Tracing Antisemitic Language Through Diachronic Embedding Projections: France 1789-1914

Rocco Tripodi **Massimo Warglien** **Simon Levis Sullam** **Deborah Paci**
Ca' Foscari University of Venice
{rocco.tripodi, warglien, levissmn, deborah.paci}@unive.it

Abstract

We investigate some aspects of the history of antisemitism in France, one of the cradles of modern antisemitism, using diachronic word embeddings. We constructed a large corpus of French books and periodicals issues that contain a keyword related to Jews and performed a diachronic word embedding over the 1789-1914 period. We studied the changes over time in the semantic spaces of 4 target words and performed embedding projections over 6 streams of antisemitic discourse. This allowed us to track the evolution of antisemitic bias in the religious, economic, socio-politic, racial, ethic and conspiratorial domains. Projections show a trend of growing antisemitism, especially in the years starting in the mid-80s and culminating in the Dreyfus affair. Our analysis also allows us to highlight the peculiar adverse bias towards Judaism in the broader context of other religions.

1 Introduction

Word embeddings are widely used in many Natural Language Processing (NLP) tasks. They provide a machine-interpretable representation of lexical features. Their effectiveness in representing words semantics consists essentially in the ability of learning association patterns in the training dataset. For this reason the learned representations contain human-like biases (Caliskan et al., 2017). These biases can be detected easily and can be related to gender, ethic or racial aspects (Garg et al., 2018; Voigt et al., 2017).

Since the use of word embedding is ubiquitous in many commercial products such as search engines and machine translators, the research community has introduced different techniques to debias them (Bolukbasi et al., 2016; Zhao et al., 2018), especially under the gender dimension. Despite these efforts debiasing word embeddings seems to be harder than expected. In fact, while Bolukbasi et al. (2016) and Zhao et al. (2018) demonstrated that it is possible to debias specific gendered-words, even after the debiasing procedure, the geometry of the embeddings remains almost the same with respect to non gendered-words (Gonen and Goldberg, 2019), preserving their original bias.

In this work, we turn these biases to the historian's advantage and shed light on some aspects of the history of antisemitism in France during the so called *long XIX century*, between the French Revolution and the First Word War, using diachronic word embedding. This technique allows to capture diachronic conceptual changes and to analyse stereotyped historical biases. We tracked how historical events and publications influenced the construction of the collective imaginary related to the Jewish question.

We assume that words do not have a fixed meanings. They can be used in different contexts to evoke a great variety of meanings using different connotational nuances. These multiple meanings are acquired (or lost) over time in correspondence to specific socio-political events. For example, one of the meanings of the word *usurier* (i. e.; *money lender*), as reported by the *French Historical Dictionary*, refers to: *the financial activities of the Jews* [who since the Middle Ages were], *the only ones authorised to lend on pawns* (Rey et al., 2010). This association derives from the fact that especially between the XVI and the XIX century, this word acquired a negative connotation, nurtured by anti-Jewish prejudice and stereotyping developing from the idea of an illegitimate interest attached to this activity. This image, as the above mentioned definition explains, was also fixed in the collective imaginary by Shylock, the Jewish protagonist in Shakespeare's *Merchant of Venice* (1598).

In this work, we trace the conceptual changes of words related to the Jewish question. We collected a large corpus for this purpose, composed of thousands of books and newspapers published in France between 1789 and 1914. We used diachronic word embedding to represent the data, measures of local changes in the semantic space of different words, and embedding projections to quantify biases in different semantic spheres. The measurement of local changes is particularly suited for our study because we do not want to identify new meanings in the words related to the Jewish question, instead we want to trace how the context of their use changed and how these changes affected the representation of Jews at the time of the rise of modern antisemitism. Measuring biases over time is particularly interesting because it allows to connect them with antisemitic streams as identified by historians in the field (Wilson, 1982) and operationalised by us.

2 Related Work

Models for capturing diachronic conceptual changes are associated with the distributional hypothesis (Harris, 1954; Firth, 1957; Weaver, 1955): the semantics of a word is defined by the context in which it is used. Following this assumption, different models have been presented, based on co-occurrence vectors (Sagi et al., 2009; Gulordava and Baroni, 2011; Basile et al., 2016) or word embeddings (Kim et al., 2014; Kulkarni et al., 2015; Hamilton et al., 2016a).

These works are brought together by the idea of analysing the contexts in which a word occurs and have culminated in the measures of *semantic shift* and *cultural drift*, proposed by (Hamilton et al., 2016a) and *the law prototipicality* proposed by Dubossarsky et al. (2015). Semantic shifts are regular linguistic processes such as semantic widening (e.g., *dog*, that in Middle English was used to refer to dogs of a particular breed) (Bloomfield, 1933). This measure was used to derive two laws of semantic change: the *law of conformity*: semantic change scales with a negative power of word frequency; and the *law of innovation*: polysemous words have significantly higher rates of semantic change (Hamilton et al., 2016b). Cultural drifts involve local changes to a lexical form's use (e.g.: the changes in the meaning of the word *cell*: *prison cell* → *cell phone*) (Hamilton et al., 2016a). The *law of prototipicality* was introduced by Dubossarsky et al. (2015): it states that prototypical words, words that are near to the centroid of a cluster in a semantic space, change slower than words that are in a peripheral position. The laws of *conformity, innovation and prototipicality* have been questioned by Dubossarsky et al. (2017), who used controlled conditions to test them.

Different works that tried to measure, directly or indirectly, cultural drifts have been proposed recently. Garg et al. (2018) analysed gender and ethnic stereotypes in the United States during the 20th and 21st centuries, using word embeddings trained on the Google Books and Corpus of Historical American English (COHA) corpora. Kozlowski et al. (2018) used diachronic word embeddings to conduct macro-cultural investigation of social stereotypes. Kutuzov et al. (2017) attempted to model the dynamics of wordwide armed conflicts using word embeddings trained on a news corpus. Zhao et al. (2017) analyzed the amplification effect that learning models present on the gender dimension when trained on biased data.

3 Motivations and historical background

We have looked at linguistics representation of Jews in 19th century France, which was one of the cradles of modern antisemitism in Europe, i.e. of the mostly secularized and racial transformation of the centuries-old Christian prejudice against the Jews (Katz, 1980).

Since the entry and gradual integration of the Jews in French society after the Revolution of 1789, the appearance of anti-Jewish texts, the rise of public controversies, or the burst of cases and scandals in which Jews were supposedly involved marked the emergence and spread of the Jewish question on the French scene, in what have been called *antisemitic moments* or *episodes* (Birnbaum, 2011). Especially during the Third Republic, beginning in 1870, references to Jews entered the French public discourse in relation to a supposed growing influence of the Jews on political and economic affairs, the rise of anticlericalism in the face of Catholic France (for which Jews were considered responsible), the accusation of an alliance between Jews and Freemasonry.

This process reached its climax with the Dreyfus affair (1894), the unfounded accusation against a French army officer to have sold intelligence information to the German enemy (Dreyfus would be exonerated in 1906): the affair caused the heavy

spread of antisemitic accusations and anti-Jewish movements of opinion (Wilson, 1982). Different streams of antisemitism ran accross French society throughout this time, together with a pro-Jewish reaction driven by the supporters of Dreyfus' innocence (Kalman, 2010).

The publication in 1845 (republished in 1886) of Alfred Toussenel's *Les Juifs rois de l'epoque* caused especially the rise of the so-called economic antisemitism, which accused the Jews of an increasing economic and financial influence, of which the Rothschilds were considered the protagonists and became a symbol. This accusation was later confirmed by the supposed Jewish role in the crash of the Catholic bank Union Générale (1882) and in the Panama corruption scandal (1892), together with the revival of nationalism tied to the Boulangist crisis (Sternhell, 1998). These events generated a resurgence of antisemitism. In response to the growing secularization and anticlericalism, French Catholics revived an ancient tradition of religious antisemitism, marked in this time by the appearance of works such as Gougenot des Mousseaux's *Le Juif, le judaïsme et la judaïsation des peuples chrétiens* (1869) and by the anti-Jewish campaigns of Catholic periodicals such as *L'Univers* and *La Croix*.

In 1886 the journalist Edouard Drumont published the hugely successful *La France juive. Essai d'histoire contemporaine*, which described a French society under a greedy Jewish influence and control, painting in the style of a novelist (inspired by Balzac and by contemporary feuilletons or serialized novels) the contours of Jewish conspiracies. Although the subtitle of the work suggests an essay of contemporary history, on reading it is as if one is before an enormous cauldron of common place assumptions on Jews which includes Catholic, social, racial, economic, and conspiratorial anti-Semitism. The success of his work depended on the waves it made in the intellectual milieu of the era and its impact on the popular masses attracted by the synthesis of anti-Semitism of the right, of a church worried about laicisation, and anti-Semitism of the left, anti-capitalist and laical. This and other books by Drumont mixed Catholic, socio-political, ethic and conspirationist antisemitism, accusing Jews of all sorts of religious offenses, political machinations, moral perversions and secret plots (Kauffmann, 2008).

The combination of these streams of anti-Jewish accusations, prejudices and stereotypes would christallize - or reach its climax - in the Dreyfus affair. We suggest that the usage in print (books and periodicals) of the term *juif* or other terms related to the Jewish question, all characterised by an adverse bias, was especially connected to antisemitic tendencies. However, we should note that this vocabulary was also present at the time in Biblical and theological scholarship, art and art-historical publications, fictional and theatrical literature, medical treatises and the rising social sciences. References to Jews in the public discourse were therefore not necessarily mobilised in a political context with explicit antisemitics aims. Our investigation asks whether using diachronic word embeddings trained on a large corpus confirms the chronological development of antisemitic language which historians have described on a qualitative level (and if it sheds light on different, previously ignored, *antisemitic moments*). We also examine the relevance of the semantic areas or streams in relation to the *Jew* which we have identified based on (Wilson, 1982), and we show the trends through time of unfavourable biases towards Jews in the period considered.

4 The Corpus and the Embeddings

4.1 The corpus

The corpus[1] was constructed downloading from Gallica, the online library of the *Bibliothèque Nationale de France*[2], the raw text of all the resources that contain a keyword related to Jews (see appendix A for the complete list of keywords) and have been published between 1789-1914. The research was further restricted to those resources that have an OCR quality higher that 98%. The resulting corpus contains 54.403 books and 245.188 periodicals issues. It is important to notice here that we downloaded the full text of a book or newspaper issue even if a keyword appeared only once in it.

Figures 1a and 1b indicate the distribution of resources per year in the periodicals and books subcorpora, respectively, together with the total number of resources in Gallica. The resources distribution per year is not homogenous in neither subcorpora: publications increase significantly year

[1]The metadata of the corpus, the embeddings and the code used for the experiments can be downloaded from https://github.com/roccotrip/antisem.

[2]https://gallica.bnf.fr/

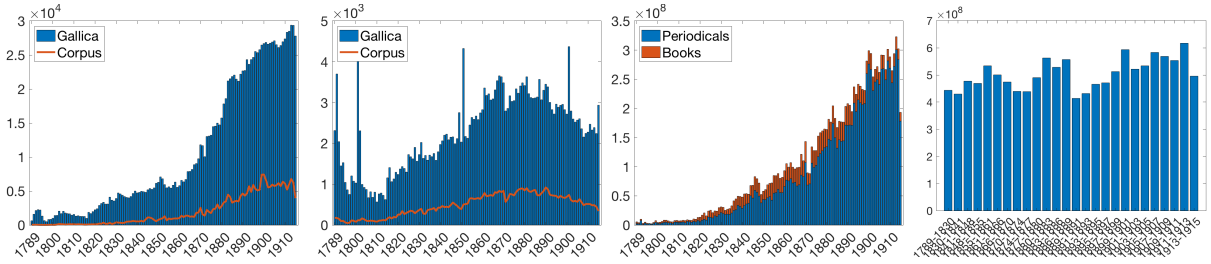

(a) Periodicals distribution (b) Books distribution (c) Tokens distribution (d) Num. of tokens in each bin

Figure 1: Distribuition of resources in the corpus and time bins division.

by year. Several hypotheses can explain this proliferation of documents over time. One straightforward hypothesis can be related to increasing importance of Jews in the French public debate with the proliferation of anti-Semitic movements and newspapers such as *La Croix*, *La libre parole*, *La Lutte antijuive* and *L'Intransigeant*, just to name a few. Yet, a second hypothesis can be related to the fact that the print industry grew over time. In fact, many newspapers and publishers were founded after 1825. For example, *Hachette*, the publisher with the largest number of books in our corpus (1558), was founded in 1826. The newspapers *Le Figaro* was founded in 1826, *L'Univers* in 1833 and *Le Temps* in 1861. Figure 1a and 1b, plotting our corpus compared to the whole Gallica one, seems to suggest that the second hypothesis is the most plausible. In fact, the quantities of resources in our corpora follow a trend similar to those observed in the whole Gallica.

4.2 The embeddings

Figure 1c shows the distribution of tokens per year distinguishing periodicals and books. The greater part of the data is from the periodicals, giving to the corpus a focus on the contemporaneity. Given this distribution it is impossible to train a model using equally sized time bins. For this reason, we decided to group the data into approximately equal bins in terms of tokens. The resulting division comprehend 26 time bins of $\approx$ 450 millions tokens each (see Figure 1d).

For each bin we trained a word2vec skip-gram model (Mikolov et al., 2013) using a window size of 5 words on both sides, a word vector of 300 dimensions and removing the words that occur less than 25 times.

5 Analysis

In this section we analyse the resulting embeddings. First we study the changes in the semantic space of 4 target words. Then we analyse the biases of the same words for 6 different dimensions, each of which corresponds to a predetermined stream.

5.1 Local changes

The first analysis that we conducted is the measurement of the changes in the semantic space of the words used to refer to Jews: *juif* (noun/adjective, masculine, singular), *juifs* (noun/adjective, masculine, plural), *juive* (noun/adjective, feminine, singular) and *juives* (noun/adjective, feminine, plural). For this measurement we used the local neighborhood measure proposed by Hamilton et al. (2016a). To compute this measure it is necessary to create a second order vector, s, according to equation 1,

$$s_i^t = \text{cos-sim}(\mathbf{w}_i^{(t)}, \mathbf{w}_j^{(t)}) \forall w_j \in N_k(w_i^{(t)}) \cup N_k(w_i^{(t+1)}), \quad (1)$$

where $N_k(w_i^{(t)})$ represents the k-nearest neighbours ($k - nn$) at time (t) (according to cosine similarity) of a target word w_i and $\mathbf{w}_*$ is the embedding corresponding to word w_*. Once these vectors are constructed we compute the cosine distance, d, among them to quantify their differences, with equation 2,

$$d(\mathbf{s}_i^{t1}, \mathbf{s}_i^{t2}) = 1 - \text{cos-sim}(\mathbf{s}_i^{t1}, \mathbf{s}_i^{t2}). \quad (2)$$

The results of this experiment are presented in Figure 2 for all the morphological variants of the word *juif*, using $k = 100$[3]. What emerges clearly

[3]We noticed that the general trend of the curves in Figure 2 does not change much using different values of k (10, 25,

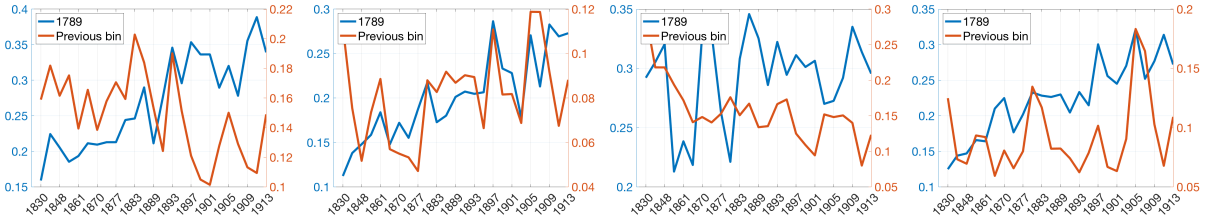

(a) juif (b) juifs (c) juive (d) juives

Figure 2: Local neighborhood measure. The y axes indicates the cosine distance of the second-order vector constructed for each time period compared to the 1789 (blu line) and the preceding time period (red line).

juif		**juifs**		**juive**		**juives**	
⤸ 1841	⤹	⤸ 1861	⤹	⤸ 1874	⤹	⤸ 1870	⤹
laquedem	juive	crucifient	juif	huguenots	judaïque	syriennes	négociantes
mécréant	judaïque	schismatiques	israëlites	favorite	musulmane	iraniennes	samaritaines
rogatons	rabin	judaïsants	juive	opera	syrienne	musulmanes	réfugiées
blasphémateur	bouddhiste	fétichistes	rabbins	rigoletto	héroine	israëlites	ascètes
⤸ 1886	⤹	⤸ 1870	⤹	⤸ 1886	⤹	⤸ 1880	⤹
ghetto	judaïque	judaïsants	juif	drumont	iranienne	israélites	épousées
déicides	rabin	hérétiques	synagogues	antisémitisme	apostasié	musulmanes	luthériennes
francmaçon	wanderghen	cabalistes	talmud	circoncis	lithuanienne	femmes	turques
aryen	anabaptiste	luciférien	sanhédrin	théâtrale	puritaine	célébrations	dissolues
⤸ 1893	⤹	⤸ 1897	⤹	⤸ 1893	⤹	⤸ 1897	⤹
déicide	talmud	antisémites	samaritains	juiverie	synagogue	juif	dissolues
youtre	bouddhiste	youtres	talmud	satanisme	héroine	youtres	baptisées
francmaçon	sodomite	youpins	idolâtres	monogamique	lapidée	antijuives	prostituaient
youpins	anabaptiste	enjuivés	pharisiens	opprimée	persécutrice	antisémitiques	ascètes
⤸ 1897	⤹	⤸ 1905	⤹	⤸ 1901	⤹	⤸ 1905	⤹
youtre	rabin	judaïsants	synagogues	stigmatisant	dragonnade	massacrées	courtisannes
sémite	usurier	hellénisants	talmud	antijuive	torturée	terrorisées	païennes
judaïsant	shylock	diaspora	pharisiens	antinationale	puritaine	diaspora	prostituaient
antisémite	anabaptiste	massacrant	ismaélites	dreyfusiste	anabaptiste	déportées	émigrées

Table 1: Words that have been introduced (left column ⤸) or eliminated (right column ⤹) for our 4 target words in time periods with a high local neighborhood distance, compared to 1789.

from these figures is that there are certain periods of time in which the relation among a target word and its local neighbourhood changed consistently. What we noticed from them is that besides changes in the relative similarity among two words what changes more is the k-nn itself, with the introduction or elimination of specific words.

Some of the words that were introduced (or eliminated) to (from) the k-nn of relevant time periods (according to local neighbourhood measure) are presented in Table 1. The words in this table are ordered according to the cosine similarity with the target word. We can see an elimination of words related to the religious domain for all the target words that we used, terms like *rabbin* (i.e., rabbi), *talmud* (i.e. the study of the Jewish law), *synagogue* and *sanhédrin* (i.e., the Jewish council) are replaced by more negatively connotated words such as *ghetto*, *déicides*, *antisémites* and *antijuives*. From the few words presented in Table 1 one can also notice a possible rise of antisemitic prejudice (or at least of antisemitic language), with the introduction of specific words in the vocabulary specifically tailored to connote Jewish people in a derogatory way. *Youtre* and *youpin* are slang racist insults negatively connoting the *Jew*. They appear increasingly during the period 1880-1900.

Other terms with a negative connotation that entered the semantic area of our target words are *judaïsants* (i.e., judaizers), *enjuivés* (i.e., strongly influenced by the Jewish spirit) and *francmaçon* (i.e., freemason). These terms, as we will see in the next section, are related to the idea of a Jewish conspiracy against the world. This is a clear example of the growth of the antisemitic vocabulary.

50, 100) and that fixing $k = 100$ gives a good representation of the variations over time. Increasing this value gives high fluctuations and introduces many irrelevant words.

The analysis of the word *juive* is especially interesting. The word *drumont* entered its space in the time period 1886-1889. It refers to Éduard Drumont, a well known antisemite who published one of the bestsellers of the antisemitism (*La France juive*), in 1886, and was the editor of an antisemitic newspaper (*La libre Parole*), founded in 1892. We can also notice that in the semantic space of the word *juive* there are different words related to theatre. This probably derives from literary and theatre representations of Jewish female characters, as well as references to supposed Jewish inappropriate moral and sexual behaviours. Among the theatre representations we may recall that of *La Juive*, first shown in 1835, one of the most popular French operas of the 19th century, which tells the story of an impossible love affair between a Christian man and a Jewish woman. The fictional Jew, invariably seen as an outsider, provides a mirror for the phobias and obsessions of French society at a time when old Jew hatred becomes politicised, when anti-Semitism begins to permeate French ideology (Weinberg, 1983; Hallman, 2007; Samuels, 2009).

We can also see the introduction of the word *aryen* (i.e.: aryan) in 1886. This word entered the semantic space in a syntagmatic relation with the word *juif* and, as we will see in the next section, the period in which it entered is characterised by a strong antisemitism characterised by an intensification of racial and socio-political stereotypes.

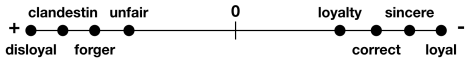

Figure 3: Semantic axis and projections.

5.2 Embedding projections

5.2.1 The streams

To quantify biases in word embeddings semantic spaces it is common to project a specific word vector on a semantic axis (Bolukbasi et al., 2016; Caliskan et al., 2017). The semantic axis can be computed as $\mathbf{g} = \mathbf{w}_i - \mathbf{w}_j$ and its projection as the dot product $\hat{b} = \mathbf{w} \cdot \mathbf{g}$, assuming that the vectors are normalised, the projection is equal to the cosine similarity. The higher the values of the projection, the more biased the word is toward that direction.

In previous literature (Bolukbasi et al., 2016) the gender direction (e.g., $\vec{he} - \vec{she}$) was used to project words related to occupations in order to quantify if these words embed information about gender. In this work we do not want to project words only according to a single direction but we want to analyse different adverse and (or) favourable biases, comparing them over time. For this reason, we defined six different semantic axes, that correspond to six antisemitic streams (S) (Wilson, 1982).

For each stream, $s \in S$, we identified a set of n antonyms pairs, $z_s = \{(a_1^{neg}, a_1^{pos}), ..., (a_n^{neg}, a_n^{pos})\}$ to construct the bias subspace in the embedding. To avoid selection biases we selected the antonyms pairs starting from a positive seed word, that is highly representative for the stream, and used a knowledge base to collect its synonyms and the corresponding antonyms (see appendix B for the complete list of antonyms used). We noticed that computing the PCA of each subspaces the corresponding explained variance is concentrated on the first component and that it is stable over time. For example, the first component of the racial stream has an explained variance of 0.34 (mean) with a standard deviation of 0.012.

The six different semantic areas, which may correspond to related antisemitic discourses are:

1. religious: antisemitism based on theological doctrines or narratives, and on religious prejudices and accusations. The seed word is *believer* (*unbeliever*);

2. economic: antisemitism based on a supposed Jewish role in the economy or on stereotypes concerning Jews' economic behaviours. The seed word is *generosity* (*greed*);

3. socio-political: antisemitism based on malevolent, e.g. anti-national, political behaviours or on supposedly threatening Jewish actions. The seed word is *honor* (*shame*);

4. racial: antisemitism based on the definition of Jews as a race, considered inferior. The seed word is *pure* (*impure*);

5. conspiratorial: antisemitism based on conspiracy theories. The seed word is *loyal* (*disloyal*);

6. ethic: antisemitism based on Jewish supposed unethical or perverse morals or behaviours. The seed word is *moral* (*immoral*);

To quantify the biases for all the time we computed the mean bias, b, for each stream as the

Figure 4: Projections of our 4 target words to the 6 semantic axes. Positive values indicates the adverse bias.

arithmetic mean of the individual biases, $\hat{b}$ on each axis, according to equation 3:

$$b(w_i, s) = \frac{1}{n}\sum_{j=1}^{n} \mathbf{w}_i \cdot (\mathbf{w}_{a_j^{neg}} - \mathbf{w}_{a_j^{pos}}), \quad (3)$$

where n is the number of antonyms pairs in stream s, Given the ordering of the antonyms in the computation of the bias axis ($\mathbf{g} = \mathbf{w}_{a_j^{neg}} - \mathbf{w}_{a_j^{pos}}$) we define an adverse bias when b is positive and a favourable bias when b is negative.

An example of semantic axis constructed with the pair *disloyal* as negative word and *loyal* as positive, is presented in Figure 3 ($\vec{disloyal} - \vec{disloyal}$). From this figure we can see that words that have a high projection value are words very similar to the negative word, on the other hand, words with a low projection are very similar to the word on the other side. The projection tells us if a word is closer to one extreme or the other. Unbiased words should have a projection close to 0.

5.2.2 Biases related to *Jews*

The results of this experiment are presented in Figure 4. The adverse bias is always high for the words *juif*, *juifs* and *juive*. For the word *juives* only on few cases it is negative. Adverse and favourable biases are measured with positive and negative measures respectively.

Our analysis confirms the chronological development of *antisemitic moments* identified by historians, with a steady increase of adverse bias starting in the 1880s, before the Dreyfus affair. We also notice an unexpected peak in adverse bias between 1855 and 1866, in connection with the French Second Empire (1851-1870). The semantic areas or streams in relation to the Jew identified on the basis of (Wilson, 1982) seem relevant for the description of adverse bias in *antisemitic moments*. The highest adverse bias characterises the religious semantic area, followed by the economic and ethic areas. The religious adverse bias shows a peak starting in 1855, after the establishment of Napoleon III's Second Empire, a time of renewed allegiance to the Catholic Church and in 1895 at the beginning of the Dreyfus affair. Also the economic adverse bias shows a peak starting in 1855, perhaps because of the increase of economic discourse on Jews following the publication of Toussenel's *Les Juifs rois de l'époque*, and again coinciding with the establishment of the Second Empire. Another peak comes with the Dreyfus affair. The ethic adverse bias peaks in the period 1830-1855, diminishes afterward and peaks again toward the end of the Dreyfus affair.

Racial, conspiratorial and especially sociopolitical semantic areas show a steady adverse bias and an increase mostly after 1886, i.e. after the publication of Drumont's La France juive (1886). The conspiratorial adverse bias also peaks – like the religious, the economic and the ethic adverse bias – in 1855.

The singular *juif* prevails in the conspiratorial

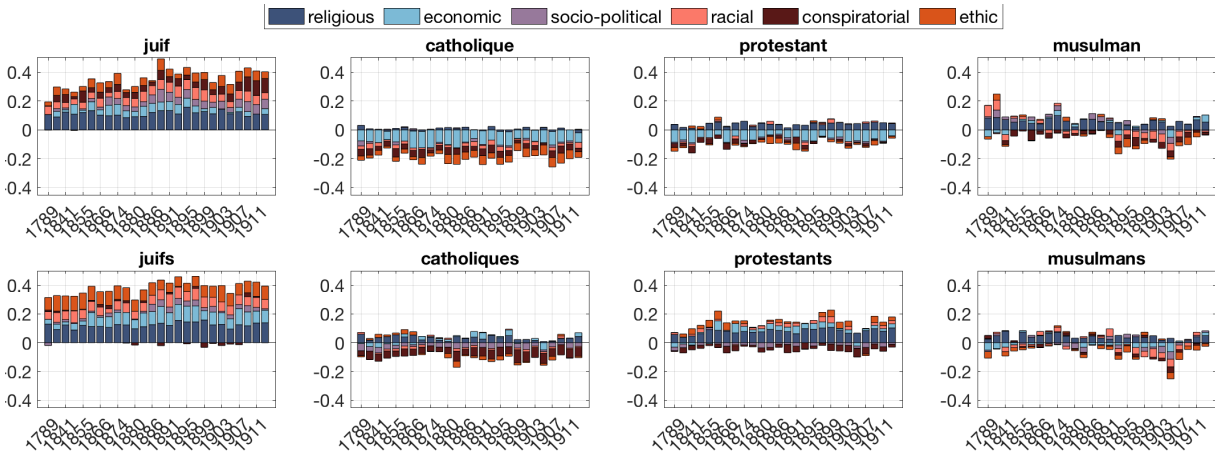

Figure 5: Cumulative bias projections compared to different religious groups.

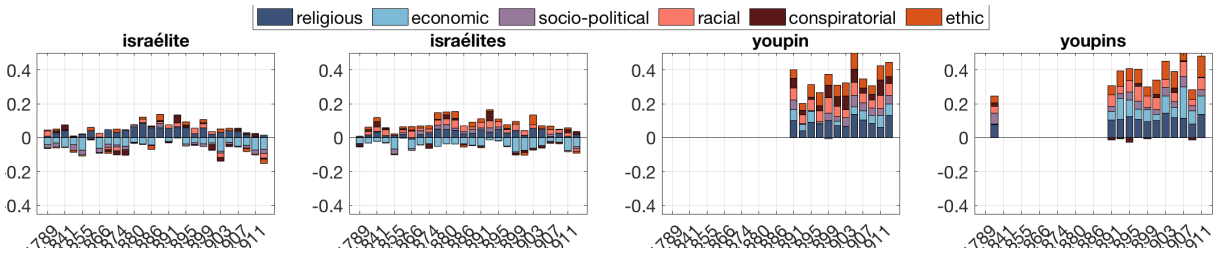

Figure 6: Cumulative bias projections for other words used to refer to Jews.

and socio-political semantic areas, which seem to entail general statements about *the Jew*. This tendency has been noticed by historians as typical of modern antisemitism and has been called *singularisation* (Miccoli, 2003). This underlines that there are features *common to all [Jews], because in all and every one there emerges something which constitutes a common and exclusive feature* of the Jew as *the enemy to be defeated* (Miccoli, 2003). On the other hand, the plural *juifs* prevails in the economic and ethic areas, as implying collective behaviours of Jews.

Racial, sociopolitical and conspiratorial semantic areas show a steady adverse bias and increase especially after 1886. As the racist vision of the Jew increases, it is turned increasingly into a political vision, and it is also nourished by a conspirationist worldview, which will culminate in the Dreyfus affair.

5.2.3 Comparative biases concerning different religious groups

The results of this experiment are presented in Figure 5. They show a comparison with three different religious groups: Catholic, (*catholique*), Protestant (*protestante*) and Muslim (*musulman*). The plots sum positive and negative biases to give a general picture of the biases at each time step.

Juif and *catholic* have a completely opposite bias: exclusively adverse in the first case, entirely favorable in the latter. Confronting *juif* and *protestant* we notice a similar bias, adverse in the first case, favorable in the latter. But the favorable bias of Protestant is much more reduced than that of *catholic*. Confronting *juifs* and *protestantes*, both show an adverse bias (lower in the case of Protestants). The adverse bias concerns *protestantes* especially in relation to the religious domain. *Musulman* and *musulmans* also show an adverse bias concentrated in the religious sphere. If we look at racial stream, this grows for *juif(s)* reference to *protestants* is absent; while there is an occasional emergence of *musulman*, with an adverse bias between 1789 and 1840, when questions of citizenship are being defined (France conquers Egypt in 1798 and in 1834 Algeria is annexed to France; in 1870 the Crémieux Decree granted French citizenship to Algerian Jews but not to Muslims), and a favourable bias in 1891-95 (in 1890 a bill is proposed for the granting of French citizenship to Algerian muslims, see Weill, 2005). The last increase is probably also connected with the availability of a larger quantity of digitised North-African press in the corpus.

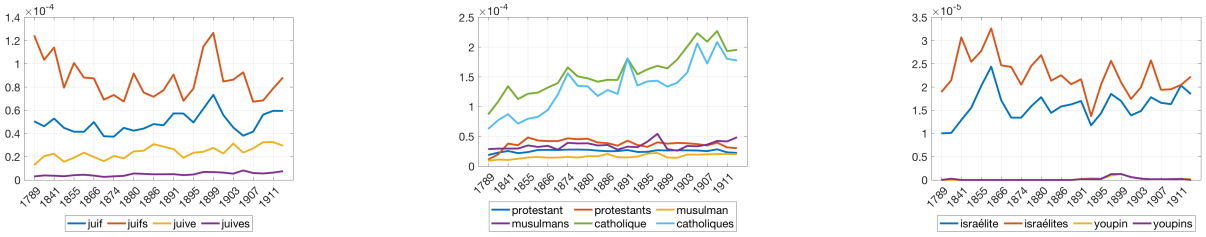

(a) (b) (c)

Figure 7: Target words frequency.

5.2.4 Comparative bias concerning different words used to refer to Jews

The results of this experiment are presented in Figure 6. *Israelite* and *israelites* do not show a particular bias as the terms are often used euphemistically (including by Jews themselves), i.e. preferred to the more direct and connotated *juif* and *juifs*. These terms refers to the cultural assimilation and social integration of Jews into French society, as described by Honoré (1981).

The slang and derogatory *youpin* spread starting around 1886; its shows an exclusively adverse bias and a trend similar to *juif*, as if the terms *juif* and *youpin* were interchangeable.

5.3 Target words frequency

Even if the corpus has been constructed selecting documents containing words related to the Jewish question, we noticed that the frequencies of words related to other religious groups is higher for *catholique* and *catholiques* and slightly lower for the words *protestant*, *protestantes*, *musulman* and *musulmanes*. The frequencies of all the target wordsare reported in Figure 7a, 7b and 7c.

6 Conclusions

References to Jews increase throughout the 19th century, as Jews were integrated within French society and these references appear to be mostly associated with an adverse bias in all semantic areas. The adverse bias grows starting in the mid-1880s, i.e. in the second half of the Third Republic, when the rise of anticlericalism and socialism was associated with Jews by the conservative and catholic public opinion. Around this time the publication of Drumont's *La France juive* provokes an adverse bias towards Jews clearly associated to antisemitic discourse in all semantic areas, which prepares the outburst of the Dreyfus affair, and it remains steady during and after the affair.

The highest adverse bias characterises the religious semantic area, followed by the economic and ethic spheres. The conspiratorial and sociopolitical areas show an adverse bias more often associated with the singular *juif*, as if they provoked categorical statements. Adverse bias in the economic and ethic areas is expressed through the plural *juifs* as describing collective behaviours.

The confrontation between *juif* and *catholic* shows an entirely adverse bias in the first case and an entirely favorable bias in the latter case. The adverse bias towards other minorities, i.e. Protestants and Muslims concerns the religious semantic area. No bias concerning protestant emerges in the racial semantic area, while a negative and positive bias emerge in relation to Muslims at times when the question of French citizenship is being defined.

As one evaluates the presence of the word *juif*, and the semantic areas surrounding it, one should also consider that these may emerge in texts which are not antisemitic per se, but still contribute to the spread of images of Jews, with specific biases. We refer here especially to literary texts.

We suggest that the adverse bias in various semantic areas may be associated with antisemitic discourses, but this association should be further explored though an examination of the historical context (for example that of *antisemitic moments*) or an analysis of the textual sources which spread the words associated with *the Jew*.

Acknowledgments

The authors of this work have received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 732942. The experiments have been run on the SCSCF cluster of Ca' Foscari University.

References

Pierpaolo Basile, Annalina Caputo, Roberta Luisi, and Giovanni Semeraro. 2016. Diachronic analysis of the italian language exploiting google ngram. *CLiC it*, page 56.

Pierre Birnbaum. 2011. *The anti-semitic moment: A tour of France in 1898*. Chicago University Press.

Leonard Bloomfield. 1933. *Language*. George Allen and Unwin LTD.

Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in neural information processing systems*, pages 4349–4357.

Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.

Haim Dubossarsky, Yulia Tsvetkov, Chris Dyer, and Eitan Grossman. 2015. A bottom up approach to category mapping and meaning change. In *NetWordS*, pages 66–70.

Haim Dubossarsky, Daphna Weinshall, and Eitan Grossman. 2017. Outta control: Laws of semantic change and inherent biases in word representation models. In *Proceedings of the 2017 conference on empirical methods in natural language processing*, pages 1136–1145.

John R Firth. 1957. A synopsis of linguistic theory, 1930-1955. *Studies in Linguistic Analysis*.

Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16):E3635–E3644.

Hila Gonen and Yoav Goldberg. 2019. Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them. *arXiv preprint arXiv:1903.03862*.

Kristina Gulordava and Marco Baroni. 2011. A distributional similarity approach to the detection of semantic change in the google books ngram corpus. In *Proceedings of the GEMS 2011 workshop on geometrical models of natural language semantics*, pages 67–71.

Diana R Hallman. 2007. *Opera, Liberalism, and Antisemitism in Nineteenth-Century France: The Politics of Halévy's La Juive*. Cambridge University Press.

William L Hamilton, Jure Leskovec, and Dan Jurafsky. 2016a. Cultural shift or linguistic drift? comparing two computational measures of semantic change. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing*, volume 2016, page 2116. NIH Public Access.

William L Hamilton, Jure Leskovec, and Dan Jurafsky. 2016b. Diachronic word embeddings reveal statistical laws of semantic change. *arXiv preprint arXiv:1605.09096*.

Zellig S Harris. 1954. Distributional structure. *Word*, 10(2-3):146–162.

Jean-Paul Honoré. 1981. Le vocabulaire de l'antisémitisme en france pendant l'affaire dreyfus. *Mots. Les langages du politique*, 2(1):73–92.

Julie Kalman. 2010. *Rethinking antisemitism in nineteenth-century France*. Cambridge University Press.

Jacob Katz. 1980. *From prejudice to destruction: anti-Semitism, 1700-1933*. Harvard University Press.

Grégoire Kauffmann. 2008. *Édouard Drumont*. Perrin.

Yoon Kim, Yi-I Chiu, Kentaro Hanaki, Darshan Hegde, and Slav Petrov. 2014. Temporal analysis of language through neural language models. *arXiv preprint arXiv:1405.3515*.

Austin C Kozlowski, Matt Taddy, and James A Evans. 2018. The geometry of culture: Analyzing meaning through word embeddings. *arXiv preprint arXiv:1803.09288*.

Vivek Kulkarni, Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2015. Statistically significant detection of linguistic change. In *Proceedings of the 24th International Conference on World Wide Web*, pages 625–635. International World Wide Web Conferences Steering Committee.

Andrey Kutuzov, Erik Velldal, and Lilja Øvrelid. 2017. Temporal dynamics of semantic relations in word embeddings: an application to predicting armed conflict participants. *arXiv preprint arXiv:1707.08660*.

Giovanni Miccoli. 2003. Antiebraismo, antisemitismo: un nesso fluttuante. In *Les racines chrétiennes de l'antisémitisme politique*, pages 1000–1021. École française de Rome.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.

Alain Rey, Chantal Tanet, and Marianne Tomi. 2010. Dictionnaire historique de la langue française.

Eyal Sagi, Stefan Kaufmann, and Brady Clark. 2009. Semantic density analysis: Comparing word meaning across time and phonetic space. In *Proceedings of the Workshop on Geometrical Models of Natural Language Semantics*, pages 104–111. Association for Computational Linguistics.

Maurice Samuels. 2009. *Inventing the Israelite: Jewish Fiction in Nineteenth-Century France*. Stanford University Press.

Zeev Sternhell. 1998. *La droite révolutionnaire: les origines francaises du fascisme 1885-1914*. Gallimard.

Rob Voigt, Nicholas P Camp, Vinodkumar Prabhakaran, William L Hamilton, Rebecca C Hetey, Camilla M Griffiths, David Jurgens, Dan Jurafsky, and Jennifer L Eberhardt. 2017. Language from police body camera footage shows racial disparities in officer respect. *Proceedings of the National Academy of Sciences*, 114(25):6521–6526.

Warren Weaver. 1955. Translation. *Machine Translation of Languages*, 14:15–23.

Henry H Weinberg. 1983. The image of the jew in late nineteenth-century french literature. *Jewish Social Studies*, 45(3/4):241–250.

Stephen Wilson. 1982. *Ideology and experience: antisemitism in France at the time of the Dreyfus affair*. Fairleigh Dickinson University Press; London: Associated University Presses.

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2017. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*.

Jieyu Zhao, Yichao Zhou, Zeyu Li, Wei Wang, and Kai-Wei Chang. 2018. Learning gender-neutral word embeddings. *arXiv preprint arXiv:1809.01496*.

A Keywords

- Juif (i.e: Jew - masculine, singular)
- Juive (i.e: Jew - feminine, singular)
- Judaisme (i.e: Judaism)
- Israëlite (i.e: Israelite)
- Israël (i.e: Israel)
- Israëlitisme (i.e: Israelitism)
- Mosaïsme (i.e: religions referred to the message of Moses)
- Talmud (i.e: Talmud)
- Judas (i.e: Judass)
- Moloch (i.e: the biblical name of a Canaanite god associated with child sacrifice)
- Ahasverus (i.e: a mythical immortal man whose legend began to spread in Europe in the 13th century. The original legend concerns a Jew who taunted Jesus on the way to the Crucifixion and was then cursed to walk the earth until the Second Coming.)

B Bias axes

The list of antonyms used to compute the bias axes. Note that the translation of the antonyms pairs is provided only for the singular. We used a public resource (http://www.synonyms-fr.com) to collect antonyms relations.

Religious angel, devil; sacred, profane; pious, atheist; pious, pagan; pious, idolater; pious, impious; sacred, cursed; venerable, abject; faithful, unfaithful; believer, unbeliever; religious, irreligious; dedicated, atheist.

Economic give, appropriate; generosity, greed; generous, greedy; generous, miserly; generous, stingy.

Socio-political prodigal, greedy; honest, rabble; honor, shame; friendly, hostile; loyal, deceitful; socialist, capitalist; friend, enemy; ally, antagonist; conservative, progressive.

Racial normal, strange; superiority, inferiority; equality, inequality; pleasant, unpleasant; benign, wicked; worthy, infamous; sympathy, hate; accepted, refused, better, worse; national, foreign; pure, impure; upper, lower; pure, filthy; clean, dirty.

Conspiratorial loyal; spy; honesty, treason; loyal, disloyal; clear, mysterious; obvious, occult; sincere, deceitful; sincere, unfair; benefactor, criminal; clear, secret; friendly, threatening; clear, dark.

Ethic chastity, lust; modest, intriguing; decent, indecent; virtuous, lascivious; faithful, unfaithful; moral, immoral; honest, dishonest; chaste, depraved; chaste, fleshly; pure, degenerate.