

1. Policy Evaluation (Online Bellman Residual)

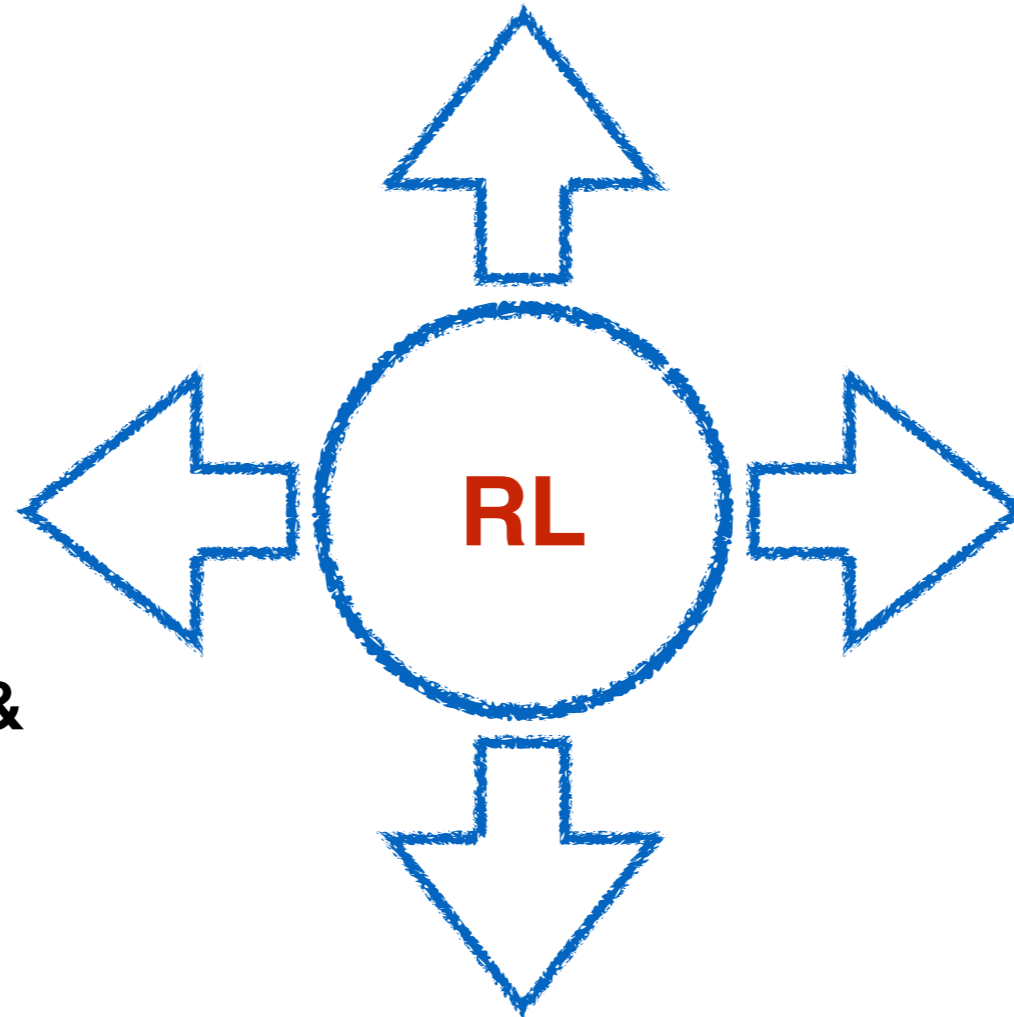
[Sun & Bagnell, 15, UAI (Best Student Paper)]

Function Approximation

2. RL via Imitation (Imitation Learning)

[Sun et.al 17, ICML; 18, ICLR]

**Function Approximation &
Imitation**



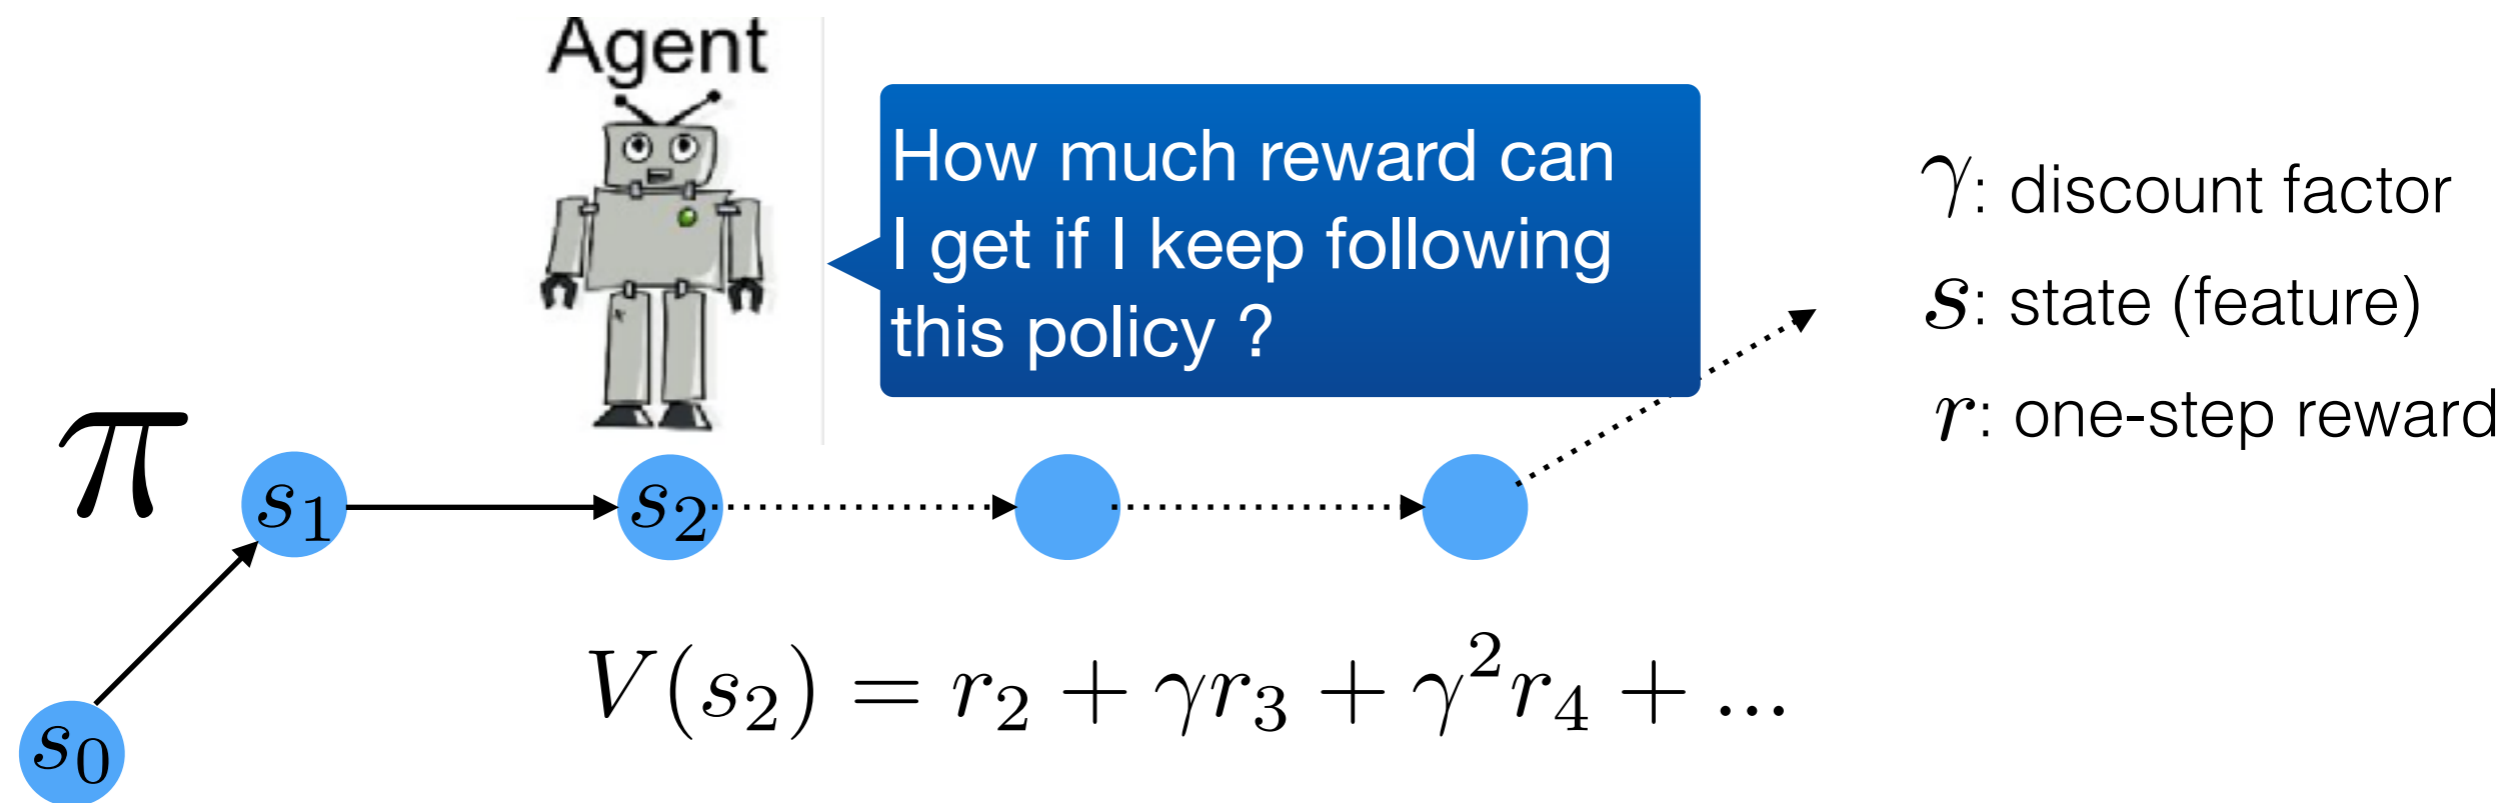
3. RL via Indirect Imitation (Dual Policy Iteration)

[Sun et.al, 18, submitted to ICML]

**Function Approximation
Optimal Control**

4. Proposed Work: Temporal Difference Learning & Apprenticeship Learning

Policy Evaluation



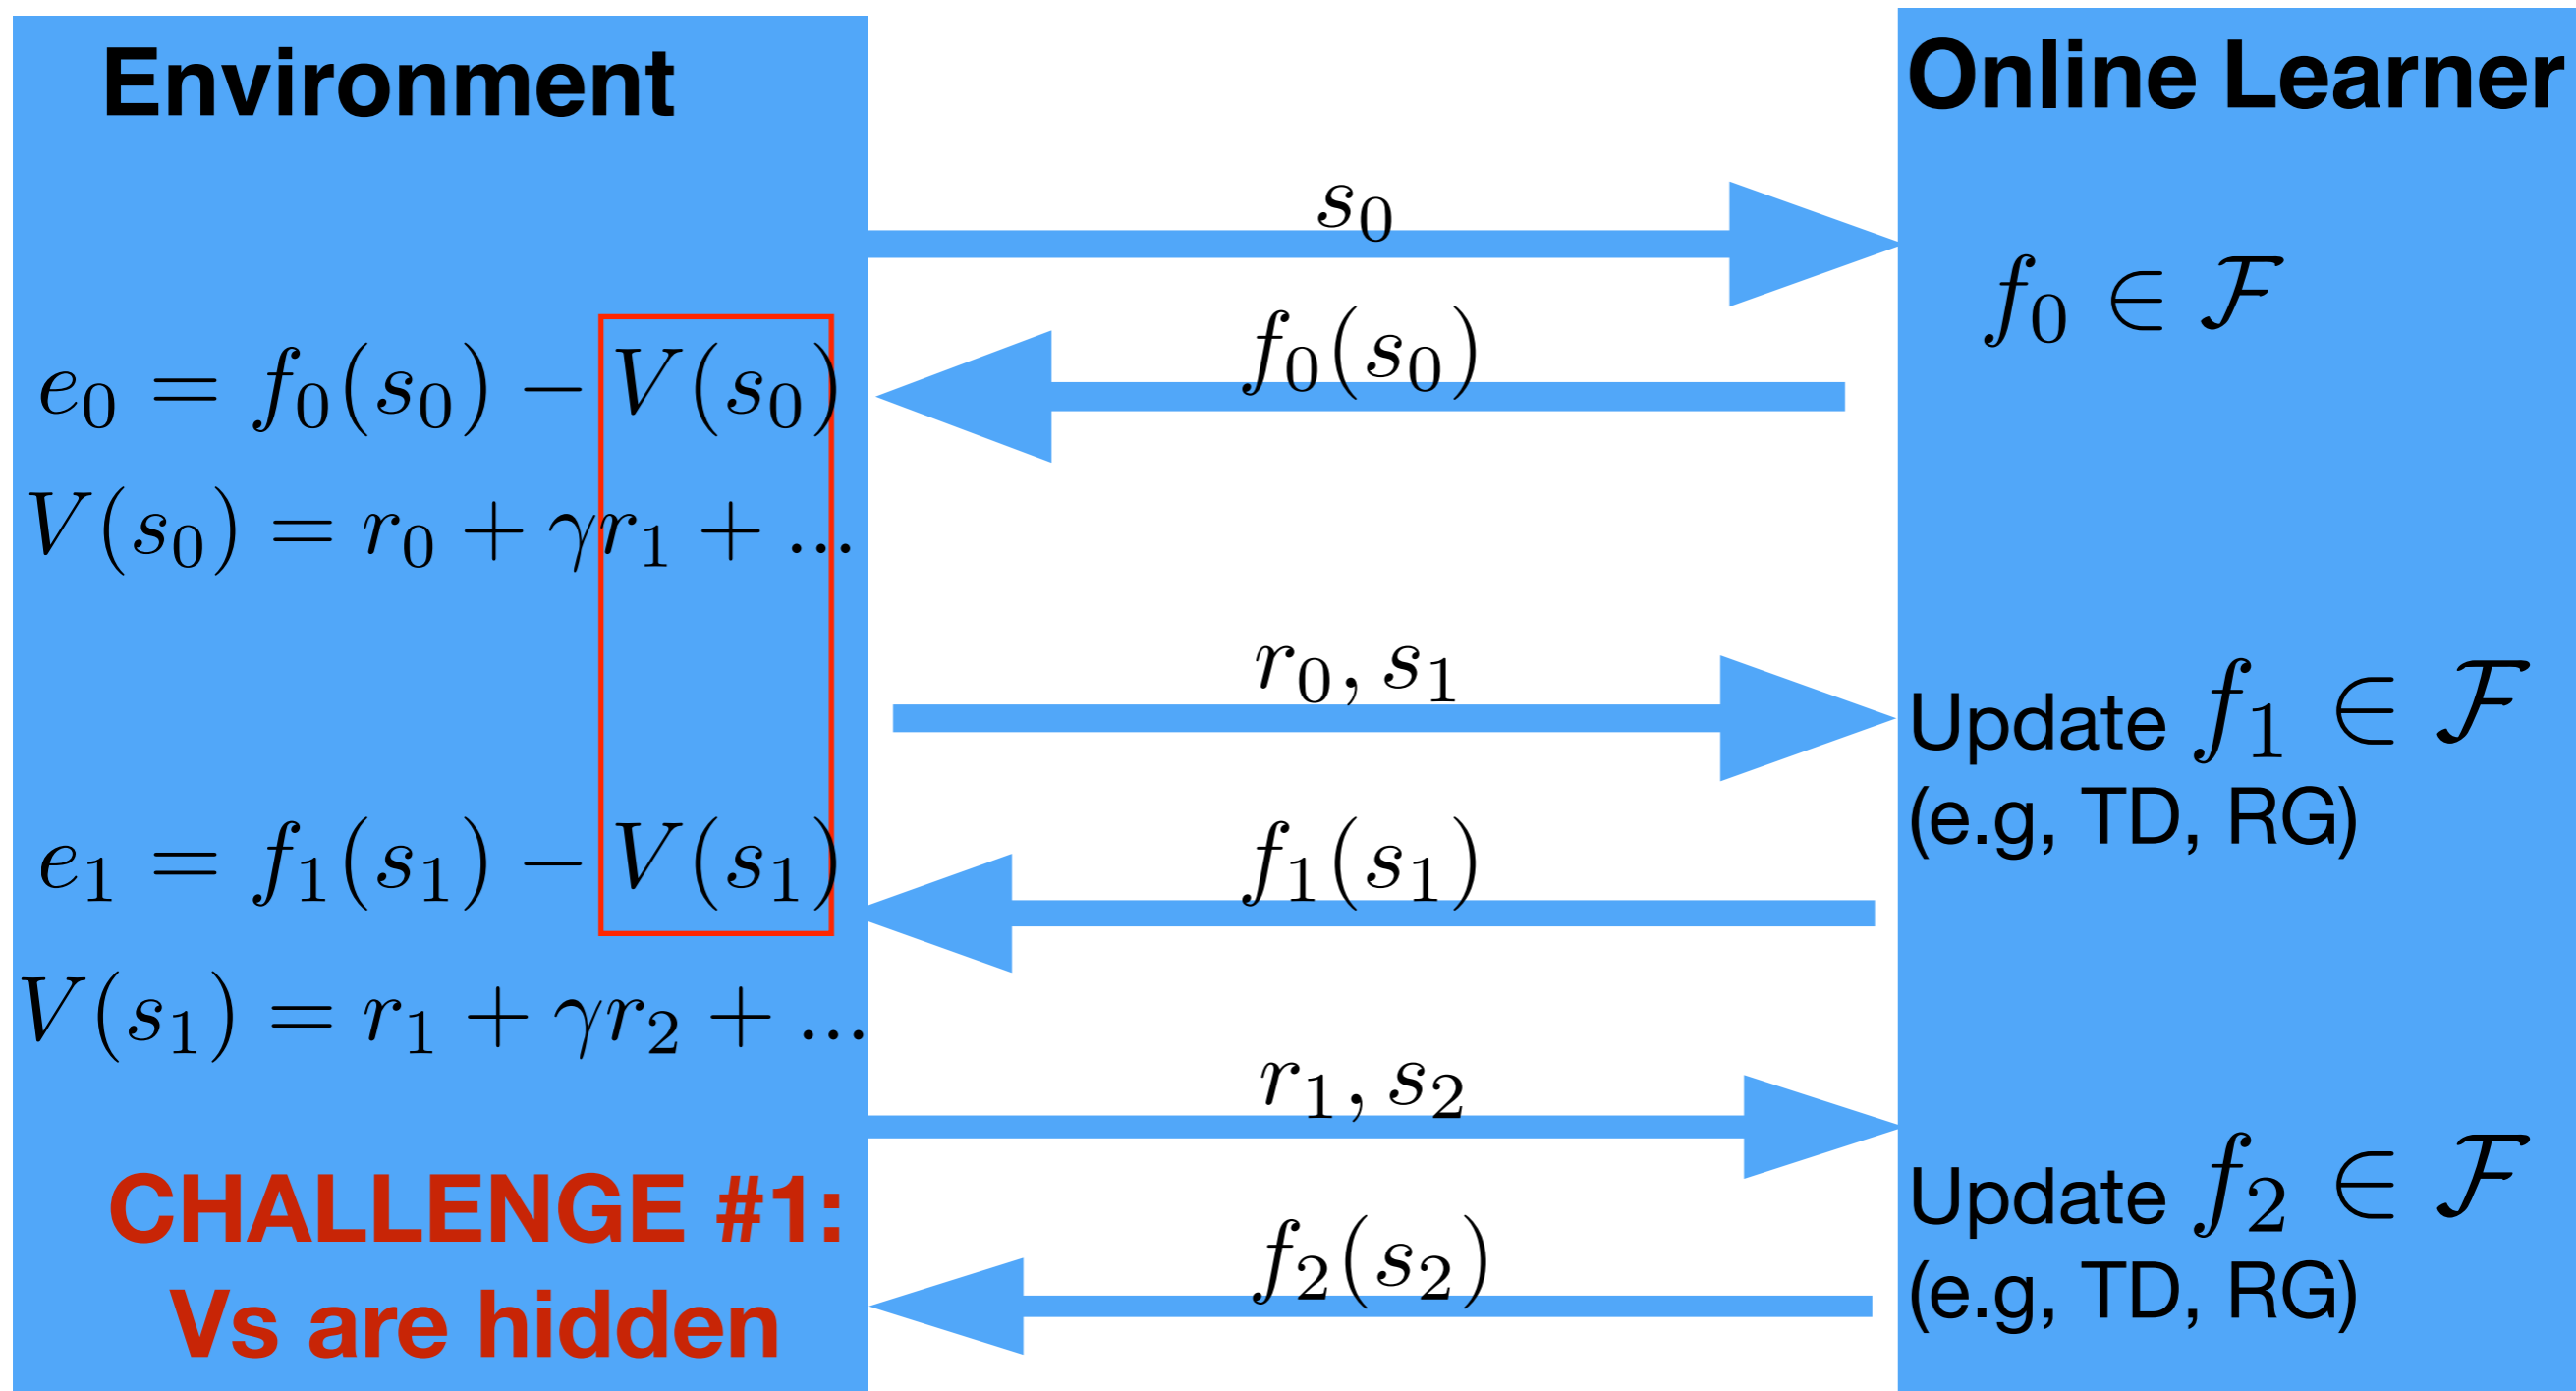
Predict **Reward-to-go** $\sum_t \gamma^t r_t$

$$f(s) \approx \sum_t \gamma^t r_t$$

Temporal Difference (TD) [Sutton, 1988]
Residual Gradient (RG) [Baird, 1995]

Sequential Online Learning Setting

[Schapire & Warmuth 96, Li 2008]



CHALLENGE #2: No statistical assumption
(e.g., Non-Markovian)

Goal

Goal: minimize the Online **Prediction Error (PE)**:

$$\sum_t e_t^2 = \sum_t (f_t(s_t) - V(s_t))^2$$

Batch **PE**:

$$\sum_t e_t^{*2} = \sum_t (f^*(s_t) - V(s_t))^2 \quad f^* \in \mathcal{F}$$

Best Hypothesis in hindsight

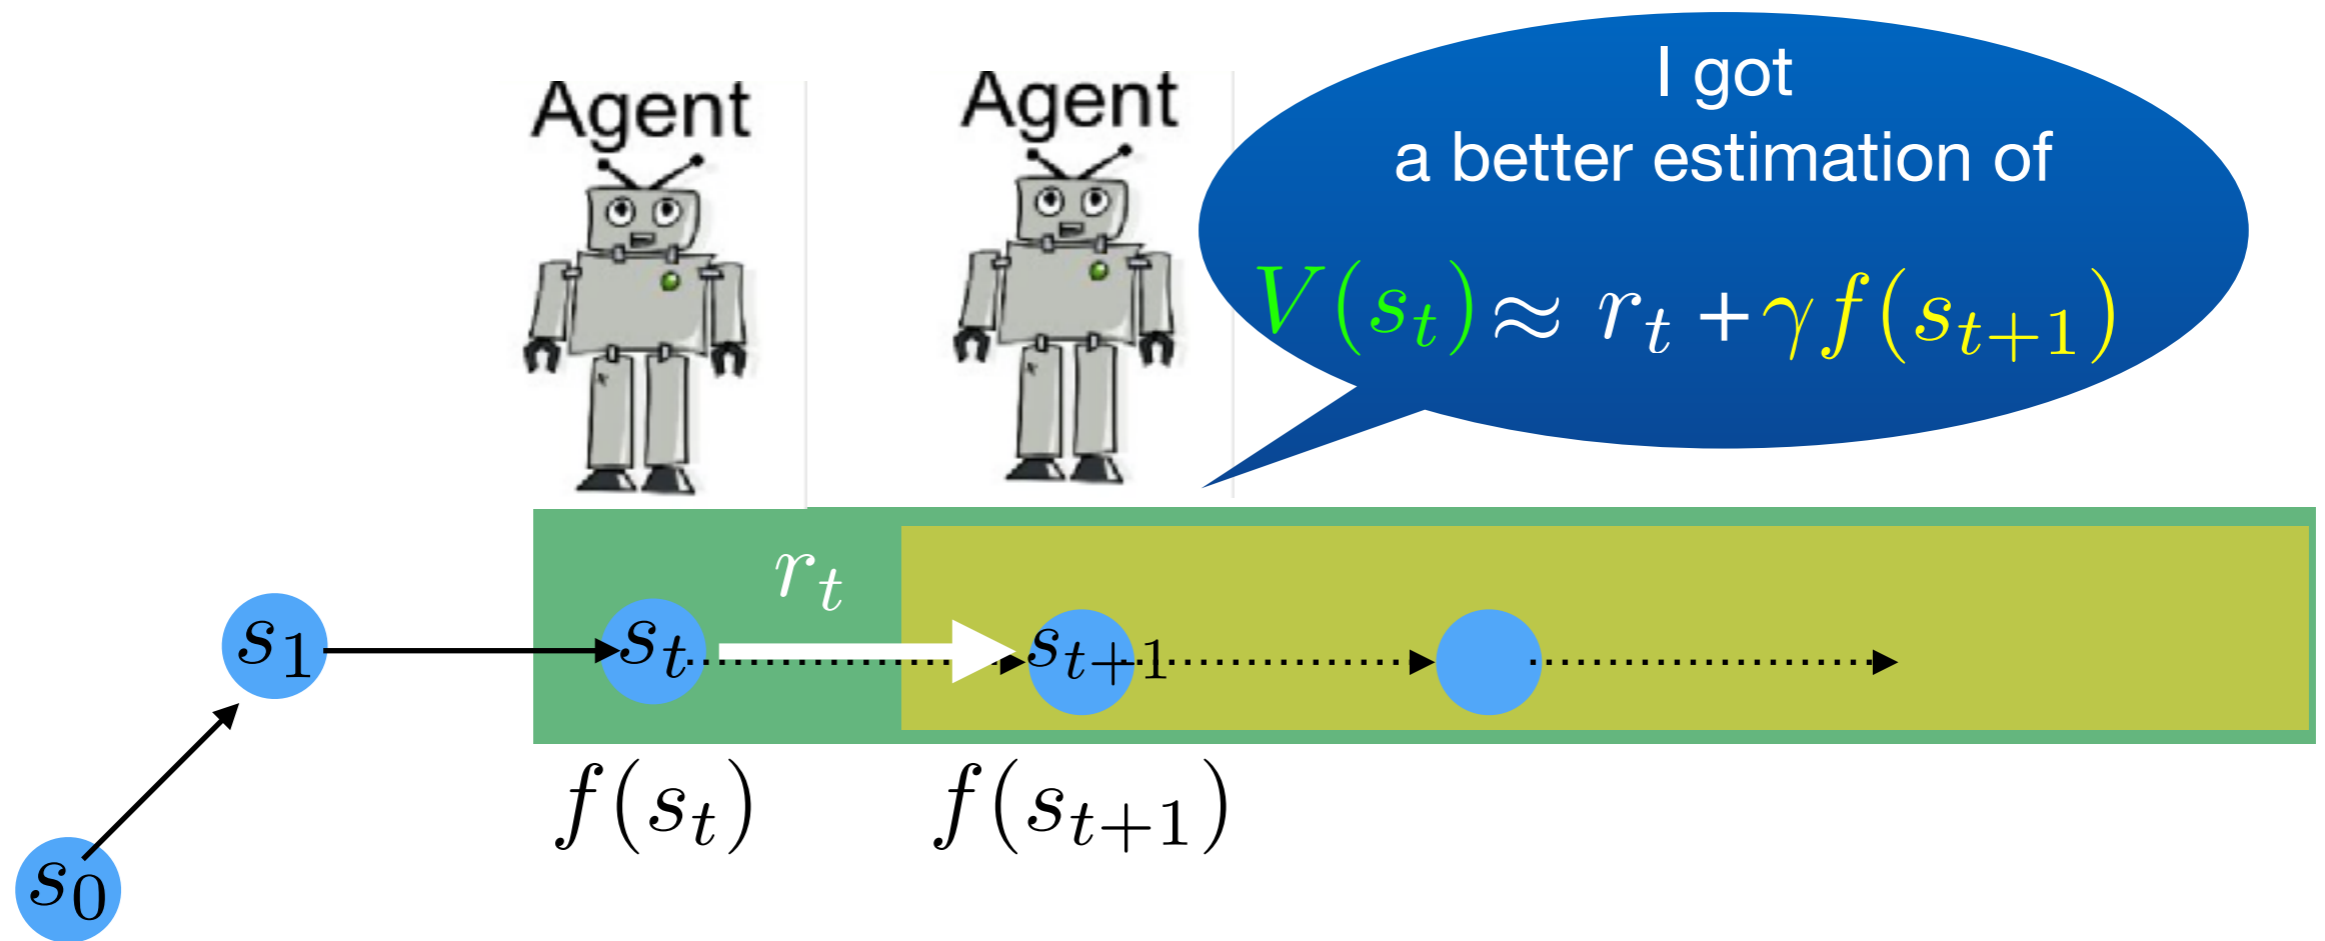
Average Online PE

$$\frac{1}{T} \sum_t e_t^2 \leq c \frac{1}{T} \sum_t e_t^{*2}, \quad T \rightarrow \infty$$

Smallest possible Batch PE

Bellman Loss

Bellman Loss: $l_t(f) = (f(s_t) - r_t - \gamma f(s_{t+1}))^2$

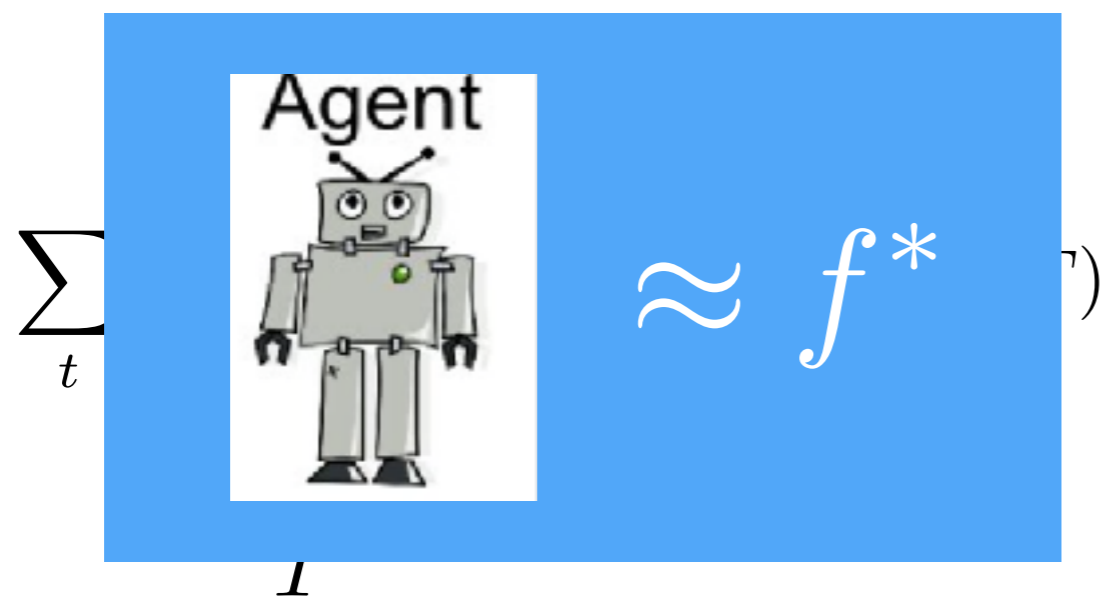
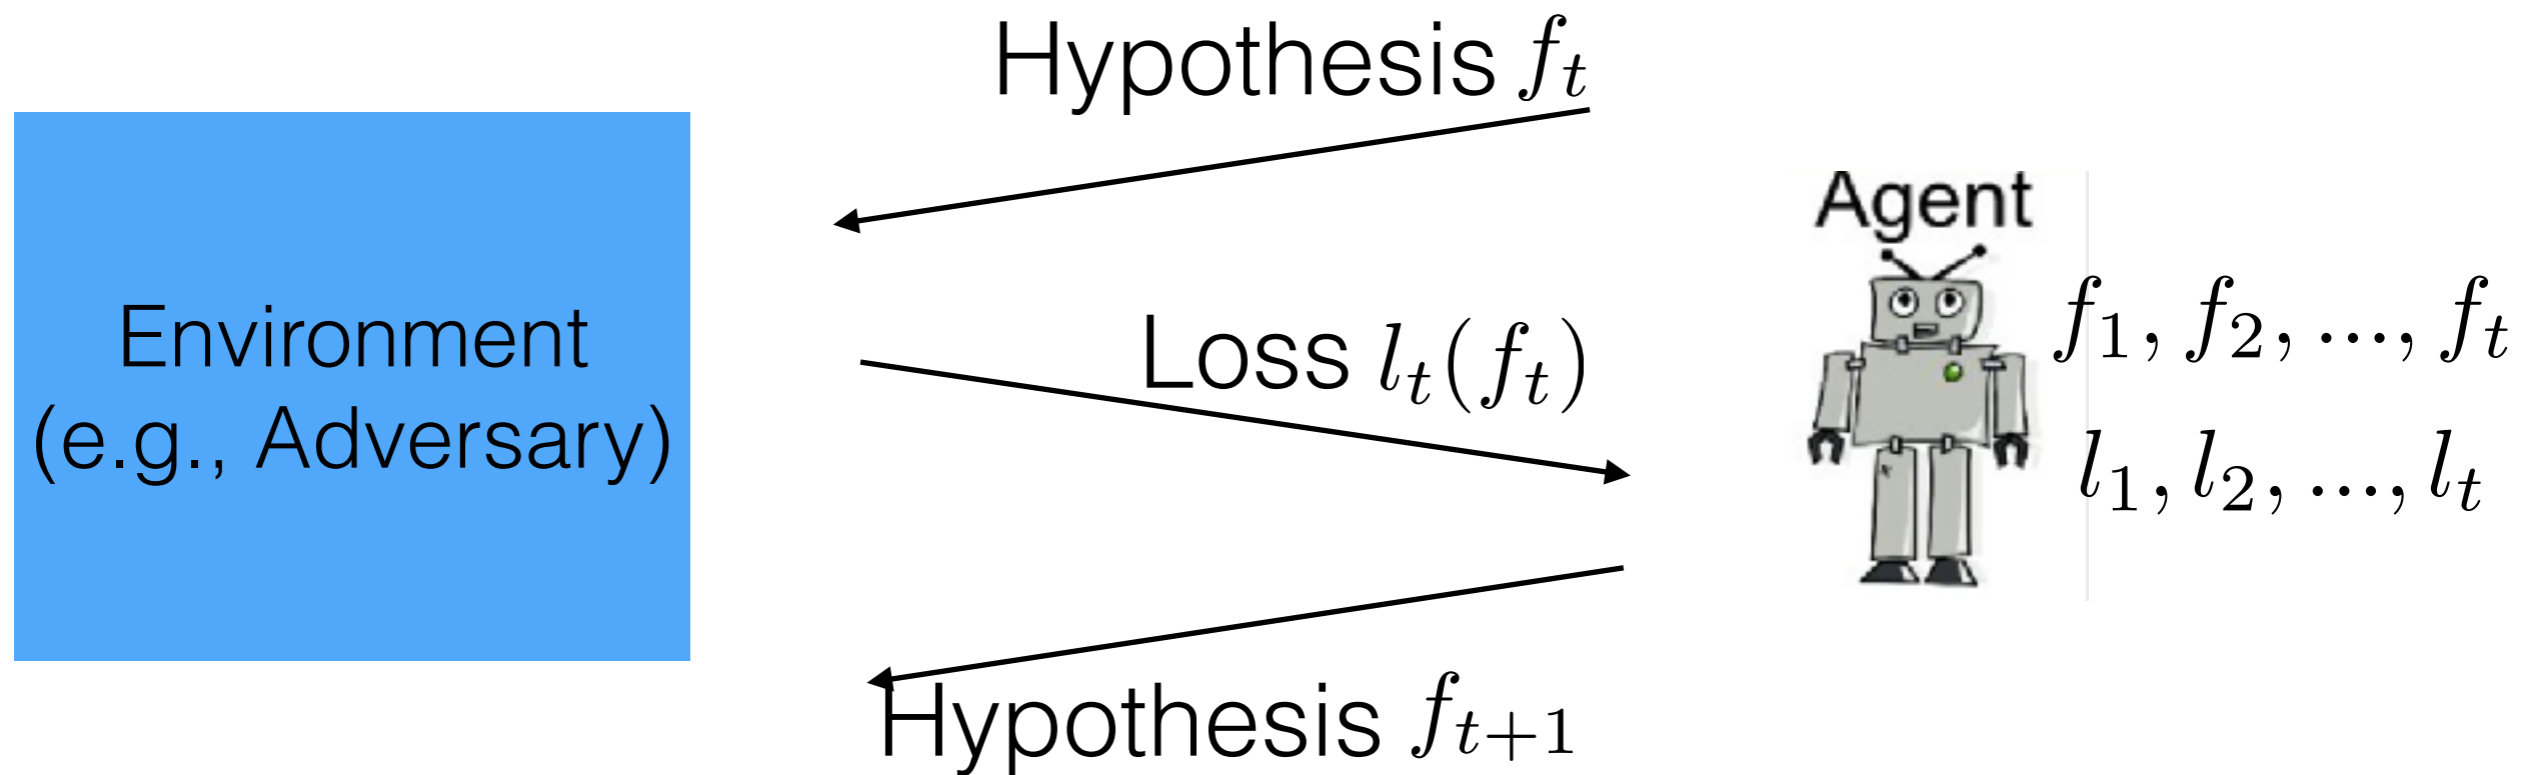


$$f(s_t) - V(s_t) \approx f(s_t) - (r_t + \gamma f(s_{t+1}))$$

Bootstrap

No-regret Online Learning

[Gordon, 99, COLT; Zinkevich, 03, ICML; Shalve-Schwartz, 12]



Online Stability:

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_t \|f_{t+1} - f_t\|^2 = 0$$

[Ross & Bagnell 2011, Saha, 2012]

Reduction to No-Regret and Stable Online Learning

Recall Bellman Loss at time step t

$$l_t(f) = (f(s_t) - r_t - \gamma f(s_{t+1}))^2$$

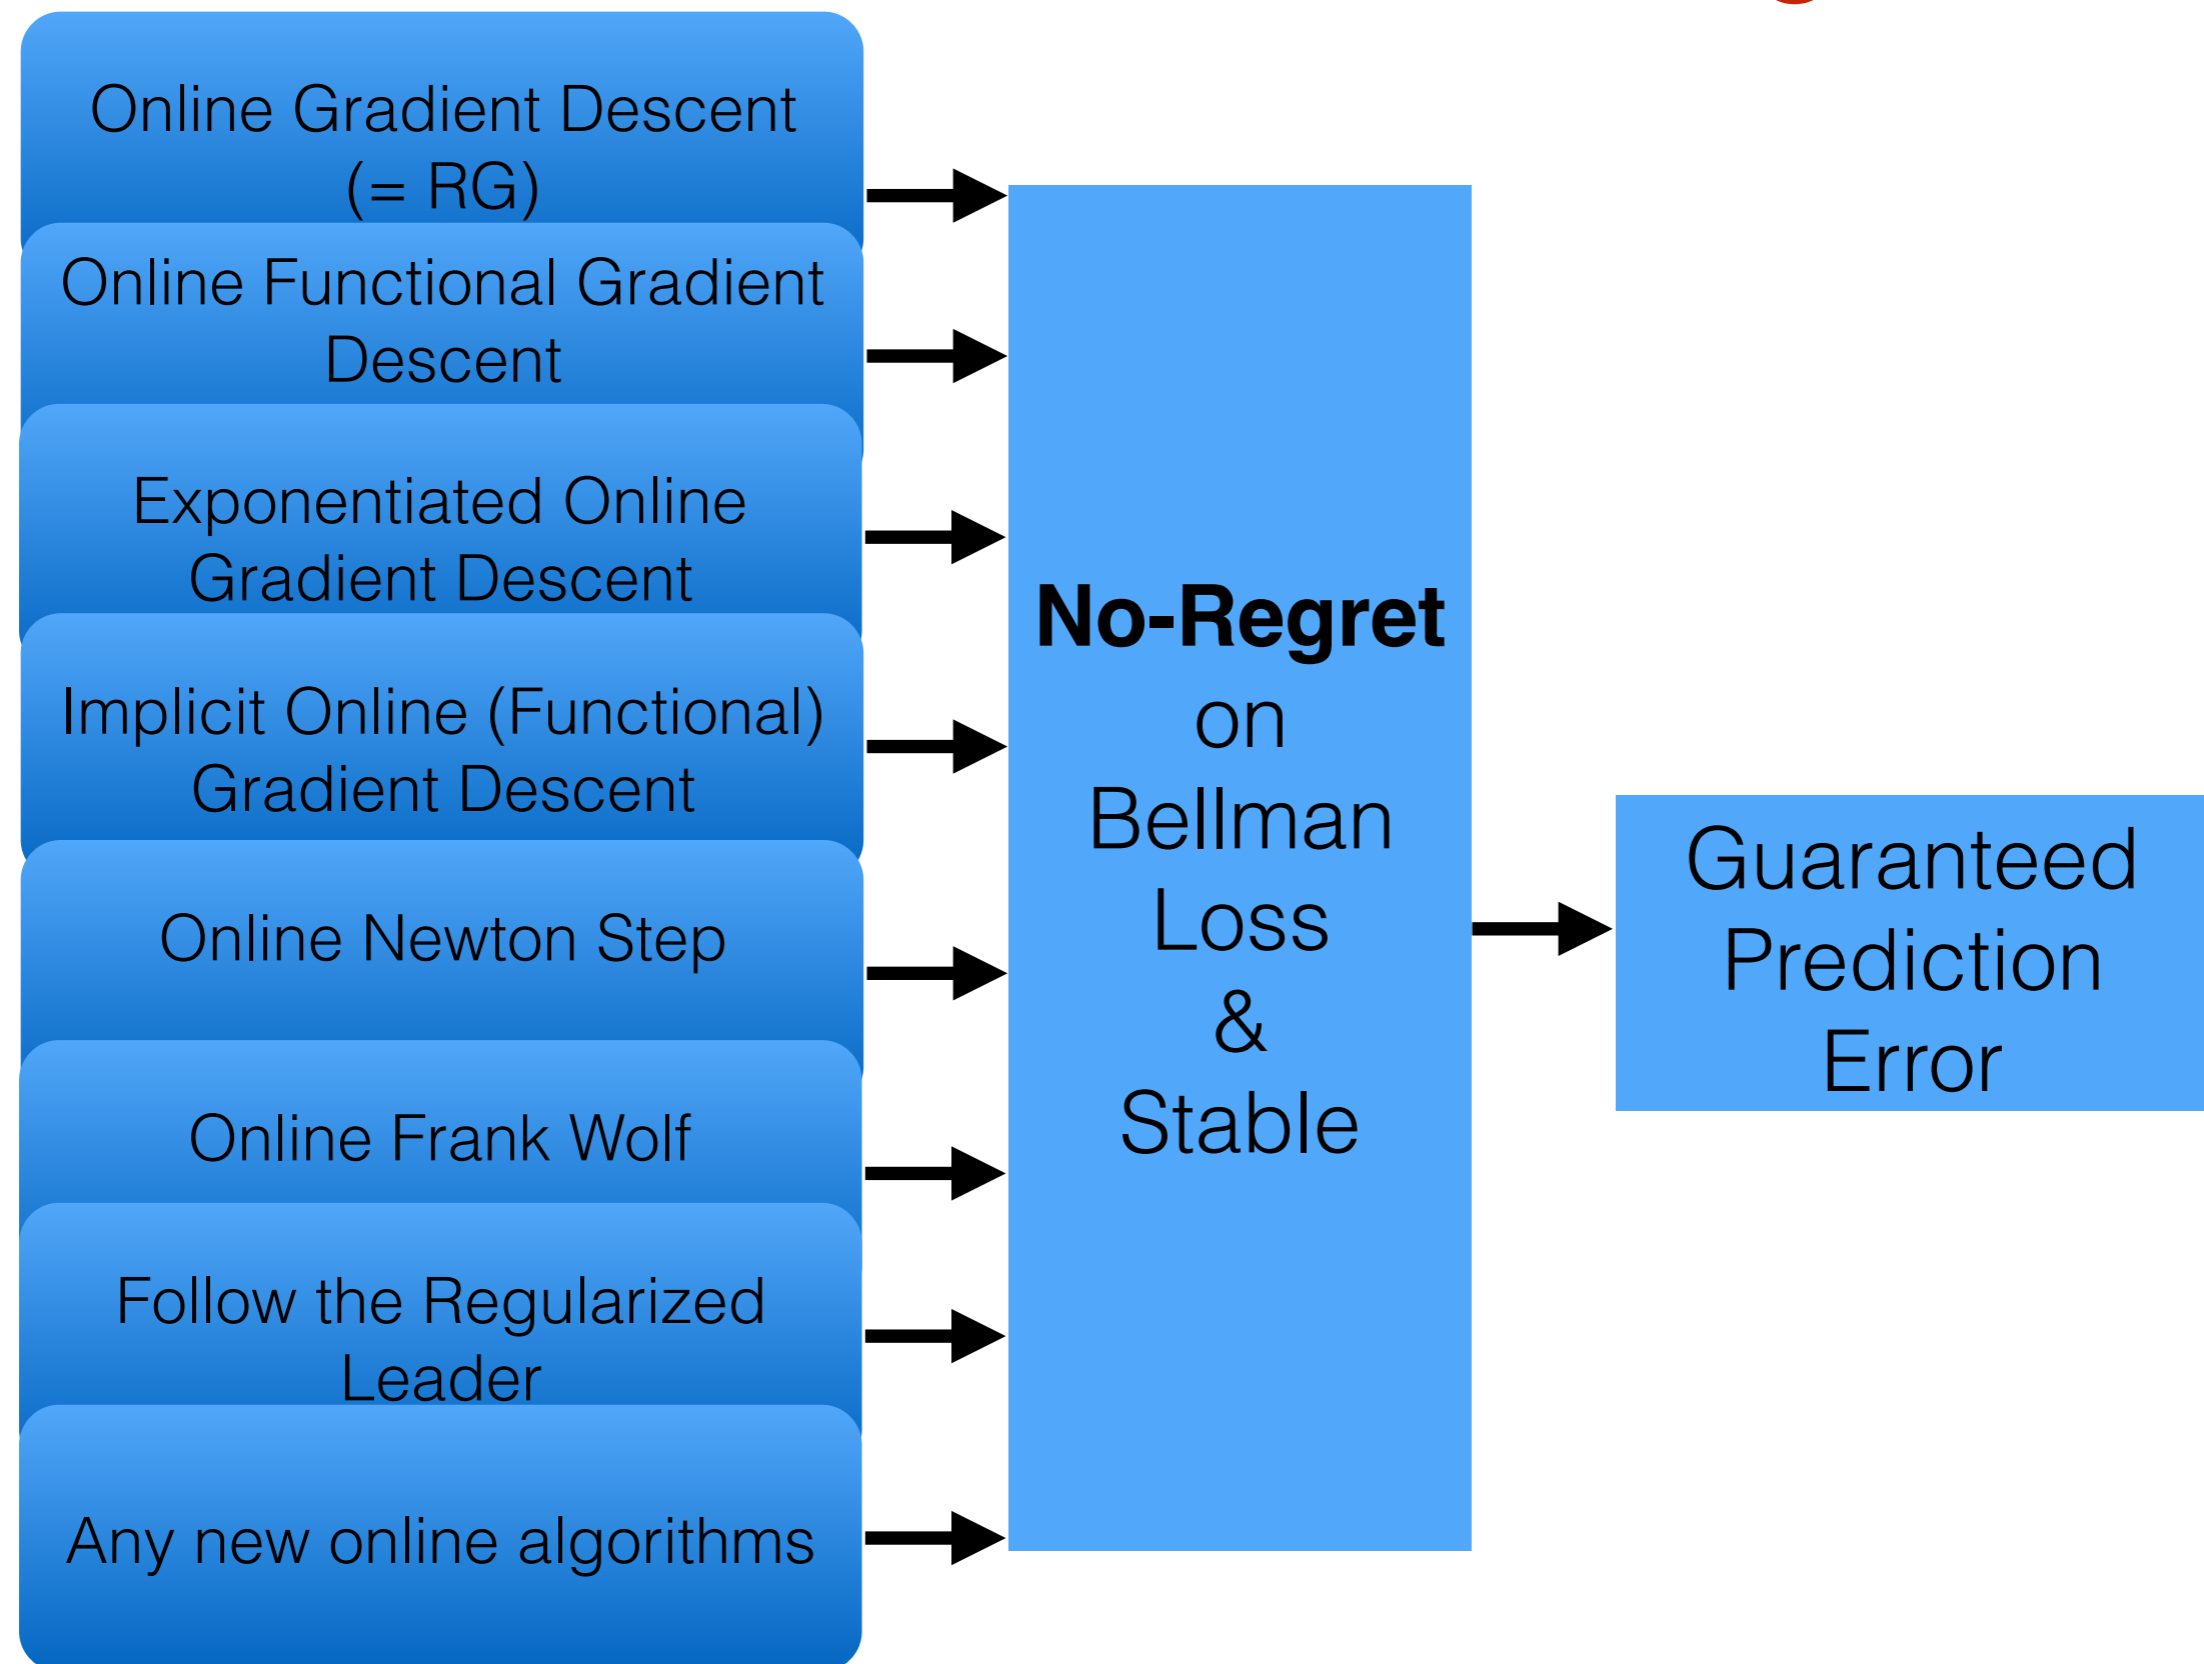
No-Regret & Stable

$$l_1(f), l_2(f), \dots, l_T(f)$$

Lead to

$$\frac{1}{T} \sum_t e_t^2 \leq \frac{1}{(1-\gamma)^2} \frac{1}{T} \sum_t e_t^{*2}, \quad T \rightarrow \infty$$

Reduction Leads to a Set of Algorithms



Summary

Message #1:

Agnostic Performance Guarantee with function approximation

Message #2:

Generalization and Efficiency of Policy Evaluation
via Reduction to No-Regret Online Learning