Computing Fixpoints of Learned Functions
Chaotic Iteration and Simple Stochastic Games

Florian Wittbold
Universität Duisburg-Essen

Joint work with Paolo Baldan, Sebastian Gurke, Barbara König

TACAS 2026

April 14, 2026

Fixpoints in Computer Science

General Task

Given a function $f\colon L \to L$ over a partially ordered set $(L, \sqsubseteq)$, compute its least fixpoint μf, i.e., the least element $x \in L$ fulfilling

$$f(x) = x.$$

Fixpoints in Computer Science

General Task

Given a function $f\colon L \to L$ over a partially ordered set $(L, \sqsubseteq)$, compute its least fixpoint μf, i.e., the least element $x \in L$ fulfilling

$$f(x) = x.$$

Applications in

- concurrency theory (behavioural equivalences and metrics)

Fixpoints in Computer Science

General Task

Given a function $f\colon L \to L$ over a partially ordered set $(L, \sqsubseteq)$, compute its least fixpoint μf, i.e., the least element $x \in L$ fulfilling

$$f(x) = x.$$

Applications in

- concurrency theory (behavioural equivalences and metrics)
- model checking (μ-calculus)

Fixpoints in Computer Science

General Task

Given a function $f\colon L \to L$ over a partially ordered set $(L, \sqsubseteq)$, compute its least fixpoint μf, i.e., the least element $x \in L$ fulfilling

$$f(x) = x.$$

Applications in

- concurrency theory (behavioural equivalences and metrics)
- model checking (μ-calculus)
- program analysis (dataflow analysis)
- ...

Fixpoints in Computer Science

General Task

Given a function $f\colon L \to L$ over a partially ordered set $(L, \sqsubseteq)$, compute its least fixpoint μf, i.e., the least element $x \in L$ fulfilling

$$f(x) = x.$$

Applications in

- concurrency theory (behavioural equivalences and metrics)
- model checking (μ-calculus)
- program analysis (dataflow analysis)
- ...

Question: What happens if the function f is not known but only approximated?

Case Study: Simple Stochastic Games

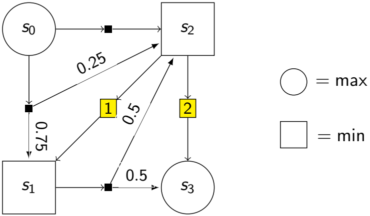

Case Study: Simple Stochastic Games

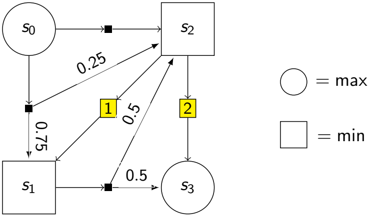

Optimal behavior through state-action values of SSG as least fixpoint of Bellman optimality operator $f : \mathbb{R}_{\geq 0}^{S \times A} \to \mathbb{R}_{\geq 0}^{S \times A}$

$$f(q)(s,a) = R(s,a) + \sum_{s' \in S} T(s' \mid s, a) \operatorname*{opt}_{a' \in A(s')} q(s', a')$$

where opt $\in \{\max, \min\}$ depending on active player at state s'.

Case Study: Simple Stochastic Games

New Task: Compute μf of a system when the transition probabilities (given by T) and the rewards (given by R) are not known precisely but can only be approximated?

Case Study: Simple Stochastic Games

New Task: Compute μf of a system when the transition probabilities (given by T) and the rewards (given by R) are not known precisely but can only be approximated?

Sources of approximation:

- Unknown / partially known system
- Changes to the system
- Approximated partial results
- ...

Case Study: Simple Stochastic Games

New Task: Compute μf of a system when the transition probabilities (given by T) and the rewards (given by R) are not known precisely but can only be approximated?

Sources of approximation:

- Unknown / partially known system
- Changes to the system
- Approximated partial results
- ...

$\Rightarrow$ e.g., (model-based) Reinforcement Learning

Non-expansive Setting

Task

Given a sequence of monotone and non-expansive functions $f_0, f_1, \dots : \mathbb{R}_+^d \to \mathbb{R}_+^d$ that converges (pointwise) to $f : \mathbb{R}_+^d \to \mathbb{R}_+^d$:

Compute a sequence $x_0, x_1, x_2, \dots \in \mathbb{R}_+^d$ that converges to μf where x_n may only depend on $f_0, \dots, f_{n-1}$.

Non-expansive Setting

Task

Given a sequence of monotone and non-expansive functions $f_0, f_1, \dots : \mathbb{R}_+^d \to \mathbb{R}_+^d$ that converges (pointwise) to $f : \mathbb{R}_+^d \to \mathbb{R}_+^d$:

Compute a sequence $x_0, x_1, x_2, \dots \in \mathbb{R}_+^d$ that converges to μf where x_n may only depend on $f_0, \dots, f_{n-1}$.

Remark: Non-expansivity is w.r.t. the supremum norm $\| \cdot \|_\infty$, covering many interesting cases: termination probabilities and non-discounted returns in general simple stochastic games, behavioural metrics,. . .

Non-expansive Setting

Task

Given a sequence of monotone and non-expansive functions $f_0, f_1, \dots : \mathbb{R}_+^d \to \mathbb{R}_+^d$ that converges (pointwise) to $f : \mathbb{R}_+^d \to \mathbb{R}_+^d$:

Compute a sequence $x_0, x_1, x_2, \dots \in \mathbb{R}_+^d$ that converges to μf where x_n may only depend on $f_0, \dots, f_{n-1}$.

Remark: Non-expansivity is w.r.t. the supremum norm $\|\cdot\|_\infty$, covering many interesting cases: termination probabilities and non-discounted returns in general simple stochastic games, behavioural metrics,...

Will a standard Kleene iteration ($x_{n+1} = f_n(x_n), x_0 = 0$) still work? Or do we need new techniques?

The Problem with Over-Approximations

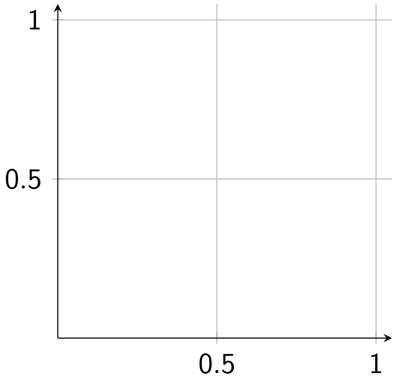

$$f(x) = \max(x, 0.5)$$

The Problem with Over-Approximations

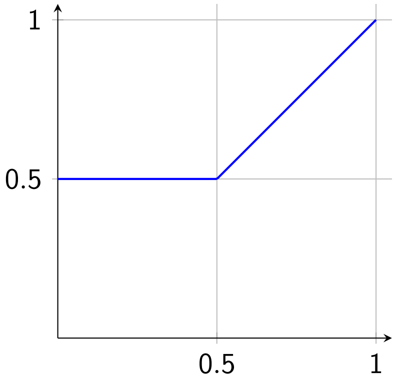

$$f(x) = \max(x, 0.5)$$

The Problem with Over-Approximations

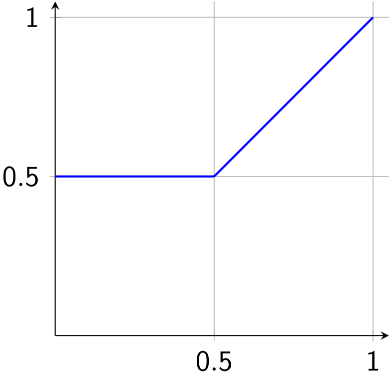

$$f(x) = \max(x, 0.5)$$

The Problem with Over-Approximations

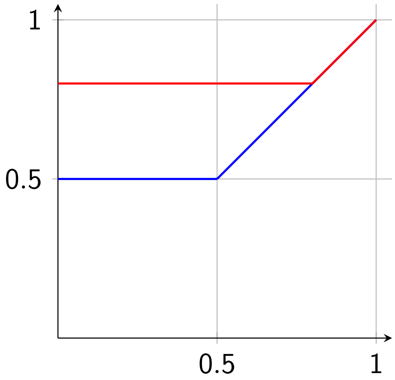

$$f(x) = \max(x, 0.8)$$

The Problem with Over-Approximations

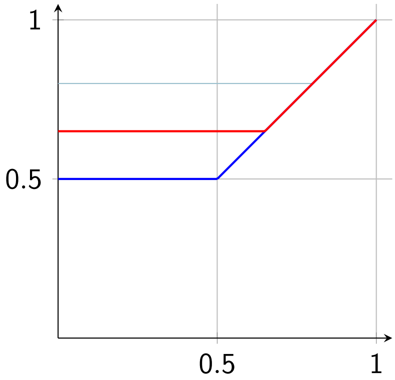

$$f(x) = \max(x, 0.65)$$

The Problem with Over-Approximations

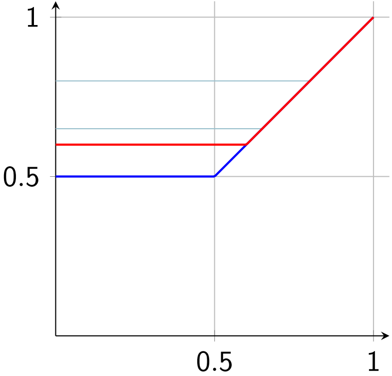

$$f(x) = \max(x, 0.6)$$

The Problem with Over-Approximations

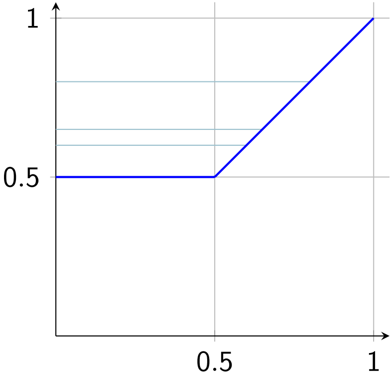

$$f_n(x) = \max(x, 0.5 + 0.3/n)$$

The Problem with Over-Approximations

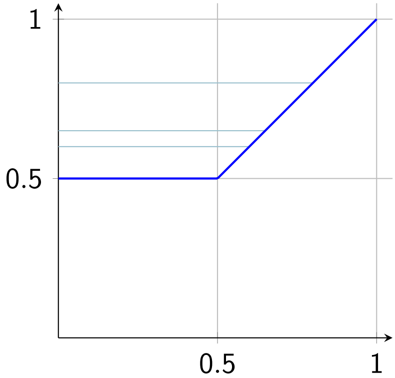

$$f_n(x) = \max(x, 0.5 + 0.3/n)$$

We have $\mu f_n = 0.5 + 0.3/n \to 0.5 = \mu f$, but iterating with any f_n will over-estimate μf and further iterations will never decrease it.

Introducing the Dampened Mann Iteration

(Kleene) Iteration

$$x_{n+1} = f_n(x_n)$$

Introducing the Dampened Mann Iteration

Dampened Mann Iteration

$$x_{n+1} = (1 - \beta_n)(x_n + \alpha_n(f_n(x_n) - x_n))$$

Introducing the Dampened Mann Iteration

Dampened Mann Iteration

$$x_{n+1} = (1 - \beta_n)(x_n + \alpha_n(f_n(x_n) - x_n))$$

The parameters $\beta_n \in [0, 1]$ are *dampening factors*, the $\alpha_n \in [0, 1]$ *learning rates* fulfilling

$$(1) \lim_{n\to\infty} {}^{\beta_n}/_{\alpha_n} = 0 \qquad (2) \sum_{n\to\infty} \beta_n = \infty$$

Introducing the Dampened Mann Iteration

Dampened Mann Iteration

$$x_{n+1} = (1 - \beta_n)(x_n + \alpha_n(f_n(x_n) - x_n))$$

The parameters $\beta_n \in [0,1]$ are *dampening factors*, the $\alpha_n \in [0,1]$ *learning rates* fulfilling

$$(1) \lim_{n\to\infty} \beta_n/\alpha_n = 0 \qquad (2) \sum_{n\to\infty} \beta_n = \infty$$

Intuition:

(1) Keep update large enough to counteract dampening

(2) Ensure enough 'dampening power' left to decrease over-estimation

Main Contributions

$$x_{n+1} = (1 - \beta_n)(x_n + \alpha_n(f_n(x_n) - x_n))$$

Main Contributions

$$x_{n+1} = (1 - \beta_n)(x_n + \alpha_n(f_n(x_n) - x_n))$$

Proofs of Convergence

- Case $\alpha_n \to 0$

Main Contributions

$$x_{n+1} = (1 - \beta_n)(x_n + \alpha_n(f_n(x_n) - x_n))$$

Proofs of Convergence

- Case $\alpha_n \to 0$
  - $\Rightarrow$ not a Kleene iteration in the limit

Main Contributions

$$x_{n+1} = (1 - \beta_n)(x_n + \alpha_n(f_n(x_n) - x_n))$$

Proofs of Convergence

- Case $\alpha_n \to 0$
  - $\Rightarrow$ not a Kleene iteration in the limit
  - $\Rightarrow$ required completely new proofs

Main Contributions

$$x_{n+1} = (1 - \beta_n)(x_n + \alpha_n(f_n(x_n) - x_n))$$

Proofs of Convergence

- Case $\alpha_n \to 0$
  - $\Rightarrow$ not a Kleene iteration in the limit
  - $\Rightarrow$ required completely new proofs
- For chaotic iteration

Main Contributions

$$x_{n+1} = (1 - \beta_n)(x_n + \alpha_n(f_n(x_n) - x_n))$$

Proofs of Convergence

- Case $\alpha_n \to 0$
  - $\Rightarrow$ not a Kleene iteration in the limit
  - $\Rightarrow$ required completely new proofs
- For chaotic iteration
  - $\Rightarrow$ even with component-dependent parameters

Main Contributions

$$x_{n+1} = (1 - \beta_n)(x_n + \alpha_n(f_n(x_n) - x_n))$$

Proofs of Convergence

- Case $\alpha_n \to 0$
  - $\Rightarrow$ not a Kleene iteration in the limit
  - $\Rightarrow$ required completely new proofs
- For chaotic iteration
  - $\Rightarrow$ even with component-dependent parameters
- For simple stochastic games

Main Contributions

$$x_{n+1} = (1 - \beta_n)(x_n + \alpha_n(f_n(x_n) - x_n))$$

Proofs of Convergence

- Case $\alpha_n \to 0$
  - $\Rightarrow$ not a Kleene iteration in the limit
  - $\Rightarrow$ required completely new proofs
- For chaotic iteration
  - $\Rightarrow$ even with component-dependent parameters
- For simple stochastic games
  - $\Rightarrow$ without using the concept of end components

Comparison to Previous Results [CAV'25]

Convergence of the dampened Mann iteration to μf holds if

- the parameter sequences fulfill

$$\lim_{n\to\infty} \alpha_n \in (0,1], \quad \lim_{n\to\infty} \beta_n = 0, \text{ and } \quad \sum_{n\in\mathbb{N}_0} \beta_n = \infty$$

- the least fixpoint $\mu f < \infty$ is bounded
- one of the following properties hold:
  - (f_n) converges monotonically and $\mu f = \lim_{n\to\infty} \mu f_n$
  - (f_n) converges normally to f, i.e., $\sum_{n\in\mathbb{N}_0} \|f - f_n\|_\infty < \infty$
  - $\alpha_n \to 1$ and (f_n) is a (faultless) sampling of an MDP

Comparison to Previous Results [CAV'25]

Convergence of the dampened Mann iteration to μf holds if

- the parameter sequences fulfill

$$\cancel{\lim_{n\to\infty} \alpha_n \in (0,1],} \quad \lim_{n\to\infty} {}^{\beta_n}/_{\alpha_n} = 0, \text{ and } \quad \sum_{n\in\mathbb{N}_0} \beta_n = \infty$$

- the least fixpoint $\mu f < \infty$ is bounded
- one of the following properties hold:
  - (f_n) converges monotonically and $\mu f = \lim_{n\to\infty} \mu f_n$
  - (f_n) converges normally to f, i.e., $\sum_{n\in\mathbb{N}_0} \|f - f_n\|_\infty < \infty$
  - $\alpha_n \to 1$ and (f_n) is a (faultless) sampling of an MDP

Comparison to Previous Results [CAV'25]

Convergence of the dampened Mann iteration to μf holds if

- the parameter sequences fulfill

$$\lim_{n\to\infty} {}^{\beta_n}/\alpha_n = 0 \quad \text{and} \quad \sum_{n\in\mathbb{N}_0} \beta_n = \infty$$

- the least fixpoint $\mu f < \infty$ is bounded
- one of the following properties hold:
  - (f_n) converges monotonically and $\mu f = \lim_{n\to\infty} \mu f_n$
  - (f_n) converges normally to f, i.e., $\sum_{n\in\mathbb{N}_0} \|f - f_n\|_\infty < \infty$
  - $\alpha_n \to 1$ and (f_n) is a (faultless) sampling of an MDP

Comparison to Previous Results [CAV'25]

Convergence of the dampened Mann iteration to μf holds if

- the parameter sequences fulfill

$$\lim_{n\to\infty} {}^{\beta_n}/\alpha_n = 0 \quad \text{and} \quad \sum_{n\in\mathbb{N}_0} \beta_n = \infty$$

- ~~the least fixpoint $\mu f < \infty$ is bounded~~
- one of the following properties hold:
  - (f_n) converges monotonically and $\mu f = \lim_{n\to\infty} \mu f_n$
  - (f_n) converges normally to f, i.e., $\sum_{n\in\mathbb{N}_0} \|f - f_n\|_\infty < \infty$
  - $\alpha_n \to 1$ and (f_n) is a (faultless) sampling of an MDP

Comparison to Previous Results [CAV'25]

Convergence of the dampened Mann iteration to μf holds if

- the parameter sequences fulfill

$$\lim_{n\to\infty} {}^{\beta_n}/_{\alpha_n} = 0 \quad \text{and} \quad \sum_{n\in\mathbb{N}_0} \beta_n = \infty$$

- one of the following properties hold:
  - (f_n) converges monotonically and $\mu f = \lim_{n\to\infty} \mu f_n$
  - (f_n) converges normally to f, i.e., $\sum_{n\in\mathbb{N}_0} \|f - f_n\|_\infty < \infty$
  - $\alpha_n \to 1$ and (f_n) is a (faultless) sampling of an MDP

Comparison to Previous Results [CAV'25]

Convergence of the dampened Mann iteration to μf holds if

- the parameter sequences fulfill

$$\lim_{n\to\infty} {}^{\beta_n}/\alpha_n = 0 \quad \text{and} \quad \sum_{n\in\mathbb{N}_0} \beta_n = \infty$$

- one of the following properties hold:
  - ~~(f_n) converges monotonically and~~ $\mu f = \lim_{n\to\infty} \mu \sup_{k>n} f_k$
  - (f_n) converges normally to f, i.e., $\sum_{n\in\mathbb{N}_0} \|f - f_n\|_\infty < \infty$
  - $\alpha_n \to 1$ and (f_n) is a (faultless) sampling of an MDP

Comparison to Previous Results [CAV'25]

Convergence of the dampened Mann iteration to μf holds if

- the parameter sequences fulfill

$$\lim_{n\to\infty} {}^{\beta_n}\!/\alpha_n = 0 \quad \text{and} \quad \sum_{n\in\mathbb{N}_0} \beta_n = \infty$$

- one of the following properties hold:
  - $\mu f = \lim_{n\to\infty} \mu \sup_{k>n} f_k$
  - (f_n) converges normally to f, i.e., $\sum_{n\in\mathbb{N}_0} \|f - f_n\|_\infty < \infty$
  - $\alpha_n \to 1$ and (f_n) is a (faultless) sampling of an MDP

Comparison to Previous Results [CAV'25]

Convergence of the dampened Mann iteration to μf holds if

- the parameter sequences fulfill

$$\lim_{n\to\infty} {}^{\beta_n}/_{\alpha_n} = 0 \quad \text{and} \quad \sum_{n\in\mathbb{N}_0} \beta_n = \infty$$

- one of the following properties hold:
  - $\mu f = \lim_{n\to\infty} \mu \sup_{k>n} f_k$
  - (f_n) converges α-normally to f, i.e., $\sum_{n\in\mathbb{N}_0} \alpha_n \|f - f_n\|_\infty < \infty$
  - $\alpha_n \to 1$ and (f_n) is a (faultless) sampling of an MDP

Comparison to Previous Results [CAV'25]

Convergence of the dampened Mann iteration to μf holds if

- the parameter sequences fulfill

$$\lim_{n\to\infty} {}^{\beta_n}\!/_{\alpha_n} = 0 \quad \text{and} \quad \sum_{n\in\mathbb{N}_0} \beta_n = \infty$$

- one of the following properties hold:
  - $\mu f = \lim_{n\to\infty} \mu \sup_{k>n} f_k$
  - (f_n) converges α-normally to f, i.e., $\sum_{n\in\mathbb{N}_0} \alpha_n \|f - f_n\|_\infty < \infty$
  - $\alpha_n \to 1$ and (f_n) is a (faultless) sampling of an MDP

Comparison to Previous Results [CAV'25]

Convergence of the dampened Mann iteration to μf holds if

- the parameter sequences fulfill

$$\lim_{n \to \infty} {}^{\beta_n}\!/\alpha_n = 0 \quad \text{and} \quad \sum_{n \in \mathbb{N}_0} \beta_n = \infty$$

- one of the following properties hold:
  - $\mu f = \lim_{n \to \infty} \mu \sup_{k > n} f_k$
  - (f_n) converges α-normally to f, i.e., $\sum_{n \in \mathbb{N}_0} \alpha_n \|f - f_n\|_\infty < \infty$
  - $\cancel{\alpha_n \to 1 \text{ and}}$ (f_n) is a (faultless) sampling of an SSG

Comparison to Previous Results [CAV'25]

Convergence of the dampened Mann iteration to μf holds if

- the parameter sequences fulfill

$$\lim_{n\to\infty} {}^{\beta_n}\!/\alpha_n = 0 \quad \text{and} \quad \sum_{n\in\mathbb{N}_0} \beta_n = \infty$$

- one of the following properties hold:
  - $\mu f = \lim_{n\to\infty} \mu \sup_{k>n} f_k$
  - (f_n) converges α-normally to f, i.e., $\sum_{n\in\mathbb{N}_0} \alpha_n \|f - f_n\|_\infty < \infty$
  - (f_n) is a (faultless) sampling of an SSG

Comparison to Previous Results [CAV'25]

Convergence of the dampened Mann iteration to μf holds if

- the parameter sequences fulfill

$$\lim_{n\to\infty} {}^{\beta_n}\!/\alpha_n = 0 \quad \text{and} \quad \sum_{n\in\mathbb{N}_0} \beta_n = \infty$$

- one of the following properties hold:
  - $\mu f = \lim_{n\to\infty} \mu \sup_{k>n} f_k$
  - (f_n) converges α-normally to f, i.e., $\sum_{n\in\mathbb{N}_0} \alpha_n \|f - f_n\|_\infty < \infty$
  - (f_n) is a (faultless) sampling of an SSG

Still holds with component-dependent parameters and chaotic iteration (under appropriate fairness conditions).

Comparison to Previous Results [CAV'25]

Convergence of the dampened Mann iteration to μf holds if

- the parameter sequences fulfill

$$\lim_{n\to\infty} {}^{\beta_n}/_{\alpha_n} = 0 \quad \text{and} \quad \sum_{n\in\mathbb{N}_0} \beta_n = \infty$$

- one of the following properties hold:
  - $\mu f = \lim_{n\to\infty} \mu \sup_{k>n} f_k$
  - (f_n) converges α-normally to f, i.e., $\sum_{n\in\mathbb{N}_0} \alpha_n \|f - f_n\|_\infty < \infty$
  - (f_n) is a (faultless) sampling of an SSG

Still holds with component-dependent parameters and chaotic iteration (under appropriate fairness conditions).

Experiments show no significant difference between full and (random) chaotic updates.

Summary

Task: Compute least fixpoints of (monotone and non-expansive) functions $f : \mathbb{R}^d_{\geq 0} \to \mathbb{R}^d_{\geq 0}$ using only approximations $f_n \to f$.

Approach: A (chaotic) dampened Mann iteration

$$\begin{cases} x_0 \in \mathbb{R}^d_{\geq 0} \\ x_{n+1} = (1 - \beta_n)(x_n + \alpha_n(f_n(x_n) - x_n)) \end{cases}$$

Results: Convergence whenever

- the target function is sufficiently 'regular' (Picard operator)
  - $\Rightarrow$ e.g., discounted settings or with termination guarantees
- the convergence of the approximations is sufficiently 'regular'
  - $\Rightarrow$ e.g., approximations are generated through 'faultless' sampling
- the convergence of the approximations is sufficiently 'fast' with respect to the learning rates

Future Work

Arbitrary rewards / functions over $\mathbb{R}^d$

- Computing the least fixpoint
- Computing the fixpoint closest to 0

Infinite state systems / functions over infinite-dimensional domains

- Theoretical results
- Practical algorithms

Model-free learning / Césaro-convergent function sequences

$\Rightarrow$ Instead of $f_n \to f$ only $1/N \sum_{n=1}^{N} f_n \to f$

Results for chaotic iteration

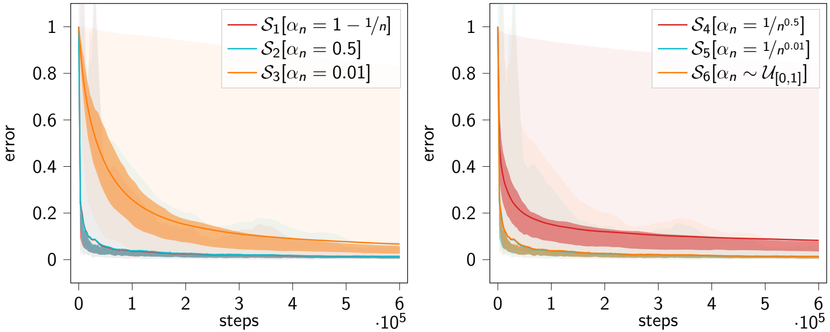

Figure: Results for the experiments of chaotic iterations on 50 randomly generated SSGs.

Results for full iteration

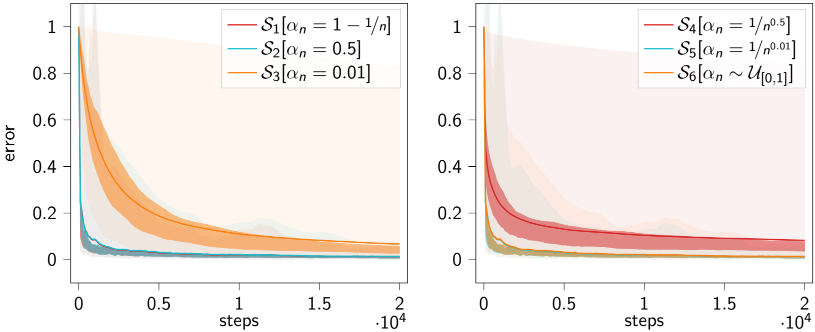

Figure: Results for the experiments of full sweep iterations on 50 randomly generated SSGs.