University of Padova

Department of Mathematics

Master Thesis in Data Science

Talk like a populist, think like a populist?

A text-based evaluation of populist discourse on social media

Supervisor
Professor Tomaso Erseghe
University of Padova

Master Candidate
Elisa Tremolada

Academic Year

2022-2023

To all those who struggle to know how to be here,
in a world that asks us to know where we want to be next.

Abstract

Few concepts in social science are as elusive as that of populism: is it an ideology, a rhetorical style, or just a winning strategy for elections? Due to its enigmatic nature, populism has been assigned as many definitions as the forms it can take in the real world. This has been a long-standing challenge for researchers: how can we aim to study populism on a cross-cultural, multinational scale, if there is no consensus about what it actually is? The last decade has opened up new research possibilities, particularly with respect to quantitative approaches to the social sciences: the widespread use of social media, particularly as a means of establishing and reinforcing relations between political leaders and their electorate, has meant a previously unimaginable amount of text data is now available for analysis. Moreover, the surge in the development of machine learning techniques – in particular, Natural Language Processing methods – has given experts in the field a chance to inspect huge amounts data with a precision and speed which were never possible before.

This dissertation takes on an interdisciplinary journey through political science, discourse analysis, psychology of language and text analysis techniques. State-of-the-art NLP methods are combined with strategies coming from different fields in order to extract features pertaining to populist communication logic from 3+million Tweets posted by politicians of 23 countries in the last 10 years. Using Laclau's [1] discursive-performative framework as a theoretical starting point, four dimensions of populism are associated to indexes made up of variables extracted purely from text. These dimensions are then compared with existing survey data about populism [2], in order to evaluate their usability and robustness. Far from being able to reach its natural conclusion, this dissertation has the hope to inspire more researchers to move freely within and between fields in order to develop new methodological frameworks that better suit the modern challenges the social sciences are facing.

Contents

Listing of figures

Listing of tables

Listing of acronyms

NLP . . . . Natural Language Processing

ML . . . . Machine Learning

NMT . . . . Neural Machine Translation

EDA . . . . Exploratory Data Analysis

SA . . . . Sentiment Analysis

BERT . . . . Bidirectional Econded Representations from Transformers

MFD . . . . Moral Foundations Dictionary

NMF . . . . Non-Negative Matrix Factorization

UMAP . . . . Uniform Manifold Approximation and Projection for Dimension Reduction

GPS . . . . Global Party Survey

LIWC . . . . Linguistic Inquiry and Word Count

0
Introduction

There is virtually no field of science that has not benefited, in the last 30 years, from what we call the fourth industrial revolution, i.e. *the fusion of technologies that is blurring the lines between the physical, digital, and biological spheres* [3].

The social sciences in particular have been – and perhaps are still – living their own revolution: spoken, fleeting word has been substituted, at least partially, by easily traceable and analyzable tweets, posts, stories, likes and views. Perhaps for the first time in human history, researchers get first-hand access to records of human behaviour at a previously unthinkable scale and, more surprisingly, at no cost - if not that of learning the necessary skills to collect them.

The concept of populism, in particular, has sparked interest in academics and journalists alike all along the digital revolution: it was nominated Word of the Year by the Cambridge English Dictionary in 2017, the New York Times used the term 2537 times that same year [4] and the number of books employing the expression has more than doubled since the 2000s [5]. Simply put, *Populism is sexy*. [4].

Research on populism has defined it as an ideology, a political strategy, a folkloric style of politics, a tool for popular mobilization and even an emancipatory force [6]. In the last decade or so, however, a stronger consensus has emerged in the scientific community over a definition of populism as a set of ideas that concerns the antagonistic relationship between the corrupt elite and the pure people, the latter constituting the core of populist (thin-centered) ideology [4]. Populist actors typically promise to act as representatives, advocates, and mouthpieces of the people. As such, they seek a fast, direct and unmediated connection with the electorate, one that does not entail the interference of mainstream media and democratic checks and balances alike [7]. Moreover, the Internet and populism have been regarded both as correctives of democracy and as one of the causes of the deterioration of democratic life in Western societies [6][7][1].

Therefore, which environment could be more well-suited to host populist actors than the Web?

As the academic agreement around a conceptualization solidified, researchers shifted their focus from defining to measuring populism, mostly by means of quantitative content analysis methods [4]. However, the countless empirical studies produced in recent years tend to share two common fallacies: the first is the lack of a strict theoretical framework to support the research design and data collection, which can often result in these papers being more relevant to the study of similar or overlapping phenomena (e.g. the radical right) rather than to the study of populism itself; the second is the failure to recognize that populism, by its own enigmatic nature, is "*one of the privileged places of emergence of [...] contamination*" between different literatures. As such, it also the perfect ground to allow fruitful fertilization across apparently disjointed scientific fields to take place [1][4].

This thesis has the ambition of overcoming both of these challenges.

Chapter 1 constitutes an attempt at linking and reviewing different subfields related to populism, from communication theory, to discourse analysis, to populism studies themselves, with the aim of laying a solid theoretical foundation for the rest of the dissertation. Section 1.1 will present the idea that institutional contexts, such as the representative democracies where populism has emerged, "*are enacted through media technologies and perspectives*" [8]. Moreover, this section discusses how a populist communication logic – the content of a populist message, its form and the aims and motives behind its diffusion – perfectly fits social media as a distributed, non-hierarchical and democratic media channel, such that studying a contemporary political phenomenon necessarily entails approaching studies of online, user-generated communication [9]. Section 1.2 will construct a rigorous theoretical framework for defining populism from a discursive-performative perspective and motivate this choice of approach. In turn, Section 1.3 will operationalize the theoretical framework outlined in the preceding sections by borrowing notions coming from social psychology, linguistics and critical discourse analysis. Thus, this section begins to enact that beneficial multidisciplinary contagion which constitutes one of the peculiarities of this work. The section is concluded by presenting the four discursive dimensions this dissertation's operational definition of populism spans, which correspond to the empirical indexes presented at the end of Chapter 2 and constitute a further step in the fruitful contamination process mentioned above.

The data collection performed as part of this research is presented in Chapter 2. Section 2.1 will present and motivate the data collection performed by researchers at the Ing. Rodolfo De Benedetti Foundation, who selected a collection of around 3 million tweets from almost 300 political leaders, posted over 12 years in 23 countries. Section 2.2 is dedicated to a preliminary EDA on the raw dataset downloaded from Twitter. Section 2.3 will discuss in detail all the methods used for data preparation, including translation into English and all feature extraction methods (manual data collection, topic analysis, sentiment analysis, LIWC-based feature extraction, survey-based feature extraction) used in this research. Further, this section also contains the description and motivation of the methods used for construction of populist online communication logic indexes in Section 2.3.5. Finally, the methods used for obtaining the results of preliminary analyses presented in Chapter 3, i.e. XGBoost and linear regression, are described in Section 2.4. It must be noted that the real aim of this work was to construct text-based variables which reflect as closely as possible the four dimensions of populism described in Section 1.3, while also building all those control variables that could be useful for testing the similarities and differences between the discursive-performative framework and other frameworks under which populism has been studied, such as the ideational one. Therefore, Chapter 2 contains many of the tasks conducted as part of this dissertation.

As stated above, the reader should keep in mind that Chapter 3 constitutes an exploration and validation of the proposed text-based empirical measures of populism, rather than a proper data analysis, which would be out of the computational and time resources of this author. That being said, Section 3.1 will present the results of the XGBoost procedure applied for a first feature selection, while Section 3.2 discusses the outcomes of a preliminary regression aimed at capturing selected variables' effects on our cumulative populism index.

Lastly, Chapter 4, in Section 4.1, will conclude that this paper provides a useful blueprint for bridging the gap between populist studies and adjacent fields, with the help of powerful analytical tools such as the ones offered today by NLP and related practices. Further research directions, such as the need for more in-depth analyses and for a similar study on the demand side of populism, will be presented in Section 4.2.

1

Populism in text: towards a definition

1.1 Social media and political life

"*Talking to friends, exchanging gossip, showing holiday pictures, scribbling notes, checking on a friend's well-being, or watching a neighbor's home video used to be casual, evanescent (speech) acts, commonly shared only with selected individuals [...] through social media, these casual speech acts have turned into formalized inscriptions, which [...] are now released into a public domain where they can have far-reaching and long-lasting effects*"[10].

Social media platforms have become one of the main tools for human communication. In 2019, a mere 30 years since the invention of the World Wide Web in 1991, around a third of world population was active on social media [11]. Yet, these services did not just change the way we communicate in private; they also altered the nature of public communication: by 2015, more than three quarters of world leaders had a social media account [10][12].

As many pivotal movements of our history, such as the Arab Spring, the election of Donald Trump and Brexit, have proven, the increased availability of social media services is part of a societal change which goes way beyond a simple technological transition. If we accept the notion that "*social order is a communicated order, and the rules and logics of the underlying formats of communication have reshaped many activities, and have initiated numerous others*" [8], we must concede that the conversations happening on these platforms are indeed changing the way politics is being made, from mass surveillance methods to electoral campaign design.

Those conversations are a good example of *media logic*, i.e. the ensemble of the *form* of communication and of the *process* through which media transmit and communicate information. As remarked in the previous paragraph, "*media logic reflexively shapes interaction processes, routines and institutional orders, while institutional orders reflect and reify a communication order operating with media logic*" [8].

Hence, media logic is a helpful theoretical concept when trying to account for social media's role in politics. As Altheide argues, the increased academic interest for this research area is not only a consequence of the popularity social media platforms have acquired, but also of the impact that the shift to this new medium and its formats has had on the features of our political discourse [8][13]. Therefore, this thesis embraces the idea that "*changes in the way we communicate – from the Gutenberg printing press to the first televised US presidential debate - always have an effect on politics*" [13].

Many expected that the shift to online communication, where no gatekeeping power could be exercised by mass media, would foster a direct link between politics and citizens, thus strengthening political engagement and participation [9]. The optimism of these observers hasn't - yet - been rewarded: in 2014, the number of European citizens who logged on to a social media platform at least once a day surpassed the number of those who voted in European elections [13]. In fact, active involvement in politics and utilization of social media platforms seem to be following opposite trends: as smartphones have started becoming easily accessible to the public and social media use has been rapidly growing, political engagement has continued declining globally. This observation is regardless of whether we measure it as trust in politicians, level of formal party membership or voter turnout [13][14].

Whether one sees them as a cause or a consequence of the decline in active political participation, social media platforms and online news outlets now constitute mainstream arenas of political life; populist actors, unlike many conventional ones, have been swift to exploit these arenas for electoral purposes [13]. Some may argue that the real protagonist of this growing electoral momentum are right-wing populist actors, rather than populist actors more generally [15]. Yet, while in Europe the majority of populist actors that have enjoyed electoral success in recent years are placed at the far-right of the political spectrum, in the rest of the world populist actors with very different ideological bases have also made notable electoral strides [2].

From this short introduction, it should be apparent that the parallel "rise" of populism and of social media precludes any investigation of these two concepts that does not consider their interplay. This thesis will go on, in Section 1.2, to define what we mean when we talk of populism and its actors, building on the existing literature.

1.2 What is populism?

The answer to the title of this section could fill – or, rather, has filled - entire books [6] [1]. This section constitutes a brief attempt to sketch the current approaches popular in academia and their link to discourse and communication studies, with the aim of laying the basis for an operational definition of populist communication, which will be outlined in Section 1.2.

Cas Mudde and Rovira Kaltwasser, perhaps the most influential researchers in the field alongside Ernesto Laclau, have said populism has been described as a political strategy, a folkloric style of politics, a way to restore popular agency, an emancipatory force and even a consequence of socio-economic changes in society [6]. In fact, a glance at the size of the Oxford Handbook of Populism should already be a convincing proof of how debated and complex the phenomenon is; in the subheadings, the word "populism" is combined with pretty much every ideology, country, political theory or praxis one could ever think of [16].

That being said, one approach to the study of populism seems to have become the *new mainstream* in the scientific community during recent years: the ideational approach [17]. According to this framework, populism can be classified as "*a thin-centered ideology that considers society to be separated in two homogeneous and antagonistic camps: the pure people and the corrupt elite; and which argues that politics should be an expression of the volonté génerale of the people*" [6].

A *thin* ideology lacks the capacity - which is instead typical of *thick* or *full* ideologies, such as Marxism or fascism - to put forward a wide-ranging and coherent programme for the solution to critical political questions [18]; to overcome this limitation, thin ideologies such as populism can be *hosted* by other *thick* ideologies such as the ones quoted above [6]. However, if one desires to study populism in general, rather than specific variants of it such as far-right populism, it is imperative to ask the question: who is *not* a populist?[4]

Under the theoretical framework of the ideational approach, answering the question at the end of the previous paragraph is not an easy task. Consider the basic tenets of populism as proposed by researchers who adhere to this approach:

- society is made of electors – *the people* – and elected representatives – *the elite*
- sovereignty ultimately rests with the people
- sovereignty can be temporarily given by the people to the elite, with the aim of giving the elite the task of exercising the general will of the people

After all, aren't these so-called populist ideas the core beliefs which constitute the very basis of modern representative democracy? Isn't the lack of clear parties' and leaders' accountability in today's complex, globally interconnected democratic systems one of the most objective reasons for the progressive decline in political participation [13]? And isn't this (real or perceived) unaccountability of elected representatives a good reason for citizens to believe that the *corrupt elite* is taking advantage of them, i.e. of the *pure people*? Therefore, aren't all and any parties, provided they exist in a democratic environment, technically populist parties? [1]

The discursive-performative approach to populism describes fundamental populist attitudes in a way that, at first glance, could seem identical to that proposed by ideationalists:

- populist attitudes are constructed around the nodal point *the people*
- these attitudes reflect a perception/representation of society as divided between two hostile camps : *the people* against *the elite*

As is often the case in political science, the devil is in the detail: while the ideational approach sees populism as a particular world view, ideology or set of ideas that citizens, parties and politicians can *hold*, the discursive-performative approach sees populism as something that is *done* and expressed by political actors[17].

Why, one may ask, is this nuance in understanding so crucial to this dissertation?

Firstly, one of the main fallacies of recent empirical studies of populism has been that of undervaluing the importance of constructing a strong and specific theoretical basis on which to build the empirical investigation [4]. While a broad definition, such as the one typically used by empirical studies of populism conducted under the ideational framework (e.g. [9] and [12]), allows for wide geographical, chronological and ideological comparisons, it does not allow for any resolution when it comes to the challenge presented in Section 2.1: what role did, do and will the Internet and social media play in the spreading of populist messages? What does this new medium do to the form and content of these people-centric, anti-elitist messages? More specifically, could the network media logic (i.e. the media logic performed on social media platforms, as defined by Klinger and Svensson in [19]) through which populist messages are communicated be what distinguishes modern populist communication from people-centric, anti-elitist communication more generally, which has itself existed since the birth of democracy? [1] This dissertation contends that these questions are impossible to answer under a purely ideational framework, which assumes that populist beliefs have to be *held* by a politician before they can be communicated to the public: firstly, it is common knowledge that ideological beliefs are rarely correctly self-reported. Secondly, no expert classification exists, to this author's knowledge, of all worldwide leaders from the perspective of populism studies - and, even if it did, the rapid changes in balances of power within and between parties would rapidly make it outdated.

Instead, the discursive-performative approach shifts the emphasis from the *content* to the *form* of populism, focusing primarily on the shared populist logic, i.e. on "*the particular way in which the various discursive elements are organized and articulated in a given discourse*" [17], thus allowing for the consideration of the role of media logic in spreading and even modifying populist messages (see Section 1.1). Under this theoretical umbrella, leaders who use populist rhetoric as a tool for furthering their agenda or popularity may, or may not, *hold* populist beliefs from an ideological point of view; populism is rather seen as a *discursive practice* that seeks to pitch *the people* against *the elite*, with a wide variety of options for how this discourse construction takes place and how conflicts around the meaning of the *people* and *elite* signifiers play out [20].

It should be evident now, after this discussion, that this author has several reasons to adopt the performative-discursive approach over the ideational one as the backbone for this research project. Keeping in mind that the aim of this dissertation is to lay the foundations for constructing a generalizable and reliable measure of populism in text, regardless of geographical, ideological or chronological variety, and that the dataset employed has been mined from Twitter (see Chapter 2 for a full description), here is a summary of the main challenges faced while conducting this research, which ultimately led this author to adopt the peformative-discursive theoretical perspective:

(i) The online arena dedicated to the spreading of populist messages is far from neutral, and reflexively shapes the form and content of these appeals by virtue of its (network) media logic [19][8]

(ii) While, to a good extent, we can ascertain a leaders' political ideology from their party affiliation, the ideational approach itself argues that populism is not attached to any *thick*, recognizable, ideological content, empirical research conducted under this framework tends to produce analyses that focus on one or a few types of populism, based on their *host* ideology [6], which often tend to result more relevant to the study of the host ideology rather than to populist studies

(iii) Twitter and other social media data are large-scale, easily collectible collections of leaders' personal appeals to the electorate or to their adversaries, i.e. they are large collections of performed acts of online communication. As such, they should be studied by taking into account the impact of discourse performativity and of online mediation and mediatization [1][20][8]

The solution to challenge (i) is a somewhat brute-force argument: by using a dataset homogeneously mined from the same social media platform, the hope is to do away with any potential differences in content, form and logic that could be present when comparing data coming from heterogeneous social media services. Moreover, the translation of the whole dataset into English, performed through Google's context-aware NMT method, should in theory ensure that the form of messages in maintained homogeneous across different countries and languages [21]. Regardless, the time span of the dataset (12 years) and the geographical variety (23 countries) means this dataset is bound to suffer from problems due to cultural differences, which are not fully accounted for. Further research should aim to comprehend more specifically the features of online communication as described in [7], as to be able to capture the effect these structures have on online political discourse in various cultures and languages, thus allowing for a more structured response to this challenge.

Conversely, there is no attempt here to actively find a solution to challenge (ii); quite simply, ideological labels will be included in the analyses as control (see Chapter 3) and have been carefully manually coded, in order to reflect changes in ideologies over time and ensure maximum comparability between the ideational and discursive-performative approaches (see Section 2.3.3 for more details).

Challenge (iii), which is of course closely connected to challenges (i) and (ii), is the real turning point for understanding the operative definition sketched in Section 1.3. As argued by Chatterje-Doody and colleagues in [22], a discursive-performative approach to populism allows for identifying a populist (online) communication logic which is in no way exclusive to political actors, but can also been employed by other actors that gravitate around the political arena, including private economic actors, leaders of international organizations, opinion leaders and the media, sometimes in unexpected ways. The authors make the example of online news articles about US former President Donald Trump posted on the website *i100*, an online, left-leaning news outlet founded by the owners of The Independent newspaper. In many of these articles, critical stances about the billionaire-turned-president are conveyed using a populist communication logic that positions Trumps as "*an entitled, elitist leader out of touch and at odds with the people*". Moreover, they point out that these kinds of articles often include content originally shared by political leaders on social media, further corroborating the idea that online social platforms, far from just mirroring mass media behaviour in the digital sphere, are actually changing the way we live and perform the political [22].

From the example above, but also from the discussion of the discursive-performative approach to populism, it should be clear that viewing populism as a type of performed, discursive communication, i.e. as "*a type of language that has important effects on how political identity – and politics more broadly - operates*" [20], allows for both more precision in specifying what populism is, as we can now base this definition on empirically observable discourse rather than on party-based ideologies, and for more flexibility in data collection and preparation, as researchers can choose and collect their own data, rather than relying on data coming from expert studies, which are extremely costly to conduct [4].

1.3 Populist online communication logic: operational definition

While it is true that research on populism, particularly on the empirical side, has blossomed in the past decade [4], much still has to be discovered, particularly from the point of view of quantitative, text-based methods. As Meijers and Zaslove argue, "*we lack data that measures populism in political parties in a valid and precise manner, that recognizes that populism is constituted by multiple but distinct dimensions, and that ensures full coverage of all parties in Europe*" [23]. Moreover, there are currently no studies, to this author's knowledge, that study populism with text-based methods at the leader-level of inquiry to the geographical and chronological extent covered in this dissertation.

The operational definition outlined in this section builds on the discursive-performative approach to populism described in Section 1.2. This approach, as any Critical Discourse Analysis approach, contends that "*the situational, institutional and social context shape and affect the discourse and, in turn, the discourse influences the social and political reality*" [24]. As a consequence, this thesis is focused on the various dimensions of what we could call *populist online communication logic*. While this concept has been employed in the literature before (see [7]), so far it seems that it hasn't been applied under the discursive-performative theoretical framework for studying populism.

Keeping these premises in mind, *populist online communication logic* is defined and operationalized as a concept spanning across the following four dimensions:

(i) *People-centrism*: putting *the people*, however identified (through nationality, ethnicity, etc.) at the center of the discourse

(ii) *Anti-elitism*: expressing negative feelings towards *the elite*, however identified (political, economic, etc.)

(iii) *Identification of leader with the people*: framing the discourse in such a way that the speaker, while they are technically part of the (political) elite, is actually perceived as a member of *the people*

(iv) *Antagonism*: using discourse to create feelings of antagonism and conflict, from *the people* towards *the elite* but also from *the people* towards an undefined *other* who is not included in the in-group (e.g. immigrants, economic elites...)

The rationale underlying these shorts definitions for the four dimensions has been extensively discussed in Section 1.1 and Section 1.2.

This definition is operationalized by means of feature extraction, index generation and comparison of the constructed variables with other ones obtained under different research designs. Most variables have been constructed directly from textual data, as not to make any assumptions about *who* might be using a populist online communication logic. Methodological choices for textual variable construction are motivated and discussed in Section 2.3.3.

Moreover, in order to capture *how* this populist online communication logic plays out in terms of electoral gains/losses, government/opposition roles and rhetoric style, four indexes (plus one cumulative index) have been constructed on the basis of the operational definition above. The reader will find a detailed description of index construction in Section 2.3.3.

Finally, data about party ideology, party attitudes about divisive issues (e.g. migration, climate change, etc.) has been collected by exploiting the 2019 Global Party Survey [25] and Manifesto Project [26] variables: by controlling for the effect of ideological differences at a fine-grained level, this dissertation aims to enable future studies to measure the extent of the use of populist communication logic for each leader, *regardless* of their declared political affiliation and policy choices. Additionally, it will be interesting to discuss the different conclusions about the nature and effects of populist communication that a text-based approach generates vis-à-vis a survey-based approach, and any similarities between the two approaches.

It must be noted that, under the approach chosen to frame populism in this thesis, populist attitudes to online communication are not interpreted as privileged entry points into the real and supposedly stable identity of social actors (such as they would be, for example, under an ideational approach), but rather as discursive units like many others [17]. This means that we are measuring how populist someone is *while* they communicate online, rather than aiming at measuring their underlying populist attitudes. In Chapter 3, which reports the results of a preliminary data analysis on the variables constructed, the difference between the two frameworks will become explicit: while not all political leaders are ideologically populist, any politician has the potential - and, in terms of electoral gain, the incentive - to use a populist communication logic on social media.

By showing how all textual variables used in the dissertation are constructed and validating their use on such a large dataset in terms of time and space, the hope is to produce a generalizable measure of populism in text, so that future works can use this as a starting point for more in-depth studies of populism in text.

2
Methodology

This dissertation was developed entirely with Python programming language v3.10. Part of the work was carried out locally, part was carried out on the Bocconi University BIDSA server and most of it on a Colab Premium Pro account.

This chapter is divided in three sections: Section 2.1 describes the method for data collection and the criteria for data selection; Section 2.2 is dedicated to a preliminary EDA of the collected data; finally, Section 2.3 describes in detail the methods used for feature extraction and engineering, including the generation of indexes relative to the operational definition of populism presented in Section 1.3.

2.1 Data collection and selection criteria

The data collection portion of this project has been performed by researchers at the Ing. Rodolfo De Benedetti Foundation (fRDB) [27] through Twitter Academic API, version 2 [28]. Researchers collected these tweets in the context of the project *Populist rhetoric and representative democracy*, which I have joined since July 18th, 2022.

In 2019, when the fRDB inaugurated the project, it represented one of the first attempts to study online populist rhetoric and behaviour with quantitative, ML-based methods. The main objective is that of constructing a measure of populism which will allow researchers to study it at the level of single politicians, rather than at the party level as has been done in existing research (e.g. by Pippa Norris in [2]).

The fRDB project focuses on examining the *supply* side of populism, specifically concentrating on individual politicians, their rhetoric, and their behaviour. In Chapter 4, this thesis will delve deeper into this decision and highlight the necessity for a complementary project addressing the *demand* side of this phenomenon. Nevertheless, to the best of this author's knowledge, there is currently no study of populist attitudes that encompasses the same temporal and geographical scope as this project, at least at the leader-level of inquiry [4].

Twitter's Academic API has been an invaluable tool to researchers around the world who are interested in text and network analysis, since it facilitates data collection by embedding metadata, user interactions and other features in each downloaded [28]. While is it is true that some political conversations, particularly those within groups of radical supporters, tend to take place more often on Facebook, Twitter has always been a privileged place of political discussion, as it is used daily by journalists and other - real or self-defined - "watchdogs" of those in power, including so-called populist critics sitting in government oppositions [29]. Hence, the fRDB researchers chose Twitter as their text mining environment, perhaps also because of the notorious usability and openness of this social media platform's API.

The political leaders covered in the data mining on Twitter were selected according to the following criteria:

(i) all candidates for Prime Minister/President in recent general elections (between $2001 - 2020$)

(ii) leaders of main political parties with vote share above 5% in at least one general election over the observed period ($2010 - 2022$)

(iii) selected influential political leaders or leaders of political minorities, even if the vote share of their party is below 5% (e.g., (e.g. Giorgia Meloni in Italy, Nicola Sturgeon in UK, Basque and Catalan political leaders in Spain)

These criteria isolated 404 political leaders for the 23 countries selected. Next, the researchers restricted the selection to leaders with an active Twitter account. At the beginning of this author's experience as a fRDB researcher, the raw dataset had just been mined and no further steps had been taken in data cleaning. A considerable amount of time was spent on manually checking the username-name correspondence for each leader and verifying if there were any problems during data retrieval (e.g. some non-verified accounts were removed, as they were presumably run by someone else). After removing unreliable data and duplicated tweets, the dataset contains 3, 724, 837 tweets, posted between January 1st, 2010 and June 13th, 2022 by 392 leaders. Of these, 2, 078, 962 are tweets and 1, 645, 875 are retweets.

A complete description of the countries and leaders included in the full dataset can be found in Table 2.1, while you can find the complete description of features downloaded alongside each tweet during the data mining process in Table 2.2. The general characteristics of the dataset are described in Section 2.2.

Table 2.1: Countries and corresponding political leaders included in the dataset

Selected leaders	Country
Andrew Bartlett, Bill Shorten, Bob Katter, Campbell Newman, Christine Milne, Clive Palmer, Deb Frecklington, John Anderson, John-Paul Langbroek, Julia Gillard, Kevin Rudd, Kim Beazley, Malcolm Turnbull, Mark Latham, Natasha Stott Despoja, Nick Xenophon, Pauline Hanson, Richard Di Natale, Scott Morrison, Tim Nicholls, Tony Abbott	Australia
Alexander Van der Bellen, Beate Meinl-Reisinger, Christian Kern, Heinz-Christian Strache, Josef Bucher, Matthias Strolz, Norbert Hofer, Pamela Rendi-Wagner, Peter Pilz, Sebastian Kurz, Ulrike Lunacek, Werner Kogler	Austria
Alexander De Croo, Bart De Wever, Benoît Lutgen, Caroline Gennez, Charles Michel, Didier Reynders, Elio Di Rupo, Guy Verhofstadt, Gwendolyn Rutten, Herman Van Rompuy, Jean-Marc Nollet, John Crombez, Joëlle Milquet, Juliette Boulet, Marianne Thyssen, Meyrem Almaci, Peter Mertens, Philippe Henry, Philippe Lamberts, Stefaan De Clerck, Tom Van Grieken, Wouter Beke, Wouter Van Besien, Yves Leterme, Zoé Genot	Belgium
Anthony Garotinho, Aécio Neves, Cabo Daciolo, Ciro Gomes, Cristovam Buarque, Dilma Rousseff, Fernando Haddad, Geraldo Alckmin, Heloísa Helena, Henrique Meirelles, Jair Bolsonaro, José Serra, João Amoêdo, Luciana Genro, Luiz Inácio Lula da Silva, Marina Silva, Rui Costa Pimenta	Brazil
Andrew Scheer, Claude Carignan, Don Plett, Elizabeth May, Gilles Duceppe, Jagmeet Singh, Jim Harris, Justin Trudeau, Leona Alleslev, Mario Beaulieu, Martine Ouellet, Michael Ignatieff, Peter Julian, Stephen Harper, Stéphane Dion, Suzanne Cowan, Tom Mulcair, Yves Perron, Yves-François Blanchet	Canada
Andrej Babiš, Ivan Bartoš, Jan Farský, Jiří Paroubek, Jiří Rusnok, Karel Schwarzenberg, Lubomír Zaorálek, Mirek Topolánek, Miroslav Kalousek, Miroslava Němcová, Pavel Bělobrádek, Petr Fiala, Tomio Okamura, Vladimír Špidla, Vojtěch Filip	Czech Republic
Anders Fogh Rasmussen, Anders Samuelsen, Bendt Bendtsen, Helle Thorning-Schmidt, Holger K. Nielsen, Kristian Thulesen Dahl, Lars Barfoed, Lars Løkke Rasmussen, Margrethe Vestager, Marianne Jelved, Mogens Lykketoft, Morten Østergaard, Naser Khader, Pernille Skipper, Pernille Vermund, Pia Kjærsgaard, Pia Olsen Dyhr, Søren Pape Poulsen, Uffe Elbæk	Denmark
Alexander Stubb, Anna-Maja Henriksson, Anneli Jäätteenmäki, Anni Sinnemäki, Antti Rinne, Carl Haglund, Eero Heinäluoma, Harry Harkimo, Juha Sipilä, Jussi Halla-aho, Jutta Urpilainen, Jyrki Katainen, Li Andersson, Mari Kiviniemi, Matti Vanhanen, Osmo Soininvaara, Paavo Arhinmäki, Pekka Haavisto, Petteri Orpo, Päivi Räsänen, Sari Essayah, Suvi-Anne Siimes, Tarja Cronberg, Ville Niinistö	Finland
Benoît Hamon, Bernard Cazeneuve, Dominique Voynet, Dominique de Villepin, Edouard Philippe, Emmanuel Macron, Eva Joly, François Bayrou, François Fillon, François Hollande, Jean-Luc Mélenchon, Jean-Marc Ayrault, Jean-Marie Le Pen, Jean-Pierre Raffarin, Manuel Valls, Marine Le Pen, Nicolas Dupont-Aignan, Nicolas Sarkozy, Noël Mamère, Ségolène Royal	France
Alice Weidel, Bernd Lucke, Cem Özdemir, Christian Lindner, Dietmar Bartsch, Frank?Walter-Steinmeier, Gabi Zimmer, Gregor Gysi, Jürgen Trittin, Katrin Göring-Eckardt, Martin Schulz, Peer Steinbrück, Renate Künast, Sahra Wagenknecht	Germany
Bocskor Andrea, Deli Andor, Deutsch Tamás, Edina Toth, Gergely Karácsony, Gordon Bajnai, István Ujhelyi, Judit Varga, Katalin Novák, László Andor, Márton Gyöngyösi, Péter Niedermüller, Tibor Navracsic, Zoltan Kovacs	Hungary

Table 2.1 Countries and corresponding political leaders included in the dataset - continued

Selected leaders	Country
Alessandro Di Battista, Beppe Grillo, Carlo Calenda, Daniela Santanché, Emma Bonino, Enrico Letta, Francesco Rutelli, Giorgia Meloni, Giuseppe Conte, Luigi Di Maio, Mario Monti, Matteo Renzi, Matteo Salvini, Paolo Gentiloni, Pier Ferdinando Casini, Pier Luigi Bersani, Pietro Grasso, Roberto Maroni, Silvio Berlusconi, Walter Veltroni	Italy
Alberto Anaya Gutiérrez, Alejandro Moreno, Alfonso Ramírez Cuéllar, Andrés Manuel López Obrador, Beatriz Mojica Morga, Carlos Puente Salas, Carolina Viggiano, Clemente Castañeda Hoeflich, Dante Delgado, Enrique Peña Nieto, Felipe Calderón, Gabriel Quadri de la Torre, Hugo Eric Flores Cervantes, Héctor Larios, Jaime Rodríguez Calderón, Josefina Vázquez Mota, José Antonio Meade, Luis Castro Obregón, Manuel Granados Covarrubias, Marko Cortés, Martí Batres, Patricia Mercado Castro, Reginaldo Sandoval Flores, Ricardo Anaya, Roberto Campa Cifrián, Yeidckol Polevnsky	Mexico
Alexander Pechtold, André Rouvoet, Arie Slob, Diederik Samsom, Emile Roemer, Geert Wilders, Gert-Jan Segers, Henk Krol, Jan Peter Balkenende, Jesse Klaver, Job Cohen, Kees van der Staaij, Lodewijk Asscher, Marianne Thieme, Mark Rutte, Paul Rosenmöller, Sybrand van Haersma Buma, Thierry Baudet, Thom de Graaf, Tunahan Kuzu	Netherlands
Audun Lysbakken, Bjørnar Moxnes, Dagfinn Høybråten, Erna Solberg, Hanna E. Marcussen, Jens Stoltenberg, Jonas Gahr Støre, Knut Arild Hareide, Kristin Halvorsen, Liv Signe Navarsete, Rasmus Hansson, Siv Jensen, Thorbjørn Jagland, Trine Skei Grande, Une Aina Bastholm, Åslaug Haga	Norway
Barbara Nowacka, Beata Maria Kusińska, Donald Tusk, Ewa Kopacz, Grzegorz Napieralski, Janusz Korwin-Mikke, Janusz Palikot, Janusz Piechociński, Jarosław Aleksander Kaczyński, Leszek Miller, Mateusz Morawiecki, Małgorzata Maria Kidawa-Błońska, Paweł Piotr Kukiz, Roman Giertych, Ryszard Petru, Waldemar Pawlak, Wojciech Olejniczak, Władysław Marcin Kosiniak-Kamysz, Włodzimierz Czarzasty	Poland
André Ventura, António Costa, Assunção Cristas, Carlos Guimarães Pinto, Catarina Martins, José Manuel Barroso, José Sócrates, Mário Centeno, Pedro Passos Coelho, Rui Rio	Portugal
Andrej Danko, Andrej Kiska, Igor Matovič, Ján Fige?, Mikuláš Dzurinda, Peter Pellegrini, Richard Sulík	Slovakia
Alenka Bratušek, Borut Pahor, Dejan Židan, Franc Bogovič, Gregor Golobič, Gregor Virant, Janez Janša, Janez Podobnik, Karl Erjavec, Katarina Kresal, Ljudmila Novak, Luka Mesec, Marjan Šarec, Matej Tonin, Miro Cerar, Violeta Tomic, Zmago Jelinčič Plemeniti, Zoran Janković	Slovenia
Albert Rivera, Cayo Lara, Francesc Homs, Gabriel Rufián, Gaspar Llamazares, Joan Ridao, Josu Erkoreka, Mariano Rajoy Brey, Oriol Junqueras, Pablo Casado, Pablo Iglesias, Pedro Sánchez, Rosa Díez, Santiago Abascal	Spain
Alf Svensson, Annie Lööf, Ebba Busch Thor, Gudrun Schyman, Göran Hägglund, Göran Persson, Isabella Lövin, Jan Björklund, Jimmie Åkesson, Jonas Sjöstedt, Lars Ohly, Maria Wetterstrand, Peter Eriksson, Stefan Löfven, Åsa Romson	Sweden
Arlene Foster, Boris Johnson, Charles Kennedy, David Cameron, Ed Miliband, Gordon Brown, Jeremy Corbyn, Jo Swinson, John Swinney, Nick Clegg, Nicola Sturgeon, Nigel Farage, Paddy Ashdown, Theresa May, Tim Farron, William Hague	United Kingdom
Alexandria Ocasio-Cortez, Barack Obama, Bernie Sanders, Bill Frist, Chuck Schumer, Donald Trump, Elizabeth Warren, George Bush, Harry Reid, Hillary Clinton, Joe Biden, John Boehner, John Kerry, John McCain, José E. Serrano, Kevin McCarthy, Mike Pence, Mitch McConnell, Mitt Romney, Nancy Pelosi, Paul Ryan, Steny Hoyer	United States

2.2 Exploratory Data Analysis

This section will explore the basic characteristics of the raw dataset. Figure 2.1 depicts the number of tweets collected for each of the 23 countries included in the dataset. Spain is the country with most tweets extracted (358,347), followed by Mexico (341,955) and Canada (307,340).

Variable	Description	Included in dataset
attachments.media_keys	Media associated with the tweet	No
author_id	User ID of the tweet author	Yes
context_annotations	Contextual information about entities in the tweet	Yes
created_at	Time the tweet was created	Yes
entities.annotations	Annotations indicating entities in the tweet	No
entities.hashtags	Hashtags present in the tweet	Yes
entities.mentions	Users mentioned in the tweet	Yes
entities.urls	URLs present in the tweet	Yes
id	Unique ID of the tweet	Yes
lang	Language of the tweet	Yes
public_metrics.like_count	Number of likes the tweet received	Yes
public_metrics.quote_count	Number of times the tweet was quoted	Yes
public_metrics.reply_count	Number of replies to the tweet	Yes
public_metrics.retweet_count	Number of times the tweet was retweeted	Yes
reply_settings	Reply settings for the tweet	Yes
text	Content of the tweet	Yes
attachments.poll_ids	Polls associated with the tweet	No
entities.cashtags	Cashtags present in the tweet	No
withheld.copyright	Copyright information if withheld	No
withheld.country_codes	Country codes if content is withheld	No
withheld.scope	Scope of content withholding	No
domain_idx	Domain index	No

Table 2.2: Twitter API metadata and inclusion in dataset

The countries with the least tweets extracted are Portugal (41,040), Hungary (29,791) and Slovakia (17,759).

Figure 2.2 describes the number of tweets extracted for each of the years present in the dataset (2010-2022). It is apparent that Twitter use was not very popular with politicians until 2011, but we are still interested in the evolution of politicians' behaviour, so the whole time series has been preserved for future analysis.

It is clear from Figure 2.3, which depicts the distribution of tweet length (measured as the number of words per tweet) across the whole dataset, that tweets with less than 2 word tokens don't really fit into the left-skewed normal distribution. Therefore, these tweets were dropped before data preparation, leaving us with $3,704,333$ tweets, of which $2,058,458$ are actual tweets and $1,645,875$ are retweets.

Moving on, Section 2.3 will now delve into the details of how this raw dataset was manipulated in order to create both variables connected to dimensions of populist rhetoric and control variables. In particular, the next section deals with: data preparation, including cleaning; methods used for feature extraction and engineering; methods used for validation of the constructed features through preliminary analyses.

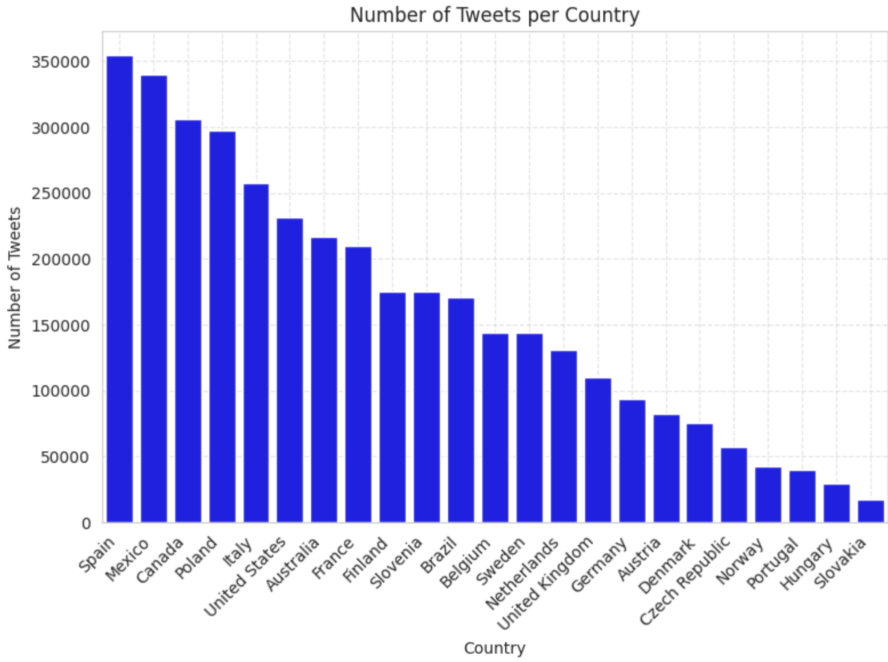

Figure 2.1: Number of tweets per leader's country of origin.

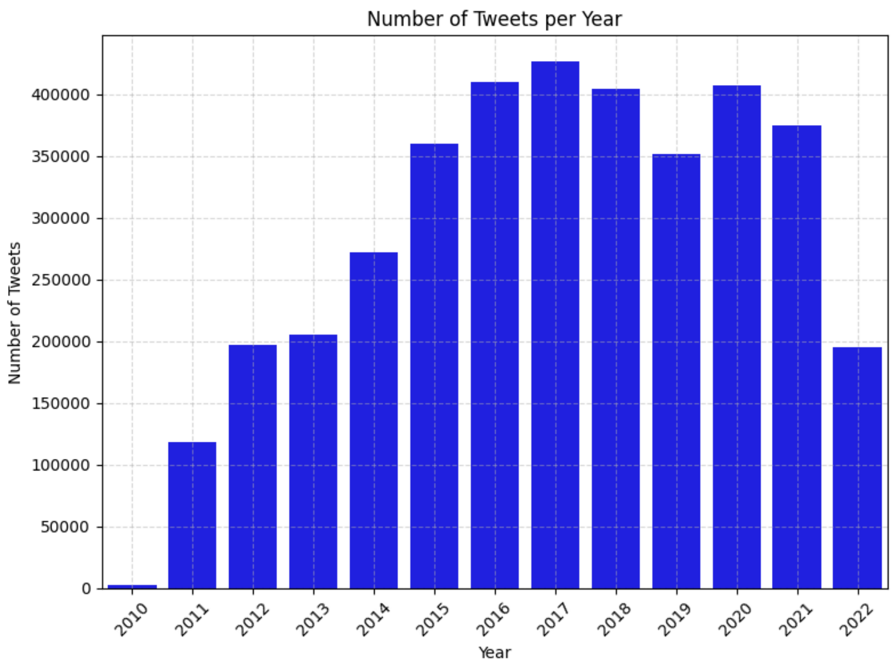

Figure 2.2: Number of tweets per year.

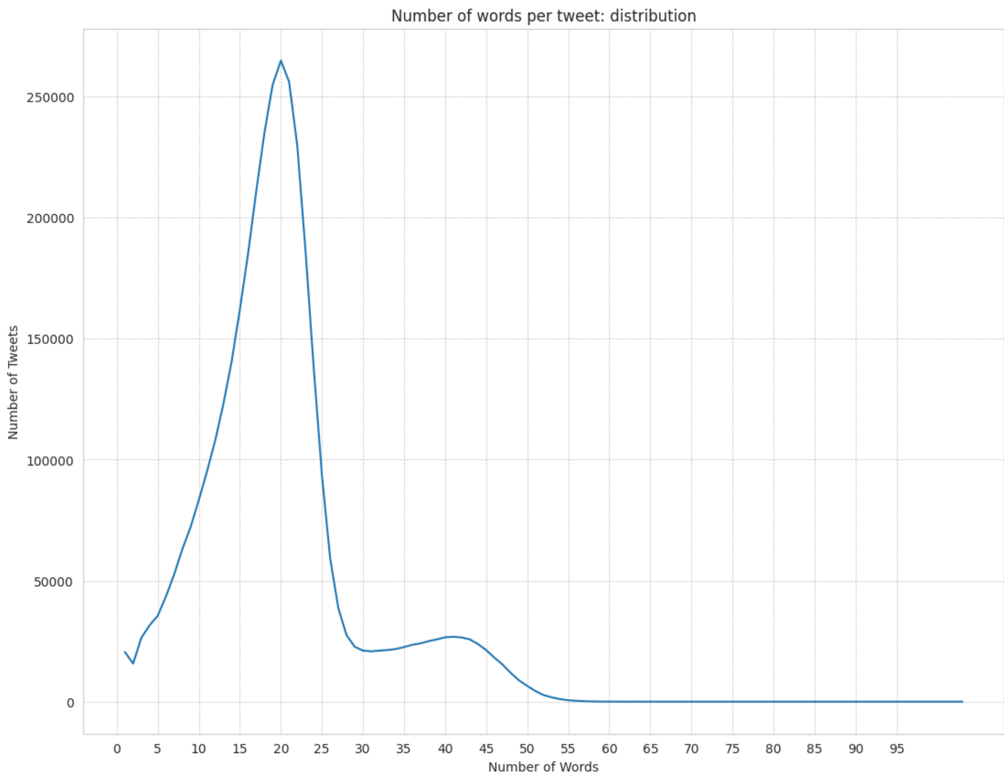

Figure 2.3: Distribution of number of words per tweet

2.3 Data preparation

2.3.1 Basic text cleaning

Before applying any of the techniques described in the rest of the section, the following basic text cleaning steps were performed:

- Lowercasing
- Special characters and emojis removal
- User mentions (@'s) removal
- URL and other hyperlinks removal

Hashtags were kept as they may contain useful information for some processes, such as topic analysis. If a technique required further cleaning steps, this is specified in the appropriate subsection.

2.3.2 Translation into English

When downloading data from Twitter, the API itself analyses the language of the tweet and an ISO code is attached as metadata to each tweet for language identification [28]. In Figure 2.4, you can observe the distribution of the number of tweets for each of the 53 languages present in the dataset. Most likely, some of the most uncommon languages may have been erroneously detected and assigned. However, as the verification of this fact went beyond the scope of this dissertation, those tweets were removed from the dataset, leaving us with $2,725,552$ data points.

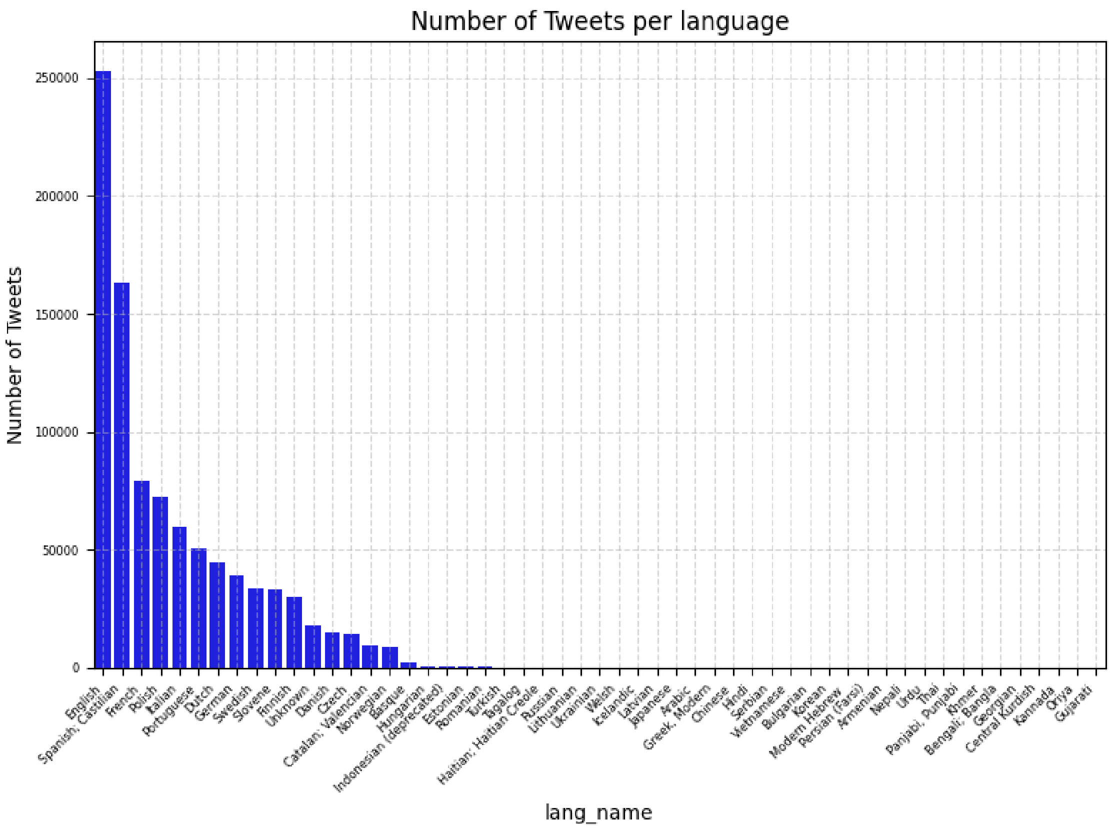

Figure 2.4: Number of tweets per language

Keeping in mind that the subsequent steps of data preparation involve deep learning methods, dictionary methods and other complex NLP tasks, it was decided that all tweets shall be translated into English. While it is acknowledged that the translation process may introduce errors and overlook certain language-specific nuances, the benefits of translating the tweets into English outweigh these potential drawbacks. The advantages, such as ensuring consistency, facilitating cross-cultural comparability, leveraging readily available models, libraries, and dictionaries, as well as simplifying text pre-processing, make translation into English the optimal choice [30].

The translation was carried out using the open-access Python library *googletrans* [21]. Translating the whole dataset into English took around 14 days on the Bocconi University BIDSA server.

2.3.3 Feature extraction

Manual data entry

Several variables had to be coded by hand, since no databases were available that contained all the information needed across the 23 countries of our dataset and the 12 years it spans timewise.

In particular, the challenge for this paper was to create a leader/party classification that reflected changes in party affiliation, ideology and electoral success over time (e.g. a leader once belonging to a party in opposition may very well change affiliations and switch to a party with a different ideology that is in government during the time span our dataset includes, i.e. 12 years). The reasons behind the necessity for a fine-grained assignment of leader's ideologies - which are assumed to be reflected by party affiliation - should be clear from the discussion in Section 1.3.

All variables, sources and criteria used for the assignment are summarized in Table 2.3. In order to ensure consistency, the criteria assignment was performed by two researchers separately and reviewed by a third party.

Dictionary-based variable construction

While it is true that deep-learning methods, particularly transformer-based ones such as the now renowned OpenAI's GPT-3, on which the chatbot *ChatGPT* has been based, seem to have forever revolutionized the world of text analysis, dictionary methods are still widely used in many social sciences applications, particularly those dedicated to political discourse analysis [31][32][33][34].

There are several reasons behind the persistence of the use of these models in the social sciences. Firstly, humanities and social science researchers often follow very diverse paths toward digital literacy (i.e. the development of skills related to digital tools such as data mining, data

Variable	Variable name	Description	Source	Criteria for assignment	Data type
Dynamic Party Label	party_final	Party affiliation of leader, considering changes in affiliations over time, if relevant	Wikipedia and reliable news articles	• Party affiliation is considered only for nationally relevant parties • Party affiliation is re-evaluated at each national election, but never in between • The affiliation of a leader from January, 2010 (earliest month in dataset) to the time of the earliest national election in the dataset is assumed identical to the leader's party affiliation during the first electoral period included • The affiliation of a leader from the time of the latest national election in the dataset up to June, 2022 (latest month in dataset) is assumed identical to the leader's party affiliation during the last electoral period included, unless an explicit change in affiliation is found in the sources	string
Political career start	entered	Time of entrance into politics for each leader, if later than January 1st, 2010	Wikipedia and reliable news articles	• No record of previous political activity (European and local political activities are included) • Explicit statement in the sources that specifies activities carried out in a different professional sector before entering politics	date
Political career end	retired	Time of retirement from politics for each leader, if earlier than June 13th, 2022	Wikipedia and reliable news articles	• No record of subsequent political activity (European and local political activities are included) • Explicit statement in the sources that specifies activities carried out in a different professional sector after leaving politics • Explicit retirement statement in sources • No public declaration after supposed retirement date for at least 2 years	date
Gender	male	Gender of leader	Wikipedia and reliable news articles	• If leader is male, dummy=1; if leader is female, dummy=0	integer
Political role	opposition	leader's inclusion in government	Wikipedia	• If leader's party is not in government for the electoral period the tweet belongs to, opposition=1; else, opposition=0 • Leaders who have run with no party affiliation in elections are accounted for when possible	integer

Table 2.3: Manually coded variables' description and assignment criteria

cleaning, programming skills, etc.) [35]; secondly, these methods are, by their own nature, highly interpretable, as we can directly observe the words included in each dictionary in order to evaluate and ameliorate performance; lastly, their long-standing utilization in the social science scientific community means researchers can perform comparisons with results coming from existing studies [36]. This dissertation employs dictionary-based methods for two different feature extraction tasks:

(i) **Topic analysis pre-processing**: as anticipated when discussing the methods used for topic analysis, three topic variables, associated to what we presume to be the most divisive issues in contemporary politics (migration, climate change and the labour market), were first created with a very simple keyword-matching approach that is the equivalent of a low-level dictionary method. The lists of words used for each topic's dictionary can be found in Table 2.6. The subset of tweets that were assigned a topic ($432,881$ tweets) was used as the training set for guided topic modeling using BERTopic in order to avoid issues in unsupervised topic analysis due to the size of the dataset used and diversity of the topics it contains (refer to topic analysis subsection for more details). This task was carried out in a Google Colab Pro Premium environment with A100 GPU and took around 5 minutes to complete.

(ii) **Discursive features extraction through LIWC-2022**: discursive variables extracted from text (e.g. all-or-nothing discourse dynamic, use of first-person plural pronouns to create a sense of community, etc.) were extracted using the LIWC-2022 software and dictionary, a software developed by social psychologists in the 1990s which has enjoyed great popularity among social scientists ever since[31]. The LIWC-2022 dictionary is made up of over 12,000 words, word stems, phrases, and emoticons grouped into categories designed to assess psychosocial constructs[36]. However, it is also possible to use other external dictionaries or to make up your own dictionary, which allows for great flexibility. The LIWC software was used to extract discursive and linguistic features using both the standard LIWC-2022 dictionary (see [36][37] for more details) and the Moral Foundations dictionary [38]; this particular dictionary has been developed specifically to identify moral foundations underpinning discourse, such as loyalty to one's social group, or the moral evaluation certain values, e.g. fairness, authority and so on. It has led to interesting results in social and political science [39], as it identifies discursive features highly relative to extremist and exclusionary attitudes. In particular, the empirical study about racist attitudes by Faulkner and Bliuc [40] was a great inspiration when choosing which moral and textual features to include in the indexes used to measure populist communication attitudes, especially those around in-group/out-group dynamics. Included features and their characteristics are described in Table 2.4, while index construction is illustrated in detail in Section 2.3.3. This task was carried out using LIWC-2022 software under an academic licence and took around 6 hours total to complete.

Topic analysis

As previously discussed in Chapter 1, this dissertation analyses populism from a discursive-performative perspective; unlike the ideational one, on which most empirical studies available nowadays have been based, this approach does not assume that expressing populist attitudes in discourse necessarily correlates with holding populist views at the ideological level [17]. Put plainly, this means that whatever is mentioned in the analysed populist text (policy issues, social issues, etc.) is not anymore the focus of analysis: the discursive-performative approach moves the focus of populism studies from the *content* to the *form* of populism. In order to correctly employ the theoretical framework under which this dissertation was developed, topic analysis was conducted on the whole dataset; this will enable researchers to control for the effect of the topics of tweets during the analyses and analyse more precisely the role and weight of discursive features in producing populist messages (see Chapter3).

Several methods were considered during the preliminary phases of topic analysis (e.g. classic LDA, NMF...). However, further research pointed quite clearly at the fact that unsupervised models based on neural networks, and particularly on BERT-generated embeddings, are now outperforming all other methods on a number of topic analysis tasks [44], to the point that they are also being used as a starting point for using "older" models, such as LDA. The reason BERT-generated embeddings work so well in a variety of NLP-related processes (online reviews analysis, machine translation, etc.) is that the embeddings generated with these techniques capture the contextual information of tokens (i.e. the smaller units of the text processed by BERT) by allowing the transformer-based neural-network architecture to consider words surrounding the tokens in both directions. Thus, BERT embeddings are designed to capture the semantic and syntactic relationships between words and their contextual meanings, allowing downstream natural language processing tasks to benefit from this rich representation.

Amongst all the BERT-based tools and variants, one that has sparked much interest and has been tested on a great variety of tasks is BERTopic [45]. This model exploits SBERT, a variant of BERT used for generating contextual sentence vector representations [46], for the embedding step of the topic analysis process. As you can observe in Figure 2.5, this is just the first of $5 + 1$ steps that BERTopic takes when generating topic representations, the others being: dimensionality reduction of embeddings (usually through UMAP [47]); clustering (by deafult with HDBSCAN, which allows noise to be modeled as outliers [48]); tokenization of words in the clusters (not in the single documents) through a Bag-of-Words approach, such as the *CountVectorizer* method from the *scikit-learn* Python library [49]; finally, topic representations generation for each cluster through a class-based TF-IDF model, depicted in Equation 2.1. However, the true strength of BERTopic lies in the fact that is it *modular*, i.e. all the "building blocks" depicted in Figure 2.5 could be swapped for other models, allowing for this model to evolve as technological advancements are made [45]. Moreover, additional steps can be taken to fine-tune the model according to the specific characteristics of the dataset, such as hierarchical topic merging, outlier reduction, and many others.

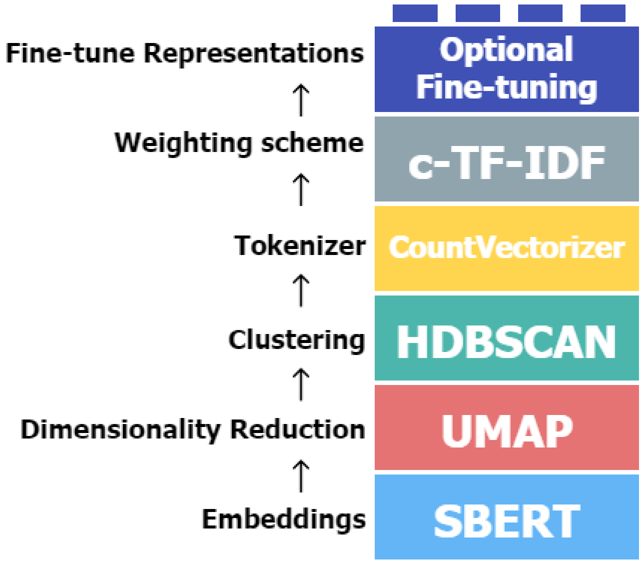

Figure 2.5: BERTopic step sequence to create topic representations

$$
\begin{gathered}
For\ a\ term\ x\ within\ class\ c:\\
W_{x,c} = \|tf_{x,c}\| \cdot \log\left(1 + \frac{A}{f_x}\right)\\
tf_{x,c} = frequency\ of\ word\ x\ in\ class\ c\\
f_x = frequency\ of\ word\ x\ across\ all\ classes\\
A = average\ number\ of\ words\ per\ class
\end{gathered}
\tag{2.1}
$$

For what regards preprocessing, it is typically not necessary as all parts of a document are important in understanding the general topic of the document. Therefore, the only text cleaning steps taken before applying BERTopic are the ones mentioned in Section 2.3.1.

The experiments performed with BERTopic were conducted between March and April 2023 on a Colab Pro Premium account with A100 GPU. Despite the fact that this environment offers a 83GB RAM option, using the whole dataset for training proved to be impossible. Moreover, the online learning implementation for BERTopic didn't support HDBSCAN clustering but just K-means; this was not a small issue, since predetermining the number of topics contained in $2,725,552$ tweets is an unfeasible task. Therefore, a subset of tweets was used for training. This subset was selected, as previously mentioned, through a simple keyword-matching approach for isolating tweets regarding the most debated political topics: (illegal) migration, climate change and the labour market. The words included in the dictionaries and sources used for building them can be found in Table 2.6.

Re-tweets where excluded before applying keyword matching and therefore before training BERTopic, as we don't expect them to be able to carry much weight in topic analysis (retweets are used to reply to other users, so their text is often shorter than that of the average tweet). The selected subset contained $432,881$ tweets.

Unfortunately, the default version of BERTopic didn't work very well on our dataset. Regardless of hyperparameter settings, for which a good number of combinations were tried, BERTopic tended to generate a very high number of topics (above 500), requiring - even when using hierarchical topic reduction as in [44] - "*labor-intensive inspection of each topic*" in order to merge them [44]. However, manually merging the topics proved difficult and yielded almost no results: it is hard to objectively determine which topics should be merged together, since there aren't any evaluation metrics that truly reflect the nuances of topics available, particularly for BERTopic (e.g. is a topic called 'redistribution' more correctly merged with one named 'wages' or with one named 'pensions'?).

Therefore, a different approach was eventually used for generating a smaller number of topics despite our dataset being very large and diverse, as well as maintaining the topics relevant to our analysis without having to manually merge them for weeks, potentially introducing errors in the process: **Guided Topic Modeling**. In this version of BERTopic, the topic modeling approach is guided by setting several seed topics to which the model will converge to - in these case, the topic representations correspond to the dictionaries in Table 2.6. This means that we set a predefined number of topic which we are sure will be in the documents analysed, i.e. in our dataset. After creating embeddings for the seeded topics, these are compared to existing document embeddings through cosine similarity, such that topic creation is nudged towards the seeded topics. The full pipeline of the process is represented in Figure 2.6

The guided topic model was trained on the 432k tweets subset generated through keyword matching, in order to avoid RAM issues occuring. It took around 3 hours in a Colab Pro Premium environment with A100 GPU. A summary of the parameter settings chosen for each submodel after careful experimentation is reported in Table 2.7.

After training, 41 topics were assigned to the subset of tweets, of which 250k were assigned to the outliers cluster. In order to check whether these were actually outliers, the outlier reduc-

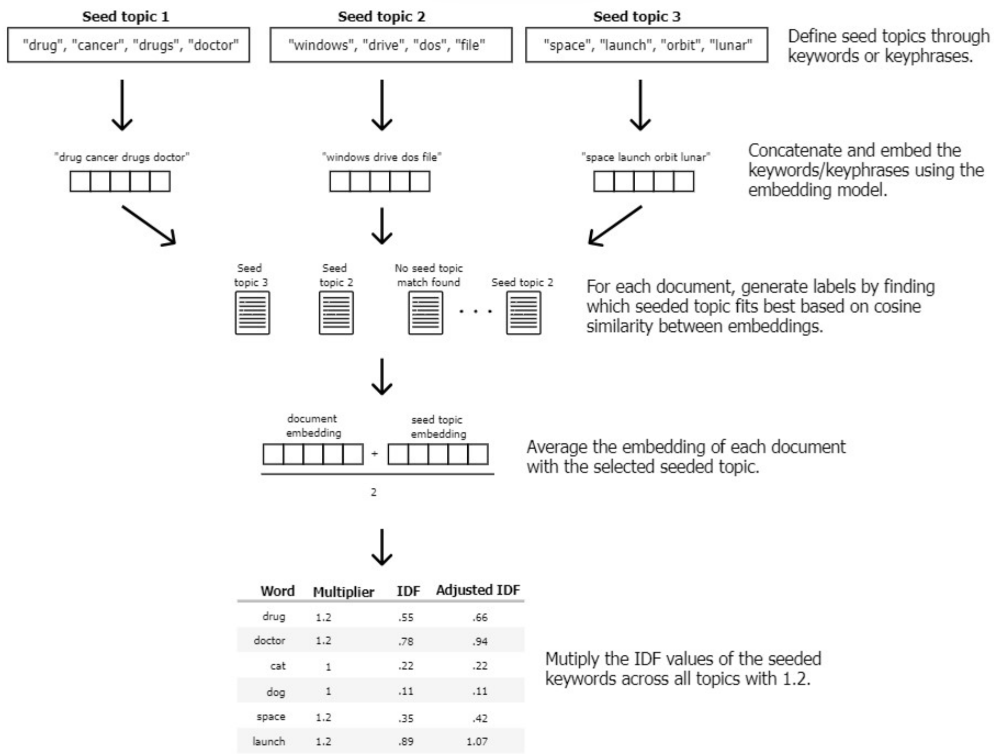

Figure 2.6: Guided BERTopic pipeline

tion method provided in the BERTopic library "*reduce_outliers*" was applied, using the 'embeddings' similarity measure with thresholds as low as $= 0.6$. While this method reduced the outliers by more than 90%, regardless of the chosen threshold, it also deteriorated the term score of the topics, i.e. the score which reflects the importance of the most frequent words in each topic for the topic itself, meaning that reducing the outliers actually "pollutes" our topic representations. Therefore, a choice was made not to apply outlier reduction. Instead, in order to have better results when transforming the dataset, several similar topics were manually merged.

After training, the guided topic model was used to assign a topic to the remaining tweets in the dataset. This process assigned topics to a much larger portion of the dataset; however, many tweets were either classified into topics of no interest for our analysis, or as outliers. Therefore, after isolating only tweets that were assigned one of the topics of interest, we are left with a subset of $425,584$ tweets containing a BERTopic-generated topic of interest. The selected topics were assigned the following labels after careful inspections of the topic representations: *'redistribution', 'job_market', 'migration', 'environment', 'eu', 'gender_equality', 'edu-*

cation', 'ukraine_russia_war', 'govt_crisis', 'police', 'healthcare'. This seleciton of topics should already be sufficient to compare the ideational and discursive-performative approaches by using topic dummies as controls, which was the ultimate goal of this task (see Section 2.3.4 for dummy variable encoding). You can observe a distribution of assigned topics to the 425k subset in Figure 2.7.

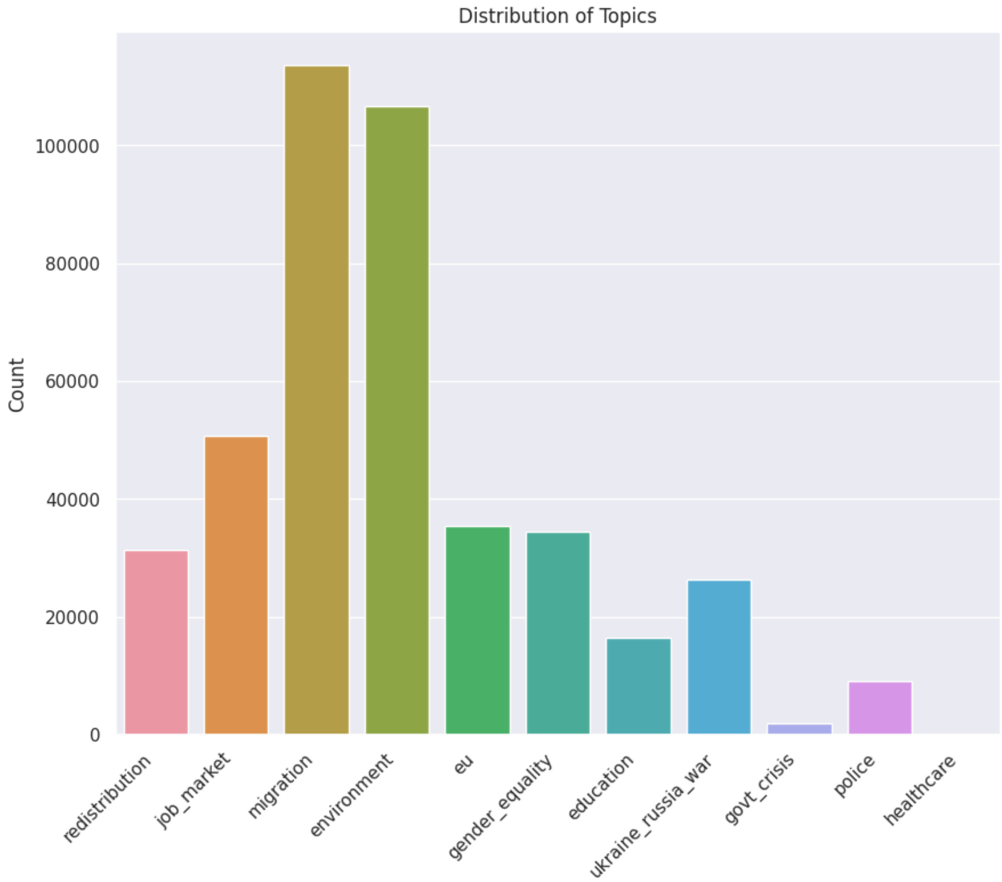

Figure 2.7: Selected topics distribution on 425k tweets subset

In case of a strong preference for assigning a larger number of tweets a topic, rather than for the quality of the topics, one could easily employ the topics obtained with the dictionaries depicted in Table 2.6, which as stated before associate a topic to a further $432,881$ tweets, i.e. the subset of tweets used for training the BERTopic model employed in this analysis.

Sentiment analysis

As already stated in the previous subsection, the introduction of BERT [50] was a pivotal moment for many fields in NLP, and sentiment analysis (SA) is no exception.

Deep-learning based SA models had already become popular in the field [51], but BERT-based models have been proven to outperform them on various SA tasks, especially when we are dealing with large datasets and short unstructured text, as is the case for this dissertation [52].

Therefore, several BERT-based methods were evaluated for the SA task and BERTsent, developed in 2022, was chosen as the most usable state-of-the-art method [53]. BERTsent is a SA-specific version of BERTweet [54] which was fine-tuned, using the *transformers* library, [55] on the SemEval 2017 corpus (39k plus tweets) .

In an effort to maintain consistency, the method was tested on the pre-labeled Twitter U.S. airline sentiment dataset [56]. A 10k sample was taken from the dataset for testing: it resulted in a 71.05% accuracy level on this dataset, with a slight preference towards negative labels. This was considered a good enough result, considering airline reviews tend to be even shorter than real tweets and usually carry a less emotional tone. Moreover, an interesting feature of this method was the fact that it outputted probabilities for each of the labels (negative, neutral, positive), thus allowing for the creation of continuous sentiment variables, as discussed in Section 2.3.3.

The method was applied on the translated English text (see Section 2.3.1) with no further pre-processing, as advised in the source paper [53]. Prediction was run in batches in a Colab Premium Pro environment with A100 GPU and took around 3 days to complete.

Variable construction from expert-based studies

Two well-known expert studies about parties' political attitudes, the Global Party Survey [57] and the Manifesto Project [26], were employed for the construction of several variables. The intended use of the constructed variables is to control for ideological, policy and other *ideational* dimensions, vis-à-vis the *discursive* dimensions we are actually trying to isolate the effects of.

Below you can find a brief description of the two expert studies the variables have been extracted from.

- **GPS Survey**: the Global Party Survey is an expert study which examines and compares ideological values, policy positions and rhetoric styles of political parties around the world [25]. In particular, the GPS 2019 was designed by Pippa Norris in order to include questions about populist rhetoric, as she explains in the seminal paper "Measuring populism worldwide" [2]. The questionnaire was carefully designed to also include questions about each expert's nationality and personal data, in order to assess whether they were qualified for classifying a specific group of parties. It is one of the few existing studies of populism spanning such a wide range of countries (163) and parties (1043), building on the knowledge of more than 1500 experts. It measures populist attitudes from an *ideational* perspective, and yet Norris clearly rejects the notion that "*populism, in itself, makes other ideological claims about [...] what should be done; instead it is a rhetoric about the rightful location of governance in this society*" [2]. Therefore, it will be interesting to study whether Norris' classification of populist parties, which is very close in theory to a discursive-performative framework, overlaps (or not) with the text-based measures proposed in this dissertation.

- **Manifesto Project**: the Manifesto Project [26] is an expert study which systematically analyses the contents of parties' political manifestos, in order to classify them according to the nature and salience of their position on certain political issues [58]. Citing from the 2023 version Codebook, the Manifesto Project includes "*election programmes of all those parties that have won one or two seats in the respective national elections to the lower house*" [59], and more generally it tries to include as many single parties as possible if they were relevant throughout a countries' electoral history. Experts from the 67 countries included have to assign codes to sentences/tokens of the manifestos, which correspond to positions on single issues. These scores are then manipulated to get cumulative statistics for each of the 1373 parties included in the dataset.

The GPS Survey variables included in the dissertation analyses are primarily those that measure policy and rhetorical attitudes (see Table 2.8 for a full list). This is because the assignment of variables' values to each tweet have been done by assigning leaders to parties in the GPS survey; many of these parties are indeed merged into coalitions by Norris (presumably for simplicity), making her ideological measures not very reliable for what regards single politicians, who may belong to a smaller party precisely due to slight ideological disagreements. Instead, it is assumed that policy and rhetorical attitudes tend to be more homogeneous across coalitions - indeed, coalitions are often formed in order to forward one or more policy issues -, and therefore we include those variables in the analyses.

The Manifesto Project database [26], on the other hand, was particularly useful for coding ideology and opposition variables. In fact, to this author's knowledge no other project has the same precision for single parties (vs coalitions) and over time, despite the fact that a lengthy manual coding of party names was required in order to adapt the *party_final* original variable labels to the labels in the dataset (see Table 2.3 for details on manually coded variables). The project is constantly updated, so it also includes parties of recent formation [59].

Additionally, the variables in the dataset are also more specific to a party's declared positions in the manifesto rather than to the perception of that party's policy choices over the years, which may be more subjective to the expert's opinion following the party historical behavior. You may find a full list of the Manifesto Project variables included in the dissertation analyses in Table 2.8.

Time-based variable construction

In order to assess how political leaders may (or may not) change their behaviour in proximity of elections, a continuous time variable was constructed that denotes the closeness of a tweet's date to the election date. The variable has been specifically coded to adapt to each country's election dates, according to the country each leader has been operating in.

The variable, which has been named *continuous_value*, behaves in the following way: it increases continually from 0 to 1 from 90 days before the election date up until the day of elections, where *continuous_value*$= 1$; afterwards, it decreases continually from 1 to 0, up until $90+1$ days after the election, where *continuous_value*$= 0$. Note that when we say the variable's value increases continually, the meaning is that the value increases by $1/90$, i.e. approximately by 0.01 each day, and vice-versa for the post-electoral period.

This was an empirical choice made to monitor a tweet's closeness to elections and as such is not substantiated by any specific references. You can observe a graphical representation of the behavior of the variable for Italy in Figure 2.8 below.

2.3.4 Feature engineering

This section is dedicated to the description of all further steps taken to engineer features after extraction, for all those cases where it was necessary in order to conduct the analyses described in Chapter 3.

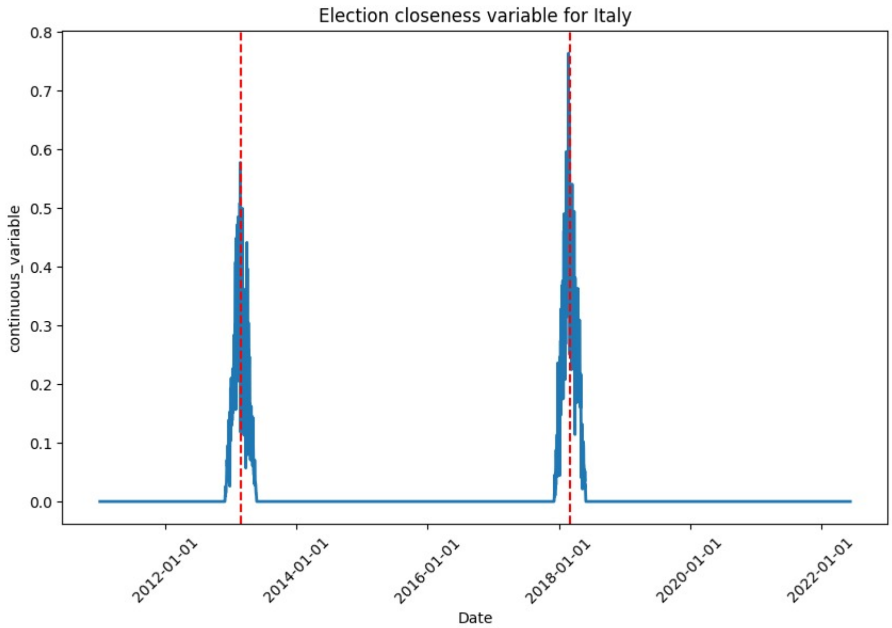

Figure 2.8: Graphical depiction of *continuous_value* behaviour around elections in Italy (elections corresponding to vertical red lines)

One-hot enconding

In the process of *one-hot encoding* for categorical variables, each unique category in the original variable is transformed into a separate binary variable (dummy variable). For each observation, the dummy variable corresponding to its category is set to 1, while the value of dummy variables corresponding to other categories are set to 0. This encoding allows categorical variables to be included in mathematical models and algorithms that require numerical inputs.

In our case, some categorical variables (e.g. topic variables) only had one category, but the process is still the same.

You can find a list of all variables that have been converted into dummies, as well as a brief description, below.

- Topic variables:
  - *redistribution*: the tweet talks about systems of wealth redistribution, such as taxes, subsidies, etc.
  - *job_market*: the tweet talks about the job market, i.e. about wages, unemployment, etc.
  - *migration*: the tweet talks about (illegal) migration, arrivals, and the like.
  - *environment*: the tweet talks about topics related to climate change and environmental protection

- *eu*: the tweet talks about EU policies and institutions
- *gender_equality*: the tweet talks about topics surrounding gender equality, such as the gender wage gap, parenting roles, violence against women, etc.
- *education*: the tweet talks about the education system and related issues
- *ukraine_russia_war*: the tweet talks about conflicts between Russia and Ukraine, including the most recent full-scale war between the two countries
- *govt_crisis*: the tweet talks about problems of sovereignty and the (supposed) crisis of democracy in Western countries
- *police*: the tweet talks about the police system and security issues
- *healthcare*: the tweet talks about the healthcare system and related issues

- Ideology variables
  - *RL*: radical left dummy variable, corresponds to LEF in Manifesto Project 2023
  - *L*: left dummy variable, corresponds to SOC in Manifesto Project 2023
  - *C*: centre dummy variable, corresponds to CHR + LIB in Manifesto Project 2023
  - *R*: right dummy variable, corresponds to CON in Manifesto Project 2023
  - *RR*: radical right variable (1=RR), corresponds to NAT in Manifesto Project 2023
  - *ECO*: environmental party variable, corresponds to ECO in Manifesto Project 2023
- Country variables
  - *Included countries*: Australia, Austria, Belgium, Brazil, Canada, Czech Republic, Denmark, Finland, France, Germany, Hungary, Italy, Mexico, Netherlands, Norway, Poland, Portugal, Slovakia, Slovenia, Spain, Sweden, United Kingdom, United States
  - *NB*: The corresponding dummy variable for each country is equal to 1 when the tweet was posted by a leader coming from that country in the dataset, 0 otherwise.

Feature scaling

In order to ensure consistency in the interpretation of results, all continuous variables have been normalized to be between 0 and 1, using the *scikit-learn* library's tool *MinMaxScaler* for scaling [49].

Since the only variables that weren't already in this range (or dummies) are the variables extracted thorough the GPS-2019 database [25] and the 2023 version of the Manifesto Project related to policy dimensions [26], the reader can refer to Table 2.8 for a full list of normalized continuous variables and their corresponding description.

Continuous sentiment labels generation

In order to measure sentiment in a continuous way, an empirical method was applied to the output of BERTsent (see Section 2.3.3 for more details). Very simply, the label with the highest probability (i.e. the label the model would have assigned to the tweet) has been substituted with its corresponding probability, as shown in Equation 2.2 below.

$$\begin{aligned}
&For\ x: \\
&\qquad continuous - sentiment_x = |pred_x| \cdot labelix_x \\
&\qquad x = tweet\ analyzed \\
&\qquad pred_x = predicted\ probability\ for\ label \\
&\qquad labelix_x = index\ of\ assigned\ sentiment\ label,\ one - hot\ encoded
\end{aligned} \tag{2.2}$$

The method was applied only to negative and positive sentiment labels, as neutral sentiment does not appear to be relevant to populist rhetoric under the chosen framework, generating the two variables:

- *pos_sent_cont*: sentiment probability for positive label, when label was assigned
- *neg_sent_cont*: sentiment probability for negative label, when label was assigned

This empirical method allows sentiment labels to be included in feature importance studies in a more precise way with respect to simple one-hot encoded labels (see Section 2.3.3 for more details on SA).

GPS populist rhetoric variables aggregation

In order to evaluate the similarities and differences between party-level and leader-level studies of populism, as well as the nuances in variables' effects due to the different (expert-based vs text-based) methods for measuring populism, two variables regarding populist rhetoric from the GPS 2019 study [25] have been aggregated by taking their mean into a single one, which has been named *Norris_pop*.

These variables are, specifically:

- V_8: party strongly favors populist rhetoric, i.e. a language that typically challenges the legitimacy of established political institutions and emphasizes that the will of the people should prevail, measured on a $0 - 10$ scale (10 = strong preference for populist rhetoric)

- V_9: party gives great importance to the use of populist rhetoric, measured on a $0 - 10$ scale (10 = great importance given to the use of populist rhetoric)

2.3.5 Index construction with AHP: populist online communication logic

This dissertation's main peculiarity, other than the size of the dataset constructed, is the proposition to use methods coming from different (sometimes apparently unrelated) fields in order to capture the enigmatic and complex nature of populist online communication logic from a text analysis perspective. Thus, it should come as no surprise that the final step in data preparation, i.e. the construction of indexes relative to each of the four dimensions of populist online communication logic described in Section 1.3, was inspired and constructed through conversations with knowledgeable colleagues coming from the most diverse fields (economics, marketing, engineering...).

In fact, despite an extensive literature review of empirical studies of populism, it doesn't seem like a work similar to that carried out in this dissertation has been conducted before. As previously argued, the reason probably lies in the fact most studies conducted under the discursive-performative approach tend to have used survey research designs or other expert-based designs, rather than NLP techniques and text analysis [17]. Therefore, keeping in mind the time and computational constraints a master's thesis has with respect to longer research projects, a methodological choice was made to construct the indexes empirically, rather than trying to reproduce other studies that employ completely different processes for feature extraction.

A first attempt was made by summing up variables of interest regarding the four dimensions and simply taking their mean. Despite the normalization procedure described in the previous paragraphs, high variability is still present in the variables' distributions and general nature, due to the very diverse methods used for feature extraction. As was likely, the method proved to yield very unsatisfactory results.

The Analytic Hierarchical Process (AHP), a decision-making process invented in the 1980s by Thomas Saaty [60], was eventually selected for constructing weighted indexes. While this choice may seem arbitrary, the results shown in Chapter 3 paint a different picture: not only the results are incomparably better than the ones previously obtained; by being consistent with the theoretical discussion of populism conducted in Chapter 1, they actually point to the idea that this method could be fruitfully employed as a decision strategy by social science researchers.

AHP is one of the most popular and widely employed multicriteria methods. It is used for integrating the processes of rating alternatives and of finding the most relevant one, when designing a decision process including multiple criteria with respect to an overall goal.

The application of the methodology consists of establishing the importance weights to be associated to the criteria in defining the overall goal. This is done by comparing the criteria pairwise. Let us consider two criteria, C_j and C_k. The researcher expresses a graded comparative judgment about the pair in terms of the relative importance of C_j over C_k with respect to the goal. The comparative judgement is captured on a semantic scale (equally important/moderately more important/strongly important and so on, according to Saaty's scale [60], and is converted into a numerical integer value a_{jk}. The relative importance of C_k over C_j is defined as its reciprocal, i.e., $a_{kj} = 1/a_{jk}$. A reciprocal pairwise comparison matrix A is then formed using a_{jk}, for all j and k. Note that $a_{jj} = 1$. It has been generally agreed that the weights of criteria can be estimated by finding the principal eigenvector w of the matrix A, that we will call w. When the vector w is normalized, it becomes the vector of priorities of the criteria with respect to the goal; λ_{max} is the largest eigenvalue of the matrix A, and the corresponding eigenvector w contains only positive entries. Using similar procedures, the weights of alternatives with respect to each criterion are computed. Then, the overall weights of alternatives are computed using the weighted summation:

$$Overall\ weight\ of\ alternative\ i = \sum_j (Weight\ of\ alternative\ i\ with\ respect\ to\ C) \tag{2.3}$$

AHP is popular for its simplicity, flexibility, intuitive appeal, and ability to mix quantitative and qualitative criteria in the same decision framework[61]. Considering these were actually the characteristics needed in the case of this dissertation's index construction and that no specific text-analysis literature existed to help in the decision, this seemed like a robust way of measuring criteria weights in the various dimensions.

The variables used for constructing the indexes of the four dimensions of populist communication logic, and the corresponding index weights constructed through AHP, are described in Table 2.9. An example of the decision matrix A, its normalized version and the extracted weights for the *people_centrism _index* is depicted in Figure 2.9.

Original importance matrix:

	I	WE	INGROUP	PURITY	PEOPLE
I	1.00	0.17	0.33	0.50	0.25
WE	6.00	1.00	2.00	3.00	3.00
INGROUP	3.00	0.50	1.00	3.00	2.00
PURITY	2.00	0.33	0.33	1.00	0.33
PEOPLE	4.00	0.33	0.50	3.00	1.00
SUM	16.00	2.33	4.17	10.50	6.58

Normalized matrix for weight assignment:

	I	WE	INGROUP	PURITY	PEOPLE	weight
I	0.06	0.07	0.08	0.05	0.04	0.06
WE	0.38	0.43	0.48	0.29	0.46	0.40
INGROUP	0.19	0.21	0.24	0.29	0.30	0.25
PURITY	0.13	0.14	0.08	0.10	0.05	0.10
PEOPLE	0.25	0.14	0.12	0.29	0.15	0.19

Figure 2.9: Graphical depiction of AHP weight assignment process for index *people_centrism_IX*

Finally, here is a summary of the four dimensions of populist online communication described in Section1.3 and the variables used to construct their data-based counterpart:

(i) *People-centrism*: putting *the people*, however identified (through nationality, ethnicity, etc.) at the center of the discourse; **variables included**: we, i, IngroupVirtue, PurityVirtue, contains_peopl* (NB: this last one is a very simple variable that assigns value 1 to any tweet containing regex "peopl")

(ii) *Anti-elitism*: expressing negative feelings towards *the elite*, however identified (political, economic, etc.); **variables included**: they, power, shehe, IngroupVice, PurityVice, contains_elit* (NB: this last one is a very simple variable that assigns value 1 to any tweet containing regex "elit")

(iii) *Identification of leader with the people*: framing the discourse in such a way that the speaker, while they are technically part of the (political) elite, is actually perceived as a member of *the people*; **variables included**: Authentic, swear, we

(iv) *Antagonism*: using discourse to create feelings of antagonism and conflict, from *the people* towards *the elite* but also from *the people* towards an undefined *other* who is not included in the in-group (e.g. immigrants, economic elites...); **variables included**: allnone, conflict, emo_anger

2.4 Index validation and analysis: methods

2.4.1 XGBoost

XGBoost [62] is a powerful and well-suited technique for identifying feature importances in high-dimensional datasets characterized by a mixture of continuous variables and dummy variables.

XGBoost operates by constructing an ensemble of decision trees through an iterative boosting process. It optimizes a loss function by sequentially adding trees, where each subsequent tree corrects the errors made by the previous ones; it then uses gradient descent to minimize the loss function and performs regularization to prevent overfitting, resulting in efficient and accurate predictions. By assigning weights to individual features and analyzing their contribution to the overall model performance, XGBoost provides insights into the relative importances of different features in any high-dimensional dataset.

With a dataset consisting of 2.4 million rows and many different variables, XGBoost's ability to handle large-scale datasets and diverse data types makes it an ideal choice. By leveraging gradient boosting algorithms and ensemble learning, XGBoost can effectively capture complex relationships and interactions between predictors, enabling the identification of the most influential features.

2.4.2 Linear regression

A simple regression with 10-fold cross validation procedure was conducted on the variables listed at the end of the previous paragraph, using as the dependent variable the AHP-weighted index made up of all four populist dimensions. The regression shown in Table3.2 was corroborated by also performing the same 10-fold cross-validated regression on all four dimensions separately, which yielded almost identical results.

Variable	Description	Dictionary used
IngroupVirtue	Extent to which text contains words related to moral foundation of loyalty and group identity, which includes the valuing of loyalty, patriotism, and the protection of one's group	MFD
IngroupVice	Extent to which text contains words related to disregard for group identity and loyalty, such as betraying or harming other members of the same group	MFD
PurityVirtue	Extent to which text contains words related to moral foundation of sanctity and purity, which includes the valuing of cleanliness, orderliness, and the avoidance of impurity or contamination in both physical and moral senses	MFD
PurityVice	Extent to which text contains words related to disregard for cleanliness and purity, as well as the promotion of impurity and contamination	MFD
we	Text variable measuring amount of first person plural pronouns in text	LIWC-2022
i	Text variable measuring amount of first person singular pronouns in text	LIWC-2022
they	Text variable measuring amount of third person plural pronouns in text	LIWC-2022
power	Text variable measuring amount of words that are associated with power, such as ”lead,” ”dominate,” or ”control”	LIWC-2022
shehe	Text variable measuring amount of third person singular pronouns in text	LIWC-2022
Authentic	Text variable measuring amount of words that are associated with authenticity, such as ”real,” ”genuine,” or ”honest”	LIWC-2022
swear	Text variable measuring amount of words that are considered vulgar or offensive	LIWC-2022
allnone	Text variable measuring amount of words that are associated with allness or none, such as ”always,” ”never,” or ”everything”	LIWC-2022
conflict	Text variable measuring amount of words that are associated with conflict or disagreement, such as ”argument,” ”dispute,” or ”fight”	LIWC-2022
emo_anger	Text variable measuring amount of words that are associated with anger or frustration, such as ”angry,” ”frustrated,” or ”irritated”	LIWC-2022
achieve	Text variable measuring amount of words in text reflecting sense of achievement (proxy for group efficacy)	LIWC-2022
Clout	Text variable measuring amount of words that are associated with power and authority, such as ”leader,” ”expert,” or ”influential”	LIWC-2022
comm	Text variable measuring amount of words that are associated with communication, such as”‘ ”speak,” ”listen,” or ”communicate”	LIWC-2022
prosocial	Text variable measuring amount of words that are associated with positive social behavior, such as helping or cooperating with others	LIWC-2022
lack	Text variable measuring amount of words that are associated with lacking, such as ”deficient,” ”deprived,” or ”missing”	LIWC-2022
Affect	Text variable measuring amount of words that are associated with affective experiences, such as ”excited,” ”calm,” ”tense,” or ”relaxed”	LIWC-2022
emotion	Text variable measuring amount of words that are associated with specific emotions or feelings, such as ”happy,” ”sad,” or ”angry”	LIWC-2022
want	Text variable measuring amount of words that are associated with wants, such as ”crave,” ”desire,” or ”long for”	LIWC-2022
feeling	Text variable measuring amount of words that are associated with sensory experiences and bodily sensations, such as ”touch,” ”pain,” and ”warmth”	LIWC-2022
cause	Text variable measuring amount of words that are associated with explaining why things happen, such as ”cause,” ”effect,” or ”because”	LIWC-2022
emo_anx	Text variable measuring amount of words that are associated with worry or fear, such as ”anxious,” ”worried,” or ”nervous”	LIWC-2022
need	Text variable measuring amount of words that are associated with needs, such as ”essential,” ”necessary,” or ”must-have”	LIWC-2022
risk	Text variable measuring amount of words that are associated with risk, such as ”danger,” ”hazard,” or ”chance”	LIWC-2022

Table 2.4: Variables extracted with dictionary-based method through LIWC-2022 software

Variable	Description	Dictionary used
HarmVirtue	Extent to which text contains words related to moral foundation of care and compassion, which includes the valuing of kindness, gentleness, and the prevention of harm to oneself and others	MFD
HarmVice	Extent to which text contains words related to disregard for the well-being of oneself and others, as well as the intentional infliction of harm	MFD
FairnessVirtue	Extent to which text contains words related to moral foundation of justice and fairness, which includes the valuing of equality, impartiality, and the protection of individual rights	MFD
FairnessVice	Extent to which text contains words related to disregard for fairness and justice, as well as the exploitation of others for personal gain	MFD
AuthorityVirtue	Extent to which text contains words related to moral foundation of respect for authority and tradition, which includes valuing of obedience, respect for authority, and maintenance of social order	MFD
AuthorityVice	Extent to which text contains words related to disregard for authority and tradition, as well as the promotion of chaos and anarchy	MFD
MoralityGeneral	Overall level of moral language used in a text, reflects the degree to which moral concepts are present and salient in the text	MFD

Table 2.5: Variables extracted with dictionary-based method through LIWC-2022 software - Continued

Topic name	Description	Dictionary	Source
'migration_topic'	tweets talking about migration, with a focus on illegal migration phenomena and adverse reactions	'immigrant', 'immigrants', 'illegal', 'million', 'thousands', 'stay', 'terrorist', 'migrant', 'migrants', 'economic migrant', 'economic migrants', 'asylum seeker', 'asylum seekers', 'refugee', 'refugees', 'influx', 'wave'	list of frequently used terms related to migration based on analysis of news articles and social media data obtained from the Migration Observatory at the University of Oxford [41]
'environment_topic'	tweets talking about environmental protection, climate change and related issues	'climate', 'climate change', 'plastic', 'plastic pollution', 'pollution', 'air pollution', 'climate crisis', 'crisis', 'palm oil', 'single use plastic', 'plastic waste', 'waste', 'renewable', 'renewables', 'renewable energy', 'energy', 'fossil fuel', 'fossil fuels', 'fuel', 'fuels', 'clean energy', 'greenhouse gas', 'plastic packaging', 'emissions', 'extreme weather', 'offshore drilling'	selection of most common words in dictionaries contained in a recent publication on the subject (see [42])
'work_topic'	tweets talking about the labour market and related issues	'work', 'income', 'basic income', 'universal income', 'taxes', 'tax', 'job', 'jobs', 'job market', 'labour market', 'welfare', 'policies', 'welfare policies', 'poor', 'rich', 'progressive taxation', 'taxation', 'redistribution', 'trickle down', 'lifestyle', 'poverty', 'pension', 'social security', 'retirement', 'retire'	as no specific source existed for dictionaries on this topic, I took inspiration from observing trends on social media through the Keyhole tool [43]

Table 2.6: Dictionaries for topic analysis pre-processing

Sub-model	Chosen parameter setting
UMAP	• n_components $= 5$ • min_dist $= 0.2$ • n_neighbours $= 15$ • metric = cosine
CountVectorizer	• min_df $= 15$ • ngram_range $= (1, 2)$
C-TFIDF	• default parameter setting maintained
HDBSCAN	• min_cluster_size $= 40$ • min_samples $= 50$ • gen_min_span_tree = True • metric = euclidean
BERTopic	• min_topic _size $= 500$

Table 2.7: Guided BERTopic model: final parameters settings for training

Table 2.8: Expert-based variables description and sources

Variable Name	Variable Description	Source
V4_Scale	Party economic left/right, transformed into a 0-1 scale* (1 = extreme right)	GPS 2019
V6_Scale	Party social liberalism/conservatism, transformed into a 0-1 scale* (1 = extreme conservatism)	GPS 2019
V8_Bin	Party's favoring of populist vs pluralist rhetoric, transformed into 0/1 dummy* (1 = strongly favors populist rhetoric)	GPS 2019
V8_Scale	Party's favoring of populist vs pluralist rhetoric, transformed into a 0-1 scale* (1 = strongly favors populist rhetoric)	GPS 2019
V9	Importance placed by the party on populist rhetoric, transformed into a 0-1 scale* (1 = great importance)	GPS 2019
V10	Party's opposition to environmental protection, transformed into a 0-1 scale* (1 = strongly opposes environmental protection)	GPS 2019
V11	Party's favoring of public spending vs reduced taxation, transformed into a 0-1 scale* (1 = strongly favors reduced taxation)	GPS 2019
V12	Party's opposition to environmental protection, transformed into a 0-1 scale* (1 = strongly opposes environmental protection)	GPS 2019
V13	Party's favoring of multilateralism over nationalism, transformed into a 0-1 scale* (1 = strongly favors multilateralism)	GPS 2019
V14	Party's opposition to women's rights, transformed into a 0-1 scale* (1 = strongly opposes women's rights)	GPS 2019
V15	Party's opposition to ethnic minority rights, transformed into a 0-1 scale* (1 = strongly opposes ethnic minority rights)	GPS 2019
V18	Emphasis in party rhetoric on politicians following vs leading public opinion, transformed into 0-1 scale* (1 = strongly emphasizes that politicians should lead public opinion)	GPS 2019
V19	Emphasis in party rhetoric on people vs leaders deciding on important issues, transformed into 0-1 scale* (1 = strongly emphasizes that leaders should decide on important issues)	GPS 2019
V20	Emphasis in party rhetoric on politicians being honest vs corrupt, transformed into 0-1 scale* (1 = strongly emphasizes that most politicians are corrupt)	GPS 2019
V21	Party's opposition to checks and balances on executive power, transformed into 0-1 scale* (1 = strongly opposes checks and balances on executive power)	GPS 2019
welfare	Mean value of per503 + per304	Manifesto Project 2023
per503	Equality - Concept of social justice and the need for fair treatment of all people. This may include: Special protection for underprivileged social groups; Removal of class barriers; Need for fair distribution of resources; The end of discrimination (e.g., racial or sexual discrimination)	Manifesto Project 2023
per504	Welfare State Expansion - Favourable mentions of need to introduce, maintain or expand any public social service or social security scheme. This includes, for example, government funding of: Health care, Child care, Elder care and pensions, Social housing	Manifesto Project 2023
per304	Political Corruption - Need to eliminate political corruption and associated abuses of political and/or bureaucratic power. Need to abolish clientelist structures and practices	Manifesto Project 2023
per501	Environmental Protection - General policies in favor of protecting the environment, fighting climate change, and other "green" policies. For instance: General preservation of natural resources; Preservation of countryside, forests, etc.; Protection of national parks; Animal rights. May include a great variance of policies that have the unified goal of environmental protection	Manifesto Project 2023
per606	Civic Mindedness: Positive - Appeals for national solidarity and the need for society to see itself as united. Calls for solidarity with and help for fellow people, familiar and unfamiliar. May include: Favorable mention of the civil society; Decrying anti-social attitudes in times of crisis; Appeal for public spiritedness; Support for the public interest	Manifesto Project 2023
per305_3	Political Authority - References to the manifesto party's competence to govern and/or other party's lack of such competence. Also includes favorable mentions of the desirability of a strong and/or stable government in general	Manifesto Project 2023
per202_4	Direct Democracy: Positive - Favorable mentions of the system of direct democracy, in particular in contrast to representative democracy. This includes the call for the introduction and/or extension of referenda, participatory budgets and other forms of direct democracy	Manifesto Project 2023

people_centrism_IX	we	0.2
	i	0.1
	IngroupVirtue	0.3
	PurityVirtue	0.1
	contains_peopl*	0.3
pop_as_ppl_IX	Authentic	0.4
	swear	0.2
	we	0.4
anti_elitism_IX	they	0.2
	power	0.2
	IngroupVice	0.133
	PurityVice	0.067
	contains_elit*	0.3
antagonize_IX	allnone	0.389
	emo_anger	0.222
	conflict	0.389
total_pop_pop_IX	pop_as_ppl_IX	0.329
	pop_as_ppl_IX	0.263
	anti_elitism_IX	0.219
	antagonize_IX	0.188

Table 2.9: AHP-generated weights of internal populist communication logic indexes and of the total index combining all four dimensions

3 Results

3.1 Feature selection with XGBoost

In our specific case, we apply XGBoost to the dataset after clearing it of any rows containing missing values, in order to select which features should be kept for the final linear regression.

The dataset used for XGBoost contains $2,398,207$ tweets and retweets. All variables previously described are included, and XGBoost is ran with a grid-search on the baseline "ideational populism" variable constructed from GPS (*Norris_pop*), on the "total" populist online communication logic index (*total_pop_IX*), and on the other four single dimensions this index has been constructed with (*people_centrism_IX, anti_elitism_IX, pop_as_ppl_IX, antagonize_IX*).

Care is taken not to include any variable that are included in the indexes being regressed on in order to avoid polluting the results.

Table 3.1 shows the optimal grid-search parameters for each of these XGBoost regressive processes, found through 5-fold cross-validation procedure. Python library *xgboost* [63] was used, alongside *scikit-learn*'s implementation of grid-search with cross validation *GridSearchCV* [49]. The processes were ran in Colab Pro Premium environment with A100 GPU and took around 6 hours to complete.

As regressions were subsequently undertaken to further corroborate the four indexes and the cumulative index, we don't report here the feature importance plots for XGBoost. However, you can find a list of all features selected for regression by pooling the results of the 6 XGBoost procedures at the beginning of Section 3.2.

Dependent Variable	**Hyperparameters**	**Values**
Norrispop	learning_rate max_depth n_estimators	0.05, 0.1, **0.15**, 0.2 **5**, 6, 7 75, 100, 125, **130**
people_centrism_IX	learning_rate max_depth n_estimators	**0.05**, 0.1, 0.15, 0.2 5, 6, **7** 75, 100, 125, **130**
pop_as_ppl_IX	learning_rate max_depth n_estimators	**0.05**, 0.1, 0.15, 0.2 5, 6, **7** 75, 100, 125, **130**
antagonize_IX	learning_rate max_depth n_estimators	0.05, **0.1**, 0.15, 0.2 5, 6, **7** **75**, 100, 125, 130
anti_elitism_IX	learning_rate max_depth n_estimators	0.05, 0.1, **0.15**, 0.2 **5**, 6, 7 **75**, 100, 125, 130
total_pop_IX	learning_rate max_depth n_estimators	**0.05**, 0.1, 0.15, 0.2 5, 6, **7** 75, 100, 125, **130**

Table 3.1: Grid-search hyperparameters for XGBoost feature importance and selection processes: optimal choices in bold

3.2 Preliminary linear regressions

Due to time constraints, it was preferred to use a simple method in order to validate the indexes, as described in Section2.4.2. You can find a list of all included variables and controls below.

- **Topic controls**: all BERTopic-generated topic variables, i.e. redistribution, job_market, migration, environment, eu, gender_equality, education, ukraine_russia _war, govt _crisis, police, healthcare

- **Country controls**: dummy variables corresponding to each of the 23 countries in the dataset

- **Policy controls**: V10, V11, V12, V13, V14, V15, V18, V19, V20, V21 (see Table2.8 for descriptions)
- **Ideology controls**: RL, L, C, R, RR, ECO (see Section2.3.3 for mroe details)
- **Variables of interest**: focuspast, focuspresent, focusfuture, emo_anx, HarmVirtue, HarmVice, FairnessVirtue, FairnessVice, AuthorityVirtue, AuthorityVice, MoralityGeneral, pos_sent_cont, neg_sent_cont, male, opposition (see Table 2.4 and Table 2.5 for descriptions)

Results from the regression, shown in Table3.2 are commented briefly, since this author is convinced a much deeper analysis would be required to draw any final conclusions about the interactions between our variables of interest and the constructed indexes. However, the 5 highest coefficients in terms of absolute value are discussed below, in the hope this discussion could be useful for future works.

- **FairnessVice**: the coefficient related to this variable signals that its effect on the cumulative populist online communication logic index is rather large (0.32 unit change). This means that, if the text contains words that indicate a disregard for fairness and justice, the index is likely to be higher for that tweet.
- **focuspast**: the coefficient related to this variable signals that its effect on the cumulative populist online communication logic index is, again, rather large (0.31 unit change). This means that, if the text contains words that indicate a focus on past times, the index is likely to be higher for that tweet.
- **HarmVice**: the coefficient related to this variable signals that its effect on the cumulative populist online communication logic index is, again, rather large (0.25 unit change). This means that, if the text contains words that indicate a disregard for the wellbeing of others and a preference for the infliction of harm, the index is likely to be higher for that tweet.
- **focusfuture**: the coefficient related to this variable signals that its effect on the cumulative populist online communication logic index is moderate (0.21 unit change). This means that, if the text contains words that indicate a focus on past times, the index is likely to be higher for that tweet.

- **AuthorityVirtue**: the coefficient related to this variable signals that its effect on the cumulative populist online communication logic index is, again, rmoderate (-0.19 unit change). This means that, if the text contains words that indicate a preference for the respect of authority and tradition, the index is less likely to be higher for that tweet.

Table 3.2: Linear Regression Output with Y corresponding to total index of populism *total_pop_IX*. Topic, ideology and country controls have been included in the regression, but are not printed in the output. The same regression, conducted on the four single-dimension indexes that make up *total_pop_IX*, yielded almost identical results. Variables contained in the indexes (in this case, also single-dimension indexes themselves) have been excluded to avoid collinearity. ** symbol corresponds to significance with $\alpha = 0.005$

Variable	Coefficient	Absolute Coefficient	Standard Error	t-value	p-value	Significance
V18	-0.022786	0.022786	1.379854×10^{-7}	-165134.983911	0.0	**
V19	0.025285	0.025285	1.271942×10^{-7}	198789.219276	0.0	**
V20	-0.013708	0.013708	1.096916×10^{-7}	-124969.225238	0.0	**
V21	0.002684	0.002684	1.258547×10^{-7}	21328.304522	0.0	**
focuspast	0.312348	0.312348	8.504597×10^{-7}	367269.970846	0.0	**
focuspresent	0.185807	0.185807	6.603565×10^{-7}	281374.424501	0.0	**
focusfuture	0.214016	0.214016	9.948681×10^{-7}	215120.076206	0.0	**
emo_anx	0.055466	0.055466	4.502428×10^{-6}	12319.127301	0.0	**
HarmVirtue	-0.160929	0.160929	2.098004×10^{-6}	-76705.741933	0.0	**
HarmVice	0.252207	0.252207	2.378484×10^{-6}	106036.879150	0.0	**
FairnessVirtue	0.036609	0.036609	2.904150×10^{-6}	12605.855572	0.0	**
FairnessVice	0.319783	0.319783	7.271974×10^{-6}	43974.684268	0.0	**
AuthorityVirtue	-0.192402	0.192402	1.969737×10^{-6}	-97679.031239	0.0	**
AuthorityVice	0.126292	0.126292	5.234979×10^{-6}	24124.639223	0.0	**
MoralityGeneral	0.021802	0.021802	1.801912×10^{-6}	12099.451079	0.0	**
pos_sent_cont	0.009146	0.009146	1.630905×10^{-7}	56077.351048	0.0	**
neg_sent_cont	0.006494	0.006494	2.016911×10^{-7}	32197.891571	0.0	**
male	0.002811	0.002811	7.424266×10^{-8}	37868.411277	0.0	**
continuous_value	0.010654	0.010654	2.672748×10^{-7}	39861.711859	0.0	**
tone_pos	-0.047240	0.047240	5.329732×10^{-7}	-88634.875293	0.0	**
tone_neg	-0.034844	0.034844	9.247757×10^{-7}	-37678.632241	0.0	**
polite	-0.054975	0.054975	1.206755×10^{-6}	-45555.720146	0.0	**
tweet_len	0.001547	0.001547	3.397455×10^{-9}	455331.025824	0.0	**
Norris_pop	0.004691	0.004691	1.034049×10^{-7}	45362.598769	0.0	**
opposition	-0.005352	0.005352	7.314836×10^{-8}	-73164.245134	0.0	**

4

Conclusion

4.1 Conclusions

This dissertation has aimed to set a baseline for researchers who are interested in exploring political science and discourse analysis questions while exploiting the new, exciting possibilities that emerged during the *fourth industrial revolution*, particularly those offered by the combined use of readily-available social media data and NLP techniques. The dissertation is built upon a 3+million tweets dataset collected by researchers at the Ing. Rodolfo De Benedetti Foundation in 2022 [27], making this an exciting project also from the point of view of high-dimensional data handling. A mixture of methods (manual coding, expert surveys, dictionary methods, BERT-based methods...) was employed in feature extraction, to allow as much information as possible to be captured from the dataset and to perform that fruitful fertilization within/between research fields that has been proposed by field experts as an optimal direction for researchers of populism [4]. While no solid conclusion has yet been reached, due to the limited time and computational resources a Master's thesis is doomed to have, the preliminary regression results shown in Section3.2 do point towards some success in terms of correspondence with literature about populism under the discursive-performative approach.

4.2 Further directions

Many directions could be taken to improve the anlaysis here contained, but for the sake of synthesis I will only list a couple. The first, obvious, one is to add a time dimension to the regression or to directly employ a time-series model to analyze the data, such as ARIMAX. Moreover, the textual variables relative to the LIWC-2022 dictionary and to the Moral Foundations Dictionary, instead of being directly extracted through the LIWC software, could be extracted by exploiting the guided Bertopic methodology described in Section2.3.3, as to capture more information by means of BERT-based embeddings of the dictionary word vectors. Furthermore, including or excluding countries from analyses, conducting a different data mining procedure - for example one that regards the *demand-side* of populism, i.e. its supporters -, could leade to interesting results.

References

[1] E. Laclau, *On populist reason.*, 1st ed. Verso, 2005.

[2] P. Norris, "Measuring populism worldwide." *Party politics*, vol. 26, no. 6, pp. 697–717, 2020.

[3] K. Schwab, "The fourth industrial revolution and the future of jobs." *Foreign Affairs*, vol. 94, no. 4, pp. 80–90, 2015.

[4] M. Rooduijn, "State of the field: How to study populism and adjacent topics? a plea for both more and less focus." *European Journal of Political Research*, vol. 58, no. 1, pp. 362–372, 2019.

[5] Elisa Tremolada. Frequency of the word 'populism' in english books, 1800-2019. generated with google ngram viewer. [Online]. Available: https://books.google.com/ngrams/graph?content=populism&year_start=1800&year_end=2019&corpus=en-2019&smoothing=3

[6] C. Mudde and C. R. Kaltwasser, *Populism: A very short introduction.*, 1st ed. Oxford University Press., 2017.

[7] F. N. Engesser, S. and A. O. Larsson, "Populist online communication: Introduction to the special issue." *Information, communication society*, vol. 20, no. 9, pp. 1279–1292, 2017.

[8] D. L. Altheide, "Media logic, social control, and fear," *Communication theory*, vol. 23, no. 3, pp. 223–238, 2013.

[9] R. Bracciale and A. Martella, "Define the populist political communication style: the case of italian political leaders on twitter." *Information, communication society*, vol. 20, no. 9, pp. 1310–1329, 2017.

[10] J. Van Dijck, *The culture of connectivity: A critical history of social media.*, 1st ed. Oxford University Press., 2013.

[11] E. Ortiz-Ospina, "The rise of social media," *Our World in Data*, 2019, https://ourworldindata.org/rise-of-social-media.

[12] S. Waisbord and A. Amado, "Populist communication by digital means: presidential twitter in latin america." *Information, communication society*, vol. 20, no. 9, pp. 1330–1346, 2017.

[13] J. Bartlett, "Populism, social media and democratic strain." in *European populism and winning the immigration debate*, C. N. Sandelind, Ed. Fores, 2014, pp. 99–114.

[14] F. Kostelka and A. Blais, "The generational and institutional sources of the global decline in voter turnout." *World politics*, vol. 73, no. 4, pp. 629–667, 2021.

[15] B. Krämer, "Populist online practices: The function of the internet in right-wing populism." *Information, Communication Society*, vol. 20, no. 9, pp. 1293–1309, 2017.

[16] T. P. A. E. P. O. . O. P. Kaltwasser, C. R., *The Oxford handbook of populism.*, 1st ed. Oxford University Press., 2017.

[17] A. I. . K. G. Stavrakakis, Y., "A new populism index at work: identifying populist candidates and parties in the contemporary greek context." *European Politics and Society*, vol. 18, no. 4, pp. 446–464, 2017.

[18] B. Stanley, "The thin ideology of populism." *Journal of political ideologies*, vol. 13, no. 1, pp. 95–110, 2008.

[19] U. Klinger and J. Svensson, "The emergence of network media logic in political communication: A theoretical approach." *New media society*, vol. 17, no. 8, pp. 1241–1257, 2015.

[20] M. B. . W. A. Kefford, G., "Populist attitudes: Bringing together ideational and communicative approaches." *Political Studies*, vol. 70, no. 4, pp. 1006–1027, 2022.

[21] Suhun Han. Googletrans: Free and unlimited google translate api for python. v 3.0.0. [Online]. Available: https://py-googletrans.readthedocs.io/en/latest/#

[22] . C. R. Chatterje-Doody, P. N., "Populism and contemporary global media: Populist communication logics and the co-construction of transnational identities." in *Populism and world politics: Exploring inter-and transnational dimensions.*, D. B. M. Stengel, Frank A. and D. Nabers, Eds. Springer, 2019, pp. 73–99.

[23] M. J. Meijers and A. Zaslove, "Measuring populism in political parties: appraisal of a new approach." *Comparative political studies*, vol. 54, no. 2, pp. 372–407, 2021.

[24] . M. O. Biancalana, C., "Unifying and multiplying the people: The strategy of ambiguity in national-populist discourse within a cross-border area." *Nationalism and Ethnic Politics*, vol. 26, no. 3, pp. 279–298, 2020.

[25] P. Norris, "Global Party Survey, 2019," 2020. [Online]. Available: https://doi.org/10.7910/DVN/WMGTNS

[26] e. a. Volkens, Andrea, "Manifesto project database," 2022. [Online]. Available: https://manifesto-project.wzb.eu/

[27] Ing. Rodolfo De Benedetti Foundation. Populist rhetoric and representative democracy. [Online]. Available: https://frdb.org

[28] Twitter API v2. Twitter academic api. [Online]. Available: https://developer.twitter.com/en/docs/twitter-api

[29] S. L. Jacobs, K. and N. Spierings, "Twitter and facebook: Populists' double-barreled gun?." *New Media Society*, vol. 22, no. 4, pp. 611–633, 2020.

[30] A. Balahur and M. Turchi, "Comparative experiments using supervised learning and machine translation for multilingual sentiment analysis." *Computer Speech Language*, vol. 28, no. 1, pp. 56–75, 2014.

[31] L. H. Smirnova, A. and N. Kolenda, "Ideology through sentiment analysis: A changing perspective on russia and islam in nyt." *Discourse Communication*, vol. 11, no. 3, p. 296–313, 2017.

[32] X. Chen and J. Hu, "Evolution of u.s. presidential discourse over 230 years: A psycholinguistic perspective." *International Journal of English Linguistics.*, vol. 9, no. 4, pp. 28–41, 2019.

[33] H. Batra, N. S. Punn, S. K. Sonbhadra, and S. Agarwal, "An intervening intermediary: making political sense of media influence." in *How political actors use the media: a functional analysis of the media's role in politics*, 2017, pp. 21–37.

[34] F. J. M. L. J. D. L. S. D. V. P. J. C. Bucy, E. P. and C. Wells, "Performing populism: Trump's transgressive debate style and the dynamics of twitter response." *New Media Society*, vol. 22, no. 4, pp. 634–658, 2020.

[35] P. Tsatsou, "Literacy and training in digital research: Researchers' views in five social science and humanities disciplines." *New Media Society*, vol. 20, no. 3, p. 1240–1259, 2018.

[36] A. A. S. S. Boyd, R. L. and J. W. Pennebaker, "Liwc-22: Descriptive statistics and norms." 2022. [Online]. Available: https://www.liwc.app

[37] ——, "The development and psychometric properties of liwc-22." pp. 1–47, 2022.

[38] B. A. N. Jonathan Haidt, Jesse Graham and S. J. Peterson, "Moral foundations dictionary," 2008. [Online]. Available: https://moralfoundations.org/other-materials/

[39] W. I. C.-R. C. F. S. Alizadeh, M. and M. Macy, "Psychology and morality of political extremists: evidence from twitter language analysis of alt-right and antifa." *EPJ Data Science.*, vol. 8, no. 1, pp. 1–35, 2019.

[40] N. Faulkner and A.-M. Bliuc, "Breaking down the language of online racism: A comparison of the psychological dimensions of communication in racist, anti-racist, and non-activist groups," *Analyses of Social Issues and Public Policy*, vol. 18, no. 1, pp. 307–322, 2018.

[41] The Migration Observatory at the University of Oxford. What terms do we use to talk about migration? [Online]. Available: https://migrationobservatory.ox.ac.uk/resources/briefings/what-terms-do-we-use-to-talk-about-migration/

[42] M. Cameletti, S. Fabris, S. Schlosser, and D. Toninelli, "Dictionary-based classification of tweets about environment," 2020.

[43] Keyhole. Keyhole. [Online]. Available: https://keyhole.co/

[44] R. Egger and J. Yu, "A topic modeling comparison between lda, nmf, top2vec, and bertopic to demystify twitter posts," *Frontiers in sociology*, vol. 7, p. 886498, 2022.

[45] M. Grootendorst, "Bertopic: Neural topic modeling with a class-based tf-idf procedure," *arXiv preprint arXiv:2203.05794*, 2022.

[46] N. Reimers and I. Gurevych, "Sentence-bert: Sentence embeddings using siamese bert-networks," *arXiv preprint arXiv:1908.10084*, 2019.

[47] L. McInnes, J. Healy, and J. Melville, "Umap: Uniform manifold approximation and projection for dimension reduction," 2020.

[48] L. McInnes, J. Healy, and S. Astels, "Hdbscan: Hierarchical density based clustering." *J. Open Source Softw.*, vol. 2, no. 11, p. 205, 2017.

[49] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, R. Prettenhofer, P.and Weiss, V. Dubourg, J. Vanderplas, D. Passos, A. andCournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[50] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," 2019.

[51] M.-G.-M. N. Dang, N. C. and F. De la Prieta, "Sentiment analysis based on deep learning: A comparative study." *Electronics*, vol. 9, no. 3, p. 483, 2020.

[52] H. Batra, N. S. Punn, S. K. Sonbhadra, and S. Agarwal, "Bert-based sentiment analysis: A software engineering perspective," in *Lecture Notes in Computer Science*. Springer International Publishing, 2021, pp. 138–148. [Online]. Available: https://doi.org/10.1007%2F978-3-030-86472-9_13

[53] R. Lamsal, A. Harwood, and M. R. Read, "Twitter conversations predict the daily confirmed covid-19 cases," *arXiv preprint arXiv:2206.10471*, 2022.

[54] D. Q. Nguyen, T. Vu, and A. T. Nguyen, "Bertweet: A pre-trained language model for english tweets," 2020.

[55] V. G.-G.-A. M. V. T. B. G. O. Pedregosa, F. and E. Duchesnay, "Scikit-learn: Machine learning in python." *Journal of machine Learning research*, vol. 12, pp. 2825–2830, 2011.

[56] Crowdflower's Data for Everyone library. Twitter us airline sentiment. [Online]. Available: https://www.kaggle.com/datasets/crowdflower/twitter-airline-sentiment

[57] P. Norris. The global party survey 2019. [Online]. Available: https://www.globalpartysurvey.org/

[58] M. A. Eger and S. Valdez, "From radical right to neo-nationalist," *European Political Science*, pp. 1–21, 2019.

[59] P. Lehmann, S. Franzmann, T. Burst, S. Regel, F. Riethmüller, A. Volkens, B. Weßels, and L. Zehnter, "The manifesto project dataset - codebook," Berlin: Wissenschaftszentrum Berlin für Sozialforschung (WZB) / Göttingen: Institut für Demokratieforschung (IfDem), 2023.

[60] T. L. Saaty, *The analytic hierarchy process: Planning, priority setting, resource allocation*. North-Holland, 1982.

[61] R. Ramanathan, "Multicriteria analysis of energy," in *Encyclopedia of Energy*, 2004, pp. 77–88.

[62] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 2016, pp. 785–794.

[63] ——, "XGBoost: A scalable tree boosting system," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '16. New York, NY, USA: ACM, 2016, pp. 785–794. [Online]. Available: http://doi.acm.org/10.1145/2939672.2939785

Acknowledgments

I want to acknowledge the whole team at the fRDB Foundation for letting me use their data and research to build this dissertation, for providing good advice, good company and a very interesting new experience. Motivated as I was to continue to study populism, I doubt I would have looked so hard for a connection between the world of data and political science, hadn’t it been for the sequence of fortuitous events that landed me in your office.

I wanna thank my supervisor, prof. Tomaso Erseghe, for putting up with my utter disorganization and constant changes of plans with a smile: it was more appreciated than he might ever know.

I also want to thank my colleague and friend Carlo De Dominicis for his invaluable, although brief, contribution to the project, for the endless debugging Whatsapp conversations and for his patience and positive attitude in the face of adversities. I will feel very lucky to work (and not work) with you in the future.

Aside from those who put their time into the lengthy research that led to the dissertation you’re currenlty holding in your hands, many people contributed indirectly to this work, as many as I have encountered in my quarter of a lifetime. Whoever met me for more than five minutes will know that anyone not in this list has not been forgotten in my heart, only in my slightly crazy head.

First of all, I wanna thank my mum, Monica: she’s the strongest, most resilient person I know; despite our differences, I don’t think I would have ever gotten this far without her constant support, challenge and - more importantly - care.

To my brother, Franco and my dad, Carlo: in your own, very different, ways, you both managed to make a bookworm into a lover of numbers and science. I may not be a writer (yet), but I could not be happier about the way my path is leading me to unexpected places. I really hope yours does the same.

To Elena, Pietro, Dodo, Susy, Valeria, Ludovica, and the whole family: thank you for being the craziest family on the planet, with all the good and the bad this brings with it. It took me 25 whole years to realize it, but I wouldn’t want it any other way.

Now that the relatives are not going to get jealous, I can thank my other, chosen, family: my friends. The order could not be more sparse, as are my brain cells after writing this dissertation for 12 months.

To Ivan: every bit of this is as yours as it is mine. Thank you for being the best support system a girl could ever wish for, the best study buddy, the best 80s outfit inspiration. I'm so excited for the rest of our lives I can barely type it.

To all the people I met in Padova through living, playing music, and just aimlessly walking around: a page is not enough to thank you all, but I will forever be in awe of how incredibly loving, sweet and kind everyone in this city has been to me from the moment I arrived until the moment I left.

To Arghavan, Fatima, and all the lost souls walking the libraries of Piovego at night: I'm so glad you were there to share a cigarette and a laugh during the most desperate life and academia situations, and ever more glad about all the amazing milestones we'll share in the future.

To Vanni, my forever housemate: you will literally never get rid of me - and yes, it's a threat. Thank you to Max, who was always there when I needed to vent; and a big, massive thank you to my wonderful group of am3: Giorgia, Beatrice, Letizia, Caterina and Nic. To me, you're "home" now.

One final Padova shout out to Lara, who used her magic to heal my heart, even from thousands of miles away. Obrigada, amiga.

To Alessandro, the optimization magician: thank you for listening to my endless monologues behind a cigarettes, for teaching me loads, for being a good friend despite the limited time we got to know each other. Go, my friend, and crash all the Swedish university servers!

To my oldest friends Sofia, Carolina, Francesca, Irene, Lorenzo, Teresa, Marta, Anna Laura, Miriam, and all those who I forgot to list because I am writing this, as always, a few minutes before the deadline: we may not see each other as much as we used to, but if I lasted long enough to a Master's degree is also thanks to your care and love throughout the years.

To my friends Vicky, Richa and Oliver: you are missed every day, and since I would never have gotten my first degree without you, for sure I wouldn't be getting a second one. I hope life brings us as close in geography as we are in spirit.

To Lillian: I don't think I need to repeat here what you already know. Just move to this country already, because I don't think I can continue living without your chaotic joy, complete lack of judgement and huge heart.

I also want to thank Giovanni, Antonino, Giulia, and all the other wonderful people I met during this chaotic year, for filling it with love even under the worst circumstances. You are more appreciated and loved that I can express in words, but you already knew this.

A special thanks to Maria, who when I had to isolate in order to kick-start the writing of this

dissertation, took care of me, fed me (a lot) and made me feel at home outside of home.

The truth is, this list could go on forever: every human, animal, plant and organic element I met on the way to this day has somewhat contributed to this work, which is nothing if not the proof of all the loving interconnections that sustained me til its end.

None of this would have been possible in isolation; and even if it had been, it wouldn't have been half as exilarating as it was.

Thank you all, for everything.