



A Neuro-Symbolic System for Querying and Monitoring Copyright Compliance Across National Legislations

Team: John Brüne, Michele Corazza, Generoso Longo, Salvatore Sapienza, Simone Vagnoni
Team Member and Supervisor: Prof. Monica Palmirani



Content

01 Setting the Ground

Context, background and related work

02 Dataset

Legal corpus

03 Methodology

Pipeline and technical implementation

04 Results

Live demo, evaluation

05 Discussion

Limitations and outlook

Brüne, Corazza, Longo, Sapienza, Vagnoni

IRIS 2026 Hackathon

A Neuro-Symbolic System for Querying and Monitoring Copyright Compliance Across National Legislations

John Frederic Brüne Michele Corazza Generoso Longo
Salvatore Sapienza Simone Vagnoni

LAST-JD Programme, Università di Bologna, Italy

Supervision: Prof. Monica Palmirani

May 2026

Abstract

Checking whether national copyright law actually complies with WIPO treaty obligations is still largely manual work. We built a neuro-symbolic web application that does this automatically for nine jurisdictions. A large language model handles the multilingual question, and a knowledge graph built from Akoma Ntoso XML legal documents supplies the answer. A user can ask in any language; the system finds the relevant provision, links it back to the official source, and decides Berne Convention compliance with no LLM in the loop. Beyond compliance assessment, the system provides two features aimed at practical legal research. First, every answer includes authenticated source attribution through links derived from FRBR metadata contained in the original legal documents. Second, the system supports point-in-time navigation across 68 historical versions of the Italian Copyright Act, allowing users to retrieve the version in force on a selected date. We manually validated all encoded copyright-duration values and resulting compliance classifications across the nine jurisdictions. Source code is at https://github.com/Stocastico96/ReMeP_Hackathon_26.

Keywords: neuro-symbolic AI, knowledge graph, Akoma Ntoso, copyright law, treaty compliance

Team

Monica Palmirani

Legal Informatics



Michele Corazza

Computational Linguistics



Generoso Longo

Computer Engineering



Salvatore Sapienza

Legal Informatics



Simone Vagnoni

Computer Engineering



John F. Brüne

Applied Statistics



Setting the Ground

Supranational institutions must monitor and compare legislation across +190 member states to evaluate compliance and consistency



Current bottlenecks:



- Manual comparison is extremely slow (requires legal experts for each jurisdiction)
- Knowledge sharing across jurisdictions is limited
- Heterogeneous legal sources make automation difficult:
 - Different legal tradition (common law, civil law)
 - Different formats (PDF, HTML, XML, scanned documents)
 - Different structures (articles, rules, paragraphs, sections)
 - ...

→ We need a **reliable, explainable, scalable** method




WIPO members  Members  Non-members

Why not?



Problems to consider

- Semantics (legal language and multilingualism) [Ajani et al., 2017, Zilli et. al. 2025]
 - Needle in a haystack [Goebel et al., 2024, Corazza et. al 2024]
 - Missing normative references ("see section 2") [Sannier et al., 2017, Corazza et. al. 2025]
- Temporal information and temporal dynamicity of the provisions
[Palmirani et. al. 2006, de Oliveira Lima et al., 2008]



Goals

1. Compare the WIPO regulation with national legislation
2. Extract legal knowledge from the text and model it as a Knowledge Graph
3. Visualise the similarities and the discrepancies
4. Deliver a temporal modelling of the provisions




Dataset

- **Nine jurisdictions, all encoded as Akoma Ntoso XML**
 - DE, FR, IT, CH, UK, US, CA, NZ, and the EU Directive 2006/116/EC
- **Italy: 68 historical versions of Legge 633/1941, 1941 to 2025, from Normattiva**
 - example jurisdiction with full version history
- **Ontology: Custom Version of IPROnto (DMAG, Universitat Politècnica de Catalunya)**
 - extended with a legal: property vocabulary for SPARQL querying
- **Knowledge graph:**
 - built from a curated CSV, one row per jurisdiction; three rule types across nine jurisdictions, 27 core nodes


Methodology

Legal Data Science Pipeline inspired by
Hybrid Artificial Intelligence Methodology for Legal Analysis
(HAIMLA)

[2024, Palmirani, Sapienza and Ashley]



Pipeline



Neuro-symbolic steps

1. Query enrichment

(neural) Gemini 2.0 Flash reads a question in any language and extracts jurisdiction, legal topic, and time frame.

3. AKN passage extraction

(symbolic) The article reference from the graph locates the provision in the source XML, which is rendered with the supporting span highlighted.

5. Point-in-time navigation

(symbolic) For Italian documents, the version in force on a chosen date gets detected and reloaded.

2. KG retrieval

(symbolic) The three parameters select one of three SPARQL SELECT templates over an in-memory rdflib graph.

4. Source attribution


(symbolic) FRBR Work URI, applicability date, ELI identifier, and authenticity flag, are read from the AKN metadata.

6. Answer generation or compliance check


(neural): the LLM phrases an answer from KG evidence.
(symbolic): the verdict is computed/reasoned with no LLM in the loop.

Ontology

Example Ontology



Example Graph



Ontology, what is difficult?








Demo 1: Q&A Mode


- Ask in any language, for example “quando scade il diritto d'autore in Francia?”
- Answer cites the article that came from KG evidence
- The source AKN article is rendered alongside, with the supporting span highlighted
- A FRBR meta bar shows the authenticity badge, the applicability date, and a link to the official source.
- For Italian documents a date picker resolves the version in force on the selected date.



The screenshot displays the ReMeP — IP Law Assistant interface. The top navigation bar includes the title and links to Neuro-Symbolic KG, AKN Documents, and WIPO Compliance. The main interface is divided into two panels. The left panel, titled 'Q&A', shows a user question: 'What is the copyright term under the Berne Convention?'. Below the question, an 'ANSWER' section provides the response: 'Life of author + 50 years'. This answer is supported by a snippet from the 'Berne Convention for the Protection of Literary and Artistic Works', Article 7(1), which states: 'The term of protection lasts for the life of the author plus 50 years after their death (Art. 7 Berne Convention)'. The right panel, titled 'Berne Convention for the Protection of Liter...', shows the full text of Article 7, with the relevant paragraph (1) highlighted in yellow: '(1) The term of protection granted by this Convention shall be the life of the author and fifty years after his death.' The interface also includes a 'Compliance Check' status indicator, a 'Full Document' link, and an 'Ask' button at the bottom.

Demo 2: Compliance Check


- Pick a jurisdiction; the system checks its copyright duration against the Berne floor of 50 years post mortem auctoris (Art. 7).
- Purely symbolic: no LLM in the loop, so the result is reproducible and traceable back to the encoded facts.
- Knowledge Graphs are built on the basis of the underlying (and legally evaluated) ontology.
- All nine jurisdictions return the correct verdict. New Zealand, at exactly 50 years, is compliant with a margin of zero.




The screenshot displays the ReMeP — IP Law Assistant interface. The top navigation bar includes the title "ReMeP — IP Law Assistant" and links to "Neuro-Symbolic KG", "AKN Documents", and "WIPO Compliance". The main interface is divided into two panels. The left panel, titled "Q&A", features a "Compliance Check" tab. It contains dropdown menus for "COUNTRY" (set to "Germany") and "TREATY" (set to "Berne Convention (Art. 7) — copyright d"). Below these is a "Check Compliance" button. A list of countries with their ISO codes is shown: Germany (DE), United Kingdom (UK), United States (US), France (FR), Italy (IT), Switzerland (CH), European Union (EU), Canada (CA), and New Zealand (NZ). A green banner indicates "GERMANY — BERNE CONVENTION COMPLIANCE". Below this, a comparison table shows "GERMANY" with a "70 years" duration (70 years pma, triggered by death of author, § 64 UrhG) and "BERNE CONVENTION" with a "50 years" duration (minimum required, triggered by death of author). The right panel displays the "German Urheberrechtsgesetz (UrhG)" with a link to "Article" and "Full Document". It shows "§ 64 UrhG" with the text: "Das Urheberrecht erlischt siebenzig Jahre nach dem Tode des Urhebers."

Some preliminary insights: Knowledge Graphs

Knowledge Graph: Swiss Copyright Law (URG)



Knowledge Graph: German Copyright Law (UrhG)



Node type

- Article
- Country
- Duration
- Law
- Right
- Treaty

Relation

- compliesWith
- contains
- expresses
- governs
- hasDuration

Future Work



KG Embedding

Compare jurisdictions structurally, instead of one predicate at a time

Richer IPR ontology

Beyond duration: fair use (US), fair dealing (UK, CA, AU, SG), EU exception regimes


Explicability via legal design

how the reasoning path: which triples, which provision, which comparison

Further Validation with legal experts

A substantially larger query set, beyond the nine hand-checked jurisdictions

THANK YOU FOR LISTENING!



FutureData4EE (Grant. Agreement n. 101126733) Co-Funded by the European Union. Views and opinions expressed are however those of the author only and do not necessarily reflect those of the European Union or REA. Neither the European Union nor the granting authority can be held responsible for them.