

# Prediction of Unconstitutional Index of Italian Bills

Monica Palmirani<sup>1,\*</sup>, Michele Corazza<sup>1</sup>, Generoso Longo<sup>1</sup>, Salvatore Sapienza<sup>1</sup>,  
Pier Fancesco Bresciani<sup>1</sup> and Faria Ferooz<sup>1</sup>

<sup>1</sup>University of Bologna, ALMA-AI, via Galliera 3, 40121, Bologna, Italy

## Abstract

In modern democracy, the Constitutional Court guarantees that the legislative corpus is compliant with the Constitutional norms and principles. For this reason, during the drafting of the legislative bill, the institutions (Chamber of Deputies, Senate, Government, Ministries, and Staff of the President) monitor this phenomenon in order to avoid clashing with the Constitution and also to prevent institutional conflict of competences. This paper presents a proof of concept that aims to calculate the index of presumption of unconstitutionality of the bill of the Chamber of Deputies of Italy. We use a hybrid AI pipeline using different datasets in LegalDocML format (Akoma Ntoso XML) and fostering the embedding AI techniques.

## Keywords

Akoma Ntoso, AI and Law, Natural Language Processing, Hybrid AI, Legislative quality

## 1. Introduction

In modern democracy, the Constitutional Court guarantees that the legislative corpus is compliant with the Constitutional norms and principles. For this reason, during the drafting of the legislative bill, the institutions (Chamber of Deputies, Senate, Government, Ministries, and Staff of the President) monitor this phenomenon in order to avoid clashing with the Constitution and also to prevent institutional conflict of competences. This paper presents a proof of concept that aims to calculate the index of presumption of unconstitutionality of the bill of the Chamber of Deputies of Italy (now called PDL). We apply a Hybrid AI pipeline to different datasets in LegalDocML format (Akoma Ntoso XML) and leveraging the embedding AI techniques. Some criteria have been considered coming from the legal theory analysis. The similarity approach alone is not enough, and scholars of constitutional and legal theory provide a robust framework to implement the algorithm capable of predicting the tentative index of unconstitutionality. Many other works in Legal Informatics domain simply calculate the similarity index between the bill and the judgments of the Constitutional Court. We introduce a hybrid methodology based on multi-factors apart the text of the Constitutional Court decisions. We consider the following:

- the temporal aspects of the Constitutional decisions and of the laws cited in the judgments;
- the qualified legal annotation made by other judges (e.g., summary and catchwords);
- the content of the legal acts cited in the bill;
- the relevance of the citations in the judgements;

*Ital-IA 2025: 5th National Conference on Artificial Intelligence, organized by CINI, June 23-24, 2025, Trieste, Italy*

\*Corresponding author.

✉ monica.palmirani@unibo.it (M. Palmirani); michele.corazza2@unibo.it (M. Corazza); generoso.longo2@unibo.it (G. Longo); salvatore.sapienza@unibo.it (S. Sapienza); pierfrancesco.bresciani@unibo.it (P. F. Bresciani); faria.ferooz2@unibo.it (F. Ferooz)

🌐 <https://www.unibo.it/sitoweb/monica.palmirani> (M. Palmirani); <https://www.unibo.it/sitoweb/michele.corazza2> (M. Corazza); <https://www.unibo.it/sitoweb/generoso.longo2/> (G. Longo); <https://www.unibo.it/sitoweb/salvatore.sapienza> (S. Sapienza); <https://www.unibo.it/sitoweb/pierfrancesco.bresciani> (P. F. Bresciani); <https://www.unibo.it/sitoweb/faria.ferooz2> (F. Ferooz)

🆔 0000-0002-8557-8084 (M. Palmirani); 0000-0002-7288-6635 (M. Corazza); 0009-0003-2687-5884 (G. Longo); 0000-0002-5429-5217 (S. Sapienza); 0000-0002-6889-1618 (P. F. Bresciani); 0000-0002-8150-7852 (F. Ferooz)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

- the explicability of the results to guarantee the application of the AI Act <sup>1</sup> because this use-case is classified at high-risk.

## 2. Related Work

In recent years the application of NLP techniques to the legal domain has been one of the main trends in the literature. In the context of Italian legal documents, Italian-Legal-BERT [1] Italian is a family of models based on BERT which are either trained from scratch or fine-tuned on Italian legal documents.

In terms of contributions on the Italian Constitutional Court, in [2] the authors use semantic similarity to produce a classification of court decisions according to the commissions used in the Italian Chamber of Deputies, in addition to performing the same task for other types of legal documents (Italian bills, Italian laws).

In [3] a dataset of summaries of decisions from the Italian Constitutional Court is described, and it is used to perform experiments that show how segmentation-based techniques achieve better results than sequence-to-sequence approaches.

## 3. Methodology and Dataset

In order to cope with the goal, we have considered different parameters and datasets, all transformed in Akoma Ntoso to foster the descriptive information of the XML (e.g., temporal data, normative references). We consider the following criteria coming from the legal theory analysis made with constitutional law scholars:

- The Constitutional Court decisions, including the metadata concerning the types and the results of the decision. Some types of decisions (e.g., conflict between State and Regions) have been filtered because not relevant for this use-case (see Table 1 for the list of relevant types of judgments).
- We annotated the types of results as either “positive” (no violation of the constitution) or “negative” (some type of violation). Then, some typologies of results (e.g., we have excluded all the decisions “inadmissible” or “ceased topic”) which could not be categorized have also been deleted because unnecessary for the goal (see Table 2 for a list of the relevant results);
- The abstract (now called *massima*) drafted by the Constitutional Court judges that includes the main legal arguments of the court decision;
- The so called “neretti” or “catchwords” that are in brief the main legal concepts used by the Court for the legal argumentation and reasoning (*ration decidendi*);
- The quality of the normative and case-law references, if they are playing a positive or negative role in the judgment argumentation;
- The link to the normative references, and according with the point-in-time principles, we have selected the laws that are now abrogated. This information is offered to the end-user to evaluate the relevance or not of the case;
- We have classified the PDLs using the Committees descriptions and competences of the Chamber of Deputies. The Committees are 14 and defined by law with topics. <sup>2</sup>.

To compute a meaningful similarity between PDLs and judgments from the Italian Constitutional Courts, we use a multi-factored similarity, based on embeddings:

- PDL content and judgment content similarity;
- Committee classification and “neretti” classification similarity;
- PDL content and *massime* content similarity.

In order to produce these similarities, we first have to generate embeddings that encode the PDLs and the Corte Costituzionale judgment. In our preliminary experiments, we used the BGE-M3 model [4] to create all our embeddings.

<sup>1</sup><http://data.europa.eu/eli/reg/2024/1689/oj>

<sup>2</sup><https://www.camera.it/leg18/48>

GIUDIZIO DI ACCUSA	GIUDIZIO DI LEGITTIMITÀ COSTITUZIONALE IN VIA INCIDENTALE
GIUDIZIO DI LEGITTIMITÀ COSTITUZIONALE IN VIA PRINCIPALE	GIUDIZIO DI LEGITTIMITÀ COSTITUZIONALE IN VIA PRINCIPALE + GIUDIZIO DI LEGITTIMITÀ COSTITUZIONALE IN VIA INCIDENTALE
GIUDIZIO PER CONFLITTO DI ATTRIBUZIONE TRA POTERI DELLO STATO	GIUDIZIO SU CONFLITTO DI ATTRIBUZIONE TRA ENTI + GIUDIZIO DI LEGITTIMITÀ COSTITUZIONALE IN VIA INCIDENTALE
GIUDIZIO SU CONFLITTO DI ATTRIBUZIONE TRA ENTI + GIUDIZIO DI LEGITTIMITÀ COSTITUZIONALE IN VIA PRINCIPALE	GIUDIZIO SULL'AMMISSIBILITÀ DEI REFERENDUM
QUESTIONE INCIDENTALE DI LEGITTIMITÀ COSTITUZIONALE	

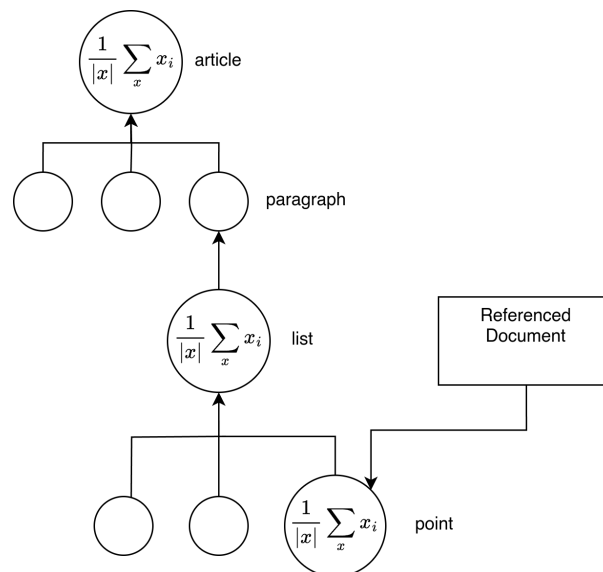
**Table 1**

The types of judgments which are relevant for our analysis

Result	Annotation
ammissibilità referendum	positive
ill. cost. conseg. ex art. 27 legge n. 87/1953	negative
ill. cost. parziale conseg. ex art. 27 legge n. 87/1953	negative
illegittimità costituzionale	negative
illegittimità costituzionale parziale	negative
manifesta infondatezza	positive
non fondatezza	positive

**Table 2**

The relevant results of the judgments and their annotation



**Figure 1:** The tree aggregation strategy for PDLs

### 3.1. Embeddings

In order to generate embeddings for the PDLs, we use a hierarchical approach which leverages the structure of these legal documents as represented in Akoma Ntoso XML to produce a vector representation for each article in the document. This approach works by finding “leaf” nodes inside the XML tree and producing embeddings for their textual content, then their parent’s representation is derived from the mean of their children, and so on.

Additionally, this approach uses normative references as well, meaning that the vector representation of a given leaf node is averaged with a mean vector obtained from all the referenced documents. A more formal description of this approach can be found in [2]. The same approach adopted for the PDLs is used for the Corte Costituzionale judgments, where we substitute the articles with the decisions found inside each document.

In addition to the articles and decisions from PDLs and judgments, we also produce sentence embeddings for the title of each PDL and introduction of each judgment. For both of these cases, we simply apply an embedding model to their text. For neretti, we modify them to be semicolon-separated and feed them to our embedding model. A similar approach was adopted for the Chamber of Deputies commissions, which are represented as a string describing them, which contains their titles (shown in Table 3), as well as their semicolon-separated description as presented in the Circolare del Presidente della Camera (16 ottobre 1996, n. 3), the official document that regulates the matters of competence for each of them.

Affari esteri e comunitari
Difesa
Bilancio, tesoro e programmazione
Finanze
Cultura, scienza ed istruzione
Ambiente, territorio e lavori pubblici
Trasporti, poste e telecomunicazioni
Attività produttive, commercio e turismo
Lavoro pubblico e privato
Affari sociali
Agricoltura
Politiche dell'Unione Europea

**Table 3**  
Italian Chamber Committees

Finally, the massime embeddings are obtained from their unaltered textual content.

### 3.2. Classifying PDLs by Committee

Since not all PDLs are assigned a Committee, we opted to produce an unsupervised classification which assigns each PDL a Committee using their semantic similarity. In order to obtain this result, we produce an average between the cosine similarity of the title of each PDL and each Committee, with the mean article - committee similarity. Then, we select the highest value and classify each document with its Committee. This allows us to associate the committee with the document, and to produce a comparison with the neretti.

### 3.3. Distances

The semantic similarity between the PDLs and the judgments is produced from two values:

- The mean cosine similarity of each article embedding from the PDLs with each decision in the judgments;
- The similarity between the introduction of the judgments and the title of the PDLs

These values are then averaged in order to obtain a quantity  $s_{pj}$  which represents the similarity between a PDL and a judgment. In order to compare the committee with the neretti, we just produce the cosine similarity between their embeddings to produce  $s_{cn}$ . Finally, the similarity between the PDL and the massime is obtained from the mean cosine similarity between each article embedding of a PDL and each massima embedding associated with a judgment. We call this value  $s_{pm}$ . Finally, the overall similarity between a PDL and a judgment is given by the following equation:

$$S = \frac{\lambda_1 s_{pj} + \lambda_2 s_{cn} + \lambda_3 s_{pm}}{\lambda_1 + \lambda_2 + \lambda_3} \quad (1)$$



**Figure 2:** The user interface that returns the results.

Where  $\lambda_1, \lambda_2, \lambda_3$  are used to tune the relative importance of all three similarities. We are still in the process of tuning these values.

### 3.4. Prediction

Since our goal is to predict whether a given PDL might be unconstitutional, we then use a list of the 10 most similar judgments for a given PDL and use their outcome to inform users about potential issues with the bill. In particular, we have compiled a categorization of types of judgments (“tipo procedimento”) and the outcomes (“dispositivo”) which can be used to produce a meaningful categorization of the judgments. Additionally, we categorized the outcomes as either “positive” (the judgment did not find issues of unconstitutionality) or “negative” (there were issues). Finally, we produce a measure of risk as the fraction of judgments between the 10 most similar which contain at least one “negative” outcome.

Another issue when considering precedents is the presence of laws and principles that have been

abrogated since the judgment referenced them. In our predictive approach, we do not exclude these judgments from the computation, but we do flag the appropriate judgments so that the user can assess whether the final conclusions of each judgments are still relevant or whether the fact that some of the normative references point to abrogated portions of documents might, even partially, invalidate the conclusions.

## 4. Future Work

We have installed this preliminary work inside of the portal of testing <sup>3</sup> and we have started the internal evaluation. The user interface has been designed to permit legal experts to perform query, evaluate the results, see the relevant Constitutional Court decisions offered in support of the prediction of unconstitutionality index.

A set of specific test-cases has been elaborated by the legal expert team to evaluate the approach, the methodology, and the algorithm results. The test-case is composed by:

1. PDL that has particular disputable topic in the legal literature;
2. list of the Constitutional Court decisions related to the case divided in positive and negative;
3. analysis of the output of the system (FP and TN).

## Acknowledgment

This project is conducted with the support of the European Commission funds within ERC HyperModelLex. Grant agreement ID: 101055185. This project is co-funded by the European Union - NextGenerationEU under the National Recovery and Resilience Plan (PNRR) - Mission 4 Education and research - Component 2 From research to business - Investment 1.1 Notice PRIN 2022 (DD N. 104 del 02/02/2022), title Smart Legal Order in DigiTal Society (SLOTS), proposal code 2022LRL2C2 - CUP J53D23005610006.

## Declaration on Generative AI

The authors have not employed any Generative AI tools.

## References

- [1] D. Licari, G. Comandè, ITALIAN-LEGAL-BERT: A Pre-trained Transformer Language Model for Italian Law, in: D. Symeonidou, R. Yu, D. Ceolin, M. Poveda-Villalón, D. Audrito, L. D. Caro, F. Grasso, R. Nai, E. Sulis, F. J. Ekaputra, O. Kutz, N. Troquard (Eds.), Companion Proceedings of the 23rd International Conference on Knowledge Engineering and Knowledge Management, volume 3256 of *CEUR Workshop Proceedings*, CEUR, Bozen-Bolzano, Italy, 2022. URL: <https://ceur-ws.org/Vol-3256/#km4law3>, ISSN: 1613-0073.
- [2] M. Corazza, L. Zilli, M. Palmirani, Topic similarity of heterogeneous legal sources supporting the legislative process, in: F. Dell’Orletta, A. Lenci, S. Montemagni, R. Sprugnoli (Eds.), Proceedings of the Tenth Italian Conference on Computational Linguistics (CLiC-it 2024), Pisa, Italy, December 4-6, 2024, volume 3878 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2024. URL: [https://ceur-ws.org/Vol-3878/28\\_main\\_long.pdf](https://ceur-ws.org/Vol-3878/28_main_long.pdf).
- [3] L. Ragazzi, G. Moro, S. Guidi, G. Frisoni, Lawsuit: a large expert-written summarization dataset of italian constitutional court verdicts, *Artificial Intelligence and Law* (2024) 1–37.
- [4] J. Chen, S. Xiao, P. Zhang, K. Luo, D. Lian, Z. Liu, Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation, 2024. URL: <https://arxiv.org/abs/2402.03216>. arXiv: 2402.03216.

---

<sup>3</sup><http://u2.cirsfid.unibo.it/portale-camera/>