Approximate or non-continuous satisfaction of identities

W. Taylor

UH, Manoa Valley, Hawai‘i

March 3, 2012

Six sections of the talk.

Basics

Some calculations on approximate satisfaction. $\lambda(A, \Sigma)$.

Precise definitions for discontinuous satisfaction. $\mu_n(A, \Sigma)$.

An example: $\mu_1(A, \Sigma) = 0$; $\mu_2(A, \Sigma) = \mathrm{diam}(A)$.

Some further results

Algorithmic questions about λ and μ_n.

Compatibility: $A \models \Sigma$

Compatibility: $A \models \Sigma$

Given a *topological space* A and a set Σ of equations in operation symbols F_t, we write

$$A \models \Sigma,$$

and say that A and Σ are *compatible*, iff there exist *continuous* operations $\overline{F_t}$ on A satisfying Σ.

Compatibility: $A \models \Sigma$

Given a *topological space* A and a set Σ of equations in operation symbols F_t, we write

$$A \models \Sigma,$$

and say that A and Σ are *compatible*, iff there exist *continuous* operations $\overline{F_t}$ on A satisfying Σ.

Examples: Groups on S^1, S^3 and $\mathbb{R}$,

Compatibility: $A \models \Sigma$

Given a *topological space* A and a set Σ of equations in operation symbols F_t, we write

$$A \models \Sigma,$$

and say that A and Σ are *compatible*, iff there exist *continuous* operations $\overline{F_t}$ on A satisfying Σ.

Examples: Groups on S^1, S^3 and $\mathbb{R}$, various matrix groups,

Compatibility: $A \models \Sigma$

Given a *topological space* A and a set Σ of equations in operation symbols F_t, we write

$$A \models \Sigma,$$

and say that A and Σ are *compatible*, iff there exist *continuous* operations $\overline{F_t}$ on A satisfying Σ.

Examples: Groups on S^1, S^3 and $\mathbb{R}$, various matrix groups, many H-spaces,

Compatibility: $A \models \Sigma$

Given a *topological space* A and a set Σ of equations in operation symbols F_t, we write

$$A \models \Sigma,$$

and say that A and Σ are *compatible*, iff there exist *continuous* operations $\overline{F_t}$ on A satisfying Σ.

Examples: Groups on S^1, S^3 and $\mathbb{R}$, various matrix groups, many H-spaces, a lattice on $[0, 1]$,

Compatibility: $A \models \Sigma$

Given a *topological space* A and a set Σ of equations in operation symbols F_t, we write

$$A \models \Sigma,$$

and say that A and Σ are *compatible*, iff there exist *continuous* operations $\overline{F_t}$ on A satisfying Σ.

Examples: Groups on S^1, S^3 and $\mathbb{R}$, various matrix groups, many H-spaces, a lattice on $[0, 1]$, a ternary median operation on Y,

Compatibility: $A \models \Sigma$

Given a *topological space* A and a set Σ of equations in operation symbols F_t, we write

$$A \models \Sigma,$$

and say that A and Σ are *compatible*, iff there exist *continuous* operations $\overline{F_t}$ on A satisfying Σ.

Examples: Groups on S^1, S^3 and $\mathbb{R}$, various matrix groups, many H-spaces, a lattice on $[0, 1]$, a ternary median operation on Y, simple Σ on absolute-retract A,

Compatibility: $A \models \Sigma$

Given a *topological space* A and a set Σ of equations in operation symbols F_t, we write

$$A \models \Sigma,$$

and say that A and Σ are *compatible*, iff there exist *continuous* operations $\overline{F_t}$ on A satisfying Σ.

Examples: Groups on S^1, S^3 and $\mathbb{R}$, various matrix groups, many H-spaces, a lattice on $[0, 1]$, a ternary median operation on Y, simple Σ on absolute-retract A, $\mathrm{Sets}^{[n]}$ on any space A^n,

Compatibility: $A \models \Sigma$

Given a *topological space* A and a set Σ of equations in operation symbols F_t, we write

$$A \models \Sigma,$$

and say that A and Σ are *compatible*, iff there exist *continuous* operations $\overline{F_t}$ on A satisfying Σ.

Examples: Groups on S^1, S^3 and $\mathbb{R}$, various matrix groups, many H-spaces, a lattice on $[0,1]$, a ternary median operation on Y, simple Σ on absolute-retract A, $\mathrm{Sets}^{[n]}$ on any space A^n, a unital ring on $S^1 \times \mathbb{Z}$,

Compatibility: $A \models \Sigma$

Given a *topological space* A and a set Σ of equations in operation symbols F_t, we write

$$A \models \Sigma,$$

and say that A and Σ are *compatible*, iff there exist *continuous* operations $\overline{F_t}$ on A satisfying Σ.

Examples: Groups on S^1, S^3 and $\mathbb{R}$, various matrix groups, many H-spaces, a lattice on $[0,1]$, a ternary median operation on Y, simple Σ on absolute-retract A, $\mathrm{Sets}^{[n]}$ on any space A^n, a unital ring on $S^1 \times \mathbb{Z}$, a Boolean algebra on $\{0,1\}^{\aleph_0}$.

Compatibility: $A \models \Sigma$

Given a *topological space* A and a set Σ of equations in operation symbols F_t, we write

$$A \models \Sigma,$$

and say that A and Σ are *compatible*, iff there exist *continuous* operations $\overline{F_t}$ on A satisfying Σ.

Examples: Groups on S^1, S^3 and $\mathbb{R}$, various matrix groups, many H-spaces, a lattice on $[0,1]$, a ternary median operation on Y, simple Σ on absolute-retract A, $\mathrm{Sets}^{[n]}$ on any space A^n, a unital ring on $S^1 \times \mathbb{Z}$, a Boolean algebra on $\{0,1\}^{\aleph_0}$. Even, for any A and Σ, the Świerczkowski free algebra $\mathbf{F}_A(\Sigma)$ (whose universe is a superset of A).

$A \models \Sigma$ is a mysterious, sparse relation

$A \models \Sigma$ is a mysterious, sparse relation

- Experientially, $A \models \Sigma$ occurs only sporadically, whereas it can in many cases be proved false.

$A \models \Sigma$ is a mysterious, sparse relation

- Experientially, $A \models \Sigma$ occurs only sporadically, whereas it can in many cases be proved false.
- But $A \not\models \Sigma$ does not seem to have any uniform method of proof.

$A \models \Sigma$ is a mysterious, sparse relation

- Experientially, $A \models \Sigma$ occurs only sporadically, whereas it can in many cases be proved false.
- But $A \not\models \Sigma$ does not seem to have any uniform method of proof.
- E.g. the sphere $S^n \models \Sigma$ only for trivial Σ or for $n = 1, 3, 7$. (Hard algebraic topology to prove this.)

$A \models \Sigma$ is a mysterious, sparse relation

- Experientially, $A \models \Sigma$ occurs only sporadically, whereas it can in many cases be proved false.
- But $A \not\models \Sigma$ does not seem to have any uniform method of proof.
- E.g. the sphere $S^n \models \Sigma$ only for trivial Σ or for $n = 1, 3, 7$. (Hard algebraic topology to prove this.)
- There is no algorithm that settles $\mathbb{R} \models \Sigma$ for finite Σ. (Uses Matiasevich solution of Hilbert's Tenth Problem.) Thus $\models$ is not *too* sparse.

Relaxing the demands of $A \models \Sigma$.

Relaxing the demands of $A \models \Sigma$.

"$A \models \Sigma$" demands operations $\overline{F}_t$ that **satisfy** Σ **exactly** and are **continuous**.

Relaxing the demands of $A \models \Sigma$.

"$A \models \Sigma$" demands operations $\overline{F}_t$ that **satisfy** Σ **exactly** and are **continuous**. We can relax those demands in two ways:

Relaxing the demands of $A \models \Sigma$.

"$A \models \Sigma$" demands operations $\overline{F}_t$ that **satisfy** Σ **exactly** and are **continuous**. We can relax those demands in two ways: we can consider approximate satisfaction, and

Relaxing the demands of $A \models \Sigma$.

"$A \models \Sigma$" demands operations $\overline{F}_t$ that **satisfy** Σ **exactly** and are **continuous**. We can relax those demands in two ways: we can consider approximate satisfaction, and we can consider approximate continuity.

Approximate replacements for $A \models \Sigma$

Approximate replacements for $A \models \Sigma$

For (A, d) a metric space, and $\eta > 0$

$$A \models_\eta \Sigma,$$

will mean that there exist *continuous* operations $\overline{F_t}$ on A satisfying Σ **within** η.

Approximate replacements for $A \models \Sigma$

For (A, d) a metric space, and $\eta > 0$

$$A \models_\eta \Sigma,$$

will mean that there exist *continuous* operations $\overline{F_t}$ on A satisfying Σ **within** η. For $\varepsilon > 0$,

$$A \models^\varepsilon \Sigma,$$

will mean that there exist operations $\overline{F_t}$ on A satisfying Σ and **whose discontinuities are no greater than** ε. (Precise definition later.)

Approximate replacements for $A \models \Sigma$

For (A, d) a metric space, and $\eta > 0$

$$A \models_\eta \Sigma,$$

will mean that there exist *continuous* operations $\overline{F_t}$ on A satisfying Σ **within** η. For $\varepsilon > 0$,

$$A \models^\varepsilon \Sigma,$$

will mean that there exist operations $\overline{F_t}$ on A satisfying Σ and **whose discontinuities are no greater than** ε. (Precise definition later.)

(Of course we may also study $A \models_\eta^\varepsilon \Sigma$!)

Six sections of the talk.

Basics

Some calculations on approximate satisfaction. $\lambda(A, \Sigma)$.

Precise definitions for discontinuous satisfaction. $\mu_n(A, \Sigma)$.

An example: $\mu_1(A, \Sigma) = 0$; $\mu_2(A, \Sigma) = \mathrm{diam}(A)$.

Some further results

Algorithmic questions about λ and μ_n.

Some review of $A \models_\eta \Sigma$

Some review of $A \models_\eta \Sigma$

Our long paper, "Approximate satisfaction of identities," deals with $\models_\eta$, and we have spoken on it before, but we will sketch one result for later comparison with $\models^\varepsilon$.

Some review of $A \models_\eta \Sigma$

Our long paper, "Approximate satisfaction of identities," deals with $\models_\eta$, and we have spoken on it before, but we will sketch one result for later comparison with $\models^\varepsilon$.

For arbitrary metric space A and equations Σ, we define the real number (or $+\infty$):

$$\lambda(A, \Sigma) \;=\; \inf\{\eta \,:\, A \models_\eta \Sigma\}.$$

Σ saying, "$\overline{G}$ is a one-one binary operation on A."

Σ saying, "$\overline{G}$ is a one-one binary operation on A."

Σ will be this pair of equations:

$$F_0(G(x_0, x_1)) \approx x_0, \qquad F_1(G(x_0, x_1)) \approx x_1.$$

Σ saying, "$\overline{G}$ is a one-one binary operation on A."

Σ will be this pair of equations:

$$F_0(G(x_0, x_1)) \approx x_0, \qquad F_1(G(x_0, x_1)) \approx x_1.$$

A model of Σ has

$$A^2 \xrightarrow{\overline{G}} A \xrightarrow{\overline{F}} A^2 = \text{identity},$$

Σ saying, "$\overline{G}$ is a one-one binary operation on A."

Σ will be this pair of equations:

$$F_0(G(x_0, x_1)) \approx x_0, \qquad F_1(G(x_0, x_1)) \approx x_1.$$

A model of Σ has

$$A^2 \xrightarrow{\overline{G}} A \xrightarrow{\overline{F}} A^2 = \text{identity},$$

where $\overline{F}$ has $\overline{F}_0$ and $\overline{F}_1$ as its component functions.

Σ saying, "$\overline{G}$ is a one-one binary operation on A."

Σ will be this pair of equations:

$$F_0(G(x_0, x_1)) \approx x_0, \qquad F_1(G(x_0, x_1)) \approx x_1.$$

A model of Σ has

$$A^2 \xrightarrow{\overline{G}} A \xrightarrow{\overline{F}} A^2 = \text{identity},$$

where $\overline{F}$ has $\overline{F}_0$ and $\overline{F}_1$ as its component functions. Thus under exact satisfaction, $\overline{G}$ must be one-one and $\overline{F}$ must be onto.

$\lambda([0,1],\Sigma) \;=\; 0.5$

$\lambda([0,1],\Sigma) \;=\; 0.5$

First, an easy exercise yields $\lambda([0,1],\Sigma) \leq 0.5$.

For the opposite inequality, we consider a topolgical algebra $(A; \overline{G}, \overline{F}_0, \overline{F}_1)$ modeling Σ within K; we will show that $K \geq 0.5$.

$\lambda([0,1],\Sigma) \;=\; 0.5$

First, an easy exercise yields $\lambda([0,1],\Sigma) \leq 0.5$.

For the opposite inequality, we consider a topolgical algebra $(A; \overline{G}, \overline{F}_0, \overline{F}_1)$ modeling Σ within K; we will show that $K \geq 0.5$.

Let $\{a_0, a_1\} = \{b_0, b_1\} = \{0, 1\}$. Compare the four values $\overline{G}(a_i, b_j)$; w.l.o.g. $\overline{G}(a_0, b_0)$ is the smallest.

$\lambda([0,1],\Sigma) \;=\; 0.5$

First, an easy exercise yields $\lambda([0,1],\Sigma) \leq 0.5$.

For the opposite inequality, we consider a topolgical algebra $(A; \overline{G}, \overline{F}_0, \overline{F}_1)$ modeling Σ within K; we will show that $K \geq 0.5$.

Let $\{a_0, a_1\} = \{b_0, b_1\} = \{0, 1\}$. Compare the four values $\overline{G}(a_i, b_j)$; w.l.o.g. $\overline{G}(a_0, b_0)$ is the smallest. Again w.l.o.g. we have $\overline{G}(a_1, b_0) \leq \overline{G}(a_0, b_1)$. In other words,

$$\overline{G}(a_0, b_0) \;\leq\; \overline{G}(a_1, b_0) \;\leq\; \overline{G}(a_0, b_1).$$

$\lambda([0,1],\Sigma) \;=\; 0.5$

First, an easy exercise yields $\lambda([0,1],\Sigma) \leq 0.5$.

For the opposite inequality, we consider a topolgical algebra $(A; \overline{G}, \overline{F}_0, \overline{F}_1)$ modeling Σ within K; we will show that $K \geq 0.5$.

Let $\{a_0, a_1\} = \{b_0, b_1\} = \{0, 1\}$. Compare the four values $\overline{G}(a_i, b_j)$; w.l.o.g. $\overline{G}(a_0, b_0)$ is the smallest. Again w.l.o.g. we have $\overline{G}(a_1, b_0) \leq \overline{G}(a_0, b_1)$. In other words,

$$\overline{G}(a_0, b_0) \;\leq\; \overline{G}(a_1, b_0) \;\leq\; \overline{G}(a_0, b_1).$$

We consider the real function $\overline{H}(x) \;=\; \overline{G}(a_0, x)$.

$\lambda([0,1],\Sigma) \;=\; 0.5$

First, an easy exercise yields $\lambda([0,1],\Sigma) \leq 0.5$.

For the opposite inequality, we consider a topolgical algebra $(A; \overline{G}, \overline{F}_0, \overline{F}_1)$ modeling Σ within K; we will show that $K \geq 0.5$.

Let $\{a_0, a_1\} = \{b_0, b_1\} = \{0, 1\}$. Compare the four values $\overline{G}(a_i, b_j)$; w.l.o.g. $\overline{G}(a_0, b_0)$ is the smallest. Again w.l.o.g. we have $\overline{G}(a_1, b_0) \leq \overline{G}(a_0, b_1)$. In other words,

$$\overline{G}(a_0, b_0) \;\leq\; \overline{G}(a_1, b_0) \;\leq\; \overline{G}(a_0, b_1).$$

We consider the real function $\overline{H}(x) \;=\; \overline{G}(a_0, x)$. By the IVT there exists e with

$$\overline{G}(a_0, e) \;=\; \overline{G}(a_1, b_0).$$

$\lambda([0,1],\Sigma) \;=\; 0.5$, continued

To repeat, we now have e with

$$\overline{G}(a_0, e) \;=\; \overline{G}(a_1, b_0).$$

From the approximate validity of Σ, we now calculate

$$a_0 \;\approx\; \overline{F}_0(\overline{G}(a_0, e)) \;=\; \overline{F}_0(\overline{G}(a_1, b_0)) \;\approx\; a_1.$$

$\lambda([0,1],\Sigma) \;=\; 0.5$, continued

To repeat, we now have e with

$$\overline{G}(a_0, e) \;=\; \overline{G}(a_1, b_0).$$

From the approximate validity of Σ, we now calculate

$$a_0 \;\approx\; \overline{F}_0(\overline{G}(a_0, e)) \;=\; \overline{F}_0(\overline{G}(a_1, b_0)) \;\approx\; a_1.$$

Since $\{a_0, a_1\} = \{0, 1\}$, we now have the desired conclusion that

$$1 \;=\; d(a_0, a_1) \;\leq\; 2K.$$

$\lambda([0,1]^2, \Sigma) = 0.5$; not a topological invariant

$\lambda([0,1]^2, \Sigma) = 0.5$; not a topological invariant

$\lambda([0,1]^2, \Sigma) = 0.5$, not much more difficult (omitted).

$\lambda([0,1]^2, \Sigma) = 0.5$; not a topological invariant

$\lambda([0,1]^2, \Sigma) = 0.5$, not much more difficult (omitted).

Given real $\eta > 0$, we construct a unit-diameter metric d for $[0,1]^2$ such that $\lambda(([0,1]^2, d), \Sigma) \leq \eta$.

$\lambda([0,1]^2, \Sigma) = 0.5$; not a topological invariant

$\lambda([0,1]^2, \Sigma) = 0.5$, not much more difficult (omitted).

Given real $\eta > 0$, we construct a unit-diameter metric d for $[0,1]^2$ such that $\lambda(([0,1]^2, d), \Sigma) \leq \eta$.

We replace $[0,1]^2$ by $B = [0, \sqrt{1-\eta^2}] \times [0, \eta]$, while taking d to be the Euclidean metric. Clearly $B \cong [0,1]^2$, and moreover B has unit diameter.

$\lambda([0,1]^2, \Sigma) = 0.5$; not a topological invariant

$\lambda([0,1]^2, \Sigma) = 0.5$, not much more difficult (omitted).

Given real $\eta > 0$, we construct a unit-diameter metric d for $[0,1]^2$ such that $\lambda(([0,1]^2, d), \Sigma) \leq \eta$.

We replace $[0,1]^2$ by $B = [0, \sqrt{1-\eta^2}] \times [0, \eta]$, while taking d to be the Euclidean metric. Clearly $B \cong [0,1]^2$, and moreover B has unit diameter. For an upper estimate on $\lambda((B,d), \Sigma)$, we define these three operations on B:

$$\begin{aligned} \overline{G}((a_0, a_1), (b_0, b_1)) &= (a_0, b_0/K) \\ \overline{F}_0(a_0, a_1) &= (a_0, 0) \\ \overline{F}_1(a_0, a_1) &= (Ka_1, 0). \end{aligned}$$

where $K = \sqrt{1-\eta^2}/\eta$.

$\lambda([0,1]^2, \Sigma)$ not invariant, continued

We now calculate

$$\begin{aligned} \overline{F}_0(\overline{G}((a_0,a_1),(b_0,b_1))) &= (a_0,0) \\ d((a_0,a_1),\overline{F}_0(\overline{G}((a_0,a_1),(b_0,b_1)))) &= d((a_0,a_1),(a_0,0)) \\ &\leq \eta. \end{aligned}$$

Thus $\overline{F}_0$, $\overline{F}_1$ and $\overline{G}$ satisfy $\overline{F}_0(\overline{G}(x_0,x_1)) \approx x_0$ within η on (B,d).

$\lambda([0,1]^2, \Sigma)$ not invariant, continued

We now calculate

$$\begin{aligned} \overline{F}_0(\overline{G}((a_0,a_1),(b_0,b_1))) &= (a_0,0) \\ d((a_0,a_1),\overline{F}_0(\overline{G}((a_0,a_1),(b_0,b_1)))) &= d((a_0,a_1),(a_0,0)) \\ &\leq \eta. \end{aligned}$$

Thus $\overline{F}_0$, $\overline{F}_1$ and $\overline{G}$ satisfy $\overline{F}_0(\overline{G}(x_0,x_1)) \approx x_0$ within η on (B,d). The approximate satisfaction of $\overline{F}_1(\overline{G}(x_0,x_1)) \approx x_1$ is handled similarly. Thus

$$\lambda((B,d),\Sigma) \leq \eta.$$

Six sections of the talk.

Basics

Some calculations on approximate satisfaction. $\lambda(A, \Sigma)$.

Precise definitions for discontinuous satisfaction. $\mu_n(A, \Sigma)$.

An example: $\mu_1(A, \Sigma) = 0$; $\mu_2(A, \Sigma) = \mathrm{diam}(A)$.

Some further results

Algorithmic questions about λ and μ_n.

(δ, ε)-constraints.

(δ, ε)-constraints.

Let (A, d) and (B, e) be metric spaces.

(δ, ε)-constraints.

Let (A, d) and (B, e) be metric spaces.

Given $\overline{F} : B \longrightarrow A$ (not necessarily continuous) and $\delta, \varepsilon > 0$, we say that $\overline{F}$ is (δ, ε)-**constrained** it satisfies: for all $b, b' \in B$, if $e(b, b') < \delta$, then $d(\overline{F}(b), \overline{F}(b')) < \varepsilon$.

(δ, ε)-constraints.

Let (A, d) and (B, e) be metric spaces.

Given $\overline{F} : B \longrightarrow A$ (not necessarily continuous) and $\delta, \varepsilon > 0$, we say that $\overline{F}$ is (δ, ε)-**constrained** it satisfies: for all $b, b' \in B$, if $e(b, b') < \delta$, then $d(\overline{F}(b), \overline{F}(b')) < \varepsilon$.

(Uniform continuity rephrased: for every $\varepsilon > 0$ there exists $\delta > 0$ so that $\overline{F}$ is constrained by (δ, ε).)

(δ, ε)-constraints.

Let (A, d) and (B, e) be metric spaces.

Given $\overline{F} : B \longrightarrow A$ (not necessarily continuous) and $\delta, \varepsilon > 0$, we say that $\overline{F}$ is (δ, ε)-**constrained** it satisfies: for all $b, b' \in B$, if $e(b, b') < \delta$, then $d(\overline{F}(b), \overline{F}(b')) < \varepsilon$.

(Uniform continuity rephrased: for every $\varepsilon > 0$ there exists $\delta > 0$ so that $\overline{F}$ is constrained by (δ, ε).)

We say that $\overline{F}$ is n-**constrained by** (δ_0, δ_n) iff there exist $0 < \delta_0 \leq \delta_1 \leq \cdots \leq \delta_n$ such that $\overline{F}$ is (δ_0, δ_1)-constrained and (δ_1, δ_2)-constrained, and so on, up to (δ_{n-1}, δ_n)-constrained.

The intermediate-value theorem, revisited.

Lemma

Suppose that f maps a convex subset of $\mathbb{R}$ into $\mathbb{R}$, and that f is (δ, ε)-constrained with $\delta, \varepsilon > 0$. If $a < c$ and s is between $f(a)$ and $f(c)$, then there exists b with $a \leq b \leq c$ and with $d(f(b), s) < \varepsilon/2$.

Definitions of $\models_n^\varepsilon$ and $\mu_n(A, \Sigma)$

$$A \models_n^\varepsilon \Sigma$$

Definitions of $\models_n^\varepsilon$ and $\mu_n(A, \Sigma)$

$$A \models_n^\varepsilon \Sigma$$

means that *there exists an algebra* $\mathbf{A} = (A, \overline{F}_t)_{t \in T}$ *modeling* Σ *and a real number* $\delta_0 > 0$ *such that each* $\overline{F}_t$ *is n-constrained by* (δ_0, ε).

Definitions of $\models_n^\varepsilon$ and $\mu_n(A, \Sigma)$

$$A \models_n^\varepsilon \Sigma$$

means that *there exists an algebra* $\mathbf{A} = (A, \overline{F}_t)_{t \in T}$ *modeling* Σ *and a real number* $\delta_0 > 0$ *such that each* $\overline{F}_t$ *is n-constrained by* (δ_0, ε).

We define

$$\mu_n(A, \Sigma) \;=\; \inf \{\varepsilon : A \models_n^\varepsilon \Sigma\}.$$

Definitions of $\models_n^\varepsilon$ and $\mu_n(A, \Sigma)$

$$A \models_n^\varepsilon \Sigma$$

means that *there exists an algebra* $\mathbf{A} = (A, \overline{F}_t)_{t \in T}$ *modeling* Σ *and a real number* $\delta_0 > 0$ *such that each* $\overline{F}_t$ *is n-constrained by* (δ_0, ε).

We define

$$\mu_n(A, \Sigma) \;=\; \inf \{\varepsilon : A \models_n^\varepsilon \Sigma\}.$$

It is not hard to see that

$$0 \;\leq\; \mu_1(A, \Sigma) \;\leq\; \mu_2(A, \Sigma) \;\leq\; \cdots \;\leq\; \mathrm{diam}(A).$$

Connection of $\models_n^\varepsilon$ and $\mu_n(A, \Sigma)$ with $A \models \Sigma$

We repeat the definition:

$$A \models_n^\varepsilon \Sigma$$

means that *there exists an algebra* $\mathbf{A} = (A, \overline{F}_t)_{t \in T}$ *modeling* Σ *and a real number* $\delta_0 > 0$ *such that each* $\overline{F}_t$ *is n-constrained by* (δ_0, ε).

Connection of $\models_n^\varepsilon$ and $\mu_n(A, \Sigma)$ with $A \models \Sigma$

We repeat the definition:

$$A \models_n^\varepsilon \Sigma$$

means that *there exists an algebra* $\mathbf{A} = (A, \overline{F}_t)_{t \in T}$ *modeling* Σ *and a real number* $\delta_0 > 0$ *such that each* $\overline{F}_t$ *is n-constrained by* (δ_0, ε).

Thus if $A \models \Sigma$, then $A \models_n^\varepsilon \Sigma$ for every n and every $\varepsilon > 0$,

Connection of $\models_n^\varepsilon$ and $\mu_n(A, \Sigma)$ with $A \models \Sigma$

We repeat the definition:

$$A \models_n^\varepsilon \Sigma$$

means that *there exists an algebra* $\mathbf{A} = (A, \overline{F}_t)_{t \in T}$ *modeling* Σ *and a real number* $\delta_0 > 0$ *such that each* $\overline{F}_t$ *is n-constrained by* (δ_0, ε).

Thus if $A \models \Sigma$, then $A \models_n^\varepsilon \Sigma$ for every n and every $\varepsilon > 0$, and hence

$$\mu_n(A, \Sigma) = \inf \{\varepsilon : A \models_n^\varepsilon \Sigma\} = 0$$

for every n.

Six sections of the talk.

Basics

Some calculations on approximate satisfaction. $\lambda(A, \Sigma)$.

Precise definitions for discontinuous satisfaction. $\mu_n(A, \Sigma)$.

An example: $\mu_1(A, \Sigma) = 0$; $\mu_2(A, \Sigma) = \mathrm{diam}(A)$.

Some further results

Algorithmic questions about λ and μ_n.

$\mu_1([0,1],\Sigma) = 0$ for $\Sigma =$ injective binary

$\mu_1([0,1],\Sigma) = 0$ for $\Sigma =$ injective binary

Our first Σ is taken as before: succinctly, it says that $A^2 \xrightarrow{\overline{G}} A \xrightarrow{\overline{F}} A^2 \;=\;$ identity.

$\mu_1([0,1],\Sigma) = 0$ for Σ = injective binary

Our first Σ is taken as before: succinctly, it says that
$A^2 \xrightarrow{\overline{G}} A \xrightarrow{\overline{F}} A^2 \;=\;$ identity.

So let $\overline{F}$ be a Peano curve: continuous from $[0,1]$ **onto** $[0,1]^2$, and let $\overline{G}$ be any left-inverse to $\overline{F}$.

$\mu_1([0,1],\Sigma) = 0$ for Σ = injective binary

Our first Σ is taken as before: succinctly, it says that $A^2 \xrightarrow{\overline{G}} A \xrightarrow{\overline{F}} A^2 \;=\;$ identity.

So let $\overline{F}$ be a Peano curve: continuous from $[0,1]$ **onto** $[0,1]^2$, and let $\overline{G}$ be any left-inverse to $\overline{F}$. ($\overline{G}$ is perforce discontinuous.) For arbitrary $\varepsilon > 0$, define functions F' and G' via

$$G'(a_0, b_0) \;=\; \varepsilon\, \overline{G}(a_0, b_0); \quad F'(a) \;=\; \overline{F}(1 \wedge (a/\varepsilon)).$$

$\mu_1([0,1],\Sigma) = 0$ for Σ = injective binary

Our first Σ is taken as before: succinctly, it says that $A^2 \xrightarrow{\overline{G}} A \xrightarrow{\overline{F}} A^2 \;=\;$ identity.

So let $\overline{F}$ be a Peano curve: continuous from $[0,1]$ **onto** $[0,1]^2$, and let $\overline{G}$ be any left-inverse to $\overline{F}$. ($\overline{G}$ is perforce discontinuous.) For arbitrary $\varepsilon > 0$, define functions F' and G' via

$$G'(a_0, b_0) \;=\; \varepsilon\,\overline{G}(a_0, b_0); \quad F'(a) \;=\; \overline{F}(1 \wedge (a/\varepsilon)).$$

Now the discontinuities of G' are no larger than ε, and F' remains continuous, while F' and G' still satisfy Σ. Thus $A \models_1^{\varepsilon}$ for every $\varepsilon > 0$; hence $\mu_1([0,1],\Sigma) = 0$. ∎

$\mu_2([0,1],\Sigma) = 1$

$\mu_2([0,1],\Sigma) = 1$

We consider $(A; \overline{G}, \overline{F}_0, \overline{F}_1)$ modeling Σ, with the operations (δ_0, δ_1)-constrained and (δ_1, δ_2)-constrained. We will show that $\delta_2 \geq 1$.

$\mu_2([0,1],\Sigma) = 1$

We consider $(A; \overline{G}, \overline{F}_0, \overline{F}_1)$ modeling Σ, with the operations (δ_0, δ_1)-constrained and (δ_1, δ_2)-constrained. We will show that $\delta_2 \geq 1$.

Let $\{a_0, a_1\} = \{b_0, b_1\} = \{0, 1\}$. Compare the four values $\overline{G}(a_i, b_j)$; w.l.o.g. $\overline{G}(a_0, b_0)$ is the smallest.

$\mu_2([0,1],\Sigma) = 1$

We consider $(A; \overline{G}, \overline{F}_0, \overline{F}_1)$ modeling Σ, with the operations (δ_0, δ_1)-constrained and (δ_1, δ_2)-constrained. We will show that $\delta_2 \geq 1$.

Let $\{a_0, a_1\} = \{b_0, b_1\} = \{0, 1\}$. Compare the four values $\overline{G}(a_i, b_j)$; w.l.o.g. $\overline{G}(a_0, b_0)$ is the smallest. Again w.l.o.g. we have $\overline{G}(a_1, b_0) \leq \overline{G}(a_0, b_1)$. In other words,

$$\overline{G}(a_0, b_0) \;\leq\; \overline{G}(a_1, b_0) \;\leq\; \overline{G}(a_0, b_1).$$

$\mu_2([0,1],\Sigma) = 1$

We consider $(A; \overline{G}, \overline{F}_0, \overline{F}_1)$ modeling Σ, with the operations (δ_0, δ_1)-constrained and (δ_1, δ_2)-constrained. We will show that $\delta_2 \geq 1$.

Let $\{a_0, a_1\} = \{b_0, b_1\} = \{0, 1\}$. Compare the four values $\overline{G}(a_i, b_j)$; w.l.o.g. $\overline{G}(a_0, b_0)$ is the smallest. Again w.l.o.g. we have $\overline{G}(a_1, b_0) \leq \overline{G}(a_0, b_1)$. In other words,

$$\overline{G}(a_0, b_0) \;\leq\; \overline{G}(a_1, b_0) \;\leq\; \overline{G}(a_0, b_1).$$

We consider the real function $\overline{H}(x) = \overline{G}(a_0, x)$.

$\mu_2([0,1],\Sigma) = 1$

We consider $(A; \overline{G}, \overline{F}_0, \overline{F}_1)$ modeling Σ, with the operations (δ_0, δ_1)-constrained and (δ_1, δ_2)-constrained. We will show that $\delta_2 \geq 1$.

Let $\{a_0, a_1\} = \{b_0, b_1\} = \{0, 1\}$. Compare the four values $\overline{G}(a_i, b_j)$; w.l.o.g. $\overline{G}(a_0, b_0)$ is the smallest. Again w.l.o.g. we have $\overline{G}(a_1, b_0) \leq \overline{G}(a_0, b_1)$. In other words,

$$\overline{G}(a_0, b_0) \;\leq\; \overline{G}(a_1, b_0) \;\leq\; \overline{G}(a_0, b_1).$$

We consider the real function $\overline{H}(x) = \overline{G}(a_0, x)$. By the Lemma (IVT), there exists $e \in [0, 1]$ with

$$d\big(\overline{G}(a_0, e), \overline{G}(a_1, b_0)\big)\big) \;<\; \delta_1.$$

$\mu_2([0,1],\Sigma) = 1$, concluded

Repeat: $\quad d\big(\overline{G}(a_0, e), \overline{G}(a_1, b_0)\big)\big) \;<\; \delta_1.$

$\mu_2([0,1],\Sigma) = 1$, concluded

Repeat: $\quad d\big(\overline{G}(a_0,e), \overline{G}(a_1,b_0)\big)\big) \;<\; \delta_1.$

Because $\{a_0,a_1\} = \{0,1\}$, because of Σ, and because the function $\overline{F}_0$ is (δ_1,δ_2)-constrained, we now have:

$$1 \;=\; d(a_0,a_1) \;=\; d\big(\overline{F}_0(\overline{G}(a_0,e)),\, \overline{F}_0(\overline{G}(a_1,b_0))\big) \;\leq\; \delta_2.$$

$\mu_2([0,1],\Sigma) = 1$, concluded

Repeat: $\quad d\big(\overline{G}(a_0,e), \overline{G}(a_1,b_0)\big)\big) \;<\; \delta_1.$

Because $\{a_0,a_1\} = \{0,1\}$, because of Σ, and because the function $\overline{F}_0$ is (δ_1,δ_2)-constrained, we now have:

$$1 \;=\; d(a_0,a_1) \;=\; d\big(\overline{F}_0(\overline{G}(a_0,e)),\, \overline{F}_0(\overline{G}(a_1,b_0))\big) \;\leq\; \delta_2.$$

Thus $\mu_2([0,1],\Sigma)$, being the infimum of such δ_2's, must be 1.

∎

Six sections of the talk.

Basics

Some calculations on approximate satisfaction. $\lambda(A, \Sigma)$.

Precise definitions for discontinuous satisfaction. $\mu_n(A, \Sigma)$.

An example: $\mu_1(A, \Sigma) = 0$; $\mu_2(A, \Sigma) = \mathrm{diam}(A)$.

Some further results

Algorithmic questions about λ and μ_n.

Some further results ...

- $\mu_2([0,1]^2, \text{same } \Sigma) = 1$ — uses Borsuk-Ulam Theorem.

Some further results . . .

- $\mu_2([0,1]^2, \text{same } \Sigma) = 1$ — uses Borsuk-Ulam Theorem.
- $\mu_2(\mathbb{R}, \text{Groups} + x^n \approx 1) \geq \text{radius}(\mathbb{R})$.

Some further results . . .

- $\mu_2([0,1]^2, \text{same } \Sigma) = 1$ — uses Borsuk-Ulam Theorem.
- $\mu_2(\mathbb{R}, \text{Groups} + x^n \approx 1) \ \geq \ \text{radius}(\mathbb{R})$.
- $\mu_2([0,1]^n, \text{Groups}) \ = \ \text{diameter}([0,1]^n)$.

Some further results . . .

- $\mu_2([0,1]^2, \text{same } \Sigma) = 1$ — uses Borsuk-Ulam Theorem.
- $\mu_2(\mathbb{R}, \text{Groups} + x^n \approx 1) \geq \text{radius}(\mathbb{R})$.
- $\mu_2([0,1]^n, \text{Groups}) = \text{diameter}([0,1]^n)$.
- $\mu_3(Y, \text{Lattices}) \geq 0.5$.

(Here Y stands for a Y-shaped one-dimensional space with each arm of unit length.)

. . . and two more

Spheres will be metrized proportional to great-circle distance, in such a way that antipodal points have distance 1.

. . . and two more

Spheres will be metrized proportional to great-circle distance, in such a way that antipodal points have distance 1.

Here Σ stands for the equations defining a binary operation with left-zero and left-one, or defining a commutative idempotent binary operation, or defining a ternary majority operation.

. . . and two more

Spheres will be metrized proportional to great-circle distance, in such a way that antipodal points have distance 1.

Here Σ stands for the equations defining a binary operation with left-zero and left-one, or defining a commutative idempotent binary operation, or defining a ternary majority operation.

- $\mu_1(S^1, \Sigma) \;=\; 2/3.$

. . . and two more

Spheres will be metrized proportional to great-circle distance, in such a way that antipodal points have distance 1.

Here Σ stands for the equations defining a binary operation with left-zero and left-one, or defining a commutative idempotent binary operation, or defining a ternary majority operation.

- $\mu_1(S^1, \Sigma) \;=\; 2/3.$
- $\mu_1(S^k, \Sigma) \;\geq\; 2/3.$

. . . and two more

Spheres will be metrized proportional to great-circle distance, in such a way that antipodal points have distance 1.

Here Σ stands for the equations defining a binary operation with left-zero and left-one, or defining a commutative idempotent binary operation, or defining a ternary majority operation.

- $\mu_1(S^1, \Sigma) \;=\; 2/3.$
- $\mu_1(S^k, \Sigma) \;\geq\; 2/3.$

We will sketch the proof of this last result $\longrightarrow$

$\mu_1(S^k, \Sigma) \;\geq\; 2/3; \quad \Sigma =$ ternary majority.

$\mu_1(S^k, \Sigma) \geq 2/3; \quad \Sigma =$ ternary majority.

We show that if $\overline{F}$ is a majority operation on S^k that is $(\delta, 2/3)$-constrained for some $\delta > 0$, then $\overline{F}$ can be deformed into a continuous majority operation $\overline{G}$. This would contradict the fact that $S^k \not\models \Sigma$.

$\mu_1(S^k, \Sigma) \ \geq \ 2/3; \quad \Sigma =$ ternary majority.

We show that if $\overline{F}$ is a majority operation on S^k that is $(\delta, 2/3)$-constrained for some $\delta > 0$, then $\overline{F}$ can be deformed into a continuous majority operation $\overline{G}$. This would contradict the fact that $S^k \not\models \Sigma$. The non-existence of such a constrained operation is tantamount to the inequality $\mu_1(S^k, \Sigma) \ \geq \ 2/3$.

$\mu_1(S^k, \Sigma) \ \geq \ 2/3; \quad \Sigma =$ ternary majority.

We show that if $\overline{F}$ is a majority operation on S^k that is $(\delta, 2/3)$-constrained for some $\delta > 0$, then $\overline{F}$ can be deformed into a continuous majority operation $\overline{G}$. This would contradict the fact that $S^k \not\models \Sigma$. The non-existence of such a constrained operation is tantamount to the inequality $\mu_1(S^k, \Sigma) \ \geq \ 2/3$.

Lemma

If F is a finite subset of S^k with diameter$(F) < 2/3$, then there is a convex subset A of S^k such that $F \subseteq A$. ∎

$\mu_1(S^k, \text{Ternary majority}) \geq 2/3$, continued

Recall that $\overline{F}$ is a majority operation on S^k that is $(\delta, 2/3)$-constrained.

$\mu_1(S^k, \text{Ternary majority}) \geq 2/3$, continued

Recall that $\overline{F}$ is a majority operation on S^k that is $(\delta, 2/3)$-constrained.

We begin with a triangulation of $(S^k)^3$ that contains sub-complexes triangulating $\{(x, x, z) : x, z \in S^k\}$ and so on, and such that each simplex has diameter $< \delta$.

$\mu_1(S^k, \text{Ternary majority}) \geq 2/3$, continued

Recall that $\overline{F}$ is a majority operation on S^k that is $(\delta, 2/3)$-constrained.

We begin with a triangulation of $(S^k)^3$ that contains sub-complexes triangulating $\{(x, x, z) : x, z \in S^k\}$ and so on, and such that each simplex has diameter $< \delta$.

By the lemma, $\overline{F}$ maps each simplex into a convex set K, which is a universal extensor in topology; i.e. maps into K can be extended continuously to larger sets. The desired map $\overline{G}$ is then constructed one simplex at a time, by recursion on the simplex dimension.

Six sections of the talk.

Basics

Some calculations on approximate satisfaction. $\lambda(A, \Sigma)$.

Precise definitions for discontinuous satisfaction. $\mu_n(A, \Sigma)$.

An example: $\mu_1(A, \Sigma) = 0$; $\mu_2(A, \Sigma) = \mathrm{diam}(A)$.

Some further results

Algorithmic questions about λ and μ_n.

Intractability of an algorithmic approach to compatibility.

Intractability of an algorithmic approach to compatibility.

Recall that $\mathbb{R} \models \Sigma$ is not a recursive (algorithmic) property of finite Σ.

Intractability of an algorithmic approach to compatibility.

Recall that $\mathbb{R} \models \Sigma$ is not a recursive (algorithmic) property of finite Σ.

In fact there is no space A for which $A \models \Sigma$ is known to be algorithmic (besides a few known spaces A where $A \models \Sigma$ iff Σ is essentially trivial).

Intractability of an algorithmic approach to compatibility.

Recall that $\mathbb{R} \models \Sigma$ is not a recursive (algorithmic) property of finite Σ.

In fact there is no space A for which $A \models \Sigma$ is known to be algorithmic (besides a few known spaces A where $A \models \Sigma$ iff Σ is essentially trivial).

E.g. it remains unknown whether $[0, 1] \models \Sigma$ is algorithmic.

Intractability of an algorithmic approach to compatibility.

Recall that $\mathbb{R} \models \Sigma$ is not a recursive (algorithmic) property of finite Σ.

In fact there is no space A for which $A \models \Sigma$ is known to be algorithmic (besides a few known spaces A where $A \models \Sigma$ iff Σ is essentially trivial).

E.g. it remains unknown whether $[0, 1] \models \Sigma$ is algorithmic.

Similarly, recursive enumerability remains unavailable.

Intractability of an algorithmic approach to compatibility.

Recall that $\mathbb{R} \models \Sigma$ is not a recursive (algorithmic) property of finite Σ.

In fact there is no space A for which $A \models \Sigma$ is known to be algorithmic (besides a few known spaces A where $A \models \Sigma$ iff Σ is essentially trivial).

E.g. it remains unknown whether $[0, 1] \models \Sigma$ is algorithmic.

Similarly, recursive enumerability remains unavailable.

$\ldots$ Looking at λ, we gain traction $\longrightarrow$

Recursive enumeration of approximate satisfaction.

We fix a list of operation symbols F_i ($i \in \omega$), which includes each arity infinitely often. Clearly every finite set of equations is definitionally equivalent to a finite set involving only the operation symbols F_i.

Recursive enumeration of approximate satisfaction.

We fix a list of operation symbols F_i ($i \in \omega$), which includes each arity infinitely often. Clearly every finite set of equations is definitionally equivalent to a finite set involving only the operation symbols F_i.

Theorem

Let K be a finite simplicial complex, with $|K|$ its geometric realization (as a topological space), and let $\alpha > 0$ be a computable real number. There is an algorithm $\mathcal{A}_{K,\alpha}$ whose output consists of those finite sets Σ of equations in F_i ($i \in \omega$) for which $\lambda(|K|, \Sigma) < \alpha$.

Informal description of the algorithm $\mathcal{A}_{K,\alpha}$.

Fix for the moment a single finite Σ, for simplicity having a single operation F_0, which is p-ary, while also fixing $M, N \in \omega^+$. We consider the complexes $C = (K^p)^{(M)}$ and $D = K^{(N)}$, where $^{(M)}$ denotes the M-th subdivided complex. And let us consider the case where α is rational.

Informal description of the algorithm $\mathcal{A}_{K,\alpha}$.

Fix for the moment a single finite Σ, for simplicity having a single operation F_0, which is p-ary, while also fixing $M, N \in \omega^+$. We consider the complexes $C = (K^p)^{(M)}$ and $D = K^{(N)}$, where $^{(M)}$ denotes the M-th subdivided complex. And let us consider the case where α is rational.

For each simplicial map $C \longrightarrow D$ (of which there are finitely many), we may check whether the corresponding continuous map, let us call it $\overline{F}_0 : |K|^p \longrightarrow K$, satisfies Σ within $< \alpha$. In fact this proposition lies in the first-order theory of reals (using barycentric co-ordinates), and so Tarski's algorithm will yield an algorithm to check this.

The algorithm $\mathcal{A}_{K,\alpha}$, continued

In case we discover a simplicial map satisfying Σ within α, we allow this fragment of the algorithm $\mathcal{A}_{K,\alpha}$ to output Σ.

The algorithm $\mathcal{A}_{K,\alpha}$, continued

In case we discover a simplicial map satisfying Σ within α, we allow this fragment of the algorithm $\mathcal{A}_{K,\alpha}$ to output Σ.

The Simplicial Approximation Theorem ultimately says that if $A \models_\beta \Sigma$ for some $\beta < \alpha$, then one of these simplicial maps will in fact satisfy Σ within α. Therefore if we loop the above subroutine through all M, all N, and all appropriate simplicial maps, we will sooner or later output Σ if and only if $A \models_\alpha \Sigma$.

The algorithm $\mathcal{A}_{K,\alpha}$, continued

In case we discover a simplicial map satisfying Σ within α, we allow this fragment of the algorithm $\mathcal{A}_{K,\alpha}$ to output Σ.

The Simplicial Approximation Theorem ultimately says that if $A \models_\beta \Sigma$ for some $\beta < \alpha$, then one of these simplicial maps will in fact satisfy Σ within α. Therefore if we loop the above subroutine through all M, all N, and all appropriate simplicial maps, we will sooner or later output Σ if and only if $A \models_\alpha \Sigma$.

All that remains is to consider all Σ. By a suitable interlacing, we may run the above process concurrently for all of them, thereby yielding an algorithm that outputs exactly those Σ with $A \models_\alpha \Sigma$.This completes our description of $\mathcal{A}_{K,\alpha}$.

An algorithm related to $\lambda(|K|, \Sigma) = 0$.

An algorithm related to $\lambda(|K|, \Sigma) = 0$.

Corollary

There is an algorithm $\mathcal{F}$ that takes no input, and whose output is an infinite stream of triples (K, Σ, s), with each K a finite complex, with each Σ a finite set of equations in the symbols F_i $(i \in \omega)$, and with each $s \in \mathbb{Z}^+$; such that the following condition holds: arbitrary K and Σ satisfy $\lambda(|K|, \Sigma) = 0$ iff (K, Σ, s) occurs in the output stream of $\mathcal{F}_K$ for arbitrarily large s.

Proof.

Interlace the algorithms $\mathcal{A}_{K,1/s}$ for $s = 1, 2, 3, \cdots$, and for all finite K. When $\mathcal{A}_{K,1/s}$ outputs Σ, then $\mathcal{F}$ will output (K, Σ, s). $\square$

Arithmetical character of $K \models \Sigma$, $\lambda(|K|, \Sigma) = 0$ and $\mu_n(|K|, \Sigma) = 0$.

The above corollary is tantamount to saying that the set of all Σ with $\lambda(|K|, \Sigma) = 0$ is Π^0_2 in the arithmetical hierarchy.

Arithmetical character of $K \models \Sigma$, $\lambda(|K|, \Sigma) = 0$ and $\mu_n(|K|, \Sigma) = 0$.

The above corollary is tantamount to saying that the set of all Σ with $\lambda(|K|, \Sigma) = 0$ is Π_2^0 in the arithmetical hierarchy.

Note that we do not have such a result for $|K| \models \Sigma$. Nor do we know how to obtain such a result for $\mu_n(|K|, \Sigma) = 0$. The problem here would be that the validity of $|K| \models_n^\varepsilon \Sigma$ requires some operations on $|K|$ that exactly model Σ. And clearly simplicial maps will not generally satisfy Σ exactly. (Example: groups on S^1).