Fuzzy and Probabilistic van Benthem Theorems

Paul Wild

Friedrich-Alexander-Universität Erlangen-Nürnberg

June 4, 2018

A fuzzy modal logic (1)

Syntax:

$$\phi, \psi ::= c \mid p \mid \neg\phi \mid \phi \ominus c \mid \phi \wedge \psi \mid \Diamond\phi$$

where $c \in \mathbb{Q} \cap [0, 1]$ and p ranges over the set At of propositional atoms.

We denote the set of formulas of rank at most n by $\mathcal{L}_n$.

A fuzzy modal logic (1)

Syntax:

$$\phi, \psi ::= c \mid p \mid \neg\phi \mid \phi \ominus c \mid \phi \wedge \psi \mid \Diamond\phi$$

where $c \in \mathbb{Q} \cap [0,1]$ and p ranges over the set At of propositional atoms.

We denote the set of formulas of rank at most n by $\mathcal{L}_n$.

Semantics: given over *fuzzy relational models*

$$\mathcal{A} = (A, (p^{\mathcal{A}})_{p \in \mathsf{At}}, R^{\mathcal{A}})$$

- A is the set of states.
- $p^{\mathcal{A}} : A \to [0,1]$ is the interpretation for $p \in$ At.
- $R^{\mathcal{A}} : A \times A \to [0,1]$ is the transition relation.

A fuzzy modal logic (2)

Every formula ϕ has an interpretation $\phi_{\mathcal{A}} : A \to [0, 1]$, which is defined inductively (we omit the subscript $\mathcal{A}$ if clear from context).

A fuzzy modal logic (2)

Every formula ϕ has an interpretation $\phi_{\mathcal{A}} : A \to [0,1]$, which is defined inductively (we omit the subscript $\mathcal{A}$ if clear from context).

Let $\mathcal{A} = (A, (p^{\mathcal{A}})_{p \in \mathsf{At}}, R^{\mathcal{A}})$ and $a \in A$.

- constants: $c(a) = c$
- propositions: $p(a) = p^{\mathcal{A}}(a)$
- negation: $(\neg\phi)(a) = 1 - \phi(a)$
- modified subtraction: $(\phi \ominus c)(a) = \max(\phi(a) - c, 0)$
- conjunction: $(\phi \wedge \psi)(a) = \min(\phi(a), \psi(a))$
- modality: $(\Diamond\phi)(a) = \sup_{a' \in A} \min(R^{\mathcal{A}}(a, a'), \phi(a'))$

Notations: $a \wedge b = \min(a,b), a \vee b = \max(a,b), \bigvee = \sup, \bigwedge = \inf$.

Fuzzy first order logic

Syntax:

$$\phi, \psi ::= c \mid p(x) \mid R(x, y) \mid x = y \mid \neg\phi \mid \phi \ominus c \mid \phi \wedge \psi \mid \exists x.\phi$$

where $c \in \mathbb{Q} \cap [0, 1]$, $p \in \mathsf{At}$, and x, y range over a fixed countably infinite set of variables.

Fuzzy first order logic

Syntax:

$$\phi, \psi ::= c \mid p(x) \mid R(x,y) \mid x = y \mid \neg\phi \mid \phi \ominus c \mid \phi \wedge \psi \mid \exists x.\phi$$

where $c \in \mathbb{Q} \cap [0,1]$, $p \in \mathsf{At}$, and x, y range over a fixed countably infinite set of variables.

Semantics over a fuzzy relational model $\mathcal{A} = (A, (p^{\mathcal{A}})_{p \in \mathsf{At}}, R^{\mathcal{A}})$:

- Let η be a valuation mapping variables to elements of A
- $\phi_{\mathcal{A}}(\eta) \in [0,1]$ (or just $\phi(\eta)$) is defined inductively:
  - Boolean connectives and equality as expected
  - $p(x)(\eta) = p^{\mathcal{A}}(\eta(x)), R(x,y)(\eta) = R^{\mathcal{A}}(\eta(x), \eta(y))$
  - existential quantification: $(\exists x.\phi)(\eta) = \bigvee_{a \in A} \phi(\eta[x \mapsto a])$

Fuzzy first order logic

Syntax:

$$\phi, \psi ::= c \mid p(x) \mid R(x,y) \mid x = y \mid \neg\phi \mid \phi \ominus c \mid \phi \wedge \psi \mid \exists x.\phi$$

where $c \in \mathbb{Q} \cap [0,1]$, $p \in \mathsf{At}$, and x, y range over a fixed countably infinite set of variables.

Semantics over a fuzzy relational model $\mathcal{A} = (A, (p^{\mathcal{A}})_{p \in \mathsf{At}}, R^{\mathcal{A}})$:

- Let η be a valuation mapping variables to elements of A
- $\phi_{\mathcal{A}}(\eta) \in [0,1]$ (or just $\phi(\eta)$) is defined inductively:
  - Boolean connectives and equality as expected
  - $p(x)(\eta) = p^{\mathcal{A}}(\eta(x)), R(x,y)(\eta) = R^{\mathcal{A}}(\eta(x), \eta(y))$
  - existential quantification: $(\exists x.\phi)(\eta) = \bigvee_{a \in A} \phi(\eta[x \mapsto a])$

Fuzzy modal logic is a fragment of fuzzy FOL via the usual standard translation:

$$\mathsf{ST}_x(\Diamond\phi) = \exists y. R(x,y) \wedge \mathsf{ST}_y(\phi).$$

Coalgebraic view

Consider the endofunctors F and G on Set:

$$\mathsf{F}X = [0,1]^X, \quad \mathsf{F}f(g)(y) = \bigvee_{f(x)=y} g(x)$$

where $f\colon X \to Y, g \in [0,1]^X, y \in Y$.

$$\mathsf{G}X = [0,1]^{\mathsf{At}} \times \mathsf{F}X$$

Every fuzzy relational model $\mathcal{A} = (A, (p^{\mathcal{A}})_{p \in \mathsf{At}}, R^{\mathcal{A}})$ can be seen as a coalgebra $\alpha\colon A \to \mathsf{G}A$:

$$\alpha(a) = (\lambda p.p^{\mathcal{A}}(a),\ \lambda a'.R^{\mathcal{A}}(a,a')).$$

Bisimulation games

The ε-*bisimulation game* is played between the spoiler S and the duplicator D on two models $\mathcal{A}, \mathcal{B}$.

- Configurations: pairs of states $a \in A$ and $b \in B$.
- Moves:
  - S picks some $a' \in A$ such that $R^{\mathcal{A}}(a, a') > \varepsilon$
  - D picks some $b' \in B$ such that $R^{\mathcal{B}}(b, b') \geq R^{\mathcal{A}}(a, a') - \varepsilon$.
  - New configuration: (a', b').

  S may also swap the roles of $\mathcal{A}$ and $\mathcal{B}$.
- Whoever is unable to move loses.
- Winning condition for D *before* every round: $|p^{\mathcal{A}}(a) - p^{\mathcal{B}}(b)| \leq \varepsilon$ for all $p \in \mathsf{At}$.

The game can be played either indefinitely or over a fixed number of rounds n.

Pseudometric spaces

A *pseudometric space* (X, d) consists of a set X, together with a map $d : X \times X \to [0, \infty)$ satisfying:

- $d(x, x) = 0$ for all $x \in X$;
- $d(x, y) = d(y, x)$ for all $x, y \in X$;
- $d(x, y) + d(y, z) \geq d(x, z)$ for all $x, y, z \in X$.

Some important notations:

- d_e denotes the Euclidean metric on $\mathbb{R}$: $d_e(x, y) = |x - y|$.
- The space of nonexpansive maps from (X, d_1) to (Y, d_2) is denoted by $(X, d_1) \to_1 (Y, d_2)$.
- For $f, g \colon X \to \mathbb{R}$ put $\|f - g\|_\infty = \sup_{x \in X} |f(x) - g(x)|$.

Behavioural distance

Given a model $\mathcal{A}$, we obtain several notions of behavioural distance on A:

- via the game:

$$d_n^G(a,b) = \bigwedge\{\varepsilon \mid D \text{ wins the } n\text{-round } \varepsilon\text{-bisim. game for } a,b\}$$

- via the modal formulas up to some rank:

$$d_n^L(a,b) = \bigvee_{\phi \in \mathcal{L}_n} |\phi(a) - \phi(b)|$$

$\mathcal{L}_n$ denotes the set of modal formulas of rank at most n.

- via the so-called *Kantorovich lifting*:

$$d_0^K(a,b) = 0$$

$$d_{n+1}^K(a,b) = \bigvee_{p \in \mathrm{At}} |p(a) - p(b)| \vee \bigvee_{f:(A,d_n^K) \to_1 ([0,1],d_e)} |(\Diamond f)(a) - (\Diamond f)(b)|$$

where $(\Diamond f)(a) = \bigvee_{a' \in A} R^{\mathcal{A}}(a,a') \wedge f(a')$.

Modal approximation

Theorem

Let $\mathcal{A}$ be a model. For all $n \geq 0$:

1. $d_n^G = d_n^K = d_n^L =: d_n$ *on* $\mathcal{A}$.
2. (A, d_n) *is totally bounded.*
3. $\mathcal{L}_n$ *is a dense subset of* $(A, d_n) \to_1 ([0,1], d_e)$.

Bisimulation invariance

Consider a model $\mathcal{A}$ together with the pseudometric d^G arising from the unbounded bisimulation game:

$$d^G(a,b) = \bigwedge\{\varepsilon \mid D \text{ wins the unbounded } \varepsilon\text{-bisim. game for } a,b\}.$$

Bisimulation invariance

Consider a model $\mathcal{A}$ together with the pseudometric d^G arising from the unbounded bisimulation game:

$$d^G(a,b) = \bigwedge\{\varepsilon \mid D \text{ wins the unbounded } \varepsilon\text{-bisim. game for } a, b\}.$$

A formula ϕ (modal or first order) is *bisimulation invariant* if its interpretation $\phi_{\mathcal{A}}$ is nonexpansive with respect to d^G.

Lemma

Every fuzzy modal formula is bisimulation invariant.

Bisimulation invariance

Consider a model $\mathcal{A}$ together with the pseudometric d^G arising from the unbounded bisimulation game:

$$d^G(a,b) = \bigwedge\{\varepsilon \mid D \text{ wins the unbounded } \varepsilon\text{-bisim. game for } a,b\}.$$

A formula ϕ (modal or first order) is *bisimulation invariant* if its interpretation $\phi_{\mathcal{A}}$ is nonexpansive with respect to d^G.

Lemma

Every fuzzy modal formula is bisimulation invariant.

Lemma

Every bisimulation invariant formula is also depth n bisimulation invariant for some n, i.e. nonexpansive wrt. depth n behavioural distance d_n^G.

The fuzzy van-Benthem Theorem

This is the version of the van Benthem Theorem from the last talk:

Theorem (van Benthem Theorem, local)

Every bisimulation invariant formula ϕ of fuzzy FOL is modally approximable on every model $\mathcal{A}$, i.e. for every $\varepsilon > 0$ there exists a fuzzy modal formula ψ such that $\|\phi_{\mathcal{A}} - \psi_{\mathcal{A}}\|_\infty \leq \varepsilon$.

The fuzzy van-Benthem Theorem

This is the version of the van Benthem Theorem from the last talk:

Theorem (van Benthem Theorem, local)

Every bisimulation invariant formula ϕ of fuzzy FOL is modally approximable on every model $\mathcal{A}$, i.e. for every $\varepsilon > 0$ there exists a fuzzy modal formula ψ such that $\|\phi_{\mathcal{A}} - \psi_{\mathcal{A}}\|_\infty \leq \varepsilon$.

However, this is the version we actually want to prove:

Theorem (van Benthem Theorem)

Every bisimulation invariant formula ϕ of fuzzy FOL is modally approximable, i.e. for every $\varepsilon > 0$, there exists a fuzzy modal formula ψ such that for every fuzzy model $\mathcal{A}$ $\|\phi_{\mathcal{A}} - \psi_{\mathcal{A}}\|_\infty \leq \varepsilon$.

The final chain (1)

Consider the final chain $(F_n)_{n \geq 0}$ of G:

$$F_0 = \{*\}, \quad F_{n+1} = \mathsf{G}F_n = [0,1]^{\mathsf{At}} \times F_n^{[0,1]}, \quad F = \bigcup_{n \geq 0} F_n.$$

We can build a model $\mathcal{F}$ on F as follows: for $(h,g) \in F_{n+1}$ we put

$$p^{\mathcal{F}}(h,g) = h(p), \quad R^{\mathcal{F}}((h,g),y) = \begin{cases} g(y) & \text{if } y \in F_n, \\ 0 & \text{otherwise.} \end{cases}$$

For $* \in F_0$, we put $p^{\mathcal{F}}(*) = R^{\mathcal{F}}(*,y) = 0$.

We can think of $\mathcal{F}$ as a canonical model capturing all finite depth behaviours.

The final chain (2)

Given a model $\mathcal{A}$, seen as a coalgebra $\alpha\colon A \to \mathsf{G}A$, define a sequence of projections $\pi_n : A \to F_n$.

$$\pi_0 =! \quad \text{and} \quad \pi_{n+1} = \mathsf{G}\pi_n \circ \alpha,$$

Explicitly,

$$\pi_{n+1}(a) = (\lambda p.p^{\mathcal{A}}(a), \lambda y. \textstyle\bigvee_{\pi_n(a')=y} R^{\mathcal{A}}(a, a')).$$

Lemma

Let $\mathcal{A}$ be a model. Then $d_n^G(a, \pi_n(a)) = 0$ for all $a \in A$.

The fuzzy van-Benthem Theorem (cont.)

With the previous lemma, the van Benthem theorem now follows from the local version, applied to the final chain model $\mathcal{F}$:

Theorem (van Benthem Theorem)

Every bisimulation invariant formula ϕ *of fuzzy FOL is modally approximable, i.e. for every* $\varepsilon > 0$*, there exists a fuzzy modal formula* ψ *such that for every fuzzy model* $\mathcal{A}$ $\|\phi_{\mathcal{A}} - \psi_{\mathcal{A}}\|_\infty \leq \varepsilon$*.*

The fuzzy van-Benthem Theorem (cont.)

With the previous lemma, the van Benthem theorem now follows from the local version, applied to the final chain model $\mathcal{F}$:

Theorem (van Benthem Theorem)

Every bisimulation invariant formula ϕ of fuzzy FOL is modally approximable, i.e. for every $\varepsilon > 0$, there exists a fuzzy modal formula ψ such that for every fuzzy model $\mathcal{A}$ $\|\phi_{\mathcal{A}} - \psi_{\mathcal{A}}\|_\infty \leq \varepsilon$.

However, this does come at a cost: while the Rosen theorem still holds in the local version, the final chain technique cannot be used to prove a uniform version of that theorem.

A probabilistic modal logic (1)

Syntax:

$$\phi, \psi ::= c \mid p \mid \phi \ominus c \mid \neg\phi \mid \phi \wedge \psi \mid \Diamond\phi$$

where $c \in \mathbb{Q} \cap [0, 1]$ and p ranges over propositional atoms.

A probabilistic modal logic (1)

Syntax:

$$\phi, \psi ::= c \mid p \mid \phi \ominus c \mid \neg\phi \mid \phi \wedge \psi \mid \Diamond\phi$$

where $c \in \mathbb{Q} \cap [0,1]$ and p ranges over propositional atoms.
Semantics: given over models

$$\mathcal{A} = (A, (p^{\mathcal{A}})_{p \in \mathsf{At}}, \pi^{\mathcal{A}})$$

- A is the set of states
- $p^{\mathcal{A}} \colon A \to [0,1]$ is the interpretation of $p \in \mathsf{At}$.
- $\pi^{\mathcal{A}} \colon A \times A \to [0,1]$ is the transition relation

such that for each $a \in A$, the map

$$\pi_a \colon A \to [0,1], \quad \pi_a(a') = \pi^{\mathcal{A}}(a, a')$$

either vanishes or is a discrete probability measure on A, i.e.

$$\textstyle\int 1 \, \mathrm{d}\pi_a = \sum_{a' \in A} \pi_a(a') \in \{0, 1\}.$$

In the first case, we call a *terminating*, in the second case *transient*.

A probabilistic modal logic (2)

Every formula ϕ has an interpretation $\phi_{\mathcal{A}} : A \to [0,1]$, which is defined inductively.
$\phi(a)$ is defined as in the fuzzy case, except for the modality:

$$(\Diamond\phi)(a) = \int \phi \, \mathrm{d}\pi_a = \sum_{a' \in A} \pi_a(a') \cdot \phi(a').$$

Probabilistic first order logic

Syntax:

$$\phi, \psi ::= c \mid p(x) \mid x = y \mid \phi \ominus c \mid \neg\phi \mid \phi \wedge \psi \mid \exists x.\, \phi \mid x\Diamond\lceil y : \phi\rceil.$$

Semantics over a model $\mathcal{A}$: Let η be a valuation mapping variables to elements of A. The clauses are as in the fuzzy case, with

$$(x\Diamond\lceil y : \phi\rceil)(\eta) = \textstyle\int \phi(\eta[y \mapsto \cdot])\, \mathrm{d}\pi_{\eta(x)} = \sum_{a \in A} \pi(\eta(x), a) \cdot \phi(\eta[y \mapsto a]).$$

Probabilistic first order logic

Syntax:

$$\phi, \psi ::= c \mid p(x) \mid x = y \mid \phi \ominus c \mid \neg\phi \mid \phi \wedge \psi \mid \exists x.\, \phi \mid x\Diamond\lceil y : \phi\rceil.$$

Semantics over a model $\mathcal{A}$: Let η be a valuation mapping variables to elements of A. The clauses are as in the fuzzy case, with

$$(x\Diamond\lceil y : \phi\rceil)(\eta) = \int \phi(\eta[y \mapsto \cdot])\, \mathrm{d}\pi_{\eta(x)} = \sum_{a \in A} \pi(\eta(x), a) \cdot \phi(\eta[y \mapsto a]).$$

Again, we have the standard translation, which now takes the form

$$\mathsf{ST}_x(\Diamond\phi) = x\Diamond\lceil y : \mathsf{ST}_y(\phi)\rceil.$$

Pairings of probability measures

Let (X, d) be a pseudometric space and π_1, π_2 be discrete probability measures on (X, d).

Denote by $\pi_1 \otimes \pi_2$ the set of probability measures μ on $X \times X$ such that π_1 and π_2 are *marginals* of μ, i.e.

- for all $x \in X$, $\sum_{y \in X} \mu(x, y) = \pi_1(x)$;
- for all $y \in X$, $\sum_{x \in X} \mu(x, y) = \pi_2(y)$.

$\pi_1 \otimes \pi_2$ is called the set of *pairings* of π_1 and π_2.

Kantorovich and Wasserstein distance

There are two ways to define a pseudometric on the space $\mathrm{D}X$ of discrete probability measures on (X, d).

Let π_1, π_2 be discrete probability measures on X.

Definition (Kantorovich distance)

$$d^{\uparrow}(\pi_1, \pi_2) = \bigvee_{f\colon (X,d) \to_1 ([0,1], d_e)} |\textstyle\int f \,\mathrm{d}\pi_1 - \int f \,\mathrm{d}\pi_2|.$$

Kantorovich and Wasserstein distance

There are two ways to define a pseudometric on the space $\mathrm{D}X$ of discrete probability measures on (X, d).

Let π_1, π_2 be discrete probability measures on X.

Definition (Kantorovich distance)

$$d^{\uparrow}(\pi_1, \pi_2) = \bigvee_{f\colon (X,d) \to_1 ([0,1], d_e)} \left| \int f \,\mathrm{d}\pi_1 - \int f \,\mathrm{d}\pi_2 \right|.$$

Definition (Wasserstein distance)

$$d^{\downarrow}(\pi_1, \pi_2) = \bigwedge_{\mu \in \pi_1 \otimes \pi_2} \int d \,\mathrm{d}\mu.$$

Kantorovich and Wasserstein distance

There are two ways to define a pseudometric on the space $\mathrm{D}X$ of discrete probability measures on (X, d).

Let π_1, π_2 be discrete probability measures on X.

Definition (Kantorovich distance)

$$d^{\uparrow}(\pi_1, \pi_2) = \bigvee_{f\colon (X,d) \to_1 ([0,1], d_e)} \left| \int f \, \mathrm{d}\pi_1 - \int f \, \mathrm{d}\pi_2 \right|.$$

Definition (Wasserstein distance)

$$d^{\downarrow}(\pi_1, \pi_2) = \bigwedge_{\mu \in \pi_1 \otimes \pi_2} \int d \, \mathrm{d}\mu.$$

Theorem (Kantorovich-Rubinstein duality)

If the space (X, d) is separable, then $d^{\uparrow} = d^{\downarrow}$.

Bisimulation game for probabilistic modal logic

The following bisimulation game naturally arises from the Wasserstein distance. The ε-*bisimulation game* is played between the spoiler S and the duplicator D on two models $\mathcal{A}, \mathcal{B}$.

Bisimulation game for probabilistic modal logic

The following bisimulation game naturally arises from the Wasserstein distance. The ε-*bisimulation game* is played between the spoiler S and the duplicator D on two models $\mathcal{A}, \mathcal{B}$.

- Configurations: triples (a, b, ε) with $a \in A$, $b \in B$ and $\varepsilon \in [0, 1]$.

Bisimulation game for probabilistic modal logic

The following bisimulation game naturally arises from the Wasserstein distance. The ε-*bisimulation game* is played between the spoiler S and the duplicator D on two models $\mathcal{A}, \mathcal{B}$.

- Configurations: triples (a, b, ε) with $a \in A$, $b \in B$ and $\varepsilon \in [0, 1]$.
- Moves:
  - D picks $\mu \in \pi_a \otimes \pi_b$ and $\varepsilon' \colon A \times B \to [0, 1]$ such that $\int \varepsilon' \, \mathrm{d}\mu \leq \varepsilon$.

Bisimulation game for probabilistic modal logic

The following bisimulation game naturally arises from the Wasserstein distance. The ε-*bisimulation game* is played between the spoiler S and the duplicator D on two models $\mathcal{A}, \mathcal{B}$.

- Configurations: triples (a, b, ε) with $a \in A$, $b \in B$ and $\varepsilon \in [0, 1]$.
- Moves:
  - D picks $\mu \in \pi_a \otimes \pi_b$ and $\varepsilon' \colon A \times B \to [0, 1]$ such that $\int \varepsilon' \, \mathrm{d}\mu \leq \varepsilon$.
  - S picks a pair (a', b') with $\mu(a', b') > 0$.
  - New configuration: $(a', b', \varepsilon'(a', b'))$.

Bisimulation game for probabilistic modal logic

The following bisimulation game naturally arises from the Wasserstein distance. The ε-*bisimulation game* is played between the spoiler S and the duplicator D on two models $\mathcal{A}, \mathcal{B}$.

- Configurations: triples (a, b, ε) with $a \in A$, $b \in B$ and $\varepsilon \in [0, 1]$.
- Moves:
  - D picks $\mu \in \pi_a \otimes \pi_b$ and $\varepsilon' \colon A \times B \to [0, 1]$ such that $\int \varepsilon' \, \mathrm{d}\mu \leq \varepsilon$.
  - S picks a pair (a', b') with $\mu(a', b') > 0$.
  - New configuration: $(a', b', \varepsilon'(a', b'))$.
- Winning conditions:
  - If either player cannot make a move, they lose.
  - D wins if a and b are both terminating.
  - S wins if exactly one of a, b is terminating.
  - S wins if $|p(a) - p(b)| > \varepsilon$ for some atom p.

Behavioural distance

We can now define several notions of behavioural distance:

Behavioural distance

We can now define several notions of behavioural distance:

- via the modal formulas up to some rank:

$$d_n^L(a,b) = \bigvee_{\phi \in \mathcal{L}_n} |\phi(a) - \phi(b)|$$

Behavioural distance

We can now define several notions of behavioural distance:

- via the modal formulas up to some rank:

$$d_n^L(a,b) = \bigvee_{\phi \in \mathcal{L}_n} |\phi(a) - \phi(b)|$$

- via the Kantorovich distance:

$$d_0^K(a,b) = 0, \quad d_{n+1}^K(a,b) = \bigvee_{p \in \mathsf{At}} |p(a) - p(b)| \vee (d_n^K)^{\uparrow}(\pi_a, \pi_b)$$

Behavioural distance

We can now define several notions of behavioural distance:

- via the modal formulas up to some rank:

$$d_n^L(a,b) = \bigvee_{\phi \in \mathcal{L}_n} |\phi(a) - \phi(b)|$$

- via the Kantorovich distance:

$$d_0^K(a,b) = 0, \quad d_{n+1}^K(a,b) = \bigvee_{p \in \mathsf{At}} |p(a) - p(b)| \vee (d_n^K)^{\uparrow}(\pi_a, \pi_b)$$

- via the Wasserstein distance:

$$d_0^W(a,b) = 0, \quad d_{n+1}^W(a,b) = \bigvee_{p \in \mathsf{At}} |p(a) - p(b)| \vee (d_n^W)^{\downarrow}(\pi_a, \pi_b)$$

Behavioural distance

We can now define several notions of behavioural distance:

- via the modal formulas up to some rank:

$$d_n^L(a,b) = \bigvee_{\phi \in \mathcal{L}_n} |\phi(a) - \phi(b)|$$

- via the Kantorovich distance:

$$d_0^K(a,b) = 0, \quad d_{n+1}^K(a,b) = \bigvee_{p \in \mathsf{At}} |p(a) - p(b)| \vee (d_n^K)^{\uparrow}(\pi_a, \pi_b)$$

- via the Wasserstein distance:

$$d_0^W(a,b) = 0, \quad d_{n+1}^W(a,b) = \bigvee_{p \in \mathsf{At}} |p(a) - p(b)| \vee (d_n^W)^{\downarrow}(\pi_a, \pi_b)$$

- via the game:

$$d_n^G(a,b) = \bigwedge \{\varepsilon \mid D \text{ wins the } n\text{-round bisim. game on } (a,b,\varepsilon)\}$$

Behavioural distance

We can now define several notions of behavioural distance:

- via the modal formulas up to some rank:

$$d_n^L(a,b) = \bigvee_{\phi \in \mathcal{L}_n} |\phi(a) - \phi(b)|$$

- via the Kantorovich distance:

$$d_0^K(a,b) = 0, \quad d_{n+1}^K(a,b) = \bigvee_{p \in \mathsf{At}} |p(a) - p(b)| \vee (d_n^K)^{\uparrow}(\pi_a, \pi_b)$$

- via the Wasserstein distance:

$$d_0^W(a,b) = 0, \quad d_{n+1}^W(a,b) = \bigvee_{p \in \mathsf{At}} |p(a) - p(b)| \vee (d_n^W)^{\downarrow}(\pi_a, \pi_b)$$

- via the game:

$$d_n^G(a,b) = \bigwedge \{\varepsilon \mid D \text{ wins the } n\text{-round bisim. game on } (a,b,\varepsilon)\}$$

Note that we need to extend $d^{\uparrow}$ and $d^{\downarrow}$ suitably for the cases where a or b is terminating.

Modal approximation

Theorem

Let $\mathcal{A}$ be a model. For all $n \geq 0$:

1. $d_n^G = d_n^W = d_n^K = d_n^L =: d_n$ *on* $\mathcal{A}$.
2. (A, d_n) *is totally bounded.*
3. $\mathcal{L}_n$ *is a dense subset of* $(A, d_n) \to_1 ([0,1], d_e)$.

Note that we use the Kantorovich-Rubinstein duality to show $d_n^W = d_n^K$. The duality holds because every totally bounded state is also separable.

Bisimulation invariance

Consider a model $\mathcal{A}$ together with the pseudometric d^G arising from the unbounded bisimulation game:

$$d^G(a,b) = \bigwedge\{\varepsilon \mid D \text{ wins the unbounded bisim. game for } (a,b,\varepsilon)\}.$$

Bisimulation invariance

Consider a model $\mathcal{A}$ together with the pseudometric d^G arising from the unbounded bisimulation game:

$$d^G(a,b) = \bigwedge\{\varepsilon \mid D \text{ wins the unbounded bisim. game for } (a,b,\varepsilon)\}.$$

A formula ϕ (modal or first order) is *bisimulation invariant* if its interpretation $\phi_{\mathcal{A}}$ is nonexpansive with respect to d^G.

Lemma

Every probabilistic modal formula is bisimulation invariant.

Bisimulation invariance

Consider a model $\mathcal{A}$ together with the pseudometric d^G arising from the unbounded bisimulation game:

$$d^G(a,b) = \bigwedge\{\varepsilon \mid D \text{ wins the unbounded bisim. game for } (a,b,\varepsilon)\}.$$

A formula ϕ (modal or first order) is *bisimulation invariant* if its interpretation $\phi_{\mathcal{A}}$ is nonexpansive with respect to d^G.

Lemma

Every probabilistic modal formula is bisimulation invariant.

Lemma

Every bisimulation invariant formula is also depth n bisimulation invariant for some n, i.e. nonexpansive wrt. depth n behavioural distance d^G_n.

Locality

For a model $\mathcal{A}$, the *Gaifman graph* is the undirected graph with vertex set $V = A$ and edge set $E = \{\{x, y\} \mid \pi^{\mathcal{A}}(x, y) > 0\}$.

The *Gaifman distance* $D : V \times V \to \mathbb{N} \cup \{\infty\}$ is the canonical distance function for that graph, i.e. $D(a, b)$ is the length of the shortest path between a and b in the Gaifman graph.

For $\ell \geq 0$, the *radius ℓ neighbourhood* of $x \in X$ is

$$U^{\ell}_{\mathcal{A}}(x) = \{y \in X \mid D(x, y) \leq \ell\}.$$

A formula $\phi = \phi(z)$ is ℓ*-local* if for all models $\mathcal{A}$ and $x \in X$,

$$\phi_{\mathcal{A}}(a) = \phi_{\mathcal{A}^{\ell}_{a}}(a),$$

where $\mathcal{A}^{\ell}_{a}$ is the model arising by restricting A, $p^{\mathcal{A}}$ and $\pi^{\mathcal{A}}$ to $U^{\ell}(a)$, making states at distance ℓ from a terminating.

Finite depth bisimulation invariance

We use locality to prove that bisimulation invariant formulas are also bisimulation invariant at some finite depth n.

Lemma

Let $\phi = \phi(x)$ be a probabilistic first order formula of rank k and let $n = 3^k$. Then

1. *ϕ is n-local;*
2. *ϕ is bisimulation invariant at depth n.*

As was already the case in my last talk, these proofs build on those presented in [Otto03] for the classical case.

The probabilistic van-Benthem Theorem

Finally, we obtain the following:

Theorem (van Benthem Theorem)

Every bisimulation invariant formula ϕ *of probabilistic FOL is modally approximable, i.e. for every* $\varepsilon > 0$*, there exists a probabilistic modal formula* ψ *such that for every probabilistic model* $\mathcal{A}$*,* $\|\phi_{\mathcal{A}} - \psi_{\mathcal{A}}\|_\infty \leq \varepsilon$*.*

Future work

Both the Kantorovich and Wasserstein distances generalize to other set functors [BaldanEA14], so it is natural to ask if the same is true for other parts of this work. Possible future work includes:

- Formalizing a coalgebraic version of the bisimulation game, that ideally instantiates to both the fuzzy and the probabilistic games presented earlier.
- Finding a coalgebraic generalization of these van Benthem Theorems. Which conditions on the functor are needed?
- Proving or disproving the Rosen version of the theorem.
- Finding out whether the approximation can be strengthened to state that every bisimulation invariant formula of fuzzy/probabilistic FOL is in fact equivalent to a modal formula.

References

F. van Breugel and J. Worrell.
A behavioural pseudometric for probabilistic transition systems.

M. Otto.
Elementary proof of the van Benthem-Rosen characterisation theorem.

P. Baldan, F. Bonchi, H. Kerstan, B. König.
Behavioral metrics via functor lifting.

P. Wild, L. Schröder, D. Pattinson, and B. König.
A van Benthem theorem for fuzzy modal logic.

P. Wild, L. Schröder.
A Characterization Theorem for a Modal Description Logic.