

525 A Appendix

526 A.1 Tasks

527 Here we give more details about the tasks, including the performance functions, teacher dataset, and
528 sample images. Fig. 6 shows images all of simulation environments used for SOTA comparisons and
529 generalizability, with one end-effector. In each environment, the end-effectors are pickers (white
530 spheres).

- 531 1. CLOTH FOLD: Fold a square cloth in half, along a specified line. The performance metric
532 is the distance of the cloth particles left of the folding line, to those on the right of the
533 folding line. A fully folded cloth should have these two halves virtually overlap. Teacher
534 demonstrations are from an agent with two pickers (*i.e.*, $\mathcal{D}_{Teacher} = \mathcal{D}_{Demo}^{2p}$); we solve
535 the task on a student agent with one picker. Task variations are in cloth rotation.
536
- 537 2. DRY CLOTH: Pick up a square cloth from the ground and hang it on a plank to dry, variant
538 of [46]. The performance metric is the number of cloth particles (in simulation) on either
539 side of the plank and above the ground. Teacher demonstrations are from an agent with
540 two pickers (*i.e.*, $\mathcal{D}_{Teacher} = \mathcal{D}_{Demo}^{2p}$); we solve the task on a student agent with one
541 picker. Task variations are in cloth rotations and translations with respect to the plank.
542
- 543 3. THREE BOXES: A simple 2D environment where three boxes of different sizes are ran-
544 domly placed and need to be moved to designated goal locations. Teacher demonstrations
545 are from an agent with three pickers (*i.e.*, $\mathcal{D}_{Teacher} = \mathcal{D}_{Demo}^{3p}$); we solve the task on student
546 agents with one picker and two pickers. Performance is measured by the distance of each
547 object from its goal location. This task is used to illustrate the generalizability of **MAIL**
548 with various n -to- m end-effector transfers, and is not used in the SOTA comparisons.

549 A.2 Ablations

550 A.2.1 Ablate the method for creating optimized dataset $\mathcal{D}_{Student}$

551 We answer the question: how do different methods perform in creating optimized dataset $\mathcal{D}_{Student}$?
552 We ablate the optimizer used to create $\mathcal{D}_{Student}$ from the demonstrations, labeled ABL1 in Fig. 7,
553 and compare the following methods, given state inputs from $\mathcal{D}_{Teacher}$.

- 554 • Random: A trivial random guesser, that serves as a lower benchmark.
- 555 • SAC: An RL algorithm that tries to reach the goal states of the demonstrations.
- 556 • Covariant Matrix Adaption Evolution Strategy (CMA-ES): An evolutionary strategy that
557 samples optimization parameters from a multi-variate Gaussian, and updates the mean and
558 covariance at each iteration.
- 559 • Cross-Entropy Method (CEM, ours): A well-known gradient-free optimizer, where we
560 assume a Gaussian distribution for optimization parameters.

561 We did not use gradient-based trajectory optimizers since the contact-rich simulation will give rise
562 to discontinuous dynamics and noisy gradients. As shown in Table 1a, SAC is unable to improve
563 upon the random baseline, likely because of the very large state-space of our environment (> 15000
564 states for > 5000 cloth particles) and error accumulations from the imprecision of learned dynamics
565 model. Trajectory optimizers achieve the highest performance, and we chose CEM as the best
566 optimizer based on the performance of the optimized trajectory.

567 A.2.2 Ablate the dynamics model

568 We answer the question: what is the best architecture to learn the task dynamics? We ablate the
569 learned dynamics model \mathcal{T}_ψ , labeled ABL2 in Fig. 7. The environment state is the state from

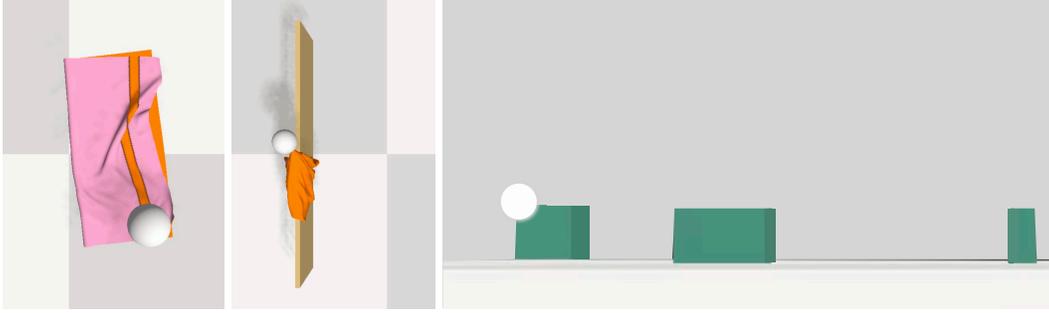


Figure 6: **Environments** used in our experiments, with one end-effector. The end-effectors are pickers (white spheres). In CLOTH FOLD (left) the robot has to fold the cloth (orange and pink) along an edge (inspired by the SoftGym [45] two-picker cloth fold task). In DRY CLOTH (middle) the robot has to hang the cloth (orange and pink) on the drying rack (brown plank). In THREE BOXES (right), the robot has to move three rigid boxes in a 2D environment.

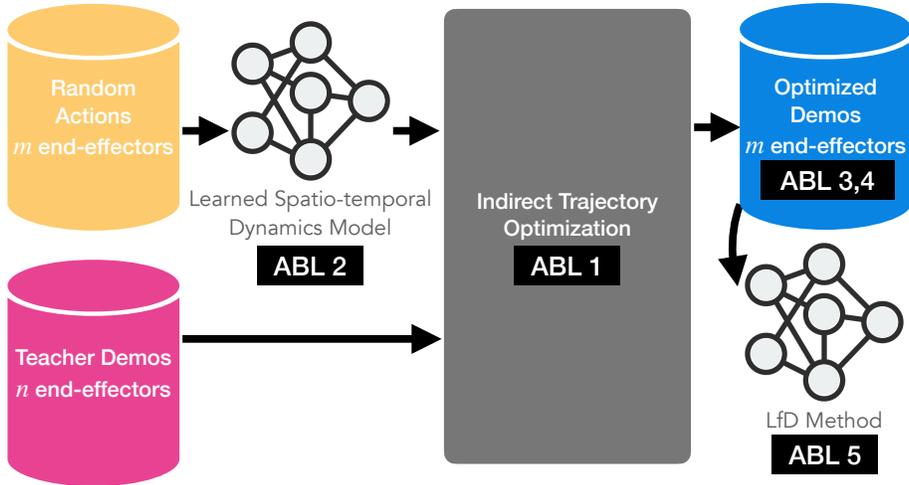


Figure 7: **Ablations** to MAIL components.

570 $\mathcal{D}_{Teacher}$, *i.e.*, positions of cloth particles. This is a structured but large state space since the cloth is
 571 discretized into > 5000 particles.

572 Table 1b shows the performance of trajectories achieved by using the dynamics models. We see that
 573 CNN-LSTM models work better than models that contain only CNNs, graph networks (GNS), or
 574 LSTMs. We hypothesize that this is the case since we need to capture the spatial structure of cloth
 575 and capture a temporal element across the whole trajectory since particle velocity is not captured in
 576 the state. Further, a 1D CNN works better because the cloth state can be simply represented as a 2D
 577 vector ($N \times 3$ which represents the xyz for N particles). This is easier to learn with than the 3D
 578 state vector fed into 2D CNNs.

579 GNS performs poorly also due to the reasons of error accumulation from large displacements, dis-
 580 cussed in Sec. 4.2. Our learned dynamics model \mathcal{T}_ψ was significantly faster than the simulator.
 581 We tested it on a simple training run of SAC [5], without parallelization. Our learned dynamics
 582 gave 162 fps, about $50x$ faster than the 3.4 fps with the simulator. The accuracy was tolerable for
 583 trajectory optimization, as shown in Fig. 8.

584 A.2.3 Compare performance of optimized dataset \mathcal{D}_{Optim}^{1p}

585 We answer the question: how good is $\mathcal{D}_{Student}$ compared to the recorded demonstrations? This
 586 ablation gauges the performance of the optimized dataset that we used as the student dataset for

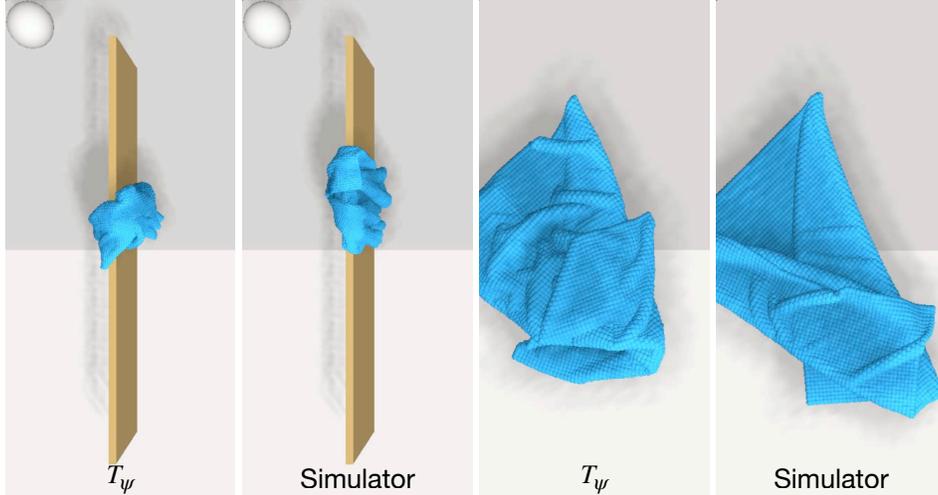


Figure 8: **Predictions of the learned spatio-temporal dynamics model T_ψ and the FleX simulator.** Predictions are made for the same state and action, shown for both cloth tasks. The learned model supports optimization approximately $50x$ faster than the simulator, albeit at the cost of accuracy.

587 LfD, $\mathcal{D}_{Student} = \mathcal{D}_{Optim}^{1p}$. We compare this to other relevant datasets to solve the task, as shown
 588 in Table 1c. It is labeled ABL3 in Fig. 7. The two-picker demonstrations \mathcal{D}_{Demo}^{2p} are recorded
 589 for an agent with two pickers as end-effectors. This is used as the teacher demonstrations in our
 590 experiment $\mathcal{D}_{Teacher} = \mathcal{D}_{Demo}^{2p}$. The one-picker demonstrations \mathcal{D}_{Demo}^{1p} are recorded for an agent
 591 with one picker as an end-effector. This is to contrast against the optimized demonstrations in the
 592 same morphology, \mathcal{D}_{Optim}^{1p} . The random action trajectories are with a one-picker agent, added as
 593 a lower performance benchmark. They are the same random trajectories used to train the spatio-
 594 temporal dynamics model T_ψ . Naturally, the teacher dataset is the best, as it is trivial to do this task
 595 with two pickers. The one-picker dataset has about the same performance as the optimized dataset
 596 \mathcal{D}_{Optim}^{1p} , both of which are suboptimal, as it is not trivial to manipulate cloth with one hand. *This*
 597 *is the kind of task we wish to unlock with this work: tasks that are easy to do for teachers in one*
 598 *morphology but difficult to program or record demonstrations for in the student’s morphology.* Note
 599 that \mathcal{D}_{Optim}^{1p} has been optimized on the fast but inaccurate learned dynamics model, which is one
 600 reason for the reduced performance. This is why the downstream LfD method uses the simulator, as
 601 accuracy is very important in the final policy.

602 A.2.4 Ablate modality of demonstrations

603 We answer the question: how well does the downstream LfD method handle different kinds of
 604 demonstrations? This ablates the composition of the student dataset fed into LfD, and is labeled
 605 ABL4 in Fig. 7. We compare the following datasets for $\mathcal{D}_{Student}$, using the notation for datasets
 606 explained in Sec. 3.1:

- 607 • Demonstrations in one-picker morphology, \mathcal{D}_{Demo}^{1p} : These are non-trivial to create and are
 608 thus not as performant, discussed above. Creating these is increasingly difficult as the task
 609 becomes more challenging.
- 610 • Optimized demos, \mathcal{D}_{Optim}^{1p} : This is optimized from the two-picker teacher demonstrations
 611 ($\mathcal{D}_{Teacher} = \mathcal{D}_{Demo}^{2p}$), which are easy to collect as the task is trivial with two pickers.
- 612 • 50% \mathcal{D}_{Demo}^{1p} and 50% \mathcal{D}_{Optim}^{1p} : A mix of trajectories from the two cases above. This is an
 613 example of handling multiple demonstrators with different morphologies.

Method	25 th %	$\mu \pm \sigma$	median	75 th %
Random	0.000	0.003±0.088	0.000	0.000
SAC	0.000	0.000±0.006	0.000	0.000
CMA-ES	0.104	0.270±0.258	0.286	0.489
CEM	0.351	0.502±0.242	0.501	0.702

(a) Ablation on the method chosen for creating demonstrations.

Method	25 th %	$\mu \pm \sigma$	median	75 th %
GNS	-0.182	0.002±0.223	-0.042	0.149
2D CNN, LSTM	0.157	0.376±0.305	0.382	0.602
No CNN, LSTM	0.327	0.465±0.213	0.463	0.595
1D CNN, No LSTM	0.202	0.407±0.237	0.387	0.587
1D CNN, LSTM (ours)	0.351	0.502±0.242	0.501	0.702

(b) Ablation on the dynamics network architecture.

Dataset	25 th %	$\mu \pm \sigma$	median	75 th %
\mathcal{D}_{Random}	0.000	0.003±0.088	0.000	0.000
\mathcal{D}_{Demo}^{1p}	0.344	0.484±0.169	0.446	0.641
\mathcal{D}_{Demo}^{2p}	0.696	0.744±0.068	0.724	0.785
\mathcal{D}_{Optim}^{1p}	0.351	0.502±0.242	0.501	0.702

(c) Compare the performance of the optimized dataset.

Table 1: Ablation results for MAIL

614 Fig. 9 illustrates that all three variants achieve similar final performance. This demonstrates that the
615 downstream LfD method is capable of solving the task with a variety of suboptimal demonstrations.
616 This could be from one dataset of demonstrations, or even a combination of datasets obtained from
617 a heterogeneous set of teachers.

618 An interesting observation here is that by comparing Fig. 9 and Table 1c, we see that the final
619 policy is better than the suboptimal demonstrations by a considerable margin, and also slightly im-
620 proves upon the performance of the teacher demonstrations. This improvement comes from the LfD
621 method’s ability to effectively utilize demonstrations and generalize across task variations. This re-
622 sult, combined with the ablation that we need demonstrations in Sec. 4.2, shows that our downstream
623 LfD method is well adapted to work with suboptimal demonstrations to solve a task.

624 A.2.5 Ablate Reference State Initialization in DMfD

625 We answer the question: how does the use of demonstration state matching affect the downstream
626 LfD? An improvement we made over the original DMfD algorithm is to disable matching with
627 expert states, known as RSI-IR, first proposed in [42]. We justify this improvement in this ablation,
628 labeled ABL5 in Fig. 7.

629 As shown in Fig. 10, removing RSI and IR has a net positive effect throughout training, and around
630 10% on the final policy performance. This means that matching expert states exactly via imitation
631 reward does not help, even during the initial stages of training when the policy is randomly initial-
632 ized. We believe this is because RSI helps when there are hard-to-reach intermediate states that the
633 policy cannot reach during the initial stages of training. This is true for dynamic or long-horizon
634 tasks, such as karate chops and roundhouse kicks. However, our tasks are quasi-static, and also have
635 a short horizon of 3 for the cloth tasks. In other words, removing this technique allows the policy to

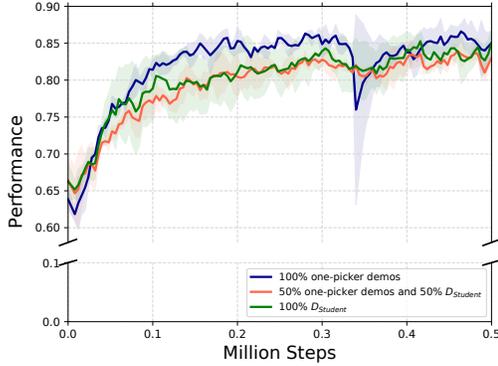


Figure 9: **Ablation on the modality of demonstrations on LfD performance.** Similar performance shows that MAIL can learn from a wide variety of demonstrations, or even a mixture of them, without loss in performance. See Sec. A.2.4.

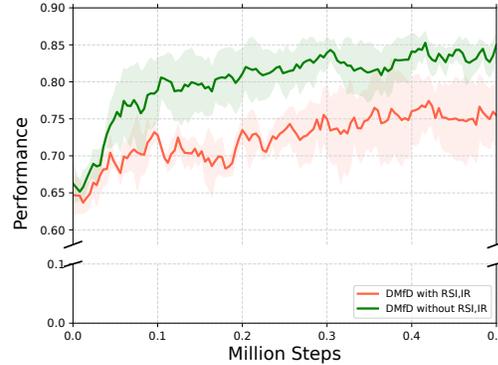


Figure 10: **Ablation on the effect of reference state initialization (RSI) and imitation reward (IR) on LfD performance.** RSI is not helpful here because our tasks are not as dynamic or long horizon as DeepMimic [42]. See Sec. A.2.5.

636 freely explore the state space while the demonstrations can still guide the RL policy learning via the
 637 advantage-weighted loss from DMfD.

638 A.2.6 Ablate the effect of cross-morphology on SOTA

639 We answer the question: how do established LfD baselines perform across morphologies? We
 640 studied the effect of why baselines such as GAIfo and GPIL performed so poorly on our tasks.
 641 In our experiments, we noticed a number of factors (such as variations in the task, diversity of
 642 demonstrations, etc.). This ablation studies the effect of cross-morphology in the demonstrations,
 643 where we compare the performance of GAIfo, when provided demonstrations from the teacher
 644 dataset $\mathcal{D}_{Teacher}$ and student dataset $\mathcal{D}_{Student}$.

645 As we can see in Table 2, there is a 36% performance improvement when using the (suboptimal)
 646 student dataset. Obviously, since the demonstration actions are not available to learn from, the pri-
 647 mary difference that the agent sees during training is the richness of demonstration states. Thus,
 648 improvement is because of the demonstration states seen in the student dataset. Since the student
 649 morphology has only one picker, any demonstration for the task (DryCloth) includes multiple inter-
 650 mediate states of the cloth in various conditions of being partially hung for drying. By contrast, the
 651 teacher requires fewer pick-place steps to complete the task, and thus there are fewer intermediate
 652 states in the demonstrations.

653 A.2.7 Ablate the effect of environment difficulty on LfD baselines

654 We answer the question: how do established LfD baselines perform across environments? Given
 655 the subpar performance of the LfD baselines GAIfo and GPIL on our SOTA environments, we
 656 ablated the effect of environment difficulty. We took the easy cloth environment (CLOTH FOLD)
 657 and used an easier variant of it, CLOTH FOLD DIAGONAL PINNED [41]. In this variant, the agent
 658 has to perform an easier fold, but one corner of the cloth is pinned to prevent sliding. Moreover, the
 659 desired fold is across the diagonal of the cloth, which can be done by manipulating only one corner
 660 of the cloth. We used the state-based observations, and the action space is the small-displacement
 661 action space, where the agent outputs incremental picker displacements instead of pick-and-place
 662 locations. This action space is similar to those seen in the experiments of GNS, GAIfo and GPIL,
 663 where they worked with rigid objects in simulation. This is an easy version of our CLOTH FOLD
 664 environment. We can see in Table 3 that the same baselines are able to perform significantly better
 665 in this environment. Hence, we believe manipulating with long-horizon pick-place actions, with an

666 image observation, makes it challenging for the baselines to work in challenging cloth environments
 667 described in Sec. 4.1.

Method	25 th %	$\mu \pm \sigma$	median	75 th %
D _{Teacher}	-0.198	-0.055±0.183	-0.043	0.078
D _{Student}	0.199	0.363±0.245	0.409	0.528

Table 2: Ablation of GAIfO on the effect of cross-morphology. We compare the normalized performance, measured at the end of the task.

Method	25 th %	$\mu \pm \sigma$	median	75 th %
GPIL	0.356	0.427±0.162	0.487	0.553
GAIfO	0.115	0.374±0.267	0.471	0.592

Table 3: Measuring performance on the easy cloth task, CLOTH FOLD DIAGONAL PINNED. We compare the normalized performance, measured at the end of the task.

668 A.3 List of environments we tried for LfD baseline ablations

669 Two LfD baselines, GAIfO and GPIL, seemed to perform quite poorly, although we expected better
 670 performance. In an effort to understand why these fail, we performed a host of studies with different
 671 varieties of easier environments, to isolate the properties of the environment that make it the most
 672 challenging to succeed. A list of the different task variants we tried are given below. The ones
 673 with the most striking difference in performance are described in further detail in Sec. A.2.6 and
 674 Sec. A.2.7.

- 675 1. Used the easier CLOTH FOLD environment instead of DRY CLOTH.
- 676 2. Used state-based environments instead of image-based environments.
- 677 3. Reduced the number of variations of the task distribution \mathcal{V} .
- 678 4. Used the small-displacement action-space that is used in GNS and GAIfO experiments,
 679 instead of the large-displacement pick-place action spaces.
- 680 5. Removed the effect of cross-morphology, by providing demonstrations in the students mor-
 681 phology.

682 A.4 Hyperparameter choices for MAIL

683 In this section, Table 4 shows the hyperparameters chosen for training the inverse dynamics model
 684 \mathcal{T}_ψ . Table 5 shows the details of CEM hyperparameter choices. Table 6 shows the hyperparameters
 685 for our chosen LfD method (DMfD).

Parameter	Description
CNN	4 layers, 32 channels, 3x3 kernel, leaky ReLU activation. stride = 2 for the first layer, stride = 1 for subsequent layers
LSTM	One layer Hidden size = 32
Other Parameters	Learning rate $\alpha = 1e-5$ Batch size = 128

Table 4: Hyper-parameters for training the forward dynamics model.

	Planning Horizon	Number of optimization iterations	Number of env interactions
1	1	2	21,000
2	2	2	15,000
3	2	2	21,000
4	2	2	31,000
5	2	2	34,000
6	2	10	21,000
7	2	1	21,000
8	2	1	15,000
9	2	1	32,000
10	3	2	21,000
11	3	10	21,000
12	4	2	21,000
13	4	10	21,000

Table 5: CEM hyper-parameters tested for tuning the trajectory optimization. We conducted ten rollouts for each parameter set and used the set with the highest average normalized performance on the teacher demonstrations. Population size is determined by the number of environment interactions. The number of elites for each CEM iteration is 10% of population size.

Parameter	Description
State encoding	Fully connected network (FCN) 2 hidden layers of 1024, ReLU activation
Image encoding	32x32 RGB input, with random crops. CNN: 4 layers, 32 channels, stride 1, 3x3 kernel, leaky ReLU activation FCN: 1 layer of 1024 neurons, <i>tanh</i> activation
Actor	Fully connected network 2 hidden layers of 1024, leaky ReLU activation
Critic	Fully connected network 2 hidden layers of 1024, leaky ReLU activation
Other parameters	Discount factor: $\gamma = 0.9$ Entropy loss weight: $w_E = 0.1$ Entropy regularizer coefficient: $\alpha = 0.5$ Batch size = 256 Replay buffer size = 600,000 RSI-IR probability = 0 (disabled)

Table 6: Hyper-parameters used in the LfD method (DMfD).

686 A.5 Performance metrics for real-world cloth experiments

687 In this section, we explain the metrics for measuring performance of the cloth, to explain the
688 sim2real results discussed in Fig. 4.2

689 For CLOTH FOLD task, we measure performance at time t by the number of pixels of the top color
690 $pix_{top,t}$ and bottom color $pix_{bot,t}$ of the flattened cloth, compared to the maximum number of pixels,
691 pix_{max} (Fig. 11).

692 For DRY CLOTH task, it is challenging to measure pixels on the sides and top of the plank. Moreover,
693 we could be double counting pixels if they are visible in both side and top views. Hence, we measure
694 the cloth to determine whether the length of the cloth *on top of* the plank is equal to or greater than
695 the side of the square cloth. We call this the spread metric.

696 The policies achieve $\sim 80\%$ performance, which is about the average performance of our method in
697 simulation, for both tasks. However, since these performance metrics are different in the simulation
698 and real world, we cannot *quantify* the sim2real gap through these numbers.

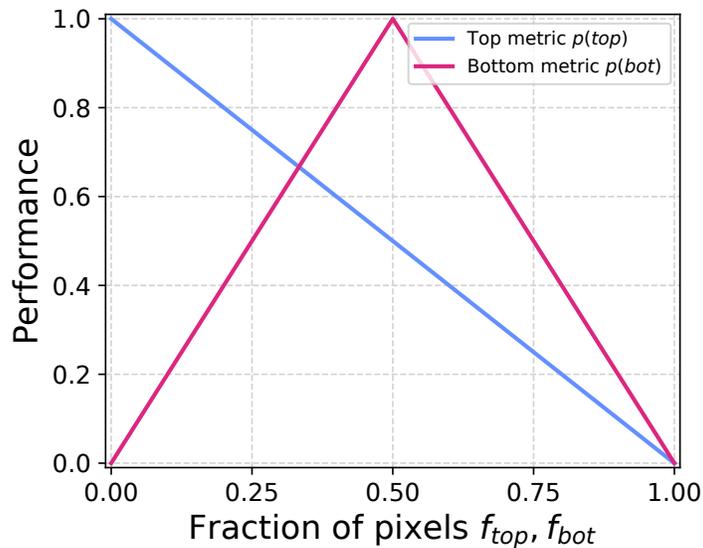


Figure 11: **Performance function for CLOTH FOLD on the real robot.** At time t , we measure the fraction of pixels visible to the maximum number of pixels visible $f_{top} = pix_{top,t}/pix_{max}$ and $f_{bot} = pix_{bot,t}/pix_{max}$. Performance for the top of the cloth should be 1 when it is not visible, $p(top) = 1 - f_{top}$. Performance for the bottom of the cloth should be 1 when it is exactly half-folded on top of the top side, $p(bot) = \min[2(1 - f_{bot}), 2f_{bot}]$. Final performance is an average of both metrics, $p(s_t) = (p(top) + p(bottom))/2$. Note that the cloth is flattened at the start, thus $p_{max} = p_{top,0}$.