Implication-Space Semantics for RDF

Towards a formal semantics of predictable inference

Bradley P. Allen
b.p.allen@uva.nl

CSAI group, Vrije Universiteit Amsterdam
11 November 2026

Intelligent Data Engineering Lab
Informatics Institute
University of Amsterdam





van Harmelen: rethinking semantics for an agentic world



- Hitzler & van Harmelen (2010): A reasonable Semantic Web
  - Agent *A* sending the claim *(NL hasCapital Amsterdam)* to agent *B* carries with it unspoken inferences
    - *(Amsterdam a City)* follows; *(NL hasCapital TheHague)* is ruled out; ...
  - Agents *A* and *B* are interoperable when they share inferences
  - Model theory is compatible with shared inference but not required by it
  - Defeasible approaches exist but are largely unadopted
  - We need a semantics of shared inference
- van Harmelen (2026): "Semantics = predictable inference"
  - Since 2000 knowledge graphs have solved data interoperability
  - Communication between agents is the new interoperability
  - Interoperability depends on predictable inference

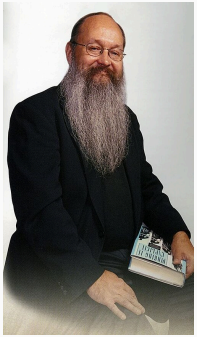
This talk: inferentialist semantics = semantics for an agentic world

- "Predictable" means understanding what an agent is committed to
- Commitment has a theory: inferentialism
- This theory is formalized by implication-space semantics
- Applied to RDF, implication-space semantics recovers RDF entailment regimes exactly
- RDF entailment regimes are a monotone corner of implication space
- Achieving full agent interoperability requires the (nonmonotonic) rest of it

What does “predictable inference” mean?

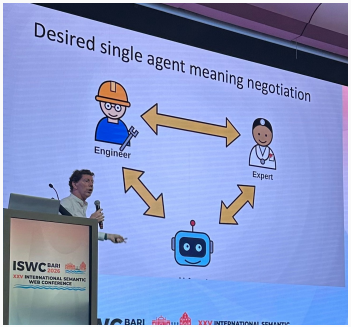
- *A* sends *B* a claim; *B* subscribes; what can *A* expect *B* to infer?
  - The claim: `(NL hasCapital Amsterdam)`
- Two readings of “predictable”:
  - **Alethic**: what follows from the claim in every model; *B* is not in view; the prediction is just correctness
  - **Normative**: what *B* is now committed to `(Amsterdam a City)`, and precluded from `(NL hasCapital TheHague)`, against the base both keep score by
- *B* subscribes by deferral: *B* takes *A* to be entitled
- What *B* will infer and what *B* is committed to can differ
- Only a normative reading has room for that gap

Inferentialism: a theory of meaning that accounts for predictable inference



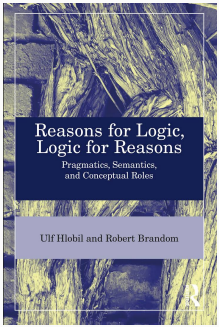
- Inferentialism (Brandom 1994): the philosophical position that the conceptual content of a claim is its inferential role within a discursive practice
  - *Discursive practice*: any social practice with the normative structure of commitment and entitlement
  - *Inferential role of a claim*: what it follows from, what follows from it, and what it is incompatible with
  - *Conceptual content*: what is grasped when one has practical mastery of inferential role
- Agents who have mastered the same role draw the same inferences from it
- That is what makes *B*'s inferences predictable to *A*

GOGAR: a game of giving and asking for reasons where meaning is negotiated



- Games of giving and asking for reasons (Brandom 1994): agents making claims, challenging them, and giving reasons for them
- Communication can change what an agent is committed and entitled to, with agents scorekeeping one another's commitments, entitlements, and incompatibilities
- Scorekeeping makes predictable inference and meaning negotiation possible

Implication-space semantics: formalizing meaning using reasons for claims



- Meaning from a dyadic reason relation between sets of claims, not a unary truth property of claims
- The arity difference is where agents come in
  - A unary property of a sentence says nothing about anyone
  - A relation between what has been asserted and what follows is the shape of a GOGAR move and its consequences
- The normative reading of “predictable” needs the dyadic primitive
- The game supplies it: pragmatics first, then semantics

Implication-space semantics: terminology

- $\mathbb{B}$: the claims agents can make (the *bearers*)
- $\mathcal{S} = \mathcal{P}(\mathbb{B}) \times \mathcal{P}(\mathbb{B})$: the implication space
- $\langle \Gamma, \Delta \rangle \in \mathcal{S}$: an implication with premises Γ and conclusions Δ
- $\mathbb{I} \subseteq \mathcal{S}$: what a discursive practice counts as good implications
- $\langle \mathbb{B}, \mathbb{I} \rangle$: a frame
- $\mathfrak{B} = \langle \mathbb{B}, \mathbb{I} \rangle$: a base where $\mathbb{B}$ consists of atomic sentences

Implication-space semantics: reading implications and bases

- $\langle \Gamma, \Delta \rangle \in \mathbb{I}$: asserting all of Γ is committing to at least one of Δ
- $\langle \Gamma, \emptyset \rangle \in \mathbb{I}$: asserting all of Γ is out of bounds
- A base $\mathfrak{B}$ is a knowledge base where "held good" replaces "held true"
  - A fact is $\langle \emptyset, \{t\} \rangle$: $\langle \emptyset, \{\texttt{(NL hasCapital Amsterdam)}\} \rangle$
  - A rule is $\langle \Gamma, \{t\} \rangle$:
    $\langle \{\texttt{(NL hasCapital Amsterdam)}\}, \{\texttt{(Amsterdam a City)}\} \rangle$
  - A constraint is $\langle \Gamma, \emptyset \rangle$:
    $\langle \{\texttt{(NL hasCapital Amsterdam)}, \texttt{(NL hasCapital TheHague)}\}, \emptyset \rangle$

Implication-space semantics: inferential role, made precise

- $\mathrm{RSR}(\langle \Gamma, \Delta \rangle) = \{\langle x, y \rangle \mid \langle \Gamma \cup x,\ \Delta \cup y \rangle \in \mathbb{I}\}$: the range of subjunctive robustness
- Further commitments an agent could add with the implication still good
- Same range, same role, same meaning
- *A*'s prediction about *B* holds across the commitments of *B* inside the range
- In the monotone case, every good implication has full range
- A defeasible inference is one with a limited range
  - $\langle \{\texttt{(Tweety a Bird)}\}, \{\texttt{(Tweety flies)}\} \rangle \in \mathbb{I}$
  - But $\langle \{\texttt{(Tweety a Bird)}, \texttt{(Tweety a Penguin)}\}, \{\texttt{(Tweety flies)}\} \rangle \notin \mathbb{I}$
  - Adding `(Tweety a Penguin)` is outside its range

Implication-space semantics: from implications to conceptual content

- The content of a claim a is its role: the implications it takes part in, as premise, generated by $\langle \{a\}, \emptyset \rangle$; as conclusion, by $\langle \emptyset, \{a\} \rangle$
- Roles compose: asserting several claims together, concluding to either of several
- $G \models D$ iff every implication the roles of G and D generate is in $\mathbb{I}$
- A frame is fit to a base $\mathfrak{B}$ (a *model* of $\mathfrak{B}$) if every implication of $\mathfrak{B}$ is good in it
- $G \models_{\mathfrak{B}} D$: $G \models D$ in *every* frame fit to $\mathfrak{B}$

Implication-space semantics for RDF: triples and graphs

- $\mathbb{B}$: generalized RDF triples over a finite vocabulary N of IRIs and literals
- A ground graph G is read conjunctively on both sides
  - As premise, assert *all* of it: $\langle \mathrm{B}(G), \emptyset \rangle$, with $\mathrm{B}(G)$ the triples of G as bearers
  - As conclusion, deny *any part*, jointly: $\langle \emptyset, S \rangle$, $\emptyset \neq S \subseteq \mathrm{B}(G)$
- A graph G with blank nodes is read through its ground instances over N, which includes fresh IRIs for the blank nodes of G
  - As premise, assert some instance
  - As conclusion, deny every instance
- **In GOGAR:** publishing a graph is asserting all of it

Implication-space semantics for RDF: entailment regimes

The base a regime generates

$\Gamma \mid\!\sim_{\mathfrak{B}_R} \Delta$ iff $\bot \in \mathrm{cl}_R(\Gamma)$ or $\Delta \cap \mathrm{cl}_R(\Gamma) \neq \emptyset$

- A regime R: all instances $\langle A, c \rangle$ of a set of rule patterns, c a triple or $\bot$; conclusions introduce no new terms
- $\mathrm{cl}_R(\Gamma)$: apply the patterns to a fixpoint; RDF, RDFS, ρdf, OWL 2 RL/RDF all fit
- Schema triples are bearers in Γ; the rule patterns turn them into licenses
- Given a context Γ, a triple t is committed if $t \in \mathrm{cl}_R(\Gamma)$, precluded if $\bot \in \mathrm{cl}_R(\Gamma \cup \{t\})$, else open
  - *(Amsterdam a City)* committed
  - *(NL hasCapital TheHague)* precluded (given the constraint on slide 8)
- **In GOGAR:** the regime is the shared scorekeeping standard

Implication-space semantics for RDF: recovery results

Theorem 1 (recovery). For any base $\mathfrak{B}$ over RDF triples

$G \models_{\mathfrak{B}} H$ iff it holds in the single fit Herbrand model $M_{\mathfrak{B}}^{\mathrm{RDF}}$

Theorem 2 (regimes). For every rule-defined regime R

$G \models_{\mathfrak{B}_R} H$ iff $G \vdash_R H$ (G inconsistent, or $\mathrm{cl}_R(G)$ simply entails H)

- Regimes: simple; RDFS (consistent G, $D = \{$`rdf:langString`, `xsd:string`$\}$); OWL 2 RL/RDF (w/o `dt-type2`, `dt-eq`, `dt-diff`)
- The base alone fixes consequence; Monotonicity and Cut are not assumed
- Checked in Lean 4 through Corollary 1, and by enumerating ∼8M graph pairs; zero disagreements
- **In GOGAR:** publishing G commits you to exactly what the regime's rules derive from it

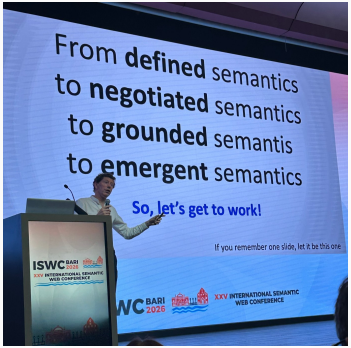
Next steps in the near term: the nonmonotonic case

- Recovery is of the rule characterization and matches the model theory wherever the rules are complete
- That is expected; the value is in what it locates
  - RDF entailment regimes occupy an austere corner of implication space: monotone, closed under Cut, explosive
  - Traditional model-theoretic RDF semantics: the corner is all there is
  - Sufficient for data interoperability, but insufficient for agent interoperability
  - Implication-space semantics: the corner is but one region
- To be proven, for bases with defeat and incompatibility:
  - Consequence is non-monotone by defeat, without explosion
  - Incompatibility is cumulative, so materialization stays sound
  - Skolemizing blank nodes gives the correct verdicts (Hogan et al. 2014)

Next steps in the near term: a reasoner implementation

- NonMonotonic MultiSuccedent logic (NMMS): a sequent calculus for supraclassical reasoning over reason relations (Hlobil & Brandom 2025)
  - Soundness and completeness for RDF is the next theorem once the nonmonotonic case is addressed
- pyNMMS: materializes a regime's closure into an RDF store (Oxigraph), then runs NMMS proof search over it, with prototypes of defeaters and constraints
  - Query cost is flat in the size of the graph and exponential only in the query (the decision problem is coNP-hard)
  - Agrees with SPARQL *ASK* on every query over a 10^7-triple RDFS graph at 14–30 μs per query; scorekeeping queries on 6M triples in under 10 ms each
  - Python package: *https://pypi.org/project/pyNMMS/*

Next steps in the long run: van Harmelen's research programme



semantics	inferentialist reading
defined	recovered, as commitments (Theorems 1 and 2)
negotiated	given a mechanism: scorekeeping and repair
grounded	language entries and exits (Sellars 1953) ground the *practice*
emergent	roles converging under repair

Summary

- Communication between agents is the new interoperability
- Achieving agent interoperability requires new pragmatics and semantics
  - Pragmatics: normative, based on the game of giving and asking for reasons
  - Semantics: inferential, based on reason relations instead of a truth property
- For RDF, the new semantics recovers the existing entailment regimes exactly, so nothing in current practice changes
- The work ahead is to prove we can handle exceptions and disagreement using the same machinery

Thank you

and thanks to Frank van Harmelen, for the challenge

b.p.allen@uva.nl
bradleypallen.org

Preprint: https://doi.org/10.5281/zenodo.24116157

