INSTITUTO
UNIVERSITÁRIO
DE LISBOA

De'Coding Concepts: Automatic Characterization of populist moments on Twitter

Vasco Lopo Hipólito Branco Ribeiro Mateus

Master in Political Economy

Supervisor:
João Miguel Duarte de Carvalho, Invited Principal Investigator,
Iscte - Instituto Universitário de Lisboa

Co-Supervisor:
Ricardo Daniel Santos Faro Marques Ribeiro, Associate Professor,
Iscte - Instituto Universitário de Lisboa

September, 2023

CIÊNCIAS SOCIAIS
E HUMANAS

Department of Political Economy

De'Coding Concepts: Automatic Characterization of populist moments on Twitter

Vasco Lopo Hipólito Branco Ribeiro Mateus

Master in Political Economy

Supervisor:
João Miguel Duarte de Carvalho, Invited Principal Investigator,
Iscte - Instituto Universitário de Lisboa

Co-Supervisor:
Ricardo Daniel Santos Faro Marques Ribeiro, Associate Professor,
Iscte - Instituto Universitário de Lisboa

September, 2023

Acknowledgments

To my incredible parents, who faced their own monumental challenges over the past two years, of which they were far more arduous than my own. Yet, they emerged victorious with unwavering strength, graciously extendeding their time, guidance, and support to nurture my endeavors. Without your boundless patience, countless proofreading sessions, and your willingness and selfless dedication to delve into unfamiliar concepts and fields, this dissertation would never have come to fruition.

I extend my heartfelt gratitude to my supervisors, Doctor João Miguel Carvalho and Professor Ricardo Ribeiro. To the former, from the very outset, undertook the task of deciphering my jumble of ideas and plans, and shaping them, through direct and stern feedback, into a smaller, coherent and viable framework. Despite his best efforts, I was unable to uphuld his ideals of pragmatism and took a lot longer to realize this dissertation than first anticipated.

To the latter, your timely entry into the trajectory of this dissertation was nothing short of transformative. With your meticulous and methodical approach, you breathed new life into a project that had been adrift for months. Your weekly discussions provided the compass that steered this metaphorical ship toward the shore of accomplishment. Your unwavering dedication and expertise have been crucial in ensuring the safe passage of my ambitions throughs through the treacherous waters of inexperience.

To my colleague and mentor Tiago Fernandes. You were there as I took my tentative first steps into the realm of computer programming, guiding me through that Python programming course, one that only I officially partook in, but your contributions surely more than merited a diploma, if not more. Your willingness to engage in countless debates and brainstorming sessions, from the inception of my ideas on dissemination and network theories to the final paragraphs that had eluded me, has been invaluable. Your steadfast support and insights were a driving force. I must acknowledge and honor your significant role in this journey, just like I had promised you.

I extend my deepest appreciation to both my extended family, Susana Franco and Rafael Saturnini, all of whom played a part in shaping this dissertation. I would also like to express my sincere gratitude to the Samsung Innovation Campus team for granting me the opportunity to work with artificial intelligence and machine learning. Their willingness to consider my application, despite my non-traditional background, played a pivotal role in enabling the interdisciplinary nature of this dissertation. My heartfelt thanks goes out to each and every one of you, your contributions, whether through annotations, guidance, or unwavering support, has been instrumental in bringing this endeavor to fruition.

Resumo

Em 1995, McChesney afirmou que a economia política da comunicação estava num ponto de ruptura. Nas décadas seguintes, novos meios de comunicação e uma intensificação da comercialização dos sistemas mediáticos lançaram a disciplina para o centro dos debates políticos e do interesse público. Portugal, outrora visto como resiliente aos movimentos populistas da década de 2010, tem rapidamente transitado em direção da sua adoção nas mais recentes eleições. Paralelamente, o país tem aproximando-se às tendências internacionais na sua recente adoção das redes sociais para estratégias políticas. Esta "recuperação" espelha também um fenômeno mais amplo presente nas suas ciências sociais, que resistem à adoção de metodologias contemporâneas e técnicas de processamento de informação mais avançadas, apesar do crescente universo de dados disponível.

Desenvolvendo um modelo de aprendizagem automática através de técnicas de “deep learning” para classificação binária, cria-mos um método capaz de caracterizar conceitos complexos, como populismo, de forma a tratá-los como variáveis. O bom desempenho do modelo na classificação de tweets de vários atores políticos portugueses, ao longo dos dois últimos ciclos eleitorais, permitiu-nos, através de uma análise complementar, identificar algumas características primárias dos momentos populistas. Este modelo serve, não só, como uma base para futuros estudos sobre populismo, como também demonstra uma forte adaptabilidade a outros conceitos dentro das ciências sociais. Com pequenas alterações nas diretrizes e objetivos do processo de rotulagem, o nosso modelo é um primeiro passo para se uniformizar este tipo de pesquisa, reduzindo as necessidades de mão-de-obra atualmente associadas às análises de grandes conjuntos de dados.

Palavras-chave: Populismo, Texto-como-variáveis, Máquina dos media, Políticos Portugueses, processamento de linguagem natural, BERT.

Abstract

In 1995, McChesney stated that the political economy of communication studies was at a breaking point. In the following two decades, new communication platforms and the heightened commercialization of media systems brought political discourse to the forefront of political debates and public interests. Portugal, which was once seen as resilient against the populist movements throughout the 2010s, has exhibited a noteworthy shift towards embracing these trends in recent election cycles. Simultaneously, the country appears to be aligning with foreign trends in its adoption of social media for political strategies. This parallel "catching-up" echoes a broader phenomenon that is also prevalent in social sciences' inertia to adopt contemporary methodologies and more advanced data techniques, despite the ever-expanding realm of information available.

In our research, we developed a machine-learning model with the objective of using deep learning techniques for binary classification to create a robust method capable of discerning complex concepts, like populism and treating them as variables. The model's accurate performance on classifying tweets from Portuguese political actors, over the last two election cycles, allowed us to establish some primary features of populist moments through a complementary fundamental analysis. The model employed not only serves as a strong foundation for future research within populism but also demonstrates notable adaptability to other possible subjects within social sciences. Through different guidelines and objectives in the labeling process, our model's versatility represents an initial stride towards streamlining this kind of research, effectively reducing the labor-intensive demands currently associated with analyzing large datasets.

Keywords: Populism, Text-as-Data, Hype Machine, Portuguese Political Actors, Natural Language Processing, BERT.

Contents

List of Figures

List of Tables

List of Abbreviations

AI - Artificial Intelligence

ANN - Artificial Neural Networks

API - Application Programming Interface

BERT - Bidirectional Encoder Representations from Transformers

BERTimbau - BERT-base-portuguese-cased

CDS - Centro Democrático Social

CH - Chega

CRISP-DM - Cross Industry Standard Process for Data Mining

DNN - Deep Neural Network

FN - False Negative

FP - False

IAA - Inter-Annotator Agreement

ID – Identifier

IL - Iniciativa Liberal

JSON - JavaScript Object Notation

L - Livre

MLM - masked language modelling

NLP - Natural Language Processing

NSP - next sentence prediction

PEV - Partido Ecologista os Verdes

POS - part-of-speech

PP - Partido Popular

PS - Partido Socialista

PSD - Partido Social Democrata

TN - True Negative

TP - True Positive

Introduction

“Portugal is the country where populism does not get votes. It is where no political power uses questions like immigration or others as a valid argument, which today, in the European sphere, shows immense relevance and I am extremely proud”, said António Guterres on the day he took office as Secretary-General of the United Nations, at the wake of the new year in 2017 (Diário de Notícias, 2017).

It was less than 10 years ago that Portugal appeared to have the most resilient party system in Southern Europe (Ferreira da Silva, 2019). With Europe recovering from the 2008 Global Financial Crisis, and with aftershocks being felt throughout the world with Brexit, the election of Donald Trump in the United States, Jair Bolsonaro in Brazil and Erdogan in Turkey, that the stability of the traditional powers in the Portuguese party systems were seen as standing strong. The centre-left party *Partido Socialista* (PS) had just recaptured power in 2015, away from the previous coalition of the center-right party *Partido Social Democrata* (PSD) and *Partido Popular* (PP) and, unlike the rest of Europe, no outside challenger party had appeared during the major years of the crisis, with established left-wing parties, such as *Bloco de Esquerda* (BE) and *Partido Comunista Português* (PCP), functioning as important aggregators of popular discontent (Mendes, 2021), while maintaining their own identities.

In Portugal, to political scientists like Mariana Costa Lobo, André Freire and Carlos Jalali, the question of a possible emergence of a far-right power was “inexistent and irrelevant” (Lopes & Henriques, 2016) and, in the aftermath of the European Parliament elections of May 2019, Portugal was proclaimed as the “Country without far-right presence” (Sapage, 2019). The perceived “Portuguese exceptionalism” (Zúquete, 2022), in the face of the new global "wave" of populism led to hasty analysis being produced against the rise of populism in Portugal. Unlike other cases in Southern Europe, where the impact of the 2008 crisis was paramount in explaining the emergence of populist forces, Portugal had experienced an improvement in democratic satisfaction in the years leading up to the 2019 election (Mendes, 2021). This, combined with the burden of the past Salazar dictatorship, some of the lowest levels of immigration in Europe and with unchanging levels of concern with corruption, made a strong case for the blockade that the contextual factors made to the emergence of populism in political discussions in the country.

Another prominent factor of the 2019 reality of the portuguese political sphere was the new and emerging use of social media by political parties. After playing a prominent role on the Bolsonaro and Trump elections in previous years, 2019 was the first time that any proper attention was paid to how much political parties were spending on social media marketing, with values being significantly lower than in other European countries (Pereira, 2019). This combined with the emergence of new, more internet savvy parties that saw in this medium an alternative to tradition mainstream media, where "For a relatively low cost, often just developing good content is enough to reach tens or hundreds of thousands of people" (Guimarães Pinto, 2019). "We were born a digital party. The most important work within the party is probably the digital communication group" emphasized Carlos Guimarães Pinto, at the time the leader of the *Iniciativa Liberal* (IL) party. With no deontology code for social media behaviour, the appeal made by BE in late 2018 for its creation was quietly ignored by all parties, and the amateurish level of the field (Barros & Espírito Santo, 2019), social media was still taking its first steps into this sphere.

Much like many other services available online, social media are mostly free to use. However, they do require big investments of time, which is time not spent doing something else, making time a scarce resource. In a free-of-charge environment, which characterizes much of the online business models, companies compete for people’s attention, in what has been coined as the “attention economy” (Goldhaber, 1997). The same way this new kind of economy impacted business models, as running on engagement means maintain attention spans to maximize profits, which in hand, profoundly impacted the dynamics in politics. Within this new age emerged a whole new way of doing politics, a transition in how socialization works, and profound changes to how people and political actors participate in political discourse. With new algorithms came new political strategies, new forms of communication, campaigning, voter behaviour, and the overall functioning of democratic processes.

In the realm of political economy, engagement metrics have played a pivotal role in shaping the landscape of emerging social media platforms and how political strategies change to fit them. The resurgence of populism from classic literature to becoming one of the most predominently used buzzwords of late 2010s politics can be widely attributed to its strong adaptation to the "attention economy”. Today, the term is prominently attached to any political actor who actively participates in leveraging the dynamics of engagement to gain traction and influence on these digital platforms, regardless of the politician beliefs.

However being a relatively untapped phenomenon, one that is highly complex and malleable, there is still a big theoretical focus around finding the best way to classify populism in a universal and standardized way. In the context of Portugal, research resources are still predominantly directed towards the challenging task of how to classifying populism within the country's reality. This process is highly labor-intensive, making it difficult to effectively integrate and cross analyse it with other crucial themes and variables, without extensive effort. The demanding nature of this endeavor still poses significant obstacles in leveraging it to its full potential across political economy and other academic disciplines.

Taking this into consideration, the overarching goal of this dissertation is to develop an reliable model using state-of-the-art machine learning and deep learning techniques. The objective is to create an automated system capable of accurately classifying text as either populist or non-populist, one that is specifically tailored to the portuguese context. This way we aim to overcome the problems of the resource intensive nature of always having to manually classify populism every time we want to use it as a variable, providing the groundwork for a more efficient and effective use of this concept in the future.

To do this, we explored the dynamics of political communication and how it intersected with populism in the context of social media. The choice to focus specifically on Twitter as the primary data source for this study stemed from its abundant availability of data and its unique characteristics as a social media platform, which made it an excellent mirror of the deeper portuguese reality. With that in mind, the study will be divided into four main chapters, each addressing essential aspects of the research.

Chapter 1 will provide the theoretical framework by firstly delving into the background of political communication, defining key concepts such as political actors, audience, media, and its relations to power. Secondly, we will also charecterize the different ages of political communication, leading up to the current “Fourth Age” defining its principle componenst, such as the Attention Economy and Hype Machine. Lastly, this chapter will offer a framework into the concept of populism, its ideational approach, ideological oppositions, and categorizability potentials.

Chapter 2 will present the methodology employed in the study. Expanding on the dissertation’s objectives and research questions, along with the design framework and analytical model that was used. The chapter will then further detail the resources utilized, particularly why and how Twitter was used as the primary data source. Finally, it will explain how the machine learning and Natural Language Processing techniques for data cleaning, pre-processing, and feature extraction, deep learning models, including Sentence Transformers and BERT, will be employed and how to assess their overall performance.

Chapter 3 will then focus on the analysis and discussion of the data. The process of data retrieval, annotation, and preparation will be described, followed by the modeling stage, including training and evaluation steps. We will then present the experiments conducted to address data imbalances and explore various models in search of the best performing one. Through a final discussion of the results and findings of the study, we aim to provide answers for the proposed research questions.

Lastly, chapter 4 will present our final remarks, concluding and summarizing the key insights obtained throughout our research, as well as the limitations felt and possible future research potentials.

CHAPTER 1

Theoretical Framework

1.1. Background on Political Communication

At the dusk of the 20th century, McChesney (1995) wrote that the political economy of communication studies was at a breaking point. He stated it would either flourish or fall further into the fringes of marginalized disciplines of other, more successful, social sciences. McChesney (2000), established a special relationship between political economy and communication, by placing them amid both capitalism and democracy, directly involved with commercial and material issues and ultimately concerned with social problems and governance. He would draw attention to the broader composition of society and how communication systems guide and influence existing social structures and relations, heavily prioritizing the study of communication in capitalist societies and commercial media systems. Following that first text, in the next 25 years, he would write more than a hundred successful articles and books on this topic. Both his books and the trajectory of his publishing career reinforced the idea that pollical economy of communication studies did not step aside and left "the rest of the social sciences sitting atop Mount Olympus, pondering the fate of society without him" (McChesney, 2000), but instead, climbed up and asserted its rightful place among them.

1.1.1. Defining Political Actors and the Audience

When delving into the realm of political communication, it is essential to establish a clear understanding of the term itself. Political communication should be viewed as an overarching concept that encompasses not only political rhetoric but also paralinguistic signs, such as body language, timing, articulation, intonation, and even visual means of signification (McNair, 2017). This type of rhetoric encompasses all forms of communication employed by political actors with the aim of achieving specific objectives. It encompasses both the communication initiated by these actors and the communication directed towards them, whether it takes place in public or private settings. In this context, political actors can be defined as individuals who seek to influence the decision-making process through the utilization of institutionalized political power, employing organizational and institutional means (McNair, 2017), be it individual actors, political parties, political organisations, or pressure groups.

The second thing that needs to be defined is the other side of the communication line, the one who is the target of this persuasion, the audience (McNair, 2017). The audience is the key element that gives any political message its relevance, as whatever its size, all communication is intended to achieve some desired effect on its receivers. The audience's role is highly debated, as they can actively decide what to do with the information that it is being passed to them. Scholars would either believe it to be a passive process, where they are mere puppets of groups who control the media, or a more active role, where they have free capacity to decode and interpret what is being told and create their own meanings (Carah, 2021). The debate over their role is also time-frame specific, evolving alongside political strategies and media advancements. Which leads to the last element of the political communication process, the media organisations.

1.1.2. Defining Media and Power

Media, as described by the Oxford Advanced Learner's Dictionary (2022), is "a way of communicating information". Stemming from the Latin plural of the word medium, where it had various meanings, all centred on the idea of something used as an intermediate or "between" two points. The study of media bases itself on the distinctive factor that humans create systems of meaning to orchestrate a shared life with each other which goes beyond just the capacity to communicate. The term media referring to mass communication was first coined by Marshall McLuhan (1969), in his book Counterblast, where he claimed media is an extension of the mechanisms of human perception, imitating their modes of apprehension and judgement and amplifying them into the macro level. He would also set them as art forms with the "power of imposing their own assumptions on the real world" (McLuhan, 1969), where they shape our new reality, as "acts have their existence, only to the extent that they are reported and received as messages" (McNair, 2017). This debate was then furthered with the emergence of new digital technologies, setting media as a set of institutions and technologies used to create, store, and circulate meaning in society, as well as digital platforms used to create, share, and access information online (Carah, 2021). The media, however, does not simply restrict itself to reporting what is happening in the political sphere in a neutral and impartial way. Having the ability to establish what we see as reality, what media constructs in its place is full of value judgments, subjectivities, and biases (McNair, 2017). But how?

Like any other factor under capitalism, the relationship between media and power is key. John Thompson (1995), and later Manuel Castells (2013) would define power as the ability to asymmetrically control and act in pursuit of one's interests, intervening in the course of events and affecting their outcome in ways that favour the empowered. The concept of structural power would then go beyond power over decisions, into power over agendas[1]. But how could one have power in media?

---

[1] **Agenda**: Is a particular program of action, underlying ideological and often one that is not directly expressed. It outlines the key goals, initiatives, and proposed actions that a party seeks to pursue in order to shape public policy and influence decision-making processes.

Power can be accessed by having material and cultural resources in order to get consent from others, by being in strong social positions that enhance the capacity to make others comply, and by controlling the structural institutions that secure social relations (Carah, 2021) with all options being equally valid. Media power arises from the capacity to control, use and distribute information and communication in order to create a shared reality (Flew, 2007). Access to power and influence in the realm of media is not readily available to just anyone who seeks it. Control over the production and dissemination of ideas is entrusted to certified institutions, which uphold standards to ensure the quality of content (Carah, 2021). These institutions establish criteria, known as credentials, which regulate and control access to legitimacy, serving as gatekeeping mechanisms, determining who has the privilege of participating and exerting influence within the media landscape.

Mass media, which emerged in the early 1900s, with the mass circulation of the printing press, consists of various institutions and established practices, that play a significant role in advocating what is often perceived as "common sense." Through its pervasive influence, mass media shapes societal norms, values, and beliefs by presenting ideas and information that are commonly accepted or regarded as rational and reasonable. By promoting certain narratives, values, and perspectives, the media implicitly constructs a shared understanding of what is considered normal or acceptable in a given society (Flew, 2007; Carah, 2021). However, it is important to recognize that mass media is not an autonomous entity but rather a mediated and reflective system that is influenced by other systems of power, particularly the prevailing economic and political forces (Flew, 2007). The media acts as a mirror, reflecting and reinforcing the ideologies, interests, and positions of those in power. It aligns itself with the dominant narratives and agendas that serve the interests of the current economic and political establishments.

The relationship between modernization, social changes, and media advancements is complex. The evolution of communication practices, citizens' roles, and especially media formats in the post-war period led Jay Blumler & Kavanagh 1999) to divide the development of political communication into three different stages, each with its own organizational principles and characteristics for whom and what controls media. Later, different scholars like Nicholas & Clark (2013), and Tanase Tasente (2020) added a fourth stage to the later typology, based on the explosion of new media advancements in the early 2010s.

1.2. The Ages of Political Communication

1.2.1. The First Age of Political Communication

Taking place during the proclaimed "golden age of parties" of the '40s up until the early '60s (Janda & Colman, 1998), Blumler & Kavanagh (1999) defined the first age as a timeframe where the political system experienced strong public credibility and was itself the main source of debate in politics. Traditionally, people's political engagement and communication was primarily shaped by enduring long-standing party affiliations. In this context, political institutions and the partisan "herd" played a central role in disseminating and reinforcing messages through word of mouth (Blumler & Kavanagh, 1999). Most people related to politics through long-lasting affiliations and, as such, communication mostly came party-dominated by way of political institutions and the partisan "herd" who reinforced and passed the messages by word of mouth (Blumler & Kavanagh, 1999).

1.2.2. The Second Age of Political Communication

The Second Age began in the early '60s, after years of economic prosperity, cheapening costs of production, and the subsequent increase in television household ownership (TV History, 2014). With the growing prevalence of televisions in households, the potential audience for political communication expanded exponentially, even reaching those elusive "floating voters" (Tasente, 2020). With this expansion, nationwide channel television emerged as the dominant platform for political communication (Blumler, 1999). This transformative period, dubbed as the "television era" by Tasente (2020), brought about a significant shift, as television monopolized the media landscape and loosened the grip of partisan opinions on voters' loyalty. Rather than serving as a vehicle for specific party propaganda, television content evolved into multifaceted discussions during prime-time slots[2], striving to present politics with the utmost fairness, impartiality, and neutrality. Two major factors came from this increase in possible viewership (Blumler & Kavanagh, 1999). First, it led to the first traces of professionalization in political communication to mass audiences, where campaigns started being strategized, organized, and coordinated by campaign strategists. Second, the search for mass audiences led campaigners to distance their communication efforts from their core audience, in an effort not to put off a potential new, "more floating" electorate.

What at the time, Blumler & Kavanagh (1999) considered ongoing, the third age of political communication began at the end of the 20th century. Designated as the "digital era" based on the emergence of the internet, its acceleration of communication speeds, information abundance, and its proliferation of the, until then, 30-year monopoly television had on the communication medium into polarized and highly tailored communication mediums. The scattering of audiences into different channels and mediums led to a more complex marketplace than the ones seen in previous ages, forming inharmonious trends in the neutral political media sphere established by the limited channel television era (Tasente, 2020). Trends such as the rise of personalized social messages as central nodes

[2] **Prime-time slots**: specific time periods during television broadcasting that fall within the evening hours when the majority of viewers are available to watch television. These slots are considered valuable advertising real estate due to the larger viewership that they attract.

influencing political communication and opinion, the increased competitive pressures and narrowcasting led to more sophisticated methods of targeted and direct communication to key groups, using any media available (not just TV but magazines, newsletters, emails, etc.), and the new state of permanent campaigning, as people started understanding the importance of keeping constant communication processes had in maintaining popular support (Blumler & Kavanagh, 1999).

1.2.3. The Third Age of Political Communication

It was within this era, with the transformation of relationships between political communicators and the public that strong currents of populism started scarcely popping up (Blumler& Kavanagh , 1999). Derived from the expansion of media outlets, the declining importance of strong ideologies, and the growth of political marketing and commercial communication allowed for its emergence. While a predominant concept nowadays. In his time, Blumler & Kavanagh (1999) were unsure if the populist “wave” would be empowering, corrosive, or merely play a symbolic part in political communication. But we shall leave that debate for now.

With the continuous evolution and expansion of the internet and the broader digital revolution of the early days of the 21st century, the ages typology proposed defined by Blumler & Kavanagh (1999) needed to be revisited. Tasente (2020) proposed the addition of a fourth age in political communication, termed "the new media era." This new digitalized era is marked by a sheer substantial surge in the creation and sharing of digital information, as well as an exponential increase in participating actors immersed in media (McChesney, 2013). It is also characterized by a transition from a predominantly exclusive, one-dimensional, and tightly controlled communication mechanism to a two-dimensional interpersonal regime, where both the message emitter and receiver play active roles (Tasente, 2020). In certain aspects, it even evolves into a three-dimensional communication system, facilitating communication in various directions as individuals connect online and transmit messages through "many-to-many communication" (Tasente, 2020).

Absent in his theory, no specific factor that changed the political communication landscape from the third to the fourth stage was given. Tasente (2020), only presented general structural changes, pointing to three particular changes as key reasons for bringing forth this new era, the broad adoption of Internet business models, the transition into the mobile space environment of smartphones, and the hegemony of Web 2.0 interactivity and social media.

When precisely pinpointing a definitive starting point for this fourth wave, McChesney (2013) provides an insightful analysis in his book. He argues that the internet played a transformative role by turning media and communication inside out, liberating them from media’s traditional formats. Notably for classifying this timeframe, McChesney (2013) highlighted that many political economists overlooked the significant influence of capitalism in "domesticating the internet" during the period from 1995 to 2011. While in previous eras, the role of capitalism in structuring the traditional top-down communication system flow (elite political actors – media companies, and channels – the audience) was fairly understood, a novel digital technology led to a discourse where the mainstream political actors no longer had control of the information flow (Tasente, 2020). The internet was seen as an almost pure "free market" and synonymous with democracy and freedom of speech (McChesney, 2013).

1.2.4. The "new" Fourth Age of Political Communication

During the course of the third era of communication, the internet gradually emerged as the undisputed preferred medium for people to engage with media. Although it took some time for the economic powers to adjust, the first decade of the 2000s witnessed a notable shift as economic forces began to populate and colonize the digital realm. This transformation was marked by the rise and subsequent dominance of new industry giants, including Facebook and Google, which seemingly appeared out of nowhere and came to control the emerging digital communication market (McChesney, 2013).

In an article written for The Atlantic, Jonathan Haidt tackled the popular idea that the "past 10 years of life have been uniquely stupid" (Haidt, 2022). In his take, he made the analogy where the rise of the internet to a modern-day interpretation of the tower of Babel[3]. In his article, he would call back to the early days of the internet in the '90s, when Web 1.0 software developed and evolved over time. This era of web development, defined by the c reation of the world wide web and its information interconnectivity, provided limited interaction between consumers (Choudhury, 2014). Applications like chat rooms, message boards, and email were the first wave of social media (Haidt, 2022), and were relatively harmless, serving to maintain social ties unconstrained by time and distance. His description of these times coincided with McChesney's (2013) view on internet literature at the time, which he named the "celebrants" of the internet.

The celebrants were leading experts who speculated that the internet had revolutionized and inaugurated a new era of utopic cultural democracy, one irrevocably better than previous ages, where sovereignty over media fell solely on its users, killing off "old media leviathans" (McChesney, 2013). Scholars, like Clay Shirky (2010), wrote that new generations would use the medium of the internet to shrink the universe, fostering better global understanding and use, consume, produce, and share media side-by-side, completely open and freed from big corporations (Shirky, 2010). In his article, Haidt (2022) wrote that this Collective intelligence and techno-democratic optimist reached its peak

[3] **Tower of Babel**: in the book of Genesis, the descendants of Noah, in order to "make a name" for themselves, built a tower whose top pierced the heavens. God (Yahweh), so offended by the hubris of humanity, destroyed it, and took away their ability to speak the same language, leaving the surviving humans to wander in its ruins unable to comprehend each other (Hiebert, 2007).

around 2011, with events such as the Arab spring[4] and the Occupy movement[5] demonstrations around the world, signalling that moment as the year that "humanity rebuilt the Tower of Babel" (Haidt, 2022). Haidt pointed out that, from there, things ended up falling apart. In 2012, as Mark Zuckerberg prepared to make Facebook go public, writing to his investors that he would "rewire the way people spread and consumed information" (Haidt, 2022). This would mark the moment that the "tower of babel pierced the heavens", and as such, is the moment we place as the first key factor for our new fourth era.

Around the same time, our media platform of choice was also changing. Just as the computer would shift media from television, the smartphone's ability to search the web produced a new shift to a mobile space in the communication industry (Nicholas & Clark, 2013). Smartphones are portable personal "pocket computers" with multipurpose functions, and according to Statista data (Gartner, 2022), the world saw a jump in smartphone sales between the years 2009 (172.38 units) and 2014, when it finally broke 1-billion-unit sales (1,244.4 units). Just as the transition to digital transformed the way we sought, reeded, trusted, and consumed information, the transition to a mobile environment liberated information beyond physical spaces like the office or the library (Nicholas & Clark, 2013). The erasure of physical delimitations led to its embeddedness in daily life, as one could access information whenever it wishes, which in itself trickled into a time shift beyond the need for operating hours (Nicholas, 2013). Within the mobile space, people could meet their information needs at the exact time of need, without any constraints. This erasure of time and space led media communication to become a 24-hour affair, "always on, always with us, constantly interrupting our lives with status updates and messages" (Haral, 2020). This shift presents the second key factor for our fourth age.Lastly, we have to talk about the emergence of Web 2.0 and what it meant for the online space. As the second generation of web development, its evolution from Web 1.0 allowed for the rise of large-scale collaborative and participative social interactions (Choudhury, 2014). Logghe et al. (2016) described Web 2.0 as the beginning of a world beyond static web pages, as an interactive network, one which spanned all connected devices in a user-generated, content creation model. This network was an ever-expanding, continually updated model that took advantage of content created by its users, who interacted, collaborated, and dialogued in virtual communities. This networking platform not only enriched user experiences but provided the ability for these users to create their own data and

[4] **Arab Spring**: Series of pro-democracy uprisings and social movements that took place across several countries in the Arab world, predominantly in North Africa and the Middle East, starting in late 2010. Social media was heralded as playing a crucial role in enabling communication, mobilization, and coordination of protests and uprisings.

[5] **Occupy movement**: Series of protests that utilized decentralized structures and temporary encampments to protest economic inequality and corporate influence. They employed social media for coordination and emphasized nonviolent civil disobedience to draw attention to socioeconomic issues.

services, one which other users could remix, and so on, on a nearly infinite level of expansion (see Logghe et al., 2016 and Mechant, 2012).

Within this Web 2.0 space, pivotal attention has to be given to the specific concept of social media. Social media, as its name would intel work as enhancers of human networks, they embody the interactive model of Web 2.0. Platforms whose primary objective is for individuals and communities to share, create and discuss any kind of content, be it ideas, interests, or other forms of expression, no matter how niche they seem (Kietzmann et al., 2011). Central to the concept of social media is its role as a framing device for social changes in the 2010s (Serra & Gonçalves, 2016).

Going back to the previous, tower of Babel metaphor. There was a belief that corporate communication had been fully democratized (Kietzmann et al., 2011). Social media had taken power away from those in marketing and public relation positions and given them to online communities to freely play around. Communication about specific brands would happen without the need for permission from their parent companies (Kietzmann et al., 2011), nor consent from the people/celebrities being talked about. While the era of Web 2.0 had been around ever since Tim O'Reilly (2005) coined the term in 2005, it was only around the time the other two key factors were happening, that we started to see evidence that pointed to social media being taken full advantage of in a serious manner, regarding corporate and political communication (Serra & Gonçalves, 2016). With social platforms like Facebook, YouTube, and Twitter, arguably the most famous platforms of this era, experiencing their first booms in userbases between the years of 2009 and 2013 (Haidt, 2022), we have established the first part of this last key factor. The second part of this transformative "boom," similarly to the allegorical tale of Babel, was when the internet's utopic tower experienced fragmentation due to the intensification of its viral dynamics (Haidt, 2022), and the grasping of the intricacies of these emerging markets by tech companies, adapting themselves, accordingly, becoming aware of what we now recognize as the "attention economy" (Goldhaber, 1997).

1.2.4.1. The Attention Economy

The first of these changes happened with the introduction of the "like" and "retweet" buttons. The like button was introduced by Facebook in 2009 and was a feature that gave users a way to publicly "like" other individual posts, the "retweet", introduced that same year by Twitter, did the same as the like button while also publicly sharing its contents to all users in their community (Haidt, 2022). Both these features were seen as massive successes for engagement and quickly became standard features of most other platforms (Haidt, 2022). The second change was the modification concerning the core economic principles of the algorithms that these platforms utilized when showing users content. These new attention and engagement metrics[6], would forever change how algorithms were being constructed in order to bring users the content they would be most likely to generate “likes and retweets” or any other kind of interaction (Haidt, 2022).

As business began to grasp on the market-base pillars of the internet, a whole new set of economic laws were established. The attention economy rested on the assumption that the online market did not suffer from the real-life scarcity problem of producing and distributing the goods. As such it looked into rewriting its economic focus into what was now an increasingly scarce and more valuable form of currency (Kubler, 2023). Attention impacted behaviours, which influenced users decision-making process, which is a highly valuable resource. These businesses learned that collecting their user’s behaviour, storing it as data, which could then sell as a commodity was a highly profitable business model (Kubler, 2023). As the demand for attention data became increasingly more valuable in all spheres, the development and fine-tuning of new, highly personalized engagement algorithms, where value was assigned by asking “Who has attracted, and continues to attract the most eyeballs?” (Kubler, 2023), revolutionized the approach to social media (Huszár et al., 2022) into an unstoppable personalized machine, which would lead to the third key factor for this fourth age of communication.

The innate human condition of dividing ourselves into teams based on our communities that have been studied throughout history (Haidt, 2022), has been amplified by the newly tuned platforms of this fourth age. These platforms were almost perfectly designed to bring out the most moralistic and less reflective selves of the individual (Haidt, 2022), micro-targeting information to groups at a previously impossible scale, in a way that would amplify and weaponize the more frivolous and emotional capacities of individuals (Haidt, 2022). This would be the downfall of the internet’s tower of Babel (Haidt, 2022), as the economic possibilities for companies, plus the triviality of the animosity between groups would speak infinitely louder than the utopic dreams of cooperation toward a common good. When directly applying this basis to political communication, this specific era would then be characterized by the chipping-away of trust in different political groups outside of our community, their political leaders, their every decision and where “every election becomes a life-and-death struggle to save the country from the other side” (Haidt, 2022).

1.2.4.2. The Hype Machine

While Tasente (2020) referred to this era as "the new media era," an alternative and more captivating name was given by Sinan Haral's (2020) description of society in his thought-provoking book. Haral

[6] **Engagement metrics**: A metric to determine the level of affinity in a relational process that facilitates communication, interaction, involvement, and exchange between community members for a range of social and organizational outcomes.

(2020) described this social age as a "gigantic information processor," where countless ideas, concepts, and opinions flow among individuals like neurons in the brain (Haral, 2020). Within this global information network, a central force emerges, known as the "Hype Machine," that represents the driving force and defining feature that gives name to this era.

Haral (2020) theorized the Hype machine as the information processor that regulated and directed the flow of information in society, “between people, brands, governments, media outlets, and international organizations” (Haral, 2020). This processor, enabled by the established algorithms, worked as the collective outcome of experiences and decisions of individuals, which in turn had already been influenced by the previous flow of information, creating a continuous spiral (Haral, 2020).

Enabled by this new information order of a constantly evolving network, conducted by twenty-four-seven human interactions, the Hype machine deeply transformed how information is produced and consumed (Haral, 2020). In order to understand how one must first understand the parts of this machine. The Hype Machine comprises three major parts: its substrate, medium, and process (Haral, 2020).

The first component has already been previously developed. The substrate at its core consists of the social networks themselves (Haral, 2020). As its name says, the substrate is what gives shape to any organism, and as such, in this case, it’s the evolving and growing populations of social media networks that shape information flows by structuring our digital connections into interconnected networks, like invisible spider-webs of data (Haral, 2020). Sociologists Paul Lazarsfeld and Robert Merton (1954) introduced the concept of homophily as primordial patterns of human relations in the physical world, where people showed strong tendencies to become friends with those who shared common attributes or similar consistent values. These patterns showed that homophily eased communication between social actors, made it easier to predict others’ behaviours, and developed trust more easily when compared to individual’s whit whom we did not (Lawrence & Shah, 2020). This concept is extremely important once we understand that those social media algorithms, when determining who, in large part, people connect to are following this homophily principle (Haral, 2020). The algorithmic tendency for the formation of homogeneous social networks then leads to a smaller diversity of opinions, information, and ideas, as such, when two dissimilar groups meet, there is a propensity for political polarization, social gridlock, and the spread of misinformation and hate speech online (Haral, 2020). All this keeping in mind that these networks work at a never-before-seen scope and speed.

The second component has also been developed extensively before, therefore not much will be developed here. The medium through which the “Hype Machine era” learns about and influences us are the mobile space of smartphones (Haral, 2020). Its uniqueness as a medium to be always on and almost always with us is what gives the machine a way to uniquely get to know and understand its actors.

The last component is its most crucial when it comes to understanding the inner works of the "hype machine" and the political and cultural zeitgeist of this era. The "hype loop" is what Haral (2020) described as a cause-and-effect cyclical process where the interplay between machine intelligence algorithms and human behaviour is constantly and repeatably. This cyclical pattern works as a constant feedback loop, where technology analyses what is happening inside its medium, like what users consume, what they share, what they like, and whom they choose to interact with. It then optimizes engagement and viewership in its algorithms specifically for each user's specific preferences, (Haral, 2020) structuring their reality by constraining what, when and how they get to see the information. These unprecedented amounts of information available in these social networks lead to a "poverty in attention" (Haral, 2020), as people don't have the time or attention to search more broadly than the curated set of options the platform is feeding them (Haral, 2020). The second part of this loop is in how human agency shapes what information the algorithms analyse, how it interprets what we want from it, in order to make suggestions (Haral, 2020). When specifically applying this phenomenon to political communication, two related factors came up, what Haral (2020) calls the "feed algorithms" and de subsequent "consume and act loop".

Feed algorithms shape how we see by recommending the content we consume, and as such determine, to a great extent, what, when and how we know things (Haral, 2020). As the number of contents began to greatly outnumber users' cognitive abilities to consume it the algorithms had to start ranking and prioritising specific content for each individual, in order to only show them what, in essence, would be what is most relevant to them. Since all social media companies are profit corporations, money plays an important role in these feed algorithms, as more engagement drives revenues up. Their objective is also to suggest or nudge users in directions that will maximize its revenues and profitability (Haral, 2020). With a business model that represents "over 65 percent of digital advertising" (Haral, 2020), firmly planted on the monetization of highly curated audiences, the better the algorithms can find an audience that fits the profile of the specific brand, the more the audience will be susceptible to its services and products and the higher the return on brands' marketing investments (Haral, 2020). As the economy became more and more aware of this the more it began to dominate the content present in social media.

A second important factor in feed algorithms is how emotionally charged social posts are more likely to get more engagement and therefore be more disseminated and promoted by the algorithms, be it positive or negative in sentiment (Stieglitz & Dang-Xuan, 2013). This is believed to be because content that we feel strong opinions and feelings induces stronger cognitive and arousal-related effects that makes us more susceptible to engage in social media (Stieglitz & Dang-Xuan, 2013). This bias to engage is then incorporated by the loop, which learns what we are more likely to engage with, and as such begins to feed us more of it (Haral, 2020). As machine recommendation algorithms create filter bubbles of polarized content consumption, we start to mix more emotions in political debates. This mixed with the tendency for people to view information as more credible if it reaffirms their pre-existing beliefs, or to research topics already searching for their position to be proven correct and discredit information that does not (Metzger & Flanagin, 2013) leads people to be caught into what Niall Ferguson (2019) called an “emocracy”. While emocracy as a concept has hardly been defined in political economy, little to no mentions of it can be found while researching as well as Ferguson (2019) itself mentioned not finding any in his research as well. The idea, however, has been planted in mainstream political debate for a couple of years now. One would not take a long time searching around politicized social media communities, before reading someone quoting Ben Shapiro’s (2018) famous catchphrase “Facts don’t care about feelings”, a phrase, often used as a rebuttal by Shapiro, and as such adopted by his audience, that more or less means that when debating something, emotions cannot overturn scientific facts. If one agrees or not with the political views of Ben Shapiro (2018) is indifferent to this debate, what is in focus is exclusively on the perception that emotions have already become deeply rooted in political debates cannot be understated. Ferguson (2019) described “emocracy” as a state of democracy where emotions and feelings matter more in debates than reason and majority rules, where the stronger the feeling, and the better someone is at weaponizing emotions, the more influence it is able to have (Ferguson, 2019). This concept becomes especially problematic once the debate leaves the general population and joins political strategies, where political actors start employing tactics that prey on people’s emotions in order to further their political agendas (Ingruber, 2021).

The consume and act loop is then the process where we act on the (things the algorithms have recommended to us. As social media is consuming and constructing our ideologies, ideas, as well as psychological and physical needs for socialization, belonging, and social approval (Haral, 2020). Continuously being fed by an economic network that perpetuates its growth led us to the economic, social, and political climate that we are living in this fourth age. This loop is the destruction of the internet’s tower of babel, information became so large and so curated it dissolved the mortar of trust, and information and reality were no longer presented to a single “mass audience” that spoke the same language, but to the remanences of people that were left once the tower had been broken, each living in their own “bubble” of reality unable to see, speak or understand people who stood outside of it, as the algorithms did favour these interactions (Haidt, 2022).

In an era as described by Haidt (2022) as a political stage where theatrical performances crushed “competence and rationality”, where reality is no longer intangible and only constant as long as the loop has not evolved past it, one where digital media has made everything so fast, layered and more multidirectional, that “certified” gatekeepers institutions have irreversibly lost control over the information received by audiences. When media prioritizes profits, users become the fuel for the attention machine, which main aim is to prolong engagement, irrespective of the content. Ultimately it makes no difference to the social media “hype machine” what kind of content is being circulated, as long as the algorithms perceives it to be good for engagement.

One such of these topic that took advantage of the dynamics of this new era and grew into the mainstream lexicon and right into the centre of political debates, was the populism phenomenon.

1.3. Background on Populism

Populism has garnered significant attention in recent years, particularly following Brexit and the 2016 US presidential election, leading to its designation as the word of the year in 2017. However, this surge in interest is not an isolated phenomenon but rather the culmination of a growing fascination with populism over the past decade, emerging as one of the defining themes of contemporary politics. Its meaning and understanding vary significantly depending on geographical location. In Europe, populism is often characterized by its association with right-wing ideologies, which prominently feature anti-immigration and xenophobic sentiments. Conversely, in Latin America, populism is more closely linked to issues such as economic mismanagement, corruption, and clientelism.

Today, there seems to be an increasing number of prominent political figures rising to prominence within the umbrella of populism, captivating audiences with novel ideas and charming use of rhetoric. This begs the question of whether there is any actual substance behind their catchy statements, whether they only aim to amass followers or also appear to bring something genuinely innovative to the table. There is some underlying negativity associated with this phenomenon, in the way, it classifies/catalogs who uses it, in the way it is generally perceived as a threat to democracy, as well as the idea that it is ideologically too vague to be accepted as conceptually valid.

1.3.1. The Ideational Approach to Populism

Mudde (2004), triggered a shift in this perception by arguing that populism was not anomalous and a negative niche in politics but was becoming a mainstream part of politics in Western democracies of the 21st century. Although there has been vivid debate about the “correct” definition of populism, among the differing interpretations, Mudde's (2004) ideational approach to populism has emerged as the most widely accepted and consensually adopted one. His definition lies on a set of related ideas about the nature of politics and society, conceptualizing populism as a thin-centered ideology focused on a binary view of the world relayed by the dichotomy of the “good people” versus “the elite” and advocates for better representation and greater scope for the expression of the general will (vox populi).

An ideology is a body of normative ideas regarding human nature and how one should run and function in society. Unlike the strong ideologies that characterized 20th-century politics, such as socialism and neo-liberalism, which are theoretically dense in their conceptualization, a thin-centered ideology lacks this denseness, and as such, it works as an interpretative framework that usually comes attached to stronger ideological elements that are not necessary or exclusive to it (Mudde, 2004; Gidron & Bonikowski, 2013; Kaltwasser et al., 2017).

On its own, populism lacks the conceptual and theoretical strengths to answer complex questions and does not have, nor does it present, coherent solutions to the widescale social, cultural, and political problems that political parties are normally expected to be able to answer (Norris & Inglehart, 2019; Taggart, 1996; Gidron & Bonikowski, 2013; Kaltwasser at al., 2017). Populism as an ideology rarely exists on its own, it does, however, work as an apprentice to more complex ideologies, being able to attach itself to their nuclear concepts and relating them to its dichotomy based on the specific societal contexts, timelines, and realities it emerges in (Mudde, 2004; Kaltwasser et al., 2017).

In the context of a thin-centered ideology, populism is considered distinct from traditional left-right cleavages. Typically, political parties are classified along the left-right spectrum based on their policy positions, ranging from the economic left (advocating regulated markets, state intervention, wealth redistribution, and public spending) to the economic right (favouring deregulation, free markets, opposition to wealth redistribution, and tax cuts) (Guriev & Papaioannou, 2022; Mudde & Kaltwasser, 2017; Gidron & Bonikowski, 2013). Populism, however, transcends this spectrum, as it aligns with parties across the entire ideological range. Following this conceptual approach, populism operates independently of a party's position on the ideological spectrum, making its core concepts detached from traditional party ideologies (Mudde & Kaltwasser, 2017; Gidron & Bonikowski, 2013). On this note, its three nuclear components are the dichotomy between “the people”, “the elite” and the “general will”.

1.3.1.1. The People

Much of the focus on populism revolves around the notion of "the people." However, there is an ongoing debate about the concept of "the people" and whether it represents an identifiable group or a socially constructed entity that depends on context and reality (Canovan, 2015). Despite its versatility and flexibility, "the people" is generally comprised of three major consensual components. First, it encompasses the idea of popular sovereignty, wherein the people are regarded as the ultimate source of political legitimacy and power in modern democracies (Mudde & Kaltwasser, 2017; Norris &

Inglehart, 2019; Gerbaudo, 2017), one which, when ignored while ruling, can potentially lead to social unrest and revolts.

Secondly, it is the symbolic manifestation of the group of the common people of the nation, or in other words its “heartland” (Canovan, 2015). This concept refers to the combination of a broad concept of cultural traditions and values, Weberian classes, and what Bourdieu calls “status” (Gane, 2005; Flemmen, 2013). These last two provide distinct perspectives on how social stratification should be done. Classes, as conceptualized by Weber (Gane, 2005), focus specifically on the economic relationships of organized groups, emphasizing the role of economic factors in shaping social positions. On the other hand, Bourdieu's concept of "status" (Flemmen, 2013), encompasses a broader understanding of social power, which extends beyond a person’s economic standing and the mere accumulation of wealth. It recognizes that social power can be derived from various sources, such as cultural capital, social connections, and symbolic recognition, in addition to economic resources. (Gane, 2005; Flemmen, 2013). It is this imprecise classification that lies at the core of the ambiguity surrounding the definition of "the people".

Lastly, it is essential to clarify that "the people" represents a distinct concept separate from merely the natives of a specific country (Canovan, 2005). The native population can consist of diverse ethnic and cultural groups, not all of which necessarily align with the targeted group referred to by the populist. This complexity prompted Paul Taggart (2017) to introduce the term "the heartland" as a means to better encapsulate what populists often express in their rhetoric.

"The heartland" is a term detached from any particular class or status group, but rather denotes the combined self-identification of the targeted community. This linguistic shift seeks to capture the collective sentiment and shared identity emphasized by populists when referring to their support base (Mudde, 2004; Taggart, 2017). By acknowledging their demand for a better, more dignified, representation in public political discourse and by recognizing their subjective and indirect exclusion from positions of power and political views and identities, populism forges its own notion of the people by various factors, which vary based on necessities and countercultural movements[7]. It considers people’s interactions with politics and democracy and uses the concept of "the people" as a rhetorical instrument with the power to be simultaneously inclusive and divisive. This idea juxtaposes the notion of people as a silent majority against a well-defined common enemy, highlighting both unity in the people and the broader division within society and its political priorities (Mudde, 2004; Hawkins, 2009).

1.3.1.2. The Elite

[7] **Counterculture:** subcultures that emerge in opposition to mainstream societal norms, values, and ideas. They challenge the established political discourse providing platforms for resistance, dissent, and critique, in the pursuit of social and political changes.

On the opposite side of this “us versus them” dichotomy, “the elite” is less controversial in its vagueness but just as debated. Populism considers the current concentration of power in the political and economic spheres in modern democracies, grouping the political establishment, the traditional media as well as the economic elites as a homogeneous group that acts both morally and physically, to protect and extend their influence and interests (Mudde & Kaltwasser, 2017; Norris & Inglehart, 2019). This dichotomy heavily emphasizes the dimension of morality, employing rhetoric to pit the virtuous and righteous people against what it perceives as a corrupt and unethical elite in an oversimplified fashion (Mudde & Kaltwasser, 2017; Norris & Inglehart, 2019).

"The elite" is depicted as antagonistic in nature, not only due to their disregard for the better interests of the people, whom, as established, hold the true legitimacy in politics but also actively undermining that factor in order to consolidate their own power. They leverage the mechanisms of liberal democracy and trustee representation to reinforce their stranglehold over the oppressed population (Mudde & Kaltwasser, 2017; Norris & Inglehart, 2019). The framing of "the elite", similar to the principles mentioned for “the people”, highlights their distinctiveness beyond national boundaries and encompasses cultural traditions and values tied to their status and influence. This moral dichotomy reflects Manichaean[8] discourse, that fosters an oversimplified 'good-versus-evil' perception of politics. Populism paints a picture of the elite, depicting it as having hidden, ill-intentioned, and malevolent forces, that should not only be ousted from power due to the unfair inequality that arises from their exclusive decision-making but also because they symbolize and embody all of the evils in society (Castanho Silva & al., 2017; Oliver & Wood, 2014; Žižek, 2006).

This tendency is particularly pronounced among populists who have not yet attained power, using the elite as a convenient scapegoat for their political struggles. However, even when in positions of power, they often continue to employ a populist rhetoric, in order to highlight how the pervasive influence of the elite hampers their every move and their political agenda (Mudde & Kaltwasser, 2017; Norris & Inglehart, 2019; Castanho Silva et al., 2017; Oliver & Wood, 2014).

1.3.1.3. The General Will

The last nuclear concept of this ideological approach is that of the "general will.". The origins of this concept can be traced back to a political economy article by Jean-Jacques Rousseau (1762), where he defined the "general will" (la volonté générale) as the collective ability of individuals to come together in a community, engaging in legislation and governance to pursue shared interests. When seeking to define the nature of this bond by which this group is created, Rousseau (1762) distinguished the people

[8] **Manichaeism**: is an ancient dualistic religious philosophy that posited a perpetual struggle between the powers of good (light) and evil (darkness) in the universe. It rejected the idea of a single, omnipotent God and instead proposed the existence of two opposing powers who use humanity as the by-product of their continuing battle (Tardieu, 2008).

as reluctantly political, tending to exhibit reluctance towards engaging in politics, often prioritizing their immediate concerns and misconceiving what is truly beneficial for them in the broader context (Norris & Inglehart, 2019). Additionally, he noted that people tend to be disinterested and preoccupied with other aspects of life, diverting their attention away from political matters.

Within the framework of democracy, particularly in the context of trustee representation, the concept of a "contract" between the government and the people emerges, wherein the government assumes the responsibility of managing and guiding politics in the best interest of the people. This "contract of submission" signifies an agreement where the subjects entrust their direct decision-making power, consenting to be ruled (Rousseau, 1762; Norris & Inglehart, 2019). Such surrender of individualism is viewed as necessary for preserving the rule of law, albeit without neglecting their individual needs. In all societies, this unwritten contract carries a moral tone, as insubordination and unwillingness to surrender these freedoms would lead to the insecurity and isolation of life outside society (Plotke, 1997; Rousseau, 1762). However, according to populism, current governments have failed to comprehend the general will, either deliberately or inadvertently. Populists argue that state power treats the people with passivity, only affording them limited influence that is felt solely during periodic elections, leading to the urgent need for a more engaged and participatory form of governance (Mudde & Kaltwasser, 2017; Norris & Inglehart, 2019).

As a result of these deficiencies, populism recognizes the necessity for a strong popular identity that demands increased political authority, which “lies not with the government, of whatever kind, but with the people as a whole, expressing the general will, and therefore sovereign” (Rousseau, 1762, p.199). Consequently, populists often exhibit an affinity for more direct forms of democracy and seek out institutional mechanisms that foster a closer bond between themselves and the people they represent. This inclination shapes their rhetoric, as they emphasize common sense, adherence to the general will, and the need for exceptional individuals to lead and embody the aspirations of ordinary people (Rousseau, 1762; Mudde & Kaltwasser, 2017; Norris & Inglehart, 2019).

1.3.2. Ideological Concepts Opposed to Populism

These ideas can seem abstract on their own. One can argue that in any kind of conceptualization, there is a need for distinguishability through the elimination of parts. The "either-or" logic[9] plays a role in ensuring theoretical and analysis coherence, enabling the identification of populism by distinguishing it from other ideologies which are antagonistic to itself (Kaltwasser et al., 2017; Mudde & Kaltwasser,

[9] **“either-or” logic**: Also known as disjunctive syllogisms, is a logical rule that allows us to infer the truth of one alternative in a disjunctive statement by excluding the truth of the other alternative. It states that if we have a statement that presents two options (A: Populism or B: Not Populism), and we can determine that one option (A) is false, then we can conclude that the other option (B) must be true.

2017; Weyland, 2001), enabling a holistic categorization. In this framework, one can discern two distinct ideologies that stand in direct opposition to populism: pluralism and elitism.

1.3.2.1. Anti-Pluralism

Populism and pluralism represent contrasting approaches to politics. Populism tends to adopt a simplistic, black-and-white perspective, often framing issues in binary terms. In contrast, pluralism recognizes and embraces the complexity of society, acknowledging the existence of diverse groups and individuals with distinct views, needs, and wishes (Dahl, 1978; Norris & Inglehart, 2019; Kaltwasser et al., 2017). Pluralistic societies prioritize multiculturalism and social diversity as valuable assets. They establish democratic institutions emphasizing inclusive decision-making and recognizing the heterogeneity within society, minimizing the prevalence of disagreements and conflicts over a supposedly, way too complex, solution that would fix all of society's problems (Dahl, 1978; Grillo, 1998; Norris & Inglehart, 2019; Kaltwasser et al., 2017). In contrast, populism tends to single out a particular group, that embodies the nation's identity, culture, and values, as the exclusive focus of governmental attention.

In this sense, Populism is inherently anti-pluralist due to its tendency to prioritize a specific group or segment of society over others. This exclusivity and emphasis on a single homogeneous identity undermines the principles of inclusivity, diversity, and equal representation that are fundamental to pluralism as an ideology (Galston, 2017; Grillo, 1998).

1.3.2.2. Anti-Elitism

The comparison between elitism and populism highlights certain shared characteristics, such as their simplistic binary outlook on politics and their shared criticism of societal divisions (Best & Higley, 2010). Elitism tends to downplay divisions based on social groups, viewing them as distracting and less relevant to effective governance (a need to focus on the "big picture"), while populism critiques societal divisions by highlighting social groups and institutions as part of the problem, accusing them of having "special interests" that prioritize their own agendas (Mudde & Kaltwasser, 2017; Norris & Inglehart, 2019).

Where elitism distinguishes itself from populism is in its approach to political compromise. While populism often rejects the idea of compromise, viewing it as a pathway to corruption and a betrayal of the people's interests, elitism recognizes the significance of political compromise and the necessity of finding common ground (Best & Higley, 2010; Norris & Inglehart, 2019). Elitism acknowledges that addressing complex societal challenges requires balancing diverse perspectives and making concessions for the greater good. It highlights that elites, by virtue of their unique position in society, are the only ones capable of being bound by institutional rationality and democratic processes, such as elections. As such, they are the only ones capable of bearing the responsibility to navigate complex issues and make informed decisions that benefit society as a whole over their own, personal interests (Best & Higley, 2010; Norris & Inglehart, 2019; Kaltwasser et al., 2017).

While also exhibiting binary tendencies, elitism focuses mostly on the importance of expertise and experience in governing when differentiating between the elite and the non-elite. Elitism often positions the elite as the sole repository of political wisdom and downplays the contributions and perspectives of those outside this privileged group (Best & Higley, 2010; Norris & Inglehart, 2019). In this binary framework, elitism may oversimplify the complexities of societal divisions by attributing all decision-making power and knowledge to the elite, while disregarding the valuable insights and contributions of the broader population, reinforcing social hierarchies, and undermining the principles of inclusivity and equal representation that pluralistic societies strive to uphold.

In this sense, populism is inherently anti-elitist as it fully rejects the idea of a privileged elite class holding a monopoly on political power and on the decision-making process. It fundamentally disagrees with elitism, positioning the elite as a disconnected and self-serving group that not only fails to represent the interests and aspirations of the broader population but fails to do so fully intentionally (Best & Higley, 2010; Norris & Inglehart, 2019). Instead of backing these elites, populist rhetoric full heatedly criticizes and challenges the authority and legitimacy of the elite, aiming to dismantle these established power structures and establish a more direct voice of the people, one which it represents.

1.3.3. Classifying Populism

One thing that is seemingly absent from the conceptualizations made by Mudde (2004) in his ideational approach theory and is usually neglected in studies of this sort is how to measure the concept of populism. The populist label is often attached to a certain party without any justification, or lacking rigor in its arbitrary application.

In Ostiguy's (2017), characterization of populism, he questioned this holistic, "all-or-nothing" characterization. In what he described as "flaunting to the low"[10], he posed that the subjective, identity-centered, and dichotomy elements that incorporate populism's ideology theories were inseparable from the notion of antagonistic[11] rhetorics and other norms of public discourse (Kaltwasser et al., 2017). He argued that populism is not an all-encompassing identity, but is, in fact, a political strategy that populist actors choose to apply when trying to establish a connection with

[10] **Flaunting to the low**: Is a political strategy of highlighting and amplifying socio-cultural elements and sentiments that are considered markers of inferiority or marginalization by the dominant culture.

[11] **Antagonism**: in political rhetoric, it refers to the use of confrontational and divisive language and strategies that highlight conflicts with the aim to mobilize support by fostering a sense of grievance, resentment, and hostility towards perceived adversaries.

certain segments of society through the use of a specific populist rhetoric. As such, in order to measure the populist phenomenon he suggested that it would make more sense to analyze “populist interventions” rather than populist actors or regimes when conducting research on populism (Kaltwasser et al., 2017) .

While Mudde's (2004) definition of populism as an ideology and Ostiguy's (2017) conceptualization of a populist political style are viewed as two distinct theoretical approaches to studying populism, it is important to recognize that characterizing populism as a thin-centered ideology does not exclude the possibility that it features a specific style of communication as well. As such is a suitable path to explore.

Following this theoretical conceptualization, we shall study populism as a thin-centered ideology that is reflected in the discourse used by political actors, we consider populism as an attribute of a text rather than a feature of any specific politician (Ostiguy, 2017; Pauwels, 2011). While populism’s lack of a concrete definition has posed challenges in establishing clear criteria for identifying and categorizing populist phenomena as either populism or not. We shall therefore adopt and adapt the one from the quantitative text analysis approach (Pauwels, 2011). The Quantitative Text Analysis approach is a research method variant of content analysis, where through analyzing large amounts of textual data one can identify patterns, trends, and insights of complex concepts. In this approach, texts are not treated as discourse to be understood and interpreted, but rather as data in the form of words and sentences that can be used to calculate patterns (Pauwels, 2011).

By integrating the principles of this quantitative text analysis approach with machine learning models, we can adapt the analysis framework to leverage the power of machine learning in order to extract insights and uncover strict semantic and grammatical patterns in the textual data that would, otherwise, be impossible or take an unsurmountable amount of time to catch with the naked human eye, ensuring a stricter guideline for future work on the field.

CHAPTER 2

Methodology

2.1. Objectives and Research Questions

In this research, communication is considered populist if it conveys populist ideas, and as such, the ideational definition lends itself to operationalization and measurement because it identifies features that should be present in discourse for it to be populist. Plus, being a widely cited definition of populism facilitates connectivity to existing research results and questions.

One of the biggest challenges in studying populism is finding ways to measure it across a large number of cases, from multiple countries and multiple parties within these countries (Hawkins & Silva, 2018). Most studies on the subject classify populism based on literature reviews or the judgments of the opinions of country specialists (Hawkins & Silva, 2018), these kinds of methods are not reliable when applied to an approach using primary data, and struggle to provide a strong framework for easy classification in most contexts.

Jagers and Walgrave (2007) were among the first researchers to apply content analysis to study populism, employing manual coding and holistic grading (i.e. classifying whole manifestos or speeches) in order to systematically retrieve reliable insides on these features. While this kind of methodology proved quite successful, being replicated by a plethora of other researchers in the coming years (Kaltwasser et al., 2017; Pauwels & Rooduijn, 2011; Hawkins & Silva, 2018), including research done specifically on the Portuguese reality, like the one done by Salgado (2019) and Mendes (2021). However, this kind of application required extensively trained annotators with a high knowledge of populism as both concept and reality (Gründl, 2022), which is highly time and resource consuming, and not always available.

To help resolve the limitations of manual coding, especially when it comes to large amounts of data, which are needed for any kind of conceptual generalization, computerized content analysis emerged as a suitable replacement for human-coded techniques. These methods are possible due to the spread of computerized tools, where the text is treated as data to be processed and analysed using the tools of quantitative analysis (Benoit et al., 2018), allowing social sciences to research beyond boundaries previously accessible to only data scientists, through Natural Language Processing (NLP). To do this, researchers like Castanho Silva et al. (2017) and Gründl (2022), used machine learning techniques to perform supervised classification of documents, where they developed a dictionary approach.

However, these dictionary approaches to measure populism, have, their own set of problems. They were limited by their theoretical and quantitative scopes, their reliance on words that directly indicate concepts, and are contingent on words that might be quite rare in the corpus (Gründl, 2022; Mendes, 2021). Resulting in occasions where populist communication might not be recognized as populist because the words used did not appear in the dictionary. On the other hand, words that are part of the dictionary might appear in moments that are not actually populist, and as such wrongly classified. Besides this, automated approaches are almost never applied to both the Portuguese reality and trained with the Portuguese language in mind.

Another factor is that most of the studies aim at mostly party manifestos (Pauwels, 2011; Mendes, 2021), ignoring the online sphere, which, as we will discuss plays a crucial role in current political and populist communication. With different political actors using social media as vehicles of less-institutionalized, direct communication (Salgado, 2019), populist rhetoric seems to be drawn to social media as it provides a platform to promote these populist ideas to their supporters without the established barrios of traditional media outlets (Gründl, 2022). These platforms, allow political actors to produce a great wealth of content every day, thus, social media has the utmost relevance to this study, providing larger amounts of data crucial in finding a path to reliably measure populist features in social media discourse.

With this in mind, through this first framework, we establish two research questions (RQ):

RQ1: It is possible to establish a reliable framework to capture populist moments in the Portuguese reality by using computerized tools?

RQ2: Can this framework reliably capture and characterize the associated discourse in the Portuguese reality through the use of Portuguese political actors' online posts?

2.2. Design Framework and Analytical Model

With the abundance of information online, it would be impossible to assimilate and understand texts containing billions of sentences. And even if that was possible, it would be impossible to organize it in a way that would allow for any comprehensive study without the application of new methods with the potential to analyse them.

Texts need to be used in a clear, precise and effective way in order to convey points from the social sciences, and in order to take advantage of the analysis potential of large corpora to allow new and important discoveries to be made (Grimmer et al., 2022), which were previously conditioned by the information barriers of smaller datasets.

The solution is to treat "text as data" to be processed and analysed through the tools of automated analysis, even if the "text" is not necessarily to be read, digested and summarized (Grimmer et al., 2022; Karišik, 2018).

As such, Large-scale text analysis applied in social sciences mostly relies on algorithms that have been initially developed in and for computer science and are closely related to the methodologies designed for that specific field. However, the adopted knowledge from computer science literature needs to be particularly adapted to our own work. The differences between fields are in how social science sees its discovery process, as even when methods are imported from computational tools, research will still aim to answer the classical social, institutional, and economic conceptualizations.

To analyse these new data sources and utilize these new techniques using a multidisciplinary approach, one must first reconsider the standard deductive approach present in social sciences research frameworks. In summary, Quivy's (1992) research framework for a successful social science analysis always follows the same three acts: Rupture, Construction, and Constatation. Rupture is where one departs from pre-existing knowledge or assumptions through further exploration and understanding, creating his hypothesis. Following that phase, construction is where one designs and implements his research idea, selects his methods and analysis tools, and retrieves his data. Lastly, Constatation, is the process of verification, where one draws the resulting conclusions based on his findings (Quivy, 1992). Quivy's (1992) methodology adhered strongly to the idea of having an establishing theory from which to derive a set of testable hypotheses before ever viewing or beginning to collect any data, in what is called the "impasse of the hypotheses" (Jacquinet, 2021). This concern that researchers might collect data indiscriminately or without a clear direction, can end up hindering progress, as it limits the opportunity to refine concepts, develop new theories, and evaluate hypotheses based on the data available (Grimmer et al., 2022). Although a research project begins with a focused question, it often evolves to a different focus, as such it is important to maintain a balance between exploration and research objectives for a successful outcome.

As such, when designing our framework, we incorporated the key principles outlined by Grimmer et al. (2022) for adapting text-as-data methods to the social sciences to overcome the established problems.

Firstly, the three steps defined by Quivy (1992) are still valid and in the same order, but they have been adapted to involve continuous feedback loops between data analysis and data retrieval processes (Grimmer et al., 2022), being able to return to previous steps if considered beneficial to respond to our established hypothesis.

Secondly, in our hypothesis, we are not trying to find the “best” underlying model for classifying populism. Instead, we are looking for a method that most correctly measures our hypothesis, whose relevance theoretically depends on the model that accurately and reliably measures the concept of populism (Grimmer et al., 2022). Paving the way for future research to no longer be burdened by this previous constraint.

Lastly, with those pre-established ideas, we will be slightly adjusting our methodology based on the Cross Industry Standard Process for Data Mining (CRISP-DM), which offers a baseline, for data mining processes, flexible enough to fit the specific needs of social sciences.

CRISP-DM aims to achieve replicable, reliable, and fast processes to avoid overly specific solutions, whereby following six sequential phases provides a structured and iterative approach to project development (Hotz, 2023), with those steps being:

- Business Understanding – focuses on understanding the resources available, objectives and requirements of the project.
- Data Understanding – Adds the foundation by identifying, collecting, and analyzing the data to access the suitability of the delineated goals.
- Data Preparation – Applies transformation techniques. Aggregating data into new classes to deal with possible imbalances and then experiments.
- Modelling – Selects and builds the model, and then assesses its performance based on domain knowledge and a set of predefined success criteria.
- Evaluation – Reviews the performance of the models and understanding of the results. With the new insights, it loops over the preparation and modeling to test new results.
- Deployment – Documents the results of the experiments so that the costumer can move on to his own post-project operations.

Since CRISP-DM phases are just a universal baseline, mutable and adaptable to individual projects characteristics, it facilitated the adaptation of its processes and respective nomenclatures in our developed methodology, as depicted in Figure 2.1. This workflow consisted of seven interconnected steps and tasks. Although the research tasks follow a sequential order, operationally revisiting or backtracking to previous steps was a necessary process for our model to reach an acceptable state.

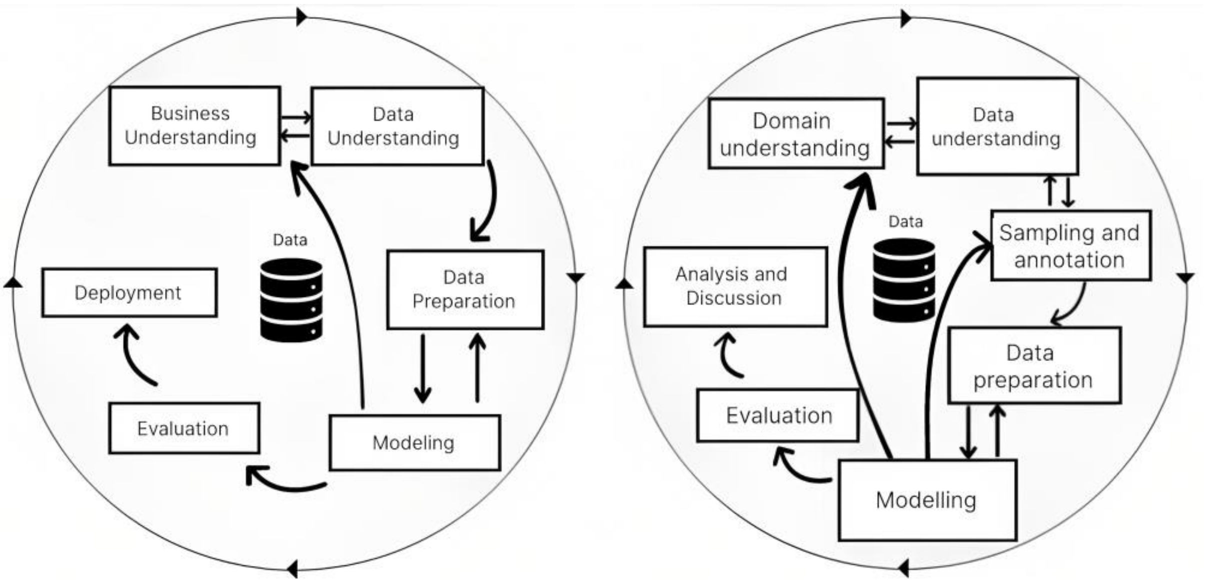

Figure 2.1. CRISP-DM (left side) and the proposed methodology (right side) (adapted from Hotz, 2023)

- Step 1: Domain understanding.

This step refers, according to the established objectives, to the stage where we decide and establish the population with which we will work. How we will be looking at the political and economic phenomenon of populism. Where we will delimit spatiotemporal criteria for the data with which we will work. This step sets itself apart from the "Business understanding" step in CRISP-DM due to its inherent nature, which is closely related to academia and research-based projects in social sciences, rather than being focused on solving specific business problems or achieving business objectives.

- Step 2: Data understanding.

This step is identical to the corresponding one in CRISP-DM, and involves the exploration and extraction of relevant attributes of the data, from the designated social media platform.

The process uses APIs for data collection to gather raw data, requiring well-defined criteria for content inclusion or exclusion and extraction methods to ensure data integrity. To this raw data, we then prepare it, through data cleaning, handling missing values, and data transformation, so it forms the basis for our subsequent analysis and research.

- Step 3: Sampling and Data annotation

In order to address our specific needs, this step was added to the CRISP-DM. Utilizing a random sampling to select a balanced and representative sample of tweets from all political actors involved in the study, ensures that all political actors are equally and fairly accounted for in the annotation step. Following this, data annotation is carried out, where binominal categories are attributed to the sample data so it can be used to train the model. These tasks work to ensure a robust foundation for our model.

- Step 4: Data preparation

In this step, the data undergoes essential preprocessing tasks using Natural Language Processing (NLP) techniques, converting the raw labeled text data into a format suitable for analysis. Feature extraction includes transforming the processed text into numerical representations and the potential addition of domain-specific features that will be used when modelling. Due to its complexity and importance in research, a big focus will be given to explaining how NLP techniques work.

- Step 5: Modeling

This step involves building and deploying the chosen deep learning model using the previously processed data and extracted features. This includes selecting an appropriate NLP-based model architecture, training the model on the training dataset, and evaluating its performance on the test dataset.

- Step 6: Evaluation

This is the crucial aspect of the analysis that was mentioned at the beginning of this chapter, as being paramount to adapt to the specific needs of social sciences research.

During this step, standard metrics based on the chosen model are generated to assess the model's performance accurately. If the results do not meet the expected criteria, it requires revisiting previous steps to experiment and refine the model or feature selection. On the other hand, if the results are satisfactory, the trained model can then be considered strong enough and deployed to make predictions on new, unseen data, which represents the remaining dataset that was not included in the initial labeling process.

- Step 7: Post-Deployment Analysis and Discussion

  In this step, the performance, effectiveness, and reliability of the deployed model are then evaluated based on its predictions of the unseen data. The analysis involves comparing the model's predictions with the grounded reality and manually processed data, assessing how well it generalizes them. The insights gained are documented, helping to understand the model's practical applicability and areas for potential improvement.

2.3. Resources

In this part, we will explain in more detail the resources used, from the platform used to collect our corpus and the dataset generated, to the choices related to the NLP. Additionally, we will go into further description of the NLP process itself.

When looking at text as data, one must start by collecting data. Given that, an annotated dataset telling apart populist moments from non-populist ones does not exist, we will have to collect one from scratch.

2.3.1. Twitter

The immense popularity and ongoing growth of the internet and its social components have made social networks a crucial part of our lives, as users rely on them to gain knowledge, explore their interests with like-minded individuals, and engage in social interactions (Amedie, 2015). These platforms serve as a virtual arena for debates and opinions, sharing habits and preferences, and showcasing customs, thereby providing a rich and diverse sample of the population in the form of saved, available and accessible data.

Within social media, Twitter emerges as a microblogging app with a high volume of communication channels used by an ever-present user base of more than 200 million daily active accounts (Curry, 2023). Twitter provides a platform for users to publicly publish and read short messages, called "tweets", in a structured fashion that allows easy navigation within its interface. The home timeline, the app's main feature, is where users will most likely see the majority of tweets. Displayed in a continuous vertical column feed, tweets are presented in reverse chronological order, where each tweet includes the username of the individual who posted it, their profile picture, the content of the tweet, any media or link included as well as the date and time of the post. Users can interact with these tweets in a plethora of different ways. They can either express their approval of the original message by liking, share the original form of the tweet with their own followers, known as retweeting, or include their own commentary while sharing the tweet via quote-retweeting, or directly respond to the original tweet feed through a reply.

There are certain characteristics that factor into Twitter being an attractive platform for text analysis. Specifically, Twitter is an incredibly rich source of data, with an estimated 500 million tweets shared daily around the world (Dixon, 2022), characterized by its publicly accessible and real-time information on a nearly universal and diverse range of topics. With all of the tweets being uploaded in real-time with their datetime saved, and with an estimated 65% of all Twitter accounts being publicly visible and searchable to anyone on the platform (Hutchinson, 2021), makes it easy for the collection of this information in a structured manor for research on various social phenomenon's.

Twitter stands out from other social media platforms for its particular restrictions on how it allows users to express themselves. The Tweet format is novel in that it constrains users to use only up to 280 characters per message posted, compelling users succinct and efficient with how clear their messages come across (Oliveira, 2022). This unique format, plus Twitter's own strengths make tweets a strong substitute unit of analysis for this kind of research when compared to traditional units of text analysis such as paragraphs or whole documents, frequently used in traditional populism analysis. Compared to traditional text sources, which are presumably longer, broader in scope and more complex and time-intensive to interpret, tweets are short, informal, and concise, while, at the same time, possibly being long and elaborate enough to extract content relating to more complex conceptual ideas such as people-centrism, anti-elitism or what can and cannot be considered populism (Ulinskaitė & Pukelis, 2021; Zhao et al. , 2011; Gunter et al., 2014).

In an effort to streamline the available data created on Twitter's platform, Twitter, in 2006, made available a set of programming tools, named the Twitter API (Application Programming Interface) v1.1, which allows for public access and interaction of Twitter data. The API provides a way, through several different programming languages, to retrieve tweets, user information, trends, and other useful information, as well as the ability to automatically create tweets, send direct messages and other actions on behalf of Twitter users outside of the app (Oliveira, 2022).

Critical for its effective usage, the Twitter API v1.1 is structured into multiple endpoints, each serving a distinct purpose. API endpoints are unique identifiers, that through HTTP requests, access a specific set of data or functions within the API, they are designed to provide a consistent and reliable "path" to an underlying more complex system (Cooksey, 2014). These endpoints are equipped with flexible and customizable search parameters that make them versatile resources for a wide range of research applications, each leveraged to support hyper-specific research goals (Twitter, 2012). For the purpose of this particular research, we will expand on three relevant endpoints, namely the Search API Endpoint, Users API Endpoint, and Timeline API Endpoint.

The Search API Endpoint is one of the most commonly used tools within the Twitter API. It allows users to specifically search for tweets that match a particular set of criteria. This can be a specific keyword, hashtags, geographic locations, date range, language, etc. Its application is particularly useful to narrow down relevance-based search results (Twitter, 2012). This kind of endpoint would be most used when researching a particular topic, tracking its popularity, or monitoring public opinion, and gathering insights.

The Users API Endpoint allows access to user-specific information. This endpoint is designed to retrieve insights into user behavior through username or user ID. It provides access to information on profiles, such as the user's bio, timeline, profile picture, location, follower and following counts, and a variety of other metadata. This endpoint is also useful for searching particular accounts that fit certain criteria, curated to what is relevant to the research.

Lastly, the Timeline API Endpoint allows for the retrieval of a collection of tweets that match a specific query. This endpoint is used to retrieve a user's home timeline, which includes both his own tweets and the tweets of the accounts they follow, the user's personal timeline, which includes just his own tweets, and it can also be used to retrieve the user's mentions timeline, which includes tweets that mention the user's Twitter username. This can then be further fine-tuned with other optional parameters to retrieve information useful for the research objectives.

With this in mind, the following metrics were extracted:

- Tweet ID: Every individual tweet has a unique numerical number that identifies it from other tweets. This is particularly useful for tracking and referencing specific tweets during the analysis.
- Tweet Text: Contains the text content of the tweet, which is the primary focus of analysis in many NLP tasks.
- Date: Contains the date when the tweet was posted, in a Y-M-D format, which provides temporal context for the analysis.
- Time: Contains the time of day when the tweet was posted, scraped in a 24-hour clock format (HH:MM:SS) in order to eliminate any mistakes possible with the use of an AM/PM system. It provides additional temporal context for the analysis.
- Likes: Presents the number of likes that the tweet has received, which can be used as a measure of engagement or popularity.
- Retweets: Presents the number of times that the tweet has been retweeted, which can also be used as a measure of engagement or popularity.
- Username: Serves to identify the Twitter account that posted the tweet.

2.3.2. Population

The pool from which data is retrieved consists of the parliamentary leaders of the relevant parties, with only content published while political actors were active in leadership roles was included.

Relevant parties are considered those that won parliamentary seats in the last two electoral cycles.

Data were collected from the party leader in the respective cycle, that is, in the case of change of party leader between cycles, both leaders were included in the analysis, with the exception that, in these situations, the tweets were collected only related to the time they were at the head of the party.

Additionally, in situations where the designated party leader does not have a Twitter account, or where the party is led by an executive body, the president of the parliamentary group was alternatively considered for selection.

In the hypothetical scenario in which this deputy also does not have a Twitter account, the analysis would fall on the vice-president of the parliamentary group. This measure ensures that all parties are represented equally during the two electoral cycles under analysis, regardless of changes that have occurred, leadership structure and the presence of the individuals on social media.

2.3.3. Natural Language Processing

With the advancement of information storage and the establishment of increasingly larger databases, how and what data we extract from the available information is a core factor of analysis in all fields of study. Artificial intelligence (AI) is the field of science and engineering of making intelligent machines capable of executing tasks that typically require human intelligence (McCarthy, 2007; Sharma et al., 2022). Within this field, machine learning is the subject that studies how to use computers to simulate human learning activities. A system starts with little to no structure, and, through incremental changes in the parameters[12], it improves its performance over time, by identifying patterns and relationships, and then, using those same patterns, to make predictions or take actions (Wang et al., 2009; Carbonell et al., 1983). Before diving further into the complexities of the subject and our respective research, let us start by providing a basic explanation of its operational framework.

A basic machine learning model operates through what is called an information cycle (Wang et al., 2009). It starts by ingesting external data, which serves as the source of information, processing that data, and converting it into valuable "knowledge," storing it in a repository (Wang et al., 2009). Next, the model utilizes this knowledge to perform specific tasks, and as a result, it receives feedbacks based on its performance. That feedback is then added to the repository to be used to refine and enhance its knowledge, creating an iterative information cycle (Wang et al., 2009).

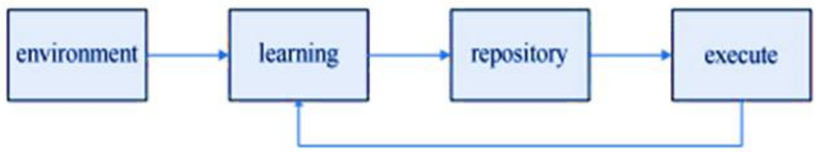

Figure 2.2. A basic machine learning model, creating an information cycle (Wang et al., 2009).

These machine learning algorithms can be classified in four basic types according to the amount and type of supervision they get while creating their "knowledge": supervised learning, unsupervised learning, semi-supervised learning, and reinforcement Learning.

In supervised learning, the algorithm is trained on labelled data, where each example in the dataset has a known target variable. The algorithm learns to make predictions by generalizing from the labelled examples it has seen during training (Géron, 2017; Dahiya et al., 2022).

With unsupervised learning the algorithm is trained on unlabelled data. Where, without any known target variable, the algorithm tries to find patterns and structure in the data by itself (Géron, 2017; Dahiya et al., 2022).

---

[12] **Parameters**: values that a machine learning algorithm uses to make predictions based on input data (like the settings on a camera that affect the quality of the picture being taken).

In the middlel, semi-supervised learning is when the algorithm is partially trained on a minority of labelled data and a majority of unlabelled data, with each serving their own purposes. The labelled data is used as a guideline for the learning process, while the unlabelled data is used to improve the generalization performance of the algorithm (Géron, 2017; Dahiya et al., 2022).

Reinforcement learning functions much like Pavlov's experiment[13]. Where the algorithm learns to make decisions based on feedbackit was given. The algorithm interacts with an environment by taking actions and then receives rewards or penalties based on its actions, with an active goal of maximizing rewards while minimizing penalties (Géron, 2017).

After introducing the basic pipeline, it is necessary to explore the strengths and weaknesses of machine learning. Machine learning has the major advantage of being able to process rich information present in large datasets without incurring significant time and cost expenses for researchers. Moreover, it can automate repetitive tasks, thereby improving time-efficiency (Zhou et al., 2017). However, one has to be careful with "bad algorithms and bad data" (Géron, 2017) that generate incorrect and unreliable predictions.

Algorithms may be deemed 'bad' when they are poorly designed, implemented, or optimized, leading to suboptimal performance. One common issue in machine learning is overfitting, where a model becomes overly complex and starts memorizing noise[14] and specific details in the training data rather than capturing general patterns (Géron, 2017), leading it to perform exceptionally well on the training data but poorly on new, unseen data. In contrast, if the model is underfitting the data, it means the model is too simple and fails to capture the underlying patterns present in the data (Géron, 2017).

It takes a lot of data for most algorithms to work properly, typically needing at least thousands of examples in order to not produce poor and biased results even if well-designed (Géron, 2017). Data with poor quality is data that is insufficient, incomplete, biased, or not representative of the overall population (Géron, 2017). A biased dataset can result from factors such as sampling bias, measurement errors, or bad labelling foundations.

On paper, it may seem naturally intuitive that using more data points would automatically result in a more accurate dataset, and, as such, the best way to determine the parameters of a model would be to use all the available data as the source of information (Gholamy et al., 2018). However, in practice, this can lead to some unperceived problems in representation. To ensure that the model is not context specific, and is able to generalize to new, unseen data, the appropriate techniques should

[13] **Pavlov's Experiment**: was a set of tests where he repeatedly rang a bell shortly before presenting food to a group of dogs. This was done in order to condition them to associate a stimulus with the arrival of food, causing the dogs to salivate at the sound of the bell alone, even when food was not present.

[14] **Noise**: in machine learning, refers to random errors and inconsistencies in the data that can negatively impact the accuracy and reliability of the model

be applied. One of these techniques involves dividing the dataset into two or three subsets: a training set, a test set, and sometimes a validation set (Moura et al., 2021).

The model is firstly trained on the training set, and its performance is evaluated on the test set (Gholamy et al., 2018; Géron, 2017). In cases where, due to over-tuning of the hyperparameters[15] to a training set, the test set seems to perform worse than expected, one can, in what is called a holdout validation, hold part of the training set as a validation set. With this, we can fine-tune the model's hyperparameters and compare the performance of different models, to evaluate the best one. After this holdout validation process, one trains the model on the full training set (including the validation set), and this gives it the final model, which is then used on the test set (Géron, 2017), to ensure that the model can generalize well to new data. Now that we have a general understanding of machine learning is, let us narrow our focus to the subset that is relevant to this research.

Natural Language Processing (NLP) is a branch of artificial intelligence, where through a set of machine learning methods it bridges human language, either written or spoken, as makes it accessible to computers (Eisenstein, 2019; Azevedo, 2020). Its diverse application is based on a common set of linguistic methods “rules, algorithms, linguistics, logic statistics and more” (Eisenstein, 2019).

In NLP, raw data comes in the form of text, which is processed to make it suitable for machine learning models. Text data is essentially discrete, made up of a millions of combined characters and words with clear boundaries, where if anything is divided or shortened it may lose a significant amount of its meaning (Eisenstein, 2019).

Additionally, language is deeply compositional, as words and letterns combine to create arrangements, like phrases and paragraphs, with an underlying hierarchical structure that is not implicitly without context. While text appears in a sequence, its meaning changes depending on its order and composition, and not just the sum of its words (e.g. while “bear cub” and “bear hug” share the word bear, they have vastly different meanings, while bear cub refers to a young bear, "bear hug" refers to an affectionate and strong embrace).

Although humans possess the innate ability to comprehend the underlying meaning of sentences from a very young age, language is actually highly complex and nuanced, with concepts like sarcasm, metaphors and irony, and machines, who operate on a set of rules and algorithms, without proper mechanisms to interpret natural language, cannot accurately understand the meaning behind sentences (Moura et al., 2021; Azevedo, 2020). In these efforts, over time, specialized techniques have been developed to process the texts presented efficiently.

[15] **Hyperparameters**: Are configurations that are manually specified by the machine learning researcher before training a model. These settings are not learned from data, unlike the model parameters.

A standard machine-learning workflow using a NLP pipeline[16] consists of three major phases, a pre-processing of the raw text data, represented in Figure 2.3, feature extraction, and then modeling.

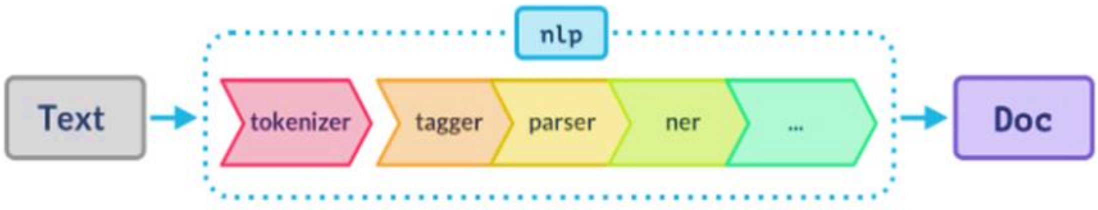

Figure 2.3. Standard machine-learning workflow and its NLP steps

2.3.3.1. Data Cleaning and Pre-processing Steps

Since, in NLP, were dealing with a lot of unstructured text with varying amounts of noise, one has to commit a big effort to processing and cleaning this data. An NLP pre-processing step breaks things into small modular pieces in order to clean, create new attributes, and transform unstructured data based on the meanings and intent of its information features in a streamlined way (Moura et al., 2021; Vasiliev, 2020). While there is no universally correct way to apply it, a basic NLP pre-processing part of a pipeline typically includes a tokenizer, either a stemmer or a lemmatizer, a tagger, and a parser.

2.3.3.1.1. Tokenization and Tokenizers

The first step in an NLP pipeline, tokenization is the process of breaking down a piece of raw text into its smallest units, called tokens, which can either be words, multi-word expressions, or punctuation, depending on the method used (Moura et al., 2021; Vasiliev, 2020; Azevedo, 2020). Its main purpose is to transform a continuous text into a sequence of discrete tokens that can actually be processed by the model, so we can extract features and meanings in later steps. Within the tokenization process, there is usually a complementary subset of operations to do: lowercasing, removal of links, and stop-word removal.

The first two techniques are quite straightforward. Lowercasing involves converting all text to lowercase to avoid treating the same word as different tokens due to capitalization. The second technique is the removal of links, which involves identifying and removing any URLs or hyperlinks present in the text data when they are not relevant to the analysis. Stop-word removal is the process of removing words that don not carry much meaning to the phrase but tend to appear recurringly in

[16] **Pipelines**: are a sequence of data processing components. Each component pulls in a large amount of data, processes it, and spits out the result in another data store, the next component in the pipeline then pulls this data, processes it, and spits out its own output, and so on, forming a chain of data transformations.

most phrases. Usually, these are the most common words in a language, mostly function words[17], as compared to lexical items (Moura et al., 2021). This technique is particularly important as it removes a lot of the text's noise, leaving only meaningful tokens for future steps in the analysis.

2.3.3.1.2. Stemmer and Lemmatization

Within linguistics, words may belong to the same word family while having a different morphology, in other cases, words share the same written form while having completely different meanings, where one can only distinguish them by the surrounding context they are placed. Altering the words in a text can make a crucial step in an NLP pipeline, whereby reducing words to their root form one can substantially reduce and simplify the vocabulary present in their data, narrowing the number of tokens analysed, reducing noise, and improving accuracy (Moura et al, 2021; Vasiliev, 2020; Azevedo, 2020). Stemming and lemmatization are two techniques that reduce words to their common root form. They do, however, operate in different forms and offer different outputs.

Stemming is the process of reducing words to their stem or form by removing affixes (Moura et al., 2021; Vasiliev, 2020). While in general, removing suffixes keeps the essential information from the word, removing prefixes can lead to the loss of important information, as it can represent the denial function in most languages. Stemming is a relatively simple and fast process based on a set of rules that does not need a prior language vocabulary to work. In contrast, this simplicity has some limitations, since its ruleset may not result in most accurate root forms for words, and in some cases, can produce stems that are not even actual words (Moura et al., 2021; Azevedo, 2020).

Lemmatization as a process involves an actual knowledge of the language being analysed, as well as an understanding of a word's context and morphological analysis in order to accurately transform it into its root form. Since this technique considers the word's context it typically needs the use of an established dictionary[18] or lexicon to map words to their base form (Vasiliev, 2020).

A practical example might be able to simplify this concept. Take the word "better" as the comparative adjective in any one sentence. Through stemming it would simply remove the suffix "-er" and return "bett" as the root form, whereas lemmatization would consider the word's morphology and return the root word as "good".

2.3.3.1.3. Tagger and POS Tagging

[17] **Function words**: Words that express grammatical relationships among other words within a sentence or specify the attitude or mood of the speaker.

[18] **Dictionary**: Also known as lexicon, is a collection of words or phrases with their corresponding meanings and other relevant information. They are valuable resources for NLP tasks, providing an already approved set of standardized, comprehensive, and contextual rich information about the words in a particular language.

Part-of-speech tagging (PoS tagger for short) is the tool in the pipeline that assigns tags, or lexical labels, to each token in a given text, based on its definition and grammatical function (Moura et al., 2021; Azevedo, 2020). We will be using the Penn Treebank tagging scheme for systactinc and part of speech information (Taylor, 2003), which goes as follow: "Today I go to school early" A PoS tagger would assign the tag "RB" (adverb) to "early", "PRP" (personal pronoun) to "I", "VBP" (verb, present) to "go", "ADP" (adposition) would be assigned to "to", and "NN" (Noun) to "school" and "today". PoS is the crucial step that provides the foundations for any, more complex, linguistic analysis beyond just simple keyword matching.

2.3.3.1.4. Parsing

Up until this point, the NLP pipeline has been processing data without worrying about the order in which words appear in a sentence. Parsing is then the process where one can analyse a sentence's grammatical structure and relationships between its components, aiming to create a structured representation of the sentence that captures all its syntactical meaning (Azevedo, 2020). Dependency parsing is used to represent sentences as a structure of interconnected words, where each word is a node, and the edges represent the direct grammatical relationships between them (Nivre, 2005), where not all words in a sentence have a direct connection. This method follows the framework of dependency grammar, a traditional linguistic analysis theory, that proposes that syntactic structures consist of "binary asymmetrical relations, called dependency relations, between two words" (De Marneffe & Nivre, 2019; Nivre & Nilsson, 2005), where the words in a sentence may or may not happen to have a direct connection, allowing for a more nuanced representation of their grammatical relationships. This is then represented visually in the form of dependency trees (Figure 2.4.).

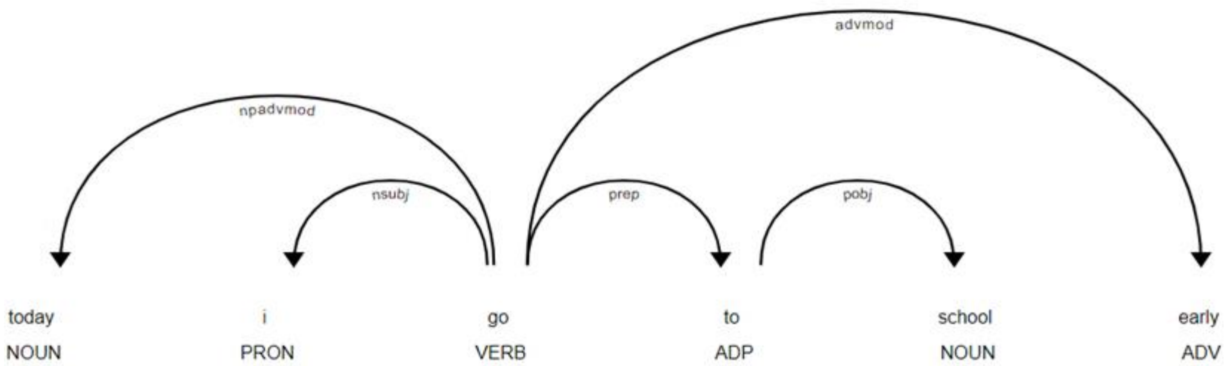

Figure 2.4. Visual representation of dependency trees

2.3.3.2. Feature Extraction

After the data is normalized and cleaned in the pre-processing phase, we can proceed to the next step in the pipeline. In this step, we are once again faced we the fact that machine learning models do not understand text, and as such we cannot use it as input before the process of extracting relevant and meaningful information from the cleaned data, by first transforming it into numerical representations (Moura et al., 2021). The feature extraction and transformation step can be done using both simple frequency-based techniques and more complex, neural network techniques, in which we will be focusing on N-grams and word embeddings.

2.3.3.2.1. N-Grams

N-grams are a frequency-based technique for extracting features from a text by breaking them down into sequences of N tokens. N-grams are essentially the count of the occurrences of each sequence of tokens within a text, with N being the number of tokens analysed, for example, a bigram (2-gram) would consist of two consecutive tokens. These items can be syllables, characters, or the most commonly used in NLP, words. A particularity of N-grams is that while being a relatively simple technique, that does not consider the relationship words have in a sentence it still preserves the order or structure of words in it to a limited extent (Moura et al., 2021).

Even though simpler techniques such as N-grams are effective most of the time, in certain tasks, their inability to consider the relationship of words inside the text can be a striking disadvantage, and end up producing disappointing results (Moura et al., 2021). For a better understanding look at the following sentences:

Sentence A: "The president endorsed the bill."

Sentence B: "The country's leader supported the legislation."

Although these two sentences essentially convey the same meaning, a frequency/count approach would fail to capture their similarity as they do not share any common words. This approach would emphasize the differences between the sentences instead of capturing their similar semantic meaning.

2.3.3.2.2. Word Embeddings

Word embeddings, like word2vec and GloVe, are context-free techniques in which feature vectors are used to represent text with the aim is to capture the meaning and relationships between words in a numerical form (Moura et al., 2021). These embedding vectors represent words as low-dimensional vectors, where the combined vectors of each word them comprise the sentence's representation. These vectors are stored in matrices, where each row corresponds to the embedding vector of a specific word in a vector space, where words with similar meanings or contexts are located close to each other. This representation of words enables mathematical operations to be performed on text, such as measuring the similarity between words or finding analogies between them (Moura et al., 2021). For instance, the cosine similarity metric can be used to measure the similarity between words. As such, when two words are similar, they will be placed close to each other in the embedding space, as such, their vectors will also be similar, in spite of how different they may be written. Since these models only output one vector for each word, combining all the different senses of the word into one vector, meaning, they would simply return the same embedding for any word written the same way, regardless of if they have different meanings (Vajpayee, 2020).

Prediction-based approaches in word embeddings are designed to analyse “what are” the intricate patterns and nuances between words’ multiple degrees of similarity. By being trained on sufficiently large datasets, these methods learn vector-space representations of words through the weights assigned to their input layers. The process involves manually adjusting the word associations to predict which words align most closely with the target word. In this way, the model captures the underlying patterns and semantic relationships in language to infer meaningful connections between words (Mikolov et al., 2013; Moura et al., 2021).

A good illustration of how these methods work is presented by Mikolov´s and his co-authors (2013) research, where they found that these dimension vector offset methods were great at understanding associations between words (Figure 2.5). Through tinkering with vectors in the embedding space, the model automatically inferred the gender relationship between the now famous "king and queen" example. Whereby identifying linguistic regularities, through the simple arithmetic operation, "king - man + woman," the resulting vector was found to be very similar to the vector representation of "queen” (Mikolov et al., 2013).

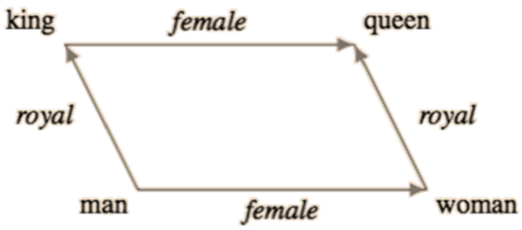

Figure 2.5. King and Queen vector example presented by Mikolov et al. (2013)

2.3.3.2.3. Artificial Neural Networks

It is important to note that the context-free word encoding techniques were already operating within the framework of a neural network, even if, in their case, they are simpler, “shallow” networks.

Artificial neural networks (ANN) are mathematical models, where the nodes are interconnected, as the human brain’s neuron networks are connected. In a network, the neuron accepts the input, processes it, and produces an output (Sharma et al., 2022). In a shallow neural network, as presented in Figure 2.6., the feature values of the input data, which are the attributes that describe an object (like all the individual divisions in a house describe what makes an actual house), are passed through a single intermediate hidden layer of nodes, also known as neurons/units. Each neuron applies a specific “calculation” to produce a response based on the input it received. These calculations are called activation functions, which decide whether the neuron's input to the network is important or not for the network using mathematical operations (Sharma et al., 2022; Yu et al., 2019). Without these activation functions, NN would be restricted to learning linear relationships between inputs and outputs, and since, most of real-life problems involve intricate relationships and non-linear dependencies among the input variables, without these functions, NN would have limited to no predictive power (Sharma et al., 2022; Yu et al., 2019). These responses then collectively form the output of the network, which can depend on the type of task the model is being applied for (Kröse & van der Smagt, 1993; Jain et al., 1996).

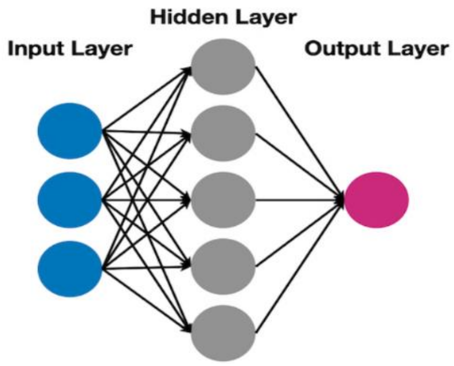

Figure 2.6. Representation of a shallow neural network

2.3.3.3. Deep Learning Models

Let us consider this in a more practical example using the word embeddings and the populist example. In this scenario, the input layer of the neural network receives labeled data, where each sentence in the text is represented as a word embedding vector. Throughout the process of knowledge assimilation, the network encounters a curated ensemble of training examples. Each of these instances is composed of a distinct input (namely, a tweet) coupled with an aspired output (which corresponds to a specific label). Engaging in a constant state of refinement, the network meticulously tunes its internal weights, rigorously minimizing the discord between its generated output and the sought-after output for each training example (Jain et al., 1996). Finally, the output layer of the network generates a prediction score that indicates whether the text is classified as either populist or non-populist. This process is repeated over many iterations until the network's performance on a validation set reaches a satisfactory level (Jain et al., 1996).

In that sense, a deep neural network (DNN) represents a natural evolution from a shallow network. By incorporating multiple hidden layers between the input and output layers, a DNN becomes capable of accurately performing a more complex hierarchy of computations. This increased depth enables the network to learn intricate representations of the data, capture higher-level features, and extract abstract patterns (Aouichaoui et al., 2021; Sharma et al., 2022). However, this increased complexity comes at the cost of higher computational needs and requires exponentially larger training datasets to mitigate overfitting and present satisfying results.

High-level, highly complex, and abstract features, such as a speaker's accent, require a sophisticated, nearly human-level understanding of the raw data. Deep learning, through the use of deep neural networks as its building blocks, enables the models to build complex concepts out of simpler nodes. Modern deep learning provides a powerful framework for supervised learning and has been applied in common text-mining tasks, such as sentiment analysis (Severyn & Moschitti, 2015), text classification (Liu et al., 2015), and text generation (Graves, 2013).

2.3.3.3.1. Sentence Transformers

Introduced by Vaswani et al. (2017), and his team of researchers at Google, Transformer models were a ground-breaking type of neural network architecture for sequence modelling, one which has since overtaken other neural network architectures/structures and became the go-to approach for most NLP tasks (Tunstall et al., 2022).

Sequence modelling is the task of predicting or generating an output sequence based on an input sequence. In the context for text, it specifically refers to predicting the next word or letter in a sentence. In order to capture the dependencies between the sequential elements effectively, one needs to apply a neural network architecture, in order to encode them into its internal representations, facilitating accurate sequence generation.

Introduced to address the challenges faced by the traditional neural network architectures in capturing dependencies between distant positions, and to overcome the slow training process, the Transformer model's key innovation lies within its self-attention mechanism (Vaswani et al., 2017), which allows the model to process all the words in a sentence simultaneously. The self-attention mechanism attends to different parts of the input sequence by assigning importance weights to each token in the sequence. These attention weights would then determine how much importance a token had in relation to other tokens in that sentence, capturing contextual relationships between all the words (Tunstall et al., 2022).

2.3.3.3.2. Transformer Model Architecture

The Transformer model architecture consists of both an encoder and a decoder working together to process sequential data (Vaswani et al., 2017). While the encoder is responsible for learning context and meaning in language, the decoder learns how to apply what was learned by the encoder to generate accurate and coherent output (Tunstall et al., 2022). Both are composed of N number of identically structured layers, with each having corresponding sub-layers, two of them identical, being the self-attention and the feedforward neural networks layer, as shown in the left and right halves of Figure 2.7.

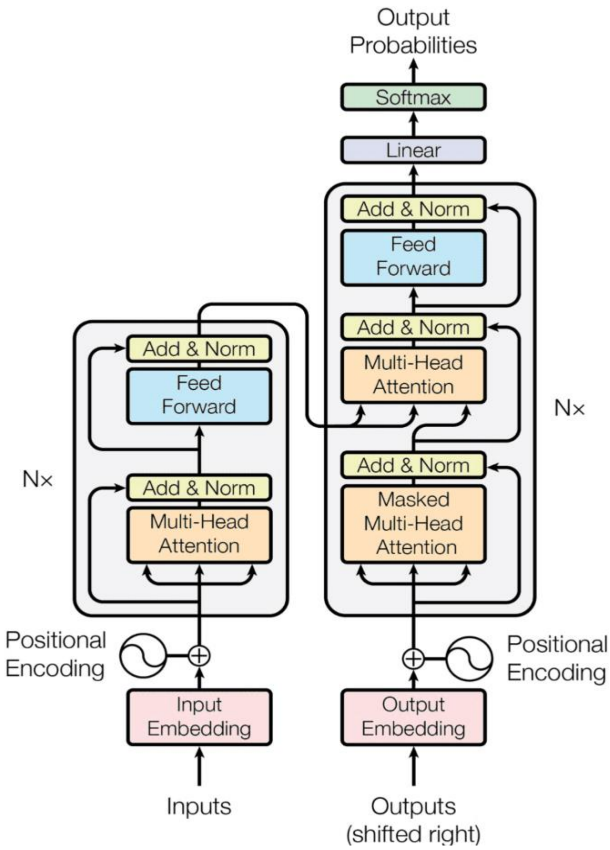

Figure 2.7. Encoder-decoder structure of the Transformer architecture

Taken from "Attention Is All You Need" (Vaswani et al., 2017)

The encoder takes an input sequence and transforms it into embeddings that capture the semantic information of each token. These embeddings are then passed through positional encodings, which provides the sequence with positional information and order, also known as the context (Braşoveanu & Andonie, 2020; Vaswani et al., 2017). These two combine into the input representation and are passed to the encoder block where it goes through a multi-headed attention layer and a feed forward layer. The attention layer passes the encoded input sequence through a stack of identical self-attention encoder layers, where it determines the tokens that require focus by utilizing an attention vector, which captures contextual relationships between words in the sequence (Vaswani et al., 2017). Following the self-attention mechanism, a feed-forward neural network is applied to each attention vector independently. This feed network introduces non-linearity and further refines them to a format that can be easily processed by the subsequent encoder or decoder blocks.

During the training phase of a model, the decoder block receives a similar input representation as the encoder one (see how their starts are the same in Figure 2.7). Additionally, the decoder also takes into account the output from the encoder as an input, where both are then passed into the decoder's block.

In addition to the two already described in the encoder, the decoder adds a third sub-layer, the Masked Self-Attention layer (Vaswani et al., 2017). The "masked" aspect refers to the fact that during training, the decoder's self-attention future positioned tokens are masked by setting their attention weights to a very small value. Masking these from the attention calculation, prevents the model from looking at future tokens during the training process, and ensures it relies solely on the information it has available up to the current position (Kierszbaum, 2020). By masking certain tokens, the model is forced to learn to predict missing or masked tokens based on the surrounding context and patterns of the data. It is like watching a movie while actively avoiding spoilers, and then making predictions on what will happen.

These masked attention vectors of the input representation, along with the encoder's representations are passed to another attention block, designated as the "Encoder-Decoder" attention layers (Vaswani et al., 2017). The output of this block attributes the attention vectors of every token from the encoder's and decoder's sentences. These vectors allow for calculating the attention weights by measuring the similarity between the decoder positions and each encoder position in order to make informed decisions during the decoding process. By attending to the relevant parts of the input sequence, the decoder can align its generation with the input and produce accurate and contextually relevant outputs (Vaswani et al., 2017). These are then passed through a feed forward unit in order to make the outputs more digestible to the next decoder block or the linear layer.

After the decoder's block steps, the output passes through additional layers in the Transformer architecture before making its predictions, specifically, the linear and SoftMax layers. The linear layers work similarly to the feed forward layers and is used to increase the dimensions of the representation space so that all tokens are mapped into the same homologous feature space (Vaswani et al., 2017). The SoftMax layer then transforms the models output into a probability distribution so that it can finally be human interpreted. It takes the raw output values from the previous layers and applies a mathematical operation that assigns a probability score to each possible output class or category, where these probability scores represent the model's confidence or belief in the presence of each class (Vaswani et al., 2017).

Let us take the case of our practical populism example. During the training phase, our transformer model learned to adjust its parameters based on the labeled dataset, where each sentence was annotated as either populist or not populist. Once the model is trained, it can take a new input sentence and pass it through its encoder-decoder network to produces probabilities scores for each category. The linear layer then adjusts these scores into a better representation of the models overall prediction. Lastly, the SoftMax layer then transforms these scores into what the model predicts as the probabilities of the sentence being populist or non-populist, ensuring that the combined probabilities add up to 1.

2.3.3.3.3. BERT

Derived from the transformers model, BERT (Bidirectional Encoder Representations from Transformers) is what we get if we stack multiple encoder blocks (Vajpayee, 2020). The bi-directional part in BERT comes from the fact that it reads all the input words simultaneously, apart from that, the sub-layers in the encoder's block, although handling different tasks and contexts, they retain functional resemblances in how they function (Vajpayee, 2020). This similarities enable BERT to to generalize its learnings from one part of language to another, making it a versatile and powerful language model. Introduced by a team at Google (Devlin et al., 2018), BERT's framework is comprised of two major steps: pre-training, and fine-tuning.

2.3.3.3.3.1. Pre-Training

The goal of the pre-training phase is to make BERT learn to understand all kinds of structures and meanings of language. Unlike traditional left-to-right learning models, BERT does this, through two simultaneous unsupervised tasks being trained on a massive amount of unlabeled text data. In Devlin et al. (2018) this data consisted of the BooksCorpus dataset (800M words) and English Wikipedia (2500M words), extracting only the text passages and ignoring lists, tables, and headers (Devlin et al., 2018; Moura et al., 2021), but BERT possesses the capacity to be trained with virtually an infinite variety of data sources.

For the first task, masked language modelling (MLM), the model takes in a sentence and randomly masks some tokens. This is done with the goal to output these masked tokens (similarly to fill-in-the-blanks exercises), helping BERT understand the bidirectional context within a sentence (Devlin et al., 2018). In the other task, next sentence prediction (NSP), the model takes in two sentences, and it determines whether the second sentence follows the first one. This helps BERT understand the context across different sentences (Devlin et al., 2018). By using both tasks together, BERT learns to form a deep understanding of language structure and semantics.

2.3.3.3.3.2. Fine-Tuning

Training a language model from scratch to achieve a high level of understanding of language is an undertaking that requires substantial computational resources, vast amounts of training data, and extensive time costs, which are highly challenging and resource intensive. Even on the occasion that one manages to successfully train such a model, the likelihood that a larger, better-trained model already exists is extremely high, undermining the potential value and practicality of such a task. Luckily fine-tuning provides a valuable alternative approach (González-Carvajal & Garrido-Merchán, 2020). A good example would be how a smartphone already comes pre-loaded with its core functionalities when we buy it. We then customize it by adding specific apps tailored to our different needs instead of having to buy new smartphones that come with those functionalities every time we need it to do something new.

Fine-tuning an already existing pre-trained model on a specific downstream task[19] is relatively straightforward. One just needs to modify or add a new set of output layers tailored to the specific task we want to do to the pre-existing model (González-Carvajal & Garrido-Merchán, 2020). This is a much simpler and less time-intensive process, as only the newly introduced output layers have to be trained from scratch, with the rest of the model's parameters being just slightly fine-tuned to fill the specificities and objectives of the task at hand (Devlin et al., 2018).

2.3.3.3.3.3. Classification Algorithms

Classification is the step of the NLP process where through the extracted features from the text data, the model learns to recognize, understand, and classify objects into preset categories (Wolff, 2020). While there are several types of classification algorithms one can use, we have chosen to focus only on logistic regression for its simplicity, ease of understanding, and practical applicability (Plummer, 2020).

Logistic regression is a widely used and well-established algorithm for binary classification tasks (ex. Populism or non-populism). The term logistics comes from the designation "logit" probability function of this model (Moura et al., 2021; Géron, 2017), where it models the relationship between the independent variables and the probability of the binary outcome, according to equation 2.1:

$$\text{odds} = P(Y=1|X) \,/\, P(Y=0|X) \tag{2.1}$$

The numbers in the equation 2.1 represent the probabilities of the dependent variable Y taking on the value that represent each binary class. The $P(Y=1|X)$ is the probability of the dependent variable Y being equal to 1 given the independent variable X, with the $P(Y=0|X)$ being the same but for value 0. As such, the value of X is able to be anything while P(X) being in the range of 0 and 1.

To achieve this, the logistic regression model applies the sigmoid function to the linear combination of the independent variables and their coefficients, mapping the linear combination to a

[19] **Downstream task**: Comes from the concept of a data flow pipeline. These tasks are considered "downstream" because they build upon and are dependent on the pre-trained base.

value between 0 and 1, ensuring that the predicted probability is always within the valid range (Géron, 2017).

$$P(X) = \frac{1}{1 + e^{-\beta X}} \tag{2.2}$$

Using the sigmoid function gives us the following Figure 2.8., represented in an S-shaped curve. Fundamentally, the sigmoid function takes a number as input and estimates the parameters so that for every sample is as close to the value of either 0 or 1 as possible, meaning, every input is "definitely part of group A" or "definitely part of group B".

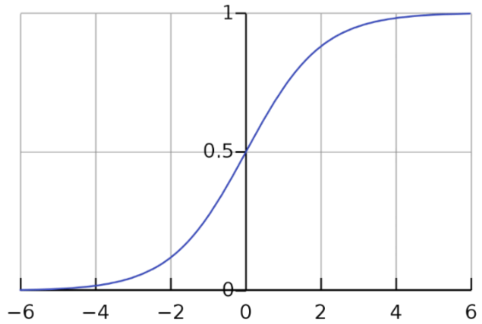

Figure 2.8. Representation of a basic sigmoid function

Since in NLP, the independent variable (embeddings) is represented as a multiple-dimension vector, an extra step is added to ensure that the logistics regression model works correctly. In this step, we reshape the vectors into a 2D arrays, organizing the values that represent each dimension into columns. This allows logistic regression to perform calculations more effectively on each feature separately when estimating the probabilities of the binary outcomes.

2.3.3.4. Performance Assessment

Having gained an understanding of how the models operate, it is now crucial to delve into the methods that measure how these trained models will perform when applied to a testing dataset. To achieve this, we employ various evaluation parameters that enable us to effectively measure and analyze the performance of the models.

2.3.3.4.1. Confusion Matrix

The confusion matrix is a valuable tool for evaluating the performance of a classification model, providing insights into the accuracy of predictions for each class. Figure 2.9. summarizes the number of correct and incorrect predictions made by the model for each class, offering a comprehensive overview of its performance (Sharma et al., 2022; Moura et al., 2021). The matrix consists of four key components:

- True Positive (TP): Represents the number of cases correctly predicted as positive. In the context of our problem, it refers to correctly identifying sentences that were labelled as populist correctly.
- True Negative (TN): Indicates the number of cases correctly predicted as negative. It signifies accurately identifying sentences that were labelled as non-populist correctly.
- False Positive (FP): Refers to the cases incorrectly predicted as positive. It represents instances where the model predicts that a sentence is populist when it is actually not. These are known as Type I errors.
- False Negative (FN): Represents the cases incorrectly predicted as negative. It represents instances where the model predicts that a sentence is non-populist when it actually belongs to the populist category. These are known as Type II errors.

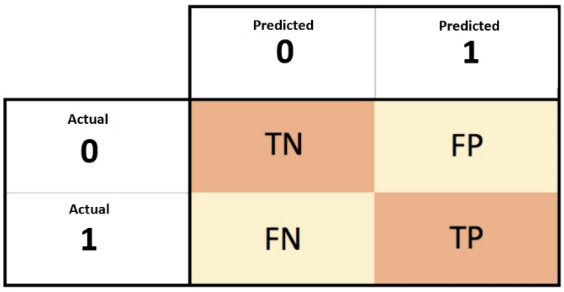

Figure 2.9. Basic representation of a confusion matrix (Sharma et al., 2022)

2.3.3.4.2. Accuracy

Accuracy is a fundamental metric for evaluating the performance of a classifier. It measures the correct classification rate among the entire population being analysed (Sharma et al., 2022).
The accuracy can be calculated using the equation 2.3:

$$\text{Accuracy} = (\text{True Positive (TP)} + \text{True Negative (TN)}) \,/\, \text{Total Population} \tag{2.3}$$

In the context of our research, accuracy represents the ability of the classifier to correctly classify sentences as either populist or non-populist.

2.3.3.4.3. Precision and Recall

Precision and recall are complementary parameters for evaluating the performance of a classifier. Precision indicates the proportion of correctly identified populist sentences among all the sentences predicted as populist. It is calculated using the equation 2.4:

$$\text{Precision} = \text{True Positive (TP)} / (\text{True Positive (TP)} + \text{False Positive (FP)}) \qquad (2.4)$$

On the other hand, Recall, also known as sensitivity or true positive rate, represents the proportion of the existing populist sentences that are correctly identified as such (Sharma et al., 2022), and is calculated using the equation 2.5:

$$\text{Recall} = \text{True Positive (TP)} / (\text{True Positive (TP)} + \text{False Negative (FN)}) \qquad (2.5)$$

A higher precision value indicates a lower rate of false positives, which means that the classifier is more accurate in correctly identifying instances of a specific class, in this case, populist sentences. On the other hand, a higher recall value indicates a lower rate of false negatives, which means that the classifier is better at capturing instances of a specific class. In this case, it implies that the classifier is better at correctly identifying populist sentences among all the available instances (Sharma et al., 2022).

2.3.3.4.4. F1 Score

The F1 score is the metric that provides a balanced measure of performance by considering both precision and recall for a classifier, being especially useful in scenarios where there is an uneven distribution of classes, as compared to the accuracy metric (Sharma et al., 2022). The formula for calculating the F1 score is as follows (equation 2.6):

$$\text{F1 score} = (2 * \text{precision} * \text{recall}) / (\text{precision} + \text{recall}) \qquad (2.6)$$

2.3.3.4.5. Cohen’s Kappa Agreement

K-agreement (Cohen’s kappa) is a statistical measure for cross analysis created by Jakob Cohen (2006), in order to evaluate the reliability or consistency between two or more annotators who are assigning qualitative (categorical) labels to a dataset. This technique was created as he recognized that simple agreement percentages could be inflated or deflated by random chance alone, and as such, it was introduced as a statistical measure that could adjust for the chance agreement was needed (Cohen, 1960) to ensure results are chance-corrected and reflect a true consensus between the labellers (Chicco et al., 2021). When employed in machine learning to evaluate binary classifications, Cohen's Kappa formula can be calculated according to equation 2.7, where P_e is determined using equation 2.8:

$$kappa(\kappa) = \frac{P_o - P_e}{1 - P_e} \qquad (2.7)$$

where

$$P_e = \sum_{i=1}^{k} P_{i+}P_{+i} \qquad (2.8)$$

The Kappa's formula is the agreement between two raters, where each annotator annotates every item. The P_0 is the relative observed agreement among raters (identical to the accuracy), and P_e is the hypothetical probability of chance agreement. P_e is calculated by taking the sum of the products of the row and column marginal proportions. Marginal proportions are the total proportions of observations in each row and column of a contingency table. By multiplying the marginal proportions of each row and column and then summing them, we get an estimate of the amount of agreement that we would expect to see by chance (Artstein & Poesio, 2008; Datanovia, 2021). The score ranges from -1 to 1, where 0 indicates an agreement that is not better than random chance. Within research deciding what counts as an adequate level of agreement is still not firmly established, where different fields and different tasks accept different levels of agreement. As such, using what Neuendorf (2017) considers acceptable reliability, a threshold was established. A reliability coefficient of 0.8 or greater would be considered acceptable in almost all situations. Any value below 0.8 would be considered as having too many disagreements (Artstein & Poesio, 2008) on the annotation scheme, and the proposed guidelines would not hold up.

CHAPTER 3

Analysis and Discussion

3.1. Data Analysis Process

3.1.1. Data Retrieval

Data extraction is a crucial component of any data analysis project. It involves collecting and retrieving data from various sources in order to transform it into a format that can be easily analysed. While it is a time-consuming process, the success of any machine learning project is dependent on the quantity and quality of its extracted data.

We applied a standard machine-learning workflow. First, we collected and prepared a dataset used to develop a machine-learning model. In order to better explain our retrieval process we will approach this step by step. To avoid potential biases, all relevant political parties in the current Portuguese political landscape were included in our observed population. Focusing only on one or two could lead to a skewed and one-sided view of populism in political communication, perpetuating biases and reinforcing incorrect stereotypes. The data selection followed specific spatiotemporal criteria, as it was decided that data would be retrieved from all represented political party leaders over the two most recent parliamentary election cycles (October 2019 up to when the data was retrieved in March 2023). This timeframe is particularly interesting, in that it saw some significant changes in Portugal's political landscape, such as the introduction of populism as a mainstream concept in the country's discourse, as well as the emergence of new, relevant political parties[20] that injected fresh energy and new ideas into Portugal's political marketing and communication scenes. The analysis considered only those parties that obtained parliamentary seats in both election cycles, in that sense, *Centro Democrático Social-Partido Popular* (CDS-PP) and *Partido Ecologista os Verdes* (PEV), whom had seats only in the 2019-2022 election cycle were disregarded.

[20] Following each party's legalization in the previous years, Iniciativa Liberal, Chega, and Livre elected parliament members for the first time in their history in the 2019 election.

For the remaining parties, tweets were collected from the party's leader in the respective cycle, meaning that, In the event that a party leader changed between cycles, both leaders were included in the analysis, with the caveat that, in these situations, the tweets were only collected in respect to the time they were in charge of the party. Furthermore, in situations where the designated party leader does not have a Twitter account, or where the party is led by an executive body, the president of the parliamentary group would then be considered for the selection. In the unlikely scenario where this member does not have a Twitter account as well, the analysis would then fall down to the vice president of the parliamentary group. This measure ensures that all parties are represented equally during both electoral cycles in the analysis, regardless of occurring changes, their leadership structure, and individuals' social media presence.

Following this retrieval hierarchy, the following parties and individual accounts were considered. From *Partido Socialista* (PS) tweets were collected from Antonio Costa (@antoniocostapm), who served as the president of the party and was elected prime minister in both election cycles. From *Partido Social-Democrata* (PSD), tweets were collected from Rui Rio (@RuiRioPT), who served as the party's president in the 2019-2022 election cycle, and Luís Montenegro (@LMontenegroPSD), who has been serving as the party's leader since 2022. From *Bloco de Esquerda* (BE), Tweets were collected from Catarina Martins (@catarina_mart), who has been serving as the party's leader in both election cycles[21]. From *Chega!* (CH), tweets were collected from André Ventura (@AndreCVentura), who has been serving as the party's leader in both election cycles. From *Iniciativa Liberal* (IL), tweets were retrieved from both João Cotrim Figueiredo (@jcf_liberal), who served as the party's leader between 2019 to January 2023, and Rui Rocha (@ruirochaliberal), who has been serving as the party's leader from that point on. From *Partido Comunista Português* (PCP), tweets were retrieved from João Oliveira (@joao_g_oliveira), who served as the party's president of the parliamentary group in the 2019-2022 election cycle, and Bruno Dias (@brunoramosdias), who is serving as the vice president of the parliamentary group in the current election cycle. Lastly, for *Livre* (L), tweets were gathered from Joacine Katar Moreira (@KatarMoreira), who served as the party's president of the parliamentary group in the 2019-2022 election cycle, and Rui Tavares (@ruitavares), who is serving as the party's president of the parliamentary group in the current election cycle.

The tweets were then obtained using Python through the Twitter API, with the aid of the Tweepy library (Roesslein, 2009), using the "user-timeline" function. This function was chosen as it allowed for the retrieval of the most recent tweets posted by any specific users (Chaudhary & Niveditha, 2021).

[21] Catarina Martins has already announced that she will leave this position, as she will not be running for the party's president position in the following elections on May 28th, 2023.

During the collection procedure, these parameters were applied to each account. In most cases, the creation of the politician's account preceded the individual's specified time criteria. To correct this problem, a line of code was added to check if the creation date of a tweet was after the specified start date, whereas through the "datetime.strptime" function and "before_date" filter, if the tweet was created before the start date, it was not included in the final dataset. For the opposite problem, where accounts published tweets after it was no longer relevant for the research, the inverse line of code was applied, where only tweets that were created on or before the specified date were collected for the final dataset.

One of the biggest limitations of the Twitter API, the "user_timeline" function is only able to retrieve the most recent 3 200 tweets from a specific account. Due to this, both Joacine Katar Moreira's tweets from 2019 to May 2021 and António Costa's tweets from 2019 were not able to be retrieved using this process. With this caviat, all individual user datasets were saved into a single file, with a repository of about 19 thousand individual tweets, distributed by politician as presented in Figure 3.1.

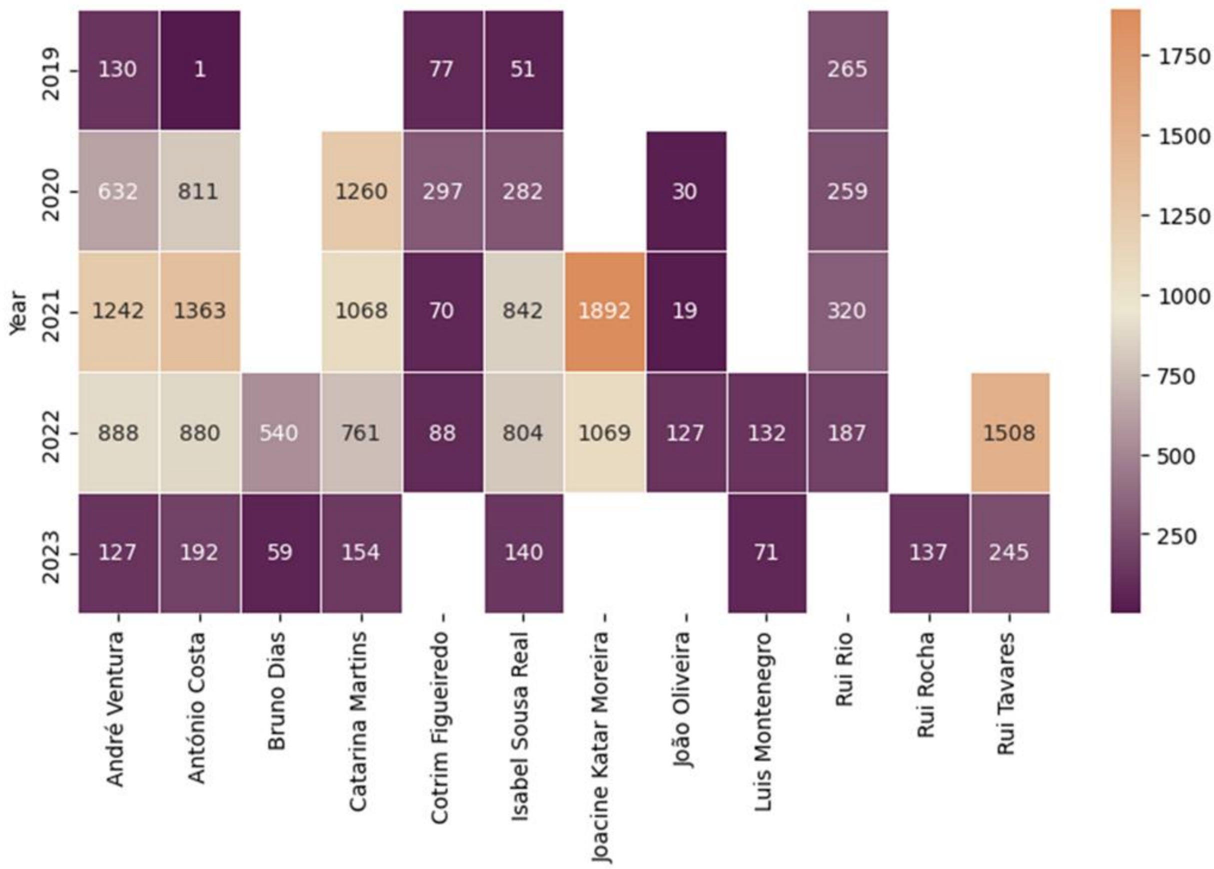

Figure 3.1. Heatmap matrix of individual tweets by political actors and year

While politicians are generally expected to uphold a formal code of conduct online, presenting themselves in a composed and dignified way, Twitter is still a social media platform that follows its own particularities when communicating online. The unstructured nature of tweets, which often include expressions in different languages, use of abbreviations, hashtags, emojis as expressions, slang words and terms, unaccounted-for spelling and grammar mistakes, as well as the misuse of capitalized charecters in sentences, makes it crucial to clean our data prior to the tokenization process. As such it is here that starts our NLP pipeline.

The cleaning techniques were applied using SpaCy (Honnibal & Montani, 2017) with some specific particularities. SpaCy (Honnibal & Montani, 2017) is a powerful open-source library for NLP in Python. It is designed to work with text in many different languages and can perform a wide range of NLP tasks. It is particularly useful for its adaptability, where SpaCy (Honnibal & Montani, 2017) allows for the customization of pre-existing models, the creation of models from scratch to meet specific requirements, as well as having the ability to integrate your own data to user-generated statistical models trained on large corpora by other popular machine learning libraries (Vasiliev, 2020).

3.1.2. Data Annotation

After cleaning the gathered dataset, we randomly selected a sample of 75 tweets from each observed politician, removed their account names, and further randomized the tweet in order to avoid bias as much as possible. The purpose of this was to label the tweets as either populist or non-populist. This labelling step would be done by two individual labellers following the same grid of theoretical guidelines, who were already previously familiar with the concept, of what is considered populism.

These theoretical guidelines on populism were conceptualized using the thin-centred ideology definition formulated by Cas Mudde (2004), defined in the previous chapter, where this is then reflected in individuals' discourse, being considered an attribute of its text rather than a feature of a political agenda (Rooduijn, 2014). For our study on identifying tweets related to populism, we decided to start by adapting the guidelines suggested by Rooduijn & Pauwels (2011) for labelling text as people-centrism or anti-elitism ideas. While we recognized the value of these guidelines, we opted for a binary classification model for its simplicity and efficiency. In this model, we exclusively labelled each tweet as either "Populism" or "Non-Populism". However, we did use the aforementioned guidelines to inform our labelling process and used elements of people-centrism, anti-elitism, and anti-pluralism as the basis for our classification of tweets as "Populism". By doing so, we aimed to capture the complexities of this highly nuanced topic while maintaining a manageable and effective classification system.

In our labelling process, we added a fourth concept that we referred to as "the emotional appeal." Inspired by the work of Keith Dowding (2018) on political persuasion, the emotional appeal involves the use of specific language that conveys emotions and connects these emotions with differing concepts to elicit stronger emotions in the audience or to reinforce their existing emotions. When categorizing tweets, we considered instances where individuals used language that was particularly vivid, moralistic, evocative, or metaphorical in order to convey a strong emotional message. In these cases, we categorized the tweets as having a populist appeal.

Binary classification is a type of supervised machine learning that requires classifying data into two mutually exclusive groups. The model uses the labelled dataset to learn and predict the possible label for new data points, by identifying keywords and patterns in the input data and finding correlations between their features and the target variables (Chzhen et al., 2019; Kumar, 2022). While the underlying mechanism in binary classification can vary depending on the model applied, the underlying principle of finding patterns to make predictions will always be the same. This kind of classification was chosen instead of multi-label and multi-class for its relative simplicity. In this research, the objective was to find patterns in grammar and semantics only needing a distinction between either "Populism" or "Non-Populism", two opposite concepts, and as such, we did not think the computing and complexity needed to further delineate if a tweet was populist because of different sub-set parameters would merit the effort it would be needed to ensure a quality model. To better exemplify we will use the following tweet as an example:

"Infelizmente, o Governo é hoje a expressão do PS e vive um espírito de desunião, de divergência, de luta do poder pelo poder. O PS está preocupado com o seu futuro e não com o futuro dos portugueses." – Tweet ID: 1636317662383882240

While this tweet could be specifically coded into having people-centrism and anti-elitism properties suggesting that the government is more focused on its own interests and power struggles rather than the needs and interests of the people, a more complex labelling system would lead labellers to focus on more technical and would exponentially increase work and time spent on it. In this case, where just identifying specific populist features such as anti-elitism, people-centrism, and anti-pluralism in the text was enough both labellers were quick to label the tweet as Populist.

Due to the complex nature of capturing populism in annotated resources, it was decided that using multiple labellers could help increase the accuracy, quality and generalizability of the results by accounting for potential individual biases or idiosyncrasies that may be present (Sheng et al., 2008).

One fundamental assumption within empirical research in computational linguistics is that labels are more reliable if the coders can be shown to be in agreement when assigning labels (Artstein & Poesio, 2008; Cohen, 1960). Inter-annotator agreement (IAA) is a measure used to assess the level of agreement or consistency among different annotators who have labelled the same dataset. If coders produce consistently similar results, then one can expect them to show a similar understanding of the annotation guidelines, and as such, perform this task consistently (Artstein & Poesio, 2008; Cohen, 1960). While it is not possible to assure 100% validity in the labels assigned by human annotators, as it is difficult to completely eliminate subjectivity and bias, reliability is certainly a prerequisite for demonstrating validity (Artstein & Poesio, 2008; Cohen, 1960; Passonneau et al., 2006). In order to calculate these, several measures of IAA have been proposed over the years. In our case, Cohen Kappa coheficient was chosen.

The data annotation process was conducted using the Label Studio platform, an open-source annotation tool designed to facilitate the creation of datasets for training machine learning models. The platform's web-based interface allowed for the creation and management of annotations online, providing support, easy collaboration, and sharing of annotation tasks, making the work flexible and simple. We exported the labelled dataset into a JSON (JavaScript Object Notation) file is a lightweight data in text format that is programming language-independent, used to transmit data between a server and a web application. This file contained information such as Tweet ID, Tweet text, Labeller ID, and annotation. Upon applying this to our dataset, we found that the counts for populism were 162 and 130 for annotators A and B, respectively. On the other hand, non-populism counts were 738 and 706 for annotators A and B, respectively.

Overall, our IAA was calculated to be 0.8104 using Cohen's Kappa, indicating a substantial agreement among our annotators, and as such we accepted it to be strong enough to follow through with our research using these annotations.

3.1.3. Data Preparation

The data pre-processing stage involves implementing various steps to prepare the previously retrieved dataset in order to train the model.

The raw dataset obtained from the Label Studio platform, was loaded and stored as a list of dictionaries into our Python environment, where an analysis on the dataframe was conducted to check for any possible missing values. This was done by printing the columns and rows that contained missing values using the appropriate python functions, where afterwards, the extended head of the dataframe was printed. Since no unforeseen gaps in the data were found, our subsequent analysis and modelling steps were ensured to be based on a complete and reliable dataset.

3.1.3.1. Pre-processing and Feature Extraction Steps

In this step, the first thing to do was to create empty lists that will be used to store the relevant features (embeddings, labels, POS tags, and parsed trees) that will be extracted for each data item. This is done so that in the future, all extracted features are easily accessible in a structured manner. Afterward, the respective model[22] was loaded with the SentenceTransformer command.

The code will then loop through each individual data entry, and during each loop it will record and store the extracted features into the respective empty lists. For each loop, the unique tweet identifier (ID), the text of the tweet, and corresponding labels are extracted. Subsequently, the text is subjected to spaCy's natural language processing pipeline.

There were a couple of particularities within this dataset that delineated from the pre-established traditional NLP pipeline. Since the retrieved tweets were in Portuguese, which, as a language, is characterized by having many inflected words (Filipe et al., 2020), stemming would be less effective as it would excessively reduce words, diminishing its overall performance. As such, lemmatization was considered the optimal choice and was consequently employed within the pipeline. Secondly, when removing stop words, the preestablished package from SpaCy for the Portuguese language ended up removing too many words that could be considered meaningful, to circumvent this problem, a personalized stop word dictionary was created from scratch (see Annexes). By implementing lemmatization, lowercasing, and excluding specific words defined in the dictionary, tokens were generated.

The processed tokens, combined to make a processed sentence, are then fed into the SentenceTransformer model, which encodes them into meaningful embeddings with 768 dimensions. Additionally, part-of-speech (POS) tags are extracted for each token, alongside generating the parsed tree representation of the processed text sentence.

Lastly, to ensure that the embeddings were consistent both in size and format, a validation mechanism was implemented to ensure uniformity. The outcome was reassuring, as no discrepancies were detected during this evaluation.

3.2. Modelling

3.2.1. Training and Evaluation

The first step of model training is to split our dataset into training and testing sets in order to implement our logistic regression model for classification. To address the cases where the annotators' labels were not in agreement, the code was divided into two separate tests. In "Test A", a tweet would be categorized as "Populism" if at least one annotator labeled it as such, in "Test B", it would only be classified as "Populism" if both annotators agreed on the category.

[22] We kept it as the "respective model" at this phase to acknowledge that multiple pre-trained models were experimented on. However, further details about the model selection process will be provided in later sections to maintain conciseness and focus.

The test predictions are then compared to the true labels to evaluate the model's performance. To assess the model's preliminary performance, the following evaluation metrics were computed on Test A (Figure 3.2):

a) Confusion matrix

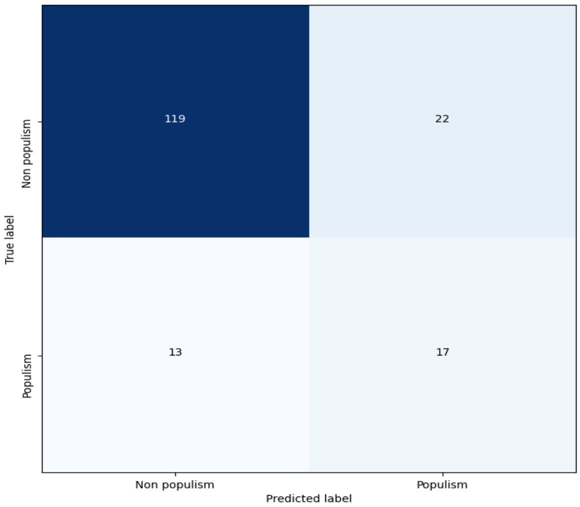

b) Classification report

	precision	recall	f1-score	support
Non-Populism	0.90	0.84	0.87	141
Populism	0.44	0.57	0.49	30
accuracy			0.79	171
Macro avg	0.67	0.71	0.68	171
weighted avg	0.82	0.80	0.81	171

Figure 3.2. Evaluation metrics computed on Test A: a) Confusion matrix; b) Classification report

3.2.2. Experiments

While on a first look, the preliminary model seems to have performed with an overall promising accuracy 0.79 (Figure 3.2b), upon closer analysis, this result can be somewhat misleading. The model exhibited bias towards predicting instances as "Non populism", struggling when it came to predicting instances as "Populism”, with a precision of less than a 50/50 chance. This could mean that the model would be over-prediction instances as “non-populist” since it had a limited understanding and grasp of the characteristics associated with the "Populism" category. The fact that the model only captures approximately 57% of the true instances of "Populism", with a higher probability for false negatives, is further support for that hypothesis. These findings reflect the imbalances in the dataset (131 non-populist for every 30 populist ones), which its detrimental impact on the model's performance should be addressed as a priority.

The experimental modifications were driven by two main objectives based on the results of the preliminary performance analysis. Firstly, at the highest priority, was addressing the issues of the dataset imbalances, which we suspected of requiring more Populism cases through additional labelling. Secondly, there was a need to improve the overall F1 score, which necessitated enhancing the overall model's consistency between categories. To address this, we explored alternative transformer models, specifically those trained for the Portuguese language, which could have variations in architectures, pre-training methods, and hyperparameters for the creation of their embeddings, which could lead to better fine-tuning. Since, each transformer model has its own strengths and weaknesses, experimenting with different models will allow for comparison between them, which will only expand possible options and increase the chances of achieving improved outcomes.

3.2.2.1. Data Imbalances Problems

To address the hypothesis on dataset imbalances, an initial step was taken to implement a resampling technique, by undersampling the majority class. This was done by randomly selecting a subset of "Non populism" cases to match the number of instances in "Populism" one. While in theory, this would help the model learn more effectively from the minority class, what ended up happening was less than ideal, as shown Figure 3.3. (Test resampling):

a) Confusion Matrix (Resampled)

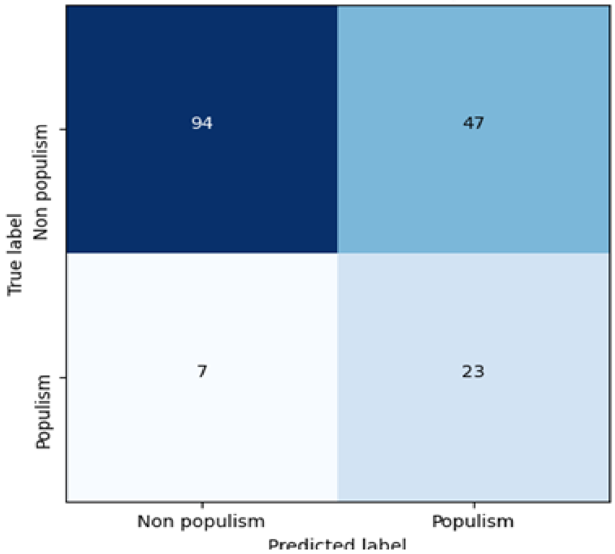

b) Classification report (Resampled)

	precision	recall	f1-score	support
Non-Populism	0.93	0.67	0.78	141
Populism	0.33	0.77	0.46	30
accuracy			0.68	171
macro avg	0.63	0.72	0.62	171
weighted avg	0.83	0.68	0.72	171

Figure 3.3. Evaluation metrics computed on Test resampling: a) Confusion matrix; b) Classification report.

Undersampling ended up confirming our previous hypothesis. It yielded a clear decline in performance when compared to the initial results. The reduction in the majority class led to a loss of valuable information, which significantly impacted accuracy when classifying "Non populism" instances. Additionally, the improvements in recall for "Populism" came at the expense precision, indicating a trade-off between sensitivity and precision in its classification. Overall, downsizing the dataset had a detrimental effect on sensitive knowledge and the overall performance of the model (Mohammed et al., 2020), as such, obtaining a larger number of instances from the minority class became the logical next step to improve our performance.

This approach led us to make two important decisions going forward. Firstly, we decided to drop Test B from future analysis. It was yielding lower scores overall, likely due to the dataset being slightly more imbalanced, having even fewer instances of "Populism" cases. Secondly, when weighing priorities, the need for a greater number of labelled instances representing "Populism" cases, coupled with the substantial agreement between labellers, led to the decision that, future classifications, would only need the judgment of a single annotator. This approach was carefully considered to ensure the research's integrity and the validity of classification would not be significantly compromised, while adopting a more practical and time effective method.

3.2.2.2. Exploring Different Models

At this juncture, the annotator proceeded to sample 250 new tweets from each politician out of the remaining 17 091 unlabelled ones. This process followed the same framework as the initial one, and afterward, the newly labeled data was again exported into a new JSON file with the same format. Considering the specific situation where certain politicians had fewer than 250 tweets available, it was decided to only label the number of available tweets for those specific individuals, to maintain the politician's representation proportions. From the resulting 2 153 labeled cases, 153 were annotated as being populist and were subsequently added. To do this, we applied the same pre-processing steps as the initial dataset to ensure consistency.

3.3. Results and Discussion

To enhance the overall model performance, we conducted experiments by applying logistic regression to four different transformer models (A-D), which involved evaluating the four models following described in Results (Table 3.1).

3.3.1. Results

- **Model A**, Xlm-r-100langs-bert-base-nli-stsb-mean-tokens model[23] is a language transformer model based on cross-lingual architecture (Reimers & Gurevych, 2019). It is built upon the BERT framework and has been trained for cross-lingual downstream tasks on a large corpus of text data from over 100 languages, including Portuguese.
- **Model B**, BERT-base-portuguese-cased[24] (BERTimbau), is a pretrained BERT model built on the transformer architecture, specifically trained for Brazilian Portuguese tasks (Souza et al.,

[23] 100langs-bert-base Is available on https://huggingface.co/papers/1908.10084

[24] bert-base-portuguese-cased is available on https://huggingface.co/neuralmind/bert-base-portuguese-cased

2020). BERTimbau was consider as a strong possibility, having gained popularity for its state-of-the-art performances for certain downstream tasks in the Portuguese language.

- **Model C**, paraphrase-xlm-r-multilingual-v1[25] model, just like model A, belongs to the Sentence Transformers library. It is based on the RoBERTa architecture, which is a multilingual variant of BERT. This model was trained for specific tasks involving paraphrase identification, which is the task of determining whether two given sentences have the same or similar meaning, even if they are expressed differently (Reimers & Gurevych, 2019). This was seen as a suitable model to apply due to its multilingual capability and expertise in identifying paraphrases.
- **Model D**, PORTULAN/albertina-ptpt[26] model, is based on the ALBERT architecture, which is a more compact variant of BERT that reduces the number of parameters while maintaining strong performance. Having been trained with over 2.2 billion tokens, the model had a version specific to the European Portuguese variant from Portugal (Rodrigues et al., 2023). We had high expectations for the performance and potential advantages of the as it emerged during the experimentation phase of our research, being very recent and having the uniqueness of being specifically trained for Portuguese from Portugal tasks.

We applied the same test to all models, reusing the same code we utilized previously, only changing the loaded model. The overall results, shown on Table 3.1, across the four models indicate varying performances. On one hand, the "PORTULAN/albertina-ptpt" model showed a disappointing performance, particularly in terms of the "Populism" label. While the other two model showing satisfying results, the "paraphrase-xlm-r-multilingual-v1" was chosen, as it had the highest precision, recall, and F1-scores for both labels, as well as best accuracy overall.

Table 3.1. Results obtained for the Models A-D

<table>
<tr><th>Model</th><th>Labels</th><th>precision</th><th>recall</th><th>f1-score</th><th>Accuracy</th><th>support</th></tr>
<tr><td rowspan="2">xlm-r-100langs-bert-base-nli-stsb-</td><td>Non-Populism</td><td>0.91</td><td>0.93</td><td>0.92</td><td rowspan="2">0.88</td><td>75</td></tr>
<tr><td>Populism</td><td>0.79</td><td>0.73</td><td>0.76</td><td>26</td></tr>
<tr><td rowspan="2">neuralmind/bert-base-</td><td>Non-Populism</td><td>0.92</td><td>0.88</td><td>0.90</td><td rowspan="2">0.85</td><td>75</td></tr>
<tr><td>Populism</td><td>0.69</td><td>0.77</td><td>0.73</td><td>26</td></tr>
<tr><td rowspan="2">paraphrase-xlm-r-</td><td>Non-Populism</td><td>0.92</td><td>0.95</td><td>0.93</td><td rowspan="2">0.90</td><td>75</td></tr>
<tr><td>Populism</td><td>0.83</td><td>0.77</td><td>0.80</td><td>26</td></tr>
<tr><td rowspan="2">PORTULAN/albertina-ptpt</td><td>Non-Populism</td><td>0.78</td><td>0.90</td><td>0.84</td><td rowspan="2">0.74</td><td>75</td></tr>
<tr><td>Populism</td><td>0.50</td><td>0.28</td><td>0.36</td><td>26</td></tr>
</table>

[25] paraphrase-xlm-r-multilingual-v1 is available on https://huggingface.co/sentence-transformers/paraphrase-xlm-r-multilingual-v1
[26] albertina-ptpt model is available on https://huggingface.co/PORTULAN/albertina-ptpt

Lastly, as part of our experimentation, we decided to skip the filtering of stopwords (Model E). This was influenced by our previous assessment, which had lead to the creation of a specific dictionary, combined with observations from other studies. They indicated that removing stopwords could disrupt the syntactic structure of Portuguese sentences (Rodrigues et al., 2023). Surprisingly, this modification appeared to enhance the overall performance of the model, particularly in terms of recall for Populism cases, which can be observed in Table 3.2.

Table 3.2 Results obtained for the Model E

Model	Labels	precision	recall	f1-score	macro avg	support
paraphrase-xlm-r-multilingual-v1	Non-Populism	0.93	0.95	0.94	0.91	75
	Populism	0.84	0.81	0.82		26

With these results, we gained confidence in the model's performance in capturing the features necessary to classify tweets as either populist or non-populist. To further explore its capabilities, we applied the trained model to classify the remaining unlabelled tweets, as it would allow us to gain insights into how the model captured the distinguishing characteristics of populist and non-populist interactions. In the next step we would compare the classification features between the manually labelled cases and those classified by the machine learning model.

The statistics presented in this next section are derived from this classifications, from both the manually labelled tweets and the machine generated ones, on the entire corpus, which resulted in exactly 19 496 tweets. Of this 19 thousand 15 662 were labeled as Non populist tweets (14 444 machine generated and 1 218 manually) and 3 358 were labeled as Populist (2 728 machine generated and 2 728 manually). As illustrated in Table 3.3 and Figure 3.4, and due to the nature of what makes something “Non-populist” being exponentially broader than what makes something “Populist”, results in the frequency of “Populist” occurrences are relatively lower in all political actors. However, from the percentage of “Populist” tweets, we can establish three tiers of populist activity from our combined labeling[27].

The low populist activity (between 0 to 15%) group is comprised of António Costa, João Oliveira, and Rui Tavares. Medium populist activity (between 16% and 25%) is comprised of Bruno Dias, Catarina Martins, Isabel Sousa Real, Joacine Katar Moreira, Luís Montenegro, Rui Rio, Rui Rocha, and Rui Tavares. Lastly, the high populist activity (over 25%) group is comprised of Cotrim Figueiredo and André Ventura.

[27] A table presenting the machine and manual division of label frequency is present in the Annexes as Table A1.

Table 3.3. Populist and Non-Populist Label Frequency among Political Actors

Name	Non Populist	Populist	
	N	N	%
André Ventura	1 969	1 050	35
António Costa	3 186	61	2
Bruno Dias	498	101	17
Catarina Martins	2 719	524	16
Cotrim Figueiredo	336	196	37
Isabel Sousa Real	1 784	335	16
Joacine Katar Moreira	2 404	557	19
João Oliveira	162	14	8
Luis Montenegro	160	43	21
Rui Rio	794	237	23
Rui Rocha	107	30	22
Rui Tavares	1 543	210	12

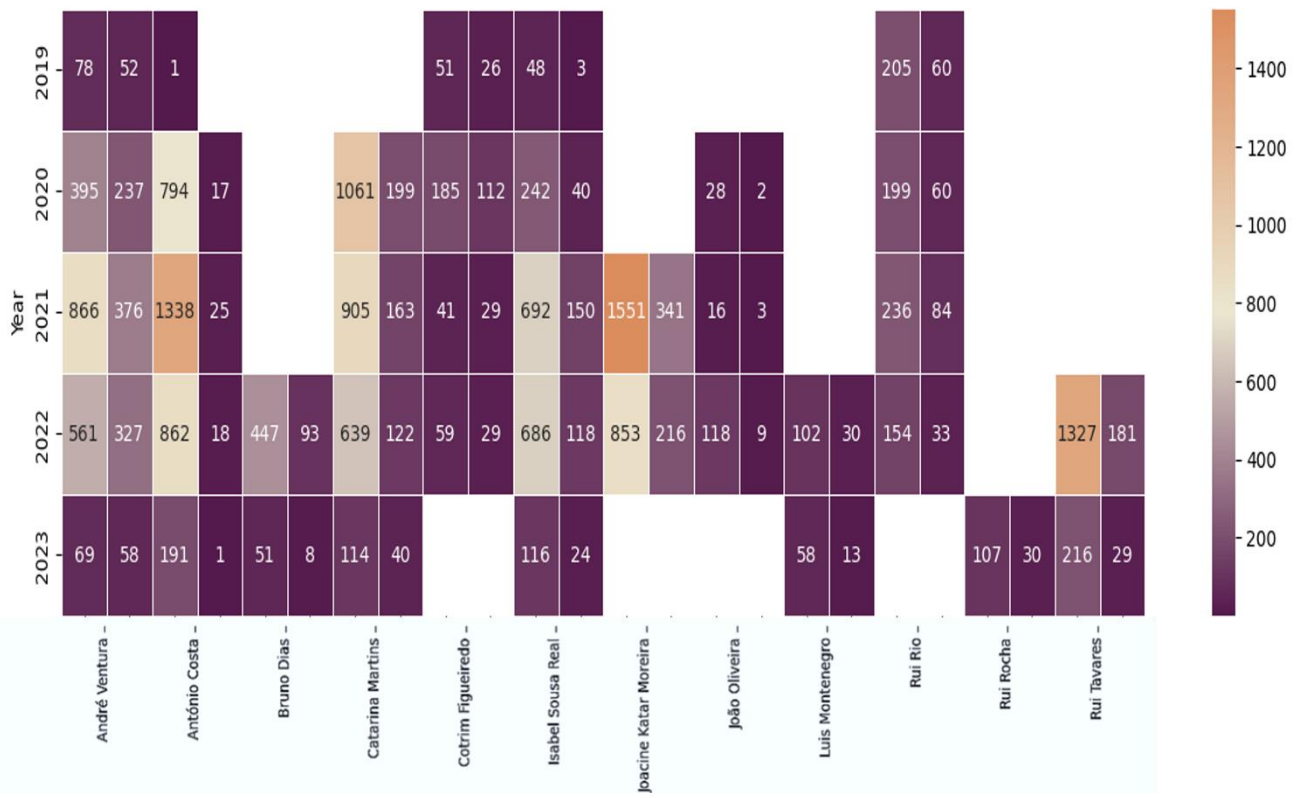

Figure 3.4. Updated Heatmap matrix of individual tweets by political actors and year

(left column represents Non-populist tweets and right column represents populist ones)

In this end, the perceived Portuguese reality seems to remain more less intact when characterized by our model. António Costa, being the prime minister during both election cycles, is synonymous with the perceived "establishment" so it would be natural for him to have by far the lowest percentage of populist moments. Rui Tavares' low amount of populist-labeled tweets, contrasted with his prolific tweet activity, sets him apart as a noteworthy outlier among other political actors. Having joined Twitter in 2008, when he was still a member of the BE party, Rui's behaviour on twitter seems more personal than most of other actors, having over 43 thousand tweets since joining, and averaging over two thousand tweets every year, consistently maintaining the highest tweet volume and longevity among all accounted-for actors, by a significant margin. This idea is further supported by the factors seen in Figure 3.5. The fact that his tweets are almost evenly distributed by all hours of the day[28], and his relatively high activity during after business hours (around 50% of his activity), and highest activity hours between 11 p.m. and 1 a.m. encouraging the idea that it is in fact a personal account.

Lastly, André Ventura's presence in the high populist activity group, and his claim to the highest number of populist labeled tweets, comes as no surprise when considering the context of Portuguese reality. Being largely considered to be the main populist actor in Portugal, and being synonymous with populist discussion in the country (Parrança, 2022), perhaps the most surprising revelation is that, to our model, he did not display the highest percentage of populist activity. Cotrim Figueiredo exhibited 2% more populist activity, albeit with the caveat of having under one-tenth of Ventura's sample size.

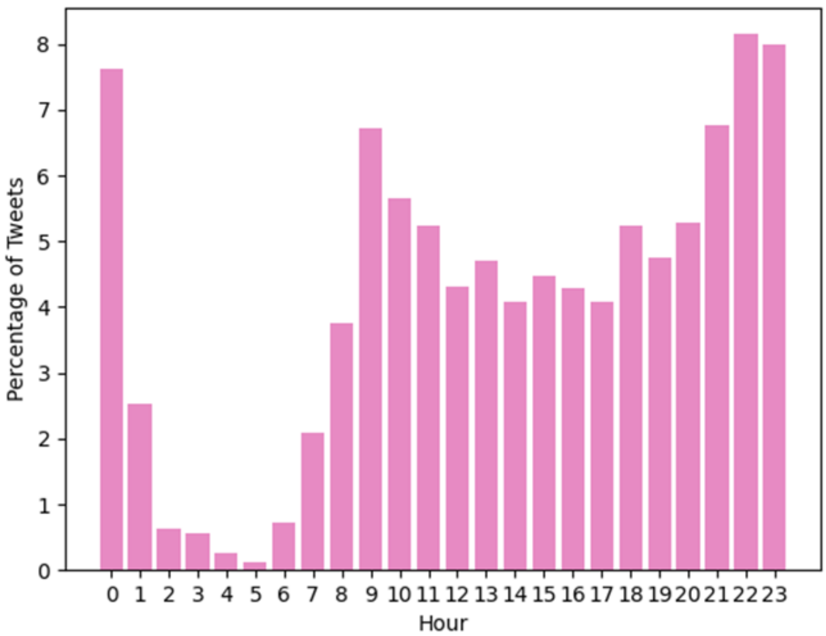

Figure 3.5. Rui tavares Tweet distribution by hour (%)

3.3.2. RQ1 Testing

With the model's performance being quite high, another metric analysed to indicate the actor's behavior in the Portuguese reality was their tweets distribution over the two election cycles (2019-2023). Figure 3.6. is a helpful visual representation to assess the reliability of the automated labels.

[28] Period between 2 a.m. and 6 a.m. is comparatively low as it corresponds to typical sleep hours.

With an R^2 of 0,65, it seems that when looking at the individual's actors populist tweets by year, the model's predictions indicates that approximately 65.13% of the variability in the machine-labeled data can be explained by the manual-labeled data. In normal conditions, it is reasonable to expect that months with a higher overall number of tweets would result in a greater amount of data available for labelling. This becomes even more apparent when looking at the tweets published by month in Figure 3.7., whereby disregarding the individual actor and label and just summing all tweets published, the R^2 jumps to 0,81.

However, in contrast to the machine-generated labels, which were applied to a very significant amount of the dataset, the manual labeling was conducted using a random sampling on a relatively short and fixed number of tweets. This random sampling approach, while effective for the task of actor representability in the annotation step, may not adequately represent the broader Portuguese reality, which somewhat limits this metric when making inferences.

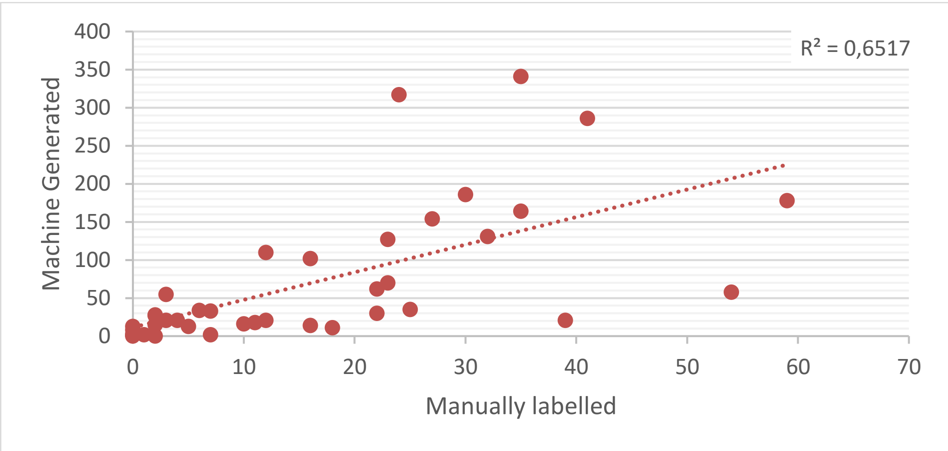

Figure 3.6. Overall Monthly Manual and Machine-Generated Populist Tweets: Scatterplot and R-squared value

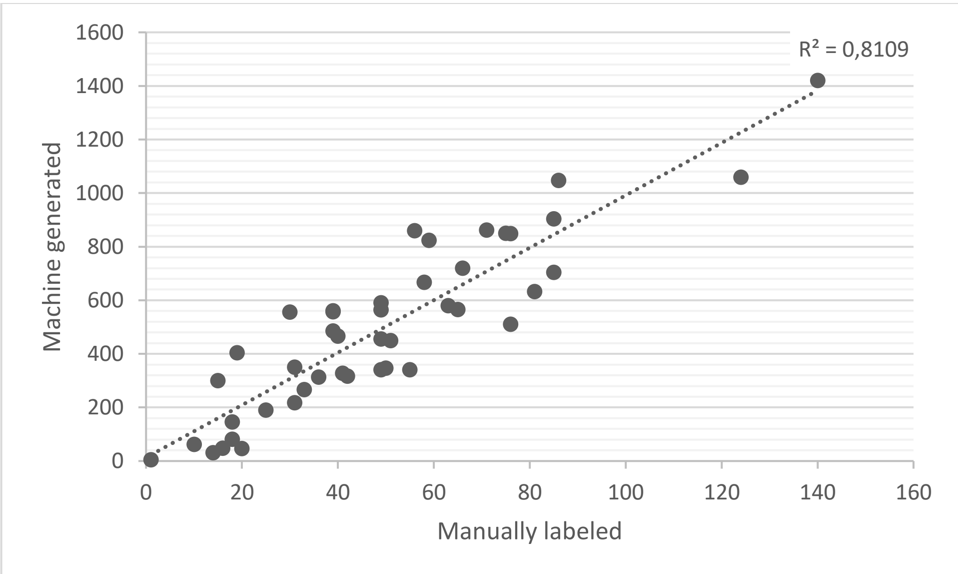

Figure 3.7. Monthly Manual and Machine-Generated Tweets: Scatterplot and R-squared value

This is more visually apparent when looking at the distribution of machine and manually labeled "Populist" tweets over the 4 year period, as present in Figure 3.8 and compare it to the overall distribution of "Populist" tweets and the behaviour of all tweets (Figure 3.9). With both machine a manual labels behaving similarly, when juxtaposing the 'Populist' tweets with our global tweet behavior, we observe that they exhibit a stronge resemblance. This finding reinforces the proposed idea that a larger volume of tweets is followed by a correspondingly larger volume of populist ones.

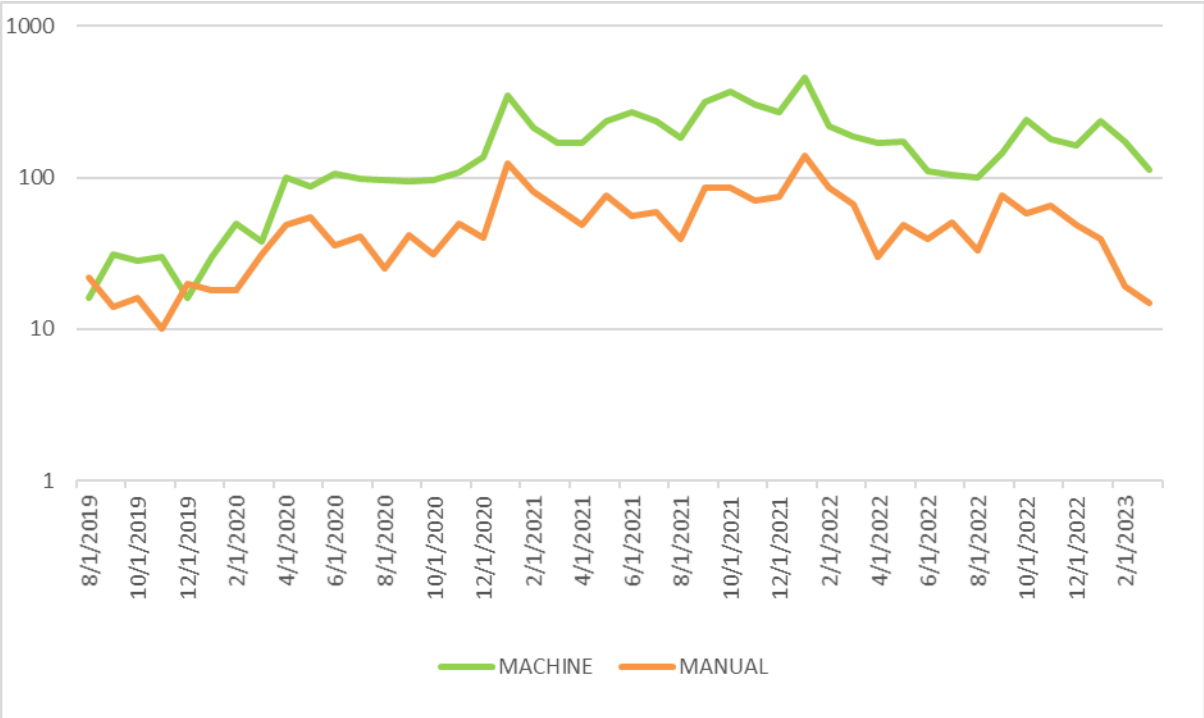

Figure 3.8. Distribution of manual and machine labelled Populist tweets over the two election cycles

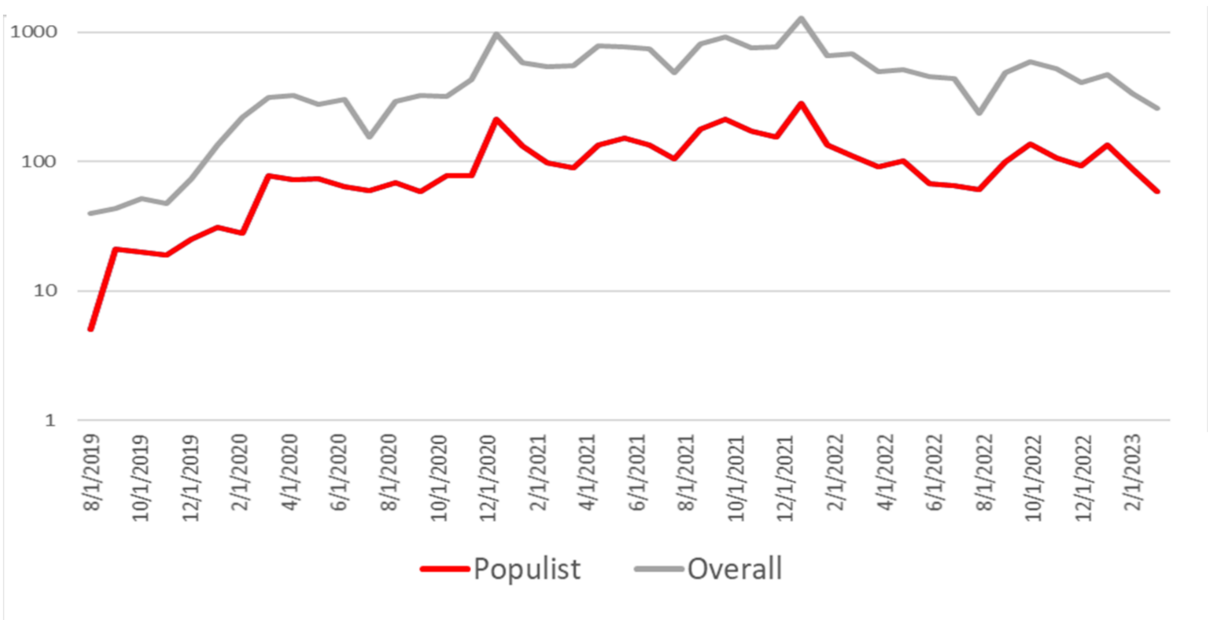

Figure 3.9. Distribution of Populist tweets and the overall distribution of tweets over the two election cycles

With this evidences we can assert that the model's behaviour in identifying and classifiying populist moments reinforces the notion that our machine learning approach is well-suited to effectively detect and analyze populist trends within the data. As a result, we can have a high level of confidence in the model's ability to generalize and capture populist moments across a broader spectrum of untrained tweets in the future, making it a reliable tool for this specific task. As such, moving forward, we will always consider both machine-generated and manually labeled tweets as one unified dataset, given our high level of confidence in their representativeness.

3.3.3. RQ2 Testing

To study the linguistic realization of the tweets, in order to capture the associated discourse in the Portuguese reality, we started by exploring the most frequent PoS tags in both “Populist” and “Non-populist”, as present in Figure 3.10. While the results show striking similarities between both categories, two notable distinctions emerge: a significantly lower occurrence of proper nouns (PRONP), and a relatively higher frequency of verb usage in populist tweets.

The unexpected aspect on the first point emerges when considering that populist communication is closely associated with personalistic rhetoric. This relationship could be connected to the inclination of populist moments to focus less on specific policies or institutions and instead prioritize their connection with the people, emphasizing their ability to represent the will of the masses. This hypothesis finds slight support in the analysis of pronoun usage, where their overall prevelence, and especially of words like "me" or "I", is higher in populist labeled tweets, with just these two appearing in approximately 36.3% of the content, compared to 28.8% in non-populist tweets.

The second point is well-aligned with existing literature on the populist phenomenon, as the observed higher average use of verbs in populist tweets could be attributed to the prevalence of emotional verbs. These emotional verbs often convey a state of emotion rather than any specific action, thus enhancing the overall emotional impact of the message.

Moreover, research suggests that highly emotional verbs are more easily integrated into neutral context sentences compared to other types of high-arousal words (Bayer, 2010). This findings could shed light on why populist tweets, which aim to evoke strong emotions and resonance with the audience, might favor emotional verbs as a powerful linguistic tool.

It is also important to stress the explicit mention of specific words, as such Table 3.4 has the most common N-grams (unigrams, bigrams and trigrams). From the most common populist tweets spaning the last two election cycles, we can gather two important insights.

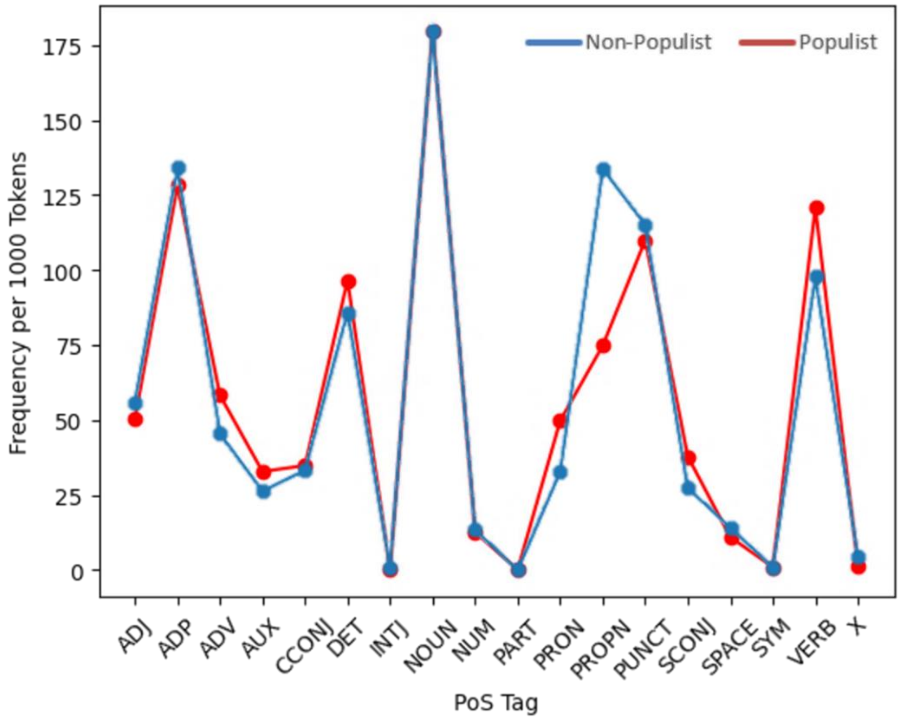

Figure 3.10. Frequencies per 1 000 tokens of PoS tags in Populist and Non-populist tweets

Firstly, unigrams depict that the most prevalent themes are political and national themes with the most common tokens, disregarding stopwords, being “governo” (government), “Portugal”, “portugueses” (Portuguese people), and “país” (country). Of the top five, “Chega”, can somewhat be grouped with the previous ones, with the particularity that it comes highly associated to the large number of populist labelled tweets its parliamentary leader André Ventura reported.

Secondly, the most common bigrams and trigrams directly name and depict prominent political figures and parties of these two election cycles, such as "António Costa", "Iniciativa Liberal", "Bloco Esquerda", "Pedro Nuno Santos" (ex Minister of Infrastructures), "Marcelo Rebelo Sousa" (the President of Portugal), “Aristides Sousa Mendes” (Portuguese Consul-General in Bourdeux during WW II), "Rui Rio", "Ana Gomes" (2nd most voted candidate of the presidential elections of 2019), "André Ventura" and "Pedro Adão Silva" (current Minister of Culture). This strong presence of names in tweets reflects the strong and direct nature of populist moments when engaging in political discussion and government actions.

When doing the same kind of introspection to non-populist tweets certain parallels can be established (as present in Table 3.4 and Table 3.5). Firstly, the same broad political and national themes are common to both types of tweets, indicated by the strong presence of the same tokens, such as “governo” (government), “Portugal”, “país” (country), “pessoas” (people), “estado” (state) also being reported in the top ten.

However, the analysis of these political debates reveals a notable enphasis on public health and environmental issues, which were not prominently registered in the populist discussions. Tokens such as "Saúde" (Health), "Plano Recuperação Resiliência" (Recovery and Resilience Plan), "Serviço Nacional Saúde" (National Health Service), and "Profissionais Saúde" (Healthcare Professionals) demonstrate an increased focus on health-related matters. Furthermore, the presence of terms like "Justa Verde Digital" (Green Digital Justice), "Combate Alterações Climáticas" (Fight Against Climate hange), and "Direitos Humanos" (Human Rights) indicates a strong connection between non-populist tweets and environmental issues.

The analysis of the overall presence of tokens in political debates reveals a significant focus on discussions related to policies, government initiatives, and sustainability, rather than heavily relying on invoking specific names of political parties and individual actors. This shift in focus suggests a move towards a more issue-oriented discourse rather than a purely personalistic approach. It's worth noting that while the names of certain prominent figures like António Costa, Joacine Katar Moreira, Aristides Sousa Mendes, and Iniciativa liberal still appear frequently, the fact that individual actors names do not dominate the non-populist table suggests the tweets have a more balanced approach, where the emphasis is on the policies and ideas being presented, rather than just personalities or party affiliations.

Table 3.4. Top 10 Ngrams on populist labelled tweets

Top 10	Unigram	N	Bigram	N	Trigram	N
1	Governo	601	António Costa	112	Pedro Nuno Santos	19
2	Portugal	501	Assembleia República	59	Luta contra Corrupção	12
3	Chega	448	25 de Abril	56	Dia 30 Janeiro	12
4	Portugueses	421	Iniciativa Liberal	46	Aristides Sousa Mendes	11
5	País	412	Bloco Esquerda	44	Mil milhões euros	9
6	todos	409	Presidente República	36	Marcelo Rebelo Sousa	9
7	PS	355	Rui Rio	36	Governo António Costa	8
8	Sempre	298	Ana Gomes	33	Presidente Assembleia República	8
9	Estado	270	Povo Português	31	Pedro Adão Silva	8
10	Política	236	André Ventura	31	Aumento Salário Mínimo	8

Table 3.5. Top 10 Ngrams on Non-populist labelled tweets

Top 10	Unigram	N	Bigram	N	Trigram	N
1	Hoje	1843	Iniciativa Liberal	219	Plano Recuperação Resiliência	44
2	Portugal	1472	António Costa	181	Aristides Sousa Mendes	41
3	Todos	1189	Direitos Humanos	167	Joacine Katar Moreira	39
4	País	1176	Assembleia República	139	Serviço Nacional Saúde	33
5	Governo	809	União Europeia	132	Combate Alterações Climáticas	31
6	Chega	725	Presidente República	117	Conferência Imprensa Conselho	31
7	Saúde	718	Alterações Climáticas	109	Governo Português Direto	31
8	Pessoas	685	Bloco Esquerda	90	Justa Verde Digital	26
9	Social	682	Rui Rio	88	Terceira Força Política	25
10	Estado	640	Profissionais Saúde	84	recuperação Economica Social	25

Lastly, the specific dates of 25th of April and January 30th were featured prominently in both populist and non-populist tweets. This notability can be attributed to the distinct significance of these dates within the political context of Portugal.

The 25th of April holds strong anti right-wing connotations and is closely associated with the fight against populism in the country, seeing how the date commemorates the Carnation Revolution of 1974, symbolizing the victory over the previous fascist rulers that paved the way to democracy. More recently, due to its strong links to left-wing parties, in particular the communist party, this date has been challenged by certain political factions that argue that the true revolution to democracy only occurred on the 25th of October, making this date even more intensively debated within the populist phenomenon, as it can be used, positively and negatively, by all political ideologies in this debate.

On the other hand, January 30th corresponds to the election day of the last parliamentary elections. The particular prominence of the mention of this date could be attributed the observed highly engagement of all political actor studied. Besides Joacine Katar Moreira, who was not in contention for reelection and had already been distanced from the party she previously represented, due to controversies during the last election cycle (Borges & Rocha, 2020), all other political actors we involved in these elections. This was reflected in the significant increase in the volume of tweets mentioning January 30th across both populist and non-populist spheres.

Having established the presence of shared and individual themes in both populist and non-populist tweets, the last crucial step is to conduct a contextual temporal analysis on the tweets, as shown in Figure 3.11, so to gain insights into how the models characterization of these themes is represented in it evolution over time. Given the similar temporal distribution of both populist and non-populist tweets, any conclusions drawn during this temporal analysis could be applicable to both types of tweets, as such there would be no point in having both lines represented in the graph.

To enrich our understanding of this temporal evolution, we complemented the analysis with a comprehensive list of notable events that occurred during the four-year period (details of which are provided in the Annexes). These events serve as essential contextual markers, helping us contextualize and interpret any fluctuations or shifts in in the populist tweets.

The most influential factor that emerges prominently from the analysis is the impact of the COVID-19 pandemic. The data clearly indicates the pandemic had a profound effect on quotidian life, as evidenced by the significant spike observed in the analysis. Moreover, the values post-March 2020 consistently remained above the pre-pandemic levels. Another recurring pattern in the number of tweets is evident during the month of August, followed by a spike in activity in september. This decline can be attributed to the parliamentary activity being closedin August, leading to a subsequent "summer holiday" break for political actors, resulting in a predictable drop in their utilization of Twitter for work-related activities, followed by a more prominent return.

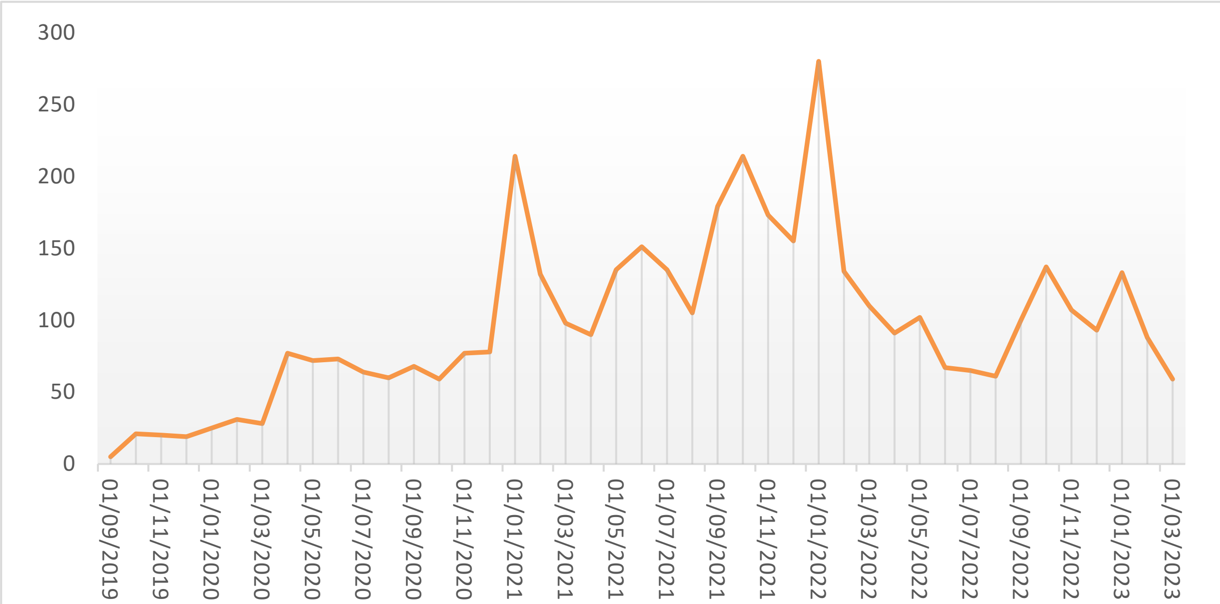

Figure 3.11. Contextual temporal analysis on the tweets

With this information established, a careful examination of Figure 3.11 allows us to visually identify five significant spikes (excluding the one in March 2020). These spikes occurred in January 2021, June 2021, October 2021, January 2022, and October 2022.

While keeping the rest of the analysis in chronological order, we will start by discussing the spike of January 2022, as we have previously dabbled in its explanaitaion when talking about the parliamentary elections and the January 30th date. With the elections being the only major highlight of the month, further context related to this event would be the PS achieving an absolute majority and Chega becoming the 3rd most voted party. When examining the most common Ngrams for January 2022, it becomes evident that there was a substantial focus on political engagement and the electoral process. As terms such as 'voto' (vote), 'votar' (to vote), '30 janeiro' (January 30th), 'legislativas' (parliamentary elections), 'maioria' (majority), and 'debate' dominated the discourse. Political parties and leaders were prominently featured in both populist and non-populist tweets, with 'Chega', 'André Ventura', 'Iniciativa Liberal', and 'PS' being recurrent themes in bost cases. Non-populist tweets also exhibited mentions of 'Livre', 'Jerónimo Sousa' (parliamentary leader of PCP at the time), 'Rui Tavares', and 'Bloco Esquerda', albeit to a slightly lesser extent than the other names mentioned.

Interestingly, TAP (Transportes Aéreos Portugueses) also emerged as a prevalent subject exclusively in populist tweets during this date, with tweets surrounding token variants of 'tap privatizada' (privatized TAP) and 'companhia aérea' (airline company) appearing in at least , the decision of privatization, its financial implications, and its future. Notably, the airline's nationalization by the Socialist government seems to be a contentious topic, having drawn considerable and continuous criticism ever since it happened, making it a significant point of debate during the election.

Receding in time to January 2021, this date is marked by the presidential elections and the sharp increased of COVID-19 cases, due to relaxation of containment policies around the Christmas time, which occurred at the peak of the pandemic. The first topic is reflected on both populist and non populist tweets, as most of the months most common ngrams are words related to the elections or, in the case of the populist ones, directly name one of running candidates. The only exception was 'Ana Gomes', whom was by far the most mentioned bigram in both populist and non-populism, likely due to her being the primary opponent of the eventual winner, Marcelo Rebelo Sousa, whose name did not make any appearance in the lists for the month, and her antagonising role and rivalry of candidate André Ventura, which was one of the main narratives of the elections.

However, the distinct focus in the COVID-19 pandemic is only apparent in the non-populist ngrams while being completely unmentioned in the populist subset. This is apparent with the topic of "saúde" (health) coming just shy of the non-populist top five, as well as the occurrence of terms such as "emergência" (state of emergency), "plano vacinação" (vaccination plan), "combate pandemia" (pandemic combat), and "vacinação contra covid19" (vaccination against COVID-19) in the top 20.

The June 2021 month spike in twitter activity is also directly related with the events of the month, namely the appointment of Pedro Adão e Silva as the organizer of the 50th year celebration of the Carnation Revolution, and the controversy around how much he would get payed for it; the hit and run incident caused by at the time Minister Eduardo Cabrita; and the leak of personal data of Ukranian protestors to the Russian Embassy by President of Lisbon Municipality Fernando Medina. The names of the individuals involved occupied the number one, two and five spots as the most mentioned bigrams por the populist tweets, and second, seventh and 18th place for non-populist, respectively.

A similar thing can be extrapolated on to the October 2021 date, a significant focus of ngrams centered around the main event of the month in Portugal, being the rejection of the yearly state budget proposed by PS by both BE and PCP. This rejection resulted in the dismantling of the parliament and triggered a subsequent political crisis, ultimately leading to the early parliamentary elections of January 2022. The prominence of "orçamento de Estado" (state budget) and other related words in the ngrams analysis demonstrates how dominant this topic was for both populist and non-populist tweets.

For October 2022 the Ngrams data reflects the presence of the normal mix of current events, political discussions, and societal issues regarding inflation and tax burdens without any particular thematic focus for both populist and non-populist tweets. As such we couldn't draw any links with what we established as the major events of the month.

While there is certanly a possibility to keep further analyzing our data from the remaining months, we feel strong enough on the evident that the framework we have designed adequately represents the Portuguese political reality. The thorough examination of this dataset has revealed distinct and noteworthy differences between the features of populist and non-populist tweets, rendering them suitable to be utilized as binary variables in future analysis. These differences are evident not only in their behaviour when presented with the same monthly event, but also in their focuses, patterns and rhetorical styles, supporting the relevance of political actors' tweets as data to gain insights into the populist phenomenon, and our models ability to capture them.

CHAPTER 4

Final Remarks

This study enabled us to address the research questions outlined. On one had, we aimed to present a model capable of automatically classifiying populist moments on Twitter. On the other hand, we sought to characterize the portuguese reality over the last two election cycles using that same algorithm. By using a deep learning model to create highly detailed embeddings, and only then applying a simpler, logistic regression for binary classification proved to be a satisfyingly effective approach as the model showed promising results. Within the test data, a macro average of 0.91 and a recall at worst of 0.81, supports the acceptance of our model to do this kind of task.

Even when applied to new and unforeseen data, our model demonstrated a notable capacity to capture the intricacies inherent to both non-populist and populist instances. The distribution of tweets categorized by label and further segmented by political actors indicated a three-tier system, hinged on the percentage of "Populist" tweets, which yielded a nuanced portrayal of each actor's engagement with populist discourse, aligning well with our initial expectations. The data analysed suggested that the model was effectively able to capture how populist tweets in the Portuguese political discourse emphasize personal connections and emotional appeals, while non-populist tweets exhibit a shift towards issue-oriented discourse, which stays in line with predominant populism theories. This observation was further substantiated by our temporal analysis, where scrutiny of populist and non-populist tweets over time unearthed significant temporal patterns that influenced both forms of rhetoric. Remarkably, events like the COVID-19 pandemic, cyclic election patterns, and seasonal variations such as summer vacations, all underscored the dynamic interplay between prominent events and shifts in rhetorical behavior, which the model was able to capture. Ultimately, the further explorations augmented our comprehension of the nuanced ways in which our model was able to capture the two different political communication rhetorics.

However, it is imperative to acknowledge certain limitations that impacted the interpretation and broader applicability of our findings for future research. The study's sample size, tailored to encompass key political actors active on Twitter, inadvertently constrained the dataset fed into the algorithm, as a trade-off had to be made between the need for a robust conceptual framework and the quantity of tweets available for analysis. Moreover, the exclusive focus on Twitter as the solitary social media platform may inadvertently omit broader dimensions of an individual's political engagement, potentially concealing distinct patterns that platforms like Facebook or Reddit might unveil. Finally, the employment of binary classification, though essential and the most suitable approach for our study, could inadvertently oversimplify the multifaceted essence of the concept when trying to differentiate between components. An alternative multi-class classification approach, tailored to encapsulate concepts such as people centrism, anti-elitism, and the general will, might provide an interesting lens capable of capturing the intricate gradations and diverse manifestations inherent in populist discourse.

Despite the acknowledged limitations, what we were able to achieve with this model presents itself as a valuable step into in integrating social sciences with computational methodologies. This approach, though currently constrained, offers a promising pathway to redefine the treatment of complex concepts as quantifiable variables. The potential amplification of this methodology through a well-resourced research and annotation team, equipped with robust computation and labor support, holds the prospect of substantially enhancing the analytical depth and scope of this kind of study. By transcending current human constraints of time and the efford needed to analyse complex concepts, like populism, and opening them to the vast amounts of data available to be studied, this fusion of disciplines has the potential to exponentially expand the horizons of social science research, which remains an avenue of considerable academic value and exploration.

Moreover, our curated dataset itself holds intrinsic value as an extensive representation of the Portuguese sociopolitical online landscape. Notably, it already incorporates two primary engagement metrics—likes and retweets—that, regrettably, remained beyond the scope of this dissertation's feasibility, but still serve as clear indicators of potential avenues for expanding our future analytical endeavors. Additionally, our methodology provides a solid foundation for future exploration. By modifying just the labeling process, the established framework can seamlessly pivot to classify alternative topics, yielding similar insightful outcomes. This adaptability offers a streamlined approach for treating diverse concepts as variables, mitigating the labor-intensive demands typically associated with such analyses. Thus, our work not only contributes to the current discourse within populist analysis but also lays the groundwork for subsequent investigations into a broader spectrum of subjects, within the Portuguese language and reality.

References

Amedie, J. 2015. "The Impact of Social Media on Society." *Advanced Writing: Pop Culture Intersections, 2*. https://scholarcommons.scu.edu/engl_176/2

Aouichaoui, A. R., R. Al, J. Abildskov & G. Sin 2021. "Comparison of Group-Contribution and Machine Learning-based Property Prediction Models with Uncertainty Quantification." *Computer Aided Chemical Engineering*, 50: 755-760. https://doi.org/10.1016/B978-0-323-88506-5.50118-2

Artstein, R. & M. Poesio 2008. "Inter-coder agreement for computational linguistics." *Computational Linguistics* 34(4): 555-596. https://doi.org/10.1162/coli.07-034-R2

Azevedo, P. J. L. 2020. *Dissecting Fact-Checking Systems: The Impact of Evidence Extraction Methods*. Master Dissertation, Faculdade de Engenharia, Universidade do Porto, https://repositorio-aberto.up.pt/bitstream/10216/129021/2/415636.pdf

Barros, J. P. & D. Espírito Santo 2019. *A Batalha das Legislativas nas Redes Sociais. O Amadorismo (ainda) é Quem Mais Ordena.* Rádio Renascença. Retrieved from https://rr.sapo.pt/extralarge/legislativas-2019/2019/10/04/a-batalha-das-legislativas-nas-redes-sociais-o-amadorismo-ainda-e-quem-mais-ordena/166953/

Bayer, M., W. Sommer & A. Schacht 2010. "Reading emotional words within sentences: the impact of arousal and valence on event-related potentials." *International Journal of Psychophysiology*, 78(3): 299-307. https://doi.org/10.1016/j.ijpsycho.2010.09.004

Benoit, K., K. Watanabe, H. Wang, P. Nulty, A. Obeng, S. Müller & A. Matsuo 2018. "quanteda: An R Package for the Quantitative Analysis of Textual Data." *Journal of Open Source Software*, 3(30): 774, https://doi.org/10.21105/joss.00774

Best, H. & J. Higley (Eds.) 2010. *Democratic Elitism: New Theoretical and Comparative Perspectives*. International Studies in Sociology and Social Anthropology, Vol. 111, BRILL, ISBN: 978-90-47-44174-8

Blumler, J. G. & D. Kavanagh 1999. "The Third Age of Political Communication: Influences and Features." *Political Communication*, 16(3): 209-230. https://doi.org/10.1080/105846099198596

Borges, L. & D. Rocha 2020. *Nova direcção do Livre – sem Joacine – eleita com 95 votos a favor e 15 brancos*. Público, January 19, 2020. Available at https://www.publico.pt/2020/01/19/politica/noticia/nova-direccao-livre-joacine-eleita-95-votos-favor-15-brancos-1900897

Braşoveanu, A. M. & R. Andonie 2020. Visualizing Transformers for NLP: A Brief Survey. In *2020 24th International Conference Information Visualisation (IV)*, Melbourne, Australia, September 7-11, 2020, pp. 270-279. IEEE. https://doi.org/10.1109/iv51561.2020.00051

Canovan, M. 2005. *The People*. Key Concepts Series, Wiley, ISBN: 978-0-745-62821-9

Carah, N. 2021. *Media and Society: Power, Platforms, and Participation*. (2nd ed.) SAGE Publications. https://www.perlego.com/book/3013385/media-and-society-power-platforms-and-participation-pdf.

Carbonell, J. G., R. S. Michalski & T. M. Mitchell 1983. "An Overview of Machine Learning." In R. S. Michalski, J. G. Carbonell & T. M. Mitchell (Eds.), *Machine Learning. Symbolic Computation*, Springer, Berlin, Heidelberg, pp. 3-23. https://doi.org/10.1007/978-3-662-12405-5

Castanho Silva, B., F. Vegetti & L. Littvay 2017. "The Elite Is Up to Something: Exploring the Relation Between Populism and Belief in Conspiracy Theories." *Swiss Political Science Review*, 23(4): 423-443, https://doi.org/10.1111/spsr.12270

Castells, M. 2013. *Communication Power*. (2nd ed.) Oxford University Press.

Chaudhary, J. & S. Niveditha 2021. "Twitter Sentiment Analysis Using Tweepy." *International Research Journal of Engineering and Technology,* 8(4): 4512-4516. https://www.irjet.net/archives/V8/i4/IRJET-V8I4864.pdf

Chicco, D., M. J. Warrens & G. Jurman 2021. "The Matthews Correlation Coefficient (MCC) is More Informative Than Cohen's Kappa and Brier Score in Binary Classification Assessment", *IEEE Access*, 9, 78368-78381.

Choudhury, N. 2014. "World Wide Web and Its Journey from Web 1.0 to Web 4.0." *International Journal of Computer Science and Information Technologie*s, 5(6): 8096-8100.

Chzhen, E., C. Denis, M. Hebiri, L. Oneto & M. Pontil 2019. Leveraging Labeled and Unlabeled Data for Consistent Fair Binary Classification. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox & R. Garnett (Eds.), *Advances in Neural Information Processing Systems, 32* (NeurIPS 2019), Vancouver, Canada.

Cohen J. 1960. "A Coefficient of Agreement for Nominal Scales." *Educational and Psychological Measurement*, 20(1): 37-46.

Cooksey B. 2014. *An Introduction to APIs*. Zapier, Inc., 23, 77.

Curry D. 2023. *Social App Report: Revenues, Usage, and Demographics on the Major Social Apps*. Business of Apps, Publisher.

Dahiya, N., S. Gupta & S. Singh 2022. "A Review Paper on Machine Learning Applications, Advantages, and Techniques". *ECS Transactions*, 107(1): 6137, https://doi.org/10.1149/10701.6137ecst

Dahl, R. A. 1978. "Pluralism Revisited." *Comparative Politics*, 10(2): 191-203.

Datanovia 2021. "Cohen's Kappa in R: For Two Categorical Variables." Retrieved May 11, 2023, from https://www.datanovia.com/en/lessons/cohens-kappa-in-r-for-two-categorical-variables/

Devlin, J., M. W. Chang, K. Lee & K. Toutanova 2018. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding." arXiv preprint arXiv:1810.04805

Diário de Notícias 2017. *Portugal é um país onde o populismo não dá votos*. Published by Diário de Notícias, January 1, 2017. https://www.dn.pt/portugal/portugal-e-um-pais-onde-o-populismo-nao-da-votos-5580231.html

Dowding, K. 2018. "Emotional Appeals in Politics and Deliberation." *Critical Review of International Social and Political Philosophy*, 21(2): 242-260, https://doi.org/10.1080/13698230.2016.1196536

Eisenstein, J. 2019. *Introduction to Natural Language Processing.* The MIT Press, 536 pp.

Ferguson, N. 2019. *In Our 'Emocracy,' Emotions Rule*. Published by Boston Globe, 28 January, 2019. https://www.bostonglobe.com/opinion/2019/01/28/opinion-niall-ferguson-our-emocracy-emotions-rule/RsgnMKr8zMC0ZSmWzss9FL/story.html

Ferreira da Silva, F., 2019. *Leaders over Parties?: The Personalization of Politics and Voting Behavior in Contemporary Western Democracies*. (Doctoral Dissertation), European University Institute. https://cadmus.eui.eu/handle/1814/63547

Filipe, S., F. Batista & R. Ribeiro 2020. "Different Lexicon-Based Approaches to Emotion Identification in Portuguese Tweets." *9th Symposium on Languages, Applications and Technologies (SLATE 2020)*, A. Simões, P. R. Henriques & R. Queirós (Eds.) Schloss Dagstuhl-Leibniz-Zentrum fur Informatik, Dagstuhl, Germany, Vol. 83, pp. 12:1-12:8. https://doi.org/10.4230/OASIcs.SLATE.2020.12

Flemmen, M. 2013. "Putting Bourdieu to work for class analysis: reflections on some recent contributions." The *British Journal of Sociology*, 64(2): 325-343. https://doi.org/10.1111/1468-4446.12020

Flew, T. 2007. *Understanding Global Media*. Palgrave Macmillan, 261 pp. http://books.google.com/books?id=g2RjAAAAMAAJ

Galston, W. A. 2017. *Anti-Pluralism: The Populist Threat to Liberal Democracy.* Yale University Press, New Haven & London.

Gane, N. 2005. "Max Weber as Social Theorist: 'Class, Status, Party'". *European Journal of Social Theory*, 8(2): 211-226. https://doi.org/10.1177/1368431005051764

Gartner, S. 2022. *Number of Smartphones Sold to End Users Worldwide from 2007 to 2021 (in million units)*. Statista 2022, available in: https://www.statista.com/statistics/263437/global-smartphone-sales-to-end-users-since-2007/

Gerbaudo, P. 2017. *The Mask and the Flag: Populism, Citizenism and Global Protest*. New York: Oxford University Press.

Géron, A. 2017. *Hands-On Machine Learning with Scikit-Learn and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems* (1st ed.) O'Reilly Media, Inc.

Gholamy, A., V. Kreinovich & O. Kosheleva 2018. "Why 70/30 or 80/20 Relation Between Training and Testing Sets: A Pedagogical Explanation." Retrieved May 12, 2023, from https://www.cs.utep.edu/vladik/2018/tr18-09.pdf

Gidron, N. & B. Bonikowski 2013. "Varieties of Populism: Literature Review and Research Agenda." Working Paper Series, Weatherhead Center for International Affairs, Harvard University, No.13-0004, http://dx.doi.org/10.2139/ssrn.2459387

Goldhaber, M. H. 1997. "The Attention Economy and the Net." *First Monday*, 2(4). https://doi.org/10.5210/fm.v2i4.519

González-Carvajal, S. & E. C. Garrido-Merchán 2020. "Comparing BERT against traditional machine learning text classification." arXiv preprint arXiv:2005.13012

Graves, A. 2013. "Generating sequences with recurrent neural networks." arXiv preprint arXiv:1308.0850

Grillo, R. D. 1998. *Pluralism and the Politics of Difference: State, Culture, and Ethnicity in Comparative Perspective*. Clarendon Press, pp. 1-27.

Grimmer, J., M. E. Roberts & B. M. Stewart 2022. *Text as Data: A New Framework for Machine Learning and the Social Sciences*. Princeton University Press. ISBN: 9780691207551

Gründl, J. 2022. “Populist ideas on social media: A dictionary-based measurement of populist communication.” *New Media & Society*, 24(6): 1481-1499. https://doi.org/10.1177/1461444820976970

Gunter, B., N. Koteyko & D. Atanasova 2014. “Sentiment Analysis: A Market-Relevant and Reliable Measure of Public Feeling?” *International Journal of Market Research*, 56(2): 231-247. https://doi.org/10.2501/IJMR-2014-014

Guriev, S. & E. Papaioannou 2022. "The Political Economy of Populism." *Journal of Economic Literature*, 60 (3): 753-832. https://doi.org/10.1257/jel.20201595

Haidt, J. 2022. *Why the Past 10 Years of American Life Have Been Uniquely Stupid*. The Atlantic, 11 April, 2022. https://www.theatlantic.com/magazine/archive/2022/05/social-media-democracy-trust-babel/629369/

Hawkins, K. A. & B. C. Silva 2018. *Textual Analysis: Big Data Approaches.* In: The Ideational Approach to Populism. (1st ed.). Routledge, pp. 27-48.

Hawkins, K. A. 2009. “Is Chávez Populist? Measuring Populist Discourse in Comparative Perspective.” *Comparative Political Studies*, 42(8): 1040-1067, https://doi.org/10.1177/0010414009331721

Hiebert, T. 2007. “The Tower of Babel and the Origin of the World's Cultures.” *Journal of Biblical Literature*, 126(1): 29-58. https://doi.org/10.2307/27638419

Honnibal, M. & Montani, I., 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing.

Hotz, N. 2023. *What is CRISP-DM?* [Web page]. Data Science Process Alliance. Retrieved March 14, 2023, from https://www.datascience-pm.com/crisp-dm-2/

Huszár, F., S. I. Ktena, C. O’Brien, L. Belli, A. Schlaikjer & M. Hardt 2022. Algorithmic Amplification of Politics on Twitter. *Proceedings of the National Academy of Sciences*, 119(1): e2025334119. https://doi.org/10.1073/pnas.2025334119

Hutchinson, A. 2021. *New Study Shows That 25% of Twitter Users Produce 97% of All Tweets*. Social Media Today. Retrieved May 12, 2023, from https://www.socialmediatoday.com/news/new-study-shows-that-25-of-twitter-users-produce-97-of-all-tweets/610146/

Ingruber, D. 2021. *Die Emokratie handelt nicht unbedingt vom Ende*. Published by austriandemocracylab. https://www.austriandemocracylab.at/die-emokratie-handelt-nicht-unbedingt-vom-ende/

Jacquinet, M. 2021. “Breve Nota de Introdução à Metodologia de Investigação”, pp. 1-9, https://repositorioaberto.uab.pt/handle/10400.2/10677

Jain, A. K., J. Mao & K. M. Mohiuddin 1996. “Artificial Neural Tetworks: A Tutorial.” *Computer*, 31-44. https://csc.lsu.edu/~jianhua/nn.pdf

Janda, K. & T. Colman 1998. “Effects of Party Organization on Performance during the 'Golden Age' of Parties.” *Political Studies,* 46(3): 611-632. https://doi.org/10.1111/1467-9248.00157

Kaltwasser, C. R., P.A. Taggart, P. O. Espejo & P. Ostiguy (Eds.) 2017. *The Oxford Handbook of Populism*. Oxford University Press, ISBN: 9780198803560

Karišik, E. 2018. *A Standardized Data Mining Method in Healthcare: A Pediatric Intensive Care Unit Case Study*. (Master Thesis) Business Informatics, Utrecht University, Information & Computing Science https://studenttheses.uu.nl/bitstream/handle/20.500.12932/39488/Final%20Thesis.pdf?sequence=2

Kierszbaum, S. 2020. “Masking in Transformers' self-attention mechanism.” *Analytics Vidhya*. Retrieved June 7, 2023, from https://medium.com/analytics-vidhya/masking-in-transformers-self-attention-mechanism-bad3c9ec235c

Kietzmann, J. H., K. Hermkens, I. P. McCarthy & B. S. Silvestre 2011. “Social media? Get serious! Understanding the functional building blocks of social media.” *Business Horizons*, 54(3): 241-251. https://doi.org/10.1016/j.bushor.2011.01.005

Kröse, B. & P. van der Smagt 1993. *An Introduction to Neural Networks*. The University of Amsterdam, Amsterdam, The Netherlands.

Kubler, K. 2023. “Influencers and the Attention Economy: the Meaning and Management of Attention on Instagram.” *Journal of Marketing Management*, 1-17, https://doi.org/10.1080/0267257X.2022.2157864

Kumar, A. 2022. “Difference: Binary, Multiclass & Multi-Label Classification.” Retrieved May 9, 2023, from https://vitalflux.com/difference-binary-multiclass-multi-label-classification/

Lawrence, B. S. & N. P. Shah 2020. “Homophily: Measures and Meaning.” *Academy of Management Annals*, 14(2): 513-597. https://doi.org/10.5465/annals.2018.0147

Lazarsfeld, P. F. & R. K. Merton 1954. *Friendship as a Social Process: A Substantive and Methodological Analysis*. In: M. Berger, T. Abel & H. Charles (Eds.), Freedom and Control in Modern Society, Van Nostrand, New York, pp. 18-66.

Liu, P., X. Qiu, X. Chen, S. Wu & X. Huang 2015. Multi-Timescale Long Short-Term Memory Neural Network for Modelling Sentences and Documents. *In Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, Lisbon, Portugal. Association for Computational Linguistics, pp. 2326–2335. https://doi.org/10.18653/v1/D15-1280

Logghe, H. J., Boeck, M. A. & Atallah, S. B. 2016. “Decoding Twitter: Understanding the History, Instruments, and Techniques for Success.” *Annals of Surgery*, 264(6): 904-908. https://doi.org/10.1097/SLA.0000000000001824

Lopes, M. J. & J. G. Henriques 2016. *A Extrema-direita em Portugal Quer Dar Nas Vistas, Mas é “Irrelevante”*. Público, 19 de novembro. Retrieved from https://www.publico.pt/2016/11/19/politica/noticia/a-extrema-direita-em-portugal-quer-dar-nas-vistas-mas-%C3%A9-irrelevante-1751557

Marneffe, M. C. de, & J. Nivre 2019. “Dependency Grammar.” *Annual Review of Linguistics*, 5: 197-218. https://doi.org/10.1146/annurev-linguistics-011718-011842

McCarthy, J. 2007. “What is artificial intelligence?” Computer Science Department Stanford University, retrieved January 7, 2023, from https://www-formal.stanford.edu/jmc/whatisai.pdf

McChesney, R. W. 1995*. Telecommunications, Mass media, and Democracy: The Battle for the Control of US Broadcasting, 1928-1935*. Oxford University Press. ISBN: 9780195093940

McChesney, R. W. 2000. “The political economy of communication and the future of the field.” *Media, Culture & Society*, 22(1): 109-116. https://doi.org/10.1177/016344300022001006

McChesney, R. W. 2013. *Digital Disconnect: How Capitalism is Turning the Internet Against Democracy*. New York, NY: New Press, 299 pp.

McLuhan, M. 1969. *Counterblast*. Harcourt, Brace & World Publishers, New York.

McNair, B. 2017. *An Introduction to Political Communication*. Routledge.

Mechant, P. 2012. "An illustrated framework for the analysis of Web2.0 interactivity." *Contemporary Social Science*, 7(3): 263-281. https://doi.org/10.1080/21582041.2012.716524

Mendes, M. S. 2021. "'Enough'of What? An Analysis of Chega's Populist Radical Right Agenda." *South European Society and Politics*, 26(3): 329-353. https://doi.org/10.1080/13608746.2022.2043073

Metzger, M. J. & A. J. Flanagin 2013. "Credibility and trust of information in online environments: The use of cognitive heuristics." *Journal of Pragmatics*, 59, 210-220. https://doi.org/10.1016/j.pragma.2013.07.012

Mikolov, T., W.T. Yih & G. Zweig 2013. Linguistic Regularities in Continuous Space Word Representations. In *Proceedings of NAACL-HLT 2013*, Atlanta, Georgia, 9-14 June 2013, pp. 746-751, North American Chapter of the Association for Computational Linguistics: Human Language Technologies.

Moffitt, B. & S. Tormey 2014. "Rethinking Populism: Politics, Mediatisation and Political Style." *Political Studies*, 62(2): 381-397, https://doi.org/10.1111/1467-9248.12032

Mohammed, R., J. Rawashdeh & M.A. Abdullah 2020. Machine Learning with Oversampling and Undersampling Techniques: Overview Study and Experimental Results. In *11th International Conference on Information and Communication Systems (ICICS)*, pp. 243-248.

Moura, R., R. Sousa-Silva & H. Lopes-Cardoso 2021. "Automated Fake News Detection Using Computational Forensic Linguistics." In *Progress in Artificial Intelligence: 20th EPIA Conference on Artificial Intelligence,* EPIA 2021, Virtual Event, 7–9 September 2021, pp. 788-800, Springer International Publishing.

Mudde, C. & C. R. Kaltwasser 2017. "*Populism: A Very Short Introduction*." Oxford University Press, ISBN: 9780190234874

Mudde, C. 2004. "The Populist Zeitgeist." *Government and Opposition*, 39(4): 541-563, http://dx.doi.org/10.1111/j.1477-7053.2004.00135.x

Neuendorf, K. A. 2017. *The Content Analysis Guidebook*. 2nd Ed., *SAGE Publications, Inc*., 456 pp.

Nicholas, D. & D. Clark 2013. "The second digital transition: to the mobile space—an analysis of Europeana." *Learned Publishing*, 26(4): 240-252. https://doi.org/10.1087/20130402

Nivre, J. & J. Nilsson 2005. Pseudo-Projective Dependency Parsing. In K. Knight, H. T. Ng, K. Oflazer (Eds.). In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL)*, Ann Arbor, Michigan, USA, June 25-30, 2005, pp. 99-106. Association for Computational Linguistics. http://dx.doi.org/10.3115/1219840.1219853

Nivre, J. 2005. "Dependency Grammar and Dependency Parsing." *MSI Report, 5133*(1959), pp. 1-32.

Norris, P. & R. Inglehart 2019. *Cultural Backlash: Trump, Brexit, and Authoritarian Populism*. Cambridge: Cambridge University Press, http://dx.doi.org/10.1017/9781108595841

O'reilly, T. 2005. *Web 2.0: compact definition*. Oreilly Radar 1 October 2025. http://radar.oreilly.com/archives/2005/10/web-20-compact-definition.html

Oliveira, M. D. L. C. 2022. Análise da Propagação e Impacto da Informação nas Redes Sociais. (Dissertação de Mestrado) Instituto Universitário de Lisboa, Repositório ISCTE, http://hdl.handle.net/10071/26486 (in Portuguese).

Oliver, J. E. & T. J. Wood 2014. “Conspiracy Theories and the Paranoid Style(s) of Mass Opinion.” *American Journal of Political Science*, 58(4): 952-966. https://doi.org/10.1111/ajps.12084

Ostiguy, P. 2017. *A socio-cultural approach*. The Oxford Handbook of Populism, pp. 73-97.

Oxford Advanced Learner's Dictionary 2022. OxfordLearnersDictionaries.com. Retrieved 1 August 2022, from https://www.oxfordlearnersdictionaries.com/definition/english/media?q=media

Parrança, A. I. M. 2022. *O Populismo de André Ventura: Uma Análise Discursiva da Campanha para as Eleições Presidenciais de 2021.* Master Dissertation, Instituto Superior de Ciências Sociais e Políticas, Universidade de Lisboa. http://hdl.handle.net/10400.5/24469

Passonneau, R., N.Y. Habash & O.C. Rambow 2006. Inter-Annotator Agreement on a Multilingual Semantic Annotation Task. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation* (LREC’06), Genoa, Italy, European Language Resources Association (ELRA). Retrieved February 9, 2023, from http://www.lrec-conf.org/proceedings/lrec2006/pdf/634_pdf.pdf

Pauwels, T. 2011. “Measuring Populism: A Quantitative Text Analysis of Party Literature in Belgium.” *Journal of Elections, Public Opinion and Parties*, 21(1): 97-119. https://doi.org/10.1080/17457289.2011.539483

Plotke, D. 1997. “Representation is Democracy.” *Constellations*, 4(1): 19-34.

Plummer, A. 2020. “The Beginner’s Guide to Logistic Regression. Towards Data Science.” Retrieved June 11, 2023, from https://towardsdatascience.com/the-beginners-guide-to-logistic-regression-part-1-920eda7aea5f#:~:text=There%20are%20a%20few%20reasons,introduction%20to%20binary%20response%20models

Quivy, R. & L. Van Campenhoudt 1992. “Manual de Investigação em Ciências Sociais.” Pp. 1-34, https://www.fep.up.pt/docentes/joao/material/manualinvestig.pdf

Reimers, N. & I. Gurevych 2019. “Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks.” arXiv preprint arXiv:1908.10084

Rodrigues, J., L. Gomes, J. Silva, A. Branco, R. Santos, H. L. Cardoso & T. Osório 2023. “Advancing Neural Encoding of Portuguese with Transformer Albertina PT-”. pp. 1-10, arXiv preprint arXiv:2305.06721

Roesslein, J., 2009. tweepy Documentation. Online] http://tweepy. readthedocs. io/en/v3, 5, p.724.

Rooduijn, M. & T. Pauwels 2011. “Measuring Populism: Comparing Two Methods of Content Analysis.” *West European Politics*, 34(6): 1272-1283, https://doi.org/10.1080/01402382.2011.616665

Rooduijn, M. 2014. “The Nucleus of Populism: In Search of the Lowest Common Denominator.” *Government and Opposition,* 49(4): 573-599, https://doi.org/10.1017/gov.2013.30

Rousseau, J. J. 1762. Social Contract (p. 134). Oxford.

Salgado, S. 2019. ‘Where’s populism? Online media and the diffusion of populist discourses and styles in Portugal.” *European Political Science*, 18(1): 53–65, https://doi.org/10.1057/s41304-017-0137-4

Sapage, S. 2019. *Portugal, o País sem Extrema-direita*. Público, 26 May. Retrieved from https://www.publico.pt/2019/05/26/politica/noticia/portugal-pais-extremadireita-1874203

Serra, J. P. & G. Gonçalves (Eds.) 2016. *Politics and Web 2.0: The Participation Gap.* Vernon Press.

Severyn, A. & A. Moschitti 2015. Twitter sentiment analysis with deep convolutional neural networks. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in information*, pp. 959-962. https://doi.org/10.1145/2766462.2767830

Shapiro, B. 2018. *Facts Don't Care About Your Feelings: Speech Is Not Violence*. Uploaded by Young Americans Foundation. https://www.youtube.com/watch?v=bGG8i1oVJhw&t=23s

Sharma, A., A. Jain, A. K. Arya & M. Ram (Eds.) 2022. *Artificial Intelligence for Signal Processing and Wireless Communication*. Published by Walter de Gruyter GmbH & Co KG, Berlin/Boston.

Sheng, V. S., F. Provost, F. & P.G. Ipeirotis 2008. Get Another Label? Improving Data Quality and Data Mining Using Multiple, Noisy Labelers. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '08).* Association for Computing Machinery, New York, USA, pp. 614-622, https://doi.org/10.1145/1401890.1401965

Shirky, C. 2010. *Cognitive Surplus: Creativity and Generosity in a Connected Age* (1st Edition). Penguin Books, UK.

Souza, F., R. Nogueira & R. Lotufo 2020. BERTimbau: Pretrained BERT Models for Brazilian Portuguese. In *Intelligent Systems: 9th Brazilian Conference*, BRACIS 2020, Rio Grande, Brazil, October 20-23, 2020, Proceedings, Part I, pp. 403-417. https://doi.org/10.1007/978-3-030-61377-8_28

Stieglitz, S. & L. Dang-Xuan 2013. "Emotions and Information Diffusion in Social Media—Sentiment of Microblogs and Sharing Behavior." *Journal of Management Information Systems*, 29(4): 217-248. https://doi.org/10.2753/MIS0742-1222290408.

Taggart, P. A. 1996. The New Populism and the New Politics. In *The New Populism and the New Politics*: New *Protest Parties in Sweden in a Comparative Perspective*. Palgrave Macmillan, London, pp. 11-46. https://doi.org/10.1007/978-1-349-13920-0_2

Tardieu, M. 2008. *Manichaeism*. University of Illinois Press.

Tasente, T. 2020. "The 4 Phases of Evolution of Political Communication Systems: from the Golden Age of the Parties to the Golden Age of the Users." *Technium Social Sciences Journal*, 2(1), 76–83. https://doi.org/10.47577/tssj.v2i1.50

Taylor, A., Marcus, M. and Santorini, B., 2003. The Penn treebank: an overview. *Treebanks: Building and using parsed corpora*, pp.5-22.

Thompson, J. B. 1995. *The Media and Modernity: A Social Theory of the Media*. Stanford University Press. http://www.sup.org/books/title/?id=2348

Tunstall, L., L. Von Werra & T. Wolf 2022. Natural language processing with transformers.". Publisher(s): O'Reilly Media, Inc. ISBN: 9781098136796

TV History. 2014 "Number of TV Households in America: 1950-1978." Accessed November 15, 2014. http://www.tvhistory.tv/Annual_TV_Households_50-78.JPG , "Number of TV Households in America 1950-1978," The American Century, accessed July 26, 2022, https://omeka.wlu.edu/americancentury/items/show/136. file:///C:/Users/abr/Downloads/number-of-tv-households-in-america-the.pdf

Twitter 2012. "Changes Coming in Version 1.1 of the Twitter API", Retrieved January 9, 2023, from https://blog.twitter.com/developer/en_us/a/2012/changes-coming-to-twitter-api

Ulinskaitė J. & L. Pukelis 2021. "Identifying Populist Paragraphs in Text: A Machine-Learning Approach." arXiv preprint arXiv:2106.03161

Vajpayee, S. 2020. "Understanding BERT — (Bidirectional Encoder Representations from Transformers)." *Towards Data Science*. Retrieved June 7, 2023 from https://towardsdatascience.com/understanding-bert-bidirectional-encoder-representations-from-transformers-45ee6cd51eef

Vasiliev, Y. 2020*. Natural Language Processing with Python and spaCy. A Practical Introduction.* No Starch Press, ISBN-13: 9781718500525

Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L. Kaiser & I. Polosukhin 2017. "Attention is all you need." *31st Conference on Neural Information Processing Systems* (NIPS 2017), Long Beach, CA, USA, December 4 - 9, 2017. Advances in Neural Information Processing Systems, 30. https://doi.org/10.48550/arXiv.1706.03762

Wang, H., C. Ma & L. Zhou 2009. "A Brief Review of Machine Learning and Its Application." In *2009 International Conference on Information Engineering and Computer Science,* Wuhan, China, December 19-20, 2009, pp. 1-4, IEEE. https://doi.org/10.1109/ICIECS.2009.5362936

Weyland, K. 2001. "Clarifying a Contested Concept: Populism in the Study of Latin American Politics." *Comparative Politics*, 34(1): 1-22. https://doi.org/10.2307/422412

Wolff, R. 2020. "5 Types of Classification Algorithms in Machine Learning." Retrieved June 9, 2023, from https://monkeylearn.com/blog/classification-algorithms/

Yu, M., Q. Huang, H. Qin, C. Scheele & C. Yang 2019. "Deep learning for real-time social media text classification for situation awareness–using Hurricanes Sandy, Harvey, and Irma as case studies." *International Journal of Digital Earth*, 12(11): 1230-1247. https://doi.org/10.1080/17538947.2019.1574316

Zhao, W. X., J. Jiang, J. Weng, J. He, E.P. Lim, H. Yan & X. Li 2011. Comparing Twitter and Traditional Media Using Topic Models. In *Advances in Information Retrieval*: *33rd European Conference on IR Research, ECIR 2011,* Dublin, Ireland, April 18-21, 2011. Proceedings 33, pp. 338-349, Springer, Berlin, Heidelberg, https://doi.org/10.1007/978-3-642-20161-5_34

Zhou, L., S. Pan, J. Wang & A.V. Vasilakos 2017. "Machine Learning on Big Data: Opportunities and Challenges", *Neurocomputing*, 237: 350-361, https://doi.org/10.1016/j.neucom.2017.01.026

Žižek, S. 2006. 'Against the Populist Temptation.' *Critical Inquiry* 32(3): 551–574, https://doi.org/10.1086/505378

Zúquete, J. P. 2022. *Populismo, Lá Fora e Cá Dentro*. Fundação Francisco Manuel dos Santos. ISBN 978-989-9064-69-0

Annexes

Figure A.1 Heatmap matrix of individual tweets by political actors and year
(left column represents Non-populist tweets and right column represents populist ones)

Table A.1 Frequency of Populist and Non-Populist Activities by political actors (Machine and Manual Classification)

Name	Non-Populist			Populist		
	Machine	Manual	Ratio	Machine	Manual	Ratio
André Ventura	1 918	51	3%	890	160	18%
António Costa	2 943	243	8%	50	11	22%
Bruno Dias	443	55	12%	78	23	29%
Catarina Martins	2 587	132	5%	439	85	19%
Cotrim Figueiredo	258	78	30%	103	93	90%
Isabel Sousa Real	1 578	206	13%	286	49	17%
Joacine Katar Moreira	2 276	128	6%	503	54	11%
João Oliveira	119	43	36%	4	10	250%
Luís Montenegro	133	27	20%	27	16	59%
Rui Rio	660	134	20%	139	98	71%
Rui Rocha	102	5	5%	28	2	7%
Rui Tavares	1 427	116	8%	181	29	16%

Timeline of events during the last two election cycles:

October 2019

- António Costa was confronted as not being present during the forest fires in Pedrogão
- Legislative Elections, PS win with Geringonça
- Racist petition to send Joacine “back to her country” with 10k signings
- Climate strikes
- Chega emerges (with 1% of the votes) every politician loses its mind
- Revolt in Catalonia after arresting independence leaders

November 2019

- Web Summit in Portugal
- Baby found in the trash
- Release of Lula from prison
- Joacine shocks Livre dinner party “they don't support me”

December 2019

- Greta Thunberg in Portugal
- Controversy surrounding "Porta dos Fundos" (Christian controversy)
- Judge Rui Rangel convicted of bribery and favors
- Corruption in Crédito Agrícola bank (woman received salary without performing any duties)

January 2020

- USA attacks Iran
- Joacine "é mentira" moment
- Outbreak of Coronavirus (anti-Chinese sentiment)
- Rui Pinto leaks information about Isabel dos Santos
- Susana Garcia and the police beating of a "black" person on a bus

February 2020

- The malnourished dogs of bullfighter João Moura (animal abuse)
- Approval of the decriminalization of euthanasia
- Racism case involving Marega (Portuguese football player)
- Government freezes Isabel dos Santos' bank accounts and assets

March 2020

- Coronavirus pandemic
- Panic buying of essential goods like toilet paper
- Dutch Finance Minister criticizes Southern countries
- Boris Johnson's parties during lockdown
- SEF workers murder a Ukrainian immigrant

April 2020

- Easter and COVID-19
- Bolsonaro and "Bolsonarism" supporters
- Protests against the prohibition of activities due to COVID-19 (ex. hair salons)
- Murder of a 21-year-old boy by two girls in algarve

May 2020

- Murder of Valentina by her father (CMTV coverage at the funeral)
- Death of George Floyd
- Black Lives Matter protests in America
- André Ventura's fight against Ricardo Quaresma (specific lockdown for the Romani community)
- Murder of Beatriz Lebre by a psychology student at ISCTE ~~University~~

June 2020

- Denialist protest against racism in Portugal
- Beach season in Portugal
- Cancellation of Santos Populares
- Vandalism of statues because of colonialism
- Phased easing of lockdown measures

July 2020

- FC Porto becomes national champion (celebrations during COVID-19)
- Easing of lockdown with bars and clubs closing at 10 p.m.
- Bruno Candé's murder by an elderly man with a hunting rifle - Bruno Caetano defends the killer
- Arson attack on a kennel (for insurance money)

- Fires in Portugal and firefighters' efforts
- Novo Banco sells assets with loans to a fund in the Cayman Islands (Operation Viriato)

August 2020

- UEFA Champions League played in Portugal (a tribute to healthcare workers)
- Lukashenko wins elections in Belarus, leading to protests
- 40% drop in the Portuguese economy compared to 2019 (due to COVID-19) reported by INE (National Institute of Statistics)

September 2020

- CP (train company) conductor sexually harasses a young woman
- Festa do Avante takes place despite covid restrictions (political festival)
- António Costa (Portuguese Prime Minister) listed as an honorary member of Luís Filipe Vieira's (Benfica president) organization
- Six young people write a letter to the UN regarding the future of the environment
- Presidential election advertisements

October 2020

- Trump tests positive for COVID-19
- Denialist protests against COVID-19 measures
- Launch of the StayAway COVID app (contact tracing) and controversy
- Associate professor murdered over a caricature of the Prophet Muhammad
- Gathering in Mafra (COVID-19 violation)

November 2020

- American elections (Trump denies defeat)
- Portugal returns to a state of emergency

December 2020

- Hungarian MP involved in a gay orgy
- "Soft" measures to combat COVID-19 during Christmas (relaxed restrictions)
- Hunger strike "A pão e água" (Ljubomir Stanisic)
- Ana Gomes receives vaccine from France (skips the line)
- Panic button installed at SEF (serviço de estrangeiros e fronteiras)
- Creation of the vaccination plan

January 2021

- Surge in COVID-19 cases due to increased freedom during the Christmas period
- Second lockdown implemented
- Capitol Hill invasion in the United States
- André Ventura calls a boy from the Jamaica neighborhood a "bandido"
- Presidential debates, and Marcelo Rebelo de Sousa reelected with a majority

February 2021

- Mamadou Ba and SOS Racism
- General Gouveia e Melo takes over the vaccination plan
- Trump banned from social media platforms for inciting anti-democratic behaviour

- Boom of cryptocurrencies and NFTs

March 2021
- Presentation of the COVID-19 situation (Quadro das cores)
- Container ship stuck in the Suez Canal
- Nuno Graciano runs for mayor of Lisbon
- Operation Marquês (José Sócrates cleared of charges due to prescription)
- European Super League in football

April 2021
- Civil requisition of ZMAR establishments
- Inhumane conditions in Odemira
- Conflict between Palestine and Israel
- Lukashenko intercepts a plane in Poland to arrest a journalist
- Rui Pinto barred from testifying in inquiry committees
- Gatherings during Sporting's celebrations

May 2021
- Palestinian journalist killed in Israel
- Texas shooter incident
- Stabbing during Porto's celebrations
- Ricardo Salgado sentenced to 16 years in prison
- Activists vandalize a Van Gogh painting with orange paint

June 2021
- Graça Freitas says "don't eat bacalhau abrás"
- Beating of a 3-year-old child over 5 days due to the mother's debts (Jéssica)
- Second United Nations Ocean Conference
- New airport plan proposed by Pedro Nuno Santos
- Polemic around the appointment of Pedro Adão e Silva as the organizer of the 50th year celebration of 25th of April
- Minister Eduardo Cabrita runs over civilian

July 2021
- Costa revokes Pedro Nuno Santos' decision
- Montenegro becomes the new President of PSD (Social Democratic Party)
- Heatwave in Portugal
- Santos Silva condemns André Ventura during the parliamentary debate
- Artists attending Avante Festival attacked as "Putin supporters"
- Articles about transgender people in newspapers
- Couple refuses to allow their children to receive sex education

August 2021
- Marta Temido resigns (Minister of Health)
- Official end of the pandemic
- Finnish Prime Minister attends a party
- Mário Ferreira becomes the first Portuguese tourist in space

- Mário Ferreira accused of money laundering
- Sérgio Figueiredo and Fernando Medina involved in a favor exchange (direct award of contracts)
- Pichardo wins gold medal in the Olympics

September 2021

- Funeral of Jorge Sampaio (former President of Portugal)
- Municipal elections
- Paulo Rangel comes out as a homosexual
- Carlos Moedas wins the Lisbon mayoral race

October 2021

- Yearly budget proposal rejected, leading to a political crisis
- Internal elections in PSD
- João Rendeiro sentenced to 10 years in prison and escapes the country
- Inquiry committees on José Sócrates and Joe Berardo
- Pandora Papers leak
- Aristides Sousa Mendes honored In the national pantheon

November 2021

- CNN launches in Portugal
- Mário Machado arrested (far-right leader)
- Dissolution of the Assembly of the Republic and elections
- Omicron variant of COVID-19 emerges
- Euthanasia rejected by President Marcelo Rebelo de Sousa

December 2021

- Portugal returns to a state of calamity
- João Rendeiro arrested in South Africa
- Marcelo Rebelo de Sousa undergoes hernia surgery
- Resignation of Eduardo Cabrita (Minister of Internal Affairs)
- New Zealand bans tobacco sales to those born after 2006

January 2022

- Legislative elections
- PS (Socialist Party) wins an absolute majority
- Chega becomes the third political force, while CDS disappears from parliament
- Incorrect poll predictions

February 2022

- Terrorist attack attempt in university
- The start of the Ukrainian war
- UE sanctions Russia over the war
- José Milhazes (journalist specializing in Russian events) comments
- Zielinski emerges as a worldwide figure
- PCP does not condone the Russian actions
- Salgado attributed amnesia

March 2022

- Mitá Ribeiro is vetoed from being the president of the Assembly of the Republic.
- Medina is the new Minister of Finance.
- Will Smith slaps Chris Rock during the Oscars.
- Fábio Guerra (Police officer) killed by a navy oficer in Lisbon's nightlife.
- Ricardo Salgado was sentenced to 6 years in prison.

April 2022

- April 25th - Chega (political party) leaves the parliament during Grandula (traditional song)
- PCP (Portuguese Communist Party) walks out of Parliament when Zelensky speaks
- Human rights in the construction of the Qatar World Cup
- Harassment at the Faculty of Law
- Elon Musk purchases Twitter
- SONAE CEO receives a salary increase supported by the state

May 2022

- Palestinian journalist assassinated in Israel
- Texas shooter incident
- Stabbing during Porto's celebrations
- Ricardo Salgado sentenced to 16 years in prison
- Environmental activists attack a Van Gogh painting with orange paint

June 2022

- Graça Freitas says "don't eat cod"
- Beating of a 3-year-old child over 5 days due to the mother's debts (Jéssica)
- Second United Nations Ocean Conference
- New airport plan proposed by Pedro Nuno Santos

July 2022

- Costa revokes Pedro Nuno Santos' decision
- Montenegro becomes the new president of PSD (Social Democratic Party)
- Heatwave in Portugal
- Santos Silva condemns André Ventura during the parliamentary debate
- Artists attending Avante Festival attacked as "Putin supporters"
- Articles about transgender individuals in newspapers
- Couple refuses to allow their children to receive sex education
- Highest inflation rate in recent years

August 2022

- Marta Temido resigns (Minister of Health)
- Official end of the pandemic
- Finnish Prime Minister attends a party
- Mário Ferreira becomes the first Portuguese tourist in space
- Mário Ferreira accused of money laundering
- Sergio Figueiredo and Fernando Medina involved in a favor exchange (direct award of contracts)

- Pichardo wins gold medal

September 2022
- Death of Queen Elizabeth II of the UK
- Anti-inflation package for energy (€125 from Costa)
- Arrest of anti-monarchy protesters
- Putin makes nuclear threats
- Minister in charge of European funds grants funds to a company created by her husband just 15 days earlier
- Mahsa Amini and Iranian women protest (cutting hair and burning hijab)

October 2022
- 17 reports of sexual abuse in the Catholic Church
- Church devalued and covered up sexual abuse crimes
- Marcelo (President) informed the Bishop's leader before the investigation into the Catholic Church became public
- Marcelo downplays the number of 400 sexual abuse victims
- Elections in Brazil, Bolsonaro loses (pro-Bolsonaro demonstrations)

November 2022
- New assistant to the Minister of the Presidency receives €4 000
- Paulo Raimundo becomes the new Secretary of the PCP (Portuguese Communist Party)
- Caminha Municipality pays €300 000 for a non-existent project to a company that has existed for 8 years but has no employees
- Minister meets activists from the Faculty of Letters (who had no proposals)
- FIFA President's "Today I Feel" speech

December 2022
- Kanye West's Hitler rant
- Floods in Lisbon
- News report: "Portuguese youth are the ones who stay with their parents for the longest time."
- New support of €240 for families

January 2023
- Scandals and resignations of Minister of Infrastructure and Housing (Pedro Nuno Santos), Secretary of State for Agriculture (Carla Alves), and Secretary of State for Treasury (Alexandra Reis)
- Floods in Porto
- Rui Rocha elected president of Iniciativa Liberal
- Elections in the Order of Physicians
- Prime Minister advisor (Pedro Ribeiro) resigns over court case

February 2023
- Housing package announced by Marina Gonçalves
- End of Golden Visa program
- Earthquake in Turkey and Syria

- European Commission supports joint armament
- Zelensky and Xi Jinping open talks about peace

March 2023

- Attack on the Ismaili Center in Lisbon
- Military refuses to carry out a mission to accompany a Russian ship north of Madeira
- Collapse of Credit Suisse bank in Switzerland
- Lisbon airport ranked as one of the worst by the Global airport Ranking
- “IVA zero” measure and its effects