CITY
ST GEORGE'S
UNIVERSITY OF LONDON

GUARD Project

Ensuring Interoperable and Trustworthy Knowledge Graphs for Defence and National Security AI

Ernesto Jiménez-Ruiz (PI), Dave Herron (RA) and Jonathon Dilworth (RA)

The Team (phase 1)

- **Dr Ernesto Jimenez-Ruiz**
  - Senior Lecturer in AI
  - Principal Investigator
- **Dr Dave Herron**
  - PhD in AI
  - Research Assistant
- **Jonathon Dilworth**
  - MSc in AI
  - Research Assistant

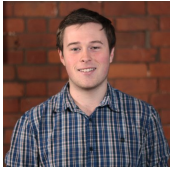

The PI

- More than **20 years** in Semantic Web and **AI research**.
- Co-organiser of the **Ontology Matching workshop** and the Ontology Alignment Evaluation Initiative (OAEI).
- Co-chair of **The Turing Interest Group on Knowledge Graphs**.
- Focus on the design and implementation of **KG-based neurosymbolic AI systems**.
- Part of the **AI Research Centre** at City St George's: `https://aicentre-csg.github.io/`

Collaborators: RAI for Ukraine Program

KG-LLM: Aligning Knowledge Graphs with (Large) Language Models

MENTORS:

Ernesto Jimenez-Ruiz
City St George's, University of London, **UK**

Artur Garcez
City St George's, University of London, **UK**

RESEARCH FELLOWS:

Sviatoslav Lushnei,
UCU,
Ukraine

Dmytro Shumskyi,
UCU,
Ukraine

Severyn Shykula,
UCU,
Ukraine

Mentoring: September 2024 - December 2025

NYU TANDON

CITY ST GEORGE'S UNIVERSITY OF LONDON

ФАКУЛЬТЕТ ПРИКЛАДНИХ НАУК

r/ai

Collaborators

- **University of Lisbon:** Pedro Cotovio (OAEI track), Daniel Faria (alignment repair)
- **Project advisors:** Catia Pesquita (Lisbon), Paul Cripps (dstl).
- **Potentially in Phase 2** (if funded): Navdeep Kaur (Alan Turing Institute).

GUARD project

Motivation: Defence, Intelligence and Security Ontologies

Ontology clusters/subdomains:

- agentic (1)
- ambient intelligence (1)
- context awareness (3)
- cyber security (15)
- digital twins (1)
- information exchange (2)
- information security (5)
- mid-level (12)
- risk management (1)
- robotics (2)
- situation awareness (4)
- smart environments (8)
- upper-level (8)

DISO: `https://github.com/city-artificial-intelligence/diso`

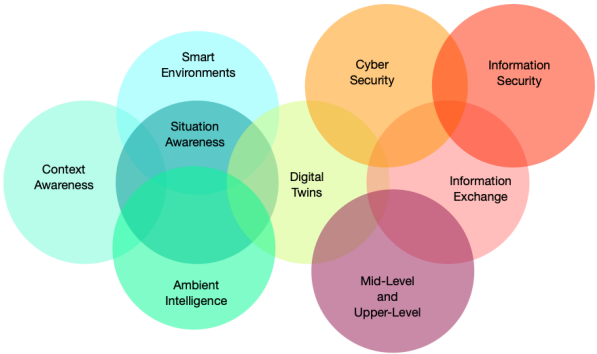

GUARD project overview

Two key challenges:

- **Phase 1: Interoperability and coverage**.
  - Ontology alignment to enhance the interoperability and coverage of ontologies in domains such as defence and national security.
- **Phase 2: Quality**.
  - Scalable reasoning to detect inconsistencies and violations of integrity constraints in a network of knowledge graphs.

Phase 1: interoperability and coverage

Background: ontology alignment (i)

- Ontology alignment critical to integrate diverse data sources, as required in a domain like Defence and National Security.
- Need to include domain experts in applications where accuracy is paramount.
- **Large language models (LLMs)** to complement a domain expert.
- Need for **cost-effective** utilisation of **LLMs as Oracles** for knowledge graph integration.
- Need for **fine-grained analysis** of the LLM contribution.

Background: ontology alignment (ii)

- Basic definition in the OM community.
- An **ontology alignment** $\mathcal{M}$ (or $\mathcal{A}$) is a set of tuples $\langle e_1, e_2, n, \rho \rangle$
  - e_1, e_2 are **entities** in the input ontologies ($e_1 \in \mathcal{O}_1$ and $e_2 \in \mathcal{O}_2$)
  - n a **confidence** value between 0 and 1
  - ρ is the **semantic relationship** between e_1 and e_2
    - OM: subsumption, equivalence, disjointness
    - Life Sciences (SKOS vocabulary): broadMatch, narrowMatch, closeMatch, relatedMatch, exactMatch.

P. Shvaiko, J. Euzenat. Ontology matching: state of the art and future challenges. IEEE Transactions on Knowledge and Data Engineering 2013

Background: ontology alignment (iii)

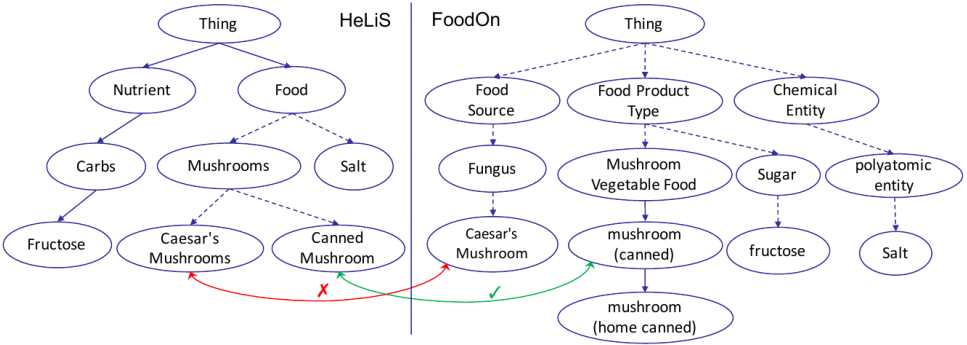

J. Chen et al. Augmenting Ontology Alignment by Semantic Embedding and Distant Supervision. ESWC 2021

Phase 1 delivery (i)

(i) **LogMapLLM**, an open-source LLM-enhanced ontology alignment system built on top of `LogMap`.

(ii) A **comprehensive evaluation** of LogMapLLM over the OAEI datasets using different prompts and large language models.

(iii) A study of how state-of-the-art **ontologies in the defence and national security** domain semantically overlap.

(iv) A new evaluation track in the **Ontology Alignment Evaluation Initiative (OAEI)** 2026 based on the analysed defence and national security domain ontologies.

Phase 1 delivery (ii)

Publication:

- **Large Language Models as Oracles for Ontology Alignment**. In 19th Conference of the European Chapter of the Association for Computational Linguistics (EACL), 2026.

GitHub repositories:

- LogMapLLM (integrated pipeline):
  `https://github.com/city-artificial-intelligence/logmap-llm`
- DISO: Defence, Intelligence and Security Ontologies:
  `https://github.com/city-artificial-intelligence/diso`
- DISO-OAEI (evaluation track):
  `https://city-artificial-intelligence.github.io/diso-oaei/`

Aligning ontologies with LogMap

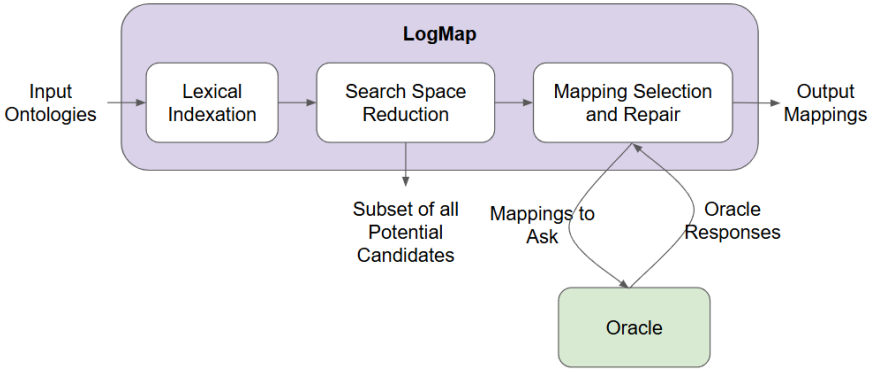

Ernesto Jiménez-Ruiz, Bernardo Cuenca Grau: LogMap: Logic-Based and Scalable Ontology Matching. ISWC (1) 2011: 273-288
Ernesto Jiménez-Ruiz, et al.: Large-scale Interactive Ontology Matching: Algorithms and Implementation. ECAI 2012: 444-449

LogMapLLM: LLMs as Oracles focusing on $\mathcal{M}_{ask}$

$\mathcal{M}_{ask}$ is a set of mappings where `LogMap` is (very) uncertain.

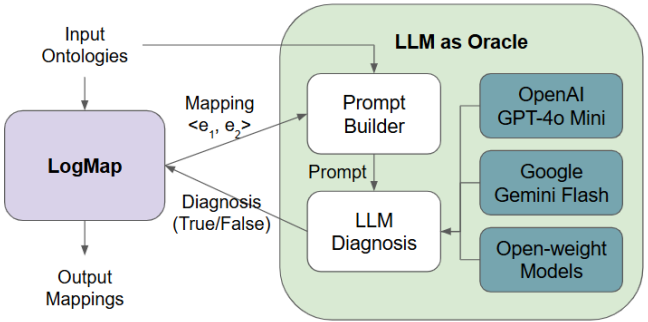

Sviatoslav Lushnei, Dmytro Shumskyi, Severyn Shykula, Ernesto Jiménez-Ruiz, Artur d'Avila Garcez: **Large Language Models as Oracles for Ontology Alignment**. EACL 2026.

LogMapLLM: Ontology-driven prompt builder

- We have evaluated a number of different types of prompts.
- Best results: natural language-friendly prompts with synonyms and only one level of context (*e.g.*, parents).
- Prompts for classes, properties and instances exploit different types of context.

Class mapping prompt builder

Class equivalence by mutual subsumption

We have two entities from different ["biomedical", "conference", ""] ["ontologies", "knowledge graph"]. The first one is "Trapezoid", also known as "Trapezoid bone", "Os trapezoideum". The second one is "Trapezoid", also known as "Trapezium", "Trapezoid"". If something is a "Trapezoid" and a "Short Bone", is it also a "Trapezoid" and a "Quadrilateral"? Respond with ["True" or "False"]["Yes" or "No"].

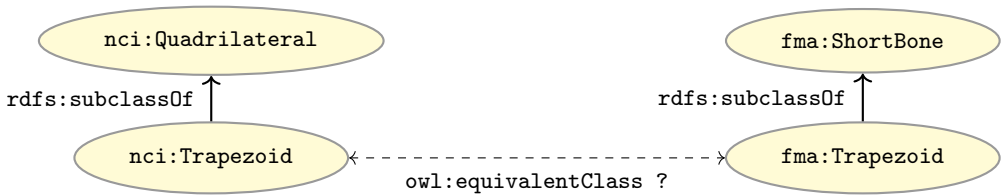

Property mapping prompt builder

Property labels with domain and range

We have two properties from different ["conference", ""] ontologies. The first property is "writePaper", connecting elements of type "Author" to elements of type "Paper". The second property is "is_author_of", connecting elements of type "Person" to elements of type "Contribution". Do these properties represent the same relationship? Respond with ["True" or "False"]["Yes" or "No"].

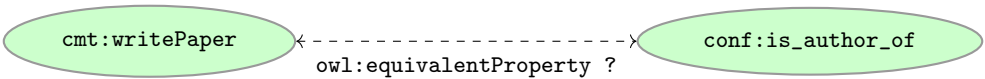

Instance mapping prompt builder

Instance labels with types and attributes

We have two entities from different knowledge graphs. The first is "Anakin Skywalker", which is a "Characters" and is "gender" "Male" and is "species" "Human" and is "homeworld" "Tatooine". The second is "Anakin Solo", which is a "Characters" and is "gender" "Male" and is "species" "Human" and is "homeworld" "New Republic". Do these refer to the same entity? Respond with ["True" or "False"]["Yes" or "No"].

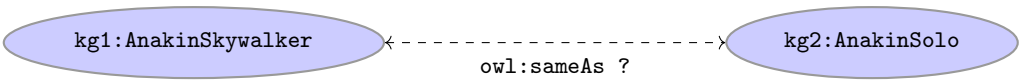

Evaluation: (subset of) OAEI tasks and metrics

OAEI track	Matching task	$\lvert\mathcal{O}_1\rvert$	$\lvert\mathcal{O}_2\rvert$	$\lvert\mathcal{M}^{RA}\rvert$
Anatomy	**Mouse-Human**	2,755	3,313	1,516
Bio-ML	**NCIT-DOID**	15,991	8,516	4,686
	OMIM-ORDO	9,662	9,320	3,721
	SNOMED-FMA.body	34,562	89,180	7,256
	SNOMED-NCIT.neoplas	23,116	20,497	3,804
	SNOMED-NCIT.pharm	29,646	22,387	5,803
Largebio	**FMA-NCI**	79,049	66,919	3,024
	FMA-SNOMED	79,049	122,521	9,008
	SNOMED-NCI	122,521	66,919	18,844

$$\text{Precision} = \frac{|\mathcal{M}^S \cap \mathcal{M}^{RA}|}{|\mathcal{M}^S|},\ \text{Recall} = \frac{|\mathcal{M}^S \cap \mathcal{M}^{RA}|}{|\mathcal{M}^{RA}|},\ \text{F-score} = 2 \cdot \frac{Pr \cdot Re}{Pr + Re}$$

$$\text{Sensitivity} = \frac{TP}{TP + FN},\ \text{Specificity} = \frac{TN}{FP + TN},\ \text{Youden's Index (YI)} = Se + Sp - 1$$

Evaluation: LLM diagnostic capabilities (ii)

- Diagnostic on the **uncertain** mappings $\mathcal{M}_{ask}$ selected by LogMap.
- LLM-based Oracle based on Gemini Flash 2.5

Matching task	$\lvert\mathcal{M}_{ask}\rvert$		LogMap **on** $\mathcal{M}_{ask}$			$\mathrm{Or}_{GF2.5}^{LLM}$ **on** $\mathcal{M}_{ask}$		
	P	**N**	**Se**	**Sp**	**YI**	**Se**	**Sp**	**YI**
Mouse-Human	165	94	1.000	0.000	0.000	0.951	0.744	**0.695**
NCIT-DOID	364	492	1.000	0.000	0.000	0.849	0.697	**0.546**
OMIM-ORDO	172	227	0.343	0.551	-0.106	0.942	0.876	**0.818**
SNOMED-FMA.body	369	619	0.881	0.149	0.029	0.884	0.669	**0.553**
SNOMED-NCIT.neoplas	704	601	0.984	0.067	0.051	0.840	0.593	**0.433**
SNOMED-NCIT.pharm	297	260	0.929	0.065	-0.005	0.848	0.669	**0.517**
FMA-NCI	410	475	0.705	0.726	0.431	0.761	0.684	**0.445**
FMA-SNOMED	831	621	0.941	0.225	0.166	0.480	0.853	**0.333**
SNOMED-NCI	1450	1128	0.887	0.395	0.281	0.846	0.763	**0.609**
Average	529	502	0.852	0.242	0.094	0.822	0.728	**0.550**

Evaluation with local open-weight models

Oracle diagnostic quality on *anatomy* and *bioml* (macro-averaged Youden's J across tasks, yes/no answer format).

Model	**Prompts**				
	syns+pars	syns+pars+sibs	subs (syns)	subs (syns+pars)	deductive
Qwen3-32B	0.489	0.493	**0.543**	0.477	0.536
Qwen3-235B	0.453	0.493	0.537	**0.567**	0.448
Qwen3.5-35B	0.495	0.504	**0.593**	0.568	0.544
Qwen3.5-122B	0.552	0.551	**0.621**	0.597	0.589
Mistral-3.2-24B	0.427	0.430	**0.529**	0.522	0.443
Gemini Flash 2.5	**0.594**	-	-	-	-

Evaluation: Participation in the OAEI 2025 (end-to-end task)

Matching task	LogMap				LogMapLLM			
	Pr	Re	F	Rank	Pr	Re	F	Rank
NCIT-DOID	0.843	0.893	0.867	#4	0.932	0.883	0.907	#2
OMIM-ORDO	0.834	0.456	0.589	#7	0.916	0.476	0.626	#4
SNOMED-FMA.body	0.760	0.569	0.651	#6	0.869	0.561	0.682	#4
SNOMED-NCIT.neoplas	0.763	0.772	0.736	#4	0.821	0.747	0.782	#1
SNOMED-NCIT.pharm	0.932	0.620	0.745	#4	0.979	0.621	0.760	#1
OVERALL	0.826	0.662	0.718	#5	0.903	0.658	0.751	#2

- LLM-based Oracle has a similar performance to an Oracle with an error rate 20%.
- An Oracle with an error rate of 30% would still be better than LogMap.
- `LogMapLLM` ranked #2 in the OAEI Bio-ML 2025 track (LogMap #5).

Few-shot prompts (i)

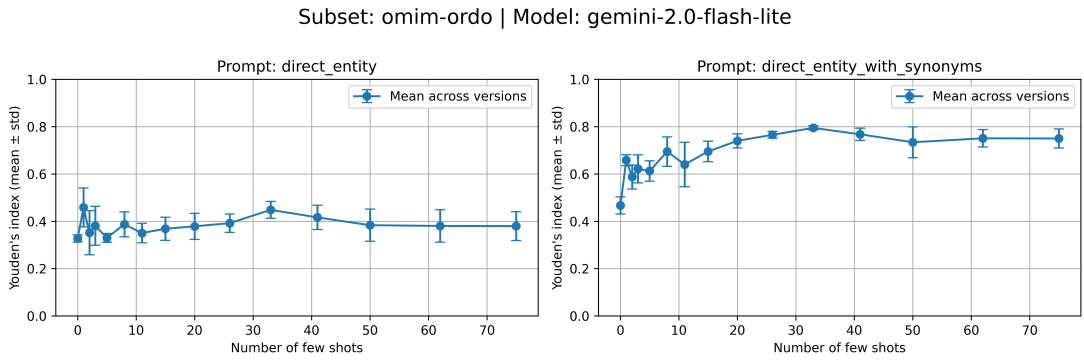

Few-shot prompts (ii)

Subset: human-mouse | Model: qwen|qwen3-vl-8b-instruct

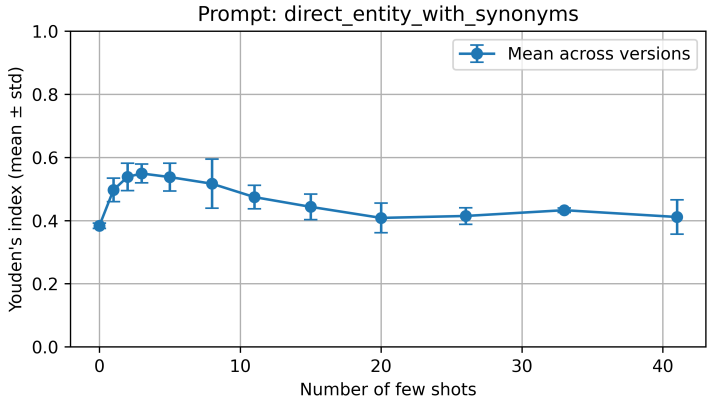

DISO: Towards a network of ontologies for Defence and Security

Defence, Intelligence and Security Ontologies (DISO)

Ontology clusters/subdomains:

- agentic (1)
- ambient intelligence (1)
- context awareness (3)
- cyber security (15)
- digital twins (1)
- information exchange (2)
- information security (5)
- mid-level (12)
- risk management (1)
- robotics (2)
- situation awareness (4)
- smart environments (8)
- upper-level (8)

DISO: `https://github.com/city-artificial-intelligence/diso`

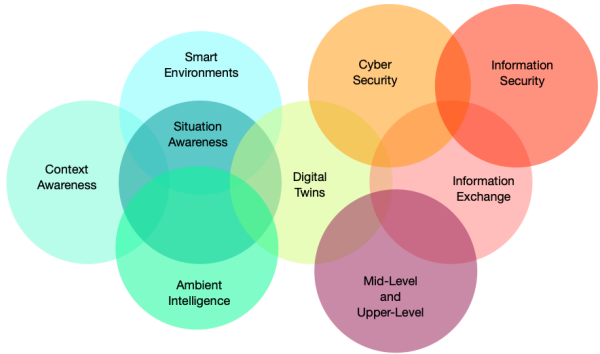

Overlapping among DISO ontologies

- `LogMap` over 58 DISO ontologies → alignments for 1,653 ontology pairs.
- 647 ontology pairs share at least 1 `LogMap` mapping.

	1653 DISO ontology pairs			105 DISO cyber-security ontology pairs		
Alignment Mapping Count	Ontology Pair Count	Ontology Pair Count Cumulative	Cumulative Pair Count As Proportion of Total Pairs	Ontology Pair Count	Ontology Pair Count Cumulative	Cumulative Pair Count As Proportion of Total Pairs
0	1006	1006	0.61	39	39	0.37
1	213	1219	0.74	17	56	0.53
2	107	1326	0.80	13	69	0.66
3	90	1416	0.86	11	80	0.76
4	48	1464	0.89	3	83	0.79
5	29	1493	0.90	3	86	0.82
6	31	1524	0.92	6	92	0.88
7	17	1541	0.93	2	94	0.90
8	14	1555	0.94	1	95	0.90
9	13	1568	0.95	0	95	0.90
10	12	1580	0.96	3	98	0.93

DISO selection for an OAEI task

- Eight pairs selected for a new OAEI track.
- The tasks include class, property, and instance mapings.
- `https://github.com/city-artificial-intelligence/diso-oaei`

Source Ontology	Target Ontology	Source DISO Cluster	Target DISO Cluster	Alignment Mappings	$\mathcal{M}_{ask}$ Mappings	OAEI Track Subdomain	OAEI Track Subdomain Anchor Ontology
UCO	STIX	cybersecurity	cybersecurity	131	43	cybersecurity	
STIX	D3FEND	cybersecurity	cybersecurity	37	31	cybersecurity	STIX
JC3IEDM	mIO!	SA	CA	219	40	SA	
JC3IEDM	Brick	SA	smart buildings	49	69	SA	JC3IEDM
JC3IEDM	FacilityOntology	SA	cco-modules	44	23	SA	
ThinkHome	Brick	smart homes	smart buildings	33	42	smart envir	
Brick	SmartEnv	smart buildings	smart homes	32	45	smart envir	Brick
CityOWL	Brick	smart cities	smart buildings	24	27	smart envir	

SA: Situation Awareness. CA: Context Awareness

Future steps for DISO

- Logic-based assessment of the compatibility of the DISO ontologies (**connection with Phase 2**).
- Synergies with DoD/IC Ontology Working Group (DIOWG): ontology foundry for defence and national security.
- Synergies with Rebecca & Paul (Dstl), and Nigel (GCHQ).
- Input from domain experts about reuse, coverage and alignment needs.
- Extension of DISO with additional relevant ontologies.
- Inclusion of (public) datasets: knowledge graphs.

Phase 2: quality (if funded)

Phase 2: where are we now?

- Scalable alignment repair techniques only on horn-propositional logic fragments.
- Constraint for KG validation defined with SHACL or ShEx languages in addition to the OWL ontology.

Phase 2: where we want to be

- Design of encodings beyond Horn-propositional logic, including the use of OWL 2 RL and Datalog reasoning, for ontology alignment repair.
  - Application to the DISO ontology network.

Phase 2: where we want to be

- Design of encodings beyond Horn-propositional logic, including the use of OWL 2 RL and Datalog reasoning, for ontology alignment repair.
  - Application to the DISO ontology network.
- Interpretation of selected OWL 2 axioms as integrity constraints and application of Datalog reasoning to identify KG violations.
  - Application to domain KGs.

Phase 2: where we want to be

- Design of encodings beyond Horn-propositional logic, including the use of OWL 2 RL and Datalog reasoning, for ontology alignment repair.
  - Application to the DISO ontology network.
- Interpretation of selected OWL 2 axioms as integrity constraints and application of Datalog reasoning to identify KG violations.
  - Application to domain KGs.
- Inclusion of LLMs within the alignment repair and KG correction processes.

Phase 2: where we want to be

- Design of encodings beyond Horn-propositional logic, including the use of OWL 2 RL and Datalog reasoning, for ontology alignment repair.
  - Application to the DISO ontology network.
- Interpretation of selected OWL 2 axioms as integrity constraints and application of Datalog reasoning to identify KG violations.
  - Application to domain KGs.
- Inclusion of LLMs within the alignment repair and KG correction processes.
- Synergies with other funded projects.

Phase 2: KG Alignment Repair.

$$\alpha : \text{PhysicalMilitaryAsset} \sqsubseteq \exists \text{deployed\_at.Location} \qquad (\mathcal{O}_1)$$

$$\mu_1 : \text{PhysicalMilitaryAsset} \sqsubseteq \text{PhysicalDefenceAsset} \qquad (\mathcal{M}, \mathcal{O}_1 \to \mathcal{O}_2)$$

$$\mu_2 : \text{deployed\_at} \equiv \text{deployed} \qquad (\mathcal{M}, \mathcal{O}_1 \to \mathcal{O}_2)$$

$$\beta : \exists \text{deployed}.\top \sqsubseteq \text{Software} \qquad (\mathcal{O}_2)$$

$$\gamma : \text{PhysicalDefenceAsset} \sqcap \text{Software} \sqsubseteq \bot \qquad (\mathcal{O}_2)$$

- This combination of axioms will lead to PhysicalMilitaryAsset $\sqsubseteq \bot$.
- Systems like `LogMap` would fail to capture this logical error.

Phase 2: KG Alignment Repair.

$$\alpha : \text{PhysicalMilitaryAsset} \sqsubseteq \exists \text{deployed\_at.Location} \qquad (\mathcal{O}_1)$$

$$\mu_1 : \text{PhysicalMilitaryAsset} \sqsubseteq \text{PhysicalDefenceAsset} \qquad (\mathcal{M}, \mathcal{O}_1 \to \mathcal{O}_2)$$

$$\mu_2 : \text{deployed\_at} \equiv \text{deployed} \qquad (\mathcal{M}, \mathcal{O}_1 \to \mathcal{O}_2)$$

$$\beta : \exists \text{deployed}.\top \sqsubseteq \text{Software} \qquad (\mathcal{O}_2)$$

$$\gamma : \text{PhysicalDefenceAsset} \sqcap \text{Software} \sqsubseteq \bot \qquad (\mathcal{O}_2)$$

- This combination of axioms will lead to PhysicalMilitaryAsset $\sqsubseteq \bot$.
- Systems like `LogMap` would fail to capture this logical error.
- **Challenge**: increasing expressiveness while preserving scalability.

Phase 2: KG Validation.

$$\alpha : \text{PhysicalMilitaryAsset} \sqsubseteq \exists \text{deployed\_at.Location}$$

- A *PhysicalMilitaryAsset* has at least an associated *Location* (OWA).
- Constraints or shapes to enforce this in the KG (CWA).
- OWL 2 axioms involving cardinalities and ranges can be interpreted as integrity constraints via Datalog rules to identify violations:

$$\text{Things\_deployed\_at\_location}(?a) \leftarrow \text{deployed\_at}(?a, ?l) \wedge \text{Location}(?l)$$

$$\text{Violation}(?a, \alpha) \leftarrow \text{PhysicalMilitaryAsset}(?a) \wedge \textbf{not } \text{Things\_deployed\_at\_location}(?a)$$

Phase 2: KG Validation.

$$\alpha : \text{PhysicalMilitaryAsset} \sqsubseteq \exists \text{deployed\_at.Location}$$

- A *PhysicalMilitaryAsset* has at least an associated *Location* (OWA).
- Constraints or shapes to enforce this in the KG (CWA).
- OWL 2 axioms involving cardinalities and ranges can be interpreted as integrity constraints via Datalog rules to identify violations:

$$\text{Things\_deployed\_at\_location}(?a) \leftarrow \text{deployed\_at}(?a, ?l) \wedge \text{Location}(?l)$$

$$\text{Violation}(?a, \alpha) \leftarrow \text{PhysicalMilitaryAsset}(?a) \wedge \textbf{not}\ \text{Things\_deployed\_at\_location}(?a)$$

- **Challenge**: (i) combining OWA and CWA. (ii) "Suggesting" LLM-based fixes according to the available context (to a human-in-the-loop). (iii) Automatically defining SHACL/ShEx constraints.

Phase 2: KG Validation (another example).

$$\alpha_2 : \text{MilitaryInstallation} \sqsubseteq \forall\, \text{hasThreatLevel}.\, \text{xsd:integer}[\geq 1 \wedge\ \leq 5]$$

– A *MilitaryInstallation* has a threat level between 1 and 5.

$$\text{Things\_with\_incorrect\_threat\_level(?m)} \leftarrow \text{hasThreatLevel(?m, ?tl)} \wedge \text{swrlb:lessThan(?tl, 1)}$$
$$\text{Things\_with\_incorrect\_threat\_level(?m)} \leftarrow \text{hasThreatLevel(?m, ?tl)} \wedge \text{swrlb:greaterThan(?tl, 5)}$$
$$\text{Violation(?m, } \alpha_2) \leftarrow \text{MilitaryInstallation(?m)} \wedge \text{Things\_with\_incorrect\_threat\_level(?m)}$$

Questions

Questions?

Acknowledgements:

- The Alan Turing Institute's Defence and Security Programme
- The RAI for Ukraine program at the Center for Responsible AI at NYU

Contact:

- ernesto.jimenez-ruiz@citystgeorges.ac.uk
- `https://ernestojimenezruiz.github.io/projects/guard/`
- `https://github.com/city-artificial-intelligence/logmap-llm`