# BRIDGING LANGUAGE MODELS AND SEMANTIC KNOWLEDGE

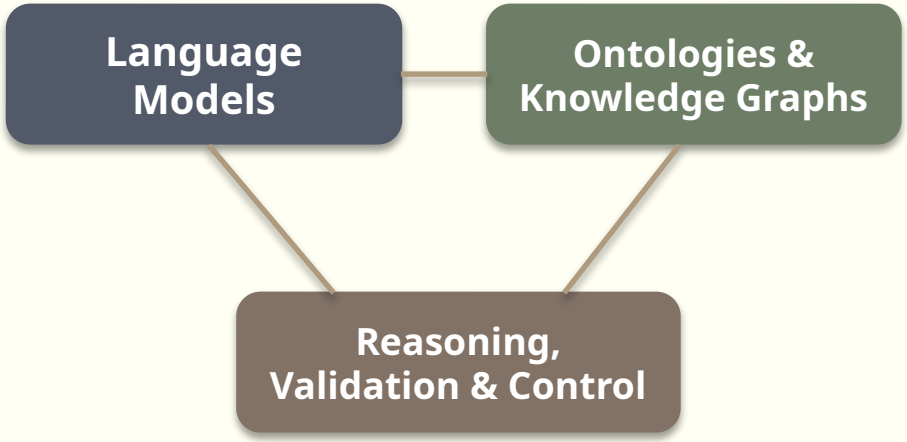


## Neuro-Symbolic Approaches to Ontology Engineering, Alignment, and Knowledge Graph Completion

Recent work with E. Skaperdas, I. Gocev, G. Meditskos and K. Cami

Nick Bassiliades · Intelligent Systems Lab · School of Informatics · Aristotle University of Thessaloniki

Olympia AI Data Infrastructure Summit 2026 · Heraklion, Crete · 4–5 November 2026

# Studies behind the talk

- Skaperdas, Cami & Bassiliades — OG-NSD: Neuro-Symbolic Ontology Drafting from Natural-Language Requirements (LLMs4KGOE @ ESWC 2026).
- Gocev, Meditskos & Bassiliades — Revisiting Ontology Matching with Language Models (WIMS 2026).
- Gocev, Meditskos & Bassiliades — Explainable Ontology Matching with Language Models and Abductive Reasoning (WIMS 2026).
- Skaperdas & Bassiliades — Ontology-Aware Relation Prediction over EurostatKG: A Unified Neural Framework (IJCKG 2025, Best Paper Award).
- Skaperdas & Bassiliades — STC: Inference-Time Semantic Target Control for Knowledge Graph Completion Rankings (ISWC 2026).

All five studies are recent work of the Aristotle University of Thessaloniki research group.

# Why semantics still matters for AI data infrastructure

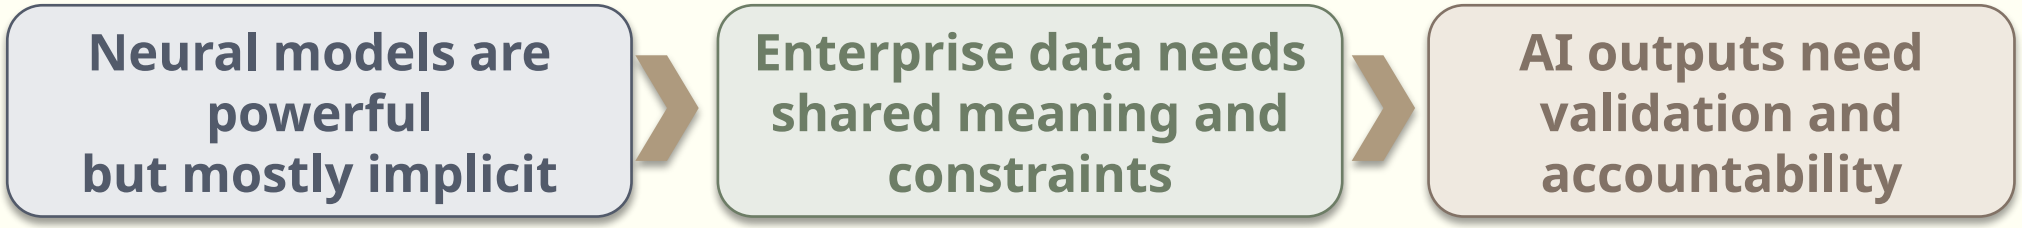


**A semantic layer can become an active interface between data, models and reasoning.**

- Ground neural models in explicit domain knowledge
- Connect heterogeneous schemas and vocabularies
- Make generation and prediction more checkable and controllable

# Five studies, one research direction

**NEURAL / LANGUAGE MODELS → SEMANTIC KNOWLEDGE**

**Draft** → **Align** → **Explain**

OG-NSD
LLM-driven ontology drafting

LM matching + abductive explanations

**SEMANTIC KNOWLEDGE → NEURAL MODELS**

**Guide learning** → **Control output**

Ontology-aware KG completion

Inference-time semantic target control

**Common question: how can neural AI and explicit semantics strengthen each other?**

# From natural-language requirements to an ontology draft

Skaperdas · Cami · Bassiliades · LLMs4KGOE 2026

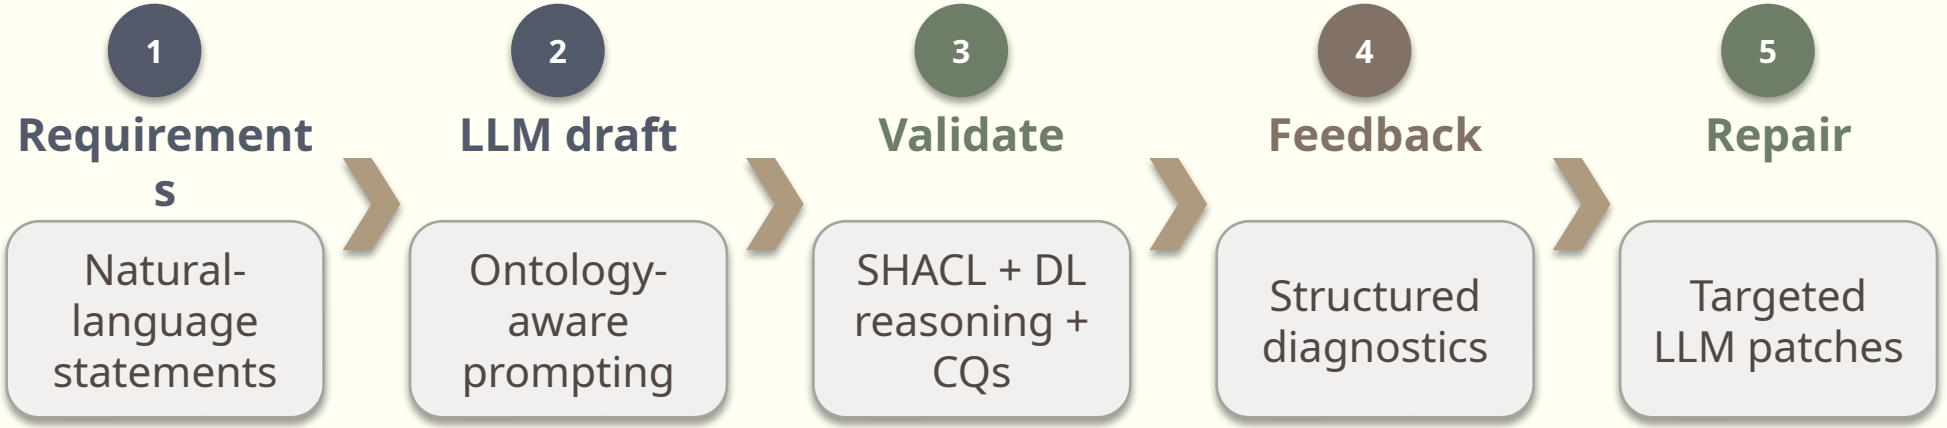


**Closed loop: symbolic violations become corrective prompts for the next draft.**

- Generation is fast, but can omit constraints or invent vocabulary
- Validation is part of generation — not only a final post-hoc check

*Based on: OG-NSD — Neuro-Symbolic Ontology Drafting from Natural-Language Requirements*

# OG-NSD: the useful part is control, not simply more repair

OG-NSD · controlled pilot study

**ONE-SHOT LLM**

**Strong low-temperature baseline**

**CONTROLLED REPAIR**

**Can improve stability and structure**

**AGGRESSIVE REPAIR**

**Can over-generate and destabilize reasoning**

**Main lesson**

**The challenge is not only generating ontologies.**
**The challenge is regulating generation with explicit semantic feedback.**

No single metric is enough: structure, logic and competency questions capture different aspects.

*Pilot result: conservative feedback is useful; aggressive CQ-driven repair can harm precision and graph stability.*

# What did symbolic feedback actually change?

OG-NSD · LLMs4KGOE 2026

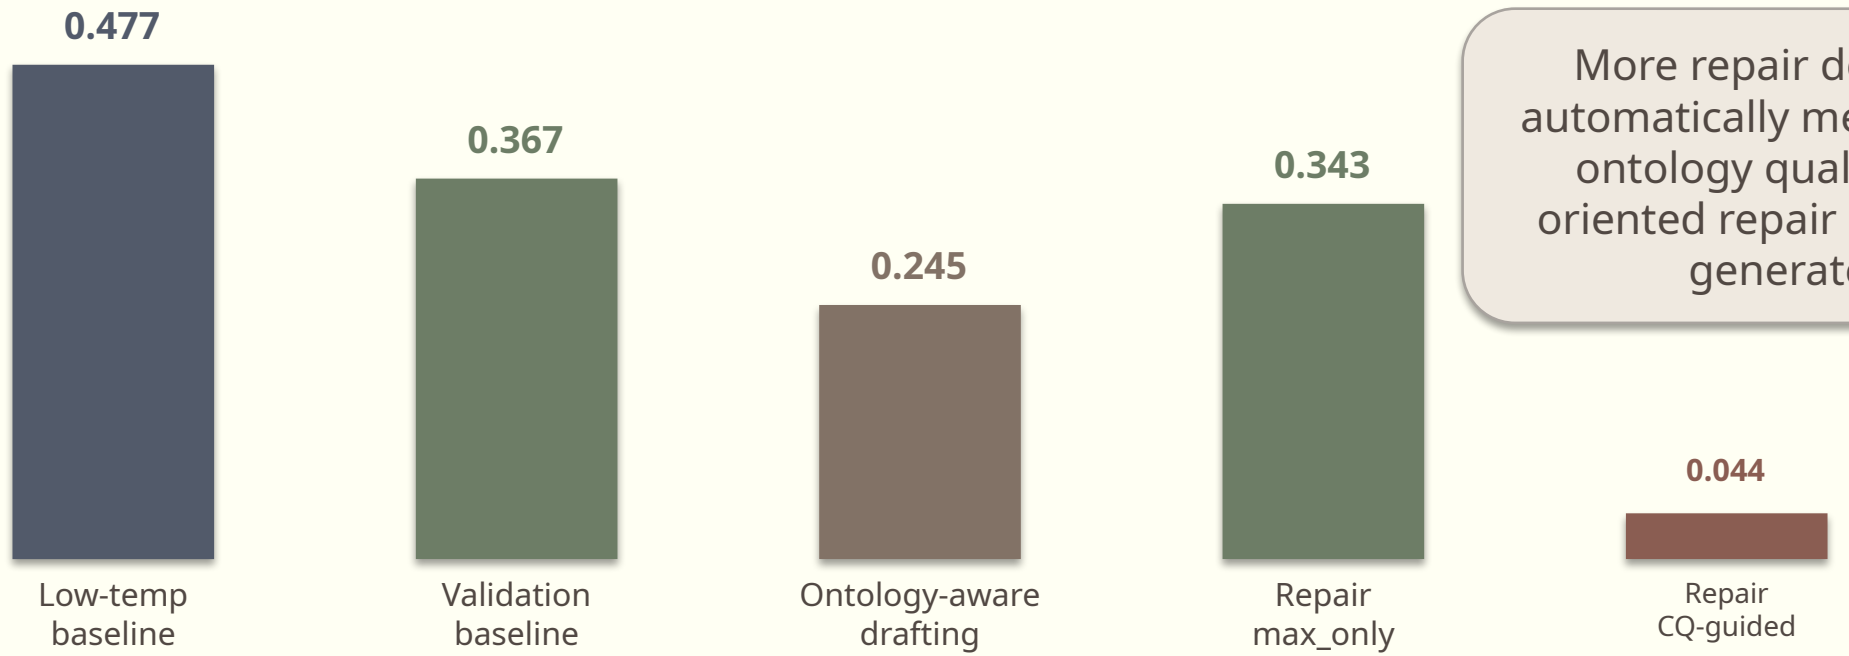


**Key caution**

More repair does not automatically mean better ontology quality. CQ-oriented repair can over-generate.

*Pilot configuration-level evidence; not variance-controlled performance claims.*

# Language models for ontology matching: context helps bridge vocabulary

Gocev · Meditskos · Bassiliades · WIMS 2026

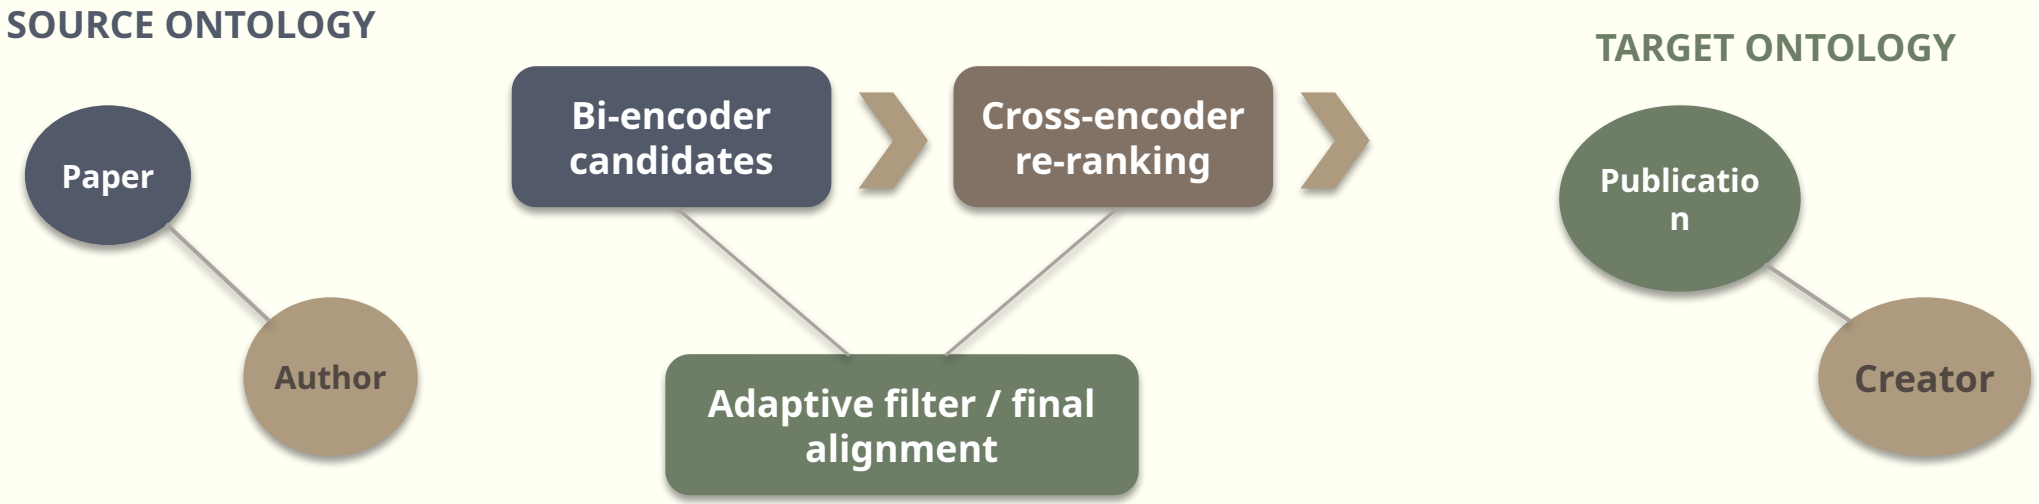


**Language models add contextual similarity when labels are different but meanings are close.**

- BERT-based encoders were evaluated as matching components
- SapBERT bi-encoder was robust and computationally attractive
  - Purely textual matching still misses structural and logical information

*Study focus: bi-encoder vs cross-encoder architectures for ontology alignment.*

# Bi-encoder or cross-encoder?

Gocev et al. · WIMS 2026

- Bi-encoder: encode entities independently and retrieve top-k candidates
  - Fast and scalable for large ontologies
- Cross-encoder: score source–target pairs jointly
  - Richer interaction, but slower
- Adaptive top-k filter: trade recall for higher precision
- Limitation: purely textual matching does not use ontology axioms or structural consistency

*Evaluated on OAEI Bio-ML ontology alignment tasks.*

# From a match score to a semantic explanation

Gocev · Meditskos · Bassiliades · WIMS 2026




**Illustrative explanation pattern**

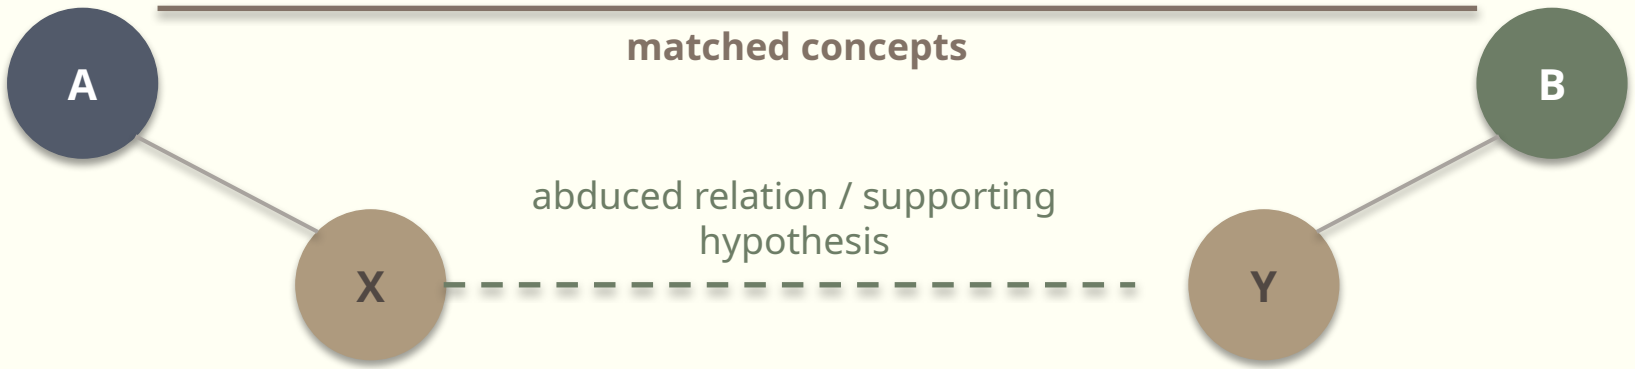


**The explanation exposes the semantic relations that would make the alignment understandable.**

*Hybrid pipeline: language models identify similar concepts; abduction generates explanations from ontology knowledge.*

# How does abduction explain an alignment?

Gocev · Meditskos · Bassiliades · WIMS 2026

$$C_{src} = \text{Adenolymphoma} \sqcap \exists \text{hasFindingSite}.\text{SalivaryGland}$$
$$C_{tgt} = \text{WarthinTumor} \sqcap \exists \text{isFoundAt}.\text{SalivaryGland} \sqcap \exists \text{isType}.\text{Benign}$$

**SNOMED CT · source description**




**after contraction**
**same tree structure**

root to root

SalivaryGland to SalivaryGland

**NCIT · target description**

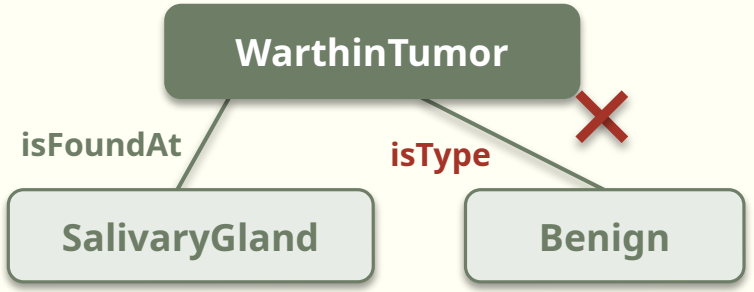


**CONTRACT / PRUNE**
**∃isType.Benign**

### 1 Structural step

Prune the unmatched isType → Benign branch.
Embed ∃isType.Benign into the target root.

### 2 Missing semantic relations

Adenolymphoma ↔ WarthinTumor ⊓ ∃isType.Benign
hasFindingSite ↔ isFoundAt

### 3 Abductive hypothesis

$$H = \{\ \text{Adenolymphoma} \equiv \text{WarthinTumor} \sqcap \exists \text{isType}.\text{Benign},\ \text{hasFindingSite} \equiv \text{isFoundAt}\ \}$$

**The pruned information is not discarded — it becomes part of H. Therefore $O_{src} \cup O_{tgt} \cup H \models C_{src} \equiv C_{tgt}$.**

# NOW REVERSE THE DIRECTION

From neural models helping semantics → semantics helping neural models

# Ontology-aware knowledge graph completion

**Skaperdas · Bassiliades · IJCKG 2025 · Best Paper Award**

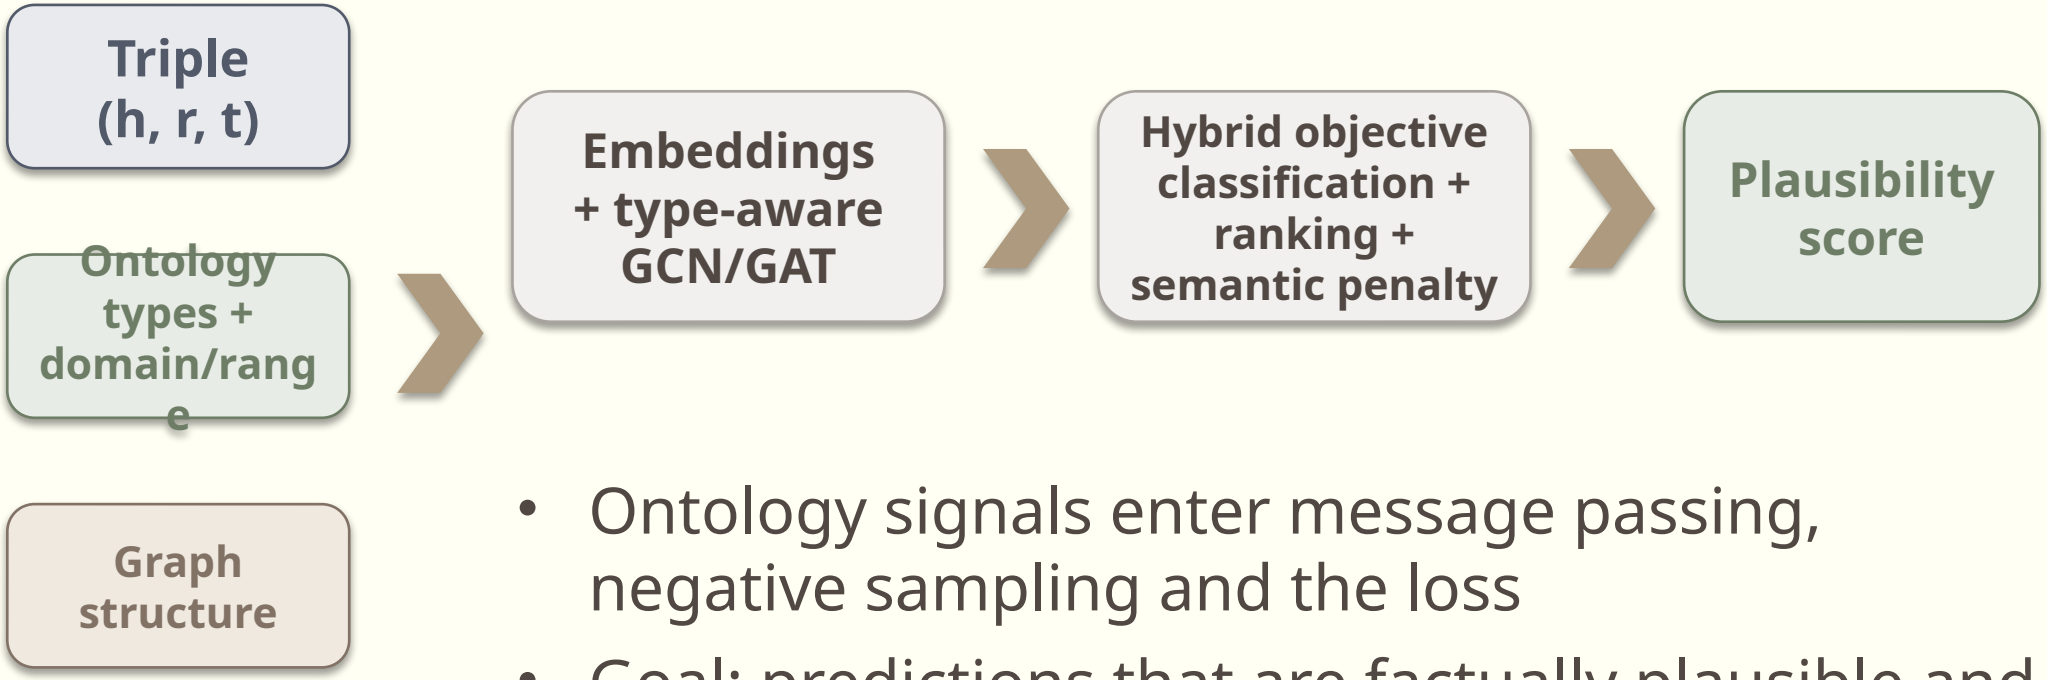


- Ontology signals enter message passing, negative sampling and the loss
- Goal: predictions that are factually plausible and semantically compliant

*EurostatKG study: symbolic knowledge is integrated into neural relation prediction during training.*

# A simple ontology constraint can change what the model learns

EurostatKG · ontology-guided relation prediction

## Illustrative schema rule

**domain(r) = StatisticalDataset**
**range(r) = StatisticalConcept**

✓ **valid types**

✗ **type violation**

Constraint-aware negatives teach the model that “plausible from embeddings” is not always “semantically valid”.

## EurostatKG result snapshot

**F1**
**0.96**

**PR-AUC**
**0.99**

**Soft H@1**
**0.97**

**Best Paper Award · IJCKG 2025**

**Takeaway: explicit semantics improves the learning signal, not only the final evaluation.**

*Reported metrics: F1 0.96, PR-AUC 0.99, Soft Hits@1 0.97 on EurostatKG.*

# Three ways to inject ontology knowledge into learning

IJCKG 2025 · Best Paper Award

### 1 Type-aware propagation

rdf:type information is embedded and propagated through GCN/GAT layers, so node representations reflect both graph structure and class membership.

### 2 Constraint-aware negatives

Training samples include deliberate domain/range violations, giving the model semantically meaningful hard negatives.

### 3 Semantic loss

A soft penalty is added when a candidate triple violates declared domain or range constraints. It guides learning without hard rejection.

**Symbolic knowledge appears in representation, supervision and evaluation.**

*EurostatKG asserted types and rdfs:domain/rdfs:range constraints are used operationally.*

# Even a good ranker can return a semantically poor top-k



**Frozen ranking**

1. **c1 admissible**
2. c2 non-admissible
3. c3 non-admissible
4. **c4 admissible**
5. c5 non-admissible
6. **c6 admissible**

**Visible top-5 contains only 2 admissible**

**Requested quota q = 3**

**But c6 is admissible and already exists in Top-M**

**This is a selection problem.**

**No retraining is needed.**
**No new candidate is needed.**
**We only need controlled re-ranking.**

**Question: what is the minimum semantic intervention that meets the quota?**

*STC separates retrieval failure from correctable selection failure in a fixed candidate window.*

# STC: semantic control after ranking, without retraining



## Simple control rule

**controlled score = model score − λ × semantic penalty**

OptQ chooses the minimum query-specific pressure λ needed to satisfy the requested quota — when the fixed window makes it feasible.

**Output is explicit and auditable: feasible, controlled, or infeasible.**

## Semantic improvement at q = 5

| EurostatKG | DBpedia | DrugBank |
|---|---|---|
| ΔViol@10 −0.269 | ΔViol@10 −0.183 | ΔViol@10 −0.488 |
| Pres@10 0.731 | Pres@10 0.823 | Pres@10 0.512 |

**Takeaway: semantics can act as a post-hoc control layer over a frozen neural ranker.**

*Canonical STC experiments: lower violation rate while preserving part of the original top-10 ranking.*

# What exactly does STC guarantee?



**Feasible window**
$|A_x| \geq q$

OptQ returns the minimum sufficient λ within the controlled-score family and guarantees at least q admissible items in the returned top-k.

**Infeasible window**
$|A_x| < q$

No re-ranking restricted to the supplied Top-M window can create missing admissible candidates. STC returns an infeasibility certificate and keeps the frozen fallback.

**Scope: guarantee is relative to the supplied semantic statuses, energies and finite candidate window.**

*STC does not claim factual truth, global OWL consistency, or retrieval beyond the supplied window.*

# One pattern across all five studies

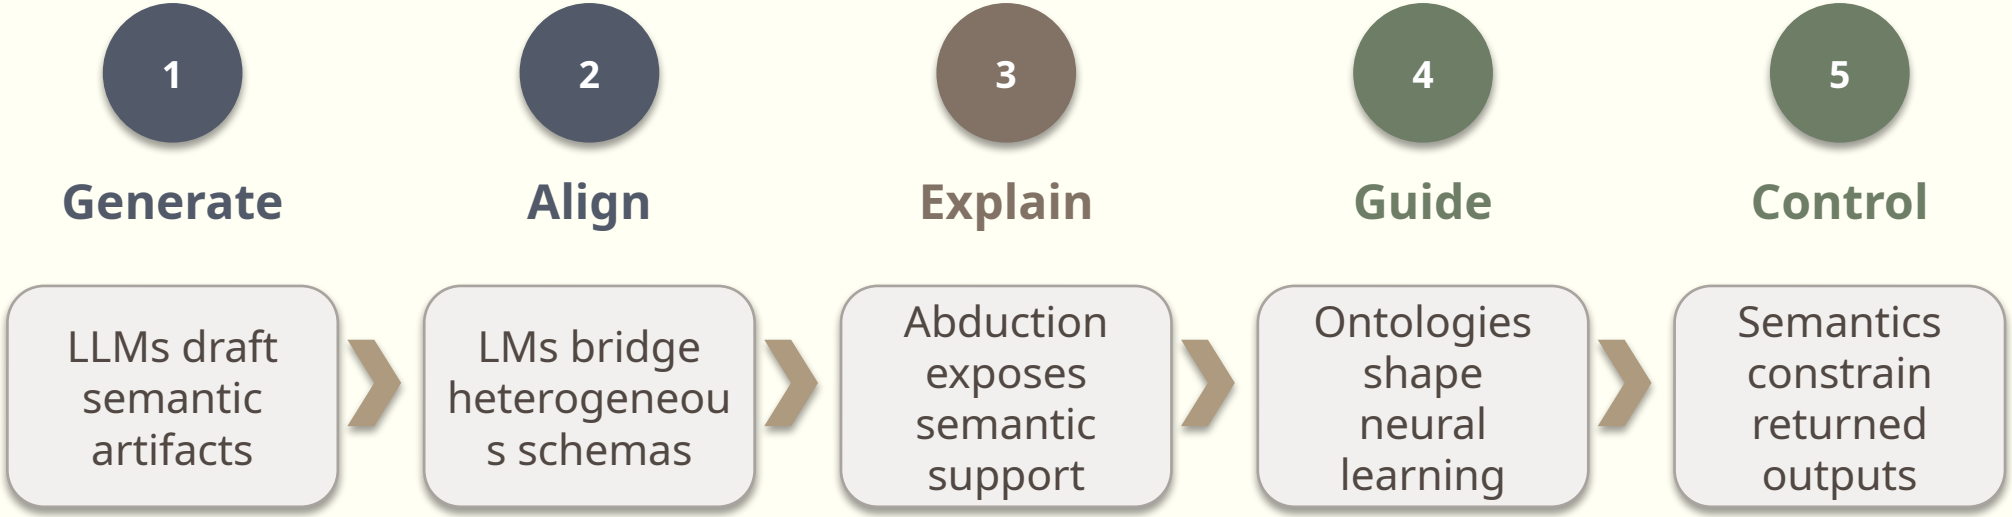


**Semantics is useful when it becomes operational — inside generation, learning, explanation, or control.**

This is the common thread, despite the different models and application settings.

# Implication for enterprise ontology and AI data infrastructure

**Enterprise ontology can be more than a shared vocabulary.**

| GROUND | CHECK | CONTROL |
| --- | --- | --- |
| Give models explicit domain concepts, types and constraints. | Validate generated knowledge and detect semantic failures. | Influence learning, ranking and downstream decisions. |

**A trusted semantic layer can connect enterprise data with AI systems in a way that is inspectable and auditable.**

# Take-home messages

- Explicit semantics and neural models are complementary — not competing approaches.
- Language models can accelerate ontology construction and alignment, but need semantic checks.
- Ontologies can actively guide learning, explanations, and inference-time control.
- The next challenge is controlled, auditable interaction between neural generation and symbolic knowledge.

**Future direction: richer semantics, human-in-the-loop control, and reusable semantic infrastructure for AI agents.**

# Questions & discussion

## Thank you

Nick Bassiliades · Aristotle University of Thessaloniki

**nbassili@csd.auth.gr**

Recent work with Efstratios Skaperdas, Ivan Gocev,
Georgios Meditskos and Kristi Cami

Olympia AI Data Infrastructure Summit 2026 · 4–5 November 2026 · Heraklion, Crete