Approximative Fixpoint Theory

and Applications to Reinforcement Learning

Florian Wittbold

University of Duisburg-Essen

July 25, 2026

Fixpoints

For an endofunction $F : X \to X$ on a set X, a point $x \in X$ is a fixpoint iff

$$F(x) = x.$$

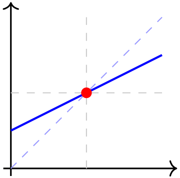

Fixpoints

For an endofunction $F : X \to X$ on a set X, a point $x \in X$ is a fixpoint iff

$$F(x) = x.$$

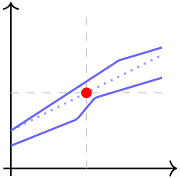

What if F is unknown?

Fixpoints

For an endofunction $F : X \to X$ on a set X, a point $x \in X$ is a fixpoint iff

$$F(x) = x.$$

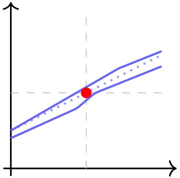

What if F is unknown?

Outline

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Fixpoints in Computer Science

Outline

Fixpoints in Computer Science

Classical Fixpoint Theory

Outline

Fixpoints in Computer Science

Classical Fixpoint Theory

Approximative Setting

Outline

Fixpoints in Computer Science

Classical Fixpoint Theory

Approximative Setting

Dampened Mann Iteration

Fixpoints in Computer Science

Examples of Fixpoints in Computer Science

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

Concurrency Theory

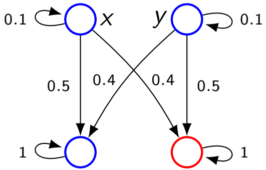

$$x \sim_\varepsilon y$$

Concurrency Theory

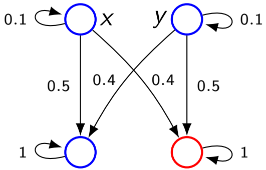

$$d(x, y) = ?$$

Concurrency Theory

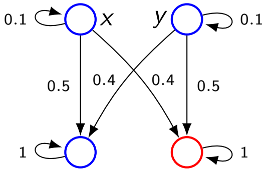

$$d(x, y) = ?$$

Model Checking

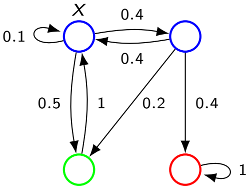

$$\Pr_x(\textcolor{blue}{blue}\,\mathcal{U}\,\textcolor{red}{red}) = ?$$

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

Concurrency Theory

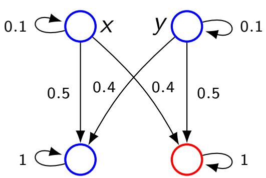

$$d(x, y) = ?$$

Model Checking

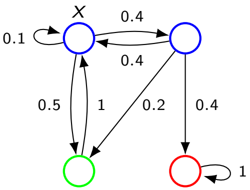

$$\mu\phi.\ \max(\textcolor{red}{red}, \min(\textcolor{blue}{blue}, \diamond\phi))$$

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

Concurrency Theory

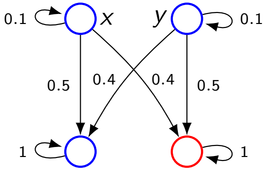

$$d(x, y) = ?$$

Model Checking

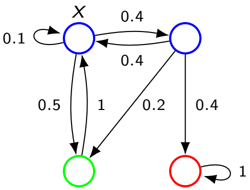

$$\mu\phi.\ \max(\textit{red}, \min(\textit{blue}, \diamond\phi))$$

Optimal Decision Making

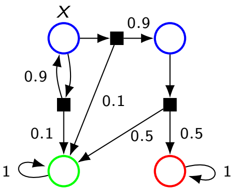

$$\arg\max_{\pi} \mathbb{V}^{\pi} = ?$$

Examples of Fixpoints in Computer Science

Concurrency Theory

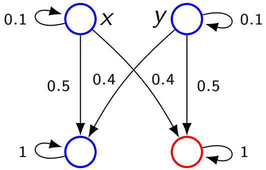

$$d(x, y) = ?$$

Model Checking

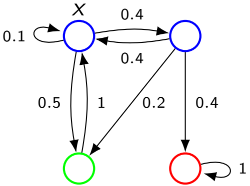

$$\mu\phi.\ \max(red, \min(blue, \diamond\phi))$$

Optimal Decision Making

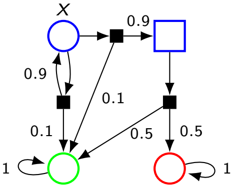

$$\arg\max_{\pi} \min_{\psi} \mathbb{V}^{(\pi,\psi)} = ?$$

Examples of Fixpoints in Computer Science

Concurrency Theory

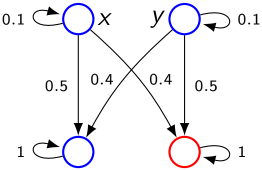

$$d(x, y) = ?$$

Model Checking

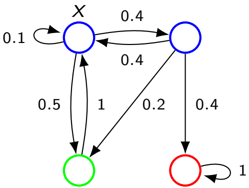

$$\mu\phi.\, \max(\textcolor{red}{\mathit{red}}, \min(\textcolor{blue}{\mathit{blue}}, \Diamond\phi))$$

Optimal Decision Making

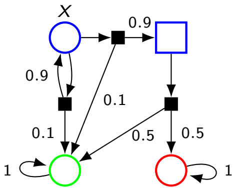

$$\arg\max_{\pi}\min_{\psi} \mathbb{V}^{(\pi,\psi)} = ?$$

Many more applications

- program analysis (dataflow analysis)
- distributed systems (consensus algorithms)
- ...

Markov decision processes (MDP) [Bellman '57]

A tuple (S, A, T, R)

- S: set of states
- A: set of actions
- $T(s' \mid s, a) \in [0, 1]$: transition probabilities
- $R(s, a) \in \mathbb{R}_+$: step-wise rewards

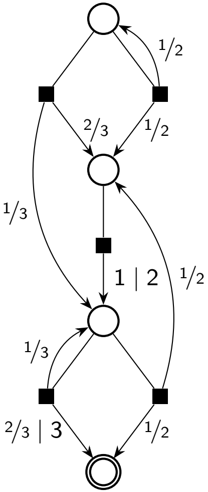

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

Markov decision processes (MDP) [Bellman '57]

A tuple (S, A, T, R)

- S: set of states
- A: set of actions
- $T(s' \mid s, a) \in [0, 1]$: transition probabilities
- $R(s, a) \in \mathbb{R}_+$: step-wise rewards

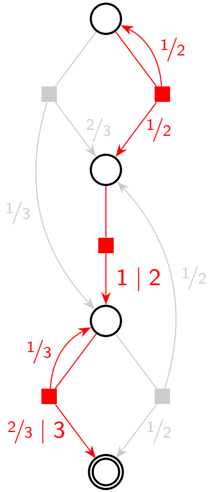

Goal: Find policy $\pi : S \to A$ optimizing expected return

$$\mathbb{V}^*_\gamma = \max_\pi \mathbb{V}^\pi_\gamma = \mathbb{E}^\pi \left[\sum_n \gamma^n R(s_n, a_n) \right]$$

for discount factor $\gamma \in [0, 1]$.

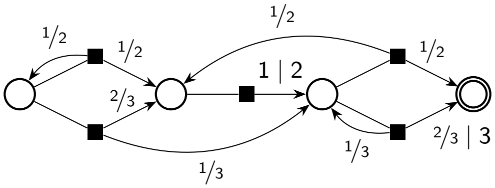

Bellman optimality operator [Bellman '57]

For an MDP $M = (S, A, T, R)$, $\mathbb{V}^*_\gamma$ can be characterized as the least fixpoint of

$$F^\gamma_M : \mathbb{R}_+^{S \times A} \to \mathbb{R}_+^{S \times A}$$

$$F^\gamma_M(\color{blue}{q}\color{black})(s, a) = R(s, a) + \gamma \sum_{s' \in S} T(s' \mid s, a) \max_{a' \in A(s')} \color{blue}{q}\color{black}(s', a')$$

Fixpoint Theory

$$F(x) = x$$

Existence

Under which conditions does a fixpoint exist?

Uniqueness

Under which conditions is the fixpoint unique?

Convergence

Under which conditions does an iteration converge to a fixpoint?

Classical Fixpoint Theory: Metric viewpoint

Setting: X is a complete metric space with metric d

Classical Fixpoint Theory: Metric viewpoint

Setting: X is a complete metric space with metric d

$\Rightarrow$ Notion of distance

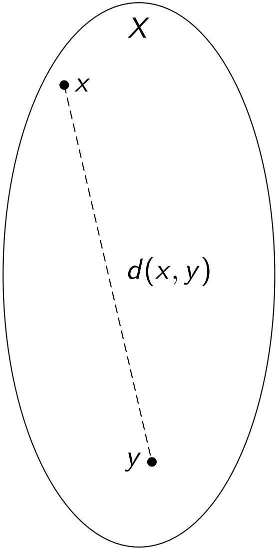

UNIVERSITÄT DUISBURG ESSEN *Offen im Denken*

Setting: X is a complete metric space with metric d

$\Rightarrow$ Notion of distance

Contraction Mapping Theorem [Banach, 1922]

Let $F : X \to X$ be a q-contraction in the complete metric space (X, d), i.e., for some $q \in (0, 1)$, we have

$$\forall x, y \in X : d(F(x), F(y)) \leq qd(x, y).$$

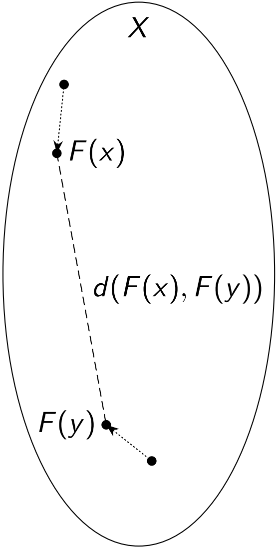

Setting: X is a complete metric space with metric d

$\Rightarrow$ Notion of distance

Contraction Mapping Theorem [Banach, 1922]

Let $F : X \to X$ be a q-contraction in the complete metric space (X, d), i.e., for some $q \in (0, 1)$, we have

$$\forall x, y \in X : d(F(x), F(y)) \leq qd(x, y).$$

Then, we have

- Existence: F has a fixpoint $x^* \in X$
- Uniqueness: x^* is the unique fixpoint of F
- Convergence: $[x_{n+1} = F(x_n)]$ converges to x^* for any $x_0 \in X$

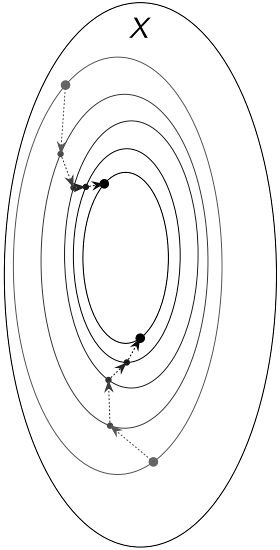

Classical Fixpoint Theory: Metric viewpoint

Setting: X is a complete metric space with metric d

$\Rightarrow$ Notion of distance

Contraction Mapping Theorem [Banach, 1922]

Let $F : X \to X$ be a q-contraction in the complete metric space (X, d), i.e., for some $q \in (0, 1)$, we have

$$\forall x, y \in X : d(F(x), F(y)) \leq qd(x, y).$$

Then, we have

- Existence: F has a fixpoint $x^* \in X$
- Uniqueness: x^* is the unique fixpoint of F
- Convergence: $[x_{n+1} = F(x_n)]$ converges to x^* for any $x_0 \in X$

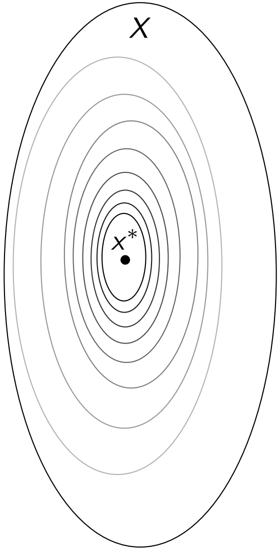

But: Contraction(-like) property required

$$F(x, y) = (y + 1, x)$$

- F is only non-expansive, i.e.,

$$\forall x, y \in X : d(F(x), F(y)) \leq d(x, y)$$

Classical Fixpoint Theory: Metric viewpoint - Contraction property

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

But: Contraction(-like) property required

$$F(x, y) = (y + 1, x)$$

- F is only non-expansive, i.e.,

$$\forall x, y \in X : d(F(x), F(y)) \leq d(x, y)$$

- No Fixpoint exists (in $\mathbb{R}^2$)
- Iteration diverges to infinity

$$F^n(x, y) \to \infty \quad \forall (x, y) \in \mathbb{R}^2$$

But: Contraction(-like) property required

$$F(x, y) = (y + 1, x)$$

- F is only non-expansive, i.e.,

$$\forall x, y \in X : d(F(x), F(y)) \leq d(x, y)$$

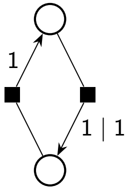

- No Fixpoint exists (in $\mathbb{R}^2$)
- Iteration diverges to infinity

$$F^n(x, y) \to \infty \quad \forall (x, y) \in \mathbb{R}^2$$

But: Contraction(-like) property required

$$F(x, y) = (y, x)$$

- Fixpoint not unique

$$F(x, x) = (x, x) \quad \forall x \in \mathbb{R}$$

- Iteration generally not convergent

$$F^n(0,1) = \begin{cases} (0,1) & \text{if } n \text{ even} \\ (1,0) & \text{if } n \text{ odd} \end{cases}$$

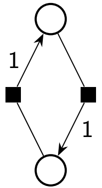

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

But: Contraction(-like) property required

End Component

A communicating subprocess, i.e., a subprocess in which every state can be reached by any other state.

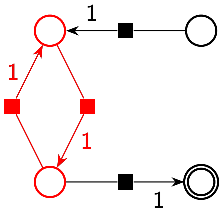

But: Contraction(-like) property required

End Component

A communicating subprocess, i.e., a subprocess in which every state can be reached by any other state.

T^1_M (power-)contraction $\Leftrightarrow$ M has no end components

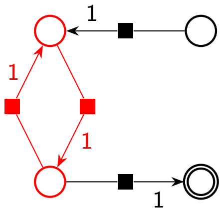

Classical Fixpoint Theory: Order-theoretic viewpoint

Setting: X is a pointed dcpo with partial order $\sqsubseteq$

Classical Fixpoint Theory: Order-theoretic viewpoint

Setting: X is a pointed dcpo with partial order $\sqsubseteq$

$\Rightarrow$ Notion of order

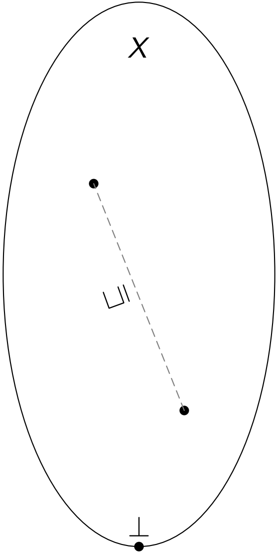

Classical Fixpoint Theory: Order-theoretic viewpoint

Setting: X is a pointed dcpo with partial order $\sqsubseteq$

$\Rightarrow$ Notion of order

Kleene Fixpoint Theorem [Scott '70]

Let $F : X \to X$ be a Scott-continuous function in the pointed dcpo $(X, \sqsubseteq)$, i.e., we have

$$\forall D \subseteq X \text{ directed} : \bigsqcup F(D) = F(\bigsqcup D).$$

Then, we have

- Existence and Uniqueness: F has a unique least fixpoint $\mu F \in X$
- Convergence: $[x_{n+1} = F(x_n)]$ fulfills $\mu F = \bigsqcup_{n \in \mathbb{N}} x_n$ for any $x_0 \in X$ with $x_0 \sqsubseteq \mu F$.

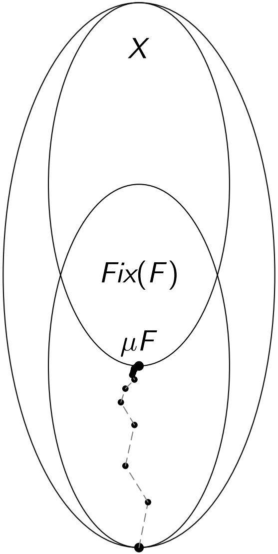

Application to Stochastic Decision Making

$$F^\gamma : \mathbb{R}_+^{S \times A} \to \mathbb{R}_+^{S \times A}$$

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

Application to Stochastic Decision Making

$$F^\gamma : \mathbb{R}_+^{S \times A} \to \mathbb{R}_+^{S \times A}$$
$$F^\gamma(q)(s,a) = R(s,a) + \gamma \sum_{s' \in S} T(s' \mid s,a) \max_{a' \in A(s')} q(s',a')$$

Metric viewpoint

- $\mathbb{R}_+^{S \times A}$ complete metric space
- F_γ γ-contractive for $\gamma < 1$

Banach's fixpoint theorem gives:

$$q_{n+1} = F_\gamma(q_n) \to \mathbb{V}_\gamma \in \mathbb{R}_+^{S \times A}$$

for any $q_0 \in \mathbb{R}_+^{S \times A}$.

Application to Stochastic Decision Making

$$F^\gamma : \mathbb{R}_+^{S\times A} \to \mathbb{R}_+^{S\times A}$$
$$F^\gamma(q)(s,a) = R(s,a) + \gamma \sum_{s'\in S} T(s' \mid s,a) \max_{a'\in A(s')} q(s',a')$$

Metric viewpoint

- $\mathbb{R}_+^{S\times A}$ complete metric space
- F_γ γ-contractive for $\gamma < 1$

Banach's fixpoint theorem gives:

$$q_{n+1} = F_\gamma(q_n) \to \mathbb{V}_\gamma \in \mathbb{R}_+^{S\times A}$$

for any $q_0 \in \mathbb{R}_+^{S\times A}$.

Order-theoretic viewpoint

- $(\mathbb{R}_+ \cup \{\infty\})^{S\times A}$ pointed dcpo
- F_γ monotone and non-expansive

Kleene's fixpoint theorem gives:

$$q_{n+1} = F_\gamma(q_n) \to \mathbb{V}_\gamma \in (\mathbb{R}_+ \cup \{\infty\})^{S\times A}$$

for any $q_0 \in \mathbb{R}_+^{S\times A}$ fulfilling $q_0 \leq \mathbb{V}_\gamma$.

Approximative Settings

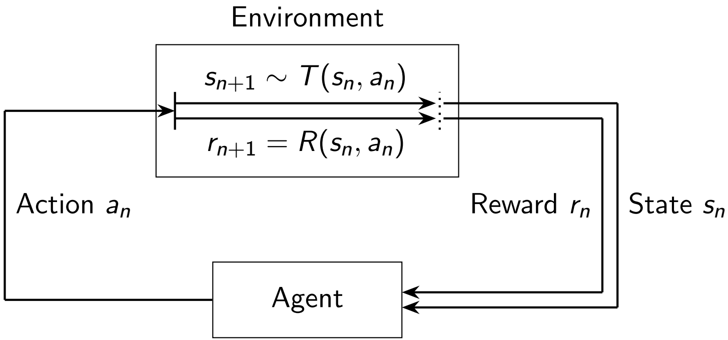
Environment
$s_{n+1} \sim T(s_n, a_n)$
$r_{n+1} = R(s_n, a_n)$
Action a_n
Reward r_n
State s_n
Agent

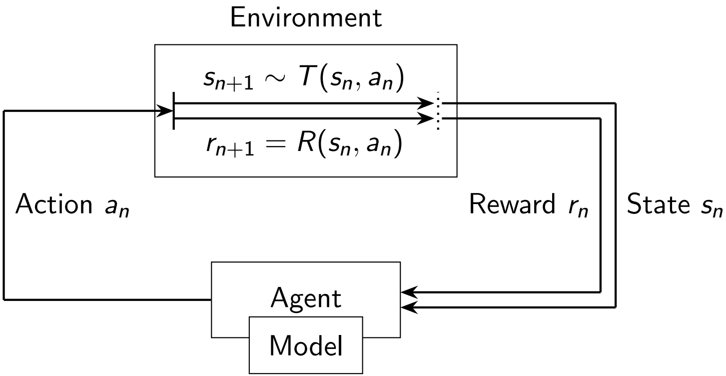

Agent learns model $M_n = (S, A, T_n, R_n)$ of (unknown) environment $M = (S, A, T, R)$

$$F^{\gamma}_{M_n} \to F^{\gamma}_{M} \text{ (pointwise)}$$

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

Dyna Algorithm [Sutton '90, Kaelbling '96]

For new experience (s, a, r, s') do:

1. Update model M based on (s, a, r, s')
2. $Q'(s, a) \leftarrow R(s, a) + \gamma \sum_{\hat{s}} T(s, a, \hat{s}) \max_{\hat{a}} Q(\hat{s}, \hat{a})$
3. Do k times:
   3.1 Sample random state-action pair (s, a)
   3.2 $Q'(s, a) \leftarrow R(s, a) + \gamma \sum_{\hat{s}} T(s, a, \hat{s}) \max_{\hat{a}} Q(\hat{s}, \hat{a})$
4. Select action a' based on Q

Dyna Algorithm [Sutton '90, Kaelbling '96]

For new experience (s, a, r, s') do:

1. Update model M based on (s, a, r, s')
2. $Q'(s, a) \leftarrow R(s, a) + \gamma \sum_{\hat{s}} T(s, a, \hat{s}) \max_{\hat{a}} Q(\hat{s}, \hat{a})$
3. Do k times:
   3.1 Sample random state-action pair (s, a)
   3.2 $Q'(s, a) \leftarrow R(s, a) + \gamma \sum_{\hat{s}} T(s, a, \hat{s}) \max_{\hat{a}} Q(\hat{s}, \hat{a})$
4. Select action a' based on Q

$\Rightarrow$ Algorithm is a chaotic version of $[x_{n+1} = F_n(x_n)]$ with $F_n \to F$ (pointwise).

What if F cannot be used?

What if F cannot be used?

Is it possible to compute fixpoints of F using only approximations $F_n \to F$?

What if F cannot be used?

Is it possible to compute fixpoints of F using only approximations $F_n \to F$?

Possible Reasons:

- System dynamics are unknown
- Exact system too complex
- Approximated partial results
- Discretization / Rounding errors

Approximative Fixpoint Theory?

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

Continuity of Fixpoint Operators

Given approximations $F_n \to F$, can we say something about the limit $\lim_{n\to\infty} x_n^*$ of (least) fixpoints of F_n?

Approximative Fixpoint Theory?

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

Continuity of Fixpoint Operators

Given approximations $F_n \to F$, can we say something about the limit $\lim_{n\to\infty} x_n^*$ of (least) fixpoints of F_n?

Convergence of Fixpoint Iterations

Given approximations $F_n \to F$, can we compute a values $x_n \in X$ converging to a (least) fixpoint of F, with x_n only depending on $F_0, \ldots, F_{n-1}$?

Related Theories

Stability of Iterations [Ostrowski '67, Rhoades '90, Berinde '02, ...]

Stability of iteration scheme $T(F,x)$:

$$x_{n+1} = T(F, x_n) \to x^* \quad \Rightarrow \quad \forall \varepsilon_n \to 0 : y_{n+1} = T(F, y_n) + \varepsilon_n \to x^*$$

Stability of Iterations [Ostrowski '67, Rhoades '90, Berinde '02, ...]

Stability of iteration scheme $T(F, x)$:

$$x_{n+1} = T(F, x_n) \to x^* \quad \Rightarrow \quad \forall \varepsilon_n \to 0 : y_{n+1} = T(F, y_n) + \varepsilon_n \to x^*$$

$\Rightarrow$ Closely related to the case where $F_n \to F$ uniformly (in smooth spaces)

Related Theories

Stability of Iterations [Ostrowski '67, Rhoades '90, Berinde '02, . . .]

Stability of iteration scheme $T(F, x)$:

$$x_{n+1} = T(F, x_n) \to x^* \quad \Rightarrow \quad \forall \varepsilon_n \to 0 : y_{n+1} = T(F, y_n) + \varepsilon_n \to x^*$$

$\Rightarrow$ Closely related to the case where $F_n \to F$ uniformly (in smooth spaces)

Stochastic Approximation [Robbins/Monro '51, Bertsekas/Tsitsiklis '96, . . .]

Convergence with stochastic error term:

$$x_{n+1} = T(F + \omega_n, x_n), \quad \mathbb{E}[\omega_n \mid \mathcal{F}_n] = 0$$

Related Theories

Stability of Iterations [Ostrowski '67, Rhoades '90, Berinde '02, ...]

Stability of iteration scheme $T(F, x)$:

$$x_{n+1} = T(F, x_n) \to x^* \quad \Rightarrow \quad \forall \varepsilon_n \to 0 : y_{n+1} = T(F, y_n) + \varepsilon_n \to x^*$$

$\Rightarrow$ Closely related to the case where $F_n \to F$ uniformly (in smooth spaces)

Stochastic Approximation [Robbins/Monro '51, Bertsekas/Tsitsiklis '96, ...]

Convergence with stochastic error term:

$$x_{n+1} = T(F + \omega_n, x_n), \quad \mathbb{E}[\omega_n \mid \mathcal{F}_n] = 0$$

But: Both focus on functions with contraction-like property / unique fixpoint

Continuity of Fixpoints [Nadler Jr '68]

Let $F_n : X \to X$ be contractions in the locally compact metric space (X, d) converging (pointwise) to the contraction $F : X \to X$. Then, the (unique) fixpoints x_n^* of F_n converge to the fixpoint x^* of F.

Continuity of Fixpoints [Nadler Jr '68]

Let $F_n : X \to X$ be contractions in the locally compact metric space (X, d) converging (pointwise) to the contraction $F : X \to X$. Then, the (unique) fixpoints x_n^* of F_n converge to the fixpoint x^* of F.

Convergence of Fixpoint Iteration [Ostrowski '67]

Let $F_n : X \to X$ be q-contractions in the complete metric space (X, d) converging (pointwise) to the (q-)contraction $F : X \to X$. Then, the iteration $[x_{n+1} = F_n(x_n)]$ converges to the fixpoint x^* of F for any $x_0 \in X$.

Non-expansive Setting

We are given

- a closed and convex subset $X \subseteq \mathbb{R}^d_+$ with $0 \in X$, attributed with the supremum distance
$$d_\infty(x, y) = \max_{i \in \{1,\ldots,d\}} |x(i) - y(i)|$$
- a (pointwise) converging sequence of monotone and non-expansive functions $F_n : X \to X$

In particular, we are not given the (pointwise) limit $F := \lim_{n\to\infty} F_n : X \to X$ of (F_n).

Non-expansive Case: Continuity of μ

$$
\begin{aligned}
F_n : [0,1] &\to [0,1] \\
x &\mapsto (1 - {}^1\!/\!_n)x + {}^1\!/\!_n
\end{aligned}
$$

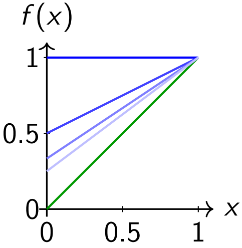

Non-expansive Case: Continuity of μ

$$\begin{aligned} F_n : [0,1] &\to [0,1] \\ x &\mapsto (1 - {}^1\!/\!_n)x + {}^1\!/\!_n \end{aligned}$$

- F_n is $(1 - {}^1\!/\!_n)$-contractive
- F_n converges (uniformly) to $F = id$

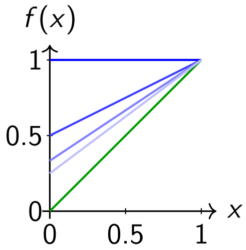

Non-expansive Case: Continuity of μ

$$\begin{aligned} F_n : [0,1] &\to [0,1] \\ x &\mapsto (1 - {}^1/_n)x + {}^1/_n \end{aligned}$$

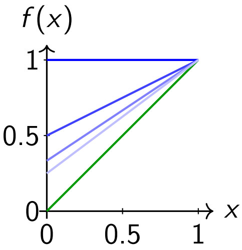

- F_n is $(1 - {}^1/_n)$-contractive
- F_n converges (uniformly) to $F = id$

But: $\mu F_n = 1 \not\to 0 = \mu F$

$$x_{n+1} = F_n(x_n) = 1 \quad \forall n \in \mathbb{N}_0, x_1 \in [0,1]$$

Non-expansive Case: Continuity of μ

$$\begin{aligned} F_n : [0,1] &\to [0,1] \\ x &\mapsto (1 - 1/n)x + 1/n \end{aligned}$$

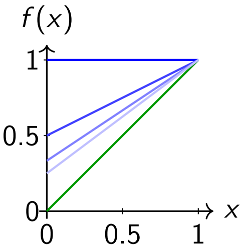

- F_n is $(1 - 1/n)$-contractive
- F_n converges (uniformly) to $F = id$

But: $\mu F_n = 1 \not\to 0 = \mu F$

$$x_{n+1} = F_n(x_n) = 1 \quad \forall n \in \mathbb{N}_0, x_1 \in [0,1]$$

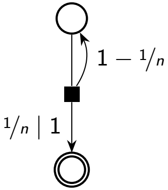

Non-expansive Case: Overapproximations

$$F_n : [0,1] \to [0,1]$$
$$x \mapsto \max(k + {}^{(1-k)}/_{n}, x)$$

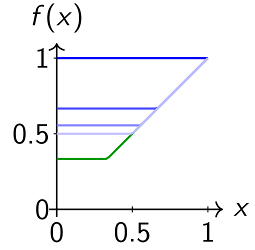

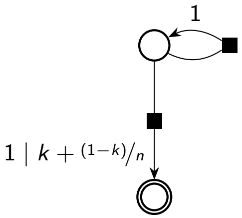

Non-expansive Case: Overapproximations

$$F_n : [0,1] \to [0,1]$$
$$x \mapsto \max(k + {}^{(1-k)}/_{n}, x)$$

- F_n converges (uniformly) to $F = \max(k, id)$
- $\mu F_n = k + {}^{(1-k)}/_{n} \to k = \mu F$

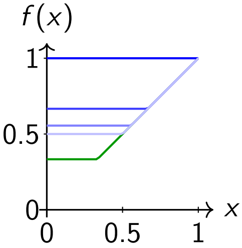

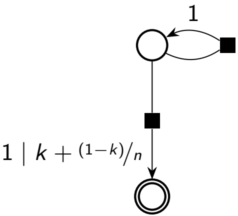

Non-expansive Case: Overapproximations

$$F_n : [0,1] \to [0,1]$$
$$x \mapsto \max(k + {}^{(1-k)}/_{n}, x)$$

- F_n converges (uniformly) to $F = \max(k, id)$
- $\mu F_n = k + {}^{(1-k)}/_{n} \to k = \mu F$

But: Iterating with any F_n over-estimates μF and further iterations cannot decrease it

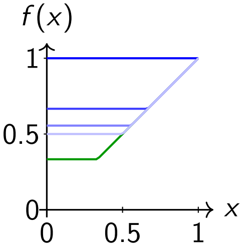

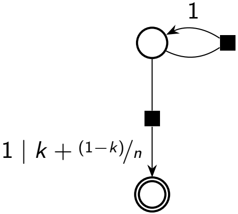

UNIVERSITÄT DUISBURG ESSEN *Offen im Denken*

$$F_n : [0,1] \to [0,1]$$
$$x \mapsto \max(k + {}^{(1-k)}/_{n}, x)$$

- F_n converges (uniformly) to $F = \max(k, id)$
- $\mu F_n = k + {}^{(1-k)}/_{n} \to k = \mu F$

But: Iterating with any F_n over-estimates μF and further iterations cannot decrease it

Related to computation of upper bounds in exact case, e.g., Bounded VI [Eisentraut et al. '22], Optimistic VI [Hartmanns/Kaminski '20]

$\Rightarrow$ depend on collapsing / deflating end components

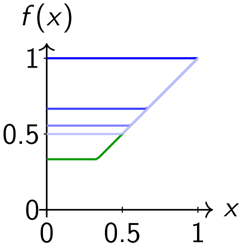

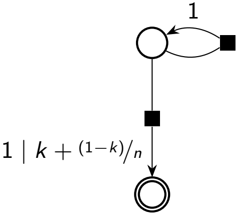

Dampened Mann Iteration

Finding Fixpoints of Non-expansive Functions: Metric viewpoint

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

Krasnoselskii-Mann iteration [Mann ’53, Krasnoselskii ’55]

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

for learning rates $\alpha_n \in (0, 1)$

Krasnoselskii-Mann iteration [Mann '53, Krasnoselskii '55]

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

for learning rates $\alpha_n \in (0, 1)$

[Ishikawa '76, Reich '79]: Convergence results for non-expansive functions in the exact case for appropriate learning rates α_n

Finding Fixpoints of Non-expansive Functions: Metric viewpoint

Krasnoselskii-Mann iteration [Mann '53, Krasnoselskii '55]

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

for learning rates $\alpha_n \in (0, 1)$

[Ishikawa '76, Reich '79]: Convergence results for non-expansive functions in the exact case for appropriate learning rates α_n

Some results in (smooth) approximative settings [Bauschke/Combettes '17]

Krasnoselskii-Mann iteration [Mann '53, Krasnoselskii '55]

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

for learning rates $\alpha_n \in (0,1)$

[Ishikawa '76, Reich '79]: Convergence results for non-expansive functions in the exact case for appropriate learning rates α_n

Some results in (smooth) approximative settings [Bauschke/Combettes '17]

But: Convergence to *some* fixpoint

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

Dampened Mann iteration

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

for learning rates $\alpha_n \in (0, 1)$ and dampening factors $\beta_n \in (0, 1)$

Dampened Mann Iteration

Dampened Mann iteration

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

for learning rates $\alpha_n \in (0,1)$ and dampening factors $\beta_n \in (0,1)$

[Kim/Xu '05]: Convergence in exact case to fixpoint closest to 0 in (uniformly smooth) Banach spaces if fixpoint exists and α_n, β_n sufficiently regular

Dampening Factors

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

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

Dampening Factors

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

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

Intuition: $(1 - \beta_n)$ as a form of ‘vanishing discount / contraction factor’.

Dampening Factors

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

Intuition: $(1 - \beta_n)$ as a form of 'vanishing discount / contraction factor'.

Required conditions:

- $\beta_n/\alpha_n \to 0$: More 'learning' than 'dampening'
- $\sum_n \beta_n = \infty$: Always enough 'power' left to decrease potential overapproximation (not needed for (power-)contractions)

Dampening Factors

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

Intuition: $(1 - \beta_n)$ as a form of 'vanishing discount / contraction factor'.

Required conditions:

- $\beta_n/\alpha_n \to 0$: More 'learning' than 'dampening'
- $\sum_n \beta_n = \infty$: Always enough 'power' left to decrease potential overapproximation (not needed for (power-)contractions)

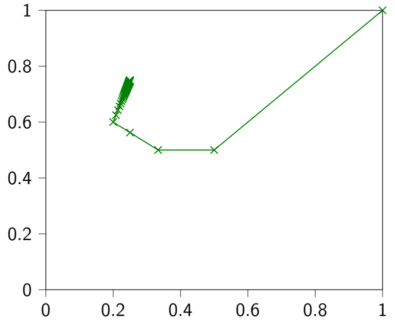

Our Convergence Results

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Convergence of Dampened Mann Iteration [Baldan et al CAV '25, TACAS '26]

For monotone and non-expansive approximations $F_n : X \to X$ of $F : X \to X$ (w.r.t. $\|\cdot\|_\infty$) and sequences $\alpha_n, \beta_n \in [0,1]$ fulfilling

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

the sequence (x_n) generated by the dampened Mann iteration converges to $\mu F \in [0,\infty]^d$ for any initial $x_0 \in X$ provided that

- the target function F is 'sufficiently regular' *or*
- the convergence $F_n \to F$ is 'sufficiently regular' *or*
- the convergence $F_n \to F$ is 'sufficiently fast'

Convergence Criteria: Regular Target Function

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

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

Convergence when

$$F \text{ is a Picard operator}$$

that means $F^n(x_0) \to x^*$ for any $x_0 \in X$.

Convergence Criteria: Regular Target Function

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

Convergence when

F is a Picard operator

that means $F^n(x_0) \to x^*$ for any $x_0 \in X$.

Application to MDPs: Iteration works for any approximation of an MDP without any end components
Note: Any (close) approximation of an MDP without end components has no end components

Convergence Criteria: Regular Convergence

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

Convergence when

$$\lim_{n \to \infty} \mu(\sup_{k > n} F_k) = \mu F$$

Convergence Criteria: Regular Convergence

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

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

Convergence when

$$\lim_{n\to\infty} \mu(\sup_{k>n} F_k) = \mu F$$

Application to MDPs: Iteration works when approximations are eventually structure-preserving, i.e., for $s, s' \in S, a \in A(s)$

- $T(s' \mid s, a) > 0 \Leftrightarrow T_n(s' \mid s, a) > 0$ eventually
- $R(s, a) > 0 \Leftrightarrow R_n(s, a) > 0$ eventually

Convergence Criteria: Fast Convergence

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

Convergence when

$$\sum_n \alpha_n d_\mu(F, F_n) < \infty$$

where $d_\mu(F, F_n) = \sup\{\|F(x) - F_n(x)\| \mid x \leq \mu F\}$.

Convergence Criteria: Fast Convergence

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

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

Convergence when

$$\sum_n \alpha_n d_\mu(F, F_n) < \infty$$

where $d_\mu(F, F_n) = \sup\{\|F(x) - F_n(x)\| \mid x \leq \mu F\}$.

Intuition: Only a finite amount of error mass is incorporated in the computation.

Convergence Criteria: Fast Convergence

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

Convergence when

$$\sum_n \alpha_n d_\mu(F, F_n) < \infty$$

where $d_\mu(F, F_n) = \sup\{\|F(x) - F_n(x)\| \mid x \leq \mu F\}$.

Intuition: Only a finite amount of error mass is incorporated in the computation.

Consequence:

- If F_n converges 'fast enough' to F, convergence of the iteration is guaranteed.
- If the errors $d_\mu(F, F_n)$ can be estimated, the parameters α_n and β_n can be chosen appropriately to guarantee convergence.

Chaotic Iteration

Similar results for

$$\begin{cases} x_{n+1}(i) = (1-\beta_n(i))(x_n(i) + \alpha_n(i)(F(x_n)(i) - x_n(i))) & i \in S_n \\ x_{n+1}(i) = x_n(i) & i \notin S_n \end{cases}$$

under sufficient fairness conditions on the sequence (S_n) of updated components [Baldan et al TACAS'26].

Similar results for

$$\begin{cases} x_{n+1}(i) = (1 - \beta_n(i))(x_n(i) + \alpha_n(i)(F(x_n)(i) - x_n(i))) & i \in S_n \\ x_{n+1}(i) = x_n(i) & i \notin S_n \end{cases}$$

under sufficient fairness conditions on the sequence (S_n) of updated components [Baldan et al TACAS'26].

Stochastic guarantees

If the updated parameters S_n are chosen randomly and there exists $k \in \mathbb{N}$ and $\varepsilon > 0$ such that in every time window $[n, n+k]$, every component is updated with probability at least ε, then $x_{n+1} \to \mu F$ almost surely.

(Model-based) Reinforcement Learning - Dyna Algorithm

Dyna Algorithm [Sutton '90, Kaelbling '96]

For new experience (s, a, r, s') do:

1. Update model M based on (s, a, r, s')
2. $Q'(s,a) \leftarrow R(s,a) + \gamma \sum_{\hat{s}} T(s,a,\hat{s}) \max_{\hat{a}} Q(\hat{s},\hat{a})$// $q_{n+1}(s,a) = F(q_n)(s,a)$
3. Do k times:
   - 3.1 Sample random state-action pair (s,a)// Ensures fair update
   - 3.2 $Q'(s,a) \leftarrow R(s,a) + \gamma \sum_{\hat{s}} T(s,a,\hat{s}) \max_{\hat{a}} Q(\hat{s},\hat{a})$// $q_{n+1}(s,a) = F(q_n)(s,a)$
4. Select action a' based on Q

$\Rightarrow$ Almost-sure convergence as long as exploration ensures correct model learning

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

Dyna algorithm for undiscounted setting

For new experience (s, a, r, s') do:

1. Update model M based on (s, a, r, s')
2. $Q'(s,a) \leftarrow (1-\beta_n)(\alpha_n R(s,a) + \sum_{\hat{s}} T(s,a,\hat{s}) \max_{\hat{a}}(\alpha_n Q(\hat{s}, \hat{a} + (1-\alpha_n)Q(s,a))))$
3. Do k times:
   - 3.1 Sample random state-action pair (s, a)
   - 3.2 $Q'(s,a) \leftarrow (1-\beta_n)(\alpha_n R(s,a) + \sum_{\hat{s}} T(s,a,\hat{s}) \max_{\hat{a}}(\alpha_n Q(\hat{s}, \hat{a} + (1-\alpha_n)Q(s,a))))$
4. Select action a' based on Q

$\Rightarrow$ Framework gives almost-sure convergence for undiscounted case

Looking ahead: Model-free Reinforcement Learning

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

Q-learning [Watkins '89]

For new experience (s, a, r, s') do:

1. $Q'(s,a) \leftarrow Q(s,a) + \alpha_n(r + \gamma \max_{a' \in A(s')} Q(s',a') - Q(s,a))$ // direct update
2. Select action a' based on Q

$\Rightarrow$ Chaotic version of $[x_{n+1} = F_n(x_n)]$ without $F_n \to F$ (even pointwise)

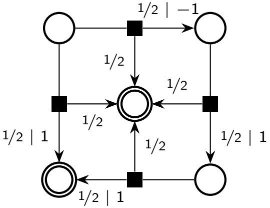

Q-learning [Watkins '89]

For new experience (s, a, r, s') do:

1. $Q'(s,a) \leftarrow Q(s,a) + \alpha_n(r + \gamma \max_{a' \in A(s')} Q(s', a') - Q(s,a))$ // direct update
2. Select action a' based on Q

$\Rightarrow$ Chaotic version of $[x_{n+1} = F_n(x_n)]$ without $F_n \to F$ (even pointwise)

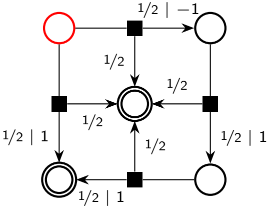

Q-learning [Watkins '89]

For new experience (s, a, r, s') do:

1. $Q'(s,a) \leftarrow Q(s,a) + \alpha_n(r + \gamma \max_{a' \in A(s')} Q(s', a') - Q(s,a))$ // direct update
2. Select action a' based on Q

$\Rightarrow$ Chaotic version of $[x_{n+1} = F_n(x_n)]$ without $F_n \to F$ (even pointwise)

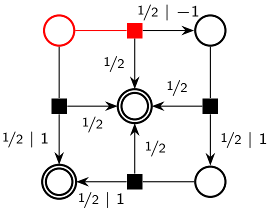

Q-learning [Watkins ’89]

For new experience (s, a, r, s') do:

1. $Q'(s,a) \leftarrow Q(s,a) + \alpha_n(r + \gamma \max_{a' \in A(s')} Q(s', a') - Q(s,a))$ // direct update
2. Select action a' based on Q

$\Rightarrow$ Chaotic version of $[x_{n+1} = F_n(x_n)]$ without $F_n \to F$ (even pointwise)

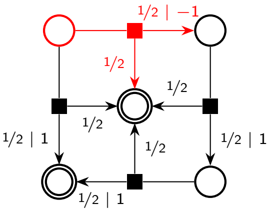

Q-learning [Watkins '89]

For new experience (s, a, r, s') do:

1. $Q'(s, a) \leftarrow Q(s, a) + \alpha_n(r + \gamma \max_{a' \in A(s')} Q(s', a') - Q(s, a))$ // direct update
2. Select action a' based on Q

$\Rightarrow$ Chaotic version of $[x_{n+1} = F_n(x_n)]$ without $F_n \to F$ (even pointwise)

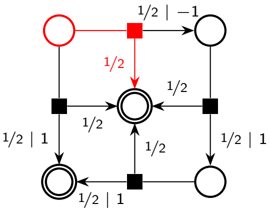

Q-learning [Watkins '89]

For new experience (s, a, r, s') do:

1. $Q'(s,a) \leftarrow Q(s,a) + \alpha_n(r + \gamma \max_{a' \in A(s')} Q(s',a') - Q(s,a))$ // direct update
2. Select action a' based on Q

$\Rightarrow$ Chaotic version of $[x_{n+1} = F_n(x_n)]$ ~~without~~ with $1/n \sum_{k=1}^{n} F_k \to F$ (Césaro limit)

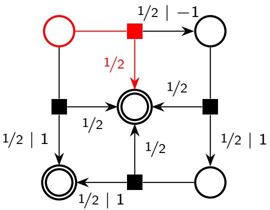

Dampened Mann iteration

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

- Compute (least) fixpoints of monotone and non-expansive functions F using only approximations $F_n \to F$
- Combining metric and order-theoretic reasoning for convergence results whenever
  - the target function is sufficiently 'regular' (Picard operator) *or*
  - the convergence $F_n \to F$ is sufficiently 'regular' ($\mu F = \lim \mu(\sup_{k>n} F_k)$) *or*
  - the convergence $F_n \to F$ is sufficiently 'fast' ($\sum \alpha_n \| F - F_n \| < \infty$)