Scent and Sensibility: Perception Shifts in the Olfactory Domain

Teresa Paccosi,[1,2] **Stefano Menini,**[1] **Elisa Leonardelli,**[1]
Ilaria Barzon,[3] **Sara Tonelli,**[1]
[1]Fondazione Bruno Kessler, Trento, Italy
[2]Dept. of Cognitive Science, University of Trento, Italy
[3]Dept. of Humanities, University of Pavia, Italy
{tpaccosi, menini, eleonardelli, satonelli}@fbk.eu

Abstract

In this work, we investigate olfactory perception shifts, analysing how the description of the smells emitted by specific sources has changed over time. We first create a benchmark of selected smell sources, relying upon existing historical studies related to olfaction. We also collect an English text corpus by retrieving large collections of documents from freely available resources, spanning from 1500 to 2000 and covering different domains. We label such corpus using a system for olfactory information extraction inspired by frame semantics, where the semantic roles around the smell sources in the benchmark are marked. We then analyse how the roles describing *Qualities* of smell sources change over time and how they can contribute to characterise perception shifts, also in comparison with more standard statistical approaches.

1 Introduction

Over the past few decades, there has been a proliferation of studies in the realm of linguistics and perception (Winter, 2019; Bagli, 2021). However, there remains a distinct shortage of research dedicated to the tracking of perceptual changes over time. Although it has been already highlighted how much sensory language can be informative in terms of cultural attitudes (Majid and Burenhult, 2014), there has been a relatively limited exploration of how perceptual experiences are linguistically encoded over an extended period of time. A first attempt concerning the diachronic analysis of the olfactory domain using NLP has been presented in Menini et al. (2023), although this study was rather exploratory and relied on an existing approach utilizing word embeddings.

One of the reasons of the scarcity of studies using automatic approaches in this area is the difficult assessment of perceptual shifts due to the limited availability of suitable evaluation benchmarks. Therefore, in this paper, we first introduce a manually created benchmark containing a list of smell sources (mainly objects) that underwent some changes in the way their odour was described over time. This benchmark is based upon existing literature in historical studies and olfactory cultural heritage. We then present some analyses of perception shifts that compare standard statistical approaches to a novel framework based on the output of a system for olfactory information extraction. Our approach involves modelling perception shifts as changes in the association between a given smell source and its description in terms of (olfactory) quality. We show that focusing the analysis on text spans that the system identifies as being smell qualities makes the output more precise and tailored to the domain of interest. The results are validated on a selected set of smell sources from the benchmark, which is available at `https://github.com/dhfbk/scent-change`.

2 Related Work

There is limited research involving the diachronic analysis of sensory language, and the use of computational methods to study this phenomenon are even scarcer. Among the few works investigating this research direction, Strik Lievers (2021) proposes an analysis of the possible variations of olfactory lexicon in the transition from Latin to Italian. The results of the study show that olfactory lexicon did not present substantial alterations in its overall size and differentiation. However, there is evidence suggesting that it did evolve towards a more negatively-oriented lexicon. In Lievers and De Felice (2019), the authors test the hypothesis of the directionality of sensory adjectives in Latin and Italian from a diachronic perspective. This study provides evidence for the fact that the primary meanings of sensory adjectives and the hierarchy of synaesthetic metaphors did not undergo variations over time.

As regards the development of structured re-

Proceedings of the 4th Workshop on Computational Approaches to Historical Language Change, pages 143–152
December 6, 2023 ©2023 Association for Computational Linguistics

sources to investigate the evolution of sensory language, Menini et al. (2022a) present a multilingual taxonomy for olfactory-related terms, which was created semi-automatically, with the goal to describe the evolution of odours and smell sources' descriptions. Furthermore, in Menini et al. (2022b), the authors present a multilingual benchmark, manually annotated with smell-related information, to support the development of olfactory information extraction systems. Nevertheless, the first exploratory analysis of shifts in olfactory descriptors based on word embeddings between two time periods is introduced in Menini et al. (2023). The approach is inspired by the method for semantic change detection in El-Ebshihy et al. (2018), which was adapted to detect *perception shifts* rather than semantic ones. The hypothesis is that methods employed to detect how the meaning of a word changes over time (i.e. *semantic shift*) (Tahmasebi et al., 2021) can be adapted to analyse possible variations in the way sensory items are perceived and therefore described over time (i.e. *perception shifts*). The work we present in this paper further investigates this phenomenon by introducing a novel benchmark to study olfactory perception shifts. We also present a comparison between a 'traditional' PMI-based approach to shift detection (Hamilton et al., 2016) and our contribution that introduces an intermediate layer focusing on specific semantic roles.

3 Benchmark of Smell Perception Shifts

Given the difficulty to evaluate shifts in language use, we first develop a benchmark with the purpose to trace the history of some selected odors over time. This resource can be used as a test set for the evaluation of systems analysing possible changes in the way specific odors have been described in the past. We rely on historical studies in the olfactory domain (Tullett, 2019; Tullett et al., 2022) and on the Online Encyclopedia of Smell History and Heritage[1] to identify 16 words that domain experts consider particularly related to smell and whose perception may have changed over time: *asphalt*, *candle*, *brewing*, *car*, *chloride of lime*, *coffee*, *(perfumed) gloves*, *incense/frankincense*, *lavender*, *ozone*, *pomander*, *plastic*, *sulphur*, *tea*, *tobacco*, *wig*. For each of the above items, we then gather information on the **perception shift** it underwent, trying to address when this happened, whether it involved some changes in smell quality, whether it is connected to a change in location, and what type of shift it was. Indeed, we identify four possible types of perception shift, manually checked by two experts in olfactory language:

(a) **appearance**: in a mainly Eurocentric perspective, an odor that was not initially mentioned and that manifests itself at a certain point either due to trades and new habits (e.g. *coffee*) or as the outcome of inventions (e.g. *asphalt*);

(b) **disappearance**: in contrast with *appearance*, an odor associated with a particular era that slowly fades away over time. For instance, the pomander, a widely used item during the 16th century for carrying and diffusing fragrances, which eventually diminishes its presence, until its disappearance;

(c) **topic shift**: a change of environment/location in which a certain smell can appear, as the conditions of use or the meaning changes from a cultural point of view (e.g. *incense*, which disappears from Protestant churches after the Reformation of Henry the VIIIth, but which has been used in houses since the 18th century);

(d) **quality shift**: a change in the perception of the olfactory quality of a given odor over time, for instance the smell of candles that changes its olfactory connotation due to the different materials used to make them.

For each item in the benchmark, we specify one of the above types of perception shift, as well as the time period when the shift happened, the bibliographic or sitographic references, and in some cases the associated places for each period. Note that for each term in the benchmark different time periods may be related to a perception shift. In Table 1, we report an example of shift related to five smell sources with a brief description.

4 Olfactory Information Extraction

Our approach to analyse perception shifts in the olfactory domain relies on two components: *i)* a system for olfactory information extraction, and *ii)* a historical corpus of English, possibly well-balanced across topics and time periods, which is processed with the above system.

[1] https://odeuropa.eu/encyclopedia/

Smell source	Type of Shift	Brief Shift Description
Candle	*Quality*	From negative to positive perception due to materials' choice
Gloves	*Disappearance*	From being an object related to olfactory domain to not
Incense	*Topic*	A shift in the locations of usage
Ozone	*Quality*	From a connotation related to electricity to an healthy one
Tobacco	*Quality*	With the rise of snuff consumption, from positive to negative

Table 1: Selected smell sources from the benchmark

4.1 System Description

We develop a system for olfactory information extraction able to recognise smell-related information in a text. In particular, we detect olfactory events, typically evoked by smell words such as 'stink', 'odour', 'stench', 'whiff', 'stink', and the two semantic roles (or *frame elements*) that are more frequently mentioned in relation to these olfactory events, i.e. *Smell source* (items from where a smell comes from) and *Quality* (how such smell is described). For instance, in the sentence 'The tobacco has a pungent smell', 'The tobacco' would be *Smell source* and 'pungent' a *Quality*, while 'smell' would be the smell word evoking the olfactory event. This annotation framework, inspired by frame semantics (Fillmore and Baker, 2001) is described in detail in Tonelli and Menini (2021) and has been adopted to annotate an English benchmark (Menini et al., 2022b), which we use to train our system for olfactory information extraction.

For the supervised classifier, we adopt a multi-task learning approach (Caruana, 1993, 1997). In this configuration each task updates the model's shared parameters, leading to a more robust representation with less over-fitting. Each task corresponds to the classification of a single olfactory element, namely *Smell Word*, *Smell Source* and *Quality*.

We adopt a multi-task approach, since it performs better than a single multiclass classifier (see Table 2 for a comparison), and because simpler tasks, as can be smell word detection, can act as auxiliary task and share information for the classification of olfactory elements, which are more challenging to detect. To fine-tune the models, we use the MaChAmp framework (van der Goot et al., 2021), a toolkit for multi-task learning. The classification of each olfactory element was configured as a BIO task. Indeed, the tokens in the frame elements (that often span over multiple words) are marked with either B-FRAME_ELEMENT (beginning of a span), I-FRAME_ELEMENT (inside of a span) or O (outside the frame element).

All the results reported in Table 2 are the average of the experiments done with 10 different data splits, with each data split having 80% of the smell words and related olfactory elements as training data, 10% for validation and 10% as test. The splits are not completely random, as we keep the same temporal and domain distribution in every run.

We run a hyperparameter search[2] on one of the data splits and the best performance was obtained with a learning rate of $1e-4$ and a batch size of 32, and all the loss weight set to 1, which yield the best performance.

We report in Table 2 the performance of the multitask classifier on each of the three olfactory elements of interest, and compare it with a baseline obtained by fine-tuning the model with a single-task approach for multiclass classification. In both the configurations the fine-tuned model is bert-base-cased[3] (Devlin et al., 2019).

	Smell Word	Smell Source	Quality
Multitask	**0.871**	**0.571**	**0.758**
Multiclass	0.821	0.461	0.652

Table 2: Results of olfactory information extraction. Each result (F1) is the average of 10 different runs on 10 different data splits

4.2 Corpus Labelling

We launch the information extraction system on a set of historical corpora of English. We focus on seven freely available corpora:

Project Gutenberg:[4] A volunteer effort to digitize and archive cultural works, it contains different repositories, mainly in the literary domain.

[2]Search space: learning rate $[1e-3, 1e-4, 1e-5]$, batch size $[16, 32]$, training epochs $range(1, 20)$.

[3]https://huggingface.co/bert-base-cased

[4]https://www.gutenberg.org/

Early English Books Online (EEBO):[5] A collection containing documents published between 1475 and 1700 in different domains such as literature, philosophy, politics, religion, geography, history, politics, mathematics.
British Library:[6]A collection of 65,227 digitised volumes from the 16th to the 19th Century.
London Pulse Medical Reports:[7]A collection of 5800 Medical Officer of Health reports from the Greater London area from 1848 to 1972.
Wikisource:[8] An online digital library of free-content textual sources managed by the Wikimedia Foundation.
Eighteenth Century Collections Online (ECCO) :[9] A collection of over 3,000 titles printed in the United Kingdom during the 18th century.
UK Medical Heritage Library:[10] A collection of books and pamphlets from 10 research libraries in the UK, focused on the 19th and early 20th century history of medicine and related disciplines.

In Table 3 we provide an overview of the *Smell Sources* and *Qualities* instances extracted from the above set of corpora. Note that we report only the instances of smell sources present also in the benchmark (Section 3) that according to the system were part of an olfactory event. Qualities are less frequent than smell sources because they may not be necessarily mentioned when describing an odour.

Frame Element	Extracted Instances
Smell Sources	40,191
Qualities	39,521

Table 3: Number of *Smell Sources* from the benchmark extracted from the corpus and associated *Qualities*.

5 Analysis of Perception Shifts

In our analysis, we aim at detecting possible variations in the way specific smell-related concepts are described over time. For the sake of brevity, we focus our investigation on five *Smell Sources* selected from the benchmark (Section 3) that undergo some sort of change in terms of perception over time: *candle*, *gloves*, *incense/frankincense*, *ozone*, *tobacco*. Nevertheless, our approach to perception shift analysis can be generalised to any *Smell Source*. To conduct our study we use the corpus presented in Section 4.2, which was processed with the system for olfactory information extraction (Section 4.1).

5.1 Frequency Analysis of Smell-related Terms

The first analysis we perform is aimed at showing when specific items have become smell-related, i.e. when they started to be considered *smell sources* in olfactory descriptions. For instance, history scholars showed that leather *gloves* in the 17th Century used to be scented with perfumes to temper their bad smell coming from compounds used to make leather softer (Marx et al., 2022). Thus they were seen as strong olfactory objects at the time, while nowadays they are not considered 'smelly' items. In order to analyse the variations in the perception of items as being smell-related or not, for each of the five smell sources we compute the percentage of mentions in our corpus that are labeled also as Smell Source. Each time point (1 time point = 1 year) is calculated as the average percentage of a time range of 20 years centered around the time point. The results are plotted in the graph reported in Figure 1. Intuitively, a peak in the graph corresponds to a time period in which a term was strongly associated with the olfactory domain.

If we compare the plots for the five terms of interest, we observe that *incense* is the item that is overall more associated with the olfactory domain, in particular around 1860 and 1970, when almost 40% of its mentions are smell-related. The graph for *candle(s)*, instead, displays a growth after 1960, probably related to the widespread use of scented candles. As regards *glove(s)*, the graph shows that it stops being perceived as an olfactory object after 1950, as already mentioned before, but that nevertheless it was characterised as smell-related only rarely before that date (less than 2% of the mentions). Finally, *tobacco* and *ozone* are more 'modern' smells, in particular the latter, which was first used to characterise the aroma resulting from experiments with electricity around 1840.

5.2 PMI-based Analysis of Smell Qualities

While the analysis displayed in Figure 1 shows when a specific item was used in relation to the olfactory domain, it does not show *how* this rela-

[5]https://textcreationpartnership.org/tcp-texts/eebo-tcp-early-english-books-online/
[6]https://data.bl.uk/digbks/
[7]https://wellcomelibrary.org/moh/about-the-reports/about-the-medical-officer-of-health-reports/
[8]https://en.wikisource.org/
[9]https://textcreationpartnership.org/faq
[10]https://ukmhl.historicaltexts.jisc.ac.uk/home

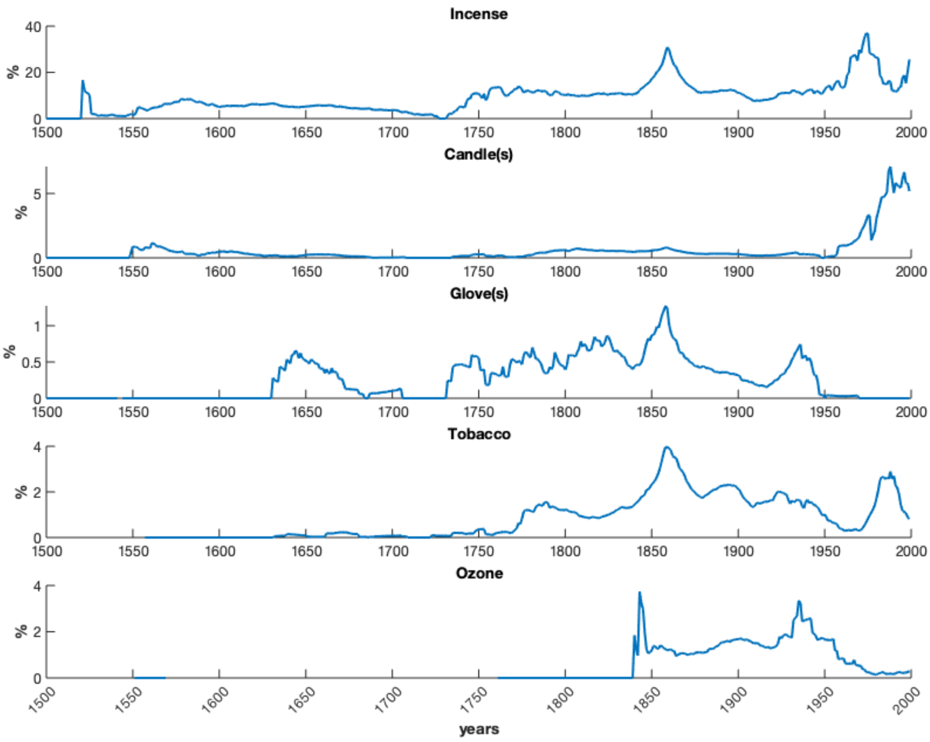

Figure 1: Percentage of term occurrences in our corpus that are also labeled as *Smell source*

tion was described, i.e. how an item's smell was characterised. To further address this aspect, we perform an analysis based on PMI with the goal to investigate more in detail the type of perception shifts of smell sources over the time. We opt for a PMI-based approach because it is a solution that can be straightforwardly combined with information from olfactory elements, usually consisting of few tokens, while other solutions like contextualized embeddings would require longer texts to be effective (Giulianelli et al., 2020). Furthermore, comparing a PMI-based analysis with and without olfactory element information gives us the possibility to assess the actual contribution of the latter to capture perception shifts.

We compute the association strength between a smell source and the words labelled as being their *Quality* by the information extraction system. The analysis is performed across different time periods marked as turning points in perception or attitudes towards these items, as identified in the benchmark. We calculate the PMI of a given smell source (w_1) and its associated qualities (w_2) in the following way:

$$PMI\,(w_1; w_2) = \log_2 \frac{P(w_1, w_2)}{P(w_1)P(w_2)}$$

where $P(w_1, w_2)$ is the probability of the smell source and a word/quality to co-occur, while $P(w_1)$ and $P(w_2)$ are their independent probabilities. We report in Table 4 the top-five qualities ranked by PMI for each smell source of interest in each time span. As a comparison, we also compute PMI for each of the five items in the whole corpus, dividing the analyses by the same time intervals, without considering the spans labeled as qualities. This comparison should highlight the difference between standard PMI-based analysis of shifts (see for example Hamilton et al. (2018)) and our approach, which targets the olfactory domain and is therefore carried out on specific text spans. We adopt the same setting as Hamilton et al. (2018) by considering a window of 4 terms before and after w_1. The top-ranked adjectives and nouns obtained without considering the olfactory annotation are reported in gray in Table 4. We report only these grammatical categories because they are prevalent also for *Qualities*.

We observe that in some cases the olfactory aspect is prevalent also if we do not consider only smell qualities, see for example the occurrences of 'perfumed/perfuming' in gray for all time periods related to *incense*. For *candle*, instead, PMI computed on raw text shows an alternation between the olfactory and the visual dimension, while focusing only on olfactory qualities allows us to capture the negative characterisation of candle smells in the past. Indeed, candles before 1800 were made from animal fats (pig tallow until 1700, followed by whale fat), resulting in a predominantly unpleasant odor (Muchembled, 2020). It wasn't until around 1830 that candles began to be fashioned from paraffin wax, leading to a likely shift in odor towards a more neutral quality. With the advent of kerosene lamps and the incandescent light bulbs, which rendered candles obsolete for illumination, these items found new purposes as decorations, ambient fragrance enhancers or votive offerings (Phillips, 1999).

As regards *ozone*, it is a peculiar element since it has no strongly associated smell qualities after 1950. Indeed, starting from 1840, the term "ozone" emerged to characterize the aroma resulting from experiments with electricity, often associated with thunder and lightning (Forster, 1813). However, as the 20th century unfolded, its connotation underwent a complete transformation and ozone was considered accountable for the healthful qualities found in mountainous and seaside air (Anonymous, 1910), while it was not perceived as an odorous element anymore. After this period, we have no data in our olfactory corpus since its primary role as a descriptor of scents diminishes until disappearing. Instead, it starts to be associated with atmospheric phenomena, particularly in relation to the ozone depletion event.

5.3 Perception Shift Analysis using PMI Vectors

We further use PMI to analyse the perception shifts involving the smell sources in different time periods. We first create vector embeddings containing the PMI value between each smell source in the benchmark and the fixed set of their context words, following an approach similar to the one presented in Hamilton et al. (2018). We consider as context only the spans labeled as *Qualities* of smell sources with a frequency higher than 3. In this way, for each item of the benchmark in each period, a vector was calculated, obtaining 56 vectors with 1,416 values. After keeping only the vectors containing more than 5 non-zero values, Pearson correlation between the vectors was used to calculate similarity/dissimilarity between them. We then utilized the correlations with the 'linkage' function within the MATLAB software to calculate the hierarchical clustering and finally represent it in a dendrogram (Figure 2 above). A high similarity between the vectors of the same smell source in two different time periods shows that the perception shift was limited. Moreover, different smell sources clustered together indicate that the qualities associated to them are similar. As a comparison, we create similar PMI-based embeddings but without considering the *Smell Source* and *Quality* information and using simple co-occurences in text in a window of 4 words between and after the occurence of the terms presented as Smell Sources in the benchmark (see approach presented in Section 5.2). This time the size of each embedding vector increased to 84,378 non-zero values and we calculate the dendogram in a similar way to what described above (Figure 2 below).

The above representation (PMI-embeddings based on *Smell Sources* and *Qualities*), shows that the vectors of the same Smell Source in different time periods tend to be more far apart and belong to different clusters, as can be observed for *gloves*, *ozone* and *incense / frankincense*. The last two terms, in particular, were considered interchangeable in the past (see yellow and green cluster), but from the beginning of the twentieth century frankincense seems to be used in different contexts (red cluster). On the contrary, the graph below tends to just group the vectors of the same smell sources across different periods, and seems therefore less suitable to capture shifts in time, see for example how *incense* and *frankincense* have all been clustered in the same group (red). This suggests that focusing the analysis only on elements that are relevant to the shift domain is beneficial to the quality of the outcome, enhancing its precision.

6 Discussion

Our analyses provide insights into the olfactory changes that were identified by domain experts, validating them from a quantitative point of view. However, we observe some differences in the outcome of our analyses. The results which better reflect the shifts manually identified in our benchmark are those whose changes were labeled as *quality shift*, namely 'candle', 'ozone' and 'tobacco'. This is not surprising considering that we focus on text spans classified as *Quality*. When it comes instead to'incense' and 'gloves', whose changes in perception are identified respectively as *topic shift*

Smell Source	Time period							
	1530 – 1600		1601 – 1800		1801 – 1900		1901 – 2000	
incense (8,310)	aromatical	*perfume*	vernal	*dragge*	noisomely	*nidorous*	somnolent	*donative*
	perfume	*odours*	breathe	*breezy*	frank	*sepulchred*	sacerdotal	*exasperate*
	sweet	*perfumed*	acceptable	*perfumed*	sanguinary	*perfuming*	frank	*wafting*
	fragrant	*fuming*	strange	*odours*	raptourous	*sweetsmelling*	sacred	*perfuming*
	odoriferous	*burnt*	holy	*perfume*	murky	*lawny*	heavenly	*enrage*
	1500 – 1700		1701 – 1829		1830 – 1900		1901 – 2000	
candle (1,186)	abominable	*lighted*	ferous	*snuffing*	salutary	*guttering*	fragrance	*fumigating*
	ill	*lighting*	offensive	*lighted*	corrupt	*arsenicated*	scented	*relighted*
	fetid	*blinking*	ill	*cerifera*	filthy	*relighted*	nauseous	*lighting*
	stink	*tallow*	odoriferous	*stationery*	snuff	*fumigating*	scent	*lighted*
	odoriferous	*cereus*	olfactory	*suppurating*	unsavoury	*sputtering*	perfume	*flickering*
	1500 – 1750		1751 – 1900		1901 – 2000			
gloves (670)	excellent	*perfumed*	perfume	*perfuming*	perfume	*gauntleted*		
	venomous	*fringed*	spanish	*pictured*	scented	*buttoning*		
	fine	*imbroidered*	remarkable	*cuticular*	scent	*boxing*		
	rich	*itchy*	costly	*worded*	odoriferous	*unbuttoning*		
	sweet	*scented*	excellent	*worshipful*	odorous	*rubber*		
	1600 – 1730		1731 – 1800		1801 – 1900		1901 – 2000	
tobacco (7,516)	hateful	*smoaked*	olfactory	*smoky*	undiminished	*pipeful*	homely	*latakia*
	fulsom	*nicotian*	perfume	*chewing*	hateful	*negrohead*	indefinable	*unmanufactured*
	ungrateful	*fulling*	peculiar	*fulling*	superficial	*unmanufactured*	spirituous	*chewing*
	offensive	*heroically*	grateful	*narcotick*	snug	*superexcellent*	stale	*carcinogenic*
	bad	*spicery*	pungent	*chewed*	vilest	*smoaking*	medicinal	*snuffing*
	1840 – 1899		1900 – 1950		1951 – 2000			
ozone (830)	restorative	*allotropique*	refresh	*ozonized*	None found	*photochemical*		
	inexhaustible	*oxidiser*	odorless	*allotropique*		*diurnal*		
	denser	*ozonized*	peculiar	*triatomic*		*antarctic*		
	electrical	*sterilizes*	fresh	*ultraviolet*		*nickelic*		
	obvious	*vigorating*	pungent	*transboundary*		*spheric*		

Table 4: Words most associated with a given smell source (left), ranked by PMI, in different time periods associated to time shifts in the benchmark. Terms in normal font belong to *Qualities*, while those in gray have been extracted regardless of olfactory information. Below each smell source the number of total occurrences in the corpus has been reported.

and *disappearance*, the results are less evident compared to the defined changes in the benchmark. Our findings suggest that different types of shifts may require distinct approaches for proper detection. Indeed, if we want to capture shifts mostly due to *disappearance*, an analysis like the one displayed in Figure 1 is probably more effective than the one based on PMI, in particular because we identify to what extent an item is considered a smell source, see for example the graph for ‘gloves’ after 1950.

Nevertheless, qualities associated with gloves in the olfactory analysis closely align with the way perfumed gloves were described during their historical use. Adjectives such as ‘venomous’ or ‘spanish’ are indeed part of the practice to perfume gloves, since venom is hidden by the perfume and has been used to kill monarchs, while ‘spanish’ recalls the origin of glove-perfuming tradition from Spain and Italy. This observation provides further confirmation that this analytical approach effectively identifies qualities exclusively related to the olfactory domain with a precision that faithfully reflects the actual historical data. On the contrary, with regards to ‘incense’, its pronounced olfactory significance, as previously observed in Section 5.1, presented a challenge in detecting noteworthy changes through the quality-based methodology. To uncover *topic shifts* in textual data, further research is needed.

7 Conclusions

In this paper, we describe a range of analyses to investigate changes in the perceptual descriptions of five selected smell-related objects in textual data. We first present a frequency-based analysis aimed at delineating the olfactory relevance of these items over time. We then perform a PMI-based analysis to identify the qualities linked to smell sources during specific time periods, with the attempt to uncover changes in descriptions that reflect actual

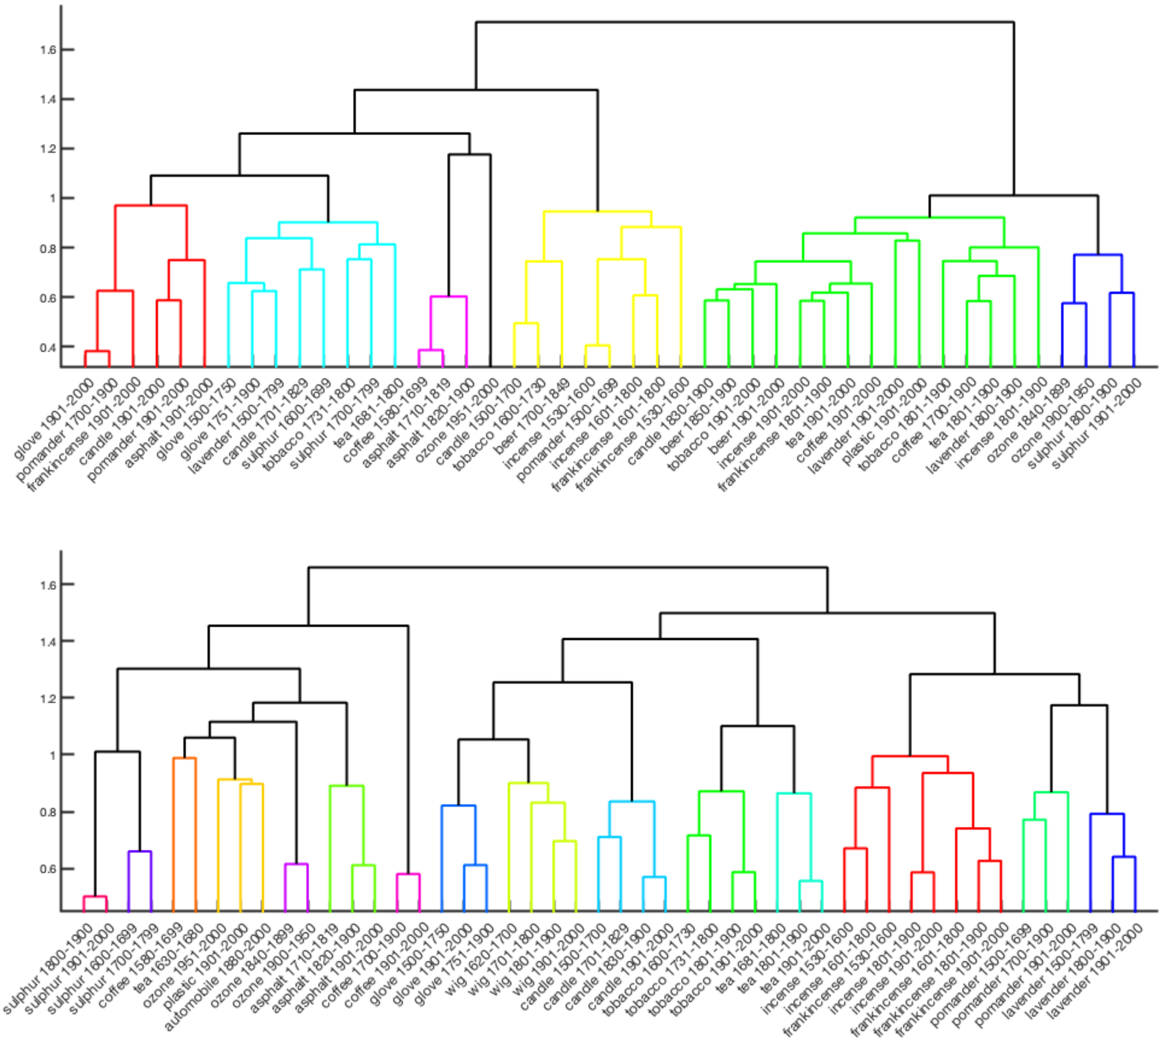

Figure 2: *Above*: Dendrogram clustering the PMI-embeddings of specific smell sources computed only on olfactory qualities for different time periods. *Below*: Dendrogram of the PMI-embeddings of the same words for the same time periods regardless of olfactory information.

shifts in perception. Additionally, we carry out a further analysis using PMI to represent the items of interest with vectors. The outcomes of these analyses support a twofold observation. On the one hand, the approaches previously used to detect diachronic semantic change prove effective in identifying variations also with regards to perceptual descriptions. On the other hand, the effectiveness of this adaptation is also due to the systematic encoding of the olfactory information offered by the frame-based approach. This work shows a novel approach which combines the power of frames in depicting semantic context and the tradition of semantic change detection to explore the evolution of olfactory language from a diachronic perspective. As previously discussed in Section 6, it would be worthwhile to expand our investigations by employing alternative frame elements to identify *topic shifts* associated with specific smell objects. Additionally, in the light of the observation made in Section 5.3, extending also the embedding-based approach to this type of shift detection could represents a promising path for prospective research. In future, we plan to further develop this methodology aiming towards a comprehensive approach for the study of perceptual shifts in texts.

Limitations

Like every corpus-based analysis, our work is strongly dependent on the corpus we were able to collect for this study. Although we tried to cover different domains and time periods, the limited availability of historical texts in good digital format is a major factor affecting our results.

Ethics Statement

No ethical issues are related to the current work.

Acknowledgements

This research has been supported by the European Union's Horizon 2020 program project ODEUROPA under grant agreement number 101004469.

References

Anonymous. 1910. Some recent applications of ozone. *Nature. 83, 47*.

Marco Bagli. 2021. *Tastes we live by: The linguistic conceptualisation of taste in English*, volume 50. Walter de Gruyter GmbH & Co KG.

R Caruana. 1993. Multitask learning: A knowledge-based source of inductive bias1. In *Proceedings of the Tenth International Conference on Machine Learning*, pages 41–48. Citeseer.

Rich Caruana. 1997. Multitask learning. *Machine learning*, 28(1):41–75.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.

Alaa El-Ebshihy, Nagwa M El-Makky, and Khaled Nagi. 2018. Using google books ngram in detecting linguistic shifts over time. In *KDIR*, pages 330–337.

Charles J. Fillmore and Collin F. Baker. 2001. Frame semantics for text understanding. In *In Proceedings of WordNet and Other Lexical Resources Workshop*.

Thomas Forster. 1813. Meterological observations made at cambridge from march 18 to april 8, 1813. *The Philosophical Magazine. 41*.

Mario Giulianelli, Marco Del Tredici, and Raquel Fernández. 2020. Analysing lexical semantic change with contextualised word representations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3960–3973, Online. Association for Computational Linguistics.

William L Hamilton, Jure Leskovec, and Dan Jurafsky. 2016. Diachronic word embeddings reveal statistical laws of semantic change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1489–1501.

William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 2018. Diachronic word embeddings reveal statistical laws of semantic change.

Francesca Strik Lievers and Irene De Felice. 2019. *Metaphors and perception in the lexicon*. John Benjamins Publishing Company.

Asifa Majid and Niclas Burenhult. 2014. Odors are expressible in language, as long as you speak the right language. *Cognition*, 130(2):266–270.

Lizzie Marx, Sofia Collette Ehrich, William Tullett, Inger Leemans, Cecilia Bembibre, IFF (www. iff. com) Odeuropa, and Museum Ulm. 2022. Making whiffstory: A contemporary re-creation of an early modern scent for perfumed gloves. *The American Historical Review*, 127(2):881–893.

Stefano Menini, Teresa Paccosi, Serra Sinem Tekiroglu, and Sara Tonelli. 2022a. Building a multilingual taxonomy of olfactory terms with timestamps. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4030–4039. European Language Resources Association.

Stefano Menini, Teresa Paccosi, Serra Sinem Tekiroğlu, and Sara Tonelli. 2023. Scent mining: Extracting olfactory events, smell sources and qualities. In *Proceedings of the 7th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, pages 135–140, Dubrovnik, Croatia. Association for Computational Linguistics.

Stefano Menini, Teresa Paccosi, Sara Tonelli, Marieke Van Erp, Inger Leemans, Pasquale Lisena, Raphael Troncy, William Tullett, Ali Hürriyetoğlu, Ger Dijkstra, Femke Gordijn, Elias Jürgens, Josephine Koopman, Aron Ouwerkerk, Sanne Steen, Inna Novalija, Janez Brank, Dunja Mladenic, and Anja Zidar. 2022b. A multilingual benchmark to capture olfactory situations over time. In *Proceedings of the 3rd Workshop on Computational Approaches to Historical Language Change*, pages 1–10, Dublin, Ireland. Association for Computational Linguistics.

Robert Muchembled. 2020. *Smells: a cultural history of odours in early modern times*. pages 53-92, John Wiley & Sons.

Gordon Phillips. 1999. *Seven Centuries of Light: The Tallow Chandlers Company*. Book Production Consultants plc. p. 74.

Francesca Strik Lievers. 2021. Smelling over time. the lexicon of olfaction from latin to italian. pages 369–397. John Benjamins, Amsterdam.

Nina Tahmasebi, Lars Borin, and Adam Jatowt. 2021. Survey of computational approaches to lexical semantic change detection. *Computational approaches to semantic change*, 6(1).

Sara Tonelli and Stefano Menini. 2021. FrameNet-like annotation of olfactory information in texts. In *Proceedings of the 5th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, pages 11–20, Punta Cana, Dominican Republic (online). Association for Computational Linguistics.

William Tullett. 2019. *Smell in Eighteenth-Century England: A Social Sense*. Oxford University Press.

William Tullett, Inger Leemans, Hsuan Hsu, Stephanie Weismann, Cecilia Bembibre, Melanie A. Kiechle, Duane Jethro, Anna Chen, Xuelei Huang, Jorge Otero-Pailos, and Mark Bradley. 2022. Smell, history, and heritage. *American historical review*, 127(1):261–309.

Rob van der Goot, Ahmet Üstün, Alan Ramponi, Ibrahim Sharaf, and Barbara Plank. 2021. Massive choice, ample tasks (MaChAmp): A toolkit for multi-task learning in NLP. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 176–197, Online. Association for Computational Linguistics.

Bodo Winter. 2019. *Sensory linguistics: Language, perception and metaphor*, volume 20. John Benjamins Publishing Company.