

Augmented Democracy: Artificial Intelligence as a Tool to Fight Disinformation

ALEXANDRE ALCOFORADO, ^I

THOMAS PALMEIRA FERRAZ, ^{II} ENZO BUSTOS, ^{III}

ANDRÉ SEIDEL OLIVEIRA, ^{IV} RODRIGO GERBER, ^V

GIAN LUCCA DU MONT SANTORO, ^{VI}

ISRAEL CAMPOS FAMA, ^{VII} BRUNO MIGUEL VELOSO, ^{VIII}

FÁBIO LEVY SIQUEIRA, ^{IX} ANNA HELENA REALI COSTA ^X

Introduction

WELL-INFORMED citizens and spaces for debate and criticism are fundamental principles of democracy. Without them, the quality of debate that societies must go through to improve themselves may degrade (Prothro; Grigg, 1960).

In the last decade, there has been a significant change in the concept of space for information and discussion: new technologies have allowed people to organize themselves to express their opinions and change their countries' political regimes. An exciting example of the phenomenon was the “Arab Spring”, in 2011, when mobilizations promoted on social networks, such as YouTube, Facebook, and Twitter caused great citizen engagement (Safranek, 2012). However, these same technologies have generated social tension and polarization, as they facilitate the dissemination of fake news and extremist speeches and may even influence election results (Tucker et al., 2018). For example, on January 8, 2023, thousands of people invaded the headquarters of the three branches of government in Brasília. In the months before the invasion, social media platforms were bombarded with fake news that not only cast doubt on the integrity of the Brazilian electoral process but also suggested a supposed constitutional legitimacy for the military forces to take power (Mota, 2023). Therefore, new technologies present us with both a challenge and an opportunity: to combat the spread of disinformation propaganda; and use this new medium as a service for democracy to bring politics closer to the citizenry and promote debate on the main questions and important issues. This challenge has become even

greater with the launch of ChatGPT¹ in 2022, a language model developed by OpenAI capable of generating texts based on a prompt provided by the user. In other words, in addition to amplifying the reach of fake news, technology is now capable of creating content (well-structured texts that are sometimes difficult to distinguish whether they were generated by humans or machines).

It is noteworthy that adopting new technologies has changed aspects of life in society. Among many reasons for this, the popularization of remotely connected devices allows people to consume information and publish their opinions extremely fast. However, although life has streamlined in many ways, most aspects of politics continue to be conducted similarly, without exploiting the power of synthesis present in most digital content. What occurs, therefore, is a mismatch between citizens – and their digital life – and the political behavior of their countries and the world, which may be one of the factors that strongly contribute to the current reduction of belief and trust in democratic institutions (Simon et al., 2017). In this context, the concept of *e-democracy* or digital democracy gains strength: it can be understood as the use of technology in the policy formulation process and citizen-state relations by creating tools that encourage direct citizen participation in the decisions and discussions societies must go through (Council of Europe, 2009). Using these tools, citizens can be much more active in public life and decision-making process (Breindl; Franco, 2008). Vedel (2003) defines three axes for digital democracy: information, discussion, and decision. We can interpret these axes as how citizens can act politically on the internet: inform themselves, be heard in the debate, and effectively participate in the decisions.

Nonetheless, misinformation can be a central problem for digital democracy, as uninformed citizens produce poor discussions and decisions. This phenomenon occurs fundamentally because not having information or having false information causes disinformation. However, information overload (Bontcheva; Gorrell; Wessels, 2013) is also a fundamental cause for disinformation, greatly aggravated by this digital context in which social networks have emerged. The information available in these networks is *noisy* because there is much more data than any user can consume. By their very nature, social networks have caused a more passive consumption of information on the internet: instead of actively seeking it out, people spend most of their time filtering and managing the enormous amount of information they receive daily.

It is worth noting that information produced by state agencies is also *noisy* in general: it often consists of extensive documentation that is difficult for citizens to interpret, facilitating malicious groups to use decontextualized parts of this data to mislead the population. This excess of data generated from state agencies' documents can be characterized as *big data* since they are information assets with high volume, high velocity and high variety (or the three “V”s) that require advanced and innovative techniques and technologies for captur-

ing, storing, and processing information for better analysis and decision-making (Gandomi; Haider, 2015). Therefore, since the set of all information generated by state agencies can be characterized as big data, the creation of services and tools for the dissemination, in a clear, direct and straightforward way of materials that encourage the participation of citizens in ongoing political discussions becomes extremely important.

In search of better ways to handle big data, the last decade provided a revolution in *Artificial Intelligence*, expanding horizons and bringing discoveries of new tasks that machines can perform. In particular, *natural language processing* techniques already have the power to automatically interpret large volumes of information (Young et al., 2018). This ability enables, for example, the extraction of relevant and straightforward information from the large number of documents produced by state agencies, many of which are composed of long and complex texts. Creating a tool that uses these techniques to simplify information and translate it for the public would have many practical applications, among which we mention: assisting news agencies in information queries and the elaboration of journalistic works based on data (*Data Journalism*), and providing an increase in transparency of the actions of the political class towards society. It presents an opportunity to employ these new natural language processing techniques for big data generated by state agencies.

In this context, this paper aims to contribute with a tool that allows the automatic collection, organization, selection, processing and simplification of information generated in political discussions. Our tool reduces data noise to get better information by processing the available public data. Therefore, we approach one of the essential axes of digital democracy information. In particular, we show a case study to validate our augmented democracy tool in this work. We turned to the Portuguese Parliament and used the minutes of the plenary meetings described in the Gazette of the Assembly of the Republic (or *Diário da Assembleia da República Portuguesa*, in Portuguese), seeking to develop a tool that helps to clarify the legislative work. Our main contribution is to demonstrate the potential and feasibility of the proposed tool that:

- 1 Processes the massive and complex amounts of data produced in the legislative body, segmenting the minutes into pieces of discussion;
- 2 Extracts knowledge from political discourse, identifying actors, topics, subjects of interest, themes, projects and citations in the discussions;
- 3 Summarizes (*i.e.*, compresses) relevant information, transforming large corpora of text into a smaller, easy-to-read representative text;
- 4 Provides a set of analysis and information dashboards for citizens and organizations.

The remainder of this paper is organized as follows. Initially, we sought to describe existing related works, seeking inspiration for features, and detecting similarities and differences with our proposal. Next, we describe our proposed

tool and summarize its functionalities, which are described in the following sections illustrated with representative examples drawn from the case study carried out with the minutes of the Portuguese Parliament. Finally, we close with our discussions and proposals for future work.

Related work

Some authors are already focusing on the textual analysis of public documents. Most applications developed in this context are in the *sentiment analysis* class, a branch of natural language processing. The sentiment analysis aims to classify the opinion expressed in texts as being positive, neutral, or negative about the items they are commenting on. Most works in the literature focused on the English language, but recent research in other languages has emerged more often (Rao, 2019).

One of the pioneering works was that developed by Mullen and Malouf (2006). They explored the sentiment analysis of informal political discourse, reporting statistical tests performed on data from an American political discourse group in English. Although the results were preliminary, it is clear that simple text classification methods only sometimes yield satisfactory results for data as complex as political discourse texts.

However, in recent times, natural language processing has undergone a significant transformation due to powerful artificial intelligence techniques, such as neural networks and transformers, enabling more exciting results. Most of these techniques, notwithstanding, follow a supervised learning paradigm that requires a large amount of labeled data to train them, which is only sometimes available.

Recently, Watanabe and Zhou (2022) applied semi-supervised text classification techniques in United Nations documents. Semi-supervised techniques aim to reduce dependence on labeled data in training. They achieved robust results in Japanese and English languages using only small seed sentiment words to train the algorithm. While the results were encouraging, seed words for classification are still needed, which can be a potential flaw. It is in our interest to explore classification methods that only need a list of possible classes, not depending on other auxiliary information provided by the user.

An inspiring example of our work is the project “Decide Madrid” (Procter et al., 2021), which, in addition to summarizing political information, allows citizens to give their opinion on the projects through the voting process. Among other goals, the “Decide Madrid” project aims to identify citizens’ political stances, group them by similar interests, and suggest proposals those users may want to support. It intensively explores a branch of artificial intelligence, *machine learning*, just as we do. However, in the current state of our tool, the platform offered by the project “Decide Madrid” is much more interactive with users, while our tool mainly offers data visualization.

In Brazil, Silva et al. (2021) applied topic modeling to a dataset provided by The Chamber of Deputies Board of Innovation and Information Technology, obtaining robust results and showing the effectiveness of applying modern techniques of natural language processing to political discourse. This work inspires us because it instigates the curiosity of seeing what would happen when not one, but many techniques are applied to political discourse.

Although not the main focus of the proposed tool, it is interesting to observe that the amount of data generated does not present a challenge solely to the average citizen. Technicians and political agents from various levels of the public sector are required to analyze ever-increasing amounts of data in their work routines. In the legal field, for instance (Carmo et al., 2023), point out the existence of approximately 77.3 million ongoing cases in the Brazilian judicial system as of 2022. With such a significant volume, the use of technologies capable of accelerating the process flow has a relevant positive impact on society. In this regard, the authors mention the Justice 4.0 Program, a governmental initiative aimed at promoting digital solutions to modernize the Brazilian Judiciary.

In the Portuguese scenario, a mobile application called *meuparlamento.pt*² allows people to verify which political party they are closest to when only the votes carried out on specific projects are analyzed. This application addresses the same domain as our tool, the Portuguese Parliament; however, there is no exploration of artificial intelligence techniques to extract more elaborate knowledge from the minutes of parliamentary sessions – and that is our goal here.

DIÁRIO

da Assembleia da República

VII LEGISLATURA

4.ª SESSÃO LEGISLATIVA (1998-1999)

REUNIÃO PLENÁRIA DE 16 DE SETEMBRO DE 1998

Presidente: Ex.^{mo} Sr. António de Almeida Santos

Secretários: Ex.^{mos} Srs. Artur Rodrigues Pereira dos Penedos
Duarte Rogério Matos Ventura Pacheco
João Cerveira Corregedor da Fonseca
Rosa Maria da Silva Bastos da Horta Albernaz

SUMÁRIO

O Sr. Presidente, que saudou os Deputados na abertura do ano parlamentar, declarou aberta a sessão às 15 horas e 25 minutos.

Antes da ordem do dia. — *Deu-se conta da entrada na Mesa da proposta de lei n.º 207/VII, da proposta de resolução n.º 119/VII e do projecto de lei n.º 556/VII, bem como de requerimentos e da resposta a alguns outros.*

Em declaração política, o Sr. Deputado Francisco de Assis (PS) elogiou a acção governativa em vários domínios e criticou o PSD pela sua postura em relação não só à regionalização, mas também ao inquérito parlamentar realizado na sequência das acusações feitas ao Governo pelo líder do Partido Social Democrata, no último Congresso do PSD realizado no Algarve, e por este partido ter apresentado um projecto de lei que tem a ver com a questão da Sport TV.

Também em declaração política, o Sr. Deputado Carlos Encarnação (PSD) acusou o Sr. Ministro da Agricultura por ter desprezado o protesto dos agricultores, assim como o Sr. Primeiro-Ministro pelo modo de regiões apresentado pelo PS.

Igualmente em declaração política, o Sr. Deputado Octávio Teixeira (PCP) chamou a atenção para o adiamento por parte do Governo do cumprimento de promessas eleitorais.

Por fim, também em declaração política, o Sr. Deputado Nuno Abecasis (CDS-PP) fez um apelo aos Deputados portugueses de solidariedade para com os representantes da UNITA no Parlamento angolano.

A requerimento do PS, a discussão e votação do voto n.º 130/VII — De protesto pela falta de resposta do Governo aos problemas da

lavoura nacional (CDS-PP) foram adiadas para a sessão plenária do dia seguinte. Usaram da palavra os Srs. Deputados Acácio Barreiros (PS) e Luís Queiró (CDS-PP).

A Câmara aprovou um relatório e parecer da Comissão de Assuntos Constitucionais, Direitos, Liberdades e Garantias sobre a retoma de um Deputado do PSD e a substituição de três Deputados, um do PS, outro do PCP e outro do PSD.

Ordem do dia. — *Procedeu-se à apreciação conjunta das propostas de resolução n.º 106/VII — Aprova, para ratificação, o Protocolo de Adesão da República da Polónia ao Tratado do Atlântico Norte, assinado em Bruxelas, em 16 de Dezembro de 1997, 107/VII — Aprova, para ratificação, o Protocolo de Adesão da República Checa ao Tratado do Atlântico Norte, assinado em Bruxelas, em 16 de Dezembro de 1997, e 108/VII — Aprova, para ratificação, o Protocolo de Adesão da República da Hungria ao Tratado do Atlântico Norte, assinado em Bruxelas, em 16 de Dezembro de 1997, tendo as mesmas sido aprovadas em votação global. Intervieram, além do Sr. Ministro dos Negócios Estrangeiros (Álvaro Gama), os Srs. Deputados Pedro Holstein Campinho (PSD), Laurentino Dias (PS), Nuno Abecasis (CDS-PP) e João Amaral (PCP).*

O projecto de lei n.º 530/VII — Privatização do notariado (PSD) foi discutido na generalidade, tendo usado da palavra, a diversos títulos, os Srs. Deputados Maria Eduarda Azevedo (PSD), Odete Santos (PCP), Nuno Baltazar Mendes (PS), Guilherme Silva (PSD) e Silvano Rui Cervan (CDS-PP).

O Sr. Presidente encerrou a sessão eram 18 horas e 40 minutos.

Source: Site of the Assembly of the Republic, DAR I Series n.001 (Accessed June 29, 2024).

Figure 1 – Gazette of the Assembly of the Republic (DAR, from “Diário da Assembleia da República”, in Portuguese)

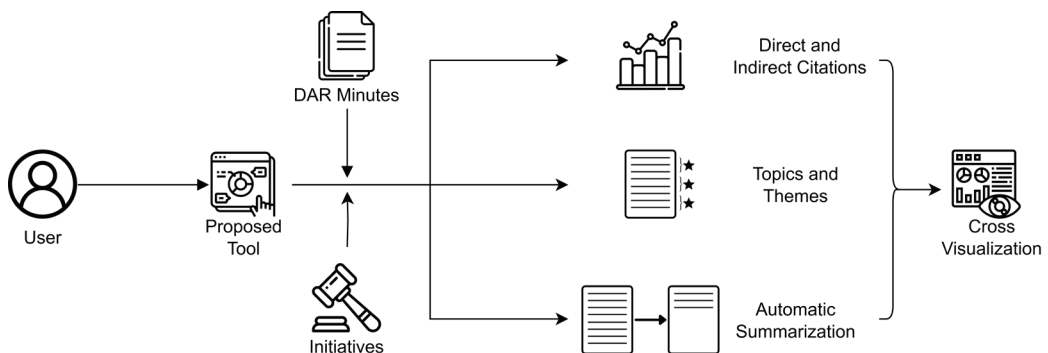
Proposed tool

Portugal has been a democracy since 1974, and in 1975, the minutes of the Portuguese Parliament started being transcribed. From 1976, on the 3rd. Republic, the transcripts of the Gazette of the Assembly of the Republic (“Diário da Assembleia da República”, DAR, in Portuguese) appeared, illustrated in Figure 1, freely available to the public (Assembleia da República, 2024). The transcripts of the DAR minutes are used by our tool to extract information of public interest.

The tool proposed in this paper presupposes determining which information is relevant to the public and presenting the results simply and visually in an easy-to-handle interface. By focusing on the Portuguese Parliament, we sought to identify which information could interest citizens and which information al-

ready produced by the DAR minutes could be summarized and exposed to the citizen in different and understandable ways.

Figure 2 illustrates the general scheme of the proposed tool and its respective functionalities. Initially, both the DAR minutes and the initiatives discussed in parliament are accessed, and transcripts are taken from the parliament website², where all minutes are publicly and freely available. The texts of the minutes are previously processed and stored in a database, which will then be available for the functionalities to access, as described in the next section. From the DAR minutes and initiatives obtained in the data collection step and stored in the database, the tool can automatically compute citations spoken by members of parliament in the sessions of interest, classify discussions in the sessions into topics and themes, and summarize parts of the minutes. Each of these functionalities and respective forms of presentation in the user interface is described in the subsequent sections.



Source: Elaborated by the authors.

Figure 2 – General scheme of the proposed tool and its respective functionalities: (i) Computation of direct or indirect citations; (ii) Classification of topics and themes covered in the minutes; and (iii) Automatic summarization of parts of the minutes of interest. These functionalities access a database generated from the DAR minutes and the initiatives discussed in the parliament sessions. Results are shown in a user interaction interface.

Data collection

Initially, we used *web crawling* (Pant; Srinivasan; Menczer, 2004) and *web scraping* (Mitchell, 2018) techniques to collect data periodically published on the Portuguese Parliament’s website, namely the minutes of parliamentary sessions (Assembleia da Republica, 2021). This data is available in the form of texts, each containing the transcripts of a session. We then structured the content of these documents in a database to facilitate processing, and building large text corpora. These data were pre-segmented in the agenda items of each meeting, which involve political statements, discussions, and votes on parliamentary initiatives (bills, draft resolutions, parliamentary inquiries, among others), in addition to other matters. We refer to these tasks as the *data collection* step.

With this database organized, we apply machine learning techniques associated with natural language processing methods to provide functionalities to the user. All results were obtained by conducting experiments on minutes from 09/16/2020 to 02/25/2021.

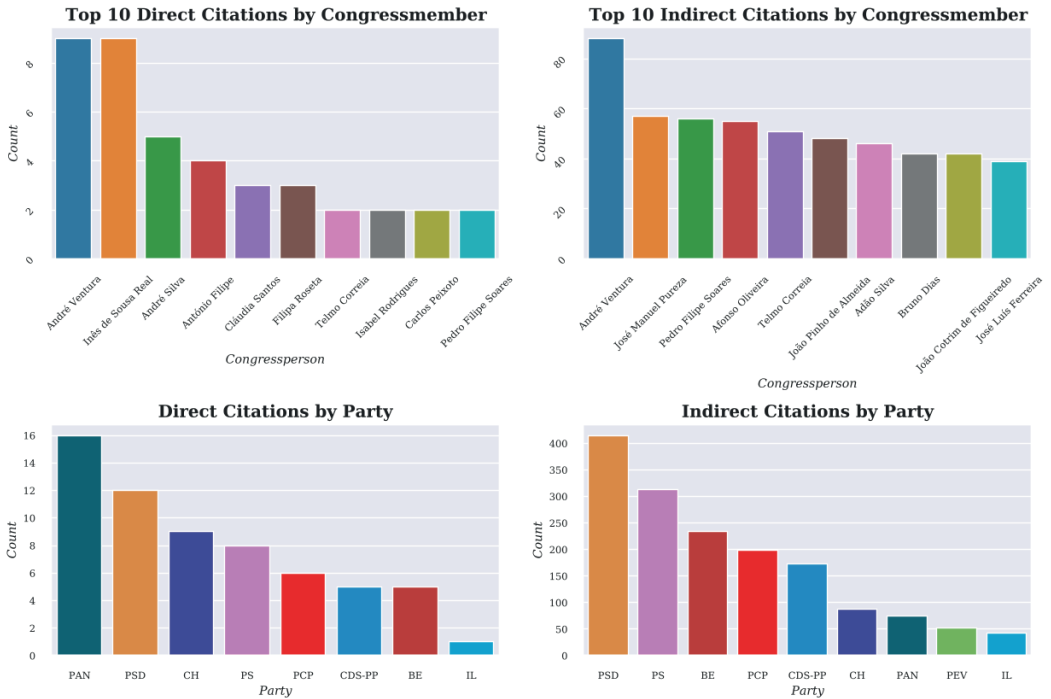
Computation of citations

We define a *subject of interest* as a set of user-defined keywords related to a subject, such as “corruption” or “education”. We define *direct citations* as parliament members’ speeches that explicitly mention the keywords. The other speeches contained in the same discussion in which there was at least one direct citation are *indirect citations*. A functionality presented by the augmented democracy tool proposed in this work uses segmented data from political statements to, according to the subject of interest, calculate the direct and indirect citations for each party and each member of parliament within the time interval defined by the user.

To make it possible to search for direct and indirect citations, the DE-BACER algorithm (Ferraz et al., 2021) was applied in the minutes to partition the sets of all political statements of each parliamentary session into *blocks of speech* that contain complete discussions (with a beginning and end), which is essential for understanding what is being said by the participants. We stored these partitions in a database and developed an analysis procedure. User-defined keywords were searched in these partitions, identifying the parliamentarian who explicitly cited them. From this data, various information could be generated. In particular, we show how often the keyword is cited by each member of parliament (or per party) in the user-defined period of interest. Indirect citations, both from a member of parliament and parties, are also computed.

We present the results of a case study for the subject of interest, “corruption”, to demonstrate the effectiveness of the proposed tool, analyzing the participation of parties and members of parliament in this subject. The importance of using corruption as an example in the Portuguese case can be substantiated by several facts (Prémio Tágides, 2021, 2021), such as: (i) the estimated cost for the Portuguese of known corruption cases is equivalent to 30 % of the national public debt; (ii) the fact that only 1 of the 15 anti-corruption acts recommended in 2016 was fully implemented in Portugal; and (iii) also the fact that the European Parliament has estimated that corruption in Portugal costs the equivalent of 8 % to 10 % of GDP.

Figure 3 presents an example of information made available to the user. The figure shows the histograms indicating the number of direct and indirect citations of the user-defined word “corruption” (“*corrupção*” in Portuguese) and its derivatives (*e.g.*, corrupt, corruptions, etc.), computed from the period between 09/16/2020 and 02/25/2021, either by each member of parliament or by the party.



Source: Elaborated by the authors.

Figure 3 – Histograms of direct and indirect citations of derivatives of the word “corruption” (“*corrupção*” in Portuguese) in the period from 09/16/2020 to 02/25/2021. Top: Citations by parliamentarians (only the ten most cited). Bottom: Citations by party.

We can combine this tool with other functionalities. For example, we could determine the dominant theme on which the keyword was cited (for example, “corruption in health”) if the citation functionality were combined with the one that discovers topics and themes described in the next section.

Topics and themes

A second functionality developed in our tool performs two steps in the minutes of parliamentary sessions or parts of them, according to the period specified by the user. The first step is topic modeling, and the other is classifying the speeches into themes.

A *topic model* is a text-mining statistical model used to discover semantic structures hidden in a collection of documents. We define a *topic* as a set of words that belong to the same semantic field, like “school, teachers, class” or “5G, science, network”. We use *topic modeling* (Blei; Ng; Jordan, 2003; Groo-tendorst, 2022) to determine topics that occur in documents automatically. Unlike user-defined subjects of interest given by keywords used in the computation of citations, topics are automatically extracted from documents by an artificial intelligence algorithm trained on a specific language (here, Portuguese).

The topic modeling algorithm found several topics and associated each speech with them. The algorithm returns the probability that indicates the relevance of each topic to the speech. Because of this, topics can be presented to the user differently. For example, the topic with the *highest association probability* to the text, *top 3* topics, *all topics* and their respective probabilities, and so on.

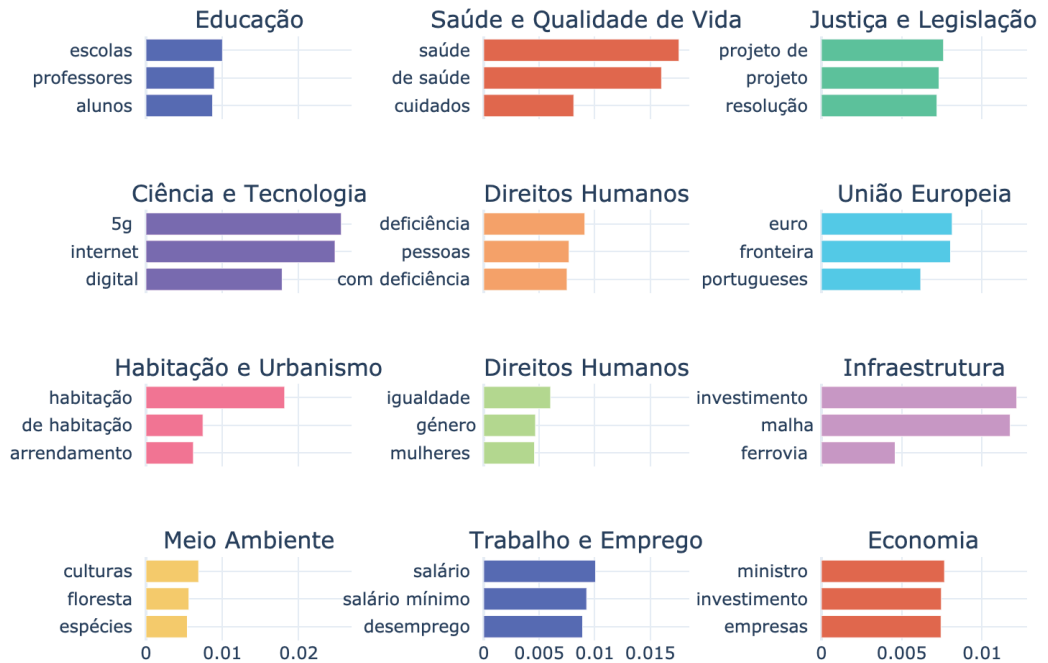
However, topics found by topic modeling algorithms are only sometimes interpretable by humans. They form clusters of words but do not have their name for each group, making it difficult for the user to infer the meaning of that cluster. So, our functionality allows users to choose whether they are satisfied with these automatically found topics or would like to use the second step of the functionality, which is to classify speeches into *themes*.

Themes are defined a priori in the tool and describe categories of interest to parliamentary activity, such as Economy, Environment, etc. To classify speeches into pre-established themes, this functionality uses Zero-Shot algorithm (Romera-Paredes; Torr, 2015; Yin; Hay; Roth, 2019). Technically, we apply the ZeroBERTo algorithm (Alcoforado et al., 2022), that considers: (i) the topics found by the topic modeling algorithm in the previous step; (ii) the distribution of probabilities of each speech belonging to each topic; (iii) the defined themes. The themes are then probabilistically associated with the speeches by a trained language model, leveraging its general knowledge of the language. Similar to the previous association (speech to topic), this step associates each speech with all previously defined themes. Thus, themes can be presented in different ways as well. For example, the theme with the *highest association probability* to the text, *top 5* themes, *all themes* and their respective probabilities, a graph with the temporal evolution of a particular theme, among others.

An important aspect of this functionality is that the same keywords may be chosen as a *subject of interest* and *theme*. However, there are fundamental differences between them:

- First, the input for the subject of interest in the first functionality is a single term (e.g., “corruption”), and from it is defined a set of keywords that derive from the given term (e.g., corrupts, corrupted, corruption, etc.), while the input for themes is necessarily a set of words related to the issues discussed in parliament (e.g.: science and technology, environment, human rights, health, infrastructure);
- Second, the subject of interest in the first functionality is searched within the time interval set by the user, looking for exact or approximate matches for those keywords, while the themes are automatically associated with the speeches by a language model;
- Third, a subject of interest may or may not be found in a speech in a given time interval, while every theme, in turn, is associated with every speech, and this association is represented by a probability, which may be similar for more than one theme.

Information about a documents' topic model and theme classification is then visually displayed to the user, as illustrated in Figure 4. In the figure, we see some examples presented to the user in the graphical interface: the 12 most frequent topics found, described by the three most representative words (ex. the pink topic is defined by the words “*investimento*”(investment), “*malha*” (network) and “*ferrovia*” (railway). Above the 12 topics presented in the figure, we see the themes on which each topic was best rated (for example, the pink topic was rated first as “*Infraestrutura*” (infrastructure). These data are related to the minutes from 09/16/2020 to 02/25/2021.



Source: Elaborated by the authors.

Figure 4 – Topics found (only the 12 most frequent topics), along with their most representative words (only the three most representative). Above the bars is the theme associated with each specific topic (only the most representative theme), from 09/16/2020 to 02/25/2021. Results are in Portuguese.


Applying this functionality in the referred minutes resulted in about 60 topics in the topic modeling step. This experiment was carried out with the following set of themes: Corruption (“*Corrupção*”), Culture (“*Cultura*”), Economy (“*Economia*”), Education (“*Educação*”), Energy (“*Energia*”), Environment (“*Meio Ambiente*”), European Union (“*União Europeia*”), Health and Quality of Life (“*Saúde e Qualidade de Vida*”), Housing and Urban Planning (“*Habituação e Urbanismo*”), Human Rights (“*Direitos Humanos*”), Industry and Agriculture and Commerce (“*Indústria e Agricultura e Comércio*”), Infrastructure

(“*Infraestrutura*”), Justice (“*Justiça*”), Legislation (“*Legislação*”), National Defense and Public Security (“*Defesa Nacional e Segurança Pública*”), Science and Technology (“*Ciência e Tecnologia*”), Tourism (“*Turismo*”), Work and Employment (“*Trabalho e Emprego*”).

With this functionality, it is also possible to calculate, for example, how many times a particular member of parliament or party spoke about a particular theme. Another option is, for example, to summarize speeches that have a high probability of being associated with a specific topic according to the user’s desire. The summarization functionality is described in the next section.

Summarization

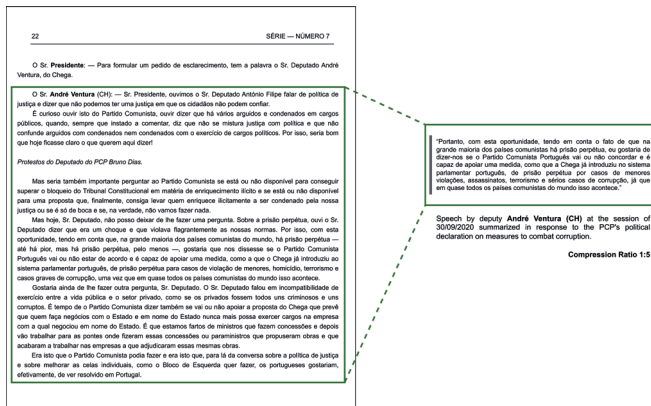
The automatic summarization functionality applies state-of-the-art neural networks trained for summarization to generate short reports of speeches related to reports of initiatives, in particular bills. These speeches are longer than the others and occur right after a discussion starts, contextualizing the discussion to other members of the parliament. More specifically, we apply the PEGASUS model (Zhang et al., 2020), a Transformer neural network (Vaswani et al., 2017) capable of generating authorial summaries in English, that is, summaries with original sentences that are not simple extracts from source texts. In order to do that, our algorithm first translates the source text of the speech to English, generates the summary with PEGASUS, and finally translates the summary back to Portuguese. We apply the M2M100 translator (Fan et al., 2021) since it is a robust multilingual Transformer-based model. Figure 5 illustrates the full summarization procedure.



Source: Elaborated by the authors.

Figure 5 – The summary of an initiative report is generated from texts initially translated from Portuguese to English by the M2M100 translator, followed by the PEGASUS English summarizer and finally translated back to Portuguese, generating short summaries.

These short summaries of the speeches of the rapporteurs of each initiative are displayed to users as a report of the main discussions that took place in the minutes. Users can also use the functionality to generate additional summaries of the full discussion about that project, or they can be redirected to the original minutes that transcribe the entire discussion of interest. An illustration of the original text and its generated summary is in Figure 6.



Source: Elaborated by the authors.

Figure 6 – Summary example: on the left is the part of a minute that was summarized in the figure on the right.

Cross-Visualization

The functionalities previously presented generate outputs such as summaries of documents, the number of direct or indirect citations of a particular subject of interest and their respective speakers, and the topics and themes covered in different parts of the minutes of the parliamentary sessions.

The interaction of the information extracted by the different functionalities of the developed tool allows many other data visualization options. As all data is also available for time interval filters and parts of minutes that users are interested in, the tool can be used in a very personalized way.

We note that external information may also be used with the information generated by our tool. One way could be, for example, to compare GDP growth with the temporal evolution of the discussion on the “Economy” theme that took place in parliament, or it could even compare the unemployment rate and the value of the minimum wage with the temporal evolution of the topic “Work and Employment”.

These different ways of visualizing data are still under development in our tool. After development, usability testing with end users will still be required.

Discussion and future work

In this paper, we proposed a tool with an interactive graphical interface that simplifies the information produced by the Portuguese Parliament in the time interval defined by the user and presents the following results:

- 1 Statistics on the participation of a member of parliament or party in a specified subject of interest, such as direct and indirect citations of the subject;
- 2 The most/least relevant topics discussed by members of parliament and parties in the defined period;
- 3 The more or less frequent themes that correspond to the topics extracted from the discussions;

4 Speech summaries;

5 Statistics on speeches of a deputy or party significantly associated with the defined themes.

Thus, our proposed tool collaborates directly with the information axis, which is the first axis of digital democracy. We collected and structured public data and processed it, extracting knowledge that would otherwise be impossible due to the massive amount of textual data. By synthesizing, summarizing, modeling, and classifying this available data with machine learning techniques, we show it is possible to provide ways to visualize information already public – but unprocessed – that is easier to understand by citizens and society in general. In addition, the interaction between the different functionalities of our tool can offer a variety of visualization options. The authors are still discussing different visualization possibilities with different audiences of interest.

The information produced by our tool allows citizens to identify the political behavior of members of parliament and parties, evaluate the adherence to their political ideologies, and the coherence between discourse and practice of the parties and members of parliament. As a result, people can encourage greater attention by political actors to the issues society deems most important.

The proposed tool is still under development. To assess its feasibility and the value of its results, we conducted a set of experiments in a modest amount of minutes. The algorithms involved in these experiments must be integrated into the tool to be made available to the public and validated by the target users.

An important point of attention is the complexity of the machine learning algorithms used in the tool. Most functionalities are very difficult to run in real-time due to their computational cost, introducing user limitations or requiring powerful hardware. We intend further to investigate the machine learning engineering aspect of our tool, aiming to make it robust and accessible to everyone.

More investigations and experiments with these techniques and data are needed. Validation with human users, for example, is essential for any artificial intelligence tool to be made publicly available (Sichman, 2021). Research in human-computer interaction is being carried out (Jaimes; Sebe, 2007), reporting robust results in very complex applications. In this context, Cascella et al. (2023) assessed the potential and limitations of GhatGPT in four scenarios related to healthcare: (1) clinical practice support, (2) scientific production, (3) misuse in medicine and research, and (4) argumentation on topics in public health. The authors point out that tools like ChatGPT can indeed accelerate scientific production as they are capable of supporting various aspects of research, such as summarizing large volumes of medical texts, patient records, and scientific articles. On the other hand, when asked to write an article about a set of fake data provided by the authors, ChatGPT generated a plausible and well-structured text. The authors conducted several experiments, and the results highlight the

importance for the scientific community to be aware of the ability of ChatGPT (and similar tools) to generate and disseminate false information.

The data our tool deals with is complex and extensive. We have a real and concrete concern about not presenting the original data in a precise and understandable way, as the lack of precision in the information shown or interpretability in the procedures can have counterproductive effects, confusing users in a scenario of information overload.

We must also pay attention to an important aspect of the proposed tool. Our objective in this paper – to increase the transparency of democratic processes – assumes that a tool is itself a comprehensive tool, bringing transparency and understandability to the applied machine learning techniques (Adadi and Berrada, 2018). However, our tool remains a set of many modules that ordinary people cannot easily understand. Some directions for future work involve investigating and conducting research on Explainable AI (Arrieta et al., 2020; Adadi; Berrada, 2018) and its use in applications involving democratic processes.

It is also important to emphasize that, for a true democratization of the information provided by our tool, the interface should allow wide accessibility for people with certain limitations, such as, for example, the visually impaired. This could be a topic of attention in the future.

Finally, this tool can serve as a basis for other tools that can develop the other two axes of digital democracy, discussion, and decision. With our tool, we demonstrate the potential of using artificial intelligence, machine learning and natural language processing techniques to improve the quality of information society has access to, which, in the long run, can increase confidence in democracy.

Acknowledgments – This research was supported in part by Itaú Unibanco S.A., with the scholarship program of Programa de Bolsas Itaú (PBI), by the Coordination of Improvement of Higher Education Personnel (Capes), Finance Code 001, and by the National Council for Scientific and Technological Development (CNPq) (grant 312360/2023-1), Brazil. Any opinions, findings, and conclusions expressed in this manuscript are those of the authors and do not necessarily reflect the views, official policy or position of the Itaú-Unibanco, Capes and CNPq.

Notes

1 Available at: <<https://chatgpt.com/>>.

2 Available at: <<https://parlamento.pai.pt/>>.

References

ADADI, A.; BERRADA, M. Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). *IEEE access*, IEEE, v.6, p.52138-60, 2018.

- ALCOFORADO, A. et al. Zeroberto: Leveraging zero-shot text classification by topic modeling. In: PINHEIRO, V. et al. (Ed.) *Computational Processing of the Portuguese Language*. Cham: Springer International Publishing, 2022. p.125-36.
- ARRIETA, A. B. et al. Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible ai. *Information Fusion*, v.58, p.82-115, 2020.
- ASSEMBLEIA DA REPÚBLICA. 2021. Available at: <<https://meuparlamento.pt/>>. Accessed June 29, 2024.
- BLEI, D. M.; NG, A. Y.; JORDAN, M. I. Latent Dirichlet Allocation. *Journal of Machine Learning Research*, v.3, n.Jan, p.993-1022, 2003.
- BONTCHEVA, K.; GORRELL, G.; WESSELS, B. Social Media and Information Overload: Survey Results. *arXiv preprint arXiv:1306.0813*, 2013.
- BREINDL, Y.; FRANCO, P. Can Web 2.0 applications save e-democracy? A study of how new internet applications may enhance citizen participation in the political process online. *International Journal of Electronic Democracy*, Inderscience Publishers, v.1, n.1, p.14-31, 2008.
- CARMO, F. A. et al. Embeddings Jurídico: Representações Orientadas à Linguagem Jurídica Brasileira. In: WORKSHOP DE COMPUTAÇÃO APLICADA EM GOVERNO ELETRÔNICO (WCGE), 11, João Pessoa/PB. *Anais...* Porto Alegre: Sociedade Brasileira de Computação, p.188-99, 2023.
- CASCELLA, M. et al. Evaluating the Feasibility of ChatGPT in Healthcare: An Analysis of Multiple Clinical and Research Scenarios. *Journal of medical systems*, v.47, n.33, 4 Mar. 2023.
- COUNCIL OF EUROPE. *Electronic democracy (“e-democracy”) – Recommendation CM/Rec(2009)1 and explanatory memorandum*. Council of Europe Publishing, 2009.
- FAN, A. et al. Beyond english-centric multilingual machine translation. *Journal of Machine Learning Research*, v.22, n.107, p.1-48, 2021.
- FERRAZ, T. P. et al. DEBACER: a method for slicing moderated debates. In: SBC. XVIII ENCONTRO NACIONAL DE INTELIGÊNCIA ARTIFICIAL E COMPUTACIONAL. p. 667-78, 2021.
- GANDOMI, A.; HAIDER, M. Beyond the hype: Big data concepts, methods, and analytics. *International Journal of Information Management*, v.35, n.2, p.137-44, April 2015.
- GROOTENDORST, M. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794*, 2022.
- JAIMES, A.; SEBE, N. Multimodal human-computer interaction: A survey. *Computer Vision and Image Understanding*, v.108, n.1-2, p.116-34, 2007.
- MITCHELL, R. *Web scraping with Python: Collecting more data from the modern web*. O’Reilly Media, Inc., 2018.
- MOTA, C. V. 7 fatores que explicam os ataques de 8 de janeiro em Brasília. BBC News - Brasil. Available at: <<https://www.bbc.com/portuguese/articles/cye7egj6y1no>>. Accessed June 29, 2024.

- MULLEN, T.; MALOUF, R. A Preliminary Investigation into Sentiment Analysis of Informal Political Discourse. In: AAAI SPRING SYMPOSIUM: Computational Approaches to Analyzing Weblogs, p.159-62, 2006.
- PANT, G.; SRINIVASAN, P.; MENCZER, F. Crawling the web. In: *Web Dynamics: Adapting to Change in Content, Size, Topology and Use*, Berlin, Heidelberg: Springer Berlin Heidelberg, 2004. p.153-77.
- PRÉMIO TÁGIDES 2021. Available at: <<https://www.all4integrity.org/premio-tagides/edicao2021/>>. Accessed June 29, 2024.
- PROCTER, R. et al. Citizen Participation and Machine Learning for a Better Democracy. *Digital Government: Research and Practice*, ACM New York, NY, USA, v.2, n.3, p.1-22, 2021.
- PROTHRO, J. W.; GRIGG, C. M. Fundamental Principles of Democracy: Bases of Agreement and Disagreement. *The Journal of Politics*, Southern Political Science Association, v.22, n.2, p.276-94, 1960.
- RAO, P. S. The role of english as a global language. *Research Journal of English*, v.4, n.1, p. 65-79, 2019.
- ROMERA-PAREDES, B.; TORR, P. An embarrassingly simple approach to zero-shot learning. In: INTERNATIONAL CONFERENCE ON MACHINE LEARNING. p.2152-61, 2015.
- SAFRANEK, R. *The emerging role of social media in political and regime change*. p.1-14. ProQuest Discovery Guides, 2012.
- SICHMAN, J. S. Inteligência artificial e sociedade: avanços e riscos. *Estudos Avançados*, v.35, p.37-50, 2021.
- SILVA, N. F. et al. Evaluating topic models in portuguese political comments about bills from Brazil's chamber of deputies. In: BRITTO, A.; VALDIVIA DELGADO, K. (Ed.) *Intelligent Systems*. BRACIS 2021. Lecture Notes in Computer Science, v.13074, p.104-20. Springer, Cham, 2021.
- SIMON, J. et al. *Digital Democracy: The tools transforming political engagement*. [S.l.]: NESTA, UK, England and Wales 1144091, 2017. Available at: <<https://www.nesta.org.uk/report/digital-democracy-the-tools-transforming-political-engagement/>>. Accessed June 29, 2024.
- TUCKER, J. A. et al. Social media, political polarization, and political disinformation: A review of the scientific literature. *SSRN Electronic Journal*, 2018.
- VASWANI, A. et al. Attention is all you need. In: 31st CONFERENCE ON NEURAL INFORMATION PROCESSING SYSTEMS, Long Beach, USA, 2017.
- VEDEL, T. L'idée de démocratie électronique: origines, visions, questions. In: PASCAL, P. (Ed.). *Le désenchantement démocratique*. Paris: Editions de l'Aube, 2003. p.243-66.
- WATANABE, K.; ZHOU, Y. Theory-driven analysis of large corpora: Semisupervised topic classification of the UN speeches. *Social Science Computer Review*, v.40, i. 2, Apr, p.346-66, 2022.
- YIN, W.; HAY, J.; ROTH, D. Benchmarking Zero-shot Text Classification: Datasets, Evaluation and Entailment Approach. In: 2019 CONFERENCE ON EMPIRICAL

METHODS IN NATURAL LANGUAGE PROCESSING and the 9th International Joint Conference On Natural Language Processing (EMNLP-IJCNLP). p.3914-23, 2019.

YOUNG, T. et al. Recent trends in deep learning based natural language processing. *IEEE Computational Intelligence Magazine*, IEEE, v.13, n.3, p.55-75, 2018.

ZHANG, J. et al. PEGASUS: Pre-training with extracted gap-sentences for abstractive summarization. In: INTERNATIONAL CONFERENCE ON MACHINE LEARNING. p.11328-339, 2020.

ABSTRACT – One of the principles of digital democracy is to actively inform citizens and mobilize them to participate in the political debate. This paper introduces a tool that processes public political documents to make information accessible to citizens and specific professional groups. In particular, we investigate and develop artificial intelligence techniques for text mining from the Portuguese *Diário da Assembleia da República* to partition, analyze, extract and synthesize information contained in the minutes of parliamentary sessions. We also developed dashboards to show the extracted information in a simple and visual way, such as summaries of speeches and topics discussed. Our main objective is to increase transparency and accountability between elected officials and voters, rather than characterizing political behavior.

KEYWORDS: Digital democracy, Natural language processing, Artificial Intelligence, Legislative information.

RESUMO – Um dos princípios da democracia digital é informar ativamente os cidadãos e mobilizá-los para participarem no debate político. Este artigo apresenta uma ferramenta de processamento de documentos políticos públicos para tornar as informações mais acessíveis aos cidadãos e grupos profissionais específicos. Em particular, investigamos e desenvolvemos técnicas de Inteligência Artificial para mineração de textos do *Diário da Assembleia da República* de Portugal para particionar, analisar, extrair e sintetizar a informação das atas das sessões parlamentares. Desenvolvemos ainda *dashboards* que mostram as informações extraídas de forma simples e visual, como resumos de falas e tópicos discutidos. O nosso objetivo principal é, mais do que caracterizar o comportamento político, aumentar a transparência e a responsabilidade dos eleitores e das autoridades eleitas.

PALAVRAS-CHAVE: Democracia digital, Processamento de linguagem natural, Inteligência Artificial, Informação legislativa.

Alexandre Alcoforado is a M.Sc. student in Computer Engineering at Escola Politécnica da Universidade de São Paulo, and is an Electrical Engineer from the same institution.
@ – alexandre.alcoforado@usp.br / <http://orcid.org/0000-0003-3184-1534>.

Thomas Palmeira Ferraz is a Ph.D. candidate in Computer Science at Télécom Paris and École Polytechnique, Institut Polytechnique de Paris. He is an Engineer from Escola Politécnica da Universidade de São Paulo (USP) and a Master of Applied Mathematics and Artificial Intelligence from École Normale Supérieure Paris-Saclay (ENS).
@ – thomas.palmeira@telecom-paris.fr / <http://orcid.org/0000-0002-5385-9164>.

Enzo Bustos is an undergraduate student in Electrical Engineering with an emphasis in Computing at Universidade de São Paulo. @ – enzobustos@usp.br / <https://orcid.org/0000-0002-3169-4469>.

André Seidel Oliveira has a M.Sc. degree in Computer Engineering at Escola Politécnica da Universidade de São Paulo. @ – andre.seidel@usp.br / <https://orcid.org/0000-0001-6551-6911>.

Rodrigo Gerber is an Electrical Engineer from Escola Politécnica da Universidade de São Paulo. @ – rodrigo.gerber@usp.br / <http://orcid.org/0000-0002-8120-7972>.

Gian Lucca du Mont Santoro is an undergraduate student in Applied and Computational Mathematics (with a qualification in Economic Statistics) at Instituto de Matemática e Estatística da Universidade de São Paulo and a Data Scientist at Itaú Unibanco. @ – gian.santoro@usp.br / <http://orcid.org/0000-0001-6731-7940>.

Israel Campos Fama is a Ph.D. candidate in Computer Engineering at Escola Politécnica da Universidade de São Paulo. He holds a degree in Mechanical Engineering from the Universidade Federal do Ceará (UFC) and a M.Sc. degree in Aeronautical Engineering from the Instituto Tecnológico de Aeronáutica (ITA). @ - israelfama@usp.br / <http://orcid.org/0000-0001-6325-4153>.

Bruno Veloso is currently an Auxiliar Professor at Faculty of Economics - University of Porto. He is a Senior Researcher at LIAAD/INESC TEC and member of the Spanish association for artificial intelligence. @ – bruno.m.veloso@inesctec.pt / <http://orcid.org/0000-0001-7980-0972>.

Fábio Levy Siqueira holds a Ph.D. in Computer Engineering at Escola Politécnica da Universidade de São Paulo, He is currently an Assistant Professor in the Department of Computer and Digital Systems Engineering at Escola Politécnica da Universidade de São Paulo. @ – levy.siqueira@usp.br / <http://orcid.org/0000-0002-7550-2000>.

Anna Helena Reali Costa is a Full Professor in Computer Engineering at Universidade de São Paulo (USP). She took her Ph.D. at USP and was a guest researcher at Karlsruhe Institute of Technology, Germany, and Carnegie Mellon University, USA. She is the Director of the Data Science Center (C2D), a partnership between USP and the Itaú-Unibanco bank, and a member of the USP-FAPESP-IBM Center for Artificial Intelligence (C4AI). Her scientific contributions lie in Artificial Intelligence and Machine Learning. @ – anna.reali@usp.br / <http://orcid.org/0000-0001-7309-4528>.

Recebido em 26.1.2023 e aceito em 16.11.2023.

^{I,III,IV,V,VII,IX,X} Universidade de São Paulo, Escola Politécnica, São Paulo, Brasil.

^{II} Télécom Paris, Institut Polytechnique de Paris, Palaiseau, França.

^{VI} Universidade de São Paulo, Instituto de Matemática e Estatística, São Paulo, Brasil.

^{VIII} Universidade do Porto e INESC TEC, Porto, Portugal.

