

18 Imitation Learning

Previous chapters have assumed either that a reward function is known or that rewards are received while interacting with the environment. For some applications, it may be easier for an expert to demonstrate the desired behavior rather than specify a reward function. This chapter discusses algorithms for *imitation learning*, where the desired behavior is learned from expert demonstration. We will cover a variety of methods ranging from very simple likelihood-maximization methods to more complicated iterative methods that involve reinforcement learning.¹

18.1 Behavioral Cloning

A simple form of imitation learning is to treat it as a supervised learning problem. This method, called *behavioral cloning*,² trains a stochastic policy π_{θ} parameterized by θ to maximize the likelihood of actions from a data set \mathcal{D} of expert state-action pairs:

$$\underset{\theta}{\text{maximize}} \quad \prod_{(s,a) \in \mathcal{D}} \pi_{\theta}(a | s) \quad (18.1)$$

As done in earlier chapters, we can transform the maximization over the product over $\pi_{\theta}(a | s)$ to a sum over $\log \pi_{\theta}(a | s)$.

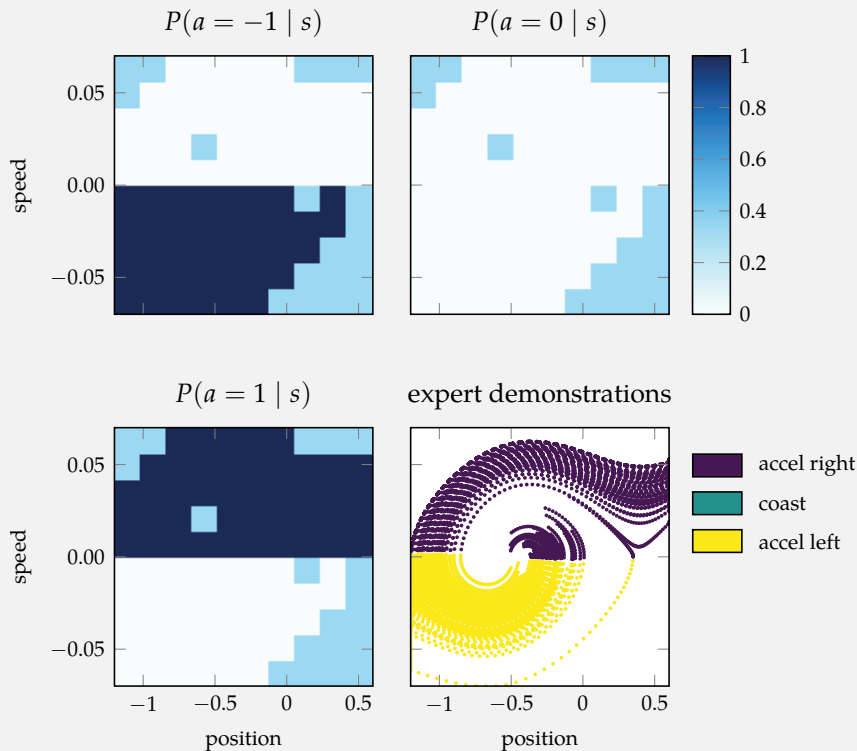
Depending on how we want to represent the conditional distribution $\pi_{\theta}(a | s)$, we may compute the maximum likelihood estimate of θ analytically. For example, if we use a discrete conditional model (section 2.4), θ would consist of the counts $N(s, a)$ from \mathcal{D} and $\pi_{\theta}(a | s) = N(s, a) / \sum_a N(s, a)$. Example 18.1 applies a discrete conditional model to data from the mountain car problem.

If we have a factored representation of our policy, we can use a Bayesian network to represent the joint distribution over our state and action variables. Figure 18.1 shows an example. We can learn both the structure (chapter 5) and

¹ Additional methods and applications are surveyed by A. Hussein, M. M. Gaber, E. Elyan, and C. Jayne, “Imitation Learning: A Survey of Learning Methods,” *ACM Computing Surveys*, vol. 50, no. 2, pp. 1–35, 2017.

² D. A. Pomerleau, “Efficient Training of Artificial Neural Networks for Autonomous Navigation,” *Neural Computation*, vol. 3, no. 1, pp. 88–97, 1991.

Consider using behavioral cloning on expert demonstrations for the mountain car problem (appendix F.4). We are given 10 rollouts from an expert policy. We fit a conditional distribution and plot the results. The continuous trajectories were discretized with 10 bins each for position and for speed.



Example 18.1. A demonstration of behavioral cloning applied to the mountain car problem. The light blue regions are areas without training data, resulting in poor policy performance when the agent encounters those states.

The state space is not fully covered by expert demonstrations, which is typical of imitation learning problems. The resulting policy may perform well when used in regions with coverage, but it assigns a uniform distribution to actions in regions without coverage. Even if we start in a region with coverage, we may transition to regions without coverage due to stochasticity in the environment.

the parameters (chapter 4) from the data \mathcal{D} . Given the current state, we can then infer the distribution over actions using one of the inference algorithms discussed earlier (chapter 3).




Figure 18.1. Bayesian networks can be used to represent a joint distribution over the state and action variables. We can apply an inference algorithm to generate a distribution over actions, given the current values of the state variables.

We can use many other representations for π_{θ} . For example, we might want to use a neural network, where the input corresponds to the values of the state variables and the output corresponds to parameters of a distribution over the action space. If our representation is differentiable, which is the case with neural networks, we can attempt to optimize equation (18.1) using gradient ascent. This approach is implemented in algorithm 18.1.

```

struct BehavioralCloning
     $\alpha$  # step size
    k_max # number of iterations
     $\nabla \log \pi$  # log likelihood gradient
end

function optimize(M::BehavioralCloning, D,  $\theta$ )
     $\alpha$ , k_max,  $\nabla \log \pi$  = M. $\alpha$ , M.k_max, M. $\nabla \log \pi$ 
    for k in 1:k_max
         $\nabla$  = mean( $\nabla \log \pi(\theta, a, s)$  for (s,a) in D)
         $\theta$  +=  $\alpha * \nabla$ 
    end
    return  $\theta$ 
end

```

Algorithm 18.1. A method for learning a parameterized stochastic policy from expert demonstrations in the form of a set of state-action tuples \mathcal{D} . The policy parameterization vector θ is iteratively improved by maximizing the log likelihood of the actions given the states. Behavioral cloning requires a step size α , an iteration count k_{\max} , and a log likelihood gradient $\nabla \log \pi$.

The closer the expert demonstrations are to optimal, the better the resulting behavioral cloning policy will perform.³ However, behavioral cloning suffers from *cascading errors*. As discussed in example 18.2, small inaccuracies compound during a rollout and eventually lead to states that are poorly represented in the training data, thereby leading to worse decisions, and ultimately to invalid or unseen situations. Although behavioral cloning is attractive due to its simplicity, cascading errors cause the method to perform poorly on many problems, especially when policies must be used for long time horizons.

³U. Syed and R.E. Schapire, “A Reduction from Apprenticeship Learning to Classification,” in *Advances in Neural Information Processing Systems (NIPS)*, 2010.

Consider applying behavioral cloning to train a policy for driving an autonomous race car. A human race car driver provides expert demonstrations. Being an expert, the driver never drifts onto the grass or too close to a railing. A model trained with behavioral cloning would have no information to use when near a railing or when drifting onto the grass, and thus it would not know how to recover.

Example 18.2. A brief example of the generalization issue inherent to behavioral cloning approaches.

18.2 Data Set Aggregation

One way to address the problem of cascading errors is to correct a trained policy using additional expert input. *Sequential interactive demonstration* methods alternate between collecting data from an expert in situations generated by a trained policy and using this data to improve this policy.

One type of sequential interactive demonstration method is called *data set aggregation* (*Dagger*) (algorithm 18.2).⁴ It starts by training a stochastic policy using behavioral cloning. The policy is then used to run several rollouts from an initial state distribution b , which are then given to an expert to provide the correct actions for each state. The new data is aggregated with the previous data set and a new policy is trained. Example 18.3 illustrates this process.

⁴S. Ross, G.J. Gordon, and J.A. Bagnell, “A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning,” in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, vol. 15, 2011.

These interactive demonstrations iteratively build a data set covering the regions of the state space that the agent is likely to encounter, based on previous learning iterations. With each iteration, newly added examples compose a smaller fraction of the data set, thereby leading to smaller policy changes. While sequential interactive demonstration can work well in practice, it is not guaranteed to converge. It can be shown that mixing in influence from the expert policy can guarantee convergence, which is the subject of the next section.

18.3 Stochastic Mixing Iterative Learning

Sequential interactive methods can also iteratively build up a policy by stochastically mixing in newly trained policies. One such method is *stochastic mixing iterative learning* (*SMILE*) (algorithm 18.3).⁵ It uses behavioral cloning in every iteration but mixes the newly trained policy with the previous ones.

⁵S. Ross and J.A. Bagnell, “Efficient Reductions for Imitation Learning,” in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2010.

We start with the expert policy, $\pi^{(1)} = \pi_E$.⁶ In each iteration, we execute the

⁶We do not have an explicit representation of π_E . Evaluating π_E requires interactively querying the expert, as done in the previous section.

```

struct DataSetAggregation
     $\mathcal{P}$  # problem with unknown reward function
    bc # behavioral cloning struct
    k_max # number of iterations
    m # number of rollouts per iteration
    d # rollout depth
    b # initial state distribution
     $\pi_E$  # expert
     $\pi_\theta$  # parameterized policy
end

function optimize(M::DataSetAggregation, D,  $\theta$ )
     $\mathcal{P}$ , bc, k_max, m = M. $\mathcal{P}$ , M.bc, M.k_max, M.m
    d, b,  $\pi_E$ ,  $\pi_\theta$  = M.d, M.b, M. $\pi_E$ , M. $\pi_\theta$ 
     $\theta$  = optimize(bc, D,  $\theta$ )
    for k in 2:k_max
        for i in 1:m
            s = rand(b)
            for j in 1:d
                push!(D, (s,  $\pi_E$ (s)))
                a = rand( $\pi_\theta$ ( $\theta$ , s))
                s = rand( $\mathcal{P}$ .T(s, a))
            end
        end
         $\theta$  = optimize(bc, D,  $\theta$ )
    end
    return  $\theta$ 
end

```

Algorithm 18.2. The DAgger method of data set aggregation for learning a stochastic parameterized policy from expert demonstrations. This method takes an initial data set of state-action tuples D , a stochastic parameterized policy $\pi_\theta(\theta, s)$, an MDP \mathcal{P} that defines a transition function, and an initial state distribution b . Behavioral cloning (algorithm 18.1) is used in each iteration to improve the policy.

An expert policy π_E labels trajectories sampled from the latest learned policy to augment the data set. The original paper generated trajectories by stochastically mixing in the expert policy. This implementation is thus the original DAgger with an extreme mixing value of zero.

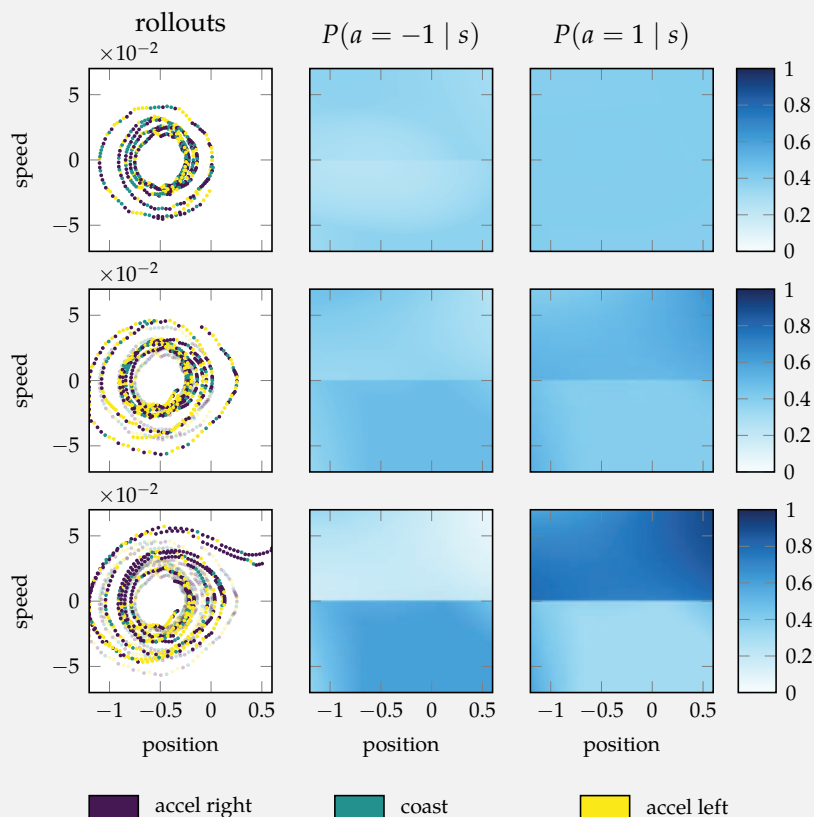
In practice, an expert policy would not exist, and calls to this policy would be replaced with queries to a human expert.

Consider using DAgger to train a policy on the mountain car problem where the reward is not observed. We use an expert policy that accelerates in the direction of travel. In this example, we train a policy using the following features:

$$\mathbf{f}(s) = [1[v > 0], 1[v < 0], x, x^2, v, v^2, xv]$$

where x and v are the position and speed of the car.

Example 18.3. DAgger applied to the mountain car problem, with iterations running from top to bottom. Trajectories accumulate in the data set over time. The behavior of the agent improves with each iteration.



Trajectories are colored according to the action. In the first iteration, the agent behaves randomly, unable to make progress toward the goal ($x \geq 0.6$). With additional iterations, the agent learns to mimic the expert policy of accelerating in the direction of travel. This behavior is apparent in the new trajectories, which spiral outward, and the policy, which assigns high likelihood to $a = 1$ when $v > 0$ and $a = -1$ when $v < 0$.

latest policy $\pi^{(k)}$ to generate a new data set, querying the expert to provide the correct actions. Behavioral cloning is applied only to this new data set to train a new *component policy* $\hat{\pi}^{(k)}$. This component policy is mixed with component policies from the previous iterations to produce a new policy $\pi^{(k+1)}$.

The mixing of component policies to generate $\pi^{(k+1)}$ is governed by a mixing scalar $\beta \in (0, 1)$. The probability of acting according to the expert policy is $(1 - \beta)^k$, and the probability of acting according to $\hat{\pi}^{(i)}$ is $\beta(1 - \beta)^{i-1}$. This scheme assigns more weight to older policies under the hypothesis that older policy components were trained on the states most likely to be encountered.⁷ With each iteration, the probability of acting according to the original expert policy decays to zero. The mixing scalar is typically small, such that the agent does not abandon the expert's policy too quickly. Example 18.4 demonstrates this approach with the mountain car problem.

⁷ In SMILe, we are acting according to our latest learned policy. We expect that this learned policy will match the expert fairly well and primarily mispredict when we deviate from the expert policy. The learned component policies generally only need to make smaller and smaller contributions with each iteration to make up the difference in what has not already been learned.

18.4 Maximum Margin Inverse Reinforcement Learning

In many application settings, we have no expert that can be interactively queried; but instead we have a batch of expert demonstration trajectories. We will assume that the expert demonstration data \mathcal{D} consists of m trajectories. Each trajectory τ in \mathcal{D} involves a rollout to depth d . In *inverse reinforcement learning*, we assume that the expert is optimizing an unknown reward function. From \mathcal{D} , we attempt to derive that reward function. With that reward function, we can use the methods discussed in prior chapters to derive an optimal policy.

There are different approaches to inverse reinforcement learning. We generally need to define a parameterization of the reward function. A common assumption is that this parameterization is linear, with $R_{\Phi}(s, a) = \Phi^{\top} \beta(s, a)$, where $\beta(s, a)$ is a feature vector and Φ is a vector of weightings. In this section, we will focus on an approach known as *maximum margin inverse reinforcement learning*,⁸ where the features are assumed to be binary. Since optimal policies remain optimal with positive scaling of the reward function, this method additionally constrains the weight vector such that $\|\Phi\|_2 \leq 1$. The expert data activates each binary feature with different frequencies, perhaps pursuing some and avoiding others. This approach attempts to learn this pattern of activation and trains an agent to mimic these activation frequencies.

⁸ P. Abbeel and A. Y. Ng, "Apprenticeship Learning via Inverse Reinforcement Learning," in *International Conference on Machine Learning (ICML)*, 2004.

```

struct SMILE
     $\mathcal{P}$  # problem with unknown reward
    bc # Behavioral cloning struct
    k_max # number of iterations
    m # number of rollouts per iteration
    d # rollout depth
    b # initial state distribution
     $\beta$  # mixing scalar (e.g.,  $d^{-3}$ )
     $\pi_E$  # expert policy
     $\pi_\theta$  # parameterized policy
end


function optimize(M::SMILE,  $\theta$ )
     $\mathcal{P}$ , bc, k_max, m = M. $\mathcal{P}$ , M.bc, M.k_max, M.m
    d, b,  $\beta$ ,  $\pi_E$ ,  $\pi_\theta$  = M.d, M.b, M. $\beta$ , M. $\pi_E$ , M. $\pi_\theta$ 
     $\mathcal{A}$ , T =  $\mathcal{P}.$  $\mathcal{A}$ ,  $\mathcal{P}.$ T
     $\theta_s$  = []
     $\pi = s \rightarrow \pi_E(s)$ 
    for k in 1:k_max
        # execute latest  $\pi$  to get new data set D
        D = []
        for i in 1:m
            s = rand(b)
            for j in 1:d
                push!(D, (s,  $\pi_E(s)$ ))
                a =  $\pi(s)$ 
                s = rand(T(s, a))
            end
        end
        # train new policy classifier
         $\theta = \text{optimize}(bc, D, \theta)$ 
        push!( $\theta_s$ ,  $\theta$ )
        # compute a new policy mixture
        P $\pi$  = Categorical(normalize([(1- $\beta$ )(i-1) for i in 1:k],1))
         $\pi = s \rightarrow \text{begin}$ 
            if rand() < (1- $\beta$ )(k-1)
                return  $\pi_E(s)$ 
            else
                return rand(Categorical( $\pi_\theta(\theta_s[\text{rand}(P\pi)], s)$ ))
            end
        end
    end
    P $s$  = normalize([(1- $\beta$ )(i-1) for i in 1:k_max],1)
    return P $s$ ,  $\theta_s$ 
end

```

Algorithm 18.3. The SMILE algorithm for training a stochastic parameterized policy from expert demonstrations for an MDP \mathcal{P} . It successively mixes in new component policies with smaller and smaller weights, while simultaneously reducing the probability of acting according to the expert policy. The method returns the probabilities P_s and parameterizations θ_s for the component policies.

Consider using SMILe to train a policy on the mountain car problem where the reward is not observed. We use the same features that were used for DAGger in example 18.3. Both DAGger and SMILe receive a new expert-labeled data set with each iteration. Instead of accumulating a larger data set of expert-labeled data, SMILe trains a new policy component using only the most recent data, mixing the new policy component with the previous policy components.

Example 18.4. Using SMILe to learn a policy for the mountain car problem. In contrast with DAGger in example 18.3, SMILe mixes the expert into the policy during rollouts. This expert component, whose influence wanes with each iteration, causes the initial rollouts to better progress toward the goal.



An important part of this algorithm involves reasoning about the expected return under a policy π for a weighting Φ and initial state distribution b :

$$\mathbb{E}_{s \sim b} [U(s)] = \mathbb{E}_{\tau} \left[\sum_{k=1}^d \gamma^{k-1} R_{\Phi}(s^{(k)}, a^{(k)}) \right] \quad (18.2)$$

$$= \mathbb{E}_{\tau} \left[\sum_{k=1}^d \gamma^{k-1} \Phi^{\top} \beta(s^{(k)}, a^{(k)}) \right] \quad (18.3)$$

$$= \Phi^{\top} \left(\mathbb{E}_{\tau} \left[\sum_{k=1}^d \gamma^{k-1} \beta(s^{(k)}, a^{(k)}) \right] \right) \quad (18.4)$$

$$= \Phi^{\top} \mu_{\pi} \quad (18.5)$$

where τ corresponds to trajectories generated by π to depth d . Here, we introduce the *feature expectations* vector μ_{π} , which is the expected discounted accumulated feature values. These feature expectations can be estimated from m rollouts, as implemented in algorithm 18.4.

```

struct InverseReinforcementLearning
   $\mathcal{P}$  # problem
   $b$  # initial state distribution
   $d$  # depth
   $m$  # number of samples
   $\pi$  # parameterized policy
   $\beta$  # binary feature mapping
   $\mu_E$  # expert feature expectations
  RL # reinforcement learning method
   $\epsilon$  # tolerance
end

function feature_expectations(M::InverseReinforcementLearning,  $\pi$ )
   $\mathcal{P}, b, m, d, \beta, \gamma = M.\mathcal{P}, M.b, M.m, M.d, M.\beta, M.\gamma$ 
   $\mu(\tau) = \text{sum}(\gamma^{k-1} * \beta(s, a) \text{ for } (k, (s, a)) \text{ in enumerate}(\tau))$ 
   $\tau_s = [\text{simulate}(\mathcal{P}, \text{rand}(b), \pi, d) \text{ for } i \text{ in } 1:m]$ 
  return mean( $\mu(\tau)$  for  $\tau$  in  $\tau_s$ )
end

```

Algorithm 18.4. A structure for inverse reinforcement learning and a method for estimating a feature expectations vector from rollouts.

We can use the expert demonstrations to estimate the expert feature expectations μ_E , and we want to find a policy that matches these feature expectations as closely as possible. At the first iteration, we begin with a randomized policy $\pi^{(1)}$ and estimate its feature expectations, denoted as $\mu^{(1)}$. At iteration k , we find a new $\Phi^{(k)}$ corresponding to a reward function $R_{\Phi^{(k)}}(s, a) = \Phi^{(k)\top} \beta(s, a)$, such

that the expert outperforms all previously found policies by the greatest *margin* t :

$$\begin{aligned} & \underset{t, \Phi}{\text{maximize}} && t \\ & \text{subject to} && \Phi^\top \mu_E \geq \Phi^\top \mu^{(i)} + t \text{ for } i = 1, \dots, k-1 \\ & && \|\Phi\|_2 \leq 1 \end{aligned} \quad (18.6)$$

Equation (18.6) is a quadratic program that can be easily solved. We then solve for a new policy $\pi^{(k)}$ using the reward function $R(s, a) = \Phi^{(k)\top} \beta(s, a)$, and produce a new vector of feature expectations. Figure 18.2 illustrates this margin maximization process.

We iterate until the margin is sufficiently small, with $t \leq \epsilon$. At convergence, we can solve for a mixed policy that attempts to have feature expectations as close as possible to that of the expert policy:

$$\begin{aligned} & \underset{\lambda}{\text{minimize}} && \|\mu_E - \mu_\lambda\|_2 \\ & \text{subject to} && \lambda \geq 0 \\ & && \|\lambda\|_1 = 1 \end{aligned} \quad (18.7)$$

where $\mu_\lambda = \sum_i \lambda_i \mu^{(i)}$. The mixture weights λ combine the policies found at each iteration. With probability λ_i , we follow policy $\pi^{(i)}$. Maximum margin inverse reinforcement learning is implemented in algorithm 18.5.

18.5 Maximum Entropy Inverse Reinforcement Learning

The inverse reinforcement learning approach from the previous section is underspecified, meaning that there are often multiple policies that can produce the same feature expectations as the expert demonstrations. This section introduces *maximum entropy inverse reinforcement learning*, which avoids this ambiguity by preferring the policy that results in the distribution over trajectories that has maximum *entropy* (appendix A.8).⁹ The problem can be transformed into one of finding the best reward function parameters Φ in a maximum likelihood estimation problem, given the expert data \mathcal{D} .

⁹ B. D. Ziebart, A. Maas, J. A. Bagnell, and A. K. Dey, "Maximum Entropy Inverse Reinforcement Learning," in *AAAI Conference on Artificial Intelligence (AAAI)*, 2008.




Figure 18.2. A geometric visualization of three example iterations of the maximum-margin inverse reinforcement learning algorithm, going top to bottom. In each iteration, the new weight vector points in the direction perpendicular to the hyperplane that separates the expert feature expectation vector from that of the previous policy with the largest possible margin. The margin decreases with each iteration.

```

function calc_weighting(M::InverseReinforcementLearning, μs)
    μE = M.μE
    k = length(μE)
    model = Model(Ipopt.Optimizer)
    @variable(model, t)
    @variable(model, φ[1:k] ≥ 0)
    @objective(model, Max, t)
    for μ in μs
        @constraint(model, φ·μE ≥ φ·μ + t)
    end
    @constraint(model, φ·φ ≤ 1)
    optimize!(model)
    return (value(t), value.(φ))
end

function calc_policy_mixture(M::InverseReinforcementLearning, μs)
    μE = M.μE
    k = length(μs)
    model = Model(Ipopt.Optimizer)
    @variable(model, λ[1:k] ≥ 0)
    @objective(model, Min, (μE - sum(λ[i]*μs[i] for i in 1:k))·
        (μE - sum(λ[i]*μs[i] for i in 1:k)))
    @constraint(model, sum(λ) == 1)
    optimize!(model)
    return value.(λ)
end

function optimize(M::InverseReinforcementLearning, θ)
    π, ε, RL = M.π, M.ε, M.RL
    θs = [θ]
    μs = [feature_expectations(M, s→π(θ,s))]
    while true
        t, φ = calc_weighting(M, μs)
        if t ≤ ε
            break
        end
        copyto!(RL.φ, φ) # R(s,a) = φ·β(s,a)
        θ = optimize(RL, π, θ)
        push!(θs, θ)
        push!(μs, feature_expectations(M, s→π(θ,s)))
    end
    λ = calc_policy_mixture(M, μs)
    return λ, θs
end

```

Algorithm 18.5. Maximum margin inverse reinforcement learning, which computes a mixed policy whose feature expectations match those of given expert demonstrations. We use JuMP.jl to solve constrained optimization problems. This implementation requires that the provided reinforcement learning struct has a weight vector ϕ that can be updated with new values. The method returns the stochastic weightings λ and parameterizations θ_s for the component policies.

Any policy π induces a distribution over trajectories¹⁰ $P_\pi(\tau)$. Different policies produce different trajectory distributions. We are free to choose any of these distributions over trajectories that match the expert feature expectations. The *principle of maximum entropy* chooses the least informative distribution, which corresponds to the one with maximum entropy.¹¹ It can be shown that the least informative trajectory distribution takes the following form:

$$P_\Phi(\tau) = \frac{1}{Z(\Phi)} \exp(R_\Phi(\tau)) \quad (18.8)$$

where $P_\Phi(\tau)$ is the likelihood of a trajectory τ given reward parameter Φ , and

$$R_\Phi(\tau) = \sum_{k=1}^d \gamma^{k-1} R_\Phi(s^{(k)}, a^{(k)}) \quad (18.9)$$

is the discounted trajectory reward. We make no assumption on the parameterization of $R_\Phi(s^{(k)}, a^{(k)})$ other than that it is differentiable, allowing representations such as neural networks. The normalization scalar $Z(\Phi)$ ensures that the probabilities sum to 1:

$$Z(\Phi) = \sum_{\tau} \exp(R_\Phi(\tau)) \quad (18.10)$$

The summation is over all possible trajectories.

We have chosen a particular class of trajectory distributions for our policy. We now fit that class to our trajectories using maximum likelihood to obtain the parameters that best describe our data:

$$\max_{\Phi} f(\Phi) = \max_{\Phi} \sum_{\tau \in \mathcal{D}} \log P_\Phi(\tau) \quad (18.11)$$

We can rewrite the objective function $f(\Phi)$ from equation (18.11):

$$f(\Phi) = \sum_{\tau \in \mathcal{D}} \log \frac{1}{Z(\Phi)} \exp(R_\Phi(\tau)) \quad (18.12)$$

$$= \left(\sum_{\tau \in \mathcal{D}} R_\Phi(\tau) \right) - |\mathcal{D}| \log Z(\Phi) \quad (18.13)$$

$$= \left(\sum_{\tau \in \mathcal{D}} R_\Phi(\tau) \right) - |\mathcal{D}| \log \sum_{\tau} \exp(R_\Phi(\tau)) \quad (18.14)$$

¹⁰ For simplicity, this section assumes a finite horizon and that the state and action spaces are discrete, making $P_\Phi(\tau)$ a probability mass. To extend maximum entropy inverse reinforcement learning both to problems with continuous state and action spaces where the dynamics may be unknown, consider guided cost learning. C. Finn, S. Levine, and P. Abbeel, "Guided Cost Learning: Deep Inverse Optimal Control via Policy Optimization," in *International Conference on Machine Learning (ICML)*, 2016.

¹¹ For an introduction to this principle, see E. T. Jaynes, "Information Theory and Statistical Mechanics," *Physical Review*, vol. 106, no. 4, pp. 620–630, 1957.

We can attempt to optimize this objective function through gradient ascent. The gradient of f is

$$\nabla_{\phi} f = \left(\sum_{\tau \in \mathcal{D}} \nabla_{\phi} R_{\phi}(\tau) \right) - \frac{|\mathcal{D}|}{\sum_{\tau} \exp(R_{\phi}(\tau))} \sum_{\tau} \exp(R_{\phi}(\tau)) \nabla_{\phi} R_{\phi}(\tau) \quad (18.15)$$

$$= \left(\sum_{\tau \in \mathcal{D}} \nabla_{\phi} R_{\phi}(\tau) \right) - |\mathcal{D}| \sum_{\tau} P_{\phi}(\tau) \nabla_{\phi} R_{\phi}(\tau) \quad (18.16)$$

$$= \left(\sum_{\tau \in \mathcal{D}} \nabla_{\phi} R_{\phi}(\tau) \right) - |\mathcal{D}| \sum_s b_{\gamma, \phi}(s) \sum_a \pi_{\phi}(a | s) \nabla_{\phi} R_{\phi}(s, a) \quad (18.17)$$

If the reward function is linear, with $R_{\phi}(s, a) = \boldsymbol{\Phi}^{\top} \boldsymbol{\beta}(s, a)$, as in the previous section, then $\nabla_{\phi} R_{\phi}(s, a)$ is simply $\boldsymbol{\beta}(s, a)$.

Updating the parameter vector $\boldsymbol{\phi}$ thus requires both the discounted state visitation frequency $b_{\gamma, \phi}$ and the optimal policy under the current parameter vector, $\pi_{\phi}(a | s)$. We can obtain the optimal policy by running reinforcement learning. To compute the discounted state visitation frequencies, we can use rollouts or take a dynamic programming approach.

If we take a dynamic programming approach to compute the discounted state visitation frequencies, we can start with the initial state distribution $b_{\gamma, \phi}^{(1)} = b(s)$ and iteratively work forward in time:

$$b_{\gamma, \phi}^{(k+1)}(s) = \gamma \sum_a \sum_{s'} b_{\gamma, \phi}^{(k)}(s') \pi(a | s') T(s | s', a) \quad (18.18)$$

This version of maximum entropy inverse reinforcement learning is implemented in algorithm 18.6.

18.6 Generative Adversarial Imitation Learning

In *generative adversarial imitation learning (GAIL)*,¹² we optimize a differentiable parameterized policy π_{θ} , often represented by a neural network. Rather than provide a reward function, we use *adversarial learning* (appendix D.7). We also train a *discriminator* $C_{\phi}(s, a)$, typically also a neural network, to return the probability that it assigns to the state-action pair coming from the learned policy. The process involves alternating between training this discriminator to become better

¹²J. Ho and S. Ermon, “Generative Adversarial Imitation Learning,” in *Advances in Neural Information Processing Systems (NIPS)*, 2016.

```

struct MaximumEntropyIRL
     $\mathcal{P}$  # problem
     $\mathbf{b}$  # initial state distribution
     $d$  # depth
     $\pi$  # parameterized policy  $\pi(\theta, s)$ 
     $P\pi$  # parameterized policy likelihood  $\pi(\theta, a, s)$ 
     $\nabla R$  # reward function gradient
     $RL$  # reinforcement learning method
     $\alpha$  # step size
     $k\_max$  # number of iterations
end

function discounted_state_visitations( $M::MaximumEntropyIRL, \theta$ )
     $\mathcal{P}, \mathbf{b}, d, P\pi = M.\mathcal{P}, M.\mathbf{b}, M.d, M.P\pi$ 
     $S, \mathcal{A}, T, \gamma = \mathcal{P}.S, \mathcal{P}.\mathcal{A}, \mathcal{P}.T, \mathcal{P}.\gamma$ 
     $\mathbf{b\_sk} = \text{zeros}(\text{length}(\mathcal{P}.S), d)$ 
     $\mathbf{b\_sk}[:,1] = [\text{pdf}(\mathbf{b}, s) \text{ for } s \text{ in } S]$ 
    for  $k$  in  $2:d$ 
        for  $(s_i', s')$  in  $\text{enumerate}(S)$ 
             $\mathbf{b\_sk}[s_i', k] = \gamma * \text{sum}(\text{sum}(\mathbf{b\_sk}[s_i, k-1] * P\pi(\theta, a, s) * T(s, a, s'))$ 
                for  $(s_i, s)$  in  $\text{enumerate}(S)$ 
                    for  $a$  in  $\mathcal{A}$ )
            end
        end
    end
    return  $\text{normalize}!(\text{vec}(\text{mean}(\mathbf{b\_sk}, \text{dims}=2)), 1)$ 
end

function optimize( $M::MaximumEntropyIRL, D, \phi, \theta$ )
     $\mathcal{P}, \pi, P\pi, \nabla R, RL, \alpha, k\_max = M.\mathcal{P}, M.\pi, M.P\pi, M.\nabla R, M.RL, M.\alpha, M.k\_max$ 
     $S, \mathcal{A}, \gamma, nD = \mathcal{P}.S, \mathcal{P}.\mathcal{A}, \mathcal{P}.\gamma, \text{length}(D)$ 
    for  $k$  in  $1:k\_max$ 
         $\text{copyto}!(RL.\phi, \phi)$  # update parameters
         $\theta = \text{optimize}(RL, \pi, \theta)$ 
         $\mathbf{b} = \text{discounted\_state\_visitations}(M, \theta)$ 
         $\nabla R\tau = \tau \rightarrow \text{sum}(\gamma^{i-1} * \nabla R(\phi, s, a) \text{ for } (i, (s, a)) \text{ in } \text{enumerate}(\tau))$ 
         $\nabla f = \text{sum}(\nabla R\tau(\tau) \text{ for } \tau \text{ in } D) - nD * \text{sum}(\mathbf{b}[s_i] * \text{sum}(P\pi(\theta, a, s) * \nabla R(\phi, s, a)$ 
            for  $(a_i, a)$  in  $\text{enumerate}(\mathcal{A})$ )
            for  $(s_i, s)$  in  $\text{enumerate}(S)$ )
         $\phi += \alpha * \nabla f$ 
    end
    return  $\phi, \theta$ 
end

```

Algorithm 18.6. Maximum entropy inverse reinforcement learning, which finds a stochastic policy that maximizes the likelihood of the expert demonstrations under a maximum-entropy trajectory distribution. This implementation computes the expected visitations using dynamic programming over all states, which requires that the problem be discrete.

at distinguishing between simulated and expert state-action pairs, and training the policy to look indistinguishable from the expert demonstrations. The process is sketched in figure 18.3.

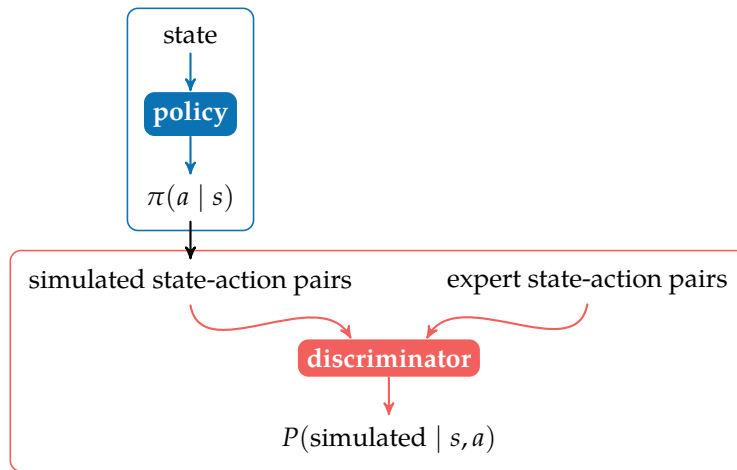


Figure 18.3. Instead of inferring a reward function, generative adversarial imitation learning optimizes a discriminator to distinguish between simulated and expert state-action pairs, and it optimizes a policy to appear indistinguishable to the discriminator. The aim is to eventually produce a policy that resembles the expert.

The discriminator and policy have opposing objectives. GAIL seeks to find a saddle point (θ, ϕ) of the negative log loss of the discriminator’s binary classification problem:¹³

$$\max_{\phi} \min_{\theta} \mathbb{E}_{(s,a) \sim \pi_{\theta}} [\log(C_{\phi}(s, a))] + \mathbb{E}_{(s,a) \sim \mathcal{D}} [\log(1 - C_{\phi}(s, a))] \quad (18.19)$$

where we use $(s, a) \sim \mathcal{D}$ to represent samples from the distribution represented by the expert data set \mathcal{D} . We can alternate between gradient ascent on ϕ to increase the objective and trust region policy optimization (section 12.4) on θ to reduce the objective, generating the necessary trajectory samples from the policy to conduct each of these steps. The discriminator provides a learning signal to the policy similar to the way that a reward signal would if it were known.

¹³ The original paper also includes the following entropy term:

$$-\lambda \mathbb{E}_{(s,a) \sim \mathcal{D}} [-\log \pi_{\theta}(a | s)]$$

18.7 Summary

- Imitation learning involves learning the desired behavior from expert demonstration without the use of a reward function.

- One type of imitation learning is behavioral cloning, which produces a stochastic policy that maximizes the conditional likelihood of the actions in the data set.
- When an expert can be queried multiple times, we can use iterative approaches like data set aggregation or stochastic mixing iterative learning.
- Inverse reinforcement learning involves inferring a reward function from expert data and then using traditional methods for finding an optimal policy.
- Maximum margin inverse reinforcement learning attempts to find a policy that matches the frequency of binary features found in the expert data set.
- Maximum entropy inverse reinforcement learning frames the problem of finding the best reward parameter as a maximum likelihood estimation problem, which it tries to solve using gradient ascent.
- Generative adversarial imitation learning iteratively optimizes a discriminator and a policy; the discriminator tries to discriminate between decisions made by the policy and decisions made by the expert, and the policy attempts to deceive the discriminator.

18.8 Exercises

Exercise 18.1. Consider applying behavioral cloning to a discrete problem where we have been given expert demonstrations. We could define a feature function $\beta(s)$ and represent the policy with a softmax distribution:

$$\pi(a | s) \propto \exp(\theta_a^\top \beta(s))$$

We would then learn the parameters θ_a for each action from the expert data. Why might we want to use this approach over one where we directly estimate a discrete distribution for each state, with one parameter per state-action pair?

Solution: In imitation learning, we are generally limited to a relatively small set of expert demonstrations. The distribution $P(a | s)$ has $(|\mathcal{A}| - 1)|\mathcal{S}|$ independent parameters that must be learned, which is often prohibitively large. Expert demonstrations typically cover only a small portion of the state space. Even if $P(a | s)$ can be reliably trained for the states covered in the provided data set, the resulting policy would be untrained in other states. Using a feature function allows generalization to unseen states.

Exercise 18.2. Section 18.1 suggested using a maximum likelihood approach for training a policy from expert data. This approach attempts to find the parameters of the policy that maximizes the likelihood assigned to the training examples. In some problems, however, we know that assigning high probability to one incorrect action is not as bad as assigning high probability to another incorrect action. For example, predicting an acceleration of -1 in the mountain car problem when the expert dictates an acceleration of 1 is worse than predicting an acceleration of 0 . How might behavioral cloning be modified to allow different penalties to be given to different misclassifications?

Solution: We can instead supply a cost function $C(s, a_{\text{true}}, a_{\text{pred}})$ that defines the cost of predicting action a_{pred} for state s when the expert's action is a_{true} . For example, with the mountain car problem, we might use

$$C(s, a_{\text{true}}, a_{\text{pred}}) = -|a_{\text{true}} - a_{\text{pred}}|$$

which penalizes greater deviations more than smaller deviations. The cost associated with the expert's action is typically zero.

If we have a stochastic policy $\pi(a | s)$, we then seek to minimize the cost over our data set:

$$\underset{\theta}{\text{minimize}} \quad \sum_{(s, a_{\text{true}}) \in \mathcal{D}} \sum_{a_{\text{pred}}} C(s, a_{\text{true}}, a_{\text{pred}}) \pi(a_{\text{pred}} | s)$$

This technique is called *cost-sensitive classification*.¹⁴ One benefit of cost-sensitive classification is that we can use a wide variety of off-the-shelf classification models, such as k -nearest neighbors, support vector machines, or decision trees, to train a policy.

¹⁴C. Elkan, "The Foundations of Cost-Sensitive Learning," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2001.

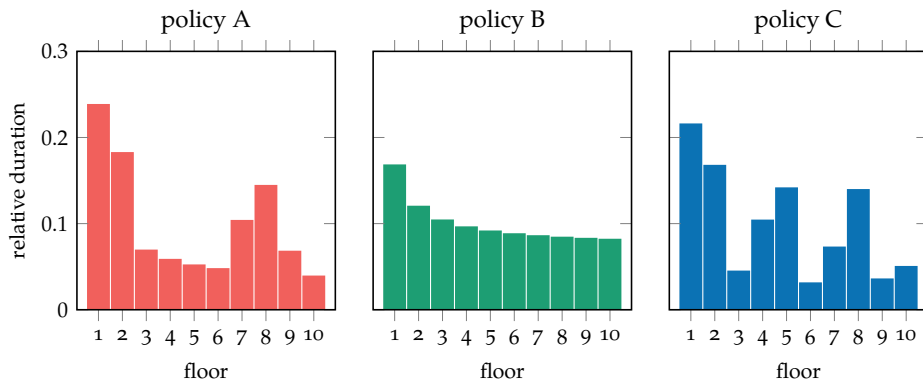
Exercise 18.3. Provide an example of where maximum margin inverse reinforcement learning does not uniquely define an optimal policy.

Solution: Maximum margin inverse reinforcement learning extracts binary features from the expert data and seeks a reward function whose optimal policy produces trajectories with the same frequencies of these binary features. There is no guarantee that multiple policies do not produce the same feature expectations. For example, an autonomous car that makes only left lane changes could have the same lane change frequencies as an autonomous car that makes only right lane changes.

Exercise 18.4. Maximum margin inverse reinforcement learning measures how similar a policy is to expert demonstrations using feature expectations. How is this similarity measure affected if nonbinary features are used?

Solution: If we use nonbinary features, then it is possible that some features can get larger than others, incentivizing the agent to match those features rather than those that tend to be smaller. Scale is not the only issue. Even if all features are constrained to lie within $[0, 1]$, then a policy that consistently produces $\phi(s, a)_1 = 0.5$ will have the same feature expectations as one that produces $\phi(s, a)_1 = 0$ half the time and $\phi(s, a)_1 = 1$ half the time. Depending on what the feature encodes, this can result in very different policies. Any set of continuous features can be discretized, and thus approximated by a set of binary features.

Exercise 18.5. Suppose we are building a system in a high-rise that must choose which floor to send an elevator. We have trained several policies to match the feature expectations of expert demonstrations, such as how long customers must wait for an elevator or how long they have to wait to get to their destinations. We run multiple rollouts for each policy and plot the relative duration spent on each floor. Which policy should we prefer according to the principle of maximum entropy, assuming that each policy matches the feature expectations equally?



Solution: These distributions over relative duration are analogous to distributions over trajectories for this elevator problem. In applying the principle of maximum entropy, we prefer the distribution with most entropy. Hence, we would choose policy B, which, in being most uniform, has the greatest entropy.

Exercise 18.6. Consider the policy optimization step in generative adversarial imitation learning. Rewrite the objective in the form of a reward function so that traditional reinforcement learning techniques can be applied.

Solution: We rewrite equation (18.19), dropping the terms dependent on the expert data set, and flip the sign to change from minimization over θ to a maximization over θ of the reward, producing the surrogate reward function:

$$\tilde{R}_\phi(s, a) = -\log C_\phi(s, a)$$

Although $\tilde{R}_\phi(s, a)$ may be quite different from the unknown true reward function, it can be used to drive the learned policy into regions of the state-action space similar to those covered by the expert.

Exercise 18.7. Explain how generative adversarial imitation learning could be changed such that the discriminator takes in trajectories rather than state-action pairs. Why might this be useful?

Solution: Changing generative adversarial imitation learning such that the discriminator takes trajectories is straightforward, especially if the trajectories are of fixed length. The expert data set is split into trajectories, and the learned policy is used to produce trajectories, just as it was before. Rather than operating on state-action pairs, the discriminator takes in trajectories using a representation such as a recurrent neural network (appendix D.5) and produces a classification probability. The objective function remains largely unchanged:

$$\max_{\phi} \min_{\theta} \mathbb{E}_{\tau \sim \pi_{\theta}} [\log(C_{\phi}(\tau))] + \mathbb{E}_{\tau \sim \mathcal{D}} [\log(1 - C_{\phi}(\tau))]$$

The advantage of running the discriminator over entire trajectories is that it can help the discriminator capture features that are not apparent from individual state-action pairs, which can result in better policies. For example, when looking at individual accelerations and turn rates for an autonomous driving policy, there is very little for a discriminator to learn. A discriminator trained to look at longer trajectories can see more of the vehicle's behavior, such as lane change aggressiveness and smoothness, to better match expert driving demonstrations.¹⁵

¹⁵This approach was used in A. Kuefler, J. Morton, T. A. Wheeler, and M. J. Kochenderfer, "Imitating Driver Behavior with Generative Adversarial Networks," in *IEEE Intelligent Vehicles Symposium (IV)*, 2017.