IYQ ELSA Online Workshops

- Europe Edition-

QSL

part of TUM THINK TANK

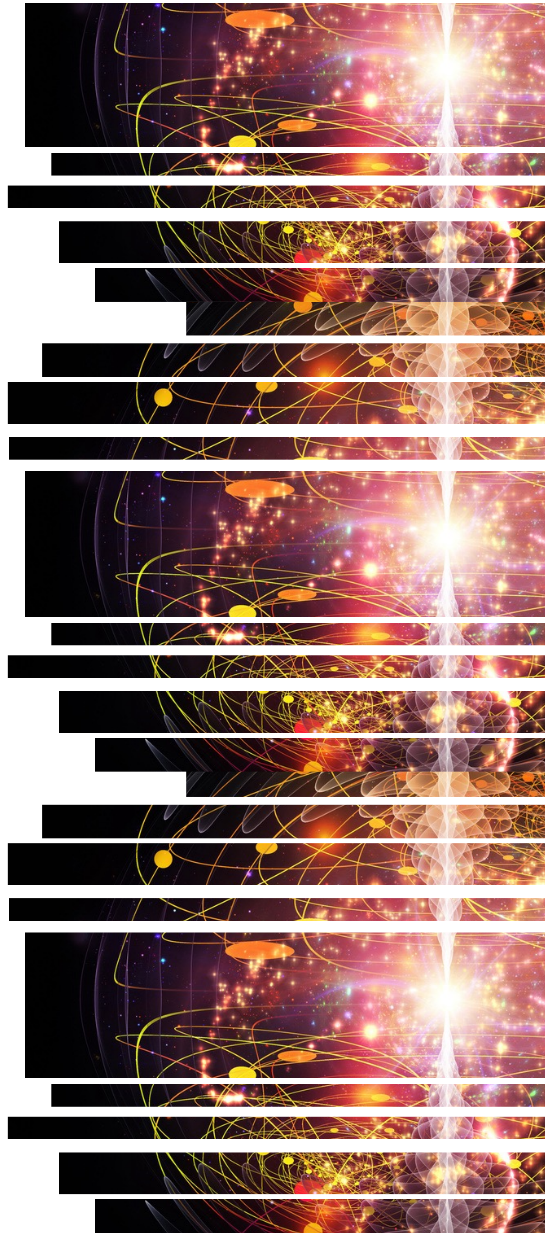

TABLE OF CONTENT

MOTIVATION

- Quantum natural language processing (QNLP) combines quantum computing with compositional models of language
- While language is in reality not just depending on the different compositions and single language tokens, it is highly influenced by the composition and its context
  - Recent big LLM models such as BERT, Transformers etc. are improving specifically on the base of a bigger context window size
- ParlaMint underscores the urgency of interpretability: parliamentary debates influence laws and public perception, yet current AI tools often operate opaquely.

→ Need of Explainability and Interpretability, specifically for sensitive data and data used for decision making as well as language-wise complex texts.

Problem: Limitations of Current Explainability Tools in Political Contexts

Post-hoc approximation:

- LIME (Local Interpretable Model-Agnostic Explanations) and SHAP (SHapley Additive exPlanations) approximate complex models through local linear regressions or Shapley value decompositions

Stability issues in LIME:

- Research shows that small, semantics-preserving perturbations (e.g., replacing *“gets”* with *“becomes”*) can drastically reorder LIME’s feature importances without changing the model’s prediction - indicating high sensitivity and low reliability and stability under perturbation

Low fidelity on complex models:

- LIME performs acceptably for simple linear models but fails on non-linear architectures, where it cannot capture intricate feature interaction - producing unstable, low-fidelity explanations

Computational cost of SHAP:

- SHAP provides more stable attributions but at a heavy computational cost, limiting its scalability for real-time or large-scale text corpora. However, specifically political texts and debates need context and additional information to be worked with.

Conceptual and pedagogical limitations:

- Both methods often generate explanations that are pedagogically irrelevant or misleading, partly due to independence assumptions that rarely hold in

→ These limitations highlight the need for **principled interpretability frameworks**—such as QNLP-based explainability—that can account for **semantic compositionality, institutional roles, and rhetorical nuance** in political discourse.

THE TESTBED I

Comprehensive documentation:

• The ParlaMint project compiles parliamentary speeches from 29 European national and regional parliaments, totaling over one billion words, enriched with linguistic annotations and metadata such as speaker role, party affiliation, and debate chronology

•Its **institutional structure** demands sensitivity to **procedural roles** (e.g., questions, amendments, committee interventions) that define legislative dialogue, but also offers to test if these characteristics correlates with the level of explainability

Legal and normative language:

•Parliamentary discourse often employs **legal and deontic vocabulary** (e.g., *must, shall, may*), whose meanings depend on **institutional and juridical context**.

•Capturing this nuance is critical for accurate semantic interpretation and explainability.

Empirical example — Slovenian Parliament:

•Classical and transformer-based models trained on Slovenian parliamentary transcripts were used to predict MPs' ideological leanings.

•**SHAP analysis** found that:

- Right-wing MPs frequently used terms linked to **illegality** and **sovereignty**.
- Left-wing MPs emphasized **human rights** and **solidarity**

THE TESTBED II

Relevance to Quantum Natural Language Processing (QNLP):

- The **compositional and hierarchical structure** of parliamentary discourse aligns naturally with **categorical and tensor-network representations** used in QNLP.
- Each debate or utterance can be represented as a **string diagram**, where grammatical composition mirrors the **monoidal composition** of morphisms in compact closed categories.
- Institutional roles (e.g., *Commission*, *Council*, *Parliament*) become interpretable **subspaces or projectors** in the semantic Hilbert space, enabling **role-traceable quantum states**.

Computational Challenges for Benchmarking:

- **Multilingual semantics:** Requires **cross-lingual functor alignment**, mapping syntactic categories across 20+ languages to a shared semantic space.
- **Procedural semantics:** Legislative actions (motions, votes, amendments) introduce **nontrivial modality operators**—natural candidates for **quantum channels** or **CP maps** acting on propositions.
- **Legal modality & deontic logic:** Frequent use of *shall, may, must* demands **explicit operator modelling** to preserve obligation vs. permission distinctions.

TESTING CRITERIA

1. Modality

“We may consider adjusting subsidies, but we must protect farmers.”

2. Nested Propositions and Semantic Transparency

“The Council insists that the Commission must report on climate targets.”

3. Institutional Role Tracability

“The Commission shall propose a directive, and Parliament may amend it.”

4. Interpretability Roubistness (across policy domains)

“Parliament urges that the Commission should ban harmful pesticides.”
&
“Parliament urges that the Commission should ban unfair tariffs.”

MATHEMATICAL OPERATIONALISATION I

Attach each token i (or phrase) a role label $r(i) \in \mathcal{R} =$ {\textsc*Commission*, \textsc*Parliament*, \textsc*Council*, ...}. Define a role subspace $H_r \subseteq H_{\mathrm{ent}}$ with projector P_r. If an entity embedding $\mathbf{e}_i \in H_{\mathrm{ent}}$ is present in the diagram via a Frobenius "copy" node, the *role salience* in the final sentence state ρ_s is:

$$\alpha_r(s) \;=\; \mathrm{Tr}\big((P_r \otimes I)\,\rho_s\big).$$

For a specific *role–predicate* pair (r, ϕ) (e.g., whether \textsc*Commission* is agent of "proposes"), define a predicate projector Q_ϕ on the appropriate predicate space and

$$\alpha_{r,\phi}(s) \;=\; \mathrm{Tr}\big((P_r \otimes Q_\phi \otimes I)\,\rho_s\big).$$

Path-sensitive attribution (diagram paths)

Let $\Pi(r \rightsquigarrow S)$ be all grammar-consistent paths from the role wire to the sentence output. Define

$$\mathrm{PathAttr}_r(s) \;=\; \sum_{\pi \in \Pi(r \rightsquigarrow S)} w(\pi)\,\mathbb{E}_\pi[M],$$

MATHEMATICAL OPERATIONALISATION II

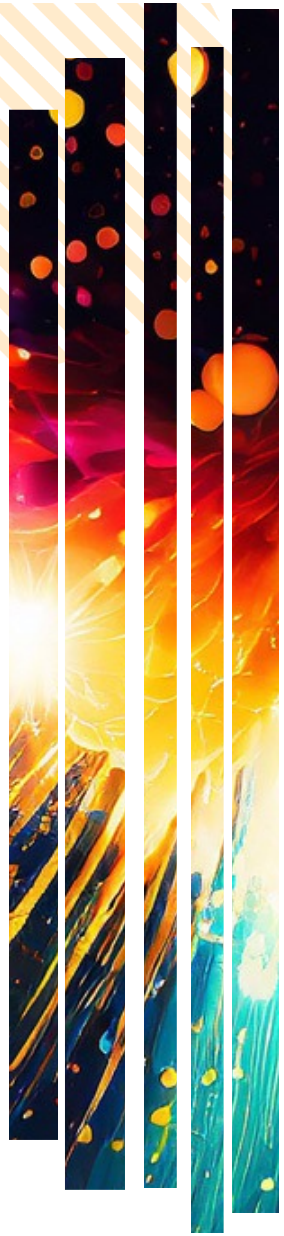

Setup

Let $\mathcal{G}$ be the pregroup grammar (or a compact closed grammar category) and $\mathcal{F} : \mathcal{G} \to \mathbf{FHilb}$ (or $\mathbf{CPM}(\mathbf{FHilb})$) a strong monoidal functor.
For a parsed sentence s with derivation (string diagram) $d_s : X_1 \otimes \cdots \otimes X_m \to S$, the sentence meaning is

$$\mathbf{m}(s) \;=\; \mathcal{F}(d_s)\big(\mathbf{v}_1 \otimes \cdots \otimes \mathbf{v}_m\big) \in \mathcal{F}(S),$$

or as a density operator ρ_s under CPM.

Suppose s contains a *nested proposition* p (e.g., an embedded clause) with sub-derivation d_p. Its local meaning is

$$\mathbf{m}(p) = \mathcal{F}(d_p)(\otimes_{i \in I_p} \mathbf{v}_i).$$

$$\widehat{\mathbf{m}}(s \mid p) \;=\; R_{p \to s}\big(\mathbf{m}(p) \otimes \otimes_{j \notin I_p} \mathbf{v}_j\big).$$

FUTURE OUTLOOK AND FIRST FINDINGS

- Future Contributions will specifically include systemc features of the political system
- Information theoretical aspects
- The role of Causality needs to be researched and taken into account
- Nested propositions work better (up to 5% difference in accuracy
- For testing even more specific annotations are needed

Framework Proposition

1. **Datasets & Parsers**
2. **Functor & Representations**
3. **Primary Scores**
4. **Quantum-Level Diagnostics**
5. **Controls & Baselines**
6. **Reporting**

SOURCES

Bankova, D., Sadrzadeh, M., & Coecke, B. (2018). *Graded entailment for compositional distributional semantics*. Journal of Logic and Computation, 28(2), 371–401.

Baroni, M., Bernardi, R., & Zamparelli, R. (2014). *Frege in space: A program for compositional distributional semantics*. LiLT, 9.

Coecke, B., & Meichanetzidis, K. (2020). *Foundations for near-term quantum natural language processing*. arXiv:2005.04147.

Coecke, B., Sadrzadeh, M., & Clark, S. (2010). *Mathematical foundations for a compositional distributional model of meaning*. Linguistic Analysis, 36, 345–384.

Coecke, B. (2024). *On quantum interpretability* (position/overview). (Use most recent version you cite.)

Meichanetzidis, K., Toumi, A., de Felice, G., Coecke, B., & Vicary, J. (2020). *Grammar-aware quantum natural language processing*. arXiv:2012.03755.

Selvaraju, R. R., et al. (2017). *Grad-CAM: Visual explanations from deep networks via gradient-based localization*. ICCV. (Classical XAI comparator.)

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). *"Why Should I Trust You?": Explaining the predictions of any classifier*. KDD (LIME).

Lundberg, S. M., & Lee, S.-I. (2017). *A unified approach to interpreting model predictions*. NeurIPS (SHAP).

von Wright, G. H. (1951). *Deontic logic*. Mind, 60(237), 1–15. (Modal grounding for MUST/SHALL/MAY.)

Entman, R. (1993). *Framing: Toward clarification of a fractured paradigm*. Journal of Communication, 43(4), 51–58. (Rhetorical framing background.)

Contact

Fabienne Marco
Head of Quantum Social Lab

Richard-Wagner-Straße 1
80333 München

E-Mail:
Fabienne.Marco@tum.de

Instagram:
@quantumsociallab
@quantumfab