, training until the validation loss converges resulted in worse performance on the physical platform.