
Smooth Imitation Learning for Online Sequence Prediction

Hoang M. Le

Andrew Kang

Yisong Yue

California Institute of Technology, Pasadena, CA, USA

HMLE@CALTECH.EDU

AKANG@CALTECH.EDU

YYUE@CALTECH.EDU

Peter Carr

Disney Research, Pittsburgh, PA, USA

PETER.CARR@DISNEYRESEARCH.COM

Abstract

We study the problem of smooth imitation learning for online sequence prediction, where the goal is to train a policy that can smoothly imitate demonstrated behavior in a dynamic and continuous environment in response to online, sequential context input. Since the mapping from context to behavior is often complex, we take a learning reduction approach to reduce smooth imitation learning to a regression problem using complex function classes that are regularized to ensure smoothness. We present a learning meta-algorithm that achieves fast and stable convergence to a good policy. Our approach enjoys several attractive properties, including being fully deterministic, employing an adaptive learning rate that can provably yield larger policy improvements compared to previous approaches, and the ability to ensure stable convergence. Our empirical results demonstrate significant performance gains over previous approaches.

1. Introduction

In many complex planning and control tasks, it can be very challenging to explicitly specify a good policy. For such tasks, the use of machine learning to automatically learn a good policy from observed expert behavior, also known as imitation learning or learning from demonstrations, has proven tremendously useful (Abbeel & Ng, 2004; Ratliff et al., 2009; Argall et al., 2009; Ross & Bagnell, 2010; Ross et al., 2011; Jain et al., 2013).

In this paper, we study the problem of imitation learning for smooth online sequence prediction in a continuous regime.

Online sequence prediction is the problem of making online decisions in response to exogenous input from the environment, and is a special case of reinforcement learning (see Section 2). We are further interested in policies that make smooth predictions in a continuous action space.

Our motivating example is the problem of learning smooth policies for automated camera planning (Chen et al., 2016): determining where a camera should look given environment information (e.g., noisy person detections) and corresponding demonstrations from a human expert.¹ It is widely accepted that a smoothly moving camera is essential for generating aesthetic video (Gaddam et al., 2015). From a problem formulation standpoint, one key difference between smooth imitation learning and conventional imitation learning is the use of a “smooth” policy class (which we formalize in Section 2), and the goal now is to mimic expert demonstrations by choosing the best smooth policy.

The conventional supervised learning approach to imitation learning is to train a classifier or regressor to predict the expert’s behavior given training data comprising input/output pairs of contexts and actions taken by the expert. However, the learned policy’s prediction affects (the distribution of) future states during the policy’s actual execution, and so violates the crucial i.i.d. assumption made by most statistical learning approaches. To address this issue, numerous learning reduction approaches have been proposed (Daumé III et al., 2009; Ross & Bagnell, 2010; Ross et al., 2011), which iteratively modify the training distribution in various ways such that any supervised learning guarantees provably lift to the sequential imitation setting (potentially at the cost of statistical or computational efficiency).

We present a learning reduction approach to smooth imitation learning for online sequence prediction, which we call SIMILE (Smooth IMitation LEarning). Building

Proceedings of the 33rd International Conference on Machine Learning, New York, NY, USA, 2016. JMLR: W&CP volume 48. Copyright 2016 by the author(s).

¹Access data at <http://www.disneyresearch.com/publication/smooth-imitation-learning/> and code at <http://github.com/hoangminhle/SIMILE>.

upon learning reductions that employ policy aggregation (Daumé III et al., 2009), we provably lift supervised learning guarantees to the smooth imitation setting and show much faster convergence behavior compared to previous work. Our contributions can be summarized as:

- We formalize the problem of smooth imitation learning for online sequence prediction, and introduce a family of smooth policy classes that is amenable to supervised learning reductions.
- We present a principled learning reduction approach, which we call SIMILE. Our approach enjoys several attractive practical properties, including learning a fully deterministic stationary policy (as opposed to SEARN (Daumé III et al., 2009)), and not requiring data aggregation (as opposed to DAgger (Ross et al., 2011)) which can lead to super-linear training time.
- We provide performance guarantees that lift the the underlying supervised learning guarantees to the smooth imitation setting. Our guarantees hold in the agnostic setting, i.e., when the supervised learner might not achieve perfect prediction.
- We show how to exploit a stability property of our smooth policy class to enable adaptive learning rates that yield provably much faster convergence compared to SEARN (Daumé III et al., 2009).
- We empirically evaluate using the setting of smooth camera planning (Chen et al., 2016), and demonstrate the performance gains of our approach.

2. Problem Formulation

Let $\mathbf{X} = \{x_1, \dots, x_T\} \subset \mathcal{X}^T$ denote a context sequence from the environment \mathcal{X} , and $\mathbf{A} = \{a_1, \dots, a_T\} \subset \mathcal{A}^T$ denote an action sequence from some action space \mathcal{A} . Context sequence is exogenous, meaning a_t does not influence future context x_{t+k} for $k \geq 1$. Let Π denote a policy class, where each $\pi \in \Pi$ generates an action sequence \mathbf{A} in response to a context sequence \mathbf{X} . Assume $\mathcal{X} \subset \mathbb{R}^m, \mathcal{A} \subset \mathbb{R}^k$ are continuous and infinite, with \mathcal{A} non-negative and bounded such that $\vec{0} \leq a \leq R\vec{1} \forall a \in \mathcal{A}$.

Predicting actions a_t may depend on recent contexts x_t, \dots, x_{t-p} and actions a_{t-1}, \dots, a_{t-q} . Without loss of generality, we define a state space \mathcal{S} as $\{s_t = [x_t, a_{t-1}]\}$.² Policies π can thus be viewed as mapping states $\mathcal{S} = \mathcal{X} \times \mathcal{A}$ to actions \mathcal{A} . A roll-out of π given context sequence $\mathbf{X} = \{x_1, \dots, x_T\}$ is the action sequence $\mathbf{A} = \{a_1, \dots, a_T\}$:

$$a_t = \pi(s_t) = \pi([x_t, a_{t-1}]),$$

$$s_{t+1} = [x_{t+1}, a_t] \quad \forall t \in [1, \dots, T].$$

Note that unlike the general reinforcement learning problem, we consider the setting where the state space splits into external and internal components (by definition, a_t influences subsequent states s_{t+k} , but not x_{t+k}). The use

of exogenous contexts $\{x_t\}$ models settings where a policy needs to take online, sequential actions based on external environmental inputs, e.g. smooth self-driving vehicles for obstacle avoidance, helicopter aerobatics in the presence of turbulence, or smart grid management for external energy demand. The technical motivation of this dichotomy is that we will enforce smoothness only on the internal state.

Consider the example of autonomous camera planning for broadcasting a sport event (Chen et al., 2016). \mathcal{X} can correspond to game information such as the locations of the players, the ball, etc., and \mathcal{A} can correspond to the pan-tilt-zoom configuration of the broadcast camera. Manually specifying a good camera policy can be very challenging due to sheer complexity involved with mapping \mathcal{X} to \mathcal{A} . It is much more natural to train $\pi \in \Pi$ to mimic observed expert demonstrations. For instance, Π can be the space of neural networks or tree-based ensembles (or both).

Following the basic setup from (Ross et al., 2011), for any policy $\pi \in \Pi$, let d_t^π denote the distribution of states at time t if π is executed for the first $t-1$ time steps. Furthermore, let $d_\pi = \frac{1}{T} \sum_{t=1}^T d_t^\pi$ be the average distribution of states if we follow π for all T steps. The goal of imitation learning is to find a policy $\hat{\pi} \in \Pi$ which minimizes the imitation loss under its own induced distribution of states:

$$\hat{\pi} = \underset{\pi \in \Pi}{\operatorname{argmin}} \ell_\pi(\pi) = \underset{\pi \in \Pi}{\operatorname{argmin}} \mathbb{E}_{s \sim d_\pi} [\ell(\pi(s))], \quad (1)$$

where the (convex) imitation loss $\ell(\pi(s))$ captures how well π imitates expert demonstrations for state s . One common ℓ is squared loss between the policy’s decision and the expert demonstration: $\ell(\pi(s)) = \|\pi(s) - \pi^*(s)\|^2$ for some norm $\|\cdot\|$. Note that computing ℓ typically requires having access to a training set of expert demonstrations π^* on some set of context sequences. We also assume an agnostic setting, where the minimizer of (1) does not necessarily achieve 0 loss (i.e. it cannot perfectly imitate the expert).

2.1. Smooth Imitation Learning & Smooth Policy Class

In addition to accuracy, a key requirement of many continuous control and planning problems is smoothness (e.g., smooth camera trajectories). Generally, “smoothness” may reflect domain knowledge about stability properties or approximate equilibria of a dynamical system. We thus formalize the problem of *smooth imitation learning* as minimizing (1) over a smooth policy class Π .

Most previous work on learning smooth policies focused on simple policy classes such as linear models (Abbeel & Ng, 2004), which can be overly restrictive. We instead define a much more general smooth policy class Π as a regularized space of complex models.

Definition 2.1 (Smooth policy class Π). *Given a complex model class \mathcal{F} and a class of smooth regularizers \mathcal{H} , we define smooth policy class $\Pi \subset \mathcal{F} \times \mathcal{H}$ as satisfying:*

²We can always concatenate consecutive contexts and actions.

$$\begin{aligned} \Pi \triangleq \{ \pi = (f, h), f \in \mathcal{F}, h \in \mathcal{H} \mid \pi(s) \text{ is close to} \\ \text{both } f(x, a) \text{ and } h(a) \\ \forall \text{ induced state } s = [x, a] \in \mathcal{S} \} \end{aligned}$$

where closeness is controlled by regularization.

For instance, \mathcal{F} can be the space of neural networks or decision trees and \mathcal{H} be the space of smooth analytic functions. Π can thus be viewed as policies that predict close to some $f \in \mathcal{F}$ but are regularized to be close to some $h \in \mathcal{H}$. For sufficiently expressive \mathcal{F} , we often have that $\Pi \subset \mathcal{F}$. Thus optimizing over Π can be viewed as constrained optimization over \mathcal{F} (by \mathcal{H}), which can be challenging. Our SIMILE approach integrates alternating optimization (between \mathcal{F} and \mathcal{H}) into the learning reduction. We provide two concrete examples of Π below.

Example 2.1 (Π_λ). Let \mathcal{F} be any complex supervised model class, and define the simplest possible $\mathcal{H} \triangleq \{h(a) = a\}$. Given $f \in \mathcal{F}$, the prediction of a policy π can be viewed as regularized optimization over the action space to ensure closeness of π to both f and h :

$$\begin{aligned} \pi(x, a) &= \operatorname{argmin}_{a' \in \mathcal{A}} \|f(x, a) - a'\|^2 + \lambda \|h(a) - a'\|^2 \\ &= \frac{f(x, a) + \lambda h(a)}{1 + \lambda} = \frac{f(x, a) + \lambda a}{1 + \lambda}, \end{aligned} \quad (2)$$

where regularization parameter λ trades-off closeness to f and to previous action. For large λ , $\pi(x, a)$ is encouraged make predictions that stays close to previous action a .

Example 2.2 (Linear auto-regressor smooth regularizers). Let \mathcal{F} be any complex supervised model class, and define \mathcal{H} using linear auto-regressors, $\mathcal{H} \triangleq \{h(a) = \theta^\top a\}$, which model actions as a linear dynamical system (Wold, 1939). We can define π analogously to (2).

In general, SIMILE requires that Π satisfies a smooth property stated below. This property, which is exploited in our theoretical analysis (see Section 5), is motivated by the observation that given a (near) constant stream of context sequence, a stable behavior policy should exhibit a corresponding action sequence with low curvature. The two examples above satisfy this property for sufficiently large λ .

Definition 2.2 (H -state-smooth imitation policy). For small constant $0 < H \ll 1$, a policy $\pi([x, a])$ is H -state-smooth if it is H -smooth w.r.t. a , i.e. for fixed $x \in \mathcal{X}$, $\forall a, a' \in \mathcal{A}$, $\forall i$: $\|\nabla \pi^i([x, a]) - \nabla \pi^i([x, a'])\|_* \leq H \|a - a'\|$ where π^i indicates the i^{th} component of vector-valued function³ $\pi(s) = [\pi^1(s), \dots, \pi^k(s)] \in \mathbb{R}^k$, and $\|\cdot\|$ and $\|\cdot\|_*$ are some norm and dual norm respectively. For twice differentiable policy π , this is equivalent to having the bound on the Hessian $\nabla^2 \pi^i([x, a]) \leq H \mathbb{I}_k \forall i$.

³This emphasizes the possibility that π is a vector-valued function of a . The gradient and Hessian are viewed as arrays of k gradient vectors and Hessian matrices of 1-d case, since we simply treat action in \mathbb{R}^k as an array of k standard functions.

3. Related Work

The most popular traditional approaches for learning from expert demonstration focused on using approximate policy iteration techniques in the MDP setting (Kakade & Langford, 2002; Bagnell et al., 2003). Most prior approaches operate in discrete and finite action space (He et al., 2012; Ratliff et al., 2009; Abbeel & Ng, 2004; Argall et al., 2009). Some focus on continuous state space (Abbeel & Ng, 2005), but requires a linear model for the system dynamics. In contrast, we focus on learning complex smooth functions within continuous action and state spaces.

One natural approach to tackle the more general setting is to reduce imitation learning to a standard supervised learning problem (Syed & Schapire, 2010; Langford & Zadrozny, 2005; Lagoudakis & Parr, 2003). However, standard supervised methods assume i.i.d. training and test examples, thus ignoring the distribution mismatch between training and rolled-out trajectories directly applied to sequential learning problems (Kakade & Langford, 2002). Thus a naive supervised learning approach normally leads to unsatisfactory results (Ross & Bagnell, 2010).

Iterative Learning Reductions. State-of-the-art learning reductions for imitation learning typically take an iterative approach, where each training round uses standard supervised learning to learn a policy (Daumé III et al., 2009; Ross et al., 2011). In each round n , the following happens:

- Given initial state s_0 drawn from the starting distribution of states, the learner executes current policy π_n , resulting in a sequence of states s_1^n, \dots, s_T^n .
- For each s_t^n , a label \hat{a}_t^n (e.g., expert feedback) is collected indicating what the expert would do given s_t^n , resulting in a new dataset $\mathcal{D}_n = \{(s_t, \hat{a}_t^n)\}$.
- The learner integrates \mathcal{D}_n to learn a policy $\hat{\pi}_n$. The learner updates the current policy to π_{n+1} based on $\hat{\pi}_n$ and π_n .

The main challenge is controlling for the cascading errors caused by the changing dynamics of the system, i.e., the distribution of states in each $\mathcal{D}_n \sim d_{\pi_n}$. A policy trained using d_{π_n} induces a different distribution of states than d_{π_n} , and so is no longer being evaluated on the same distribution as during training. A principled reduction should (approximately) preserve the i.i.d. relationship between training and test examples. Furthermore the state distribution d_π should converge to a stationary distribution.

The arguably most notable learning reduction approaches for imitation learning are SEARN (Daumé III et al., 2009) and DAGger (Ross et al., 2011). At each round, SEARN learns a new policy $\hat{\pi}_n$ and returns a distribution (or mixture) over previously learned policies: $\pi_{n+1} = \beta \hat{\pi}_n + (1 - \beta) \pi_n$ for $\beta \in (0, 1)$. For appropriately small choices of β , this stochastic mixing limits the “distribution drift” between π_n and π_{n+1} and can provably guarantee that the

performance of π_{n+1} does not degrade significantly relative to the expert demonstrations.⁴

DAgger, on the other hand, achieves stability by aggregating a new dataset at each round to learn a new policy from the combined data set $\mathcal{D} \leftarrow \mathcal{D} \cup \mathcal{D}_n$. This aggregation, however, significantly increases the computational complexity and thus is not practical for large problems that require many iterations of learning (since the training time grows super-linearly w.r.t. the number of iterations).

Both SEARN and DAgger showed that only a polynomial number of training rounds is required for convergence to a good policy, but with a dependence on the length of horizon T . In particular, to non-trivially bound the total variation distance $\|d_{\pi_{new}} - d_{\pi_{old}}\|_1$ of the state distributions between old and new policies, a learning rate $\beta < \frac{1}{T}$ is required to hold (Lemma 1 of Daumé III, Langford, and Marcu (2009) and Theorem 4.1 of Ross, Gordon, and Bagnell (2011)). As such, systems with very large time horizons might suffer from very slow convergence.

Our Contributions. Within the context of previous work, our SIMILE approach can be viewed as extending SEARN to smooth policy classes with the following improvements:

- We provide a policy improvement bound that does not depend on the time horizon T , and can thus converge much faster. In addition, SIMILE has adaptive learning rate, which can further improve convergence.
- For the smooth policy class described in Section 2, we show how to generate simulated or “virtual” expert feedback in order to guarantee stable learning. This alleviates the need to have continuous access to a dynamic oracle / expert that shows the learner what to do when it is off-track. In this regard, the way SIMILE integrates expert feedback subsumes the set-up from SEARN and DAgger.
- Unlike SEARN, SIMILE returns fully deterministic policies. Under the continuous setting, deterministic policies are strictly better than stochastic policies as (i) smoothness is critical and (ii) policy sampling requires holding more data during training, which may not be practical for infinite state and action spaces.
- Our theoretical analysis reveals a new sequential prediction setting that yields provably fast convergence, in particular for smooth policy classes on finite-horizon problems. Existing settings that enjoy such results are limited to Markovian dynamics with discounted future rewards or linear model classes.

4. Smooth Imitation Learning Algorithm

Our learning algorithm, called SIMILE (Smooth IMItation LEarning), is described in Algorithm 1. At a high level, the process can be described as:

⁴A similar approach was adopted in Conservative Policy Iteration for the MDP setting (Kakade & Langford, 2002).

Algorithm 1 SIMILE (Smooth IMItation LEarning)

Input: features $\mathbf{X} = \{x_t\}$, human trajectory $\mathbf{A}^* = \{a_t^*\}$, base routine Train , smooth regularizers $h \in \mathcal{H}$

1: Initialize $\mathbf{A}_0 \leftarrow \mathbf{A}^*$, $\mathbf{S}_0 \leftarrow \{[x_t, a_{t-1}^*]\}$,

$$h_0 = \operatorname{argmin}_{h \in \mathcal{H}} \sum_{t=1}^T \|a_t^* - h(a_{t-1}^*)\|$$

2: Initial policy $\pi_0 = \hat{\pi}_0 \leftarrow \text{Train}(\mathbf{S}_0, \mathbf{A}_0 | h_0)$

3: **for** $n = 1, \dots, N$ **do**

4: $\mathbf{A}_n = \{a_t^n\} \leftarrow \pi_{n-1}(\mathbf{S}_{n-1})$ //sequential roll-out

5: $\mathbf{S}_n \leftarrow \{s_t^n = [x_t, a_{t-1}^*]\}$ // $s_t^n = [x_{t:t-p}, a_{t-1:t-q}]$

6: $\hat{\mathbf{A}}_n = \{\hat{a}_t^n\} \forall s_t^n \in \mathbf{S}_n$ // collect smooth feedback

7: $h_n = \operatorname{argmin}_{h \in \mathcal{H}} \sum_{t=1}^T \|\hat{a}_t^n - h(\hat{a}_{t-1}^n)\|$ //new regularizer

8: $\hat{\pi}_n \leftarrow \text{Train}(\mathbf{S}_n, \hat{\mathbf{A}}_n | h_n)$ // train policy

9: $\beta \leftarrow \beta(\ell(\hat{\pi}_n), \ell(\pi_{n-1}))$ //adaptively set β

10: $\pi_n = \beta \hat{\pi}_n + (1 - \beta)\pi_{n-1}$ // update policy

11: **end for**

output Last policy π_N

1. Start with some initial policy $\hat{\pi}_0$ (Line 2).
2. At iteration n , use π_{n-1} to build a new state distribution \mathbf{S}_n and dataset $\mathcal{D}_n = \{(s_t^n, \hat{a}_t^n)\}$ (Lines 4-6).
3. Train $\hat{\pi}_n = \operatorname{argmin}_{\pi \in \Pi} \mathbb{E}_{s \sim \mathbf{S}_n} [\ell_n(\pi(s))]$, where ℓ_n is the imitation loss (Lines 7-8). Note that ℓ_n needs not be the original ℓ , but simply needs to converge to it.
4. Interpolate $\hat{\pi}_n$ and π_{n-1} to generate a new deterministic policy π_n (Lines 9-10). Repeat from Step 2 with $n \leftarrow n + 1$ until some termination condition is met.

Supervised Learning Reduction. The actual reduction is in Lines 7-8, where we follow a two-step procedure of first updating the smooth regularizer h_n , and then training $\hat{\pi}_n$ via supervised learning. In other words, Train finds the best $f \in \mathcal{F}$ possible for a fixed h_n . We discuss how to set the training targets \hat{a}_t^n below.

Policy Update. The new policy π_n is a deterministic interpolation between the previous π_{n-1} and the newly learned $\hat{\pi}_n$ (Line 10). In contrast, for SEARN, π_n is a stochastic interpolation (Daumé III et al., 2009). Lemma 5.2 and Corollary 5.3 show that deterministic interpolation converges at least as fast as stochastic for smooth policy classes.

This interpolation step plays two key roles. First, it is a form of myopic or greedy online learning. Intuitively, rolling out π_n leads to incidental exploration on the mistakes of π_n , and so each round of training is focused on refining π_n . Second, the interpolation in Line 10 ensures a slow drift in the distribution of states from round to round, which preserves an approximate i.i.d. property for the supervised regression subroutine and guarantees convergence.

However this model interpolation creates an inherent tension between maintaining approximate i.i.d. for valid su-

pervised learning and more aggressive exploration (and thus faster convergence). For example, SEARN’s guarantees only apply for small $\beta < 1/T$. SIMILE circumvents much of this tension via a policy improvement bound that allows β to adaptively increase depending on the quality of $\hat{\pi}_n$ (see Theorem 5.6), which thus guarantees a valid learning reduction while substantially speeding up convergence.

Feedback Generation. We can generate training targets \hat{a}_t^n using “virtual” feedback from simulating expert demonstrations, which has two benefits. First, we need not query the expert π^* at every iteration (as done in DAgger (Ross et al., 2011)). Continuously acquiring expert demonstrations at every round can be seen as a special case and a more expensive strategy. Second, virtual feedback ensures stable learning, i.e., every $\hat{\pi}_n$ is a feasible smooth policy.

Consider Figure 1, where our policy π_n (blue/red) made a mistake at location A, and where we have only a single expert demonstration from π^* (black). Depending on the smoothness requirements of the policy class, we can simulate virtual expert feedback as via either the red line (more smooth) or blue (less smooth) as a tradeoff between squared imitation loss and smoothness.




Figure 1.

When the roll-out of π_{n-1} (i.e. \mathbf{A}_n) differs substantially from \mathbf{A}^* , especially during early iterations, using smoother feedback (red instead of blue) can result in more stable learning. We formalize this notion for Π_λ in Proposition 5.8. Intuitively, whenever π_{n-1} makes a mistake, resulting in a “bad” state s_t^n , the feedback should recommend a smooth correction \hat{a}_t^n w.r.t. \mathbf{A}_n to make training “easier” for the learner.⁵ The virtual feedback \hat{a}_t^n should converge to the expert’s action a_t^* . In practice, we use $\hat{a}_t^n = \sigma a_t^n + (1 - \sigma)a_t^*$ with $\sigma \rightarrow 0$ as n increases (which satisfies Proposition 5.8).

5. Theoretical Results

All proofs are deferred to the supplementary material.

5.1. Stability Conditions

One natural smoothness condition is that $\pi([x, a])$ should be stable w.r.t. a if x is fixed. Consider the camera planning setting: the expert policy π^* should have very small curvature, since constant inputs should correspond to constant actions. This motivates Definition 2.2, which requires that Π has low curvature given fixed context. We also show that smooth policies per Definition 2.2 lead to stable actions, in the sense that “nearby” states are mapped to “nearby” actions. The following helper lemma is useful:

⁵A similar idea was proposed (He et al., 2012) for DAgger-type algorithm, albeit only for linear model classes.

Lemma 5.1. For a fixed x , define $\pi([x, a]) \triangleq \varphi(a)$. If φ is non-negative and H -smooth w.r.t. a , then:

$$\forall a, a' : (\varphi(a) - \varphi(a'))^2 \leq 6H (\varphi(a) + \varphi(a')) \|a - a'\|^2.$$

Writing π as $\pi([x, a]) \triangleq [\pi^1([x, a]), \dots, \pi^k([x, a])]$ with each $\pi^i([x, a])$ H -smooth, Lemma 5.1 implies $\|(\pi([x, a]) - \pi([x, a']))\| \leq \sqrt{12HR} \|a - a'\|$ for R upper bounding \mathcal{A} . Bounded action space means that a sufficiently small H leads to the following stability conditions:

Condition 1 (Stability Condition 1). Π satisfies the Stability Condition 1 if for a fixed input feature x , the actions of π in states $s = [x, a]$ and $s' = [x, a']$ satisfy $\|\pi(s) - \pi(s')\| \leq \|a - a'\|$ for all $a, a' \in \mathcal{A}$.

Condition 2 (Stability Condition 2). Π satisfies Stability Condition 2 if each π is γ -Lipschitz continuous in the action component a with $\gamma < 1$. That is, for a fixed x the actions of π in states $s = [x, a]$ and $s' = [x, a']$ satisfy $\|\pi(s) - \pi(s')\| \leq \gamma \|a - a'\|$ for all $a, a' \in \mathcal{A}$.

These two conditions directly follow from Lemma 5.1 and assuming sufficiently small H . Condition 2 is mildly stronger than Condition 1, and enables proving much stronger policy improvement compared to previous work.

5.2. Deterministic versus Stochastic

Given two policies π and $\hat{\pi}$, and interpolation parameter $\beta \in (0, 1)$, consider two ways to combine policies:

1. **stochastic:** $\pi_{sto}(s) = \hat{\pi}(s)$ with probability β , and $\pi_{sto}(s) = \pi(s)$ with probability $1 - \beta$
2. **deterministic:** $\pi_{det}(s) = \beta\hat{\pi}(s) + (1 - \beta)\pi(s)$

Previous learning reduction approaches only use stochastic interpolation (Daumé III et al., 2009; Ross et al., 2011), whereas SIMILE uses deterministic. The following result shows that deterministic and stochastic interpolation yield the same expected behavior for smooth policy classes.

Lemma 5.2. Given any starting state s_0 , sequentially execute π_{det} and π_{sto} to obtain two separate trajectories $\mathbf{A} = \{a_t\}_{t=1}^T$ and $\tilde{\mathbf{A}} = \{\tilde{a}_t\}_{t=1}^T$ such that $a_t = \pi_{det}(s_t)$ and $\tilde{a}_t = \pi_{sto}(\tilde{s}_t)$, where $s_t = [x_t, a_{t-1}]$ and $\tilde{s}_t = [x_t, \tilde{a}_{t-1}]$. Assuming the policies are stable as per Condition 1, we have $\mathbb{E}_{\tilde{\mathbf{A}}}[\tilde{a}_t] = a_t \forall t = 1, \dots, T$, where the expectation is taken over all random roll-outs of π_{sto} .

Lemma 5.2 shows that deterministic policy combination (SIMILE) yields unbiased trajectory roll-outs of stochastic policy combination (as done in SEARN & CPI). This represents a major advantage of SIMILE, since the number of stochastic roll-outs of π_{sto} to average to the deterministic trajectory of π_{det} is polynomial in the time horizon T , leading to much higher computational complexity. Furthermore, for convex imitation loss $\ell_\pi(\pi)$, Lemma 5.2 and Jensen’s inequality yield the following corollary, which states that under convex loss, deterministic policy performs at least no worse than stochastic policy in expectation:

Corollary 5.3 (Deterministic Policies Perform Better). *For deterministic π_{det} and stochastic π_{sto} interpolations of two policies π and $\hat{\pi}$, and convex loss ℓ , we have:*

$$\begin{aligned} \ell_{\pi_{det}}(\pi_{det}) &= \ell_{\pi_{sto}}(\mathbb{E}[\pi_{sto}]) \\ &\leq \mathbb{E}[\ell_{\pi_{sto}}(\pi_{sto})] \end{aligned}$$

where the expectation is over all roll-outs of π_{sto} .

Remark. We construct a simple example to show that Condition 1 may be necessary for iterative learning reductions to converge. Consider the case where contexts $\mathbf{X} \subset \mathbb{R}$ are either constant or vary negligibly. Expert demonstrations should be constant $\pi^*([x_n, a^*]) = a^*$ for all n . Consider an unstable policy π such that $\pi(s) = \pi([x, a]) = ka$ for fixed $k > 1$. The rolled-out trajectory of π diverges π^* at an exponential rate. Assume optimistically that $\hat{\pi}$ learns the correct expert behavior, which is simply $\hat{\pi}(s) = \hat{\pi}([x, a]) = a$. For any $\beta \in (0, 1)$, the updated policy $\pi' = \beta\hat{\pi} + (1-\beta)\pi$ becomes $\pi'([x, a]) = \beta a + (1-\beta)ka$. Thus the sequential roll-out of the new policy π' will also yield an exponential gap from the correct policy. By induction, the same will be true in all future iterations.

5.3. Policy Improvement

Our policy improvement guarantee builds upon the analysis from SEARN (Daumé III et al., 2009), which we extend to using adaptive learning rates β . We first restate the main policy improvement result from Daumé III et al. (2009).

Lemma 5.4 (SEARN’s policy nondegradation - Lemma 1 from Daumé III et al. (2009)). *Let $\ell_{max} = \sup_{\pi, s} \ell(\pi(s))$, π' is defined as π_{sto} in lemma 5.2. Then for $\beta \in (0, 1/T)$:*

$$\ell_{\pi'}(\pi') - \ell_{\pi}(\pi) \leq \beta T \mathbb{E}_{s \sim d_{\pi}} [\ell(\hat{\pi}(s))] + \frac{1}{2} \beta^2 T^2 \ell_{max}.$$

SEARN guarantees that the new policy π' does not degrade from the expert π^* by much only if $\beta < 1/T$. Analyses of SEARN and other previous iterative reduction methods (Ross et al., 2011; Kakade & Langford, 2002; Bagnell et al., 2003; Syed & Schapire, 2010) rely on bounding the variation distance between d_{π} and $d_{\pi'}$. Three drawbacks of this approach are: (i) non-trivial variation distance bound typically requires β to be inversely proportional to time horizon T , causing slow convergence; (ii) not easily applicable to the continuous regime; and (iii) except under MDP framework with discounted infinite horizon, previous variation distance bounds do not guarantee monotonic policy improvements (i.e. $\ell_{\pi'}(\pi') < \ell_{\pi}(\pi)$).

We provide two levels of guarantees taking advantage of Stability Conditions 1 and 2 to circumvent these drawbacks. Assuming the Condition 1 and convexity of ℓ , our first result yields a guarantee comparable with SEARN.

Theorem 5.5 (T-dependent Improvement). *Assume ℓ is convex and L-Lipschitz, and Condition 1 holds. Let $\epsilon = \max_{s \sim d_{\pi}} \|\hat{\pi}(s) - \pi(s)\|$. Then:*

$$\ell_{\pi'}(\pi') - \ell_{\pi}(\pi) \leq \beta \epsilon L T + \beta (\ell_{\pi}(\hat{\pi}) - \ell_{\pi}(\pi)). \quad (3)$$

In particular, choosing $\beta \in (0, 1/T)$ yields:

$$\ell_{\pi'}(\pi') - \ell_{\pi}(\pi) \leq \epsilon L + \beta (\ell_{\pi}(\hat{\pi}) - \ell_{\pi}(\pi)). \quad (4)$$

Similar to SEARN, Theorem 5.5 also requires $\beta \in (0, 1/T)$ to ensure the RHS of (4) stays small. However, note that the reduction term $\beta(\ell_{\pi}(\hat{\pi}) - \ell_{\pi}(\pi))$ allows the bound to be strictly negative if the policy $\hat{\pi}$ trained on d_{π} significantly improves on $\ell_{\pi}(\pi)$ (i.e., guaranteed policy improvement). We observe empirically that this often happens, especially in early iterations of training.

Under the mildly stronger Condition 2, we remove the dependency on the time horizon T , which represents a much stronger guarantee compared to previous work.

Theorem 5.6 (Policy Improvement). *Assume ℓ is convex and L-Lipschitz-continuous, and Condition 2 holds. Let $\epsilon = \max_{s \sim d_{\pi}} \|\hat{\pi}(s) - \pi(s)\|$. Then for $\beta \in (0, 1)$:*

$$\ell_{\pi'}(\pi') - \ell_{\pi}(\pi) \leq \frac{\beta \gamma \epsilon L}{(1-\beta)(1-\gamma)} + \beta (\ell_{\pi}(\hat{\pi}) - \ell_{\pi}(\pi)). \quad (5)$$

Corollary 5.7 (Monotonic Improvement). *Following the notation from Theorem 5.6, let $\Delta = \ell_{\pi}(\pi) - \ell_{\pi}(\hat{\pi})$ and $\delta = \frac{\gamma \epsilon L}{1-\gamma}$. Then choosing step size $\beta = \frac{\Delta - \delta}{2\Delta}$, we have:*

$$\ell_{\pi'}(\pi') - \ell_{\pi}(\pi) \leq -\frac{(\Delta - \delta)^2}{2(\Delta + \delta)}. \quad (6)$$

The terms ϵ and $\ell_{\pi}(\hat{\pi}) - \ell_{\pi}(\pi)$ on the RHS of (4) and (5) come from the learning reduction, as they measure the “distance” between $\hat{\pi}$ and π on the state distribution induced by π (which forms the dataset to train $\hat{\pi}$). In practice, both terms can be empirically estimated from the training round, thus allowing an estimate of β to minimize the bound.

Theorem 5.6 justifies using an adaptive and more aggressive interpolation parameter β to update policies. In the worst case, setting β close to 0 will ensure the bound from (5) to be close to 0, which is consistent with SEARN’s result. More generally, monotonic policy improvement can be guaranteed for appropriate choice of β , as seen from Corollary 5.7. This strict policy improvement was not possible under previous iterative learning reduction approaches such as SEARN and DAGger, and is enabled in our setting due to exploiting the smoothness conditions.

5.4. Smooth Feedback Analysis

Smooth Feedback Does Not Hurt: Recall from Section 4 that one way to simulate “virtual” feedback for training a new $\hat{\pi}$ is to set the target $\hat{a}_t = \sigma a_t + (1-\sigma)a_t^*$ for $\sigma \in (0, 1)$, where smooth feedback corresponds to $\sigma \rightarrow 1$. To see that simulating smooth “virtual” feedback target does not hurt the training progress, we alternatively view SIMILE as performing gradient descent in a smooth function space (Mason et al., 1999). Define the cost functional $C : \Pi \rightarrow \mathbb{R}$ over policy space to be the average imitation loss over \mathcal{S} as $C(\pi) = \int_{\mathcal{S}} \|\pi(s) - \pi^*(s)\|^2 dP(s)$. The gra-

dient (Gâteaux derivative) of $C(\pi)$ w.r.t. π is:

$$\nabla C(\pi)(s) = \frac{\partial C(\pi + \alpha \delta_s)}{\partial \alpha} \Big|_{\alpha=0} = 2(\pi(s) - \pi^*(s)),$$

where δ_s is Dirac delta function centered at s . By first order approximation $C(\pi') = C(\beta \hat{\pi} + (1 - \beta)\pi) = C(\pi + \beta(\hat{\pi} - \pi)) \approx C(\pi) + \beta \langle \nabla C(\pi), \hat{\pi} - \pi \rangle$. Like traditional gradient descent, we want to choose $\hat{\pi}$ such that the update moves the functional along the direction of negative gradient. In other words, we want to learn $\hat{\pi} \in \Pi$ such that $\langle \nabla C(\pi), \hat{\pi} - \pi \rangle \ll 0$. We can evaluate this inner product along the states induced by π . We thus have the estimate:

$$\begin{aligned} \langle \nabla C(\pi), \hat{\pi} - \pi \rangle &\approx \frac{2}{T} \sum_{t=1}^T (\pi(s_t) - \pi^*(s_t))(\hat{\pi}(s_t) - \pi(s_t)) \\ &= \frac{2}{T} \sum_{t=1}^T (a_t - a_t^*)(\hat{\pi}([x_t, a_{t-1}]) - a_t). \end{aligned}$$

Since we want $\langle \nabla C(\pi), \hat{\pi} - \pi \rangle < 0$, this motivates the construction of new data set \mathcal{D} with states $\{[x_t, a_{t-1}]\}_{t=1}^T$ and labels $\{\hat{a}_t\}_{t=1}^T$ to train a new policy $\hat{\pi}$, where we want $(a_t - a_t^*)(\hat{a}_t - a_t) < 0$. A sufficient solution is to set target $\hat{a}_t = \sigma a_t + (1 - \sigma)a_t^*$ (Section 4), as this will point the gradient in negative direction, allowing the learner to make progress.

Smooth Feedback is Sometimes Necessary: When the current policy performs poorly, smooth virtual feedback may be required to ensure stable learning, i.e. producing a feasible smooth policy at each training round. We formalize this notion of feasibility by considering the smooth policy class Π_λ in Example 2.1. Recall that smooth regularization of Π_λ via \mathcal{H} encourages the next action to be close to the previous action. Thus a natural way to measure smoothness of $\pi \in \Pi_\lambda$ is via the average first order difference of consecutive actions $\frac{1}{T} \sum_{t=1}^T \|a_t - a_{t-1}\|$. In particular, we want to explicitly constrain this difference relative to the expert trajectory $\frac{1}{T} \sum_{t=1}^T \|a_t - a_{t-1}\| \leq \eta$ at each iteration, where $\eta \propto \frac{1}{T} \sum_{t=1}^T \|a_t^* - a_{t-1}^*\|$.

When π performs poorly, i.e. the "average gap" between current trajectory $\{a_t\}$ and $\{a_t^*\}$ is large, the training target for $\hat{\pi}$ should be lowered to ensure learning a smooth policy is feasible, as inferred from the following proposition. In practice, we typically employ smooth virtual feedback in early iterations when policies tend to perform worse.

Proposition 5.8. *Let ω be the average supervised training error from \mathcal{F} , i.e. $\omega = \min_{f \in \mathcal{F}} \mathbb{E}_{x \sim \mathcal{X}} [\|f([x, 0]) - a^*\|]$.*

Let the rolled-out trajectory of current policy π be $\{a_t\}$. If the average gap between π and π^ is such that $\mathbb{E}_{t \sim \text{Uniform}[1:T]} [\|a_t^* - a_{t-1}\|] \geq 3\omega + \eta(1 + \lambda)$, then using $\{a_t^*\}$ as feedback will cause the trained policy $\hat{\pi}$ to be non-smooth, i.e.:*

$$\mathbb{E}_{t \sim \text{Uniform}[1:T]} [\|\hat{a}_t - \hat{a}_{t-1}\|] \geq \eta, \quad (7)$$

for $\{\hat{a}_t\}$ the rolled-out trajectory of $\hat{\pi}$.

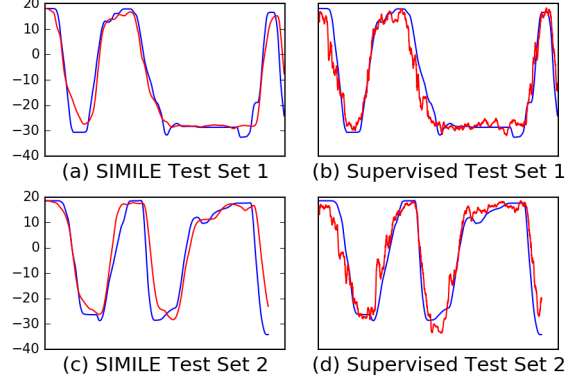


Figure 2. Expert (blue) and predicted (red) camera pan angles. Left: SIMILE with < 10 iterations. Right: non-smooth policy.




Figure 3. Adaptive versus fixed interpolation parameter β .

6. Experiments

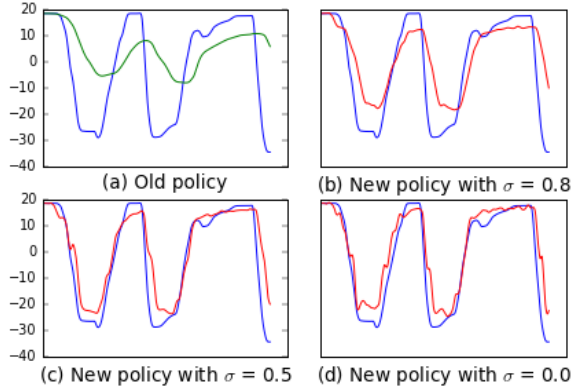
Automated Camera Planning. We evaluate SIMILE in a case study of automated camera planning for sport broadcasting (Chen & Carr, 2015; Chen et al., 2016). Given noisy tracking of players as raw input data $\{x_t\}_{t=1}^T$, and demonstrated pan camera angles from professional human operator as $\{a_t^*\}_{t=1}^T$, the goal is to learn a policy π that produces trajectory $\{a_t\}_{t=1}^T$ that is both smooth and accurate relative to $\{a_t^*\}_{t=1}^T$. Smoothness is critical in camera control: fluid movements which maintain adequate framing are preferable to jittery motions which constantly pursue perfect tracking (Gaddam et al., 2015). In this setting, time horizon T is the duration of the event multiplied by rate of sampling. Thus T tends to be very large.

Smooth Policy Class. We use a smooth policy class following Example 2.2: regression tree ensembles \mathcal{F} regularized by a class of linear autoregressor functions \mathcal{H} (Chen et al., 2016). See Appendix B for more details.

Summary of Results.

- Using our smooth policy class leads to dramatically smoother trajectories than not regularizing using \mathcal{H} .
- Using our adaptive learning rate leads to much faster convergence compared to conservative learning rates from SEARN (Daumé III et al., 2009).
- Using smooth feedback ensures stable learning of smooth policies at each iteration.
- Deterministic policy interpolation performs better than stochastic interpolation used in SEARN.

Smooth versus Non-Smooth Policy Classes. Figure 2 shows a comparison of using a smooth policy class versus a non-smooth one (e.g., not using \mathcal{H}). We see that our


 Figure 4. Comparing different values of σ .

approach can reliably learn to predict trajectories that are both smooth and accurate.

Adaptive vs. Fixed β : One can, in principle, train using SEARN, which requires a very conservative β in order to guarantee convergence. In contrast, SIMILE adaptively selects β based on relative empirical loss of π and $\hat{\pi}$ (Line 9 of Algorithm 1). Let $\text{error}(\hat{\pi})$ and $\text{error}(\pi)$ denote the mean-squared errors of rolled-out trajectories $\{\hat{a}_t\}$, $\{a_t\}$, respectively, w.r.t. ground truth $\{a_t^*\}$. We can set β as:

$$\hat{\beta} = \frac{\text{error}(\pi)}{\text{error}(\hat{\pi}) + \text{error}(\pi)}, \quad (8)$$

which encourages the learner to disregard bad policies when interpolating, thus allowing fast convergence to a good policy (see Theorem 5.6). Figure 3 compares the convergence rate of SIMILE using adaptive β versus conservative fixed values of β commonly used in SEARN (Daumé III et al., 2009). We see that adaptively choosing β enjoys substantially faster convergence. Note that very large fixed β may overshoot and worsen the combined policy after a few initial improvements.

Smooth Feedback Generation: We set the target labels to $\hat{a}_t^n = \sigma a_t^n + (1 - \sigma)a_t^*$ for $0 < \sigma < 1$ (Line 6 of Algorithm 1). Larger σ corresponds to smoother (\hat{a}_t^n is closer to a_{t-1}^n) but less accurate target (further from a_t^*), as seen in Figure 4. Figure 5 shows the trade-off between smoothness loss (blue




Figure 5.

line, measured by first order difference in Proposition 5.8) and imitation loss (red line, measured by mean squared distance) for varying σ . We navigate this trade-off by setting σ closer to 1 in early iterations, and have $\sigma \rightarrow 0$ as n increases. This “gradual increase” produces more stable policies, especially during early iterations where the learning policy tends to perform poorly (as formalized in Proposition 5.8). In Figure 4, when the initial policy (green trajectory) has poor performance, setting smooth

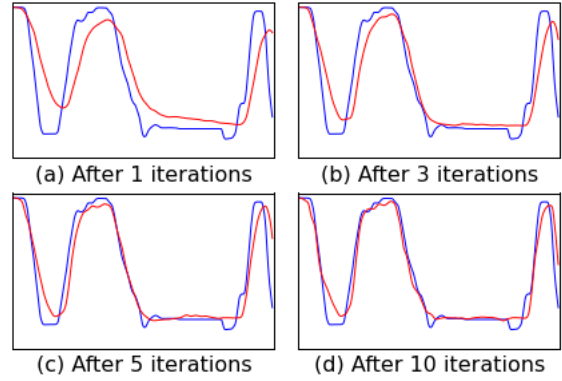



Figure 6. Performance after different number of iterations.

targets (Figure 4b) allows learning a smooth policy in the subsequent round, in contrast to more accurate but less stable performance of “difficult” targets with low σ (Figure 4c-d). Figure 6 visualizes the behavior of the intermediate policies learned by SIMILE, where we can see that each intermediate policy is a smooth policy.

Deterministic vs. Stochastic Interpolation: Finally, we evaluate the benefits of using deterministic policy averaging instead of stochastically combine different policies, as done in SEARN. To control for other factors, we set β to a fixed value of 0.5, and keep the new training dataset \mathcal{D}_n the same for each iteration n . The average imitation loss of stochastic policy sampling are evaluated after 50 stochastic roll-outs at each iterations. This average stochastic policy error tends to be higher compared to the empirical error of the deterministic trajectory, as seen from Figure 7, and confirms our finding from Corollary 5.3.


 Figure 7. Deterministic policy error vs. average stochastic policy error for $\beta = 0.5$ and 50 roll-outs of the stochastic policies.

7. Conclusion

We formalized the problem of smooth imitation learning for online sequence prediction, which is a variant of imitation learning that uses a notion of a smooth policy class. We proposed SIMILE (Smooth IMITation LEarning), which is an iterative learning reduction approach to learning smooth policies from expert demonstrations in a continuous and dynamic environment. SIMILE utilizes an adaptive learning rate that provably allows much faster convergence compared to previous learning reduction approaches, and also enjoys better sample complexity than previous work by being fully deterministic and allowing for virtual simulation of training labels. We validated the efficiency and practicality of our approach on a setting of online camera planning.

References

- Abbeel, Pieter and Ng, Andrew Y. Apprenticeship learning via inverse reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2004.
- Abbeel, Pieter and Ng, Andrew Y. Exploration and apprenticeship learning in reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2005.
- Argall, Brenna D, Chernova, Sonia, Veloso, Manuela, and Browning, Brett. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009.
- Bagnell, J Andrew, Kakade, Sham M, Schneider, Jeff G, and Ng, Andrew Y. Policy search by dynamic programming. In *Neural Information Processing Systems (NIPS)*, 2003.
- Caruana, Rich and Niculescu-Mizil, Alexandru. An empirical comparison of supervised learning algorithms. In *International Conference on Machine Learning (ICML)*, 2006.
- Chen, Jianhui and Carr, Peter. Mimicking human camera operators. In *IEEE Winter Conference Applications of Computer Vision (WACV)*, 2015.
- Chen, Jianhui, Le, Hoang M., Carr, Peter, Yue, Yisong, and Little, James J. Learning online smooth predictors for real-time camera planning using recurrent decision trees. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- Criminisi, Antonio, Shotton, Jamie, and Konukoglu, Ender. Decision forests: A unified framework for classification, regression, density estimation, manifold learning and semi-supervised learning. *Foundations and Trends in Computer Graphics and Vision*, 7(2–3):81–227, 2012.
- Daumé III, Hal, Langford, John, and Marcu, Daniel. Search-based structured prediction. *Machine learning*, 75(3):297–325, 2009.
- Gaddam, Vamsidhar Reddy, Eg, Ragnhild, Langseth, Ragnar, Griwodz, Carsten, and Halvorsen, Pål. The cameraman operating my virtual camera is artificial: Can the machine be as good as a human&quest. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 11(4):56, 2015.
- He, He, Eisner, Jason, and Daume, Hal. Imitation learning by coaching. In *Neural Information Processing Systems (NIPS)*, 2012.
- Jain, Ashesh, Wojcik, Brian, Joachims, Thorsten, and Saxena, Ashutosh. Learning trajectory preferences for manipulators via iterative improvement. In *Neural Information Processing Systems (NIPS)*, 2013.
- Kakade, Sham and Langford, John. Approximately optimal approximate reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2002.
- Lagoudakis, Michail and Parr, Ronald. Reinforcement learning as classification: Leveraging modern classifiers. In *International Conference on Machine Learning (ICML)*, 2003.
- Langford, John and Zadrozny, Bianca. Relating reinforcement learning performance to classification performance. In *International Conference on Machine Learning (ICML)*, 2005.
- Mason, Llew, Baxter, Jonathan, Bartlett, Peter L, and Frean, Marcus. Functional gradient techniques for combining hypotheses. In *Neural Information Processing Systems (NIPS)*, 1999.
- Ratliff, Nathan, Silver, David, and Bagnell, J. Andrew. Learning to search: Functional gradient techniques for imitation learning. *Autonomous Robots*, 27(1):25–53, 2009.
- Ross, Stéphane and Bagnell, Drew. Efficient reductions for imitation learning. In *Conference on Artificial Intelligence and Statistics (AISTATS)*, 2010.
- Ross, Stephane, Gordon, Geoff, and Bagnell, J. Andrew. A reduction of imitation learning and structured prediction to no-regret online learning. In *Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.
- Srebro, Nathan, Sridharan, Karthik, and Tewari, Ambuj. Smoothness, low noise and fast rates. In *Neural Information Processing Systems (NIPS)*, 2010.
- Syed, Umar and Schapire, Robert E. A reduction from apprenticeship learning to classification. In *Neural Information Processing Systems (NIPS)*, 2010.
- Wold, Herman. A study in the analysis of stationary time series, 1939.