

497 **A Proofs**

498 **A.1 Proof of Proposition 3.1**

499 Expanding the entropy in Equation 5 we obtain

$$p^*(o) \propto \exp\left(\mathbb{E}_{p^*(\mathcal{D}|o)}[R_o(\mathcal{D})/\eta - \log p^*(\mathcal{D}|o)]\right).$$

500 Using $p^*(\mathcal{D}|o) = \tilde{p}(\mathcal{D}|o)/\sum_n \tilde{p}(\mathcal{D}_n|o)$ yields

$$p^*(o) \propto \exp\left(\mathbb{E}_{p^*(\mathcal{D}|o)}[R_o(\mathcal{D})/\eta - \log \tilde{p}(\mathcal{D}|o) + \log \sum_n \tilde{p}(\mathcal{D}_n|o)]\right).$$

501 Next, leveraging that $\log \tilde{p}(\mathcal{D}|o) = R_o(\mathcal{D}_n)/\eta$ we see that

$$p^*(o) \propto \exp\left(\mathbb{E}_{p^*(\mathcal{D}|o)}[\log \sum_n \tilde{p}(\mathcal{D}_n|o)]\right) = \sum_n \tilde{p}(\mathcal{D}_n|o),$$

502 which concludes the proof. \square

503 **A.2 Proof of Corollary 3.1.2**

504 We start by rewriting the lower bound as $L(\psi, q) =$

$$\mathbb{E}_{p^*(o)}\left[\mathbb{E}_{p^*(\mathcal{D}|o)}[R_o(\mathcal{D}) - \eta \log p^*(\mathcal{D}|o)] - \eta \log p^*(o)\right].$$

505 Using $p^*(\mathcal{D}|o) \propto \tilde{p}(\mathcal{D}|o)$ and Proposition 3.1 we obtain

$$\begin{aligned} L(\psi, q) &= \mathbb{E}_{p^*(o)}\left[\mathbb{E}_{p^*(\mathcal{D}|o)}[R_o(\mathcal{D}) - \eta \log \tilde{p}(\mathcal{D}|o) \right. \\ &\quad \left. + \eta \log \sum_n \tilde{p}(\mathcal{D}_n|o)] - \eta \log \sum_n \tilde{p}(\mathcal{D}_n|o) \right. \\ &\quad \left. + \eta \log \sum_o \sum_n \tilde{p}(\mathcal{D}_n|o)\right] \end{aligned}$$

506 With $\eta \log \tilde{p}(\mathcal{D}|o) = R_o(\mathcal{D}_n)$ all most terms cancel, giving

$$\begin{aligned} L(\psi, q) &= \mathbb{E}_{p^*(o)}\left[\eta \log \sum_o \sum_n \tilde{p}(\mathcal{D}_n|o)\right] \\ &= \eta \log \sum_o \sum_n \tilde{p}(\mathcal{D}_n|o), \end{aligned}$$

507 which concludes the proof. \square

508 **A.3 Proof of Corollary 3.1.1**

509 Expanding the expected KL divergence, we get

$$\begin{aligned} &\min_{\phi} \mathbb{E}_{p(\mathcal{D})} D_{\text{KL}}(p(o|\mathcal{D}) \| g_{\phi}(o|\mathbf{x})) \\ &= \min_{\phi} \sum_n p(\mathcal{D}_n) \sum_o p(o|\mathcal{D}_n) \log \frac{p(o|\mathcal{D}_n)}{g_{\phi}(o|\mathbf{x}_n)}. \end{aligned}$$

510 Noting that $p(o|\mathcal{D}_n)$ is independent of ϕ we can rewrite the objective as

$$\max_{\phi} \sum_n p(\mathcal{D}_n) \sum_o p(o|\mathcal{D}_n) \log g_{\phi}(o|\mathbf{x}_n).$$

511 Using that $p(o|\mathcal{D}) = \tilde{p}(\mathcal{D}|o)/\sum_o \tilde{p}(\mathcal{D}|o)$ together with $p(\mathcal{D}) = \sum_o p^*(o)p^*(\mathcal{D}|o)$ yields

$$\max_{\phi} \sum_n \sum_o p^*(o)p^*(\mathcal{D}_n|o) \sum_o \frac{\tilde{p}(\mathcal{D}_n|o)}{\sum_o \tilde{p}(\mathcal{D}_n|o)} \log g_{\phi}(o|\mathbf{x}_n).$$

512 Using Proposition 3.1 we can rewrite $p^*(o)p^*(\mathcal{D}|o)$ as $\tilde{p}(\mathcal{D}|o)/\sum_o \sum_n \tilde{p}(\mathcal{D}_n|o)$. Since the constant
 513 factor $1/\sum_o \sum_n \tilde{p}(\mathcal{D}_n|o)$ does not affect the optimal value of ϕ we obtain

$$\begin{aligned} & \max_{\phi} \sum_n \sum_o \tilde{p}(\mathcal{D}_n|o) \sum_o \frac{\tilde{p}(\mathcal{D}_n|o)}{\sum_o \tilde{p}(\mathcal{D}_n|o)} \log g_{\phi}(o|\mathbf{x}_n) \\ & \max_{\phi} \sum_n \sum_o \tilde{p}(\mathcal{D}_n|o) \log g_{\phi}(o|\mathbf{x}_n), \end{aligned}$$

514 which concludes the proof. \square

515 B Derivations

516 B.1 Lower Bound Decomposition

517 To arrive at Equation 4 by marginalizing over the latent variable o for the entropy of the joint
 518 curriculum, i.e.,

$$\begin{aligned} \mathcal{H}(\mathcal{D}) &= - \sum_n p(\mathcal{D}_n) \log p(\mathcal{D}_n) \\ &= - \sum_n p(\mathcal{D}_n) \sum_o p(o|\mathcal{D}_n) \log p(\mathcal{D}_n) \end{aligned}$$

519 Next, we use Bayes' theorem, that is, $p(\mathcal{D}_n) = p(o)p(\mathcal{D}_n|o)/p(o|\mathcal{D}_n)$, giving

$$\begin{aligned} \mathcal{H}(\mathcal{D}) &= - \sum_n p(\mathcal{D}_n) \sum_o p(o|\mathcal{D}_n) (\log p(o) + \log p(\mathcal{D}_n|o) \\ &\quad - \log p(o|\mathcal{D}_n)). \end{aligned}$$

520 Moreover, we add and subtract the log auxiliary distribution $\log q(o|\mathcal{D}_n)$ which yields

$$\begin{aligned} \mathcal{H}(\mathcal{D}) &= - \sum_n p(\mathcal{D}_n) \sum_o p(o|\mathcal{D}_n) (\log p(o) + \log p(\mathcal{D}_n|o) \\ &\quad - \log p(o|\mathcal{D}_n) + \log q(o|\mathcal{D}_n) - \log q(o|\mathcal{D}_n)). \end{aligned}$$

521 Rearranging the terms leads and writing the sums in terms of expectations we arrive at

$$\begin{aligned} \mathcal{H}(\mathcal{D}) &= \mathbb{E}_{p(o)} [\mathbb{E}_{p(o|\mathcal{D})} [\log q(o|\mathcal{D})] + \mathcal{H}(\mathcal{D}|o)] + \mathcal{H}(o) \\ &\quad + D_{\text{KL}}(p(o|\mathcal{D})\|q(o|\mathcal{D})). \end{aligned}$$

522 Lastly, multiplying $\mathcal{H}(\mathcal{D})$ with η and adding $\mathbb{E}_{p(o)} \mathbb{E}_{p(\mathcal{D}|o)} [\log p_{\theta_o}(\mathbf{y}|\mathbf{x}, o)]$ we arrive at Equation 4
 523 which concludes the derivation.

524 B.2 M-Step Objectives

525 **Closed-Form Curriculum Updates.** In order to derive the closed-form solution to Equation equa-
 526 tion 5 (RHS) we solve

$$\max_{p(\mathcal{D}|z)} J_z(p(\mathcal{D}|z), \theta_z) = \max_{p(\mathcal{D}|z)} \mathbb{E}_{p(\mathcal{D}|z)} [R_z(\mathcal{D})] + \eta \mathcal{H}(\mathcal{D}|z) \quad \text{subject to} \quad \sum_n p(\mathcal{D}) = 1.$$

527 Following the procedure of constrained optimization, we write down the Lagrangian function [55] as

$$\mathcal{L}(p, \lambda) = \sum_n p(\mathcal{D}_n|z) R_z(\mathcal{D}_n) - \eta \sum_n p(\mathcal{D}_n|z) \log p(\mathcal{D}_n|z) + \lambda (\sum_n p(\mathcal{D}_n|z) - 1),$$

528 where λ is the Lagrangian multiplier. As p is discrete, we solve for the optimal entries of $p(\mathcal{D}_n|z)$,
 529 that is, $p'(\mathcal{D}_n, \lambda|z) = \arg \max_p \mathcal{L}(p, \lambda)$. Setting the partial derivative of $\mathcal{L}(p, \lambda)$ with respect to p
 530 zero, i.e.,

$$\frac{\partial}{\partial p(\mathcal{D}_n|z)} \mathcal{L}(p, \lambda) = R_z(\mathcal{D}_n) - \eta \log p(\mathcal{D}_n|z) - \eta + \lambda \stackrel{!}{=} 0.$$

531 yields $p'(\mathcal{D}_n, \lambda|z) = \exp(R_z(\mathcal{D}_n) - \eta + \lambda)/\eta$.

532 Plugging p' back in the Lagrangian gives the dual function $g(\lambda)$, that is,

$$g(\eta) = \mathcal{L}(p', \lambda) = -\eta + \eta \sum_n \exp(R_z(\mathcal{D}_n) - \eta + \lambda)/\eta.$$

533 Solving for $\lambda^* = \arg \min_{\lambda \geq 0} g(\lambda)$ equates to

$$\begin{aligned} \frac{\partial}{\partial \lambda} g(\lambda) &= -1 + \eta \sum_n \exp(R_z(\mathcal{D}_n) - \eta + \lambda)/\eta \stackrel{!}{=} 0 \\ \iff \lambda^* &= -\log\left(\eta \sum_n \exp(R_z(\mathcal{D}_n) - \eta)/\eta\right). \end{aligned}$$

534 Finally, substituting λ^* into p' we have

$$p^*(\mathcal{D}_n|z) = p'(\mathcal{D}_n, \lambda^*|z) = \frac{\exp(R_z(\mathcal{D}_n)/\eta)}{\sum_n \exp(R_z(\mathcal{D}_n)/\eta)},$$

535 which concludes the derivation. The derivation of the optimal mixture weights $p^*(z)$ works analo-
536 gously.

537 **Expert Objective.** In order to derive the expert objective of Equation 6 we solve

$$\max_{\theta_z} J_z(p(\mathcal{D}|z), \theta_z) = \max_{\theta_z} \sum_n p(\mathcal{D}_n|z) \left(\log p_{\theta_z}(\mathbf{a}_n|\mathbf{o}_n, z) + \eta \log q(z|\mathcal{D}_n) - \eta \log p(\mathcal{D}_n|z) \right).$$

538 Noting that $q(z|\mathcal{D}_n)$ and $p(\mathcal{D}_n|z)$ are independent of θ_z and $p(\mathcal{D}_n|z) = \tilde{p}(\mathcal{D}_n|z)/\sum_n \tilde{p}(\mathcal{D}_n|z)$ we
539 find that

$$\max_{\theta_z} J_z(p(\mathcal{D}|z), \theta_z) = \max_{\theta_z} \sum_n \frac{\tilde{p}(\mathcal{D}_n|z)}{\sum_n \tilde{p}(\mathcal{D}_n|z)} \log p_{\theta_z}(\mathbf{a}_n|\mathbf{o}_n, z).$$

540 Noting that $\sum_n \tilde{p}(\mathcal{D}_n|z)$ is a constant scaling factor concludes the derivation.

541 C Experiment Setup

542 C.1 Environments and Datasets

543 C.1.1 Obstacle Avoidance

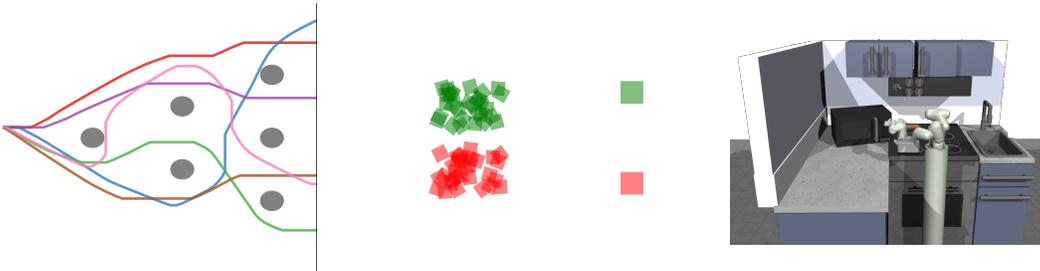


Figure 7: The left figure shows 6 out of 24 ways of completing the obstacle avoidance task. The middle figure shows 30 randomly sampled initial block configurations for the block pushing task. The right figure visualizes the Franka kitchen environment.

544 **Dataset.** The obstacle avoidance dataset contains 96 trajectories resulting in a total of 7.3k (\mathbf{o}, \mathbf{a})
545 pairs. The observations $\mathbf{o} \in \mathbb{R}^4$ contain the end-effector position and velocity in Cartesian space.
546 Please note that the height of the robot is fixed. The actions $\mathbf{a} \in \mathbb{R}^2$ represent the desired position of
547 the robot. The data is recorded such that there are an equal amount of trajectories for all 24 ways of
548 avoiding the obstacles and reaching the target line. For successful example trajectories see Figure 7.

549 **Performance Metrics.** The *success rate* indicates the number of end-effector trajectories that
 550 successfully reach the target line (indicated by green color in Figure 3). The *entropy*

$$\mathcal{H}_{24}(\tau) = - \sum_{\tau} p(\tau) \log_{24} p(\tau),$$

551 is computed for successful trajectories τ . To assess the model performance, we simulate 1000
 552 end-effector trajectories. We count the number of successful trajectories for each way of completing
 553 the task. From that, we calculate a categorical distribution $p(\tau)$ which is used to compute the entropy.
 554 By the use of \log_{24} we make sure that $\mathcal{H}_{24}(\tau) \in [0, 1]$. If a model is able to discover all modes in
 555 the data distribution with equal probability, its entropy will be close to 1. In contrast, $\mathcal{H}_{24}(\tau) = 0$ if
 556 a model only learns one solution.

557 C.1.2 Block Pushing

558 **Dataset.** The block pushing dataset contains 500 trajectories for each of the four push sequences
 559 (see Figure 4) resulting in a total of 2000 trajectories or 463k (\mathbf{o}, \mathbf{a}) pairs. The observations $\mathbf{o} \in \mathbb{R}^{16}$
 560 contain the desired position and velocity of the robot in addition to the position and orientation of the
 561 green and red block. Please note that the orientation of the blocks is represented as quaternion number
 562 system and that the height of the robot is fixed. The actions $\mathbf{a} \in \mathbb{R}^2$ represent the desired position
 563 of the robot. This task is similar to the one proposed in [17]. However, they use a deterministic
 564 controller to record the data whereas we use human demonstrators which increases the difficulty of
 565 the task significantly due to the inherent versatility in human behavior.

566 **Performance Metrics.** The *success rate* indicates the number of end-effector trajectories τ that
 567 successfully push both blocks to different target zones. To assess the model performance on non-
 568 successful trajectories, we consider the *distance error*, that is, the euclidean distance from the blocks
 569 to the target zones at the final block configuration of an end-effector trajectory. As there are a total of
 570 four push sequences (see Figure 3) we use the expected *entropy*

$$\mathbb{E}_{p(\mathbf{c})} \mathcal{H}_4(\tau|\mathbf{c}) = - \sum_{\mathbf{c}} p(\mathbf{c}) \sum_{\tau} p(\tau|\mathbf{c}) \log_4 p(\tau|\mathbf{c}),$$

571 to quantify a model’s ability to cover the modes in the data distribution. Please note that we set
 572 $p(\mathbf{c}) = 1/30$ as we sample 30 block configurations uniformly from a configuration space (see Figure
 573 7). For each \mathbf{c} we simulate 16 end-effector trajectories. For a given configuration, we count how often
 574 each of the four push-sequences is executed successfully and use the result to calculate a categorical
 575 distribution $p(\tau|\mathbf{c})$. Once repeated for all 30 configurations, we compute $\mathbb{E}_{p(\mathbf{c})} \mathcal{H}_4(\tau|\mathbf{c})$. Using \log_4
 576 we make sure that the expected entropy is upper bounded by 1. This bound is achieved if a model is
 577 able to execute each of the push sequences with equal probability for all configurations. If a model
 578 only executes one sequence successfully, the entropy is 0.

579 C.1.3 Franka Kitchen

580 **Dataset.** The Franka kitchen environment was introduced in [51]. It contains 566 human-collected
 581 trajectories resulting in a total of 128k (\mathbf{o}, \mathbf{a}) pairs. The observations $\mathbf{o} \in \mathbb{R}^{30}$ contain information
 582 about the position and orientation of the task-relevant objects in the environment. The actions $\mathbf{a} \in \mathbb{R}^9$
 583 represent the signals to control the robot and the gripper. The dataset comprises sequences that
 584 successfully solve 4 out of 7 tasks in different orders.

585 **Performance Metrics.** First, we consider the *success rate* for a different number of tasks solved. We
 586 additionally compute the *entropy* over task sequences. This is computed using 100 simulated robot
 587 trajectories. For trajectories with a single task solved, we count how frequently each of the tasks is
 588 executed. From that, we calculate a categorical distribution which is then used for computing the
 589 entropy. We generalize this concept to more successful task completions, by calculating a categorical
 590 distribution over all 7^k possible task sequences for k task completions.

591 C.1.4 Table Tennis

592 **Dataset.** The table tennis dataset contains 5000 (\mathbf{o}, \mathbf{a}) pairs. The observations $\mathbf{o} \in \mathbb{R}^4$ contain the
 593 coordinates of the initial and target ball position as projection on the table. Movement primitives
 594 (MPs) [30] are used to describe the joint space trajectories of the robot manipulator using two basis
 595 functions per joint and thus $\mathbf{a} \in \mathbb{R}^{14}$.

596 **Metrics.** To evaluate the different algorithms on the demonstrations recorded using the table tennis
 597 environment quantitatively, we employ two performance metrics: The *success rate* and the *distance*
 598 *error*. The success rate is the percentage of strikes where the ball is successfully returned to the
 599 opponent’s side. The distance error, is the distance between the target position and landing position
 600 of the ball for successful strikes.

601 C.1.5 Human Subjects for Data Collection

602 For the obstacle avoidance as well as the block pushing experiments we used data collected by
 603 humans. We note that all human subjects included in the data collection process are individuals who
 604 are collaborating on this work. The participants did, therefore, not receive any financial compensation
 605 for their involvement in the study.

606 C.2 IMC Details and Hyperparameter

607 IMC employs a parameterized inference network and conditional Gaussian distributions to represent
 608 experts. For the latter, we use a fixed variance of 1 and parameterize the means as neural networks.
 609 For both inference network and expert means we use residual MLPs [56]. For all experiments, we use
 610 batch-size $|\mathcal{B}| = |\mathcal{D}|$, number of components $N_z = 50$ and expert learning rate equal to 5×10^{-4} .
 611 Furthermore, we initialized all curriculum weights as $p(\mathcal{D}_n|z) = 1$. For the table tennis and obstacle
 612 avoidance task, we found the best results using a multi-head expert parameterization (see Section
 613 E.1) where we tested 1 – 4 layer neural networks. We found that using 1 layer with 32 neurons
 614 performs best on the table tennis task and 2 layer with 64 neurons for the obstacle avoidance task. For
 615 the block pushing and Franka kitchen experiments, we obtained the best results using a single-head
 616 parameterization of the experts. We used 6 layer MLPs with 128 neurons for both tasks. For the
 617 inference network, we used a fixed set of parameters that are listed in Table 2. For the entropy scaling
 618 factor η we performed a hyperparameter sweep using Bayesian optimization. The respective values
 619 are $\eta = 1/30$ for obstacle avoidance, $\eta = 2$ for block pushing and Franka kitchen and $\eta = 1$ for
 table tennis.

Table 2: IMC & EM Hyperparameter.

PARAMETER	VALUE
EXPERT LEARNING RATE	10^{-4}
EXPERT BATCHSIZE	1024
EXPERT VARIANCE (σ^2)	1
INFERENCE NET HIDDEN LAYER	6
INFERENCE NET HIDDEN UNITS	256
INFERENCE NET EPOCHS	800
INFERENCE NET LEARNING RATE	10^{-3}
INFERENCE NET BATCHSIZE	1024

620

621 C.3 Baselines and Hyperparameter

622 We now briefly mention the baselines and their hyperparameters. We used Bayesian optimization to
 623 tune the most important hyperparameters.

624 **Mixture of Experts trained with Expectation-Maximization (EM).** The architecture of the mixture
 625 of experts model trained with EM [48] is identical to the one optimized with IMC: We employ a
 626 parameterized inference network and conditional Gaussian distributions to represent experts with
 627 the same hyperparameters as shown in Table 2. Furthermore, we initialized all responsibilities as
 628 $p(z|\mathbf{o}) = 1/N_z$, where N_z is the number of components.

629 **Mixture Density Network (MDN).** The mixture density network [8] uses a shared backbone neural
 630 network with multiple heads for predicting component indices as well as the expert likelihood. For the
 631 experts, we employ conditional Gaussians with a fixed variance. The model likelihood is maximized
 632 in an end-to-end fashion using stochastic gradient ascent. We experimented with different backbones
 633 and expert architectures. However, we found that the MDN is not able to partition the input space in

634 a meaningful way, often resulting in sub-optimal outcomes, presumably due to mode averaging. To
 635 find an appropriate model complexity we tested up to 50 expert heads. We found that the number of
 636 experts heads did not significantly influence the results, further indicating that the MDN is not able to
 637 utilize multiple experts to solve sub-tasks. We additionally experimented with a version of the MDN
 638 that adds an entropy bonus to the objective [57] to encourage more diverse and multimodal solutions.
 639 However, we did not find significant improvements compared to the standard version of the MDN.
 For a list of hyperparameter choices see 3.

Table 3: **MDN Hyperparameter.** The ‘Value’ column indicates sweep values for the obstacle avoidance task, the block pushing task, the Franka kitchen task and the table tennis task (in this order).

PARAMETER	SWEEP	VALUE
EXPERT HIDDEN LAYER	{1, 2}	1, 1, 1, 1
EXPERT HIDDEN UNITS	{30, 50}	50, 30, 30, 50
BACKBONE HID. LAYER	{2, 3, 4, 6, 8, 10}	3, 2, 4, 3
BACKBONE HID. UNITS	{50, 100, 150, 200}	200, 200, 200, 200
LEARNING RATE $\times 10^{-3}$	[0.1, 1]	5.949, 7.748, 1.299, 2.577
EXPERT VARIANCE (σ^2)	—	1
MAX. EPOCHS	—	2000
BATCHSIZE	—	512

640

641 **Denosing Diffusion Probabilistic Models (DDPM).** We consider the denosing diffusion proba-
 642 bilistic model proposed by [24]. Following common practice we parameterize the model as residual
 643 MLP [27] with a sinusoidal positional encoding [54] for the diffusion steps. Moreover, we use the
 644 cosine-based variance scheduler proposed by [58]. For further details on hyperparameter choices see
 Table 4.

Table 4: **DDPM Hyperparameter.** The ‘Value’ column indicates sweep values for the obstacle avoidance task, the block pushing task, the Franka kitchen task, and the table tennis task (in this order).

PARAMETER	SWEEP	VALUE
HIDDEN LAYER	{4, 6, 8, 10, 12}	6, 6, 8, 6
HIDDEN UNITS	{50, 100, 150, 200}	200, 150, 200, 200
DIFFUSION STEPS	{5, 15, 25, 50}	15, 15, 15, 15
VARIANCE SCHEDULER	—	COSINE
LEARNING RATE	—	10^{-3}
MAX. EPOCHS	—	2000
BATCHSIZE	—	512

645

646 **Normalizing Flow (NF).** For all experiments, we build the normalizing flow by stacking masked
 647 autoregressive flows [59] paired with permutation layers [18]. As base distribution, we use a
 648 conditional isotropic Gaussian. Following common practice, we optimize the model parameters by
 maximizing its likelihood. See Table 5 for a list of hyperparameters.

Table 5: **NF Hyperparameter.** The ‘Value’ column indicates sweep values for the obstacle avoidance task, the block pushing task, the Franka kitchen task and the table tennis task (in this order).

PARAMETER	SWEEP	VALUE
NUM. FLOWS	{4, 6, 8, 10, 12}	6, 6, 4, 4
HIDDEN UNITS PER FLOW	{50, 100, 150, 200}	100, 150, 200, 150
LEARNING RATE $\times 10^{-4}$	[0.01, 10]	7.43, 4.5, 4.62, 7.67
MAX. EPOCHS	—	2000
BATCHSIZE	—	512

649

650 **Conditional Variational Autoencoder (CVAE).** We consider the conditional version of the autoen-
 651 coder proposed in [20]. We parameterize the encoder and decoder with a neural network with mirrored
 652 architecture. Moreover, we consider an additional scaling factor (β) for the KL regularization in the
 lower bound objective of the VAE as suggested in [60].

Table 6: **CVAE Hyperparameter.** The ‘Value’ column indicates sweep values for the obstacle avoidance task, the block pushing task, the Franka kitchen task and the table tennis task (in this order).

PARAMETER	SWEEP	VALUE
HIDDEN LAYER	{4, 6, 8, 10, 12}	8, 10, 4, 4
HIDDEN UNITS	{50, 100, 150, 200}	100, 150, 100, 100
LATENT DIMENSION	{4, 16, 32, 64}	32, 16, 16, 16
D_{KL} SCALING (β)	$[10^{-3}, 10^2]$	1.641, 1.008, 0.452, 0.698
LEARNING RATE	—	10^{-3}
MAX. EPOCHS	—	2000
BATCHSIZE	—	512

653

654 **Implicit Behavior Cloning (IBC).** IBC was proposed in [17] and uses energy-based models to learn
 655 a joint distribution over inputs and targets. Following common practice we parameterize the model as
 656 neural network. Moreover, we use the version that adds a gradient penalty to the InfoNCE loss [17].
 657 For sampling, we use gradient-based Langevin MCMC [56]. Despite our effort, we could not achieve
 658 good results with IBC. A list of hyperparameters is shown in Table 7.

Table 7: **IBC Hyperparameter.** The ‘Value’ column indicates sweep values for the obstacle avoidance task and the table tennis task (in this order). We do not get any good results for the block push task and the Franka kitchen task.

PARAMETER	SWEEP	VALUE
HIDDEN DIM	{50, 100, 150, 200, 256}	200, 256
HIDDEN LAYERS	{4, 6, 8, 10}	4, 6
NOISE SCALE	[0.1, 0.5]	0.1662, 0.1
TRAIN SAMPLES	[8, 64]	44, 8
NOISE SHRINK	—	0.5
TRAIN ITERATIONS	—	20
INFERENCE ITERATIONS	—	40
LEARNING RATE	—	10^{-4}
BATCH SIZE	—	512
EPOCHS	—	1000

659 **Behavior Transformer (BET).** Behavior transformers were recently proposed in [22]. The model
 660 employs a minGPT transformer [61] to predict targets by decomposing them into cluster centers and
 661 residual offsets. To obtain a fair comparison, we compare our method to the version with no history.
 662 A comprehensive list of hyperparameters is shown in Table 8.

663 D Connection to Expectation Maximization

664 In this section we want to highlight the commonalities and differences between our algorithm and the
 665 expectation-maximization (EM) algorithm for mixtures of experts. First, we look at the updates of
 666 the variational distribution q . Next, we compare the expert optimization. Lastly, we take a closer look
 667 at the optimization of the gating distribution.

668 The EM algorithm sets the variational distribution during the E-step to

$$q(z|\mathbf{o}_n) = p(z|\mathbf{o}_n, \mathbf{a}_n) = \frac{p_{\theta_z}(\mathbf{a}_n|\mathbf{o}_n, z)p(z|\mathbf{o}_n)}{\sum_z p_{\theta_z}(\mathbf{a}_n|\mathbf{o}_n, z)p(z|\mathbf{o}_n)}, \quad (7)$$

Table 8: **BET Hyperparameter.** The ‘Value’ column indicates sweep values for the obstacle avoidance task, the block pushing task, the Franka kitchen task and the table tennis task (in this order).

PARAMETER	SWEEP	VALUE
TRANSFORMER BLOCKS	{2, 3, 4, 6}	3, 4, 6, 2
OFFSET LOSS SCALE	{1.0, 100.0, 1000.0}	1.0, 1.0, 1.0, 1.0
EMBEDDING WIDTH	{48, 72, 96, 120}	96, 72, 120, 48
NUMBER OF BINS	{8, 10, 16, 32, 50, 64}	50, 10, 64, 64
ATTENTION HEADS	{4, 6}	4, 4, 6, 4
CONTEXT SIZE	–	1
TRAINING EPOCHS	–	500
BATCH SIZE	–	512
LEARNING RATE	–	10^{-4}

669 for all samples n and components z . In the M-step, the gating distribution $p(z|\mathbf{o})$ is updated such
670 that the KL divergence between $q(z|\mathbf{o})$ and $p(z|\mathbf{o})$ is minimized. Using the properties of the KL
671 divergence, we obtain a global optimum by setting $p(z|\mathbf{o}_n) = q(z|\mathbf{o}_n)$ for all n and all z . This allows
672 us to rewrite Equation 7 using the recursion in q , giving

$$q(z|\mathbf{o}_n)^{(i+1)} = \frac{p_{\theta_z}(\mathbf{a}_n|\mathbf{o}_n, z)q(z|\mathbf{o}_n)^{(i)}}{\sum_z p_{\theta_z}(\mathbf{a}_n|\mathbf{o}_n, z)q(z|\mathbf{o}_n)^{(i)}},$$

673 where (i) denotes the iteration of the EM algorithm. The update for the variational distribution of the
674 IMC algorithm is given by

$$\begin{aligned} q(z|\mathcal{D}_n)^{(i+1)} &= \frac{\tilde{p}(\mathcal{D}_n|z)^{(i+1)}}{\sum_z \tilde{p}(\mathcal{D}_n|z)^{(i+1)}} \\ &= \frac{p_{\theta_z}(\mathbf{a}_n|\mathbf{o}_n, z)^{1/\eta}q(z|\mathcal{D}_n)^{(i)}}{\sum_z p_{\theta_z}(\mathbf{a}_n|\mathbf{o}_n, z)^{1/\eta}q(z|\mathcal{D}_n)^{(i)}}. \end{aligned}$$

675 Consequently, we see that $q(z|\mathbf{o}) = q(z|\mathcal{D})$ for $\eta = 1$. However, the two algorithms mainly differ in
676 the M-step for the experts: The EM algorithm uses the variational distribution to assign weights to
677 samples, i.e.

$$\max_{\theta_z} \sum_{n=1}^N q(z|\mathbf{o}_n) \log p_{\theta_z}(\mathbf{a}_n|\mathbf{o}_n, z),$$

678 whereas IMC uses the curricula as weights, that is,

$$\max_{\theta_z} \sum_{n=1}^N p(\mathcal{D}_n|z) \log p_{\theta_z}(\mathbf{a}_n|\mathbf{o}_n, z).$$

679 This subtle difference shows the properties of moment and information projection: In the EM
680 algorithm each sample \mathbf{o}_n contributes to the expert optimization as $\sum_z q(z|\mathbf{o}_n) = 1$. However, if all
681 curricula ignore the n th sample, it will not have impact on the expert optimization. Assuming that
682 the curricula ignore samples that the corresponding experts are not able to represent, IMC prevents
683 experts from having to average over ‘too hard’ samples. Furthermore, this results in reduced outlier
684 sensitivity as they are likely to be ignored for expert optimization. Lastly, we highlight the difference
685 between the gating optimization: Assuming that both algorithms train a gating network $g_\phi(z|\mathbf{o})$ we
686 have

$$\max_{\phi} \sum_n \sum_z q(z|\mathbf{o}_n) \log g_\phi(z|\mathbf{o}_n),$$

687 for the EM algorithm and

$$\max_{\phi} \sum_n \sum_z \tilde{p}(\mathcal{D}_n|z) \log g_\phi(z|\mathbf{o}_n),$$

688 for IMC. Similar to the expert optimization, EM includes all samples to fit the parameters of the
689 gating network, whereas IMC ignores samples where the unnormalized curriculum weights $\tilde{p}(\mathcal{D}_n|z)$
690 are zero for all components.

691 E Algorithm Details & Ablation Studies

692 E.1 Expert Design Choices

693 **Distribution.** In our mixture of experts policy, we employ Gaussian distributions with a fixed
 694 variance to represent the individual experts. This choice offers several benefits in terms of likelihood
 695 calculation, optimization and ease of sampling:

696 To perform the M-Step for the curricula (Section 3.3), exact log-likelihood computation is necessary.
 697 This computation becomes straightforward when using Gaussian distributions. Additionally, when
 698 Gaussian distributions with fixed variances are employed to represent the experts, the M-Step for the
 699 experts simplifies to a weighted squared-error minimization. Specifically, maximizing the weighted
 700 likelihood reduces to minimizing the weighted squared error between the predicted actions and the
 701 actual actions.

702 The optimization problem for the expert update can be formulated as follows:

$$\theta_z^* = \arg \max_{\theta_z} \sum_n \tilde{p}(\mathcal{D}_n|z) \log p_{\theta_z}(\mathbf{a}_n|\mathbf{o}_n, z), = \arg \min_{\theta_z} \sum_n \tilde{p}(\mathcal{D}_n|z) \|\mu_{\theta_z}(\mathbf{o}_n) - \mathbf{a}_n\|_2^2.$$

703 This optimization problem can be efficiently solved using gradient-based methods. Lastly, sampling
 704 from Gaussian distributions is well-known to be straightforward and efficient.

705 **Parameterization.** We experimented with two different parameterizations of the Gaussian expert
 706 means μ_{θ_z} , which we dub *single-head* and *multi-head*: For single-head, there is no parameter sharing
 707 between the different experts. Each expert has its own set of parameters θ_z . As a result, we learn N_z
 708 different multi-layer perceptrons (MLPs) $\mu_{\theta_z} : \mathbb{R}^{|\mathcal{O}|} \rightarrow \mathbb{R}^{|\mathcal{A}|}$, where N_z is the number of mixture
 709 components. In contrast, the multi-head parameterization uses a global set of parameters θ for all
 710 experts and hence allows for feature sharing. We thus learn a single MLP $\mu_{\theta} : \mathbb{R}^{|\mathcal{O}|} \rightarrow \mathbb{R}^{N_z \times |\mathcal{A}|}$.

711 To compare both parameterizations, we conducted an ablation study where we evaluate the MoE
 712 policy on obstacle avoidance, table tennis and Franka kitchen. In order to have a similar number
 713 of parameters, we used smaller MLPs for single-head, that is, 1 – 4 layers whereas for multi-head
 714 we used a 6 layer MLP. The results are shown in Table 9 and are generated using 30 components
 715 for the obstacle avoidance and table tennis task. The remaining hyperparameters are equal to the
 716 ones listed in the main manuscript. For Franka kitchen, we report the cumulative success rate and
 717 entropy for a different number of completed tasks. We report the mean and standard deviation
 718 calculated across 10 different seeds. Our findings indicate that, in the majority of experiments,
 719 the single-head parameterization outperforms the multi-head alternative. Notably, we observed a
 substantial performance disparity, especially in the case of Franka kitchen.

Table 9: **Expert Parameterization Ablation:** We compare IMC with single- and multi-head expert parameterization. For further details, please refer to the accompanying text.

ARCHITECTURE	OBSTACLE AVOIDANCE		TABLE TENNIS		FRANKA KITCHEN	
	SUCCESS RATE (↑)	ENTROPY (↑)	SUCCESS RATE (↑)	DISTANCE ERR. (↓)	SUCCESS RATE (↑)	ENTROPY (↑)
SINGLE-HEAD	0.899±0.035	0.887±0.043	0.812±0.039	0.168±0.007	3.644±0.230	6.189±1.135
MULTI-HEAD	0.855±0.053	0.930±0.031	0.870±0.017	0.153±0.007	3.248±0.062	4.657±0.312

720

721 **Expert Complexity.** We conducted an ablation study to evaluate the effect of expert complexity on
 722 the performance of the IMC algorithm. The study involved varying the number of hidden layers in
 723 the single-head expert architecture while assessing the IMC algorithm’s performance on the Franka
 724 kitchen task using the cumulative success rate and entropy. The results, presented in Figure 8, were
 725 obtained using the hyperparameters specified in the main manuscript. Mean and standard deviation
 726 were calculated across 5 different seeds. Our findings demonstrate a positive correlation between
 727 expert complexity and achieved performance.

728 E.2 Curriculum Pacing Sensitivity

729 To examine the algorithm’s sensitivity to the curriculum pacing parameter η , we conducted an ablation
 730 study. Figure 9 presents the results obtained using 30 components for the obstacle avoidance and

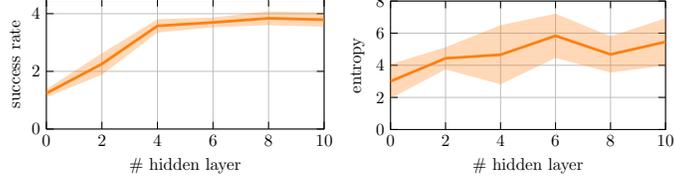


Figure 8: **Expert Complexity Ablation:** Evaluation of the IMC algorithm on the Franka kitchen task with varying numbers of hidden layers in the single-head expert architecture.

731 table tennis tasks, while maintaining the remaining hyperparameters as listed in the main manuscript.
 732 For the Franka kitchen task, we analyzed the cumulative success rate and entropy across varying
 733 numbers of completed tasks. The mean and standard deviation were calculated across 5 different
 734 seeds. Our findings reveal that the optimal value for η is dependent on the specific task. Nevertheless,
 the algorithm exhibits stable performance even when η values differ by an order of magnitude.

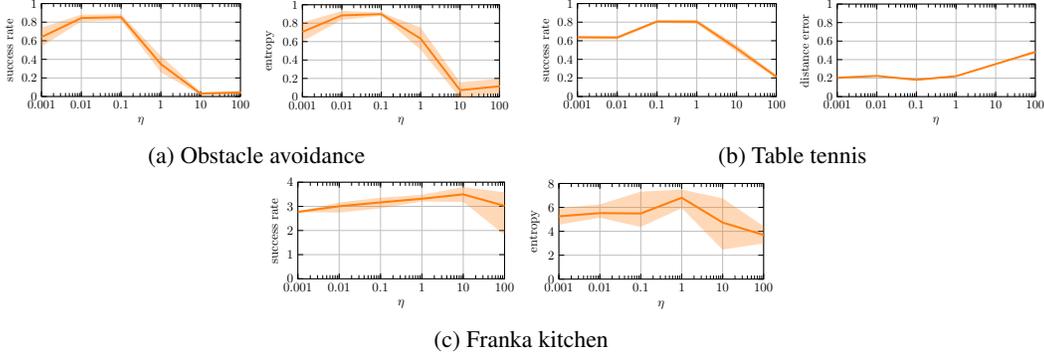


Figure 9: **Curriculum Pacing Sensitivity:** Sensitivity analysis of the IMC algorithm for performance metrics in the obstacle avoidance, table tennis, and Franka kitchen tasks, considering varying curriculum pacing (η) values. The results illustrate the mean and standard deviation across 5 different seeds.

735

736 E.3 Inference Details

737 We provide pseudocode to further clarify the inference procedure of our proposed method (see Algorithm 2).

Algorithm 2 IMC Action Generation

- 1: **Require:** Curriculum weights $\{\tilde{p}(\mathcal{D}_n|z) \mid n \in \{1, \dots, N\}, z \in \{1, \dots, N_z\}\}$
 - 2: **Require:** Expert parameter $\{\theta_z \mid z \in \{1, \dots, N_z\}\}$
 - 3: **Require:** New observation \mathbf{o}^*
 - 4: **if not** parameter_updated **then**
 - 5: $\phi^* \leftarrow \arg \max_{\phi} \sum_n \sum_z \tilde{p}(\mathcal{D}_n|z) \log g_{\phi}(z|\mathbf{o}_n)$
 - 6: parameter_updated \leftarrow True
 - 7: **end if**
 - 8: Sample $z' \sim g_{\phi^*}(z|\mathbf{o}^*)$
 - 9: Sample $\mathbf{a}' \sim p_{\theta_{z'}}(\mathbf{a}|\mathbf{o}^*, z')$
 - 10: **Return** \mathbf{a}'
-

738