Proabilistic consequence relations

Paul Égré and David Ripley

paulegre.free.fr and davewripley.rocks

Table of contents

Section 1

Setup

- The language $\mathcal{L}$ is a classical propositional language.
- An *argument* is $\Gamma \succ \Delta$, where Γ, Δ are *finite* sets of sentences.
- $\vDash_{\mathsf{CL}}$ is classical validity.
- A *probability assignment* is a function $Pr : \mathcal{L} \to [0, 1]$ such that:
  - $Pr(\top) = 1$, and
  - if $A \vDash_{\mathsf{CL}} \neg B$, then $Pr(A \vee B) = Pr(A) + Pr(B)$

The idea is to treat probability assignments like we treat logical valuations.

This has been done before (Adams, Knight, Paris, others), but not a ton.

Plus, all previous work that we know of is SET-FMLA,
which obscures some important distinctions.

Upsets:

An *upset* is some set $\alpha \subseteq [0, 1]$ such that:

- $1 \in \alpha$ and $0 \notin \alpha$, and
- if $x \in \alpha$ and $x \leq y \in [0, 1]$, then $y \in \alpha$.

A choice of upset is a choice of which probabilities are "high enough".

Open and closed:

- Every upset is either $[x, 1]$ or $(x, 1]$ for some $x \in [0, 1]$; that x is the upset's *threshold*.
- In the first case the upset is *closed*; in the second *open*.

A *counterexample notion* is a three-place relation between upsets, probability assignments, and arguments.

Parameterized by an upset, it says which assignments count as counterexamples to which arguments.

Given such a notion and an upset, we get a set of arguments to count as valid: the arguments such that no assignment is a counterexample to them.

- We consider three main counterexample notions:
  - material consequence,
  - preservation consequence,
  - symmetric consequence.

Section 2

Material consequence

Definition:

Given an upset α, let *Pr* be an α-material counterexample to $\Gamma \succ \Delta$ iff:
$Pr(\bigwedge \Gamma \supset \bigvee \Delta) \notin \alpha$.

$\Gamma \succ \Delta$ is α-materially valid (written $\Gamma \vDash^m_\alpha \Delta$)
when no *Pr* is an α-material counterexample to it.

That is, a material counterexample to an argument
is a *Pr* where the material-conditional version of the argument
does not have a high enough probability.

Iff there is no such *Pr*, then the argument is materially valid.

$\vDash^m_\alpha$ is $\vDash_{CL}$

Suppose $\Gamma \vDash_{CL} \Delta$. Then $\vDash_{CL} \bigwedge \Gamma \supset \bigvee \Delta$.
So for any Pr, α, we have $Pr(\bigwedge \Gamma \supset \bigvee \Delta) = 1 \in \alpha$.

Suppose $\Gamma \nvDash_{CL} \Delta$. Then $\nvDash_{CL} \bigwedge \Gamma \supset \bigvee \Delta$.
So there is some Pr where $Pr(\bigwedge \Gamma \supset \bigvee \Delta) = 0$,
and $0 \notin \alpha$ for any α.

Material consequence, then, gives us a picture
of how full SET-SET classical logic can fit together with uncertainty.

It does not, however, reflect this in its consequence relations:
they're all exactly classical.

Section 3

Preservation consequence

Definition:

Given an upset α, let Pr be an α-preservation counterexample to $\Gamma \succ \Delta$ iff: $Pr[\Gamma] \subseteq \alpha$ and $Pr[\Delta] \subseteq [0,1] \setminus \alpha$.

$\Gamma \succ \Delta$ is α-preservation valid (written $\Gamma \vDash^{p}_{\alpha} \Delta$) when no Pr is an α-preservation counterexample to it.

That is, a preservation counterexample to an argument
gives a high-enough probability to every premise
and a not-high-enough probability to every conclusion.

When $\Gamma \succ \Delta$ is preservation valid,
if everything in Γ has a high enough probability,
then so must something in Δ.

The SET-FMLA special case (Paris)

The SET-FMLA fragment of α-preservation consequence is well-behaved:

Strengthening to a classical limit

- If $\alpha \subseteq \beta$ and $\Gamma \vDash_{\beta}^{p} A$, then $\Gamma \vDash_{\alpha}^{p} A$
- $\Gamma \vDash_{\{1\}}^{p} A$ iff $\Gamma \vDash_{\mathsf{CL}} A$

As the upset tightens, more and more arguments become valid,
until at the limit of perfect certaintly classical logic is reached.

Consider our two extreme upsets: $(0,1]$ and $\{1\}$.

The $\textsc{Set-Set}$ preservation logics determined by these upsets are familiar, and neither is classical:

Super- and subvaluationistic consequence:

- $\models^{p}_{\{1\}}$ is supervaluationist consequence
- $\models^{p}_{(0,1]}$ is subvaluationist consequence

What it means that $\textsc{Set-Fmla}$ reaches a 'classical' limit at $\{1\}$: just that supervaluationist consequence is $\textsc{Set-Fmla}$ classical.

The neat strengthening we saw in the SET-FMLA fragment is *reversed* in FMLA-SET:

Weakening as upsets tighten:

If $\alpha \subseteq \beta$ and $A \vDash^{p}_{\alpha} \Delta$, then $A \vDash^{p}_{\beta} \Delta$

And the overall picture is nothing neat or simple:

Incomparability:

For any α, β, the two consequence relations $\vDash^{p}_{\alpha}$ and $\vDash^{p}_{\beta}$ are either identical or incomparable.

Definitions:

- Γ is α-satisfiable iff there is a *Pr* with $Pr[\Gamma] \subseteq \alpha$
- Δ is α-tautologous iff there is no *Pr* with $Pr[\Delta] \cap \alpha = \emptyset$

The role of rationals 1 (Knight):

For any finite Γ, there is some rational $x \in [0, 1]$
such that Γ is $[x, 1]$-satisfiable but not $(x, 1]$-satisfiable.

The role of rationals 2 (Fritz):

For any rational $x \in [0, 1]$, there is some finite Γ
such that Γ is $[x, 1]$-satisfiable but not $(x, 1]$-satisfiable.

Results on sameness and difference:

- If x is irrational, then $\vDash^{p}_{[x,1]} = \vDash^{p}_{(x,1]}$
- If x is rational, then $\vDash^{p}_{[x,1]}$ is incomparable to $\vDash^{p}_{(x,1]}$
- If x is the threshold of α and y is the threshold of β and $x \neq y$, then $\vDash^{p}_{\alpha}$ and $\vDash^{p}_{\beta}$ are incomparable
- There are uncountably many distinct preservation consequence relations

To build intuitions for preservation consequence,
three individually sufficient conditions for invalidity are useful:

If ____________, then $\Gamma \not\models^{p}_{\alpha} \Delta$:

- $\Gamma \not\models_{CL} \Delta$; or
- $\alpha \neq \{1\}$ and:
  - Γ is α-satisfiable,
  - $\bigvee \Delta$ is not a classical tautology, and
  - there is no $\gamma \in \Gamma$ with $\gamma \models_{CL} \bigvee \Delta$; or
- $\alpha \neq (0, 1]$ and:
  - Δ is not α-tautologous,
  - $\bigwedge \Gamma$ is classically satisfiable, and
  - there is no $\delta \in \Delta$ with $\bigwedge \Gamma \models_{CL} \delta$

Examples

It follows from these that:

- if $\alpha \neq \{1\}$, then $p, q \not\models^{p}_{\alpha} p \wedge q$;
- if $\alpha \neq \{1\}$, then $p \supset q, p \not\models^{p}_{\alpha} q$;
- for any α, $p, q \vee r \not\models^{p}_{\alpha} p \wedge q, r$

Conjecture:

We're led to a conjecture that, if true, would fully describe $\vDash^{p}_{\alpha}$ for all α other than the two extremes (which are already described):

Conjecture:

If $(0,1] \neq \alpha \neq \{1\}$, then $\Gamma \vDash^{p}_{\alpha} \Delta$ iff:

- Γ is α-unsatisfiable, or
- Δ is α-tautologous, or
- there is some $\gamma \in \Gamma$ and $\delta \in \Delta$ with $\gamma \vDash_{CL} \delta$

Certainly any of these three conditions suffices for validity; the conjecture is that *every* validity comes this way.

Section 4

Symmetric consequence

Focusing just on the SET-FMLA fragment of preservation consequence, there was a nice picture.

More and more arguments get valid as the upset narrows, until at the limit of $\{1\}$ we become classical.

Our SET-SET perspective made that fall apart, but it was a nice picture.

Mirror image

Given an upset α, its *mirror image* $\overline{\alpha} \subseteq [0,1]$ is $\{x \mid 1 - x \in \alpha\}$.

Definition:

Given an upset α, let *Pr* be an α-symmetric counterexample to $\Gamma \succ \Delta$ iff:
$Pr[\Gamma] \subseteq \alpha$ and $Pr[\Delta] \subseteq \overline{\alpha}$.

$\Gamma \succ \Delta$ is α-symmetric valid (written $\Gamma \vDash^{s}_{\alpha} \Delta$)
when no *Pr* is an α-symmetric counterexample to it.

A symmetric counterexample to an argument
gives a high-enough probability to every premise
and a low-enough probability to every conclusion.

We assume that strigent standards for what's high enough
come with stringent standards for what's low enough.

Again, we start with the extremes:

Extreme upsets:

- $\models^s_{\{1\}}$ is $\models_{CL}$
- $\Gamma \models^s_{(0,1]} \Delta$ iff either:
  - there is some $\gamma \in \Gamma$ with $\gamma \models_{CL}$, or
  - there is some $\delta \in \Delta$ with $\models_{CL} \delta$

This is something new:
unlike $\models^p$ $\models^s$ is sometimes classical,
and unlike $\models^m$ $\models^s$ isn't always classical.

We also get narrower upsets strengthening the consequence relation:

Narrower is stronger:

If $\alpha \subseteq \beta$ and $\Gamma \models^{s}_{\beta} \Delta$, then $\Gamma \models^{s}_{\alpha} \Delta$.

So the α-symmetric consequence relations form a clean linear order by strength.

Results on sameness and difference:

- If x is irrational, then $\models^{s}_{[x,1]} \; = \; \models^{s}_{(x,1]}$
- If x is rational, then $\models^{s}_{[x,1]} \; \neq \; \models^{s}_{(x,1]}$
- If x is the threshold of α and y is the threshold of β and $x \neq y$, then $\models^{s}_{\alpha} \; \neq \; \models^{s}_{\beta}$
- There are uncountably many distinct symmetric consequence relations

Relations to preservation:

- For any α, β, we have $\models^{s}_{\alpha} \; \neq \; \models^{p}_{\beta}$
- If $.5 \in \alpha$, then $\models^{s}_{\alpha} \; \subseteq \; \models^{p}_{\alpha}$;
  and if $.5 \notin \alpha$, then $\models^{p}_{\alpha} \; \subseteq \; \models^{s}_{\alpha}$

Closed upsets are limits:

If $\Gamma \models^{S}_{[x,1]} \Delta$, then there is some $\alpha \supsetneq [x,1]$ such that $\Gamma \models^{S}_{\alpha} \Delta$

- Another way to think about that: $\models^{S}_{[x,1]} = \bigcup_{\alpha \supsetneq [x,1]} \models^{S}_{\alpha}$
- Another nother way: no argument becomes valid at a closed upset.

And $\{1\}$ is a closed upset.

So symmetric consequence really gives us
what preservation consequence only appeared to:
strengthening logics as upsets narrow,
until classical logic is reached at $\{1\}$

Examples

- $p_1, \ldots, p_n \vDash^S_\alpha \bigwedge p_i$ iff $\alpha \subseteq (\frac{n}{n+1}, 1]$
- $p_1, p_1 \supset p_2, \ldots, p_{n-1} \supset p_n \vDash^S_\alpha p_n$ iff $\alpha \subseteq (\frac{n}{n+1}, 1]$
- $p, q \vee r \vDash^S_\alpha p \wedge q, r$ iff $\alpha \subseteq (\frac{3}{4}, 1]$

Indeed:

If $\Gamma \succ \Delta$ is classically valid and has no classically valid proper subargument,
then $\Gamma \vDash^S_\alpha \Delta$ iff $\alpha \subseteq (\frac{n-1}{n}, 1]$,
where n is the number of sentences in $\Gamma \succ \Delta$.

So bigger arguments can take longer to shake out,
but eventually any classically-valid argument gets caught as we narrow our upset.

Section 5

Conclusion

- There are lots of ways to define consequence relations from probability assignments; we've looked at three.
- Material consequence is always classical.
- Preservation consequence is never classical,
  but can be super- or subvaluational.
- Preservation consequence is also a bit messy in the middle.
- Symmetric consequence is classical at the limit,
  and gradually approaches that limit in a describable way.
  - (Also if $\alpha \neq \{1\}$ then $\models^{\mathbf{s}}_{\alpha}$ is nontransitive xor nonreflexive.)

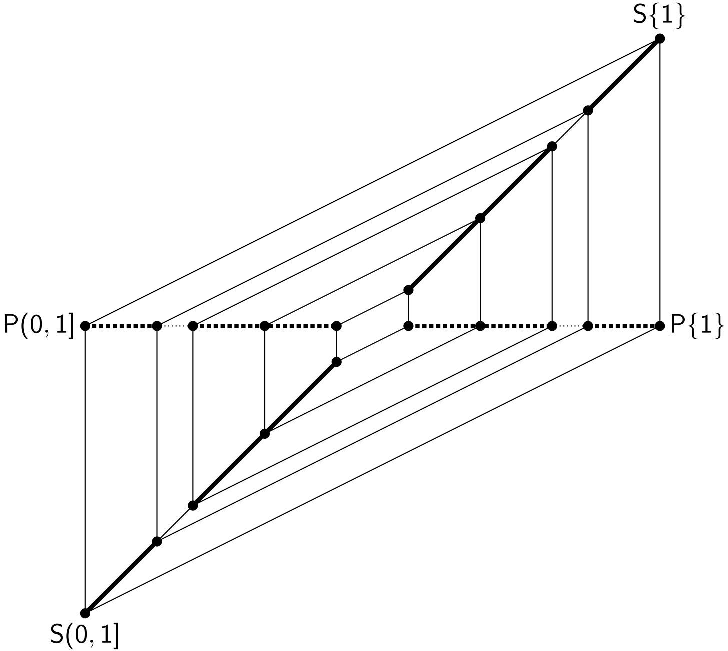
S{1}
P(0, 1]
P{1}
S(0, 1]