

CHARLES UNIVERSITY
FACULTY OF SOCIAL SCIENCES
Institute of Political Studies
Department of Political Science

Master's Thesis

2024

Bc. Štěpán Košan

CHARLES UNIVERSITY
FACULTY OF SOCIAL SCIENCES
Institute of Political Studies
Department of Political Science

**Measuring Institutional "Wokeness": Text as Data
Approach**

Master's Thesis

Author of the Thesis: Bc. Štěpán Košan

Study programme: Political Science (N0312A200011)

Supervisor: Mgr. Jakub Stauber, Ph.D.

Year of the defence: 2024

Declaration

1. I hereby declare that I have compiled this thesis using the listed literature and resources only.
2. I hereby declare that my thesis has not been used to gain any other academic title.
3. I fully agree to my work being used for study and scientific purposes.
4. During the preparation of this thesis, the author used ChatGPT4, Llama 3 & Github Copilot for code generation along with abstract and summary pre-formulation. After using this tools and services, the author reviewed and edited the content as necessary and takes full responsibility for the content of the publication.

In Prague on the 31st of August 2024

Bc. Štěpán Košan

References

KOŠAN, Štěpán. *Measuring Institutional "Wokeness": Text as Data Approach*. Praha, 2024. 81 s. Master's thesis (Mgr). Charles University, Faculty of Social Sciences, Institute of Political Studies, Department of Political Science. Supervisor Mgr. Jakub Stauber, Ph.D.

Length of the Thesis: 100995 (characters with spaces)

Abstract

The thesis explores the concepts of “wokeness” and “being woke” in connection to U.S. political institutions. It is argued that the “wokeness”, a term originating from the awareness of social injustices faced primarily by black Americans, has evolved into a multi-faceted concept influencing modern political discourse. U.S. federal government’s, along with state governments’ and U.S.-based universities’ press releases from 2001 to 2024 are classified based on predefined criteria indicative of “wokeness”, such as emphasis on the primacy of group identity, inter-group equality of outcome measures and social activism to those ends. Methodologically, the thesis utilises language models for text classification, particularly the fine-tuned DeBERTaV3-Large model, and proposes novel techniques for fine-tuning and evaluating its performance. The findings reveal temporal and institutional variance in the expression of “wokeness”, with a significant increase since the advent of the Biden administration, pointing to it being a more recent phenomenon. Statistical analysis further reveals a more significant association of “wokeness” with the Democratic party and much more so in recent years, falling in line with the assumption of “wokeness” as a left-wing ideology with a complicated relationship with multiculturalism and heavily drawing from critical theory. Generally, the research aims to contribute to political science in two ways: to validate the use of automated text classification in its various sub-fields and to, based on the empirical findings, coupled with an extensive literature review, argue for the incorporation of the concept of “wokeness” as a distinct ideological construct for further study.

Abstrakt

Práce zkoumá koncepty „wokeness“ a bytí „woke“ ve spojení s politickými institucemi USA. „Wokeness“, termín pocházející z „uvědomění si“ sociální nespravedlnosti které čelili především Afroameričané, se vyvinul v mnohostranný koncept ovlivňující moderní politický diskurz. Tiskové zprávy federální vlády USA, vlád dvou amerických států a univerzit se sídlem v USA z let 2001 až 2024 jsou klasifikovány na základě předem definovaných kritérií indikujících „wokeness“, jako je důraz na skupinové identity, rovnost výsledků mezi identitními skupiny a sociální aktivismus bojující za tyto cíle. Metodologicky práce využívá jazykové modely pro klasifikaci textu, zejména „fine-tuned“ model DeBERTaV3-Large, a navrhuje nové metody pro „fine-tuning“ a evaluaci jeho validity. Výzkum odhaluje časové a institucionální rozdíly v přítomnosti „wokeness“, s

výrazným nárůstem frekvence výskytu od nástupu administrativy presidenta Bidena. Podrobná statistická analýza odhaluje asociaci „wokeness“ s Demokratickou stranou, a to více v posledních letech, což je v souladu s předpokladem „wokeness“ jako levicové ideologie s komplikovaným vztahem k multikulturalismu a široce čerpající z kritických teorií. Práce si obecně klade za cíl přispět politologickému výzkumu dvěma způsoby: otestovat používání automatizované klasifikace textů v jejích dílčích oblastech a na základě empirických zjištění spolu s rozsáhlým přehledem literatury argumentovat pro začlenění konceptu „wokeness“ jako distinktivního ideologického konstruktů pro další studium.

Keywords

Content analysis, Wokeness, U.S. Politics, Machine learning, Identity politics, Critical theory, Language models

Klíčová slova

Obsahová analýza, Wokeness, Politika USA, Machine learning, Politika identit, Kritická teorie, Jazykové modely

Title

Measuring Institutional "Wokeness": Text as Data Approach

Název práce

Měření institucionální "Wokeness": Přístup na základě textové analýzy

Acknowledgement

I would like to express my gratitude to my family, namely my mother Martina, my father Štěpán and my grandmother Zdeňka, for their undying and patient support in my studies and without whom I would not be here. Next, my thanks go to my partner Eugenia, for whom the same applies, and also for her consultations on this work. Last but certainly not least, I must acknowledge the extraordinary goodwill and educational skills of Dr. Stauber, who guided me in marrying my two passions of programming and political science.

Table of Contents

Introduction.....	8
1 Defining “Wokeness“.....	12
1.1 Literature Review.....	12
1.2 Conceptual Overview.....	18
1.3 The Case for “Wokeness” as an Ideology.....	20
2 Methodology and Research goals.....	23
2.1 Subject of Research.....	23
2.2 Methodology.....	25
3 Research.....	35
3.1 Devisement of Classification Criteria.....	35
3.2 Data Collection and Characterisation.....	36
3.3 Training Corpus.....	42
3.4 Fine-tuning the Model.....	43
3.5 Classification and Evaluation.....	44
3.6 Statistical Analysis.....	47
3.7 Discussion.....	57
Conclusion.....	62
Summary.....	65
List of References.....	66
Index of Tables.....	79
Table of Figures.....	80
List of Appendices.....	81

Introduction

“Wokeness” as a noun or “woke” as an adjective has increasingly entered the lexicon worldwide (Jurik 2023; Google Trends 2024).

The roots of the current meaning are inextricably linked with the discrimination that black Americans were undoubtedly exposed to in the United States, at least prior to the civil rights movement, the demontage of the segregation regime in the American South, and the general multi-decade attitude shift country-wide (Gallup 2021; National Archives 2021; Public Policy Institute of California 2021). While the exact first mention of “wokeness” in this context is disputed, it was generally meant to denote the state of awareness of injustice in various respects that black Americans were experiencing at the time. (NAACP 2023)

More widespread use of “being woke” in public discourse began after the Ferguson, Missouri shooting in 2014 and other perceived unlawful killings of black men (Ng 2021), which were followed by protests leading to the advent of the Black Lives Matter movement¹ (albeit specifically the killing of Michael Brown in Ferguson was ruled justified at many levels of review (Kessler 2019; Capehart 2022)).

Contemporarily with the advent of this “positive” connotation and use of “wokeness” as something to strive for and a personal quality to be attained, more negative use in political rhetoric has evolved, where “wokeness” is meant as “progressive overreach”, negative overt focus on identitarian “justice” and authoritarian speech policies enforcing the “woke” worldview. (Rufo 2020; Smith, Boag, Keegan, Butler-Warke 2023; Alfonseca 2024)

While the public across the Western world remains largely divided on what they consider “wokeness” to be (Montanaro 2023; Smith, Boag, Keegan, Butler-Warke 2023), this thesis attempts to capture how the term is conceptualised in academia (Cammaerts 2022; Zavattaro, Bearfield 2022) and by what the author of this thesis terms “academia-adjacent” thinkers and theorists (Lukianoff 2021; McWhorter 2022; Hanania 2023).

Based on an extensive literature review, a meta-classification of “wokeness” is proposed from across the existing corpus of literature, with one end of the spectrum being

¹ Black Lives Matter as an umbrella organisation was, however, instituted a year prior in reaction to the case of George Zimmerman killing of Trayvon Martin. (BLM 2024)

“‘wokeness’ as a feature of political rhetoric”, conceptualisation most often found in works of social science authors, where “wokeness” is primarily constituted in political speech of its opponents and does not possess definition of its own, because what is termed “woke” is essentially the normatively desired state of politics and society at large (Cammaerts 2022; Pilkington 2022; Zavattaro, Bearfield 2022).

The midpoint of the spectrum is “‘wokeness’ as an individual ethic”, a concept of “wokeness” as a state of an individual aware of the “systemic and structural injustice” some identity groups, especially black Americans allegedly have to cope with and prescribes to the persons belonging to the “oppressor” group to effectively treat those individuals differently as to “equalise” their position in society (Roy 2018; Whiteout 2018; Atkins 2023a).

The other pole of the conceptual spectrum is “‘wokeness as an ideology”, a conceptual framework that identifies “wokeness” as an ideology that posits that group identity ought to be central to politics, culture and life. In government, all policy should primarily be judged on If it reduces inter-group “inequity” (inequality of outcome) and specific previously or currently “marginalised” communities should enjoy special protections and benefits as to achieve the across group equality of outcome (Pluckrose, Lindsay 2020; Sullivan 2020a; Yang 2021; 2021).

What this thesis aims to achieve theoretically is to examine the soundness of the conceptualisation of “‘wokeness’ as an ideology”, when authors outside of academia and political actors frame it as such. Along with examining empirical observations and claims that its proponents offer. Methodologically, it aims to offer an implementation of new techniques of text classification and devisement of specific classification model to aid in reasearch.

Especially relevant to the field of political science and social sciences at large, appears to be the claim that alleges that “wokeness” has taken hold in the U.S. federal government, and that primarily since the advent of the administration of President Biden (Lukianoff 2021; Yang 2021) and informs the policy choices or, at the very least, the political communication of the administration.

To investigate that hypothesis and generally map the potential presence of “wokeness” in the U.S. political landscape, the thesis sets out to utilise the methodology of content analysis, which posits that texts are objectively classifiable into categories based on determined criteria (Krippendorff 2004; Hayes, Krippendorff 2007; Grimmer, Roberts, Stewart 2022). This framework is then used to classify press releases from primarily sections of the U.S. federal government between the years 2001 and 2024, with an auxiliary text corpus of U.S. universities’ and state governments’ press releases on whether they contain notions associated with “wokeness’ as an ideology”.

Due to the considerable number of units of texts, a machine text classification technique is used. Based on the most recent trends from across social science and beyond (Eckhard, Jankauskas, Leuschner, Burton, Kerl, Sevastjanova 2023; Mahendru, Pandit 2024), the mode of supervised learning – fine-tuning a language model on a pre-classified sample is chosen as possibly the most valid method for both the size of the dataset and the research question outlined. Furthermore, to tackle the issues inherent in single-coder classification of the training sample (Hayes, Krippendorff 2007), a zero-shot classification language model is employed (Hugging Face 2023). Both the zero-shot model and the to-be-fine-tuned model are of the DeBERTaV3 variety, the most recent descendants of the original BERT model and, for now, the endpoint of the encoder-only language model evolutionary branch (He, Liu, Gao, Chen 2021; Vijay Srinivas Tida, Hsu 2022). Their inner functioning and peculiarities that have to be accounted for are further detailed in the methodology section.

Provided successful classification, as evaluated by a novel sample technique proposed, the dataset is analysed for year-on-year changes and the ratio of “woke” press releases across the researched time period is visualised. Along with these methods, logistic regression is used to ascertain what is the probability of a press release being “woke” based on the institutional attributes of its originator and the year of its release, such as control by the Democratic party or the Republican party, whether the institution is part of the federal government and others. (IBM 2021a)

The thesis treats “text as data” both in the context of automated content analysis – transformer language model classification – but also in inference from the frequency

classification and that in order to aid in answering the following primary broad research question, drawing from the conceptual and theoretical framework of “‘wokeness’ as an ideology”:

“Provided that ‘wokeness’ is an identifiable ideology, what presence has it had in sections of the U.S. federal government, as measured by their external communications, from 2001 until 2024?”

The thesis’ scientific contribution aims to be primarily two-fold. Firstly, to introduce the concept of “‘wokeness’ as an ideology” into academic literature and to test the empirical claims and hypothesis contained within it. Secondly to showcase the viability of use of advanced machine learning methodology in social sciences generally and when working in the context of research of political institutions particularly.

The thesis is divided into three main parts. In the first section “Defining ‘Wokeness’”, theoretical and conceptual review is offered, inclusive of literature review. In the second section “Methodology and Research Goals” the subject of research is justified, along with the chosen methodology and specifically automated text classification is introduced. In the third section “Research”, the empirical research is carried out.

1 Defining “Wokeness”

1.1 Literature Review

The term "wokeness" or "being woke" has become a feature of common parlance (as demonstrated by, among other things, the frequency of Google searches of the term “woke” illustrated in Figure 1) with a distinct meaning while largely resisting a singular agreed upon definition. As has been elucidated in the introduction of this thesis, used to denote concepts well defined (or at the very least discussed) in academic literature, such as critical consciousness, intersectionality, social justice, diversity, equity and others, the attempts to provide a normative definition have been sparse in peer-reviewed corpus and left mainly to academia-adjacent thinkers and political actors.

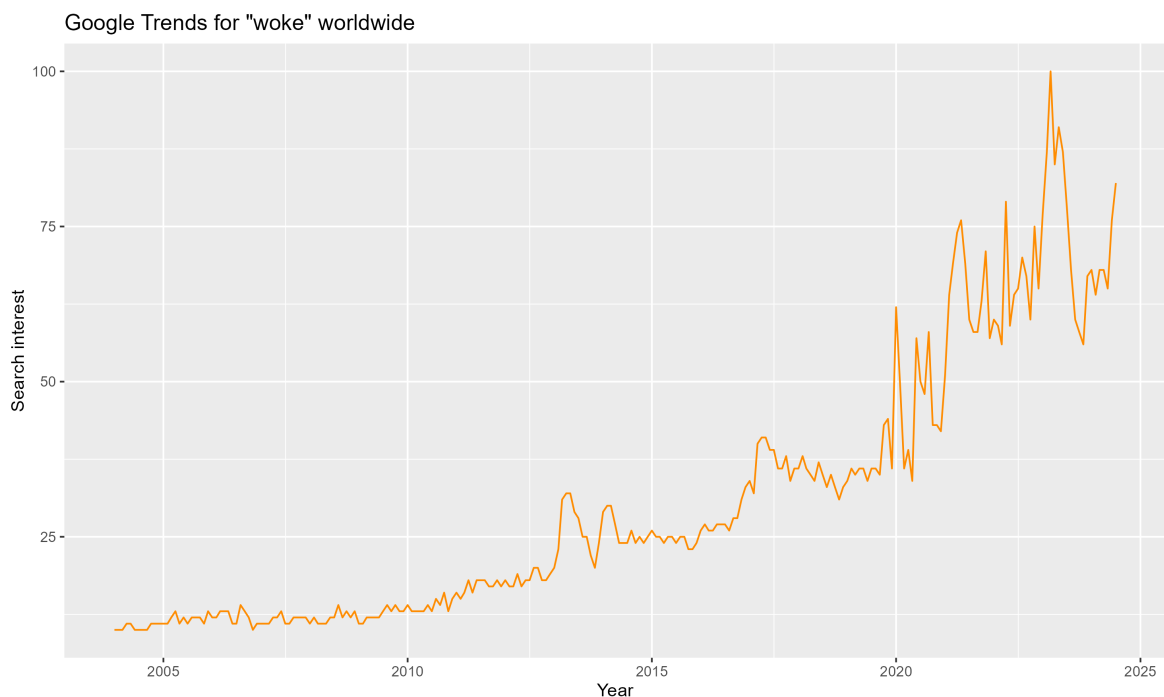


Figure 1: Google search trend for “woke” worldwide (own work, data from Google (2024))

In social sciences "wokeness" is seen predominantly as a term in political discourse rather than an ideology in its own right, employed by often populist actors in a typical “backlash” fashion against any perceived progressive or even liberal policy (Abrajano, Hajnal 2015). It usually lacks internal coherence in its application and is essentially part of

permanent campaign rhetoric or a part of a fight against weights and balances that political opposition or the permanent state bureaucracy represents. (Bennett 2023; Smith, Boag, Keegan, Butler-Warke 2023)

Specifically, according to Bart Cammaerts (2020; 2022; 2024), conservative and populist (in his view “fascist”) political actors utilise the conceptualisation of “woke” to group in broad political ideas aimed at furthering social justice by fostering the rights of “marginalised” groups. That discourse is then seen in the context of the concept of Antonio Gramsci's (1967) normalisation of the abnormal to sustain cultural and political hegemony, the normalised abnormal in this sense being the “injustice” of the social and political arrangements that actors accused of being “woke” are attempting to bring the end about to. Analogously, Schmittian distinction of friend and enemy is utilised (Schmitt 2005; 2021), with the “woke” fulfilling the role of the enemy and the “guardians” of the current order coming into the friend fold, be it right-wing media or allied political actors. While he accepts the heritage of the word as originating from the African-American civil rights struggle, i.e. being “woke” to the injustices of black Americans, the current usage standing on its own is implicitly meaningless and gains significance only when used by “anti-woke” actors in the manner outlined.

Cas Mudde (2021; 2023), while not implementing the concept fully in his academic work yet, has shown a similar attitude toward the topic while speaking on Florida's “anti-woke” act, seeing it as primarily part of anti-social justice discourse from the Republicans/right-wing in attempting to further their “hegemony” by silencing voices pointing to institutional inequalities caused by factors such as “institutional racism”.

Approximate sentiment is echoed by Andrew Pilkington (2022) when he sees “anti-wokeness” as being the definition rather than “wokeness” as such. He attempts to draw parallels between the use of the term “political correctness” that, according to the author, had also primarily exogenous and incoherent definitions originating with its opponents.

Staci M. Zavattaro (2022), on the other hand, agrees with the “woke” exclusively connected to the crusade against it post-2020; she expand upon Cammaerts’ (2020; 2022) mapping of the original usage of the word and situates it in the context of Baudrillard's (1994) theory of hyperreality, where concepts lose their connection to reality over time.

The reality in this sense being the struggle for social justice by “marginalised” groups, specifically African-Americans in the 1950s and 1960s, through the first phases of policy incorporation of “wokeness” via, according to the author, positive legal changes such as the 1964 Civil Rights Act, through a second phase of less meaningful policy actions, but fueled by the original sentiment, such as “incomplete” affirmative action only preserving “existing power structures”. Then arriving at the third phase of conceptual emptying of “wokeness”, where the substantial policy gives way to gestures, such as kneeling or gesturing participation in the “social justice struggle”. In the final and “current” stage of the concept, Zavattaro, similar to the thinkers mentioned above, find it in “counter-social justice” rhetoric, where “wokeness” is used as an umbrella term for social progress, and political actors use it to mobilise beneficiaries of current social arrangement against it. Comparatively, “anti-wokeness” here is seen more concretely as a direct discursive weapon utilised against progress, rather than a more broad feature of populist rhetoric.

Where we find more positive or rather normative articulations of what “wokeness” is, ought to be and has been is in the fields of philosophy and education, while drawing on most definitely political literature.

Kayla M Johnson (2023) points to “wokeness” being a state of Freirean “consciousness”, i.e. the development (in that case, that of students) of critical consciousness of how “systems of oppression” – primarily racism – determine their positionality in society and how that should inform their current and future actions in dismantling the “systems of oppression” (Freire, Shor, Macedo, Ramos 2018). “Wokeness” is then conceptualised as a state of individual vigilance to injustice and how “systemic injustice“ informs a particular action, for example, in the form of implicit bias. Firmly grounded in critical theory, “wokeness” is then perceived as individual praxis, ideally leading to a societal change.

The same is the case of Sam Whiteout (2018) who generalises the case for collectives (such as the Movement for Black Lives) that led the impetus for individuals to be more “woke” and perceptive to societal injustices. For Whiteout it is an aspirational state, never to be fully attained but consistently to be strived for.

Laura A. Roy (2018) expands on the notion by elaborating that in the contemporary world, it is a state specifically pertinent to members of the “oppressor” class, and for it to be the

first stage, especially for white people as individuals, in their quest for “atonement” and improving the societal conditions for people of colour, in some ways echoing the thinking of Kendi (2017) and DiAngelo (2016; 2019) and their conceptualisation of “whiteness” and its societal repercussions.

Sobande (2022) and Marshall (2023), while not citing the works of Roy explicitly, in a comparable vein develop the idea of “wokeness” as being distinct from actual social activism in the progressive direction. While being its definite prerequisite, the authors point to the risk of the so-called “performative wokeness”, such as an “anti-racist” statements on social media, that cannot supplant actual societal change. Boyce (2021). then sees “wokeness” as precisely a phenomenon of “white liberals” who, by being performative about their supposed care about the inequality in society, can allow themselves not to do the “actual” activist work.

In the field of social epistemology, the first examples of the definition of “woke” praxis, i.e. how a “woke” individual ought to behave in various contexts, either because it is ethically correct to do so in its own right or because that act ought to eventually contribute to the downfall of “systems of oppression”, are found. Based on either the theory of moral encroachment (the philosophy that a view should not be judged solely on its truthfulness or basis in reality but also on its moral merit) or the theory of group partiality, it posits that a “woke” person ought to distinguish between the opinions and acts of a member of an “oppressed” group as their actions and opinions are formed and informed by the “systems of oppression” – racist government, society, unjust laws and other systemic factors – they are forced to live under. (Atkins 2023a; 2023b)

Moving towards a somewhat broader conceptualisation of “wokeness” as a set of beliefs identifiable in culture and politics. Wesley Yang (2021) terms this the “Successor ideology”, a notion synonymous with “wokeness”, “political correctness”, and other terms, according to him, either constituent of or interchangeable with the broader concept. Fundamentally, “wokeness” is then a rejection of the liberal, individualistic Western perspective on society and a shift towards group identitarianism, along with the unequivocal acceptance of the notions that current systems of governance and economics are rife with discrimination against minority groups. Furthermore, all societal endeavour must then be aimed at its dismantling or, at the very least, its radical reform. As supporting

evidence of the existence of such ideology and its societal influence, he cites a vast corpus of governmental and other institutional actions that have explicitly moved away from the “colourblind” approach to achieve effective equality of outcome. Primarily race-based but increasingly inclusive of all possibly "marginalised" identities. He also identifies what he terms “Year Zero” being 2021 and the inauguration of President Biden as the major incursion of this thinking into the U.S. federal government, after being “brewed” in society and within the intellectual elite during the Trump years (Manhattan Institute 2023).

Drawing on the works of Lindsey and Pluckrose (2020), Andrew Sullivan (2019; 2020a; 2020b; 2020c) identifies the roots of what he terms an ideology in the new left of the 1960s and what some term neo-Marxism, where various “marginalised” groups crudely substitute the working class, that ought to fight for its liberation. In the U.S. context, he sees this outlook becoming the new orthodoxy in various institutions in the 2010s and focuses explicitly on the societal enforcement of the “dogmas” akin to what Havel identified with his work in the socialist Czechoslovakia. Analogously posits “wokeness” in opposition of broadly conceptualised liberalism, perhaps what Sartori (2005) would frame as pluralism, in “wokeness” unchangeable group identity being the determinant of one's positionality in the system and of their world-view.

McWhorter (2022) sees “wokeness” as primarily concerned with race as the defining group identity. He does not necessarily position it against liberalism in the American sense of the word but sees it as an outgrowth of it. What he terms as overreaches, in the form of “cancel culture” (what Sullivan would term orthodoxy) of dissenting voices, is to be blamed on the structure of the “wokeness” belief system that behaves de facto as a religion, with views that deny the “systematically racist” nature of the American government being actively suppressed. He primarily zeroes in on the negative effect this homogenisation of groups (African-Americans in particular) brings about, including the erasure of individual will philosophically and policies that do not in the long-term benefit “marginalised” groups, such as race-based affirmative action. Agreeing with Yang (2021), he identifies the core tenet of “wokeness” as being the focus on power differentials in-between groups that are supposed to be the determinant of who the “oppressor” and the “oppressed” is.

Goodwin (2023a) sees “wokeness” as a feature of the “new elite” divorced from the material conditions of the majority of the population. The “new elite” is concerned primarily with identitarian topics, even though those are not of concern to the majority of the governed. Rebuke to Inglehart (1977) of sorts, but not wholly as some “cultural” topics such as immigration he finds salient.

Kaufmann (2022) identifies “wokeness” as the “religious form” of what he terms “cultural socialism”, which sees the role of government and society at large in ensuring equal outcomes for identity groups (LGBT[Q], black, women etc.) along with maximalist “harm protection” instituted for these groups manifested in phenomena such as “cancel culture”, that puts it at odds with liberal-enlightenment values, such as equal treatment, objective truth, freedom of speech and due process. Further, the danger of “deculturation” is identified as the focus on “identity equity” (equality of outcome) leads to the vilification of history and institutions that form the “liberal order”.

Christopher C. Rufo (2020) focuses on what Kaufmann would consider “deculturation” and identifies the role that critical race theory (CRT) plays in that process. While many would disagree, positioning CRT only as a theory of law not to be found outside academia (Crenshaw 1991), Rufo argues that it is the CRT “praxis”, actions and modes of functioning coming out of conclusions of CRT – “systemic racism”, non-existence of race-blind justice system, “implicit bias” etc. – that serve as the propellant to the illiberal policies of preference for certain groups, which he claims serve neo-Marxist ends of equality of outcome.

The analysis of George Lukianoff (2021) eclectically points to “wokeness” as an ideology in the progressive fold that determines every (identity) unequal outcome as the direct consequence of systemic failure that ought to be remedied by policies enforcing “equity”, and those cannot be interfered with and questioned as they follow the normative “good”. In finding the root cause of what he sees as a widespread phenomenon, he largely (albeit controversially) concurs with Hanania (2023), seeing as the inceptual moment the institution of the 1964 Civil Rights Act, which, while at its face liberal, in fact led to the broadening of its scope by the U.S. Supreme court banning measures that lead to the so-

called “disparate outcomes”. Famously using IQ test equivalent for positions in civil service leading to less hiring of minority applicants was curbed (Cornell Law School 2021).

Lindsey (2021; 2022a) would however oppose that notion as he posits “wokeness” as the antithesis to the civil rights strife and legislation derived from it. He find “wokeness” to be the exact antithesis to liberalism in all its forms, and he synonymises it with “critical social justice” which he defines by the following quote taken from the relevant literature:

“An analysis of how power, privilege, and oppression impact our experience of our social identities. “Full and equal participation of all groups in a society that is mutually shaped to meet their needs. Social justice includes a vision of society in which the distribution of resources is equitable” and all members of a space, community, or institution, or society are ‘physically and psychologically safe and secure.’” (Lindsay 2021; Bell 2023, p. 23)

Furthermore, compared to the identitarian wave of the 1960s, he sees the new “wokeness” to be equally focused on what he terms “political” identities, chiefly “queer”, that he posits in opposition to other sexual minorities (LGBT), as he finds “queerness” to not be descriptive only of the nature of attraction or self-identity, but also of political leanings, i.e. being “queer” is, according to Lindsay (2022b), incompatible with anything but adhering to “wokeness” (Lancing, Lindsay 2024).

1.2 Conceptual Overview

Based on the conceptual and theoretical ideas introduced and summarized in the preceding section, “wokeness”, as found in the existing corpus of academic and non-academic literature can be broadly classified into three categories:

1. “Wokeness” as a feature of political rhetoric
2. “Wokeness” as an individual ethic
3. “Wokeness” as an ideology

“Wokeness” as a feature of political rhetoric is perhaps the most prominent conceptualisation in social sciences, where "wokeness" is predominantly seen as a feature of “anti-woke” “reactionary” political rhetoric aimed at halting or reversing social progress. Sometimes, a positioning of the invention of "culture wars" is utilised. "Wokeness" as a concept is missing a comprehensive definition or fully losing it over time being disconnected to reality. (Cammaerts 2020; 2022; 2024; Zavattaro, Bearfield 2022; Mudde 2023)

"Wokeness" as an individual ethic is found primarily in interdisciplinary academic literature, philosophy and educational studies. Drawing on a corpus of literature of critical theory, it posits that "wokeness" is an individual state of awareness of "structural injustice" and the systems that form it. It prescribes individuals from "oppressor" groups to treat individuals from "oppressed" groups differently to "equalise" the power differential. It recognises and synonymises the concept of critical consciousness with that of "wokeness", critical consciousness being the elevated awareness of world's “power structures” and their forming of individual experiences. (Roy 2018; Boyce 2021; Sobande, Kanai, Zeng 2022)

“Wokeness” as an ideology or at least a set of identifiable political notions is a concept predominantly found outside academic literature (with notable exceptions). Most authors falling into this category would agree that “wokeness” is a comprehensive framework of thought that provides a lens through which to view society, politics, economy and history. That framework posits that group identity should be central to politics, culture and life at large; policy ought to be judged on the basis of whether it minimises (and eventually abolishes) the inequality of outcome and achieves “equity” (outcome equality) based on identity groups. Furthermore, the historically and presently “oppressed” group identities ought to enjoy practical protections against potential harm that could be caused to them even by speech (Kaufmann 2022; Goodwin 2023a; Wastell 2023). The diversity, inclusion, and equity policies, along with practices stemming from the critical race theory, are considered the central political prong of this world-view (Rufo 2020; Lindsay 2022a; Hanania 2023), with "cancel culture" being the societal method of enforcing the approach (Kaufmann 2022; Manhattan Institute 2023).

The majority, but not all, authors would agree that this ideology is in fundamental disagreement with and in opposition to broadly conceptualised (process) liberalism, with

its values of freedom of speech, individual rights, blanket non-discrimination and others. There is conceptual debate whether “wokeness” is primarily concerned with race group identity (McWhorter 2022) or has come to encompass all “marginalised groups” (Sullivan 2020a; Lindsay 2021). Furthermore, where there is perhaps the most profound disagreement is on the history and origins of “wokeness” as ideology, with some authors seeing it as a natural outgrowth of American social liberalism and others as a specific extension of the thought of the new left and neo-Marxism. (Pluckrose, Lindsay 2020; Sullivan 2020a; Yang 2021; Kaufmann 2022; Hanania 2023)

1.3 The Case for “Wokeness” as an Ideology

The three-point conceptual distinction can be perceived as a three-pointed horseshoe, with one end of the horseshoe being “wokeness’ as a feature of political rhetoric”, in the middle (at the bottom) on the connective line being “wokeness’ as an individual ethic” and on the opposite end “wokeness’ as an ideology”. Both “end” concepts draw from “wokeness’ as an individual ethic” in different but fundamental ways.

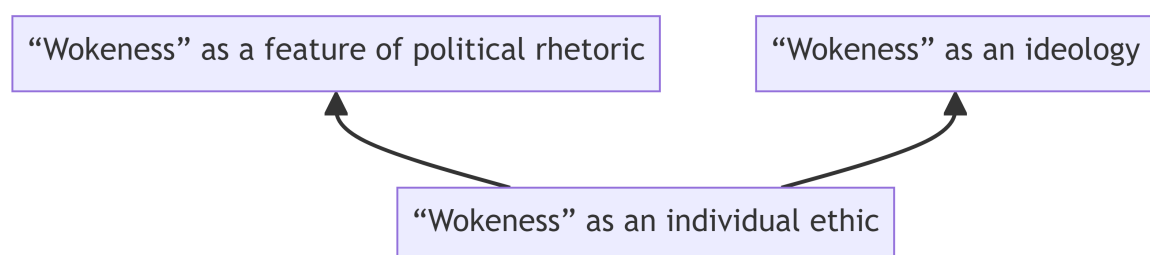


Figure 2: Conceptual map of “wokeness” (own work)

Why do authors of “‘wokeness’ as a feature of political rhetoric” see “wokeness” as nothing more? To put it simply, for them, the theories and applications of what the other side of the horseshoe would term as “woke” ideology is a normative good and the desirable standard, as demonstrated by, for example, explicit “pro-equity” policy recommendations in conclusions of Zavattaro (2022), or the framing of specifically anti-group equality of outcome speech as normalising the abnormal (inequality of outcome) by Cammaerts (2022; 2024) or the Mudde’s accusations of the United States being “institutionally racist” (2021). Essentially and most importantly, virtually all of these authors base their works on some strain of critical theory and working with its conclusions as given, inclusive of works of “‘wokeness’ as an individual ethic” authors (Freire, Shor, Macedo, Ramos 2018; Roy 2018).

While “‘wokeness’ as an ideology” concept can be undoubtedly criticised for its controversial searches of the history of the notions, along with authors at the very margins being perhaps overbroad in their conceptual branding, the view presented chiefly by Yang (2021), Kaufmann (2022), Sullivan (2020a) and partly Lindsay (2022a) and Lukianoff (2021), of “wokeness” as a fundamentally different ideology standing in opposition to liberalism, deserves its fair hearing, as it presents a comprehensive framework, firmly based in established in literature and ironically implicitly endorsed by the conceptual grounding of authors critical of it. The endpoint is that there is at least some conceptual intersection between “‘wokeness’ as a feature of political rhetoric” and “‘wokeness’ as an ideology” authors, but one set of authors would term that “the desirable state” and the other “wokeness”, and that being the reason for the Figure 2 not being a circle, but rather a horseshoe-like shape.

Somewhat superficially, but not necessarily wholly incorrectly, can “wokeness” be framed in the context of multiculturalism versus liberalism debate. While group unchangeable identity plays a role in the concept of “wokeness” of multiple authors, thus drawing parallels with authors such as Kymlicka (1995; 2018) or Taylor (1994) and in opposition to pluralistic/liberal conceptions such as of Sartori (2005) or even Rawls (1999; 2005) in some aspects, the relationship is more complicated. While group identity is at “wokeness” core, it is more focused on the systems of “oppression” and power differentials that allegedly exist between individuals belonging to “marginalised” groups, than on how the groups (ought to) form society. The focus is thus much more rooted in critical theory,

especially taking in account its repressive features very much in line with Marcuse (1970) and not to mention the advent of political identities such as queerness, which do not fit neatly into a multicultural framework. However, both “wokeness” and multiculturalism do indeed share the importance of group identity and in that respect they would be considered in opposition to liberalism, albeit authors such as Hanania (2023) would disagree.

Where “‘wokeness’ as an ideology” authors differ from ‘wokeness’ as an individual ethic” proponents is that they perceive “wokeness” to be a naturally societal phenomenon that permeates government and society at large (Sullivan 2020b; Yang 2021; Kaufmann 2022), that is an empirical statement for which they provide some evidence, but it is undoubtedly more than ripe for academic inquiry, that can not only either prove or disprove their claims, but provide further conceptual clarity based on the offered findings.

2 Methodology and Research goals

2.1 Subject of Research

Drawing on the theoretical and conceptual outlooks presented in the preceding section, with the introduction of the core tenets of the “‘wokeness’ as an ideology”, the following research question is formulated:

“Provided that ‘wokeness’ is an identifiable ideology, what presence has it had in sections of the U.S. federal government, as measured by their external communications, since 2001 until 2024?”

The subject of research—the U.S. federal government—is chosen because it is identified as the governmental institution at the centre of “wokeness” infusion, according to some in the long term (Pluckrose, Lindsay 2020; Sullivan 2020c), with cultural trends being incorporated into communication and policy or at a definite moment, as Yang (2021) posits. Furthermore, “wokeness” is almost invariably positioned as an anglosphere phenomenon that has only recently spread outside its cultural breeding ground (Kaufmann, Goodwin 2018; Kaufmann 2022; Goodwin 2023a). Thus, research on phenomenally “indigenous” institutions is warranted.

The distinction of sections is included, as the data is collected from various cabinet departments and the White House separately, along with the so-called independent agencies (Scholten 2014). The United States, being a presidential republic, has a strict delineation of constitutional powers, with the vast majority of the executive being under the direction of the president (Huntington 1981). While the theoretical extent and form of this control are sometimes debated, in practice, the president is the principal policy and direction setter, whose powers towards the executive are exercised either directly, through f.e. executive orders, specifying the execution of legislation, or indirectly through direction to him appointed (and in vast majority of cases senate confirmed) officials, who (once senate confirmed If required for the position) serve at his will, with the notable exception of in-legislation enumerated heads of independent agencies (Rozell, Barilleaux, Kelley 2010; Skowronek, Dearborn, King 2021). Some theorists propose that independent

agencies have a greater propensity to incorporate elements of “woke” ideology (Pluckrose, Lindsay 2020), as they are more insulated from direct executive oversight, especially during Republican presidencies (Crouch, Rozell, Sollenberger 2017).

Another assumption to test is If “wokeness”, being a broadly left-wing ideology, is significantly more linked with the American party of the (centre-)left, the Democratic party, and If the Republic party acts as a “safeguard” of (for some perhaps ironically) of the liberal values and general perspectives to which authors position “wokeness” in opposition to (Bobbio 2003; Ellis, Stimson 2012; Ura, Ellis 2012).

While specifically governmental output is scrutinised in the form of press releases, as opposed to campaign materials, which should allow for identification of “wokeness” in policy rather than in campaign rhetoric, due to Hatch Act exceptions (OSC 2024), presidents are allowed to conduct what would otherwise be considered campaign by governmental channels, that other, although partisan, executive appointees are not (Huntington 1981; Rozell, Barilleaux, Kelley 2010).

The time period and specific departments and agencies are chosen partly for data availability reasons and partly for theoretical justifications. Period-wise, after initial research based on the assumption of wide-spread digitisation of governmental affairs during the administration of Bill Clinton (1993-2001), it has proven that, while some material exists, they are mainly undigitised or If digitised, since inaccessible (including large parts of Clinton’s White House website). It is only since the administration of George W. Bush (2001-2009) that most agencies have publicly available records online (even if they are in difficult-to-scrape formats, as will be expanded upon). Theory-wise, most authors would position the “birth” of “wokeness” culturally somewhere in the 1990s, with a significant rise of associated sentiments at around the midpoint of the administration of Barack Obama (2009-2017), with a governmental decline under the presidency of Donald Trump (2017-2021) only to reach its apex during the Biden years (since 2021) (Pluckrose, Lindsay 2020; Yang 2021; Kaufmann 2022; Wastell 2023).

Furthermore, for control, an additional corpus of press releases of state authorities of two states, one controlled by Republicans for the whole surveyed period (Texas) and one controlled by Democrats (New York), along with the press releases of the two of the most

prestigious universities in the US (where theorists claim “wokeness” primarily originates from (Goodwin 2023b; Sullivan 2020b; 2020a)), Harvard and Princeton (Pluckrose, Lindsay 2020; Sullivan 2020b; 2020a).

To answer at least in part the broad research question outlined, mapping of the presence of “wokeness” in the external communication of the surveyed institutions in time will be undertaken, with a further focus on the originating institution, via methods outlined in the following section.

2.2 Methodology

The practical research will roughly follow steps visualised in Figure 3 that will be justified in this section:

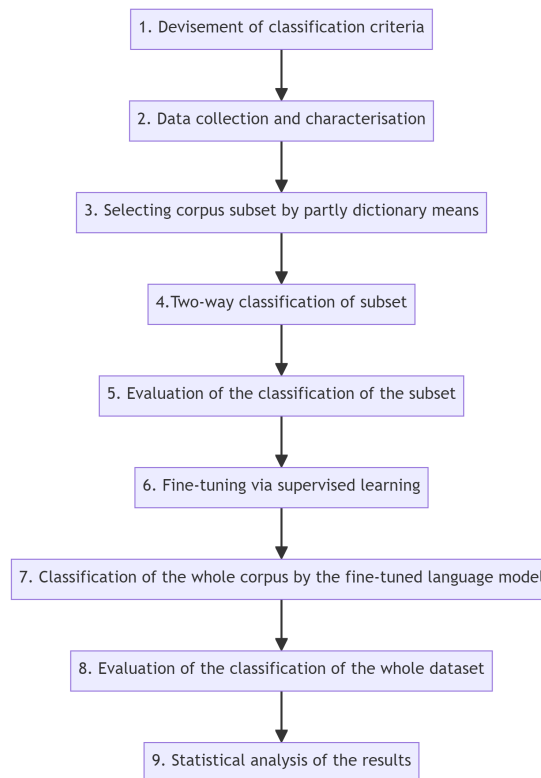


Figure 3: Methodological flowchart of the thesis (own work)

In order to be able to classify the corpus of texts, respectively, to ascertain whether one individual press release contains the notions of “wokeness”, we leverage the methodology of content analysis as formulated by Krippendorff (2004; 2007), which posits that text is

classifiable in a valid and replicable way, based on set criteria, i.e. coding scheme, with the ultimate goal of effectively converting a corpus of individual texts into a dataset, where each text is coded into either 1 or 0 (effectively boolean in the case of this thesis, but not necessarily) and then statistically computed to infer possible answers of the delineated research question (hence the “text as data” subtitle of the thesis). The coding of the majority of the dataset will be done primarily by a fine-tuned (by supervised learning) language model, as a human-coding dataset of a presumed population in the hundreds of thousands of units would not be feasible, and modern methods allow for leveraging computer-assisted classification

Utilising the Krippendorffian (2004) framework, we require the body of text on which content analysis will be carried out (*the thesis step number two*), clearly defined research question in which to “make sense” of the text corpus (*the thesis step number one*), an analytical construct to apply our understanding of the context, inferences to address the research question (*the thesis steps number three, four, six, seven and nine*), and validating evidence to justify our findings (*the thesis step number five and eight*).

As pertaining to the body of text (*the thesis step number two*), it is decided to use the individual press releases as a coding unit (or rather a unit to be coded), taking into account the research question and the fact that they are the product of a political institution presumably presenting a singular sentiment in one press release, as opposed to for example newspaper article or a record of parliamentary debate, where more granular approach would be warranted (Grimmer, Roberts, Stewart 2022; Dun, Soroka, Wlezien 2021).

Context-wise, ideally, the whole untruncated press release would be utilised to code; while that would be possible with the use of human coders for the whole dataset, that is not viable considering both the technical and financial limitations inherent in thesis work, truncation of minimally cleaned texts is introduced of each press release to 1800 characters to streamline both human and machine coding and to specifically adhere to technically feasible model ingestions limits (fitting into 512 tokens in the tokenisation stage) (Laurer 2024). Furthermore, as the primary coding will be done by either a human coder (the author) or a language model, the pre-processing of the data can be limited to simply

stripping the text of scrapping artefacts, unnecessary whitespace, line denominators, numbers and symbols as opposed to more thorough sanitation required for other text as data approaches, such as bag words approach (Eckhard, Jankauskas, Leuschner, Burton, Kerl, Sevastjanova 2023).

In data (text corpus) compilation, four main types of biases have to be taken into account. Resource bias, referring texts being more representative and frequently produced by populations of means. It can be expected that resources as such will not be a major concern with press releases of the federal government, as all sections of it undoubtedly possess the means and incentive to publicise their press releases. It can potentially be pertinent to the auxiliary corpus of state-level press releases, where the less endowed agencies could conceivably suffer from levels of funding that would not allow for fully digitised records. (Grimmer, Roberts, Stewart 2022, p. 87)

Incentive bias refers to texts being selectively published or retained, while in theory to be discarded, with data retention being prescribed in legislation, in praxis it can be expected that agencies can make (potentially selectively) materials that were created during the previous political administration more challenging to access, caution must then be exercised in automated scrapping techniques are they to capture data for the whole researched period. (Grimmer, Roberts, Stewart 2022, p. 88)

Medium bias is the bias introduced in technical or normative limitations of the medium by which text is shared. Technical limitations in the sense of text length or formatting are not expected to be of concern, as virtually all institutions use self-hosted or cloud servers with no practical limits as to the length of the text when serving press releases in the HTML format, as all institutions do. Medium bias, in a broader sense, can be present as the press release is a specific medium, where the institution can be presumed to want to inform primarily about what they expect to be considered positive “news” should they not be statutorily obliged to provide that information. (Grimmer, Roberts, Stewart 2022, p. 89)

Finally, retrieval bias is the bias that occurs when texts are selected based on statistical methods from a larger corpus of texts. That ought to be largely avoided in the research of this thesis, as the dataset is, in theory, ought to be time-delineated, and eventually, all

available texts will be classified. Where its presence is possible is in the intermediate step of the dataset that will be prepared for model supervised learning, which will be partly compiled via dictionary search, but random sampling from the whole dataset will be applied to counteract that in further populating the subset. (Grimmer, Roberts, Stewart 2022, p. 90)

For the actual classification, evaluation and validation (*the thesis steps three to eight*), due to the corpus size, potential methodological contribution and recent technological advances in natural language processing, language models are utilised with the primary goal (*the thesis steps six and seven*) of fine-tuning a language model to classify the whole corpus. For that, a BERT (Bidirectional Encoder Representations from Transformers) latest descendant, DeBERTaV3-large, is utilised. (He, Liu, Gao, Chen 2021)

Building upon the technique of machine learning, the artificial intelligence method of gradual iterative approximation of hidden statistical function based on the training data, with the goal of generalising the algorithm to outside of the training populations. A variant of machine learning, now perhaps the most widespread, is the artificial neural network (mechanism sometimes referred to by the term “black box”, however as based on probabilistic programming, can be fairly easily reproduced, thus that term can be considered controversial, but might be more appropriate for the large language model operation), which consists of de facto at least three parallel artificial neuron node series (in the case of “true” artificial neural networks) that process input data by adjusting the weights via its non-linear activation function, i.e. to close the differential between the expected output and the predicted one, however the output of the input neuron is not immediately outputted as the part of the system result, but is then reingested by the so-called hidden layer (series) of artificial neurons, that process the input from the initial layer and follow the same process of trying to adjust to the expected output, finally the output layer ingests the predictions of the hidden layer and outputs the system prediction. (MIT 2017; Oracle 2020; Farkas 2021; IBM 2021b)

The connected term deep learning, which is widely utilised in language models, simply denotes an artificial neural network that incorporates at least two hidden layers of processing – a series of neural nodes. (MIT 2017)

Specifically, language models that will be worked with in this thesis utilise the Transformer architecture, where each neural node is considered an “attention head”, with each head concerned with a different function – text pattern – when attempting to devise the ideal “process” to arrive at the desired outcome in training. Specifically, BERT is trained by text being ingested into the network with a masked word, and the network attempts to predict the mask, with evaluation based on the actual unmasked text. To “mathematise” words, to be able to use machine learning in the sense of closing the differential between predicted and actual, the so-called embeddings are utilised, a vector positional representation of individual words (at that point tokens) in relation to the corpus, effectively retaining the semantics (word meaning). Transformers also take advantage of a concept called transfer learning – a model first being trained generally, as mentioned, which then can be specialised (fine-tuned) for a specific task while taking advantage of its general pre-training. (Mikolov, Chen, Corrado, Dean 2013; Rothman, Gulli 2022)

Such is the case with the main model employed to classify the dataset (*the thesis steps six, seven and eight*) where a generally trained model is taken, (DeBERTaV3-large) and then fine-tuned to recognise the latent presence “wokeness” in a classification unit as based on the training dataset. This approach is chosen as recent research has shown that fine-tuning a language model has demonstrated somewhat superior performance to other methods, such as zero-shot classification or simple support vector machine, when attempting to extract latent and more esoteric meanings from the text. While that might change with recent advances in decoder-only large language models and the advent of a more comprehensive unsupervised “zero-shot via prompt” approach, at the time of writing, existing literature (especially in social sciences) at the cusp of science utilises some permutation of encoder-only language model classification when after similar research goals as this thesis. (Vijay Srinivas Tida, Hsu 2022; Mahendru, Pandit 2024; Nasreen, Murad Khan, Younus, Zafar, Kashif Hanif 2024)

DeBERTaV3 is the most recent model in the encoder-only branch of language models, it utilises a novel technique of pre-training, slightly different than the usual mask language modeling with purportedly more efficient technique of pre-training with replaced token detection. (He, Gao, Chen 2023)

Zero-shot classification is incorporated in the methodology of the thesis in the classification of the subset of the corpus for the fine-tuning of the final model and in the evaluation of the classification of the final model (*the thesis steps number four and eight*). Zero-shot classification essentially uses pre-trained (and usually pre-fine-tuned, but on general data) models to classify text without any specific training or prior knowledge of the dataset (Hugging Face 2023). Widely used in sentiment classification, now models with an increasingly wide spectrum of classification abilities and reasonable accuracy are available as well. A variant of the DeBERTaV3-large (DeBERTaV3-large-zeroshot-v2.0) is used to effectively substitute for what, under different circumstances, could have been a second human coder (Laurer 2024). Specifically, the zero-shot DeBERTaV3-large, as an encoder-only descendant of BERT, can offer a similar *modus operandi* by which it was initially trained, i.e. guessing a masked word in a sequence, which in the case of zero-shotting is used to construct a hypothesis in which the classification codes for “wokeness” are contained and the model decides whether the masked word is effectively 1 or 0 (more classes are possible in other applications, topics, sentiments etc.), depending on the character of the ingested text. (Mahendru, Pandit 2024)

As mentioned, this approach is used in concoction with human coding (*the thesis step number four*) getting to two classifications for each unit of text; after evaluation of inter-coder reliability and other measures of agreement, it is decided to classify text to be “woke” only when both coders are in agreement. That can be justified by “wokeness” being a novel and, most importantly, a comparatively esoteric concept; thus, giving “benefit of the doubt” to its absence can potentially improve the external validity of the study in relation to training the model on the semi-manually classified subset of the corpus. Furthermore, it can aid in achieving the external validity of the study, especially considering the coding is done primarily by the “compiler” of the theory (the author of this thesis), and accusations of non-objectivity, even considering the machine control, could be warranted. Along with the crude measure of the frequency of agreement, the measure of Krippendorff Alpha is utilised in evaluating the inter-coder agreement, which, unlike mere percent agreement can measure whether agreement is above or below chance (*thesis step number five*). (Hayes, Krippendorff 2007; Grimmer, Roberts, Stewart 2022)

Moving towards step three of the thesis, the selection of the training corpus will be made followingly: a total sample of 1000 units will be collected, which, according to the

literature on the subject, appears near the “sweet spot” for fine-tuning BERT-like models for maximising accuracy (and other measures introduced shortly) (Mahendru, Pandit 2024), while feasibly co-classified by a human coder. Considering the theoretical assumptions, it might be presumed that units to be classified “woke” would likely constitute only a marginal part (sub-ten per cent) of the sample. If it was chosen purely based on chance, it is decided to collect 50% of the sample based on chance (pseudo-random algorithm) and 50% based on dictionary search of the dataset for terms that could be theoretically associated with “wokeness” contained in Table 1.

That should create a dataset that would tend towards more equal distribution of the binary classification, which can then have positive effect on the performance of the trained model.

There is an argument for utilising more advanced means in selecting the potentially “woke” (and, in fact, the full sample) section of the sample, such as Keywords in Context, however considering that the ingestor of the classified data is a language model, the more rudimentary method ought to yield sufficient results. Finally, the final subset is checked for the presence of duplicates, as that would hinder the training process. (Dun, Soroka, Wlezien 2021)

Table 1: Dictionary capture terms, based on the outlined theory (own work)

anti-racism	anti-racist	bipoc	climate justice
colonialism	cultural appropriation	critical pedagogy	decolonize
economic justice	educational equity	epistemic injustice	environmental equity
group identity	group rights	health equity	implicit bias
inequity	institutional racism	institutionally racist	intersectional
lived experience	lgbtq	queer	racial equity
systemic bias	systemic racism	white privilege	

The training or rather fine-tuning of the final model is done by taking the pre-classified training sample of 1000, splitting it into three (usually unequal) parts: training data, evaluation data and test data, while the inclusion of the test data split section is not technically necessary in the case of the thesis, as the model then can be judged on its performance on the whole rest of the unseen dataset, this split is preserved for initial considerations. The basic reasoning for this split is that the training split is used to train the model in parallel with evaluation data in tuning the model parameters, with the test data

being unseen by the model during the training, thus providing a “real-life” benchmark on the model performance, with the goal of having as generalisable performance as possible while avoiding over-fitting (the model being too closely “taught” on the training data, learning the data noise and losing ability to generalise). Parameters of fine-tuning, such as the rate of learning and epochs (iterations), are then adjusted to maximise the model fit while avoiding the aforementioned over-fitting. (Rothman, Gulli 2022; Eckhard, Jankauskas, Leuschner, Burton, Kerl, Sevastjanova 2023)

The indicators of classification accuracy that now can be taken advantage of, taking the pre-classification as a de facto “gold standard”, are simple accuracy, precision, recall and F1-score. “Simple” accuracy is the fraction of correct predictions, and all predictions; precision is the fraction of true positives and the sum of true positives and false positives; recall (especially in the medical field, often termed sensitivity (Johns Hopkins School 2020)) is the fraction of true positives and the sum of true positives and false negatives and finally the F1-score which is a harmonic mean of precision and recall. Considering the nature of the dataset and the assumed imbalance of classification classes (“woke”/“not woke”), specifically when evaluating the whole dataset, the primary score by which the performance of a model will be judged is the macro F1-score when the F1-score is first calculated for each class separately (the F1 for “woke” classification and the F1 for “not woke” classification) and then averaged, as that ought to give the most accurate assessment of the performance of the model. (Krippendorff 2004; Hayes, Krippendorff 2007)

In the penultimate *thesis steps number seven and eight*, the fine-tuned model is run over the whole corpus of text, classifying autonomously every unit of text. Considering the technical limitations, the whole size of the dataset, and the fact that it is not wholly pre-labelled, a novel evaluation technique has been devised. Similarly to when gathering the training sample for fine-tuning the model (but now without the dictionary component), a sample of 1000 now fine-tuned-model-classified text units is pseudo-randomly sampled from the whole classified corpus; they are manually classified by a human coder and by the zero-shot model, when agreement on “woke” is reached, one is encoded when there is disagreement or agreement on “not woke” a zero is encoded. That data is then taken as another “simulated gold standard”, and metrics are computed against the fine-tuned model results, and the model as such is evaluated. (Grimmer, Roberts, Stewart 2022)

Finally, in the *thesis step number nine*, to each text unit a range of binary dummy variables is assigned to denote various institutional attributes, as outlined in Table 2.

Table 2: The presented dataset structure (own work)

Name of column	Column description
title	Title of the press release
body_text	The text unit classified
date	Date of the press release
link	URL link to the source of the press release
name	Unique name of the institution
is_federal	Binary dummy, whether originator is part of the U.S. federal gov.
is_state	Binary dummy, whether originator is part of a state gov.
is_university	Binary dummy, whether originator is part a university.
is_democrat	Binary dummy, whether originator is controlled by the Democratic party.
is_republican	Binary dummy, whether originator is controlled by the Republican party.
is_independent	Binary dummy, whether originator is an independent federal agency.
is_head	Binary dummy, whether originator is head of pertinent executive.
final_label	Final classification label by the fine-tuned model
final_probability	The certainty of the model's classification

The binary encoding is chosen to best capture the character of the institution and to have as little necessary overlap with other categories to possibly draw statistical conclusions. For example, in the case of a White House press releases of George W. Bush, the encoding would be as follows: `is_federal=1`, `is_state=0`, `is_university=0`, `is_democrat=0`, `is_republican=1`, `is_independent=0`, `is_head=1`. In regards to independent agencies, while there is a partisan oversight board and a partisan chairman of the board appointed by the president in case of FTC (2013), and at pleasure serving head of the EPA (2013), the abstraction of non-partisanship is used, as the primary focus is their at least partial policy

insulation from the executive branch sections directly under the control of the president. Other cases are believed to be self-explanatory.

Prior to the statistical analysis, the results are first broken down on a per-year basis. They are then visualised both in aggregate and separately based on the binary encodings. This process is undertaken to draw some initial conclusions in the context of the research question.

Ultimately, two generalised linear model with logistic regression link are constructed, with the response variable taking either value 1 for “woke” and 0 for “non-woke” (binary - binomial) one model of only independent variables of the binary dummy encodings and one including the year of release of the press release as a fixed effect.. This form of regression appears to be the most potentially suited for the character of the variables, along with their presumably and logically non-normal distributions and the binary character of the dependent variable. (IBM 2021a) Probit link function, while also feasible, as it is suited for binomial response variables, is not chosen, as the assumption that the binary response variable is representation of normally distributed latent continuous variable would be most likely violated.

The findings are discussed based on the results of the model, the two models are compared, and the one with better values of pseudo- R^2 and lower AIC and BIC values is chosen for further interpretation.

All programming is done either in the Python programming language, or the R programming language and all code needed for reproducibility is enclosed with this thesis.

3 Research

3.1 Devisement of Classification Criteria

As outlined in the methodology section, for proceeding with the thesis, a set of criteria for classification of individual press releases is to be constructed based on the methodology of content analysis. (Krippendorff 2004)

Drawing on the findings and conceptualization of the previous sections a four point set of criteria is devised roughly eclectically covering intersection of what authors of “wokeness’ as an Ideology” would term as “woke” especially in connection to governmental policy.

The text unit is to be classified as “woke” If:

- Emphasizes the unique importance of group identity (e.g., race, gender, sexual orientation) in politics, culture, or life, and/or
- Advocates for policies or viewpoints that aim to achieve equity (equality of outcome), and/or
- Calls for special protections for historically or presently “oppressed” groups, especially against potential harm from speech, and/or
- Endorses the implementation of diversity, inclusion, and equity policies, and measures in line with critical social justice.

In any other case the text unit is to be classified as “not woke”.

The classification criteria is conceptualised broadly, but with distinction to liberal conceptualisation of group identities as in pluralism.

3.2 Data Collection and Characterisation

Data – corpus – collection was conducted primarily in May and June 2024 with error correction and additional batch of collection at the beginning of July 2024. The bulk of collection was done via web scraping technique on the publicly accessible archives of the pertinent institution.

At first, all top-level executive departments, independent agencies, and archive white house sites were analysed to determine the presence of a press release archive, and then individual techniques for scraping were devised. From the executive departments, unfortunately, only the Department of Justice, Department of Labor and Department of Treasury offered publicly accessible (albeit with varying difficulty) archives that stretch past a single administration. With the others, such as the Department of Interior, at first, an informal request was sent out that unfortunately resulted in no help offered; furthermore, when at least a response was received (such as was the case with the Federal Reserve), it recommended lodging a freedom of information request. Considering the author of this thesis is based outside of the U.S. jurisdiction and does not possess a U.S. address, it proved effectively impossible to lodge a request based on the FOIA act, as all institutions, where it could be of use, require an address based in the United States for processing. While it would be conceivable to lodge the request via non-automated means, it automatically adds to processing times and could require an appeal; coupled with the necessity of fulfilment of necessary fees, this approach has proven, at best, impractical.

Thus taking in account medium and retrieval bias now inherent in the dataset, on the federal level, institutions that have proven to be researchable are: the Department of Justice (DoJ 2018; 2024), the Department of Labor (DoL 2024), the Department of Treasury (Treasury 2024), the Bush White House (The White House 2009), the Obama White House (The White House 2017), the Trump White House (The White House 2021a), the Biden White House (The White House 2024), the Federal Trade Commission (FTC 2024) and the United States Environmental Protection Agency (EPA 2024).

Interestingly, every section of the federal government uses radically different website format, including the HTML structure of the archive (no use of RSS or limited to most

recent press releases (Chen 2014)) and of the press releases themselves, not to mention no "neat" automated wrapper is usable, for every agency a different HTML scrapper must be individually coded and utilised. In some cases, such as that of the Department of Justice, several scraping utilities must be devised, as every year of the archive has a slightly different structure. After overcoming the issue of being blacklisted due to excessive parallelised scraping by the Department of Justice, the acquiring of the federal portion of the data was largely successful with the aforementioned caveats.

Broadly analogous is the case with the auxiliary corpus; however there is more widespread issues with completely missing datasets prior to the particular time period (habitually prior to 2015); despite these challenges, it is decided for inclusion even as the corpus does not cover the whole period in the case of state institutions, as the auxiliary corpus serves, as the name would indicate, auxiliary purpose of a partial control in answering the research question. However, the issue of time coverage is not significant with universities, as they provide comprehensive coverage.

The choice in state institutions, as was reasoned before, attempts to somewhat mimic the agenda of the federal government, as to provide a reasonable comparison. State institutions that were possible to scrap and compile are as follows: New York governor (NY Governor 2024), Texas governor (Texas Governor 2019; 2024), New York Attorney general (NY AG 2024), Texas Attorney general (Texas AG 2024), Texas Department of transportation (Texas DoT 2024), New York Department of transportation (NY DoT 2024).

Universities-wise, the news releases of Princeton (Princeton 2024) and Harvard (Harvard 2024) were successfully scrapped.

Table 3: Number of press releases scrapped by origination institution (own work)

Name	Number of press releases
epa	32182
doj	26972
wh_obama	14404
harvard	14253
dol	13569
wh_bush	12865
treasury	10806
wh_biden	10748
princeton	9293
ny_aj	8725
wh_trump	8378
ftc	8302
ny_gov	6730
tx_gov	5669
ny_dot	1820
tx_aj	1175
tx_dot	378
Total	186269

As is noticeable from Table 3, in total, 186269 units of text (press releases) were gathered from the aforementioned sources. By quantity alone, there is noticeable disproportionality towards the sections of the federal government, which is not necessarily an issue in itself but should be taken into account when further evaluating the statistical model at the very end of the research.

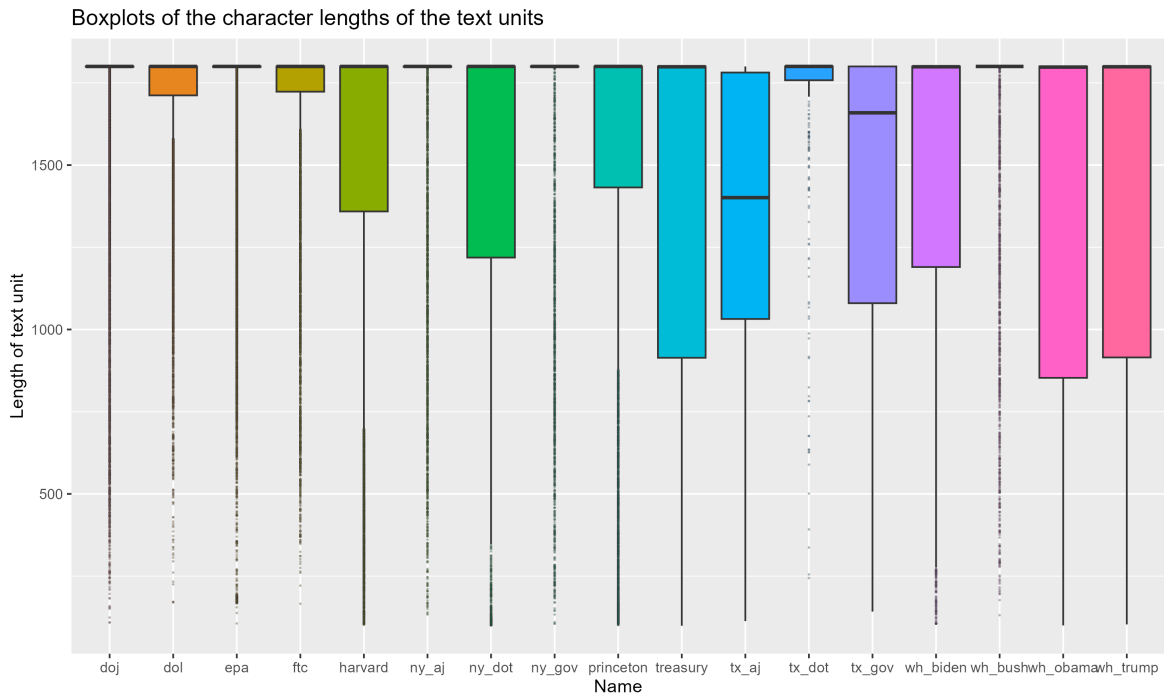


Figure 4: Boxplots of the lengths of the text units after scrapping (own work)

Regarding the length of text distribution post-truncation, it can be observed in Figure 4 that most institutions' press releases have distribution with median basically identical or very close to 1800, thus indicating they were either around or longer than 1800 characters pre-truncation and all but three have 1st quartile past 1000 thousand characters, considering the ingestion limits of the zero-shot model being around one page (1800 characters then tokenised) and of the model to be fine-tuned 512 tokens (around 2000 characters maximum) that ought to serve the purpose of classification validly.

Table 4: Temporal distribution of the press releases by year (own work)

Year	Number of press releases
2001	5717
2002	6240
2003	6920
2004	7485
2005	7025
2006	7036
2007	6387
2008	7345
2009	4876
2010	8124
2011	7452
2012	6752
2013	6332
2014	7904
2015	8320
2016	7895
2017	7740
2018	8827
2019	8548
2020	9616
2021	10823
2022	11485
2023	11589
2024	5831

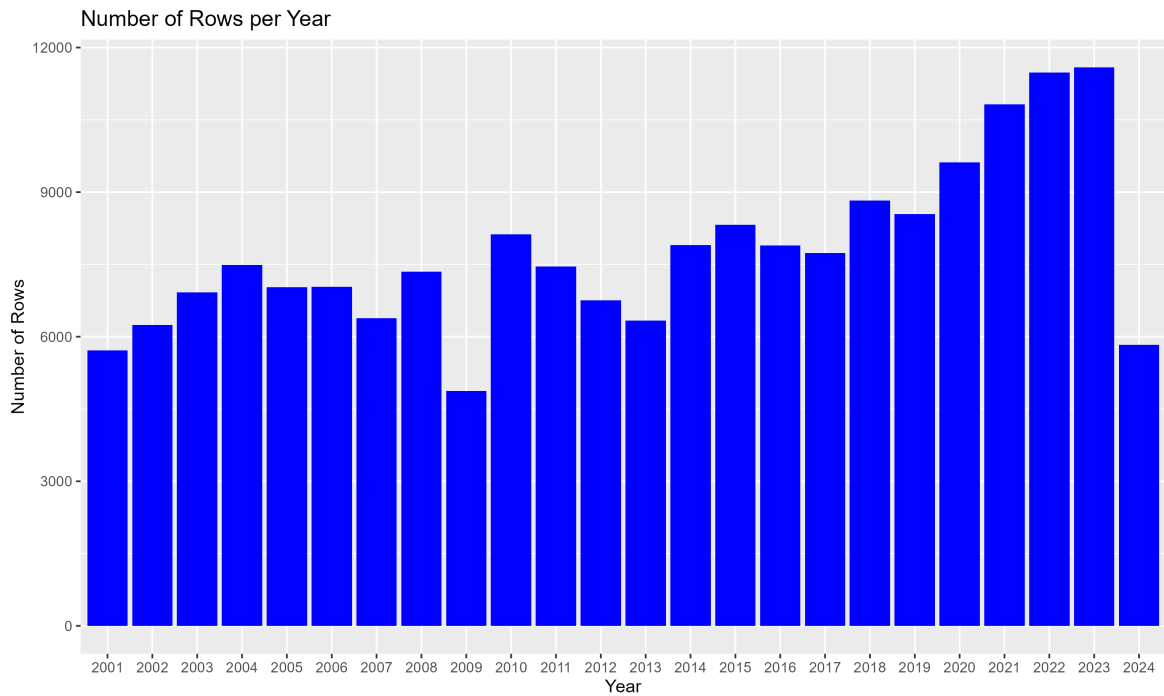


Figure 5: Temporal distribution of the press releases by year (own work)

Pertaining to the temporal distribution (Table 4 and Figure 5), an almost surprisingly uniform distribution is encountered, with a somewhat more noticeable number of units in the years 2021, 2022, and 2023 coinciding with the advent of the Biden administration. Naturally, there is a smaller number of entries for the year 2024, as only half of the year is taken into account. Oddly, the year 2009 has about a quarter fewer entries than either 2008 preceding or 2009 following it. After manual revision of the dataset for possible issues that could have been created by incomplete scraping (potentially of the Obama White House), that is discounted, and the discrepancy remains unexplained; however, it is not of the magnitude that could impede research. Mild general retrieval bias can be considered to be present as more recent years do possess comparatively more entries, but that trend could have been expected to be stronger.

3.3 Training Corpus

A training corpus of 1000 entries with half of the sample coming from random selection (500 units of text chosen pseudo-randomly based on the global random seed of 42) and the other half of the sample is coming from a dictionary capture based on theoretically relevant keywords denoted in Table 1 is composed. As the dictionary capture results exceed 500 matches, a pseudo-random sample of the desired size is taken from the matches again with the global random seed of 42. The combined sample is then first to run through DeBERTaV3-large-zeroshot-v2.0 classification with a slightly adjusted hypothesis containing the basis of the classification code in:

"Does this text exhibit {} values, including a focus on identity politics (race, gender, sexuality), systemic oppression, and a commitment to social activism as outlined by critical social justice with keywords such as diversity inclusion and especially equity meant as equality of outcome."

Experiments (prior to classification) were conducted with various versions of the hypothesis, including narrower versions. However, those were producing facially invalid results.

Consequently, the output is masked, and the same sample is given to a human coder to classify it based on the criteria. Then, at first, statistics of inter-coder reliability are produced.

Table 5: Inter-coder reliability metrics of the training sample (own work)

Metric	Value
Krippendorff's alpha	0.552
Percent agreement	77.6%

Considering both the 78.4% pure percent agreement and the for-chance-agreement accounting Krippendorff's alpha (with interval results from -1 to 1), it can be concluded that the coding is relatively reliable, and can be proceeded further.

In the next step the specific algorithm for applying “woke” label only when both coders agree (effectively an “AND” logic gate) is applied to arrive at the fine-tuning training dataset.

Table 6: Label distribution of the training sample (own work)

Label	Count
not woke	624
woke	376

While having a slight imbalance in the training data, skewed toward “not woke”, it appears a sufficient number of “woke” units are included for achieving a reasonably accurate fine-tuned model.

3.4 Fine-tuning the Model

For fine-tuning DeBERTa-v3-large itself, the training dataset is split into three parts - 700 units for training, 150 units for evaluation, and 150 units unseen for preliminary testing. A somewhat higher training proportion of the dataset is split to at least partially counter the imbalance in the whole subset, i.e. that the model “sees” enough variation of “woke” units to be able to reliably “recognise” its variants. As to the training attributes, the learning rate of $1e-5$ is chosen based on previous research and experimentation, the learning rate being the rate at which the model converges to the ideal solution and *two* training epochs - how many times the model goes through the dataset. As outlined in the methodology section, these two values are tuned to achieve the optimal fit, while avoiding over-fitting; the model is extremely performant on the training data but not able to generalise on unseen data.

Input text is tokenised, and the model is let to trained on Graphical Processing Units (GPUs) to maximise the parallelisation of the process.

Table 7: Performance metrics of the fine-tuned model on unseen test split in training (own work)

Metric	Value
Accuracy	0.933
F1 Macro	0.927
Accuracy Balanced	0.927
Precision Macro	0.927
Recall Macro	0.927

Taking the preliminary data from Table 7, it can be concluded that the fine-tuning was successful, as the model performs “just right” on the unseen units, achieving F1-score (macro) of 0.927, which is comparatively very good, considering the limitations of the dataset and other factors. Preliminarily it can be said that the model’s fit appears to be optimal. Increasing the learning rate would increase the F1-score to above 0.93, but then considering the complexity of the deep-learning model and comparatively small learning sample the F1-score around 0.91-0.92 in the pre-validation stage should be taken as optimal as to avoid over-fitting, considering this metric is taken only from performance on about 150 observations of unseen units.

3.5 Classification and Evaluation

Followingly the whole dataset is classified by the newly fine-tuned DeBERTaV3-large. The evaluation of the model performance is done via the novel method outlined in the methodological section and somewhat analogously to the preparation of the training sample. A pseudorandom sample of 1000 units of text is taken from the whole classified corpus, which is then classified by a zero-shot model and human coder separately.

Table 8: Classification distribution between “coders” (own work)

	Human classification	Fine-tuned model label	Zero-shot model classification
not woke	983	979	818
woke	17	21	182

Table 9: Inter-coder reliability metrics of the evaluation sample (between zero-shot model coder and human coder) (own work)

Metric	Value
Krippendorff's alpha	0.080
Percent agreement	83.5%

As can be observed from Table 8 and Table 9, there is a large disparity between the human coder and the zero-shot model at first glance. The zero-shot model seems to significantly “overshoot” what texts are “woke” when working with a random sample; on the other hand, the fine-tuned model appears to be much closer to the human coder. Regarding inter-coder reliability between the human coder and the zero-shot model, per cent agreement-wise, that is even higher value than there was with the training sample. However, Krippendorff’s alpha is significantly lower, just barely reaching positive numbers. That is most certainly caused by the great disparity in the zero-shot model classifying significantly more entries as “woke”, so while there is overlap between what the human coder classifies as “woke” and what the model considers “woke”, the “gung-ho” approach of the zero-shot causes the alpha to be barely above chance.

Fortunately, possible “broadness” is taken into account in the methodology, and to create a simulated “gold standard”, the “AND gate” approach is applied, thus limiting the possible issues with one of the coders being too over-broad. With the final “gold standard” measure being the evaluation standard, final metrics of performance can be applied.

Table 10: Performance metrics of the fine-tuned model on the evaluation sample (own work)

Metric	Value
accuracy	0.994
f1_macro	0.920
accuracy_balanced	0.968
precision_macro	0.880
recall_macro	0.968




Figure 6: Confusion matrix of the evaluation sample, true label is the “gold standard” and predicted label is the label predicted by the fine-tuned model. (own work)

Evaluating the final fine-tuned model’s performance lets us understand in more depth why the F1-score macro is the best-suited tool for evaluation in this case. As very imbalanced categories are encountered (according to the “gold standard”, “woke” only has 16 occurrences in the sample compared to 984 “not woke”), the model performs theoretically exceedingly accurately, getting around 0.996 in accuracy, but that tells only part of the picture, as taking the categories separately, as illustrated in Figure 4’s confusion matrix, the model “overshot” the number of “woke” cases by about 25% to 30% per cent. However, as the F1-score macro is well above 0.9, the model can be generally considered performant, especially taking the somewhat esoteric character of “wokeness” into account. A brief manual overview of incorrectly labelled cases as “woke” truly supports that notion, as virtually all are somewhat concerned with group identity and its position in the legal and political landscape, but more so from the liberal perspective, stressing the equality of opportunity rather than of outcome, which will be expanded upon in the discussion section of the thesis. Taking the classified dataset as empirically and facially broadly valid, it can be advanced to further analysis.

3.6 Statistical Analysis

First, mapping the presence of “wokeness” across all surveyed institutions in the chosen time frame:

Table 11: Yearly distribution of “woke” press releases across all institutions (own work)

Year	Number of “woke” press releases	Total number of press releases	Fraction of “woke” press releases
2001	62	5717	0.01
2002	66	6240	0.01
2003	79	6920	0.01
2004	66	7485	0.01
2005	69	7025	0.01
2006	93	7036	0.01
2007	89	6387	0.01
2008	99	7345	0.01
2009	94	4876	0.02
2010	135	8124	0.02
2011	110	7452	0.01
2012	86	6752	0.01
2013	91	6332	0.01
2014	139	7904	0.02
2015	176	8320	0.02
2016	179	7895	0.02
2017	114	7740	0.01
2018	148	8827	0.02
2019	151	8548	0.02
2020	213	9616	0.02
2021	732	10823	0.07
2022	744	11485	0.06
2023	783	11589	0.07
2024	357	5831	0.06




Figure 7: Bar chart, with local regression line, of the fraction of “woke” classifications per year across all institutions (own work)

From the information contained in Table 11 and Figure 7, a rough increasing trend can be identified over the surveyed period. Starting at around 1% of all press releases containing “woke” notions in 2001 (the bottom) and peaking in 2021 with around 7% of all press releases containing some aspect of “wokeness”. Correlation with the party in power of the federal government seems to be fairly evident, with lower values over the Bush administration, picking up somewhat after the inauguration of President Obama, peaking hitherto in his second term, only for the proportion to fall significantly when President Trump comes to power. An interesting deviation is 2020, which was during the Trump administration but reaches levels of the late Obama administration, which would be largely in line with Lindsay’s (2022a) hypothesis of the so-called George Floyd wave when even conservative (at least temporarily and selectively) institutions adopted the language of “wokeness” (Goodwin 2023b; 2023a). Nevertheless, undoubtedly, the greatest rise is seen after Biden(-Harris, as President Biden would often term it (The White House 2021b)) administration comes to power, with only a slight decline from 2023 going to 2024. To further examine the trend, a separate graph and table of only non-independent parts of the federal executive is to be scrutinised.

Table 12: Yearly distribution of “woke” press releases across non-independent federal institutions (own work)

Year	Number of “woke” press releases	Total number of press releases	Fraction of “woke” press releases
2001	10	2647	0
2002	9	3066	0
2003	12	3278	0
2004	7	3891	0
2005	9	3039	0
2006	8	3390	0
2007	11	3364	0
2008	7	3570	0
2009	13	2843	0
2010	49	4716	0.01
2011	31	4308	0.01
2012	31	3769	0.01
2013	29	3373	0.01
2014	58	4464	0.01
2015	63	4594	0.01
2016	75	4378	0.02
2017	20	4077	0
2018	17	4933	0
2019	18	4415	0
2020	34	4475	0.01
2021	443	5572	0.08
2022	433	6181	0.07
2023	432	6283	0.07
2024	199	3116	0.06




Figure 8: Bar chart, with local regression line, of the fraction of “woke” classifications per year across all non-independent federal institutions (own work)

In Figure 8 and Table 12, the trend of elevation of “wokeness” when the Democratic party is in power, or rather in control of the presidency and diminishment when a Republican is in control, is underlined. During the Bush years, the notions appeared to be especially absent; that, however, that might be due to collateral factors influencing the whole society, that race and other identities did not play a major role in the societal discourse, and geopolitical realities of the war in Iraq, Afghanistan and the war on terror as a whole played prime, along with the associated discourse of security from the side of the administration and of civil liberties of its opponents (Daalder, Lindsay 2001). It can be considered somewhat ironic that the year of the event that is most associated with questions of race during the Bush presidency, 2005, and Hurricane Katrina, and the reaction (or rather absence of reaction according to some (Walsh 2015)) contains not significantly more “woke” classified units than any other year.

Noticeable uplift can be observed in 2010, a year after President Obama came to office, perhaps somewhat surprisingly, as his candidacy and first term were widely considered “centrist” and almost “post-racial” when some theorists situate the true “outbreak” of wokeness in 2012 and Obama’s second campaign and consequent governance. Going into his second term, President Obama largely left the centrist-liberal coalition behind (James

2009), and the election of 2012 birthed the multi-ethnic/intersectional plus white college-educated coalition that came to be a feature of its campaigning and governance priorities at least until 2020 (García Bedolla, Haynie 2013). With that ought to have come an uptick in “wokeness” as some situate it; however, while there is some evidence to it, the difference between 2010/2011/2012/2013 is somewhat lesser than anticipated, and the relative rise in “woke” rhetoric seemed to have appeared later, in 2014, perhaps in connection to the case of Ferguson (Kessler 2019).

With the advent of the Trump administration, 2017 brings about ratios that are very slightly higher than in the Bush years, and that is until the year 2020 when there is something of an uptick to the level of the Obama administration in 2011. Now, with only the pertinent institutions in focus, the hypothesis of the “George Floyd” effect can be preliminarily pondered. However, the first year of Joe Biden’s presidency is when there is a truly beyond confidence interval increase, which is more so in line with Yang’s theory of the so-called “Year Zero” and the greatest advent of “wokeness” in the federal government, which appears to have peaked in 2021, but has retained a steady level since, with slight down-tick in 2024.




Figure 9: Bar chart, with local regression line, of the fraction of “woke” classifications per year across universities (own work)

Fascinating contrast is provided when attention is given to the broader presence of “wokeness” in society as measured by the press releases of US universities. There, and now in line with Sullivan’s thought (Sullivan 2020a), we see a far more uniform trend with a comparatively steadier increase (naturally) regardless of current political representation. Interestingly, there is a noticeable “dip” during the first six years of the Obama administration from the relatively uniform distribution in the Bush years, finally picking up in 2014/2015 into an almost stable year-on-year rise until 2021, with stagnation in 2022 and a final decrease in 2024 to pre-Biden levels. Compared to governmental institutions, the ratio levels are significantly higher and in peak years constituting almost 20 per cent of all universities’ press releases.

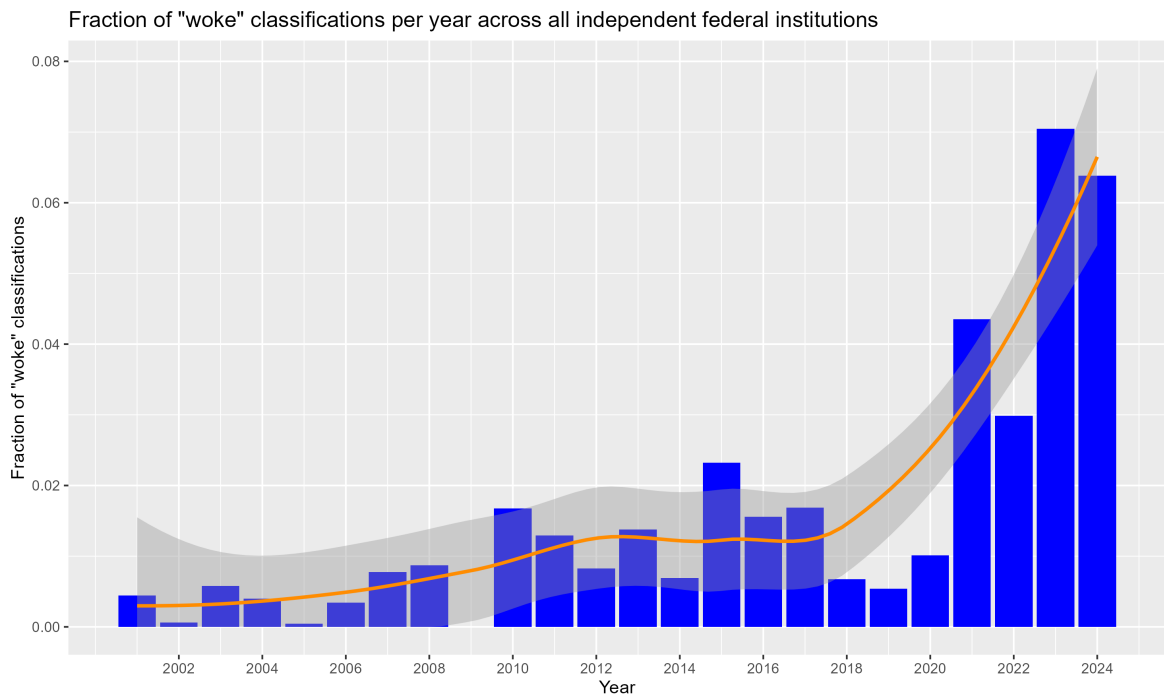


Figure 10: Bar chart, with local regression line, of the fraction of “woke” classifications per year across all independent federal institutions (own work)

The trend in independent agencies of the federal government (namely EPA & FTC) appears to broadly follow the trend of non-independent sections of the federal government, perhaps with somewhat delayed changes given their governance structure and perhaps a certain volatility can be observed, especially in the years 2021-2024, with an unexpected peak in 2023, that cannot be simply explained by missing data. In summary, it can be concluded that there is, however, no significantly higher tendency of independent agencies to display “wokeness” in their press releases than in other sections of the federal government. The t-test on the series of fractions from both independent and non-independent sections of the federal government yields a p-value over 0.8535, which points however more to the difference in level rather than in trend.

Simple logistic regression model - Dependent Variable = wokeness binary		
(Intercept)	-4.131 ***	(0.062)
is_university	1.529 ***	(0.067)
is_federal	-0.060	(0.047)
is_head	0.833 ***	(0.042)
is_democrat	0.392 ***	(0.055)
is_republican	-1.855 ***	(0.087)
N	186269	
AIC	41595.912	
BIC	41656.722	
Pseudo R2	0.088	

Standard errors are heteroskedasticity robust. *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.

Table 13: Simple logistic regression model (own work)

In Table 13, the results of the simple logistic regression model can be observed; as mentioned in the methodology section, binomial logistic regression is chosen to predict the probability that an observation is either 1 for “woke” and 0 for “not woke”, with the dependent variable being the binary classification, and independent variables being the institutional attribute dummies. While the dataset contains two further variables, *is_state* and *is_independent*, they are eventually used only for ease of programming filtration, as in modelling, they are describable by combinations of values of other variables.

The chi-square statistic (not displayed in Table 13) with a p-value under 0.001, thus the model as such being statistically significant. While the imperfect dataset disallows making broad conclusions about what factors go into the presence of “wokeness”, the model could be used to make cautious observations about what types of institutions might be more prone to its inclusion. However, before it is proceeded, the logistic model including the variable *year* as a fixed effect will be presented.

Logistic regression model with FE year - Dependent Variable = wokeness binary		
year2001	-5.380 ***	(0.143)
year2002	-5.414 ***	(0.138)
year2003	-5.244 ***	(0.130)
year2004	-5.461 ***	(0.139)
year2005	-5.460 ***	(0.134)
year2006	-5.105 ***	(0.122)
year2007	-4.954 ***	(0.125)
year2008	-4.934 ***	(0.119)
year2009	-5.170 ***	(0.124)
year2010	-5.101 ***	(0.113)
year2011	-5.212 ***	(0.120)
year2012	-5.344 ***	(0.129)
year2013	-5.229 ***	(0.128)
year2014	-4.937 ***	(0.110)
year2015	-4.647 ***	(0.104)
year2016	-4.593 ***	(0.102)
year2017	-4.613 ***	(0.116)
year2018	-4.416 ***	(0.101)
year2019	-4.359 ***	(0.102)
year2020	-4.072 ***	(0.088)
year2021	-3.303 ***	(0.073)
year2022	-3.356 ***	(0.075)
year2023	-3.273 ***	(0.076)
year2024	-3.373 ***	(0.086)
is_university	2.148 ***	(0.070)
is_federal	0.263 ***	(0.049)
is_head	0.559 ***	(0.043)
is_democrat	0.127 *	(0.057)
is_republican	-1.509 ***	(0.087)
N	186269	
AIC	39177.137	
BIC	39471.050	
Pseudo R2	0.148	

Standard errors are heteroskedasticity robust. *** p < 0.001; ** p < 0.01; * p < 0.05.

Table 14: Logistic regression model with fixed effect year variable (own work)

The model including the variable of the year of publishing of a given press release as a fixed effect presented in the Table 14, appears to have superior metric of AIC and BIC (lower than simple model) along with higher pseudo-R². Furthermore it appears to bring all predictor variables into significance at the confidence level of 95% (p-values under 0.05), and all but one (is_democrat) to significance at the 99.99% confidence level (p-values under 0.001). Thus for further analysis the model of Table 14 will be utilised.

The coefficients being in log-odds format, for an easier interpretation, conversion to odds ratio will be used. The conversion being easy powering of e by the coefficient.

In the case of utilising year as a fixed effect, there is no universal intercept and effectively the year variable (taking value in every observation) serves as the odds baseline. A similar trend that could have been observed in the Figure 7 can be noticed, with progressive years increasing the baseline odds of press release being “woke”, with the most significant coefficient “jump” between 2020 and 2021, and perhaps surprisingly the highest coefficient found in the year 2023 rather than in 2021, as could have been guessed from the mere visualisation of the data.

For the independent dummy variable is_university the conversion $e^{2.148}$ is approximately 8.564 signifying the chance for a press release originating from a university being “woke” is about 756.4% likelier than the year baseline. That is in line with both the hypothesis posited by Sullivan (2020a), Lindsey (2022a) and Kaufman (2022), and the observations in Figure 9.

The predictor is_federal’s conversion comes to about 1.30, odds thus pointing to about 30% higher chance of “woke” press release from a federal entity, which is somewhat significant, but considering the underlying dataset very “heavy” on federal entities, that might not be as persuasive.

Variable is_head has converted odds of 1.75, which largely confirms the initial assumption of “wokeness” being perhaps even more present in potential campaign rhetoric (not in the sense of Cammaerts (2024), however, rather as a positive, normative vision) as executive heads, especially in federal government (as they are not limited by Hatch Act, and presidential press releases contain among other things presidential speeches) can conduct campaign from office. However, the higher likelihood – about 75% likelier for a

presidential press release to be “woke” – probably also has a lot to do with the disproportional ratio of “woke” press releases in the Biden administration, and that is even inclusive of some controls.

The remaining two variables, `is_democrat`, have an odds conversion of about 1.135, and the variable `is_republican` of about 0.221, the only measure except year fixed effect, to have a negative log-odds and consequently sub-1 odds ratio. The affiliation to the Democratic party then means about a 13.5% higher chance of a press release being found “woke”. That is particularly interesting, as higher odds could have been expected while still above baseline; the somewhat lower than expected odds can be partly explained by the inclusion of formally non-partisan but “rich” in wokeness universities corpus, along with the (perhaps controversial) classification of EPA as independent non-partisan. Furthermore it might point to the significance of “wokeness” developing in time and not being a feature of the ideology of the Democratic party. On the other hand, the affiliation with the Republican party carries with it about 78% lower chance of a press release being “woke”, which could point to significant polarisation on the issue and “wokeness” not being the institutional consensus at least in the US, as Goodwin (2023a) sees in the UK.

3.7 Discussion

The research can be considered a success as the initial goal of developing a fine-tuned language model to classify press releases based on them containing notions associated with what has been presented as “wokeness” in the theoretical part of this thesis. The model attained a macro F1-score of 0.927 in a typical run on a test split and in an extended novel evaluation technique attaining a macro F1-score of 0.92. Generally, it can be observed that the model performs sufficiently well, however somewhat better in the “not woke” category, which can be partly explained by the somewhat specific conceptualisation of “wokeness”, where there is often a semantically rather thin line between “woke perspective” on a topic as opposed to “not woke perspective”, that can be perhaps best demonstrated on a following example from the evaluation dataset, that the fine-tuned model classified as “woke”, while the “gold standard” was “not woke”.

“Harvard University continues to make progress in strengthening the impact of its efforts to prevent and respond to sexual and gender-based harassment and assault. This is

according to the shared fiscal year annual report from the Title IX Office and the Office for Dispute Resolution (ODR), which was released today. [...], we remain committed to exploring innovative approaches and examining and implementing new and promising practices in our ongoing work toward creating a safer and more inclusive environment for everyone here at Harvard.” The Title IX Office leads the University's efforts to provide educational programming and resources to prevent sexual and gender-based harassment while also overseeing responses to disclosures through its system [...] A wealth of new initiatives was also piloted, including the University [...]” (Abridged and corrected compared to the raw data version; (Harvard University 2019))

There is a potential keyword of “inclusive” and a very theoretically relevant keyword, “gender,” present, which is most likely “pulling” the model towards the “woke” classification. However, the text contains no notion of identity-dependent enforcement nor advocates for the policy based on the disproportional impact sexual harassment has on a specific sex or another identity group. It can be conceivably argued that gender-based violence is a special protection category, but that is not realistically the case, as the formulation is specifically set in a way not to highlight the protection of a specific identity group (women). Thus, the “not woke” classification is warranted.

On the other hand, the model classifies texts such as the following correctly as “woke”.

“Black Americans with disabilities are deserving of dignity, respect, and equal opportunity to achieve the American Dream.[...] This country was founded upon the idea that all people are created equal and should be treated equally. However, it is important to acknowledge that racism and ableism together are longstanding barriers to the full participation and independence of Black people with disabilities. America has made significant progress, but there is more work to do to become an inclusive, accessible, and equitable nation.[...] the Biden-Harris Administration has taken historic action to advance equity through Executive Orders, administrative actions, and the provision of funding and resources.[...] [A]gencies are examining their policies and programs for disparities that have made it difficult for Black Americans and people with disabilities to thrive in their communities[...]. The Administration has delivered real change and will continue to work towards equitable outcomes and opportunities for Black disabled Americans.” (Abridged and corrected compared to the raw data version; (The White House 2023))

While this press release appears to be identity-focused, at the very start, it emphasises the notions of equality of opportunity and equal treatment, then it however, almost negates that notion as insufficient in praxis and advocates for and touts identity-based policy, along with the review of current policy that it does not produce “inequitable” outcomes. Coupled with the open call for “equitable outcomes” at the very end of the press release, the unit can be safely classified as “woke”, as the fine-tuned model correctly did.

Considering the model’s empirical validity as attested to by the macro F1-score and the facially valid approach as attested to by the human control in multiple stages of the classification process (“human in the loop” (Grimmer, Roberts, Stewart 2022)), but also with control against the possible bias inherent in single-coder classifications when a zero-shot model is employed. What does have to be acknowledged, however, is the bias present in the data collection, which is expanded upon in the previous sections, along with the modelling, especially pertinent, the overwhelming imbalance of the classification with the “woke” category consisting of a sub 3% of all labels. The label imbalance is partly ameliorated during the collection of the training dataset (Tanti, Van Der Plas, Borg, Gatt 2021; Dun, Soroka, Wlezien 2021).

The classification as a whole can be considered a success and produce data that can aid in answering at least in part the research question delineated.

“Provided that ‘wokeness’ is an identifiable ideology, what presence has it had in sections of the U.S. federal government, as measured by their external communications, since 2001 until 2024?”

Based on the findings presented here, it can be concluded that while there has been some “wokeness” present in press releases across all years (taking into account only non-independent sections of the U.S. federal government), there has been a marked rise since the inauguration of the Biden administration (and consequent appointment and confirmation of executive department heads) with the by far greatest year-on-year increase going from 2020 to 2021. This is in line with the hypothesis of Wesley Yang of “Year Zero”, where he claims that while there have been aspects of “woke” thinking in the Obama administration, the true rise in government he situates rightly into 2021. He posits that wider societal elites (especially educational elites) had been undergoing this attitudinal

shift into “wokeness”, as outlined in the theoretical section of this thesis, especially then in the year 2020 in what he calls “Racial Reckoning” (Yang 2021; Manhattan Institute 2023). This hypothesis is also at least partly supported by the findings presented here, as the universities’ fractions have already elevated during the Trump administration (2017-2021). While it can be drawn upon from the theory that prior to the Biden presidency, identity questions were framed in the broadly liberal fold by both parties and concepts coming from critical social justice were largely absent (Lindsay 2022a), what can be said, based on the research presented here, with certainty, is that 2021 brought with it significant change in the intellectual framework present concerning group identity and redistribution with it associated, and thus affirm the usefulness of conceptualisation of “wokeness’ as an ideology” since texts containing “woke” notions constituted around 7% of all press release output of the U.S. federal government at the peak. A very high number, considering the various departmental agendas and even potential multifaceted campaign rhetoric from the president.

The research results can also be considered as a partial rebuke to the “wokeness’ as a feature of political rhetoric” approach, as even if one can be conceivably sympathetic to “wokeness” being the desired intellectual framework or even a moral “baseline”, discounting “wokeness” as a *sui generis* (left-wing (Bobbio 2003), or even conceivably multiculturalist (Taylor, Appiah, Habermas, Rockefeller, Walzer, Wolf 1994)) ideological approach would disallow the mapping of the phenomenon in time and by actors. So even if “wokeness” is normatively positive, it has to have its non-exogenous attributes, so questions such as why it is specifically the Biden administration that introduced “equity-based” policy, while the Obama administration used that framing to a significantly lower degree can be validly asked and researched.

The logistic regression model results also fall broadly in line with what has been advocated for in theory, as there is a significant year-on-year coefficient increase. While acknowledging the limits of the model due to the imperfect dataset, in this regard, it can be cautiously taken as supporting the view of “wokeness” as a framework only appearing in general discourse fairly recently and increasingly so. The model also confirms the assumption of it being associated much more with the Democratic party than the Republican party, as the `is_democrat` variable has modest log odds of 0.127 and

is_republican having -1.509 log odds, the only negative ones (and sub-1 in odds ratio conversion) highlighting the possible polarisation on the issue (Ura, Ellis 2012), and further underlying the lack of political consensus, thus putting a further dent in conceptualising “’wokeness’ as a feature of political rhetoric” only unless one is to consider the ideological framework of the Republican party of the entire last 20 years to be “abnormal”.

Conclusion

The thesis successfully explored the issue of the presence of “wokeness” across sections of the U.S. federal government by employing the “text as data” approach, working with the innovative methodology of machine text classification, chiefly based on the content analysis methodology.

Theoretically, the thesis argues for the adoption of the concept of “wokeness’ as an ideology”, as opposed to “wokeness’ [solely] as a feature of political rhetoric”, a conceptual approach that has been hitherto dominant in social sciences, that sees “wokeness” as primarily a feature of (right-wing) political rhetoric.

It argues for the incorporation of ideas present outside academia that, however, draw heavily on academic frameworks and theories in attempts at describing reality and fundamentally, both conceptual “branches” share their theoretical pedigree in the conceptualisation of “wokeness as an individual ethic” but differ in their conclusions due to fundamental disagreement on whether “wokeness” is to be considered a normative good and the desired “normal” (Cammaerts 2020; 2022), or an approach at the very least lacking in societal consensus (Sullivan 2020b; 2020c; Yang 2021; Goodwin 2023a). It is asserted that neglecting “wokeness as an ideology” can, among other things, lead to a faulty analysis of political processes, ideological shifts and institutional behaviour, and thus testing empirical observations contained in the aforementioned conceptualisation ought to be considered valuable.

In regards to methodology, it proposes three major adjustments to utilising the supervised learning technique of fine-tuning a language model to classify a text corpus. Firstly it is argued when highly imbalanced classes of classification (in the case of the thesis “woke” and “not woke”, with “not woke” expected to have over 90% share) are expected to be encountered, but only one classification is clearly identifiable (“woke”) and the other is “everything else”, a combination of dictionary based collection of the learning sample, along with random sample is to be used, as to achieve a reasonable proportion of classes for the training data in fine-tuning the model.

Secondly, it introduces a measure to counteract the possible biases that a single-coder-coded training dataset would bring about by having the training sample co-classified by a zero-shot text classification language model. The results from both the zero-shot model and the human coder are then processed by an “AND gate” algorithm, only considering units “woke” if both coders agree (Krippendorff 2004, p. 212). This approach should aid external validity, as the concept introduced is novel and a more conservative approach to identifying “wokeness” might be warranted.

Thirdly, working with a large dataset of over 180000 units, a model evaluation on a random sample is proposed. The sample pseudo-randomly chosen then undergoes the same procedure as the training sample, zero-shot and human classified results are processed by an “AND gate” and the “simulated gold standard” is then evaluated against the fine-tuned model’s results. (Grimmer, Roberts, Stewart 2022, pp. 338–344)

Model such trained, in the case of this thesis, appears to perform well, empirically validly, achieving over 0.9 in F1-score (macro) in the evaluation. Furthermore, this methodological adjustment ensures “the human in the loop” (Grimmer, Roberts, Stewart 2022, p. 67) at multiple stages of the automated classification, thus contributing to facial validity of the output (Krippendorff 2004, p. 319).

Empirically, the research finds that there has been a significant rise in the utilisation of the “woke” framework in the external communication of the U.S. federal government since the inauguration of President Biden. There is a significant focus on “equity” across identity groups, diversity, equity and inclusion policies, and general highlighting of the significance of group identity, with the “woke” classification reaching up to 7% of all press releases in the year 2021. That is fully in line with Wesley Yang’s (2021) hypothesis of “wokeness” gaining institutional ground with the advent of the Biden administration. Surprisingly, contrary to some predictions, the levels of “wokeness” in the Obama administration are significantly lower compared to the Biden years, albeit significantly higher than both the preceding Bush years and the following Trump years (barring the deviant 2020, where there is noticeable elevation to early Obama levels, most likely connected to the death of George Floyd (Goodwin 2023b)).

The logistic regression model points to similar conclusions, “wokeness” being more associated with institutions controlled by the Democratic party than the Republican party and being fairly significantly more associated with the head of government (the president in case of the U.S. federal government, the governor in the case of a state government), along with very significant presence in the universities surveyed and the chance of press release being “woke” increasing year on year.

Based on these findings it is argued the concept of “wokeness’ as an ideology” and its operationalisation can be useful in description of ideological shift at a minimum in the U.S. political landscape. Furthermore it can be operationalised in the context of party politics study, political systems study and most certainly political communication studies, among others.

This case study thesis thus hopes to contribute in two major ways to political science research. To showcase the usability of machine text classification and the “text as data” approach as a whole in research and to make a convincing case for the introduction of “wokeness’ as an ideology”, a framework emphasising the importance of group identity in relation to “systems of oppression”, inter-group equality of outcome and extensive harm reduction for “marginalised” groups.

Specifically, further “wokeness” research should focus on better delineating its possible triangular relationship with multiculturalism and liberalism, as while “wokeness” could be simply seen as a radical extension of identitarian multiculturalism, its critical theory heritage and empirical observations point to a more complex reality. Cross-country comparative analysis could also be of interest, as to further establish or disprove the sometimes purported unique status of it in American political life.

Summary

The thesis explores the topic of “wokeness” chiefly within the U.S. federal government using a “text as data” approach leveraging machine text classification based on content analysis methodological framework.

It successfully introduces and utilises two methodological adjustments for text classification with a language model, particularly when dealing with imbalanced datasets. Firstly, the combined use of dictionary-based and random sampling for training data for fine-tuning the primary classification language model and secondly, the use of a zero-shot classification model along a human coder to simulate a “gold standard” after applying an “AND gate” approach on such classified results that are then used to fine-tune the primary model and analogously in evaluation, when a random sample is taken from the wholly classified dataset and new “gold standard” is devised.

The conceptualisation of “wokeness” is mapped in current science and public discourse, where its use can be grouped into roughly three categories, “wokeness’ as a feature of political rhetoric”, “wokeness’ as an individual ethic”, and “wokeness as an ideology”. It argues for the utility of “wokeness as an ideology”, albeit it being largely absent from academic literature, as a concept denoting ideology emphasising the importance of group identity in relation to “systems of oppression”, inter-group equality of outcome and harm reduction for “marginalised” groups.

Empirical findings indicate a significant rise in the use of “wokeness” rhetoric in U.S. federal government in the President Biden’s administration, which contrasts with the lower levels of “wokeness” mapped during the Obama administration and the Bush and Trump years. Based on statistical analysis, the thesis suggests that “wokeness” is more associated with the Democratic Party and is increasingly prevalent in communications from both government and educational institutions.

Further research and conceptualisation of “wokeness” appears to be warranted based on the findings of the thesis, inclusive of decoding its complicated relationship with multiculturalism and liberalism.

List of References

ABRAJANO, Marisa and HAJNAL, Zoltan L., 2015. *White Backlash : Immigration, Race, and American Politics*. Course Book. Princeton, NJ: Princeton University Press. ISBN 0-691-17619-1.

ALFONSECA, Kiara, 2024. What does “woke” mean and why are some conservatives using it? *ABC News*. Online. 24 January 2024. Available from: <https://abcnews.go.com/Politics/woke-conservatives/story?id=93051138> [Accessed 30 July 2024].

ATKINS, J. Spencer, 2023a. Defining Wokeness. *Social epistemology*. 2023. Vol. 37, no. 3, p. 321–338. DOI 10.1080/02691728.2022.2145857.

ATKINS, J. Spencer, 2023b. Moral Encroachment, Wokeness, and the Epistemology of Holding. *Episteme*. 2023. Vol. 20, no. 1, p. 86–100. DOI 10.1017/epi.2020.50.

BAUDRILLARD, Jean and GLASER, Sheila Faria, 1994. *Simulacra and simulation*. Michigan: University of Michigan Press. The body, in theory : histories of cultural materialism. ISBN 978-0-472-06521-9.

BELL, Lee Anne, 2023. Theoretical Foundations for Social Justice Education. In: . 4. Routledge. p. 3–25. ISBN 0-367-43120-3.

BENNETT, Bryan, 2023. Americans are Concerned By Elected Officials Using “Wokeness” as a Distraction. *Navigator*. Online. 4 April 2023. Available from: <https://navigatorresearch.org/americans-are-concerned-by-elected-officials-using-woke-as-a-distraction/> [Accessed 14 February 2024].

BLM, 2024. Our History. *Black Lives Matter*. Online. 19 March 2024. Available from: <https://blacklivesmatter.com/our-history/> [Accessed 30 July 2024].

BOBBIO, Norberto, 2003. *Pravice a levice : důvod a smysl rozdělení politické scény*. 1. vydání. Brno: Centrum pro studium demokracie a kultury. Politologická řada ; sv. 16. ISBN 80-7325-012-8.

BOYCE, Benjamin, 2021. Racist compared to what? the myth of white wokeness. *Whiteness and Education*. 30 June 2021. Vol. 6, p. 1–15. DOI 10.1080/23793406.2021.1925579.

CAMMAERTS, Bart, 2020. The neo-fascist discourse and its normalisation through mediation. *Journal of*

multicultural discourses. 2020. Vol. 15, no. 3, p. 241–256. DOI 10.1080/17447143.2020.1743296.

CAMMAERTS, Bart, 2022. The abnormalisation of social justice: The ‘anti-woke culture war’ discourse in the UK. *Discourse & society*. 2022. Vol. 33, no. 6, p. 730–743. DOI 10.1177/09579265221095407.

CAMMAERTS, Bart, 2024. Defending Democracy Against Populist Neo-Fascist Attacks: The Role and Problems of Public Sphere Theory. *Javnost (Ljubljana, Slovenia)*. 2024. Vol. 31, no. 1, p. 26–45. DOI 10.1080/13183222.2024.2310983.

CAPEHART, Jonathan, 2022. Hands Up, Don’t Shoot! Built on a Lie | Facing History & Ourselves. . Online. 28 April 2022. Available from: <https://www.facinghistory.org/resource-library/hands-dont-shoot-built-lie> [Accessed 30 July 2024].

CHEN, Andrew, 2014. The death of RSS in a single graph at andrewchen. . Online. 2014. Available from: <https://andrewchen.com/the-death-of-rss-in-a-single-graph/> [Accessed 20 July 2024].

CORNELL LAW SCHOOL, 2021. Griggs v. Duke Power Co. (1971). *LII / Legal Information Institute*. Online. December 2021. Available from: [https://www.law.cornell.edu/wex/griggs_v_duke_power_co_\(1971\)](https://www.law.cornell.edu/wex/griggs_v_duke_power_co_(1971)) [Accessed 20 July 2024].

CRENSHAW, Kimberle, 1991. Mapping the Margins: Intersectionality, Identity Politics, and Violence against Women of Color. *Stanford law review*. 1991. Vol. 43, no. 6, p. 1241–1299. DOI 10.2307/1229039.

CROUCH, Jeffrey, ROZELL, Mark J. and SOLLENBERGER, Mitchel A., 2017. The Law: The Unitary Executive Theory and President Donald J. Trump. *Presidential studies quarterly*. 2017. Vol. 47, no. 3, p. 561–573. DOI 10.1111/psq.12401.

DAALDER, Ivo H. and LINDSAY, James M., 2001. Nasty, Brutish and Long: America’s War on Terrorism. *Brookings*. Online. December 2001. Available from: <https://www.brookings.edu/articles/nasty-brutish-and-long-americas-war-on-terrorism/> [Accessed 20 July 2024].

DIANGELO, Robin, 2016. White Fragility. *Counterpoints (New York, N.Y.)*. 2016. Vol. 497, p. 245–253.

DIANGELO, Robin J., 2019. *White Fragility: Why It’s So Hard for White People to Talk About Racism*. London: Allen Lane. ISBN 978-0-14-199056-9.

DOJ, 2018. Archives | Justice News Archive. . Online. 9 October 2018. Available from: <https://www.justice.gov/archives/justice-news-archive> [Accessed 20 July 2024].

DOJ, 2024. Press Releases | United States Department of Justice. . Online. 2024. Available from: <https://www.justice.gov/news/press-releases> [Accessed 20 July 2024].

DOL, 2024. News Releases. *DOL*. Online. 2024. Available from: <https://www.dol.gov/newsroom/releases> [Accessed 20 July 2024].

DUN, Lindsay, SOROKA, Stuart and WLEZIEN, Christopher, 2021. Dictionaries, Supervised Learning, and Media Coverage of Public Policy. *Political communication*. 2021. Vol. 38, no. 1–2, p. 140–158. DOI 10.1080/10584609.2020.1763529.

ECKHARD, Steffen, JANKAUSKAS, Vytautas, LEUSCHNER, Elena, BURTON, Ian, KERL, Tilman and SEVASTJANOVA, Rita, 2023. The performance of international organizations: a new measure and dataset based on computational text analysis of evaluation reports. *Review of International Organizations*. 2023. Vol. 18, no. 4, p. 753–776. DOI 10.1007/s11558-023-09489-1.

ELLIS, Christopher and STIMSON, James A., 2012. *Ideology in America*. New York: Cambridge University Press. ISBN 978-1-107-01903-4.

EPA, 2013. Our Mission and What We Do. . Online. 29 January 2013. Available from: <https://www.epa.gov/aboutepa/our-mission-and-what-we-do> [Accessed 20 July 2024].

EPA, 2024. Search News Releases | US EPA. . Online. 15 July 2024. Available from: <https://www.epa.gov/newsreleases/search> [Accessed 20 July 2024].

FARKAS, Igor, 2021. *Artificial Neural Networks and Machine Learning – ICANN 2021 30th International Conference on Artificial Neural Networks, Bratislava, Slovakia, September 14–17, 2021, Proceedings, Part V*. 1st ed. 2021. Cham: Springer International Publishing. Theoretical Computer Science and General Issues, 12895. ISBN 3-030-86383-2.

FREIRE, Paulo, SHOR, Ira, MACEDO, Donaldo P. (Donaldo Pereira) and RAMOS, Myra Bergman, 2018. *Pedagogy of the oppressed*. Fiftieth anniversary edition. New York, New York: Bloomsbury Academic.

ISBN 1-5013-1415-7.

FTC, 2013. Commissioners. *Federal Trade Commission*. Online. 7 June 2013. Available from: <https://www.ftc.gov/about-ftc/commissioners-staff/commissioners> [Accessed 20 July 2024].

FTC, 2024. Press Releases. *Federal Trade Commission*. Online. 23 May 2024. Available from: <https://www.ftc.gov/news-events/news/press-releases> [Accessed 20 July 2024].

GALLUP, 2021. U.S. Approval of Interracial Marriage at New High of 94%. *Gallup.com*. Online. 10 September 2021. Available from: <https://news.gallup.com/poll/354638/approval-interracial-marriage-new-high.aspx> [Accessed 30 July 2024].

GARCÍA BEDOLLA, Lisa and HAYNIE, Kerry L., 2013. The Obama coalition and the future of American politics. *Politics, groups & identities*. 2013. Vol. 1, no. 1, p. 128–133. DOI 10.1080/21565503.2012.758593.

GOODWIN, Matthew, 2023a. *Values, voice and virtue*. Harlow, England: Penguin Books. ISBN 978-0-14-199909-8.

GOODWIN, Matthew, 2023b. The Human Cost of BLM. . Online. 2023. Available from: <https://www.mattgoodwin.org/p/the-human-cost-of-blm> [Accessed 20 July 2024].

GOOGLE TRENDS, 2024. Google Trends. *Google Trends*. Online. 7 January 2024. Available from: <https://trends.google.com/trends/explore?date=all&q=wokeness&hl=en-GB> [Accessed 20 July 2024].

GRAMSCI, Antonio, 1967. *Základy politiky*. Vydání 1. Praha: Mladá fronta.

GRIMMER, Justin, ROBERTS, Margaret E. and STEWART, Brandon M., 2022. *Text as data : a new framework for machine learning and the social sciences*. Princeton ; Princeton University press. ISBN 978-0-691-20755-1.

HANANIA, Richard, 2023. *The origins of woke*. New York, NY: Broadside Books. ISBN 978-0-06-323721-6.

HARVARD, 2024. Campus & Community Archives. *Harvard Gazette*. Online. July 2024. Available from: <https://news.harvard.edu/gazette/section/campus-community/> [Accessed 20 July 2024].

HARVARD UNIVERSITY, 2019. Annual report on Title IX shows goals met, but work will continue. *Harvard Gazette*. Online. 9 December 2019. Available from: <https://news.harvard.edu/gazette/story/2019/12/annual-report-on-title-ix-shows-goals-met-but-work-will-continue/> [Accessed 30 July 2024].

HAYES, Andrew F. and KRIPPENDORFF, Klaus, 2007. Answering the Call for a Standard Reliability Measure for Coding Data. *Communication methods and measures*. 2007. Vol. 1, no. 1, p. 77–89. DOI 10.1080/19312450709336664.

HE, Pengcheng, GAO, Jianfeng and CHEN, Weizhu, 2023. *DeBERTaV3: Improving DeBERTa using ELECTRA-Style Pre-Training with Gradient-Disentangled Embedding Sharing*. Online. 24 March 2023. arXiv. arXiv:2111.09543. [Accessed 20 July 2024].

HE, Pengcheng, LIU, Xiaodong, GAO, Jianfeng and CHEN, Weizhu, 2021. microsoft/deberta-v3-large · Hugging Face. . Online. 2021. Available from: <https://huggingface.co/microsoft/deberta-v3-large> [Accessed 22 January 2024].

HUGGING FACE, 2023. What is Zero-Shot Classification? - Hugging Face. . Online. 5 September 2023. Available from: <https://huggingface.co/tasks/zero-shot-classification> [Accessed 30 July 2024].

HUNTINGTON, Samuel P., 1981. *American politics: the promise of disharmony*. Cambridge: Belknap Press. ISBN 0-674-03021-4.

IBM, 2021a. What Is Logistic Regression? | IBM. . Online. 16 August 2021. Available from: <https://www.ibm.com/topics/logistic-regression> [Accessed 23 July 2024].

IBM, 2021b. What Is Machine Learning (ML)? | IBM. . Online. 22 September 2021. Available from: <https://www.ibm.com/topics/machine-learning> [Accessed 23 July 2024].

INGLEHART, Ronald, 1977. *The silent revolution: changing values and political styles among Western publics*. Princeton, New Jersey: Princeton University Press. Princeton Legacy Library. ISBN 0-691-64151-X.

JAMES, Stanlie M., 2009. Barack Obama: Coalitions of a Purple Mandate. *Souls (Boulder, Colo.)*. 2009. Vol. 11, no. 1, p. 51–59. DOI 10.1080/10999940902734895.

JOHNS HOPKINS SCHOOL, 2020. COVID-19 Testing: Understanding the “Percent Positive” | Johns Hopkins | Bloomberg School of Public Health. . Online. 10 August 2020. Available from: <https://publichealth.jhu.edu/2020/covid-19-testing-understanding-the-percent-positive> [Accessed 20 July 2024].

JOHNSON, Kayla M., 2023. Destination: Wokeness? Possibilities for exploring domestic educational travel as conscientizing praxis. *Education, citizenship and social justice*. 2023. DOI 10.1177/17461979231197410.

JURÍK, Stanislav, 2023. Slovo woke se zabydluje v našem slovníku. *Plzeň*. Online. 26 June 2023. Available from: <https://plzen.rozhlas.cz/slovo-woke-se-zabydluje-v-nasem-slovniku-9022098> [Accessed 30 July 2024].

KAUFMANN, Eric, 2022. The new culture wars: Why critical race theory matters more than cancel culture. *Social science quarterly*. 2022. Vol. 103, no. 4, p. 773–788. DOI 10.1111/ssqu.13156.

KAUFMANN, Eric and GOODWIN, Matthew J., 2018. The diversity Wave:A meta-analysis of the native-born white response to ethnic diversity. *Social science research*. 2018. Vol. 76, p. 120–131. DOI 10.1016/j.ssresearch.2018.07.008.

KENDI, Ibram X., 2017. *Stamped from the beginning : the definitive history of racist ideas in America*. First trade paperback edition. New York: Bold Type Books. Definitive history of racist ideas in America. ISBN 978-1-56858-598-7.

KESSLER, Glenn, 2019. Analysis | Harris, Warren ignore DOJ report to claim Michael Brown was ‘murdered.’ *Washington Post*. Online. 13 August 2019. Available from: <https://www.washingtonpost.com/politics/2019/08/13/harris-warren-ignore-doj-report-claim-that-michael-brown-was-murdered/> [Accessed 30 July 2024].

KRIPPENDORFF, K, 2004. Measuring the reliability of qualitative text analysis data. *Quality & quantity*. 2004. Vol. 38, no. 6, p. 787–800. DOI 10.1007/s11135-004-8107-7.

KYMLICKA, Will, 1995. *Multicultural citizenship : a liberal theory of minority rights*. Oxford: Clarendon Press. Oxford political theory. ISBN 0-19-827949-3.

- KYMLICKA, Will, 2018. Liberal Multiculturalism as a Political Theory of State–Minority Relations. *Political Theory*. 2018. Vol. 46, no. 1, p. 81–91. DOI 10.1177/0090591717696021.
- LANCING, Logan and LINDSAY, James, 2024. *The Queering of the American Child: How a New School Religious Cult Poisons the Minds and Bodies of Normal Kids*. New Discourses.
- LAURER, Moritz, 2024. MoritzLaurer/deberta-v3-large-zeroshot-v2.0 · Hugging Face. . Online. 3 April 2024. Available from: <https://huggingface.co/MoritzLaurer/deberta-v3-large-zeroshot-v2.0> [Accessed 20 July 2024].
- LINDSAY, James, 2021. Social Justice. *New Discourses*. Online. 25 November 2021. Available from: <https://newdiscourses.com/tftw-social-justice/> [Accessed 20 July 2024].
- LINDSAY, James, 2022a. *Race Marxism: The Truth About Critical Race Theory and Praxis*. Independently published. ISBN 9798795809083.
- LINDSAY, James, 2022b. Queer. *New Discourses*. Online. 11 September 2022. Available from: <https://newdiscourses.com/tftw-queer/> [Accessed 20 July 2024].
- LUKIANOFF, Greg, 2021. The Second Great Age of Political Correctness. *Reason.com*. Online. 13 December 2021. Available from: <https://reason.com/2021/12/13/the-second-great-age-of-political-correctness/> [Accessed 23 July 2024].
- MAHENDRU, Sakshi and PANDIT, Tejul, 2024. SecureNet: A Comparative Study of DeBERTa and Large Language Models for Phishing Detection. *arXiv.org*. 2024. DOI 10.48550/arxiv.2406.06663.
- MANHATTAN INSTITUTE, 2023. Institutionalized: Wokeness with Wesley Yang. *Manhattan Institute*. Online. 2023. Available from: <https://manhattan.institute/multimedia/institutionalized-wokeness-with-wesley-yang/> [Accessed 5 July 2024].
- MARSHALL, Patricia L. and WILSON, Jonee, 2023. Toward A Discourse on the Threat of Performative Wokeness to Justice Agendas in Education. *Urban education (Beverly Hills, Calif.)*. 2023. P. 4208592311534. DOI 10.1177/00420859231153406.
- MCWHORTER, John H., 2022. *Woke racism: how a new religion has betrayed Black America*. Great Britain: Forum. ISBN 978-1-80075-142-2.

MIKOLOV, Tomas, CHEN, Kai, CORRADO, Greg and DEAN, Jeffrey, 2013. Efficient Estimation of Word Representations in Vector Space. *arXiv.org*. 2013. DOI 10.48550/arxiv.1301.3781.

MIT, 2017. Explained: Neural networks. *MIT News | Massachusetts Institute of Technology*. Online. 14 April 2017. Available from: <https://news.mit.edu/2017/explained-neural-networks-deep-learning-0414> [Accessed 23 July 2024].

MONTANARO, Domenico, 2023. What does the word “woke” really mean, and where does it come from? *NPR*. Online. 19 July 2023. Available from: <https://www.npr.org/2023/07/19/1188543449/what-does-the-word-woke-really-mean-and-where-does-it-come-from> [Accessed 30 July 2024].

MUDDE, Cas, 2021. ‘Critical race theory’ is the right’s new bogeyman. The left must not fall for it. *The Guardian*. Online. 25 June 2021. Available from: <https://www.theguardian.com/commentisfree/2021/jun/25/critical-race-theory-rightwing-bogeyman-left-wing> [Accessed 23 July 2024].

MUDDE, Cas, 2023. What is behind Ron DeSantis’s Stop-Woke Act? *The Guardian*. Online. 6 February 2023. Available from: <https://www.theguardian.com/commentisfree/2023/feb/06/what-is-behind-ron-desantis-stop-woke-act> [Accessed 23 July 2024].

NAACP, 2023. Reclaiming the Word “Woke” as Part of African American Culture | NAACP. . Online. 31 December 2023. Available from: <https://naacp.org/resources/reclaiming-word-woke-part-african-american-culture> [Accessed 30 July 2024].

NASREEN, Ghazala, MURAD KHAN, Muhammad, YOUNUS, Muhammad, ZAFAR, Bushra and KASHIF HANIF, Muhammad, 2024. Email spam detection by deep learning models using novel feature selection technique and BERT. *Egyptian informatics journal*. 2024. Vol. 26, p. 100473. DOI 10.1016/j.eij.2024.100473.

NATIONAL ARCHIVES, 2021. Brown v. Board of Education (1954). *National Archives*. Online. 29 September 2021. Available from: <https://www.archives.gov/milestone-documents/brown-v-board-of-education> [Accessed 30 July 2024].

NG, Kate, 2021. What is the history of the word 'woke'? *The Independent*. Online. 22 January 2021. Available from: <https://www.independent.co.uk/news/uk/home-news/woke-meaning-word-history-b1790787.html> [Accessed 30 July 2024].

NY AG, 2024. Press Releases | New York State Attorney General. *New York Attorney General*. Online. 30 July 2024. Available from: <https://ag.ny.gov/press-releases> [Accessed 20 July 2024].

NY DOT, 2024. Press Releases. *New York Department of Transportation*. Online. 2024. Available from: <https://www.dot.ny.gov/news/press-releases/2010> [Accessed 20 July 2024].

NY GOVERNOR, 2024. News. *New York Governor*. Online. July 2024. Available from: <https://www.governor.ny.gov/news> [Accessed 20 July 2024].

ORACLE, 2020. What is Machine Learning? . Online. 2020. Available from: <https://www.oracle.com/cz/artificial-intelligence/machine-learning/what-is-machine-learning/> [Accessed 23 July 2024].

OSC, 2024. Federal Employee Hatch Act Information. *Federal Employee Hatch Act Information*. Online. 2024. Available from: <https://osc.gov/Services/Pages/HatchAct-Federal.aspx#tabGroup12|tabGroup31> [Accessed 20 July 2024].

PILKINGTON, Andrew, 2022. Political correctness in a global age: The ethical implications of a hegemonic discourse. In: . 1. Routledge. p. 15–37. ISBN 1-03-206214-2.

PINCOURT, Charles and LINDSAY, James, 2021. *Counter Wokecraft: A Field Manual for Combatting the Woke in the University and Beyond*. Independently published. ISBN 9798536815038.

PLUCKROSE, Helen and LINDSAY, James A., 2020. *Cynical theories : how activist scholarship made everything about race, gender, and identity - and why this harms everybody*. First edition. Durham: Pitchstone Publishing. ISBN 978-1-63431-202-8.

PRINCETON, 2024. All News. *Princeton University*. Online. July 2024. Available from: <https://www.princeton.edu/news> [Accessed 20 July 2024].

PUBLIC POLICY INSTITUTE OF CALIFORNIA, 2021. Fewer Californians Say Racism Is a Big Problem Today Compared to 2020. *Public Policy Institute of California*. Online. 2021. Available from: <https://www.ppic.org/blog/fewer-californians-say-racism-is-a-big-problem-today-compared-to-2020/> [Accessed 30 July 2024].

RAWLS, John, 1999. *A theory of justice*. Rev. ed. Cambridge, Mass: Belknap Press of Harvard University Press. ISBN 978-0-674-00077-3.

RAWLS, John, 2005. *Political liberalism*. Expanded ed. New York: Columbia Univ. Press. Columbia classics in philosophy. ISBN 978-0-231-13089-9.

ROTHMAN, Denis and GULLI, Antonio, 2022. *Transformers for natural language processing : build, train, and fine-tune deep neural network architectures for NLP with Python, PyTorch, TensorFlow, BERT, and GPT-3*. Second edition. Birmingham: Packt. Expert insight. ISBN 978-1-80324-733-5.

ROY, Laura A., 2018. *Teaching while white : addressing the intersections of race and immigration in the classroom*. Lanham, Maryland: Rowman & Littlefield. ISBN 1-4758-4039-X.

ROZELL, Mark J., BARILLEAUX, Ryan J. and KELLEY, Christopher S., 2010. *The Unitary Executive and the Modern Presidency*. College Station: Texas A&M University Press. Presidency and leadership (Unnumbered). ISBN 1-60344-190-5.

RUFO, Christopher C., 2020. Against Wokeness. *City Journal*. Online. 16 September 2020. Available from: <https://www.city-journal.org/article/against-wokeness> [Accessed 23 July 2024].

SARTORI, Giovanni and KŘÍŽOVÁ, Karolina, 2005. *Pluralismus, multikulturalismus a přistěhovanci : esej o multietnické společnosti*. 1. vyd. v českém jazyce. Praha: Dokořán. Bod. ISBN 80-7363-022-2.

SCHMITT, Carl, 2005. *Political theology : four chapters on the concept of sovereignty*. Chicago ; University of Chicago Press. ISBN 0-226-73889-2.

SCHMITT, Carl, 2021. *Carl Schmitt's early legal-theoretical writings : Statute and judgment and the Value of the state and the significance of the individual*. Cambridge: Cambridge University Press. Cambridge studies in constitutional law. ISBN 1-108-78616-2.

SCHOLTEN, Miroslava, 2014. *The political accountability of EU and US independent regulatory agencies*. 1st ed. Leiden, Netherlands: Brill Nijhoff. Nijhoff Studies in EU Law, Volume 6. ISBN 90-04-26299-7.

SKOWRONEK, Stephen, DEARBORN, John A. and KING, Desmond, 2021. *Phantoms of a beleaguered republic : the deep state and the unitary executive*. Oxford, UK: Oxford University Press. Oxford scholarship online. ISBN 0-19-754310-3.

SMITH, David S., BOAG, Lee, KEEGAN, Connor and BUTLER-WARKE, Alice, 2023. Land of Woke and Glory? The Conceptualisation and Framing of “Wokeness” in UK Media and Public Discourses. *Javnost - The Public*. 2 October 2023. Vol. 30, no. 4, p. 513–533. DOI 10.1080/13183222.2023.2273656.

SOBANDE, Francesca, KANAI, Akane and ZENG, Natasha, 2022. The hypervisibility and discourses of ‘wokeness’ in digital culture. *Media, culture & society*. 2022. Vol. 44, no. 8, p. 1576–1587. DOI 10.1177/01634437221117490.

SULLIVAN, Andrew, 2019. Andrew Sullivan: When the Ideologues Come for the Kids. *Intelligencer*. Online. 20 September 2019. Available from: <https://nymag.com/intelligencer/2019/09/andrew-sullivan-when-the-ideologues-come-for-the-kids.html> [Accessed 23 July 2024].

SULLIVAN, Andrew, 2020a. The Roots Of Wokeness. *The Weekly Dish*. Online. 31 July 2020. Available from: <https://andrewsullivan.substack.com/p/the-roots-of-wokeness> [Accessed 23 July 2024].

SULLIVAN, Andrew, 2020b. Andrew Sullivan: Is There Still Room for Debate? *Intelligencer*. Online. 12 June 2020. Available from: <https://nymag.com/intelligencer/2020/06/andrew-sullivan-is-there-still-room-for-debate.html> [Accessed 5 July 2024].

SULLIVAN, Andrew, 2020c. Why Is Wokeness Winning? *The Weekly Dish*. Online. 16 October 2020. Available from: <https://andrewsullivan.substack.com/p/why-is-wokeness-winning> [Accessed 23 July 2024].

TANTI, Marc, VAN DER PLAS, Lonneke, BORG, Claudia and GATT, Albert, 2021. On the Language-specificity of Multilingual BERT and the Impact of Fine-tuning. In: . Online. Punta Cana, Dominican Republic: Association for Computational Linguistics. 2021. p. 214–227. DOI 10.18653/v1/2021.blackboxnlp-1.15. [Accessed 20 July 2024].

TAYLOR, Charles, APPIAH, Kwame Anthony, HABERMAS, Jürgen, ROCKEFELLER, Stephen C., WALZER, Michael and WOLF, Susan, 1994. *Multiculturalism: Examining the Politics of Recognition*. Expanded Paperback edition. Princeton, NJ: Princeton University Press. ISBN 978-0-691-03779-0.

TEXAS AG, 2024. Press Releases | Office of the Attorney General. . Online. July 2024. Available from: https://www.texasattorneygeneral.gov/news/press_releases [Accessed 31 July 2024].

TEXAS DOT, 2024. Newsroom. *Texas Department of Transportation*. Online. 2024. Available from: <http://txdot.gov/en/home/about/newsroom.html> [Accessed 20 July 2024].

TEXAS GOVERNOR, 2019. News Archive. *Texas Governor*. Online. 2019. Available from: <https://gov.texas.gov/news/archive> [Accessed 20 July 2024].

TEXAS GOVERNOR, 2024. News. *Texas Governor*. Online. 2024. Available from: <https://gov.texas.gov/news/category/press-release> [Accessed 20 July 2024].

THE WHITE HOUSE, 2009. News Archive. . Online. 2009. Available from: <https://georgewbush-whitehouse.archives.gov/news/> [Accessed 20 July 2024].

THE WHITE HOUSE, 2017. Statements and Releases. *whitehouse.gov*. Online. 2017. Available from: <https://obamawhitehouse.archives.gov/briefing-room/statements-and-releases> [Accessed 20 July 2024].

THE WHITE HOUSE, 2021a. News – The White House. . Online. 2021. Available from: <https://trumpwhitehouse.archives.gov/news/> [Accessed 20 July 2024].

THE WHITE HOUSE, 2021b. The Biden-Harris Administration. *The White House*. Online. 2021. Available from: <https://www.whitehouse.gov/administration/> [Accessed 20 July 2024].

THE WHITE HOUSE, 2023. FACT SHEET: The Biden-Harris Administration’s Work to Support Black Disabled Americans. *The White House*. Online. 24 February 2023. Available from: <https://www.whitehouse.gov/briefing-room/statements-releases/2023/02/23/fact-sheet-the-biden-harris-administrations-work-to-support-black-disabled-americans/> [Accessed 30 July 2024].

THE WHITE HOUSE, 2024. Briefing Room. *The White House*. Online. July 2024. Available from: <https://www.whitehouse.gov/briefing-room/> [Accessed 20 July 2024].

TREASURY, 2024. Press Releases. *U.S. Department of the Treasury*. Online. 19 July 2024. Available from: <https://home.treasury.gov/news/press-releases> [Accessed 20 July 2024].

URA, Joseph Daniel and ELLIS, Christopher R., 2012. Partisan Moods: Polarization and the Dynamics of Mass Party Preferences. *The Journal of politics*. 2012. Vol. 74, no. 1, p. 277–291. DOI 10.1017/S0022381611001587.

VIJAY SRINIVAS TIDA and HSU, Sonya, 2022. Universal Spam Detection using Transfer Learning of BERT Model. *arXiv (Cornell University)*. 2022. DOI 10.48550/arxiv.2202.03480.

WALSH, Kenneth T., 2015. The Undoing of George W. Bush. *US News & World Report*. Online. 28 August 2015. Available from: [//www.usnews.com/news/the-report/articles/2015/08/28/hurricane-katrina-was-the-beginning-of-the-end-for-george-w-bush](http://www.usnews.com/news/the-report/articles/2015/08/28/hurricane-katrina-was-the-beginning-of-the-end-for-george-w-bush) [Accessed 20 July 2024].

WASTELL, Laurie, 2023. “The culture war is no sideshow; it is a fight for Western civilisation”: An Interview with Eric Kaufmann. . Online. 7 October 2023. Available from: <https://europeanconservative.com/articles/interviews/interview-with-eric-kaufmann/> [Accessed 14 February 2024].

WHITEOUT, Sam, 2018. Popularizing Wokeness. *Harvard journal of African American public policy*. 2018. P. 63–70.

WOLFF, Robert Paul, MOORE, Barrington and MARCUSE, Herbert, 1970. *A Critique of pure tolerance*. 5th printing. Boston: Beacon Press. Beacon Paperback, 328. ISBN 978-0-8070-1559-9.

YANG, Wesley, 2021. Welcome to Year Zero. *Year Zero*. Online. 14 July 2021. Available from: <https://wesleyyang.substack.com/p/welcome-to-year-zero> [Accessed 14 February 2024].

ZAVATTARO, Staci M. and BEARFIELD, Domonic, 2022. Weaponization of Wokeness: The Theater of Management and Implications for Public Administration. *Public administration review*. 2022. Vol. 82, no. 3, p. 585–593. DOI 10.1111/puar.13484.

Index of Tables

Table 1: Dictionary capture terms, based on the outlined theory (own work).....	31
Table 2: The presented dataset structure (own work).....	33
Table 3: Number of press releases scrapped by origination institution (own work).....	38
Table 4: Temporal distribution of the press releases by year (own work).....	40
Table 5: Inter-coder reliability metrics of the training sample (own work).....	42
Table 6: Label distribution of the training sample (own work).....	43
Table 7: Performance metrics of the fine-tuned model on unseen test split in training (own work).....	44
Table 8: Classification distribution between “coders” (own work).....	44
Table 9: Inter-coder reliability metrics of the evaluation sample (between zero-shot model coder and human coder) (own work).....	45
Table 10: Performance metrics of the fine-tuned model on the evaluation sample (own work).....	45
Table 11: Yearly distribution of “woke” press releases across all institutions (own work).....	47
Table 12: Yearly distribution of “woke” press releases across non-independent federal institutions (own work).....	49
Table 13: Simple logistic regression model (own work).....	54
Table 14: Logistic regression model with fixed effect year variable (own work).....	55

Table of Figures

Figure 1: Google search trend for “woke” worldwide (own work, data from Google (2024)).....	12
Figure 2: Conceptual map of “wokeness” (own work).....	20
Figure 3: Methodological flowchart of the thesis (own work).....	25
Figure 4: Boxplots of the lengths of the text units after scrapping (own work).....	39
Figure 5: Temporal distribution of the press releases by year (own work).....	41
Figure 6: Confusion matrix of the evaluation sample, true label is the “gold standard” and predicted label is the label predicted by the fine-tuned model. (own work).....	46
Figure 7: Bar chart, with local regression line, of the fraction of “woke” classifications per year across all institutions (own work).....	48
Figure 8: Bar chart, with local regression line, of the fraction of “woke” classifications per year across all non-independent federal institutions (own work).....	50
Figure 9: Bar chart, with local regression line, of the fraction of “woke” classifications per year across universities (own work).....	52
Figure 10: Bar chart, with local regression line, of the fraction of “woke” classifications per year across all independent federal institutions (own work).....	53

List of Appendices

Appendix no. 1: R script used for visualisation and statistical analysis – *1-Stepan-Kosan-Thesis-R-Script.R* (text file)

Appendix no. 2: Google search term data for “woke”- *2-Stepan-Kosan-Google-Search-Terms.csv* (csv table)

Appendix no. 3: Scraped press releases full initial dataset - *3-Stepan-Kosan-Input-Data.csv* (csv table)

Appendix no. 4: Classified data sample for fine-tuning - *4-Stepan-Kosan-Training_Sample.csv* (csv table)

Appendix no. 5: Evaluation sample - *5-Stepan-Kosan-Evaluation-Sample.csv* (csv table)

Appendix no. 6: Full classified and labeled dataset- *6-Stepan-Kosan-Full-Predicted.csv* (csv table)

Appendix no. 7: Jupyter notebook source code file- *7-Stepan-Kosan-Jupyterntb-Code.ipynb* (text file)