

435 **A Details in  $\alpha$ -MDF**

436 This section provides a detailed overview of the previously mentioned  $\alpha$ -MDF modules, and de-  
 437 scribes differentiable Ensemble Kalman filters as the underlying DFs framework for  $\alpha$ -MDF.

438 **A.1 Model Initialization and Embedding Functions**

439 An auxiliary model  $\mathcal{A}$  is supplied in the filtering process to support training by starting the filter  
 440 via projecting the actual state  $\mathbf{x}_{t-N:t-1}$  from low-dimensional space to latent space. The model is  
 441 implemented using stochastic neural networks (SNNs) [42],

$$\mathbf{x}_{t-N:t-1}^i \sim \mathcal{A}(\mathbf{x}_{t-N:t-1}^i | \mathbf{x}_{t-N:t-1}), \forall i \in E, \quad (7)$$

442 where  $\mathbf{x}_{t-N:t-1}^i$  is one latent state, the latent state ensemble is obtained by sampling  $\mathcal{A}$  for  $E$  times.  
 443 During inference, we employ the trained sensor encoders' output, which is the latent representation  
 444 of RGB, depth, or proprioception, as the initial state to initiate the filtering process.

445 Regarding the prediction step of  $\alpha$ -MDF, we apply positional embedding layers (sinusoidal func-  
 446 tions) [32] in the transformer process model (Eq. 3) to generate  $\mathbf{e}_{t-N:t-1}$  as the embedding for  
 447 time-series data,  $\mathbf{e}_{t-N:t-1} = f_{\mathcal{L}}(\mathbf{X}_{t-N:t-1}) \in \mathbb{R}^{d_e \times (N-1)}$ . The positional embedding layer is uti-  
 448 lized to label the state by index it with time  $t$ . When activating the action  $\mathbf{a}_t$  in the process model,  
 449  $\mathbf{e}_{t-N:t-1}$  is passed through a type embedding layer that indexes  $\mathbf{e}_{t-N:t-1}$  and  $\mathbf{a}_t$  with 0 and 1,  
 450 and then fed to sinusoidal functions. Subsequently, the outputs obtained from the aforementioned  
 451 procedures serve as input to the transformer process model for further processing.

452 **A.2 Differentiable Ensemble Kalman Filter**

453 Unlike prior proposals for differentiable filters, such as dEKF [9] and DPF [24], Differentiable En-  
 454 semble Kalman Filter [6] leverages recent advancements in stochastic neural networks (SNNs) [42].  
 455 Specifically, we draw inspiration from the work in [43], which established a theoretical connection  
 456 between the Dropout training algorithm and Bayesian inference in deep Gaussian processes. As a  
 457 result, we can use stochastic forward passes to produce empirical samples from the predictive poste-  
 458 rior of a neural network trained with Dropout. Hence, for the purposes of filtering, we can implicitly  
 459 model the process noise by sampling state from a neural network trained on the transition dynamics,  
 460 i.e.,  $\mathbf{x}_t \sim f_{\theta}(\mathbf{x}_{t-1})$ . In contrast to previous approaches [24, 9], the transition network  $f_{\theta}(\cdot)$  models  
 461 the system dynamics, as well as the inherent noise model in a consistent fashion without imposing  
 462 diagonality.

463 **Prediction Step:** Similar to  $\alpha$ -MDF, we use an initial ensemble of  $E$  members to represent the  
 464 initial state distribution  $\mathbf{X}_0 = [\mathbf{x}_0^1, \dots, \mathbf{x}_0^E]$ ,  $E \in \mathbb{Z}^+$ . We leverage the stochastic forward passes  
 465 from a trained state transition model to update each ensemble member:

$$\mathbf{x}_{t|t-1}^i \sim f_{\theta}(\mathbf{x}_{t|t-1}^i | \mathbf{x}_{t-1|t-1}^i), \forall i \in E. \quad (8)$$

466 Matrix  $\mathbf{X}_{t|t-1} = [\mathbf{x}_{t|t-1}^1, \dots, \mathbf{x}_{t|t-1}^E]$  holds the updated ensemble members which are propagated  
 467 one step forward through the state space. Note that sampling from the transition model  $f_{\theta}(\cdot)$  (using  
 468 the SNN methodology described above) implicitly introduces a process noise.

469 **Update step:** Given the updated ensemble members  $\mathbf{X}_{t|t-1}$ , a nonlinear observation model  $h_{\psi}(\cdot)$  is  
 470 applied to transform the ensemble members from the state space to observation space. Following our  
 471 main rationale, the observation model is realized via a neural network with weights  $\psi$ . Accordingly,  
 472 the update equations become:

$$\mathbf{H}_t \mathbf{A}_t = \mathbf{H}_t \mathbf{X}_t - \left[ \frac{1}{E} \sum_{i=1}^E h_{\psi}(\mathbf{x}_t^i), \dots, \frac{1}{E} \sum_{i=1}^E h_{\psi}(\mathbf{x}_t^i) \right], \quad (9) \quad \tilde{\mathbf{y}}_t^i \sim s(\tilde{\mathbf{y}}_t^i | \mathbf{y}_t), \forall i \in E. \quad (10)$$

474  $\mathbf{H}_t \mathbf{X}_t$  is the predicted observation, and  $\mathbf{H}_t \mathbf{A}_t$  is the sample mean of the predicted observation at  $t$ .  
 475 Traditional Ensemble Kalman Filter treats observations as random variables. Hence, the ensemble

476 can incorporate a measurement perturbed by a small stochastic noise to reflect the error covariance  
 477 of the best state estimate [6]. In differentiable Ensemble Kalman Filter, we incorporate a Bayesian  
 478 sensor encoder  $s(\cdot)$ . Sensor encoder serves to learn projections from observation space to latent  
 479 space as in Eq. 10, where  $\mathbf{y}_t$  represents the noisy sensor observation. Sampling from sensor encoder  
 480 yields latent observations  $\tilde{\mathbf{Y}}_t = [\tilde{\mathbf{y}}_t^1, \dots, \tilde{\mathbf{y}}_t^E]$ . The KF update step can then be continued by using  
 481 the learned observation and predicted observation:

$$\mathbf{K}_t = \frac{1}{E-1} \mathbf{A}_t (\mathbf{H}_t \mathbf{A}_t)^T \left( \frac{1}{E-1} (\mathbf{H}_t \mathbf{A}_t) (\mathbf{H}_t \mathbf{A}_t)^T + \mathbf{R} \right)^{-1}. \quad (11)$$

482 The measurement noise model  $\mathbf{R}$  is implemented using a multilayer perceptron (MLP), similar  
 483 to the implementation in [9]. The MLP takes a learned observation  $\tilde{\mathbf{Y}}_t$  at time  $t$  and produces  
 484 a noise covariance matrix. The final estimate of the ensemble  $\hat{\mathbf{X}}_t$  is obtained by performing the  
 485 measurement update step, given by:

$$\hat{\mathbf{X}}_t = \mathbf{X}_t + \mathbf{K}_t (\tilde{\mathbf{Y}}_t - \mathbf{H}_t \mathbf{X}_t). \quad (12)$$

486 In inference, the ensemble mean  $\bar{\mathbf{x}}_{t|t} = \frac{1}{E} \sum_{i=1}^E \mathbf{x}_{t|t}^i$  is used as the updated state.

### 487 A.3 Baselines

488 In our study, we examine two categories of baselines: (a) DFs baselines, which consist of existing  
 489 methods such as those proposed in [9, 24, 22], and (b) sensor fusion strategies, as proposed in [18].

Table 4: Dimensions pertinent to each of the robot state estimation tasks.

Method	Visual Odometry		UR5 Manipulation		Soft Robot		
	State	Observation	State	Observation	State	Observation	Action
dEKF [9]	5	2	10	10	7	7	40
DPF [24]	5	2	10	10	7	7	40
dPF-M-lrn [9]	5	2	10	10	7	7	40
Feature Fusion [18]	-	-	10/13	10/13	7	7	40
Unimodal [18]	-	-	10/13	10/13	7	7	40
Crossmodal [18]	-	-	10/13	10/13	7	7	40
$\alpha$ -MDF	256	256	256	256	256	256	40

490 **Dimensionality:** Table 4 presents the dimensions for the state, observations, and actions utilized  
 491 for each of the tasks. To ensure consistency, we opt for a dimension of 256 for  $\alpha$ -MDF in all tasks,  
 492 thus, enabling filtering over high-dimensional spaces. Unlike the baseline methods, which use low-  
 493 dimensional state definitions, we filter over higher dimension spaces with  $\alpha$ -MDF.

494 **Differentiable Filters:** To maintain consistency in the comparison of results against the DFs base-  
 495 lines, we train  $\alpha$ -MDF with a single modality. The baselines in this category include the different-  
 496 iable Extended Kalman filter (dEKF) [9], differentiable particle filter (DPF) [24], and the modified  
 497 differentiable particle filter (dPF-M-lrn) [9], which uses learned process and process noise models.  
 498 For dEKF, the Jacobian matrix in the prediction step can either be learned end-to-end or supplied if  
 499 the motion model is known. DPF employs 100 particles for both training and testing and also in-  
 500 corporates an observation likelihood estimation model  $l$ . This module takes in an image embedding  
 501 and produces a likelihood that updates each particle’s weight. Unlike DPF, dPF-M-lrn implements  
 502 a learnable process noise model. It also adopts a Gaussian Mixture Model for calculating the like-  
 503 lihood for all particles. It is worth noting that all the baseline methods perform Kalman filtering  
 504 on low-dimensional actual state space, whereas  $\alpha$ -MDF executes the filtering process in the latent  
 505 space.

506 **Sensor Fusion:** Regarding sensor fusion baselines, we use three strategies discussed in [18], namely,  
 507 Feature Fusion, Unimodal Fusion, and Crossmodal Fusion. The Feature Fusion strategy aims to  
 508 process each modality individually and subsequently merge the modalities to generate a multimodal  
 509 feature set using neural networks, which is then used for state estimation. The Unimodal Fusion

510 treats each modality  $\mathcal{N} \sim (\boldsymbol{\mu}_t^{M_1}, \boldsymbol{\Sigma}_t^{M_1})$  and  $\mathcal{N} \sim (\boldsymbol{\mu}_t^{M_2}, \boldsymbol{\Sigma}_t^{M_2})$  as distributions and fuse two uni-  
 511 modal distribution as one normally distributed multimodal distribution  $\mathcal{N} \sim (\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$ :

$$\boldsymbol{\mu}_t = \frac{(\boldsymbol{\Sigma}_t^{M_1})^{-1} \boldsymbol{\mu}_t^{M_1} + (\boldsymbol{\Sigma}_t^{M_2})^{-1} \boldsymbol{\mu}_t^{M_2}}{(\boldsymbol{\Sigma}_t^{M_1})^{-1} + (\boldsymbol{\Sigma}_t^{M_2})^{-1}}, \quad \boldsymbol{\Sigma}_t = ((\boldsymbol{\Sigma}_t^{M_1})^{-1} + (\boldsymbol{\Sigma}_t^{M_2})^{-1})^{-1}, \quad (13)$$

512 where the associative property can be used for fusing more than two modalities. For Crossmodal  
 513 Fusion, information from one modality can be used to determine the uncertainty of the other ones,  
 514 two coefficients are proposed as  $\boldsymbol{\beta}_t^{M_1}$  and  $\boldsymbol{\beta}_t^{M_2}$ , where each coefficient has the same dimension of  
 515 the state, the fused distribution is:

$$\boldsymbol{\mu}_t = \frac{\boldsymbol{\beta}_t^{M_1} \circ \boldsymbol{\mu}_t^{M_1} + \boldsymbol{\beta}_t^{M_2} \circ \boldsymbol{\mu}_t^{M_2}}{\boldsymbol{\beta}_t^{M_1} + \boldsymbol{\beta}_t^{M_2}}, \quad \boldsymbol{\Sigma}_t = \frac{\boldsymbol{B}_t^{M_1} \circ \boldsymbol{\Sigma}_t^{M_1} + \boldsymbol{B}_t^{M_2} \circ \boldsymbol{\Sigma}_t^{M_2}}{\boldsymbol{B}_t^{M_1} + \boldsymbol{B}_t^{M_2}}, \quad (14)$$

516 where  $\boldsymbol{B}_t^M = (\boldsymbol{\beta}_t^M)^T \boldsymbol{\beta}_t^M$ . As mentioned in [18], each sensor encoder was independently trained  
 517 and subsequently used for end-to-end training with DFs. We adopt a similar approach, but with a  
 518 differentiable Ensemble Kalman Filter backbone in place instead. The resampling procedure from  
 519 the fused distribution in this scenario is achieved by using the reparameterization trick [44].

## 520 B Additional Experiments

521 This section presents supplementary experimental results for each task. For (1) Visual Odometry  
 522 Tasks, we offer full detailed experiments; however, for (2) Multimodal Manipulation Tasks and (3)  
 523 Soft Robot Modeling Tasks, we concentrate mainly on ablation studies.

### 524 B.1 Visual Odometry Tasks

525 In this experiment, we investigate the performance of  $\alpha$ -MDF on the popular KITTI Visual Odom-  
 526 etry dataset [34]. We only consider RGB images as the input modality in order to make a fair com-  
 527 parison with the baselines [9, 24, 22]. Following the same evaluation procedure as our baselines, we  
 528 define the actual state of the moving vehicle as a 5-dimensional vector  $\boldsymbol{x} = [x, y, \theta, v, \dot{\theta}]^T$ , including  
 529 the position and orientation of the vehicle, and the linear and angular velocity w.r.t. the current head-  
 530 ing direction  $\theta$ . The raw observation  $\boldsymbol{y}$  corresponds to the RGB camera image of the current frame  
 531 and a difference image between the current frame and the previous frame, where  $\boldsymbol{y} \in \mathbb{R}^{150 \times 50 \times 6}$  as  
 532 shown in Fig. 8. The learned observation  $\hat{\boldsymbol{y}}$  is defined as  $\hat{\boldsymbol{y}} = [v, \dot{\theta}]^T$ , since only the relative changes  
 533 of position and orientation can be captured between two frames. We use the latent state  $\boldsymbol{x} \in \mathbb{R}^{256}$   
 534 for  $\alpha$ -MDF.

535 **Data:** The KITTI Visual Odometry dataset in-  
 536 cludes 11 trajectories capturing the ground truth  
 537 pose (translation and rotation matrices) of a ve-  
 538 hicle navigating urban areas at a data collec-  
 539 tion rate of approximately 10Hz. To facilitate  
 540 the learning process, we standardize the data by  
 541 normalizing each dimension to have a mean of  
 542 0 and a standard deviation of 1 during training.  
 543 To process the provided pose data, we convert  
 544 them to quaternions to capture the minimal changes between consecutive quaternion pairs. Subse-  
 545 quently, the results are converted back to radians to represent the angular velocity  $\dot{\theta}$ . This conversion  
 546 ensures that the angular velocity remains minimal and falls within the range of  $[-\pi, \pi]$ .

547 **Results:** The performance of state estimation is evaluated using an 11-fold cross-validation,  
 548 whereby 1 trajectory is withheld at each time. The standard KITTI benchmark metrics, namely  
 549 the translational error (m/m) and rotational error (deg/m), are reported in Table 5. The error met-  
 550 rics are computed from the test trajectory over all subsequences of 100 timesteps, as well as all  
 551 subsequences of 100, 200, 400, and 800 timesteps. Figure 9 presents the performance of  $\alpha$ -MDF  
 552 and other differentiable filtering techniques. It is important to note that incorporating domain- and

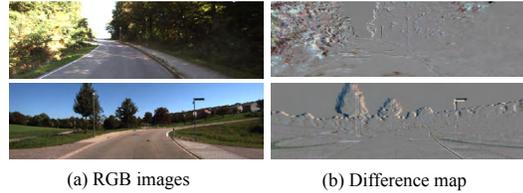


Figure 8: KITTI visual inputs.

Table 5: Result evaluations on KITTI Visual Odometry task measured in m/m and deg/m denote the translational error and the rotational error.

Method	Test 100		Test 100/200/400/800	
	m/m	deg/m	m/m	deg/m
dEKF [9]	0.2646±0.004	0.1386±0.002	0.3159±0.002	0.0923±0.005
DPF [24]	0.1344±0.002	0.1203±0.007	0.2255±0.001	0.0716±0.004
dPF-M-lrn [9]	0.1720±0.010	0.0974±0.009	0.1848±0.004	0.0611±0.003
<b><math>\alpha</math>-MDF</b>	<b>0.0718±0.001</b>	<b>0.0954±0.001</b>	<b>0.0379±0.002</b>	<b>0.0328±0.001</b>

Means±standard errors.

553 data-specific information, such as using stereo images [45], integrating LiDAR [46, 47], or applying  
 554 SLAM and loop-closure related assumptions [45, 48], can yield lower error metrics. However, to  
 555 ensure fair and comparable evaluations, we utilize the most commonly used setup when comparing  
 556 filtering techniques in a task-agnostic fashion (as performed in [9, 24, 22]).

557 Table 5 presents the outcomes of our proposed method in comparison with the existing state-of-the-art DFs, namely dEKF, DPF, and dPF-M-lrn. In order to provide a fair comparison, we do not include unstructured LSTM models as baselines since prior works [22, 9] have shown that they do not achieve comparable results. The pre-trained sensor encoder with the same visual inputs is used and integrated into all the DF frameworks evaluated. In this experiment, the motion model of the vehicle is known, and the only unknown part of the state is the velocities. In light of the above, we adopt a learnable process model to update state variables alongside an established motion model to update the  $(x, y, \theta)$  variables. While the computed Jacobian matrix is supplied in training and testing for dEKF, our  $\alpha$ -MDF demonstrates significant improvements compared to dEKF, DPF, and dPF-M-lrn. Specifically, we observed a reduction in the translational error of approximately 88%, 83%, and 79% for Test 100/200/400/800. The results also reflect a considerable reduction in rotational error of approximately 64%, 54%, and 46% as compared to each of the baselines. Our analysis of  $\alpha$ -MDF reveals that conducting filtering on high-dimensional observations in the latent space yields better results than conducting filtering on the actual state space.

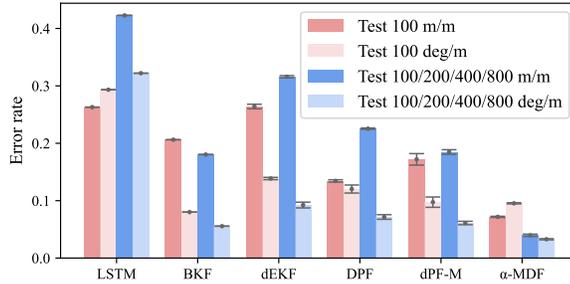


Figure 9: Visual Odometry results with different differentiable filters: the error rate for LSTM and BKF are reported from [22], dEKF, DPF, and dPF-M are reproduced.

## 579 B.2 Multimodal Manipulation Tasks

580 **Task Setup:** For  $\alpha$ -MDF, we define the latent state  $\mathbf{x} \in \mathbb{R}^{256}$  for all the manipulation tasks. The  
 581 actual state of the UR5 robot is described by  $\mathbf{x}_R$ , which consists of the seven joint angles ( $J_1$ - $J_7$ )  
 582 and the Cartesian coordinates  $(x, y, z)$  of the robot’s end-effector. This Cartesian coordinate sys-  
 583 tem is centered at the manipulation platform’s origin point  $(0, 0, 0)$ . On the other hand, the state of  
 584 the object being manipulated is represented by  $\mathbf{x}_O$ , which only includes the Cartesian coordinates  
 585  $(x, y, z)$  of the object. The input modalities for each of the three tasks differ. In task (1), input is  
 586 given through three modalities:  $\mathbf{y}^1$ ,  $\mathbf{y}^2$ , and  $\mathbf{y}^3$ . The first modality  $\mathbf{y}^1 \in \mathbb{R}^{224 \times 224 \times 3}$  is a camera im-  
 587 age captured from a frontal angle. The second modality  $\mathbf{y}^2 \in \mathbb{R}^{224 \times 224 \times 1}$  depicts depth maps from  
 588 the same camera view. Lastly,  $\mathbf{y}^3$  is a proprioceptive input source with dimensions  $\mathbb{R}^7$ , representing  
 589 the joint angles’ values. In this task, the proprioceptive input specifically refers to the joint angles  
 590 as the source. In task (2), input is given by only two modalities:  $\mathbf{y}^1$  and  $\mathbf{y}^3$ , but from a real-world  
 591 perspective. In task (3), input is received from four modalities:  $\mathbf{y}^1$ ,  $\mathbf{y}^2$ ,  $\mathbf{y}^3$ , and  $\mathbf{y}^4$ .  $\mathbf{y}^4$  contains  
 592 the Force/torque (F/T) sensor readings from the robot gripper, where  $\mathbf{y}^4 \in \mathbb{R}^6$ , while the first two  
 593 modalities are identical to task (1).

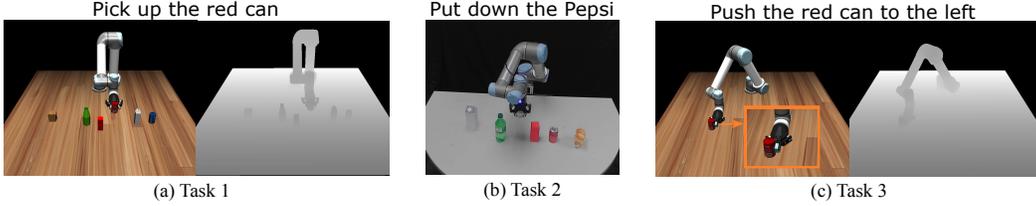


Figure 10: The multimodal manipulation experiment involves the following subtasks: (a) Task 1 utilizing RGB, depth, and joint modalities, (b) Task 2 utilizing only RGB and joint modality, and (c) Task 3 utilizing RGB, depth, joint, and Force/torque (F/T) sensor modalities. The F/T sensor is mounted on the grabber, as depicted by the orange box.

Table 6: Ablation study on UR5 manipulation task with different combination of the modalities.

	RGB	Depth	Joint	F/T	Joint (deg)	EE (cm)	Obj (cm)
Task (1)	✓				2.78±0.09	1.06±0.01	-
		✓			3.65±0.10	1.38±0.05	-
			✓		9.53±0.20	3.22±0.14	-
			✓	✓	2.39±0.11	1.01±0.02	-
	✓			✓	2.69±0.01	1.09±0.03	-
	✓	✓			<b>1.91±0.08</b>	<b>0.64±0.03</b>	-
Task (2)	✓				7.49±0.06	3.81±0.17	-
			✓		5.47±0.08	3.32±0.04	-
	✓		✓		<b>5.24±0.04</b>	<b>3.04±0.01</b>	-
Task (3)		✓	✓	✓	2.93±0.01	2.26±0.02	3.26±0.01
	✓		✓	✓	3.16±0.20	2.34±0.04	3.66±0.30
	✓	✓		✓	1.42±0.08	0.93±0.01	<b>1.47±0.02</b>
	✓	✓	✓		<b>1.37±0.02</b>	0.94±0.01	1.78±0.06
	✓	✓	✓	1.41±0.04	<b>0.90±0.01</b>	1.65±0.01	

Means±standard errors.

594 **Data:** Data collection is conducted for both simulation with MuJoCo [49] and real-world scenarios.  
 595 We record the UR5 robot operating on a random object by performing one of “pick”, “push”, and  
 596 “put down” actions. We collect 2,000 demonstrations in simulation for task (1), and 100 on the  
 597 real robot for task (2), with changing the location of each object for each demonstration. For task  
 598 (3), we collect 2,000 demonstrations in simulation with adding the tactile sensors. We use ABR  
 599 control and robosuite [50] in addition to MuJoCo to ensure rigorous dynamics in the simulator. Each  
 600 demonstration sequence has a length of approximately 350 steps with a timestep of 0.08 seconds.  
 601 An 80/20 data split is utilized for training and testing each task. It should be noted that in all tasks,  
 602 we normalize the joint modality  $\mathbf{y}^3$  and apply Gaussian noise to each joint angle, drawn from the  
 603 distribution  $\mathcal{N} \sim (0, \sigma^2 \mathbf{I})$  where  $\sigma^2 = 0.1$ . We collect the F/T sensor readings directly from  
 604 MuJoCo’s native touch sensor. Moreover, the depth maps obtained from MuJoCo are with no noise  
 605 therefore can be regarded as high-fidelity data.

606 **Ablation Study:** In addition to the findings presented in Section 4.2, we perform a comprehensive  
 607 ablation analysis for each manipulation task to address the question, “How does the use of multiple  
 608 modalities compare to a subset of modalities for state estimation with differentiable filters?”. Table 6  
 609 displays the outcome for each task with various number of modalities using MAE metric. The  
 610 highest margin of error is indicated by the red shading, while the complete modality is labeled by  
 611 green shading for each task. Interestingly, even though using all modalities can generate comparable  
 612 results, in certain tasks, utilizing all modalities does not necessarily guarantee superior performance  
 613 compared to utilizing a subset of modalities. Through our experiments in Task (1), it becomes  
 614 apparent that the optimal performance is achieved by utilizing the subset of modalities  $[\mathbf{y}^1, \mathbf{y}^2]$ ,  
 615 which yields an improvement of joint angles ( $2.19^\circ \rightarrow 1.91^\circ$ ). In Task (3), we observe that diverse  
 616 subsets of modalities lead to superior state estimation results for joint angles, EE, and the object  
 617 locations respectively. Analysis of Table 6 indicates an important role played by the depth map  
 618  $\mathbf{y}^2$  when considering all observations. This suggests that  $\mathbf{y}^2$  is treated as high-fidelity data during  
 619 training, thereby contributing the most towards the final results.

620 Henceforth, we conduct an additional  
 621 ablation analysis to ascertain whether  
 622 or not the use of a combination of  
 623 high-fidelity and low-fidelity sensor  
 624 inputs offers a potential benefit. As  
 625 noted during data collection, the pro-  
 626 prioceptive input  $\mathbf{y}^3$  comprising joint  
 627 angles is obtained via adding Gaus-  
 628 sian noise and is therefore considered  
 629 a low-fidelity input. Figure 11 illus-  
 630 trates the scenario of using  $\mathbf{y}^3$  and not  
 631 using  $\mathbf{y}^3$  while applying distinct lev-  
 632 els of Gaussian blur in the image and  
 633 depth space. Notably, without em-  
 634 ploying  $\mathbf{y}^3$ , the state estimation perfor-  
 635 mance deteriorates as the level of blur increases. On the  
 636 other hand,  $\mathbf{y}^3$  - despite being classified as a low-fidelity modality - contributes to the final state  
 637 estimation. In particular, at the highest level of blur, incorporating  $\mathbf{y}^3$  yields a 29% improvement in  
 joint angle estimation and a 17% improvement for end-effector locations.

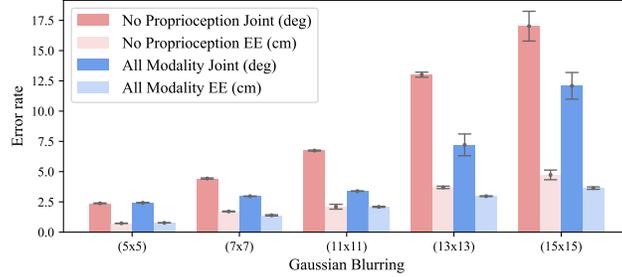


Figure 11: State estimation results are shown after introducing diverse levels of noise to  $[\mathbf{y}^1, \mathbf{y}^2]$ . The red group depicts results using  $[\mathbf{y}^1, \mathbf{y}^2]$  modality, while the blue group represents results using  $[\mathbf{y}^1, \mathbf{y}^2, \mathbf{y}^3]$  modality.

### 638 B.3 Soft Robot Modeling Tasks

639 This section presents a comprehensive analysis of the tensegrity robot structure, the bending motion  
 640 mechanism, and pertinent sensory information, followed by a description of additional experimental  
 641 outcomes related to this task.

642 **Preliminaries:** Our research  
 643 utilizes a tensegrity robot arm  
 644 (developed in [39]) that follows  
 645 a strict tensegrity structure fea-  
 646 turing struts, cables (including  
 647 spring-loaded and actuated cab-  
 648 les), and five layers of arm-like  
 649 tensegrity structures, which pro-  
 650 duce continuous bending pos-  
 651 tures when exposed to exter-  
 652 nal forces. The longitudinal  
 653 length is maintained by stiff cab-  
 654 les, while the bending direction  
 655 is solely determined by external  
 656 forces. We determine the robot’s  
 657 kinematics through data from Inertial Measurement Units (IMUs), optical motion capture (MoCap),  
 658 and proportional pressure control valves, with each of the five struts in each layer featuring an IMU.  
 659 We also record the video by placing a camera in front of the robot while collecting all sensory data.  
 660 A soft robot’s state at  $t$  is a 7-dimensional vector  $\mathbf{x}_t = [x, y, z, \mathbf{q}_x, \mathbf{q}_y, \mathbf{q}_z, \mathbf{q}_w]^T$ , indicating its po-  
 661 sition and orientation relative to the base frame (layer 1’s bottom).  $\mathbf{q}$  represents the robot’s posture.  
 662 The system’s action is the pressure vector of its 40 pneumatic cylinder actuators ( $\mathbf{a}_t \in \mathbb{R}^{40}$ ). Its raw  
 663 observation is comprised of 5 IMU readings ( $\mathbf{y}_t^3 \in \mathbb{R}^{30}$ ), with each IMU measuring a 6-dimensional  
 664 vector of accelerations and angular velocities relative to its location. Fig. 12 illustrates the locations  
 665 of the IMUs on the struts (blue cubes) in each layer.

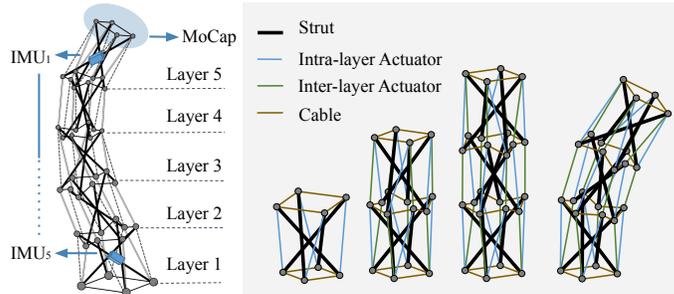


Figure 12: The tensegrity robot features 5 flexible layers, each a tensegrity module with struts, cables, and actuators. Its sensory data includes IMUs, MoCap, and pressure vector readings from pneumatic cylinders.

666 **Data:** The complete set of modalities comprises  $[\mathbf{y}^1, \mathbf{y}^2, \mathbf{y}^3]$ , where  $\mathbf{y}^1 \in \mathbb{R}^{224 \times 224 \times 3}$  represents  
 667 RGB images,  $\mathbf{y}^2 \in \mathbb{R}^{224 \times 224}$  is synthetic depth maps which we generate from DPT repo [40]  
 668 utilizing “Intel/dpt-large”, and  $\mathbf{y}^3 \in \mathbb{R}^{30}$  is proprioceptive inputs (IMUs). The dataset is generated  
 669 by performing optical motion capture on the real tensegrity robot hand tip while randomly supplying  
 670 desired pressure vectors to the pneumatic cylinder actuators. The action  $\mathbf{a}_t \in \mathbb{R}^{40}$ , 5 IMU readings  
 671  $\mathbf{y}_t^3 \in \mathbb{R}^{30}$ , and a 7-dimensional state  $\mathbf{x}_t$  are recorded, with 40-dimensional pressure vectors being

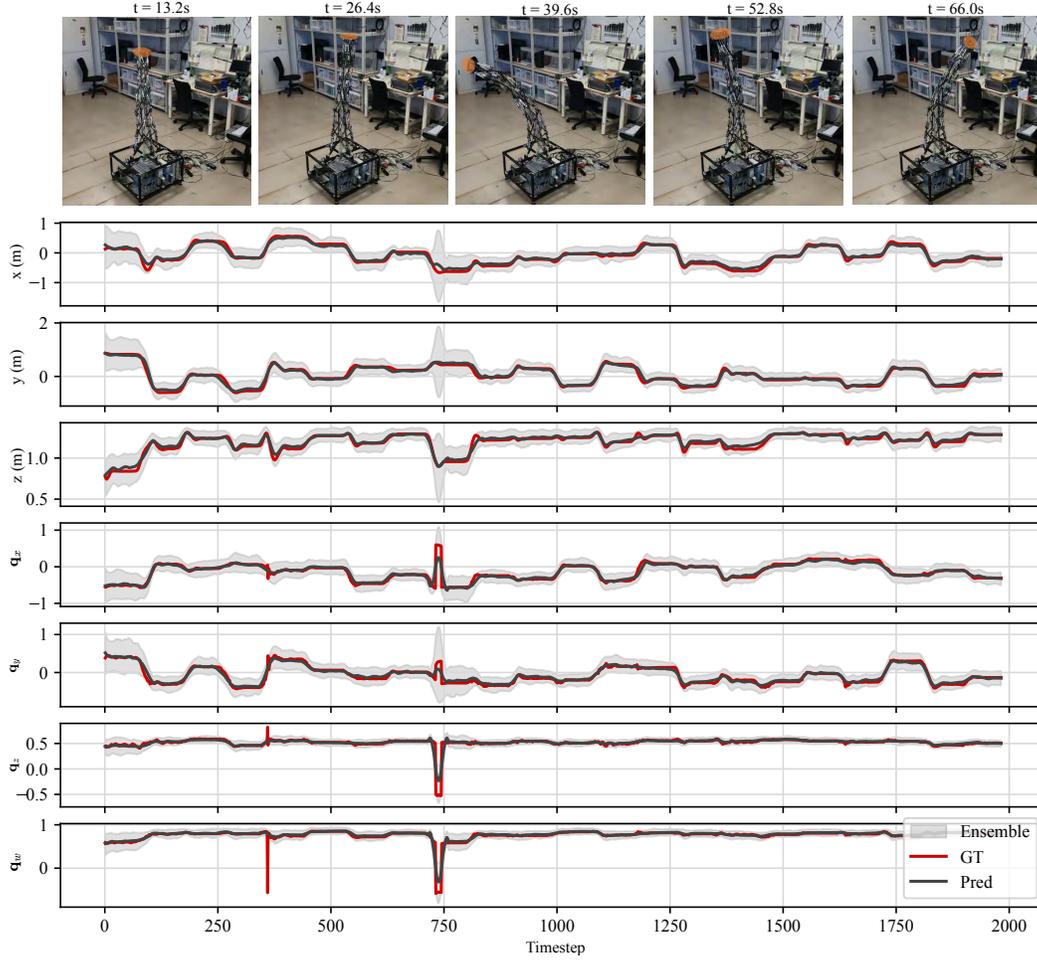


Figure 13: Predicted end-effector (EE) positions and quaternion vectors  $\mathbf{q}$  in the soft robot modeling task. The **top** row displays the actual robot posture at the corresponding time, with the orange circle indicating the EE positions, which are not included in the RGB modality input.

672 used as a control signal. A total of 12,000 trials of robot motion are collected, with each trial  
 673 involving moving the robot from its current equilibrium posture to the next equilibrium posture by  
 674 applying the new desired pressure. All data are collected via a ROS2 network with a sampling  
 675 frequency of 30Hz and are synchronized using the “message\_filters” package.

676 **Ablation Study:** In addition to the re-  
 677 sults presented in Section 4.3, we evalu-  
 678 ate various combinations of modalities  
 679 to determine whether an optimal sub-  
 680 set of modalities can be identified to at-  
 681 tain comparable outcomes without using  
 682 all modalities during the filtering opera-  
 683 tion. As demonstrated in Table7, uti-  
 684 lizing only one modality fails to achieve  
 685 comparable results, with the highest accu-  
 686 racy (2.07cm) exclusively from employ-  
 687 ing  $y^1$  (RGB). The lowest error in  
 688 posture estimation for the robot is obtained by leveraging  $[y^1, y^2]$ , showing slight improvement  
 689 (0.10 $\rightarrow$ 0.12) over leveraging the full modalities  $[y^1, y^2, y^3]$ . However, the lowest MAE error for  
 690 the EE position persists even when all modalities are employed. Interestingly, using solely  $y^3$  re-

Table 7: Ablation study on Tensegrity robot.

RGB	Depth	IMUs	EE (cm)	$\mathbf{q}(10^1)$
✓			2.07 $\pm$ 0.03	0.31 $\pm$ 0.08
	✓		2.77 $\pm$ 0.01	0.19 $\pm$ 0.05
		✓	8.99 $\pm$ 0.02	0.79 $\pm$ 0.03
	✓	✓	2.08 $\pm$ 0.03	0.14 $\pm$ 0.02
✓		✓	1.73 $\pm$ 0.05	0.12 $\pm$ 0.02
✓	✓		1.74 $\pm$ 0.06	<b>0.10<math>\pm</math>0.02</b>
✓	✓	✓	<b>1.67<math>\pm</math>0.09</b>	0.12 $\pm$ 0.01

Means $\pm$ standard errors.

691 sults in the highest state estimation error, which aligns with the lowest attention value visualized in  
 692 Fig 14. As depicted in Fig. 14, it is evident that  $\alpha$ -MDF prioritizes  $\mathbf{y}^1$  over other modalities. Inter-  
 693 estingly, the attention values change upon turning off certain modalities while the system remains  
 694 stable and functional.

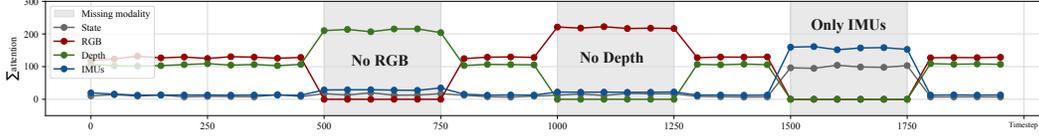


Figure 14: The corresponding accumulated attention values for each modality during testing. The gray areas show certain modalities are selected or not selected.

## 695 C Training Details

696 Table 8 provides an exhaustive enumeration of all learnable modules utilized in  $\alpha$ -MDF,  
 697 which includes three primary components: the state transition model  $f_{\theta}$ , the sensor encoders  
 698  $[s^1(\cdot), s^2(\cdot), \dots, s^M(\cdot)]$ , and the attention gain (AG) module. We adopt self-attention layers with  
 699 dimension 256 and 8 heads, denoted as “Self Attn”, in the state transition model. The cross-attention  
 700 layers, denoted as “Cross Attn”, is with dimension 32 and 4 heads in the AG module. The sensor  
 701 encoders utilized in our approach and all baseline models are identical, with  $s^1$  acting on image-  
 702 like modalities, utilizing ResNet18 [51] for learning high-dimensional observation representations,  
 703 while  $s^2$  pertains to low-dimensional modalities such as joint angles. The auxiliary model  $\mathcal{A}$  and the  
 704 decoder  $\mathcal{D}$  shares a similar structure to  $s^2$ , but with different number of neurons. Note that  $x$  is the  
 705 dimension of the actual state.

Table 8:  $\alpha$ -MDF’s learnable sub-modules.

$f_{\theta}$ :	$3 \times \text{SNN}(256, \text{ReLu}), \text{Positional Embedding}, 3 \times \text{Self Attn}(256, 8), 2 \times \text{SNN}(256, \text{ReLu}), 1 \times \text{SNN}(d_x, -)$
$s^1$ :	$1 \times \text{ResNet18}(h, w, ch), 2 \times \text{fc}(2048, \text{ReLu}), 1 \times \text{SNN}(512, \text{ReLu}), 1 \times \text{SNN}(d_x, -)$
$s^2$ :	$1 \times \text{SNN}(128, \text{ReLu}), 1 \times \text{SNN}(256, \text{ReLu}), 1 \times \text{SNN}(512, \text{ReLu}), 1 \times \text{SNN}(d_x, -)$
AG:	$\text{Positional Embedding}, 1 \times \text{Cross Attn}(32, 4, \text{mask})$
$\mathcal{A}$ :	$1 \times \text{SNN}(128, \text{ReLu}), 1 \times \text{SNN}(256, \text{ReLu}), 1 \times \text{SNN}(512, \text{ReLu}), 1 \times \text{SNN}(1024, \text{ReLu}), 1 \times \text{SNN}(d_x, -)$
$\mathcal{D}$ :	$1 \times \text{fc}(256, \text{ReLu}), 1 \times \text{SNN}(128, \text{ReLu}), 1 \times \text{SNN}(32, \text{ReLu}), 1 \times \text{SNN}(x, -)$
fc: Fully Connected, SNN: Stochastic Neural network.	

706 During  $\alpha$ -MDF training, we employ the curriculum outlined in Algorithm 1. Note that some tasks  
 707 may require pre-training the sensor encoders before performing end-to-end training the entire frame-  
 708 work. For each task, we train  $\alpha$ -MDF model with utilizing batch size of 64 on a single NVIDIA  
 709 A100 GPU for roughly 48 hours. For all the tasks, we use the Adamw [52] optimizer with a learning  
 710 rate of  $1e-4$ .

**Algorithm 1** Condition in Latent Space: training algorithm return the weights  $\omega$

---

**Input:**  $\alpha$ -MDF, dataloader  $(\{\mathbf{x}_t\}_{t-N}^{t+1}, \{\mathbf{y}_t^m\}_{m=1}^M, \{\mathbf{y}_{t+1}^m\}_{m=1}^M, \{\mathbf{a}_t\}_{t-1}^{t+1})$   
**Output:** weights  $\omega$   
**while** not converged **do**  
  Call dataloader with a random timestep  $t$ .  
  **for** timestep  $t \leftarrow t$  to  $t + 1$  **do**  
     $e_1 \leftarrow \sum_{m=1}^M \|\mathcal{D}(s^m(\mathbf{y}_t^m)) - \mathbf{x}_t\|_2^2$  according to Eq. 6  
     $e_2 \leftarrow \mathcal{L}_{f_{\theta}}(\mathbf{X}_t) + \mathcal{L}_{e_{2e}}(\hat{\mathbf{X}}_t)$  according to Eq. 6  
     $e_t \leftarrow e_1 + e_2$   
  **end for**  
   $\omega \leftarrow \text{Train}(\alpha\text{-MDF}, e_t + e_{t+1})$   
**end while**  
**return**  $\omega$

---