# Approximating Fixpoints of Approximated Functions

Florian Wittbold
Universität Duisburg-Essen

Joint work with Paolo Baldan, Sebastian Gurke, Barbara König, Tommaso Padoan

# 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 $f\colon L \to L$ 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)

# 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)
- Markov decision processes and games (computation of value functions)
- ...

# How to Compute Fixpoints?

### Kleene Iteration

Given a *monotone* and $\omega$-*continuous* function $f : L \to L$ over a *complete lattice* $(L, \sqsubseteq)$, the iteration given by

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

converges to $\mu f$ for any initial point $x_0 \leq \mu f$.

# How to Compute Fixpoints?

**Kleene Iteration**

Given a *monotone* and $\omega$-*continuous* function $f : L \to L$ over a *complete lattice* $(L, \sqsubseteq)$, the iteration given by

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

converges to $\mu f$ for any initial point $x_0 \leq \mu f$.

Corresponds to value iteration for Markov decision processes:

$$v_{n+1}(s) = f(v_n)(s) = \max_{a \in A(s)} R(s, a) + \sum_{s' \in S} T(s' \mid s, a) v_n(s')$$

where $f$ is the Bellman operator.

# Introducing Approximations

But: What if the target function $f$ can only be approximated?

# Introducing Approximations

But: What if the target function $f$ can only be approximated?

Sources of approximation:

- Unknown / partially known system
- Evolving system
- Approximated partial results
- ...

# Introducing Approximations

But: What if the target function $f$ can only be approximated?

Sources of approximation:

- Unknown / partially known system
- Evolving system
- Approximated partial results
- ...

Will Kleene iteration ($x_{n+1} = f_n(x_n)$) still work? For which types of functions? Do we need new techniques?

# Introducing Approximations in MDPs

For MDPs: How can $\mu f$ be determined if the transition probabilities (given by $T$) and the rewards (given by $R$) are not known precisely but can only be approximated, e.g., by sampling through interaction with the system?

# Introducing Approximations in MDPs

For MDPs: How can $\mu f$ be determined if the transition probabilities (given by $T$) and the rewards (given by $R$) are not known precisely but can only be approximated, e.g., by sampling through interaction with the system?

$\Longrightarrow$ (Model-based) Reinforcement Learning

# Introducing Approximations in MDPs

For MDPs: How can $\mu f$ be determined if the transition probabilities (given by $T$) and the rewards (given by $R$) are not known precisely but can only be approximated, e.g., by sampling through interaction with the system?

$\Longrightarrow$ (Model-based) Reinforcement Learning

Will a standard value iteration (with approximate $T_n$ and $R_n$) still work? For which types of MDPs? Do we need new techniques?

# Contractive vs. Non-expansive Functions

## Contractive and non-expansive functions

A function $f$ over a metric space $(X, d)$ is called

# Contractive vs. Non-expansive Functions

## Contractive and non-expansive functions

A function $f$ over a metric space $(X, d)$ is called

- a contraction if there exists $0 \leq q < 1$ such that

$$d(f(x), f(y)) \leq q \cdot d(x, y) \quad \text{for all } x, y \in X,$$

# Contractive vs. Non-expansive Functions

## Contractive and non-expansive functions

A function $f$ over a metric space $(X, d)$ is called

- a contraction if there exists $0 \leq q < 1$ such that

$$d(f(x), f(y)) \leq q \cdot d(x, y) \quad \text{for all } x, y \in X,$$

- a power contraction if there exists $n \in \mathbb{N}$ such that $f^n$ is a contraction.

# Contractive vs. Non-expansive Functions

## Contractive and non-expansive functions

A function $f$ over a metric space $(X, d)$ is called

- a contraction if there exists $0 \leq q < 1$ such that

$$d(f(x), f(y)) \leq q \cdot d(x, y) \quad \text{for all } x, y \in X,$$

- a power contraction if there exists $n \in \mathbb{N}$ such that $f^n$ is a contraction.
- non-expansive if the above holds for $q = 1$,

# 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

$$f(v)(s) = \max_{a \in A(s)} R(s,a) + \gamma \sum_{s' \in S} T(s' \mid s, a) v(s')$$

# 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

$$f(v)(s) = \max_{a \in A(s)} R(s, a) + \gamma \sum_{s' \in S} T(s' \mid s, a) v(s')$$

$\Rightarrow$ $f$ 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

$$f(v)(s) = \max_{a \in A(s)} R(s,a) + \gamma \sum_{s' \in S} T(s' \mid s, a) v(s')$$

$\Rightarrow$ $f$ is a contraction!

- the (optimal) termination probability under some termination guarantees, i.e., the (unique!) fixpoint of the Bellman operator

$$f(v)(s) = \begin{cases} 1 & \text{if } s \in F \\ \max_{a \in A(s)} \sum_{s' \in S} T(s' \mid s, a) v(s') & \text{otherwise.} \end{cases}$$

# 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

$$f(v)(s) = \max_{a \in A(s)} R(s,a) + \gamma \sum_{s' \in S} T(s' \mid s, a) v(s')$$

$\Rightarrow$ $f$ is a contraction!

- the (optimal) termination probability under some termination guarantees, i.e., the (unique!) fixpoint of the Bellman operator

$$f(v)(s) = \begin{cases} 1 & \text{if } s \in F \\ \max_{a \in A(s)} \sum_{s' \in S} T(s' \mid s, a) v(s') & \text{otherwise.} \end{cases}$$

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

# Results for (Power) Contractions

**Generalized Banach Fixpoint Theorem**

Given a sequence $(f_n)_n$ of contractions $f_n$ over a complete metric space $(X, d)$ converging (pointwise) to a contraction $f$, then $\mu f_n \to \mu f$ and the sequence

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

converges to $\mu f$ for any initial guess $x_0 \in X$.
This result even holds for the more general class of (non-expansive) power contractions.

# Non-contractive setting

## Main Goal

Given a sequence $(f_n)_n$ of *monotone* and *non-expansive* (w.r.t. the supremum distance) functions $f_n : \mathbb{R}_+^d \to \mathbb{R}_+^d$ converging (pointwise) to a function $f$.
Compute a sequence $x_0, x_1, \ldots$ that converges to $\mu f$, only using the approximations $f_n$ of $f$.

# Non-contractive setting

## Main Goal

Given a sequence $(f_n)_n$ of *monotone* and *non-expansive* (w.r.t. the supremum distance) functions $f_n : \mathbb{R}_+^d \to \mathbb{R}_+^d$ converging (pointwise) to a function $f$.
Compute a sequence $x_0, x_1, \ldots$ that converges to $\mu f$, only using the approximations $f_n$ of $f$.

Remarks:

- We also want convergence if $\mu f$ only lies in $(\mathbb{R}_+ \cup \{\infty\})^d$

# Non-contractive setting

**Main Goal**

Given a sequence $(f_n)_n$ of *monotone* and *non-expansive* (w.r.t. the supremum distance) functions $f_n : \mathbb{R}_+^d \to \mathbb{R}_+^d$ converging (pointwise) to a function $f$.
Compute a sequence $x_0, x_1, \ldots$ that converges to $\mu f$, only using the approximations $f_n$ of $f$.

Remarks:

- We also want convergence if $\mu f$ only lies in $(\mathbb{R}_+ \cup \{\infty\})^d$
- Non-expansiveness w.r.t. the supremum distance covers many interesting cases: termination probabilities and non-discounted returns (with non-negative rewards) in general Markov chains, MDPs, and stochastic games, behavioural metrics, quantitative $\mu$-calculi,...

# The problem with over-approximations

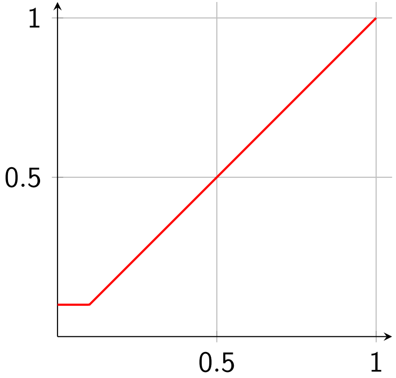

# The problem with over-approximations

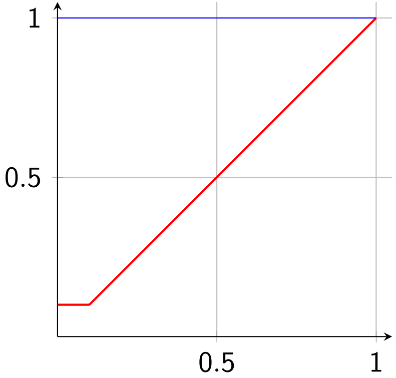

# The problem with over-approximations

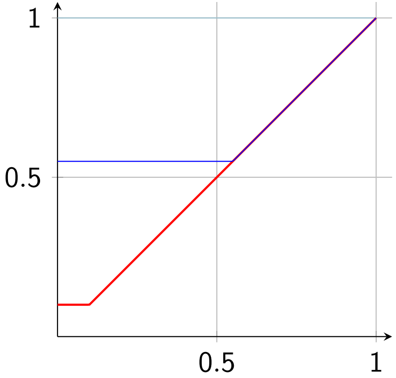

# The problem with over-approximations

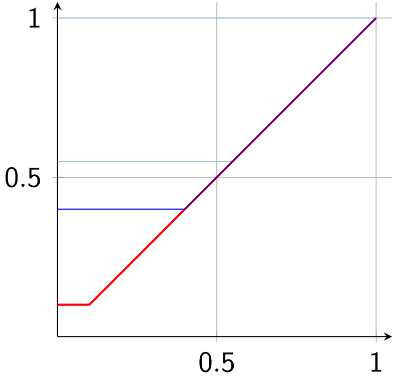

# The problem with over-approximations

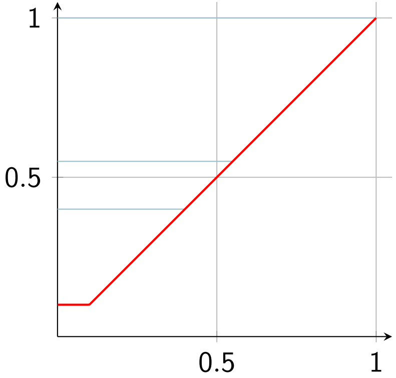


$$f_n(x) = \max(x, 0.1 + 0.9/n)$$

# The problem with over-approximations

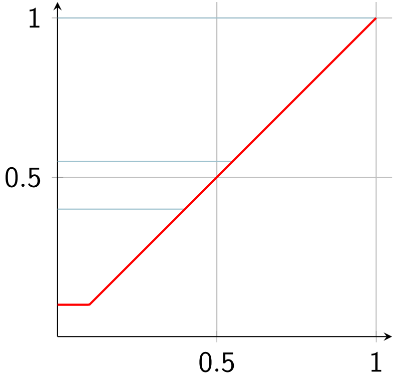


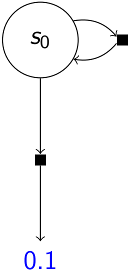


$$f_n(x) = \max(x, 0.1 + 0.9/n)$$

# The problem with over-approximations

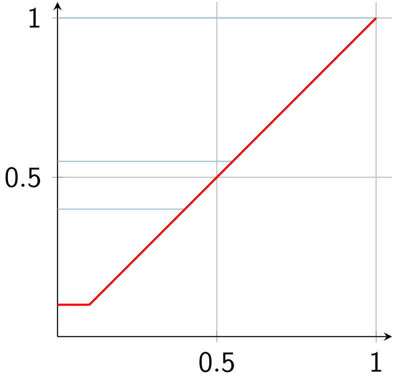


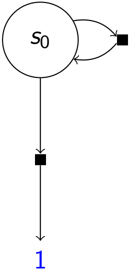


$$f_n(x) = \max(x, 0.1 + 0.9/n)$$

# The problem with over-approximations

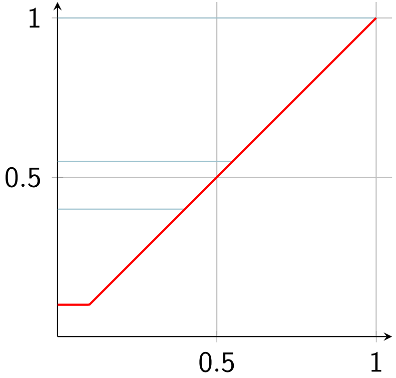


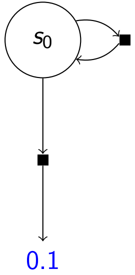


$$f_n(x) = \max(x, 0.1 + 0.9/n)$$

We have $\mu f_n = 0.1 + 0.9/n \to 0.1 = \mu f$, but iterating with any $f_n$ will over-estimate $\mu f$ and further iterations will never decrease it.

# Resetting Schemes

Idea: Approximate $f$ 'far enough' and iterate with approximation $f_n$ from bottom.

# Resetting Schemes

Idea: Approximate $f$ 'far enough' and iterate with approximation $f_n$ from bottom.

$$x_n = f_n^{k_n}(0)$$

where $k_n \to \infty$.

# Resetting Schemes

Idea: Approximate $f$ 'far enough' and iterate with approximation $f_n$ from bottom.

$$x_n = f_n^{k_n}(0)$$

where $k_n \to \infty$.

In Reinforcement Learning: Certainty Equivalent Method (CEM)

# Resetting Schemes

Idea: Approximate $f$ 'far enough' and iterate with approximation $f_n$ from bottom.

$$x_n = f_n^{k_n}(0)$$

where $k_n \to \infty$.

In Reinforcement Learning: Certainty Equivalent Method (CEM)

But:

- No *online* algorithm

# Resetting Schemes

Idea: Approximate $f$ 'far enough' and iterate with approximation $f_n$ from bottom.

$$x_n = f_n^{k_n}(0)$$

where $k_n \to \infty$.

In Reinforcement Learning: Certainty Equivalent Method (CEM)

But:

- No *online* algorithm
- When is it 'far enough'?

# Resetting Schemes

Idea: Approximate $f$ ‘far enough’ and iterate with approximation $f_n$ from bottom.

$$x_n = f_n^{k_n}(0)$$

where $k_n \to \infty$.

In Reinforcement Learning: Certainty Equivalent Method (CEM)

But:

- No *online* algorithm
- When is it ‘far enough’?
- Uses the assumption that $\mu f_n \to \mu f$

# Non-continuity of $\mu$

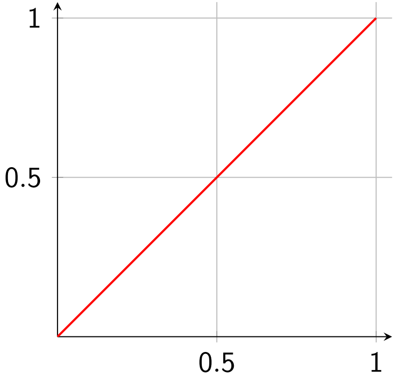

# Non-continuity of $\mu$

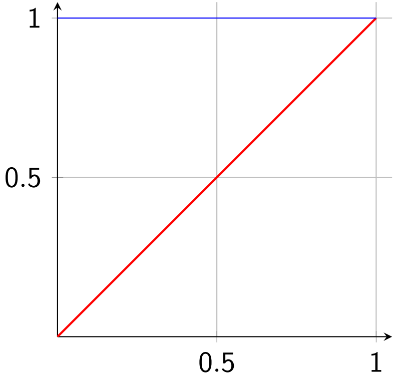

# Non-continuity of $\mu$

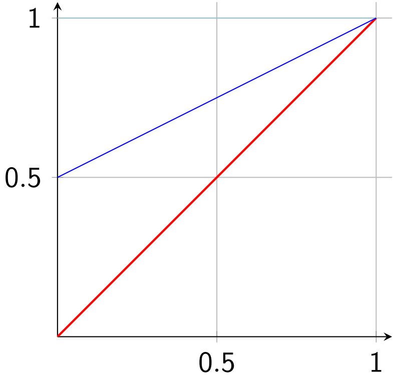

# Non-continuity of $\mu$

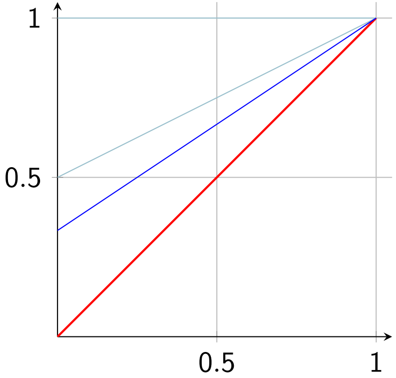

# Non-continuity of $\mu$

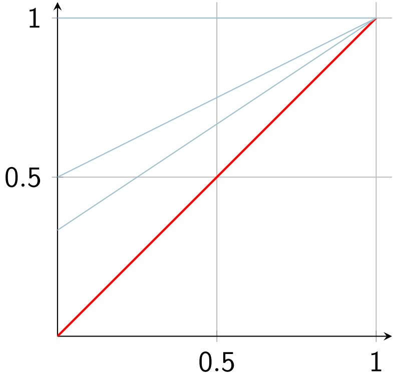


$$f_n(x) = (1 - 1/n)x + 1/n$$

# Non-continuity of $\mu$

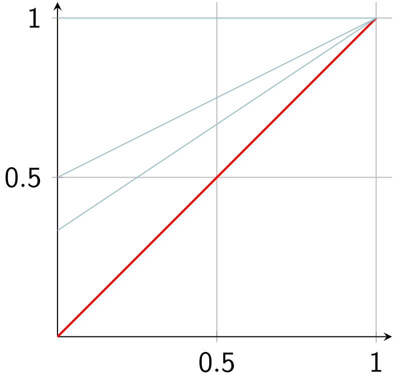


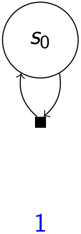


$$f_n(x) = (1 - {}^1\!/_n)x + {}^1\!/_n$$

# Non-continuity of $\mu$

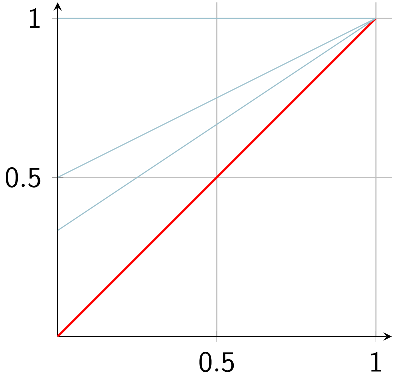


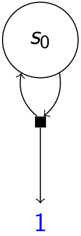


$$f_n(x) = (1 - {}^1\!/\!_n)x + {}^1\!/\!_n$$

# Non-continuity of $\mu$

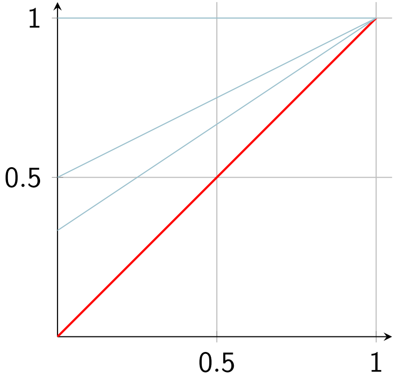


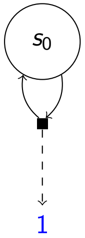


$$f_n(x) = (1 - 1/n)x + 1/n$$

# Non-continuity of $\mu$

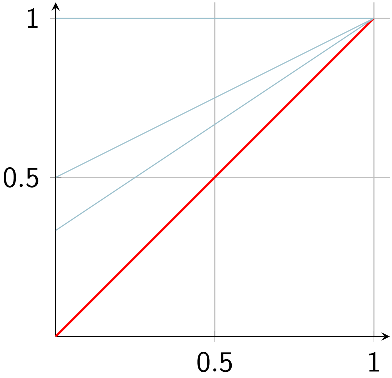


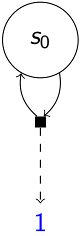


$$f_n(x) = (1 - {}^1\!/\!_n)x + {}^1\!/\!_n$$

We have $\mu f_n = 1 \nrightarrow 0 = \mu f$.

# Least fixpoints of MDPs

## Faultless Sampling

A sequence $M_n = (S, A, T_n, R_n)$ of MDPs approximating an MDP $M = (S, A, T, R)$ is called being *generated through faultless sampling* iff for all $s, s' \in S, a \in A(s)$

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

# Least fixpoints of MDPs

## Faultless Sampling

A sequence $M_n = (S, A, T_n, R_n)$ of MDPs approximating an MDP $M = (S, A, T, R)$ is called being *generated through faultless sampling* iff for all $s, s' \in S, a \in A(s)$

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

## Continuity of $\mu$ under faultless sampling

Given a sequence $(f_n)_n$ of Bellman operators of MDPs generated through faultless sampling of an MDP with Bellman operator $f$. Then, $\mu f_n \to \mu f$.

# 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.

# Main Results

**Convergence of Dampened Iteration**

The dampened (Kleene) iteration converges to $\mu f$ if

# Main Results

## Convergence of Dampened Iteration

The dampened (Kleene) iteration converges to $\mu f$ if

- $\mu f_n \to \mu f$ and $d = 1$

# Main Results

## Convergence of Dampened Iteration

The dampened (Kleene) iteration converges to $\mu f$ if

- $\mu f_n \to \mu f$ and $d = 1$
- $\mu(\sup_{k>n} f_k) \to \mu f$

# Main Results

**Convergence of Dampened Iteration**

The dampened (Kleene) iteration converges to $\mu f$ if

- $\mu f_n \to \mu f$ and $d = 1$
- $\mu(\sup_{k>n} f_k) \to \mu f$
- $\sum_{n\in\mathbb{N}} \|f_n(x_n) - f(x_n)\| < \infty$

# Diverging Iterations: Example I

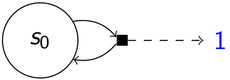


Given the sequence $(f_n)_n$ of functions $f_n : [0, 1] \to [0, 1]$ defined by

$$f_n(x) = (1 - 1/n)x + 1/n.$$

# Diverging Iterations: Example I

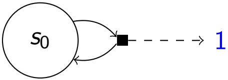


Given the sequence $(f_n)_n$ of functions $f_n : [0,1] \to [0,1]$ defined by

$$f_n(x) = (1 - {}^1\!/_n)x + {}^1\!/_n.$$

The iteration

$$x_{n+1} = (1 - {}^1\!/_n) f_n(x_n)$$

converges to ${}^1\!/_2 \neq 0 = \mu f$.

# Diverging Iterations: Example II

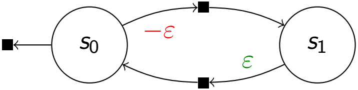

# Diverging Iterations: Example II

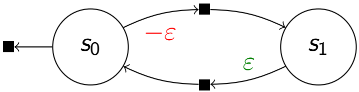


Given the sequence $(f_n)_n$ of functions $f_n : \mathbb{R}_+^2 \to \mathbb{R}_+^2$ defined by

$$f_n(x, y) = \begin{cases} (\max(0, y), x) & \text{if } n \text{ even} \\ (\max(0, y - {}^2\!/\!_n), x + {}^2\!/\!_n) & \text{if } n \text{ odd} \end{cases}$$

# Diverging Iterations: Example II

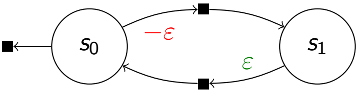


Given the sequence $(f_n)_n$ of functions $f_n : \mathbb{R}^2_+ \to \mathbb{R}^2_+$ defined by

$$f_n(x, y) = \begin{cases} (\max(0, y), x) & \text{if } n \text{ even} \\ (\max(0, y - {}^2/_n), x + {}^2/_n) & \text{if } n \text{ odd} \end{cases}$$

The iteration

$$x_{n+1} = (1 - {}^1/_{n+1}) f_n(x_n)$$

with $x_1 = (0, 1)$ diverges with cluster points $(1, 0)$ and $(0, 1)$, whereas $\mu f = (0, 0)$.

# Convergence under faultless sampling

**Convergence of Dampened Iteration (for MDPs)**

Given a sequence $(f_n)_n$ of Bellman operators of MDPs generated by faultless sampling of an MDP with Bellman operator $f$.

Then, $\mu(\sup_{k \geq n} f_k) \to \mu f$.

In particular, the dampened (Kleene) iteration converges to $\mu f$ (for any initial guess).

The same result even holds for (simple) stochastic games.

# Introducing Learning Rates

Dampened (Kleene) Iteration

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

# Introducing Learning Rates

Dampened Mann Iteration

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

# Introducing Learning Rates

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 Learning Rates

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.

# Main Results

## Convergence of Dampened Mann Iteration

The dampened Mann iteration converges to $\mu f$ if

# Main Results

## Convergence of Dampened Mann Iteration

The dampened Mann iteration converges to $\mu f$ if

- $\mu(\sup_{k>n} f_k) \to \mu f$

# Main Results

## Convergence of Dampened Mann Iteration

The dampened Mann iteration converges to $\mu f$ if

- $\mu(\sup_{k>n} f_k) \to \mu f$
- $\sum_{n\in\mathbb{N}} \alpha_n \|f_n(x_n) - f(x_n)\| < \infty$

# Main Results

## Convergence of Dampened Mann Iteration

The dampened Mann iteration converges to $\mu f$ if

- $\mu(\sup_{k>n} f_k) \to \mu f$
- $\sum_{n\in\mathbb{N}} \alpha_n \| f_n(x_n) - f(x_n)\| < \infty$

## Existence of Converging Parameters

Let $(f_n)_n$ be a sequence of monotone and non-expansive functions converging to a function $f$.
Then, there exist parameter sequences $(\alpha_n)_n$ and $(\beta_n)_n$ such that the dampened Mann iteration converges to $\mu f$.

# Main Results

## Convergence of Dampened Mann Iteration

The dampened Mann iteration converges to $\mu f$ if

- $\mu(\sup_{k>n} f_k) \to \mu f$
- $\sum_{n\in\mathbb{N}} \alpha_n \|f_n(x_n) - f(x_n)\| < \infty$

## Existence of Converging Parameters

Let $(f_n)_n$ be a sequence of monotone and non-expansive functions converging to a function $f$.
Then, there exist parameter sequences $(\alpha_n)_n$ and $(\beta_n)_n$ such that the dampened Mann iteration converges to $\mu f$.

If we can compute upper bounds $\varepsilon_n$ to $\|f_n(x_n) - f(x_n)\|$ with $\varepsilon_n \to 0$, such parameters can also be computed on the fly.

# Ongoing Work

# Ongoing Work

- Chaotic iteration

$$x_{n+1}(i) = \begin{cases} T_n(x_n)(i) & \text{if } i \in S_n \\ x_n(i) & \text{otherwise.} \end{cases}$$

# Ongoing Work

- Chaotic iteration
- Infinite-state systems

$$f : \mathbb{R}_+^{\mathbb{N}} \to \mathbb{R}_+^{\mathbb{N}}$$

# Ongoing Work

- Chaotic iteration
- Infinite-state systems
- Negative rewards

$$f : \mathbb{R}^d \to \mathbb{R}^d$$

# Ongoing Work

- Chaotic iteration
- Infinite-state systems
- Negative rewards

$$f : \mathbb{R}^d \to \mathbb{R}^d, \text{ but } \mu f_M \neq \mathbb{V}_M$$

# Ongoing Work

- Chaotic iteration
- Infinite-state systems
- Negative rewards
- Stopping criteria

$$\text{For } \varepsilon > 0, \text{ find } n \in \mathbb{N} \text{ with } \|x_n - \mu f\| < \varepsilon$$

# Ongoing Work

- Chaotic iteration
- Infinite-state systems
- Negative rewards
- Stopping criteria
- Model-free learning

$$f_n(q)(s_n, a_n) = r_n + \max_{a' \in A(s'_n)} q(s'_n, a')$$

# Ongoing Work

- Chaotic iteration
- Infinite-state systems
- Negative rewards
- Stopping criteria
- Model-free learning

$$f_n(q)(s_n, a_n) = r_n + \max_{a' \in A(s'_n)} q(s'_n, a')$$

$$\Rightarrow \quad f_n \nrightarrow f$$

# Ongoing Work

- Chaotic iteration
- Infinite-state systems
- Negative rewards
- Stopping criteria
- Model-free learning

$$f_n(q)(s_n, a_n) = r_n + \max_{a' \in A(s'_n)} q(s'_n, a')$$

$$\Rightarrow \quad f_n \nrightarrow f, \text{ but } {}^1\!/\!_n \sum_{k=1}^{n} f_n \to f$$

# Ongoing Work

- Chaotic iteration
- Infinite-state systems
- Negative rewards
- Stopping criteria
- Model-free learning