# Computing Fixpoints of Learned Functions

Florian Wittbold
Universität Duisburg-Essen

Joint work with Paolo Baldan, Sebastian Gurke, Barbara König

26.03.2026

# 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 $x$, i.e., the least element $x \in L$ fulfilling

$$f(x) = x.$$

# Fixpoints in Computer Science

## General Task

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

$$f(x) = x.$$

Applications in

- concurrency theory (behavioural equivalences and metrics)
- model checking ($\mu$-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 $x$, i.e., the least element $x \in L$ fulfilling

$$f(x) = x.$$

Applications in

- concurrency theory (behavioural equivalences and metrics)
- model checking ($\mu$-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 $x$, i.e., the least element $x \in L$ fulfilling

$$f(x) = x.$$

Applications in

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

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

# Case Study: Simple Stochastic Games

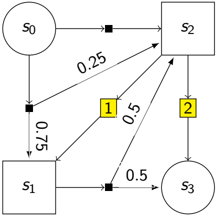

# Case Study: Simple Stochastic Games

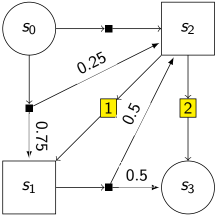


Optimal behavior can be obtained through the state-action values of the SSG, characterized as the least fixpoint of the Bellman optimality operator $f : \mathbb{R}_{\geq 0}^{S \times A} \to \mathbb{R}_{\geq 0}^{S \times A}$ given by

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

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

# Non-expansive Setting

## Functions of interest

Given a box $X = [0, C] \cap \mathbb{R}^d_{\geq 0}$ for some $C \in [0, \infty]^d$, we consider monotone functions $f : X \to X$ that are non-expansive, i.e., it holds that for all $x, y \in X$

$$\|f(x) - f(y)\|_\infty \leq \|x - y\|_\infty.$$

# Non-expansive Setting

## Functions of interest

Given a box $X = [0, C] \cap \mathbb{R}^d_{\geq 0}$ for some $C \in [0, \infty]^d$, we consider monotone functions $f : X \to X$ that are non-expansive, i.e., it holds that for all $x, y \in X$

$$\|f(x) - f(y)\|_\infty \leq \|x - y\|_\infty.$$

Here, $\| \cdot \|_\infty$ denotes the supremum norm.

# Non-expansive Setting

## Functions of interest

Given a box $X = [0, C] \cap \mathbb{R}^d_{\geq 0}$ for some $C \in [0, \infty]^d$, we consider monotone functions $f : X \to X$ that are non-expansive, i.e., it holds that for all $x, y \in X$

$$\|f(x) - f(y)\|_\infty \leq \|x - y\|_\infty.$$

Here, $\|\cdot\|_\infty$ denotes the supremum norm.

Remarks:

- This setting covers many interesting cases: termination probabilities and non-discounted returns (with non-negative rewards) in general Markov chains, MDPs, and stochastic games, behavioural metrics,...
- If $X$ is unbounded, the least fixpoint may only lie in $[0, C] \subseteq [0, \infty]^d$

# Fixpoint Theory

**Kleene and Knaster-Tarski theorems**

Given a complete lattice $(X, \sqsubseteq)$ and a (Scott-)continuous function $f : X \to X$, we have

- the set $Fix(f) \subseteq X$ of fixpoints of $f$ is a complete lattice
- $(f^n(\bot))$ is a monotonically increasing sequence and $\mu f = \bigvee_{n \in \mathbb{N}_0} f^n(\bot) = \bigwedge\{x \in X \mid x \geq f(x)\}$

# Fixpoint Theory

## Kleene and Knaster-Tarski theorems

Given a *box* $X = [0, C] \subseteq [0, \infty]^d$ for some $C \in [0, \infty]^d$, and a monotone and continuous function $f : X \to X$, we have

- the set $Fix(f) \subseteq X$ of fixpoints of $f$ is a complete lattice
- the iteration

$$\begin{cases} x_0 = 0 \\ x_{n+1} = f(x_n) \end{cases}$$

converges to $\mu f = \inf\{x \in X \mid x \geq f(x)\}$

# Case Study: Simple Stochastic Games

New Task: Compute $\mu f$ 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 $\mu f$ 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 $\mu f$ 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

# Discount Factors and Termination

# Discount Factors and Termination

Formal results in reinforcement learning and stochastic approximation typically concentrate on the case where we seek

# Discount Factors and Termination

Formal results in reinforcement learning and stochastic approximation typically concentrate on the case where we seek

- the total discounted value for a discount factor $\gamma < 1$, i.e., the (unique!) fixpoint of the discounted Bellman operator

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

# Discount Factors and Termination

Formal results in reinforcement learning and stochastic approximation typically concentrate on the case where we seek

- the total discounted value for a discount factor $\gamma < 1$, i.e., the (unique!) fixpoint of the discounted Bellman operator

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

$\Rightarrow$ $f$ is a contraction!

# Discount Factors and Termination

Formal results in reinforcement learning and stochastic approximation typically concentrate on the case where we seek

- the total discounted value for a discount factor $\gamma < 1$, i.e., the (unique!) fixpoint of the discounted Bellman operator

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

$\Rightarrow$ $f$ is a contraction!

- the value (e.g., optimal termination probability) under some termination guarantees

# Discount Factors and Termination

Formal results in reinforcement learning and stochastic approximation typically concentrate on the case where we seek

- the total discounted value for a discount factor $\gamma < 1$, i.e., the (unique!) fixpoint of the discounted Bellman operator

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

  $\Rightarrow$ $f$ is a contraction!

- the value (e.g., optimal termination probability) under some termination guarantees

  $\Rightarrow$ $f$ is a power-contraction!

# Non-expansive Setting

**Task**

Given a sequence of monotone and non-expansive functions $f_0, f_1, \cdots : X \to X$ that converges (pointwise) to $f : X \to X$:

Compute a sequence $x_0, x_1, x_2, \cdots \in X$ that converges to $\mu f$ where $x_n$ may only depend on $f_0, \ldots, f_{n-1}$.

# Non-expansive Setting

**Task**

Given a sequence of monotone and non-expansive functions $f_0, f_1, \cdots : X \to X$ that converges (pointwise) to $f : X \to X$:

Compute a sequence $x_0, x_1, x_2, \cdots \in X$ that converges to $\mu f$ where $x_n$ may only depend on $f_0, \ldots, f_{n-1}$.

Will a standard Kleene iteration

$$\begin{cases} x_0 = 0 \\ x_{n+1} = f_n(x_n) \end{cases}$$

still work? Or do we need new techniques?

# The Problem with Over-Approximations

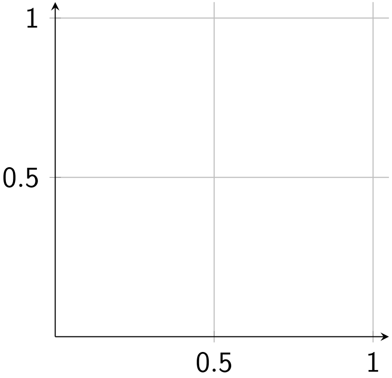


$$f(x) = \max(x, 0.5)$$

# The Problem with Over-Approximations

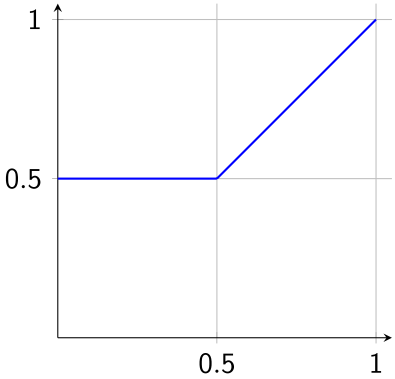


$$f(x) = \max(x, 0.5)$$

# The Problem with Over-Approximations

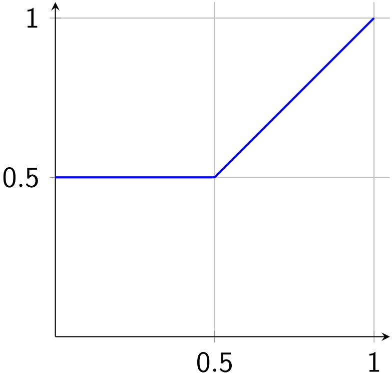


$$f(x) = \max(x, 0.5)$$

# The Problem with Over-Approximations

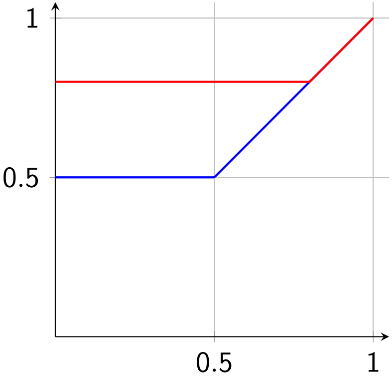


$$f(x) = \max(x, 0.8)$$

# The Problem with Over-Approximations

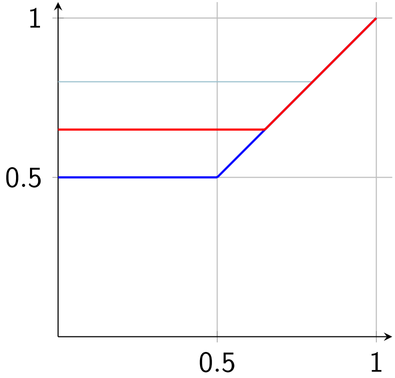


$$f(x) = \max(x, 0.65)$$

# The Problem with Over-Approximations

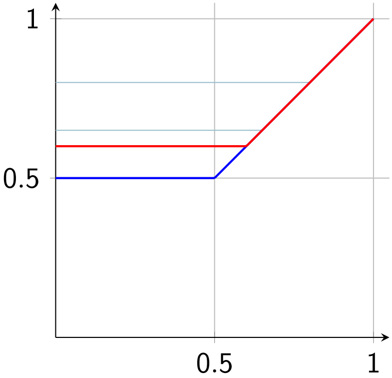


$$f(x) = \max(x, 0.6)$$

# The Problem with Over-Approximations

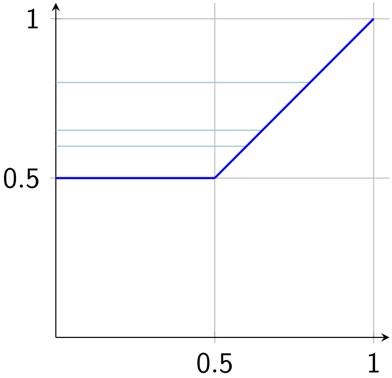


$$f_n(x) = \max(x, 0.5 + 0.3/n)$$

# The Problem with Over-Approximations

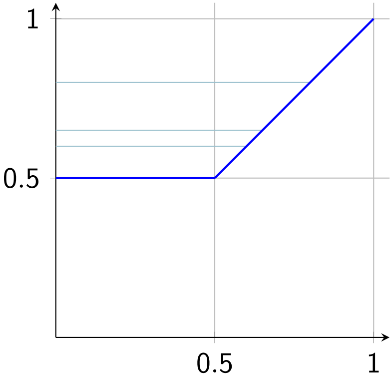


$$f_n(x) = \max(x, 0.5 + 0.3/n)$$

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

# Introducing Dampening to the Iteration

(Kleene) Iteration

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

# Introducing Dampening to the Iteration

Dampened (Kleene) Iteration

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

# Introducing Dampening to the Iteration

Dampened (Kleene) Iteration

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

The parameters $\beta_n \in [0, 1]$ are *dampening factors* that fulfill

$$\lim_{n \to \infty} \beta_n = 0 \quad \text{and} \quad \sum_{n=0}^{\infty} \beta_n = \infty$$

Standard Choice: $\beta_n = 1/n$

# Introducing Dampening to the Iteration

Dampened (Kleene) Iteration

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

The parameters $\beta_n \in [0, 1]$ are *dampening factors* that fulfill

$$\lim_{n\to\infty} \beta_n = 0 \quad \text{and} \quad \sum_{n=0}^{\infty} \beta_n = \infty$$

Standard Choice: $\beta_n = 1/n$

Intuition: a form of 'vanishing discount', but there is always enough 'power' left to decrease a current over-estimation to the true least fixpoint.

# Introducing a Learning Rate

Dampened (Kleene) Iteration

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

# Introducing a Learning Rate

Dampened Mann Iteration

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

# Introducing a Learning Rate

Dampened Mann Iteration

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

The parameters $\alpha_n \in [0, 1]$ are *learning rates* that fulfill

$$\lim_{n \to \infty} {}^{\beta_n}\!/\!_{\alpha_n} = 0$$

Standard Choices: $\alpha_n = \alpha > 0$ or $\alpha_n = \beta_n^\varepsilon, \varepsilon \in (0, 1)$

# Introducing a Learning Rate

Dampened Mann Iteration

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

The parameters $\alpha_n \in [0,1]$ are *learning rates* that fulfill

$$\lim_{n\to\infty} {}^{\beta_n}/_{\alpha_n} = 0$$

Standard Choices: $\alpha_n = \alpha > 0$ or $\alpha_n = \beta_n^{\varepsilon}, \varepsilon \in (0,1)$

Intuition: do not completely trust the current approximation, but keep the update large enough to be able to counteract the dampening.

# Convergence Results

## Convergence of Dampened Mann Iteration

Let $f_n : X \to X$ be monotone and non-expansive functions converging (pointwise) to $f$, and $(\alpha_n), (\beta_n)$ a sequence fulfilling

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

Then, the dampened Mann iteration converges to $\mu f$ for any initial point $x_0 \in X$ if

# Convergence Results

## Convergence of Dampened Mann Iteration

Let $f_n : X \to X$ be monotone and non-expansive functions converging (pointwise) to $f$, and $(\alpha_n), (\beta_n)$ a sequence fulfilling

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

Then, the dampened Mann iteration converges to $\mu f$ for any initial point $x_0 \in X$ if

- $f$ is a Picard operator

# Convergence Results

## Convergence of Dampened Mann Iteration

Let $f_n : X \to X$ be monotone and non-expansive functions converging (pointwise) to $f$, and $(\alpha_n), (\beta_n)$ a sequence fulfilling

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

Then, the dampened Mann iteration converges to $\mu f$ for any initial point $x_0 \in X$ if

- $f$ is a Picard operator
- $\mu f = \lim_{n\to\infty} \mu(\sup_{k>n} f_k)$

# Convergence Results

## Convergence of Dampened Mann Iteration

Let $f_n : X \to X$ be monotone and non-expansive functions converging (pointwise) to $f$, and $(\alpha_n), (\beta_n)$ a sequence fulfilling

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

Then, the dampened Mann iteration converges to $\mu f$ for any initial point $x_0 \in X$ if

- $f$ is a Picard operator
- $\mu f = \lim_{n\to\infty} \mu(\sup_{k>n} f_k)$
- $(f_n)$ converges $\alpha$-normally, i.e., $\sum_{n\in\mathbb{N}_0} \alpha_n \|f - f_n\|_\infty < \infty$

# Convergence Results

## Convergence of Dampened Mann Iteration

Let $f_n : X \to X$ be monotone and non-expansive functions converging (pointwise) to $f$, and $(\alpha_n), (\beta_n)$ a sequence fulfilling

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

Then, the dampened Mann iteration converges to $\mu f$ for any initial point $x_0 \in X$ if

- $f$ is a Picard operator
- $\mu f = \lim_{n\to\infty} \mu(\sup_{k>n} f_k)$
- $(f_n)$ converges $\alpha$-normally, i.e., $\sum_{n\in\mathbb{N}_0} \alpha_n \|f - f_n\|_\infty < \infty$

# Embracing the Chaos

Currently: Need for *full sweeps*!

$\Rightarrow$ Infeasible for large systems

# Embracing the Chaos

Currently: Need for *full sweeps*!

$\Rightarrow$ Infeasible for large systems

What changes if we allow for chaotic updates?

# Embracing the Chaos

Currently: Need for *full sweeps*!

$\Rightarrow$ Infeasible for large systems

> What changes if we allow for chaotic updates?

Dampened Mann Iteration

$$\begin{cases} x_0 \in X \\ x_{n+1} = (1-\beta_n)(x_n + \alpha_n(f_n(x_n) - x_n)) \end{cases}$$

# Embracing the Chaos

Currently: Need for *full sweeps*!

$\Rightarrow$ Infeasible for large systems

What changes if we allow for chaotic updates?

Chaotic Dampened Mann Iteration

$$\begin{cases} x_0 \in X \\ x_{n+1}(i) = (1-\beta_n)(x_n(i) + \alpha_n(f_n(x_n)(i) - x_n(i))) \quad \text{for all } i \in S_n \\ x_{n+1}(i) = x_n(i) \quad \text{for all } i \notin S_n \end{cases}$$

# Main Results

### Convergence of Chaotic Dampened Mann Iteration

Let $f_n : X \to X$ be monotone and non-expansive functions converging (pointwise) to $f$, and $(\alpha_n), (\beta_n), (S_n)$ sequences fulfilling

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

along with $(\beta_n)$ being non-increasing and a fairness condition on $(S_n)$. Then, the chaotic dampened Mann iteration converges to $\mu f$ for any $x_0 \in X$ if

# Main Results

## Convergence of Chaotic Dampened Mann Iteration

Let $f_n : X \to X$ be monotone and non-expansive functions converging (pointwise) to $f$, and $(\alpha_n), (\beta_n), (S_n)$ sequences fulfilling

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

along with $(\beta_n)$ being non-increasing and a fairness condition on $(S_n)$. Then, the chaotic dampened Mann iteration converges to $\mu f$ for any $x_0 \in X$ if

- $f$ is a Picard operator

# Main Results

**Convergence of Chaotic Dampened Mann Iteration**

Let $f_n : X \to X$ be monotone and non-expansive functions converging (pointwise) to $f$, and $(\alpha_n), (\beta_n), (S_n)$ sequences fulfilling

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

along with $(\beta_n)$ being non-increasing and a fairness condition on $(S_n)$. Then, the chaotic dampened Mann iteration converges to $\mu f$ for any $x_0 \in X$ if

- $f$ is a Picard operator
- $\mu f = \lim_{n\to\infty} \mu(\sup_{k>n} f_k)$

# Application to Sampled Games

## Faultless Sampling

An approximation $\mathcal{G}_n$ of $\mathcal{G}$ is called a faultless sampling iff for all $s, s' \in S, a \in A(s), n \in \mathbb{N}_0$

# Application to Sampled Games

## Faultless Sampling

An approximation $\mathcal{G}_n$ of $\mathcal{G}$ is called a faultless sampling iff for all $s, s' \in S, a \in A(s), n \in \mathbb{N}_0$

- $T(s' \mid s, a) = 0$ implies $T_n(s' \mid s, a) = 0$

# Application to Sampled Games

## Faultless Sampling

An approximation $\mathcal{G}_n$ of $\mathcal{G}$ is called a faultless sampling iff for all $s, s' \in S, a \in A(s), n \in \mathbb{N}_0$

- $T(s' \mid s, a) = 0$ implies $T_n(s' \mid s, a) = 0$
- $R(s, a) = 0$ implies $R_n(s, a) = 0$

# Application to Sampled Games

## Faultless Sampling

An approximation $\mathcal{G}_n$ of $\mathcal{G}$ is called a faultless sampling iff for all $s, s' \in S, a \in A(s), n \in \mathbb{N}_0$

- $T(s' \mid s, a) = 0$ implies $T_n(s' \mid s, a) = 0$
- $R(s, a) = 0$ implies $R_n(s, a) = 0$

## Theorem

*Let $\mathcal{G}_n$ be a faultless sampling of the SSG $\mathcal{G}$.*
*Then, we have*

# Application to Sampled Games

### Faultless Sampling

An approximation $\mathcal{G}_n$ of $\mathcal{G}$ is called a faultless sampling iff for all $s, s' \in S, a \in A(s), n \in \mathbb{N}_0$

- $T(s' \mid s, a) = 0$ implies $T_n(s' \mid s, a) = 0$
- $R(s, a) = 0$ implies $R_n(s, a) = 0$

### Theorem

*Let $\mathcal{G}_n$ be a faultless sampling of the SSG $\mathcal{G}$.*
*Then, we have*

- $\mu f_{\mathcal{G}} = \lim_{n\to\infty} \mu(\sup_{k>n} f_{\mathcal{G}_n})$

# Application to Sampled Games

## Faultless Sampling

An approximation $\mathcal{G}_n$ of $\mathcal{G}$ is called a faultless sampling iff for all $s, s' \in S, a \in A(s), n \in \mathbb{N}_0$

- $T(s' \mid s, a) = 0$ implies $T_n(s' \mid s, a) = 0$
- $R(s, a) = 0$ implies $R_n(s, a) = 0$

## Theorem

*Let $\mathcal{G}_n$ be a faultless sampling of the SSG $\mathcal{G}$.*
*Then, we have*

- $\mu f_{\mathcal{G}} = \lim_{n\to\infty} \mu(\sup_{k>n} f_{\mathcal{G}_n})$
- *the chaotic dampened Mann iteration converges to $\mu f_{\mathcal{G}}$.*

# Current 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 ${}^1/_N \sum_{n=1}^{N} f_n \to f$

# 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 $f_n \to f$.

# 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

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

# 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

# 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)

# 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

# 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' ($\mu f = \lim \mu(\sup_{k>n} f_k)$)

# 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' ($\mu f = \lim \mu(\sup_{k>n} f_k)$)
  - $\Rightarrow$ e.g., approximations are generated through 'faultless' sampling

# 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' ($\mu f = \lim \mu(\sup_{k>n} f_k)$)
  - $\Rightarrow$ e.g., approximations are generated through 'faultless' sampling
- the convergence of the approximations is sufficiently 'fast' with respect to the learning rates ($\sum \alpha_n \|f - f_n\| < \infty$)