

# Smells like Teen Spirit: An Exploration of Sensorial Style in Literary Genres

Osama Khalid, Padmini Srinivasan

University of Iowa

{osama-khalid, padmini-srinivasan}@uiowa.edu

## Abstract

It is well recognized that sensory perceptions and language have interconnections through numerous studies in psychology, neuroscience, and sensorial linguistics. Set in this rich context we ask whether the use of sensorial language in writings is part of linguistic style? This question is important from the view of stylometrics research where a rich set of language features have been explored, but with insufficient attention given to features related to sensorial language. Taking this as the goal we explore several angles about sensorial language and style in collections of lyrics, novels, and poetry. We find, for example, that individual use of sensorial language is not a random phenomenon; choice is likely involved. Also, sensorial style is generally stable over time - the shifts are extremely small. Moreover, style can be extracted from just a few hundred sentences that have sensorial terms. We also identify representative and distinctive features within each genre. For example, we observe that 4 of the top 6 representative features in novels collection involved individuals using olfactory language where we expected them to use non-olfactory language.

## 1 Introduction

Sensory perceptions shape how we use language and communicate (Paradis, 2003). When we use sensorial words (i.e. words with meanings connected to our senses) like *fuzzy* or *stinky*, besides communicating sensorial experiences these also stimulate perceptual systems in the recipient's mind (Speed and Majid, 2020).

The space of senses – sometimes called the “Aristotelian” senses (Sorabji, 1971), include the five modalities: visual, auditory, haptic, gustatory, and olfactory. Relatively recently, linguistics and psychologists have added a sixth sense — interoception (Craig, 2002). This refers to the perception of sensations from inside the body, both physical such as hunger and pain, and emotional, such as joy. This sensory space has been the basis of much

prior research.

**Sensorial Linguistics** is about studying how language relates to the senses. A key focus has been to study how different sensorial experiences and perceptions are packaged into linguistic units (Winter, 2019). Researchers have looked at how some senses dominate in language (Winter et al.), how sensorial language varies across lexical categories (Lievers and Winter, 2018) and how sensory experiences influence sensorial language (Croijmans et al., 2019; Murphy, 2019). However, the domain of sensorial linguistics is still nascent with many unexplored questions.

**Stylometrics:** As individuals grow, besides consciously learning a vocabulary they also develop a linguistic style. Some stylistic elements may be acquired subconsciously and others by choice. Stylistic choices can reflect the individual's social reality or affective state (Savoy, 2020). Several categories of stylistic features have been identified, such as the use of function words and language complexity (Holmes, 1998). Stylometrics is important for goals such as author attribution and affect classification.

**Style in sensorial language use:** A key limitation in stylometrics is that linguistic style around sensorial language has not been studied systematically. A person who wants to express being depressed has several word choices. She could use “sad” or the less frequent “downcast”. Her propensity to choose one or the other may be considered as part of her linguistic style. Consider a cloudy scene. A person may use visual language focusing on color, and say “*the clouds are white*”. Another may use haptic language focusing on texture, “*the clouds are fluffy*”. While sensorial language is clearly important for communication, we do not yet know if there are distinguishable patterns in sensory language use at the level of individuals, texts, etc. This gap in stylometrics motivates us to ask the following about sensorial language style:

- **RQ 1:** Is the notion of sensorial style meaningful or is it a product of random chance?
- **RQ 2:** How much data do we need to get a stable representation of sensorial style?
- **RQ 3:** Does sensorial style vary with time?
- **RQ 4:** Which features are representative and distinctive of the individuals within each genre?

## 2 Representing Sensorial Language Style

Sensorial style may be represented at different levels of abstractions. At the lowest level, we can represent the proportion of an individual’s language that is sensorial and also examine the frequencies of different sensorial words. At a higher level of abstraction, we can ask how frequent are different sense modalities (visual, auditory, etc.) in an individual’s language. Alternatively, we can represent style by the extent to which a person’s use of sensory modalities aligns with general expectations. This is related to synaesthesia, where one sense modality is used when another is expected - a well studied phenomenon in sensorial linguistics (Lievers, 2015; De Ullmann, 1945). As an example in their work, Lievers and Huang (2016a) developed a lexicon of perception that they used to automatically identify perception related synaesthetic metaphors. Similarly we also approach the problem of sensorial style through the lens of their synaesthetic usage.

In the 2010 animated film ‘Despicable Me’, the character of Agnes hugs a unicorn and says “It is so *fluffy*”. One reason why this quote acquired somewhat of a meme status is because it subverted audience expectations of a more visual word like “*pretty*” or “*white*” to describe the unicorn. Instead she opts for the more unexpected haptic word “*fluffy*”. This substitution of visual language for haptic, a synaesthesia, might indicate that Agnes’ perceives the world in a more tactile manner rather than in a visual way. In order to assess if this is a stylistic tendency, we can examine all of Agnes’ language use and ask the general question: to what extent does she use haptic language in contexts where we generally expect visual language? We can ask similar questions related to each combination of expected versus observed sensory modalities. Observations for all combinations, including the homogeneous non-synaesthesia ones, are accumulated to form Agnes’ (or any other individual’s, group’s or genre’s) sensory style representation.

**Terminology and notation:** More formally, we consider a sentence to be a “sensorial sentence” if

it has at least one word or phrase that appears in a sensorial lexicon. Further, we define a “sense-focused sentence” to be a sensorial sentence with a single sensorial term selected as focus term. Thus, if a sensorial sentence has  $n$  sensorial terms then we derive from it  $n$  sense-focused sentences.

Assume  $S_i = \{S_{i1}, S_{i2}, \dots, S_{in}\}$  is the set of  $n$  sense-focused sentences identified from the writings of individual  $i \in I$ . Let  $C = \{H, V, I, O, G, A\}$  represent the modalities: Haptic, Visual, Interoceptive, Olfactory, Gustatory, Auditory respectively. Using  $\bar{N}$  to represent concepts that are “not-sensorial” we define  $\tilde{C}$  as:  $\tilde{C} = C \cup \{\bar{N}\}$ .

We also define two functions.  $F(\hat{s}_{ij})$  is a sensory lexicon lookup function that returns the sensorial category  $c \in C$  for the focused sensorial term  $\hat{s}_{ij}$  in the sense focused sentence  $S_{ij}$ . E.g., given the sensorial sentence “*The unicorn is white and fluffy*”, we have two sense-focused sentences  $S_{ij}$  and  $S_{ik}$  corresponding to the two sensorial terms,  $\hat{s}_{ij}$  and  $\hat{s}_{ik}$ . For  $S_{ij}$  with the focus term “*white*”  $F(\hat{s}_{ij})$  will return  $V$ . For  $S_{ik}$  with focus term “*fluffy*” it will return  $H$ .

The second function we define is  $M(S_{ij})$  which returns the “expected” modality  $c \in \tilde{C}$  for the same focus term in  $S_{ij}$ . We describe this function in Section 2.1.

**Calculating observed to expected ratios:** We represent an individual’s sense-focused sentences as a list of length  $|S_i|$ . Each entry is a pair of expected and observed modalities of the form  $[(M(S_{ij}), (F(\hat{s}_{ij}))]$ . For observed modality  $y \in C$  and expected modality  $x \in \tilde{C}$ , the observed to expected ratio  $\alpha_i^{xy}$  for individual  $i$  is:

$$\alpha_i^{xy} = \frac{|\{S_{it} : F(\hat{s}_{it}) = y \text{ and } M(S_{it}) = x\}|}{|\{S_{it} : M(S_{it}) = x\}|} \quad (1)$$

Note that these ratios are the informative units of sensorial style. For example, if the ratio is 1 when  $x$  and  $y$  are the same modality, a homogeneous combination, then the individual’s use of that modality is highly aligned with general expectation. On the other hand, if it is close to 0 then she deviates considerably from the expected use of modality  $x$ , i.e. there is greater synaesthesia.

**Style vectors:** For each  $x \in \tilde{C}$ , we then concatenate its 6 ratios into a vector of the form:

$$s_i^x = \sqcup_{y \in C} \alpha_i^{xy} \quad (2)$$

It follows that

$$\sum_{y \in \mathcal{C}} \alpha_i^{xy} = \begin{cases} 1, & \text{if } |\{S_{it} : M(S_{it}) = x\}| \geq 1 \\ 0, & \text{otherwise} \end{cases}$$

We can now define the sensorial style vector  $u_i$  of  $i \in I$  as a concatenation of the seven vectors, one for each expected modality. The size of  $u_i$  is  $x \times y = 42$ .

$$u_i = \sqcup_{x \in \mathcal{C}} s_i^x \quad (3)$$

## 2.1 Implementing function $M(S_{ij})$

Given a sense-focused sentence, function  $M$  returns the expected modality  $c \in \tilde{\mathcal{C}}$  of the sentence’s focus sensorial term as per general expectation in English. We leverage RoBERTa-MLM<sup>1</sup> as a stand-in for general English language usage. RoBERTa is a transformer based language model pre-trained on 160 GB of data (Liu et al., 2019). Prior works like (Mosbach et al., 2020) and (Sinha et al., 2020) have shown that language models learn the norms of the language on which they are trained. This makes them ideal for our task. We mask the focus sensorial words in our sentences and input them to the model. RoBERTa returns the probabilities for all the words in its vocabulary at each masked location. Probabilities represent likelihood of appearance of the words at that location. We use these probabilities to identify the expected sense modality at each masked location as follows.

Let  $W = \{(w_1, p_1), (w_2, p_2) \dots (w_N, p_N)\}$  be the ranked set of words returned by RoBERTa for a masked location (location of focus sensorial term) in  $S_{ij}$ ; top ranked has highest probability. Using  $F(w_k)$ , we lookup the sense for each word in the top 100. We combine this information to get an aggregate probability score  $\Pi(c, S_{ij})$  for each modality  $c$  as follows:

$$\Pi(c, S_{ij}) = \sum_{\substack{k \leq 100 \\ F(w_k) = c}} p_k \quad (4)$$

For greater confidence we only include  $S_{ij}$  in our analysis if its majority modality has  $\Pi(c, S_{ij}) > 0.5$ . We then define

$$M(S_{ij}) = \operatorname{argmax}_{c \in \tilde{\mathcal{C}}} \Pi(c, S_{ij}) \quad (5)$$

In essence,  $M(S_{ij})$  returns the expected modality with the highest aggregate probability for the focused sense word in  $S_{ij}$  as determined using RoBERTa.

<sup>1</sup>We experimented with BERT as well, however, RoBERTa gave us more accurate results.

## 2.2 Sensorial Lexicon

Modality	Lexicon			
	Original		Modified	
	$N$	%	$N$	%
Visual	29552	75.0	9419	50.2
Interoceptive	3546	9.0	3449	18.4
Auditory	4528	11.5	3803	20.3
Haptic	675	1.7	972	5.2
Gustatory	890	2.3	890	4.7
Olfactory	216	0.5	216	1.2
Total	39407	100	18749	100

Table 1: Distribution of modalities in original Lynott et al. (2020) lexicon and our modified subset lexicon.

We use the sensorimotor norms lexicon published recently (Lynott et al., 2020) which has 39,954 concepts from the English Language. Brysbaert et al. estimates that the average adult lexicon is composed of approximately 42,000 words. Therefore this lexicon approximates a significant majority of everyday English.

Each concept was rated by annotators along a 0-5 scale for the six modalities (Auditory, Gustatory, Haptic, Olfactory, Visual, Interoceptive) For example “*fluffy*” is rated 4.41 for Haptic, 0.29 for Gustatory, 3.77 for Visual, 0.35 for Interoceptive and 0 for Auditory and Olfactory. The dimension with the highest rating is the dominant modality.

Dominance alone is not enough to ensure that a concept belongs to a particular sensorial modality since almost half of the concepts score less than 2.55 on any sense modality. Therefore, we filter the lexicon by ranking all concepts in a given modality by their rating and selecting only those in the top quartile. This ensures strong alignment to dominant modalities. Table 1 describes the lexicons<sup>2</sup>.

## 3 Methods

### 3.1 RQ 1: Is Sensorial Style a Product of Random Chance?

In order to be meaningful our representation of sensorial style should not be a product of random chance. If an individual chooses sensorial modalities randomly and not deliberately then we expect her observed and expected modality distributions to be independent of each other. For example in “*the clouds are white*” the expected modality may be visual but the individual randomly chooses from one of the six senses. We use this random model in our analysis.

<sup>2</sup>There are other sensory lexicons like (Liewers and Huang, 2016b) and (Lynott and Connell, 2013). Besides being the largest and most recent the Lynott et al. (2020) sensorimotor lexicon is the only one to include ratings for interoception.

As a first step, for any  $i \in I$  with a set of  $n$  sense-focused sentences  $S_i$ , we define,  $\Gamma(i)$ , the distribution of the sense modalities in  $C$  observed in  $S_i$ . For each sense-focused term  $s_{ij}$  in  $S_i$ , we use a function  $\bar{F}(s_{ij}, \Gamma(i))$  that returns a random modality  $c \in C$  with distribution  $\Gamma(i)$ .

For each  $i \in I$ , we create  $m$  random pseudo-documents  $\mathcal{R}_i = \{\mathcal{R}_i^1, \mathcal{R}_i^2 \dots \mathcal{R}_i^m\}$ . Each random pseudo-document has the same set of sense-focused sentences  $S_i$ . However, instead of using  $F(s_{ij})$  to look up the modality, we use  $\bar{F}(s_{ij}, \Gamma(i))$  to get a random modality. Equation 6, a modification of equation 1, gives us  $\bar{\alpha}_{ik}^{xy}$  which is used to calculate the style vector  $u_i^k$  for random pseudo-document  $\mathcal{R}_i^k$ .

$$\bar{\alpha}_{ik}^{xy} = \frac{|\{S_{ij} : \bar{F}(s_{ij}, \Gamma(i)) = y \text{ and } M(S_{ij}) = x\}|}{|\{S_{ij} : M(S_{ij}) = x\}|} \quad (6)$$

Thus, for each  $i \in I$  with sensorial style vector  $u_i$ , we have  $m$  random style vectors  $\{u_i^1, u_i^2 \dots u_i^m\}$  generated from the random pseudo-documents in  $\mathcal{R}_i$ .

Let  $U_i = \{u_i\} \cup \{u_i^1, u_i^2 \dots u_i^m\}$ . For each vector  $v \in U_i$ , we calculate its average cosine similarity with all other elements in  $U_i$ . Ranking the elements of  $U_i$  by decreasing order of average similarity we check whether the style vector  $u_i \in U_i$  has lower average similarity than at least 95% of the vectors in  $U_i$  (i.e.  $p$ -value  $< 0.05$ ). If so, we infer with 95% confidence that  $i$ 's style vector,  $u_i$ , is not random and therefore likely a product of an individual stylistic choice.

### 3.2 RQ 2: How much data is needed to describe sensorial style?

Given the set of sense focused sentences  $S_i$ , where  $|S_i|=n$ , we randomly sample subsets from  $S_i$  of size  $k$  and compute the style vector from each sample. We explore how increasing the values for  $k$  affect style convergence.

For a given sentence set size  $k$ , we identify  $m$  random samples (with replacement) of  $S_i$ , each of size  $k$ . Thus, we create a set of sentence sets,  $T_i^k = \{\hat{T}_i^{k1}, \hat{T}_i^{k2}, \dots, \hat{T}_i^{km}\}$ . For each sentence set  $\hat{T}_i^{kj} \in T_i^k$ , we use the method discussed in Section 2 to generate the corresponding style vector  $\hat{u}_i^{kj}$ . This gives us a set of  $m$  sensorial style vectors,  $\hat{u}_i^k = \{\hat{u}_i^{k1}, \hat{u}_i^{k2}, \dots, \hat{u}_i^{km}\}$ . We then use cosine similarity to calculate the average pairwise similarities between all elements in  $\hat{u}_i^k$ . We recompute this average self-similarity  $\text{sim}(\hat{u}_i^k)$  for different values of  $k$  in increasing steps of  $r$ . We say that the style

Genre	# Authors	# Works	#Sentences	# Sensorial Sentences	# Sensorial Expressions
Novels	130	317	1,525,894	156,570 (10%)	474,299
Lyrics	5,321	20,785	1,007,090	754,572 (75%)	1,501,501
Poetry	1,246	3,315	85,236	4,979(6%)	8,209

Table 2: Dataset details for each genre. The percentage of total sentences that are sensorial is in parentheses.

of the individual has converged for a minimum of  $k$  sensorial sentences if  $\text{sim}(\hat{u}_i^k) \approx \text{sim}(\hat{u}_i^{k+r})$ , where  $k+r$  is the next sentence set size tested.

### 3.3 RQ 3: Does sensorial style vary over time?

Here we investigate whether style vectors evolve over time spans. We segment the writings by time and consider how the average similarity in style varies with temporal distance. We first identify all pairs of time points  $t_a$  and  $t_b$  that are  $\gamma$  duration apart. We then build style vectors for each author with text anchored at  $t_a$  and for each author with text anchored at  $t_b$ . We then compute the average pairwise similarity between  $t_a$  and  $t_b$  style vectors. We repeat this for all values of  $\gamma$  that are of interest.

We use a notion of windowing around the time points ( $t_a$  and  $t_b$ ) to reduce noise. Each window is of size  $\delta$  and distributed equally around each time point. For example, for  $t_a$  we create an individual's style vector from all the sense-focused sentences that were published in the range  $t_a - \frac{\delta}{2} < \tau_a < t_a + \frac{\delta}{2}$ .

### 3.4 RQ 4: Which features are representative and distinctive of the individuals within each genre?

A genre can be represented by the set of sensorial style vectors of its members. Each style vector is composed of 42 features that explore synaesthesia. We are interested in exploring which features are representative of the members of a genre, and also features that make the members distinct.

We consider a sensorial style feature to be representative if the variation in its usage is low. This would indicate that the members use the feature in a consistent manner. Formally, a stylistic feature  $\alpha^p$  is more representative for the members of a genre than another feature  $\alpha^q$  if its standard deviation  $\sigma(\alpha^p)$ , across all the members is lower than the standard deviation  $\sigma(\alpha^q)$ . At the other end, a high variation would indicate that the feature is distinctive amongst the members.

## 4 Datasets

We analyze 3 literary genres — novels, poems, music lyrics. Compared to poems and lyrics, novels

can span tens of thousands of sentences. Additionally, novels and poetry are generally associated with a single author. Lyrics are sometimes collaborations, however, we assume an artist would not perform a song that is in a style they do not like. Thus, we assume music lyrics to be a reflection of the artist’s style.

**Novels:** We collected English language novels from the Domestic fiction genre of Project Gutenberg<sup>3</sup>. There were 317 works written by 130 authors, with the earliest by Henry Fielding from the early 18<sup>th</sup> century and the latest by Rebecca West from the mid-20<sup>th</sup> century.

**Lyrics:** We collected songs that were listed on the Billboard Hot 100 charts, 1963 to 2021 (inclusive). This weekly chart ranking of song popularity is considered the industry standard (Whitburn, 2010). We assume the first time a song is listed to be its year of production. We obtained song lyrics using the Genius API<sup>4</sup>. There are 20,785 song lyrics.

**Poetry:** Following works like (Lou et al., 2015), we used the corpus of poems available on the Poetry Foundation’s<sup>5</sup> website. To make this dataset more comparable to the lyrics dataset, we only included works published after 1963.

## 5 Results

Genre	> 95 <sup>th</sup>	N
Novels	112	123
Lyrics	701	735
Poetry	20	85

Table 3: Number of individuals with lower average similarity than 95% of random vectors.

We present our results in two parts. First, we make our general observations. Second, we present results related to our specific research questions.

### 5.1 General Observations

**Domination over lower senses:** The five Aristotelian sensorial modalities have classically been thought of as part of a hierarchy with vision and audition dominating over the three so-called “lower senses”—Touch, Taste, Smell (Howes, 2010)<sup>6</sup>. This hierarchy manifests in the frequency of language use with the visual and auditory modalities

<sup>3</sup><https://www.gutenberg.org/>

<sup>4</sup>Some lyrics were not available on <http://www.genius.com>, because they were instrumentals like the “Star Wars Theme” which hit No. 1 in 1977, or were not in the Genius database.

<sup>5</sup>The Poetry Foundation was established in 2003, and one of its goals is to make “the best poetry” accessible (<https://www.poetryfoundation.org/foundation/about>).

<sup>6</sup>Note interoception is generally not considered in discussions of this hierarchy, possibly because of its relatively recent inclusion (Connell et al., 2018).

being used more often than the lower senses (Majid et al., 2018). We have consistent results. Figure 2 shows that for all three genres visual and auditory dominate over the “lower senses”. Concepts associated with haptic, gustatory and olfactory modalities — combined, form less than 10% of the total sensorial language.

While auditory dominates the “lower senses” in all three genres, it occurs less than half as often as visual. Going beyond the classical five senses, in all cases interoception dominates the three “lower senses” surpassing audition in this regard. Additionally, in lyrics, interoception is as common as visual. Clearly interoception with its emphasis on sensations within the body, both physical and emotional, is important in language.

**Sensorial style across genres:** We investigate how sensorial style varies across genres. Using the method described in Section 2 we calculate genre-level sensorial style vectors by combining sentence sets at the genre level. We show the distribution over 42 sensorial combinations in Figure 1 for just Lyrics. Each cell represents an expected-observed modality combination.

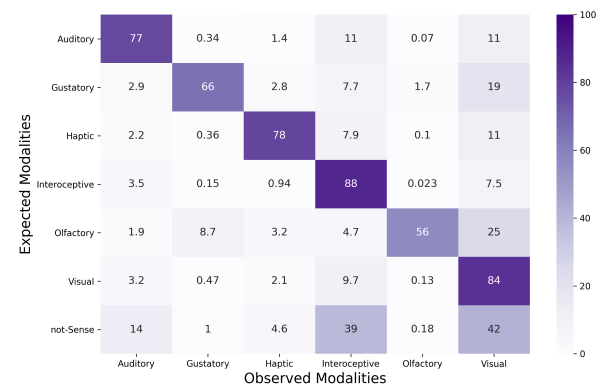


Figure 1: Distribution of expected-observed modalities in Lyrics. Note that we calculate proportions using equation 1, however, for illustrative purposes we show sensorial distribution as percentages.

**Observed modalities are largely as expected with some exceptions:** The diagonal values which are in the range 56 to 88% indicate that the observed modalities are generally consistent with expected modalities. That is, the individuals in our datasets select from the 6 sensorial modalities in a manner that is consistent with the general norms of language use. The highest consistency is for interoceptive and the lowest is for olfactory. We observe this trend across all genres (see appendix).

Looking at off diagonal values, we observe visual language used in 25% of the cases where we

	Novels	Lyrics	Poetry
	$I - O$	$I - O$	$A - G$
	$\bar{N} - O$	$\bar{N} - O$	$A - O$
	$A - O$	$V - O$	$G - I$
	$I - G$	$A - O$	$H - O$
	$V - O$	$I - G$	$O - A$
	$A - G$	$V - G$	$O - H$
<b>Range</b>	(0.00,0.01)	(0.00,0.02)	(0.00,0.00)

Table 4: The top representative features for each genre.

	Novels	Lyrics	Poetry
	$G - G$	$H - H$	$A - A$
	$O - O$	$G - G$	$I - I$
	$H - H$	$A - A$	$H - H$
	$G - V$	$H - V$	$\bar{N} - V$
	$H - V$	$G - V$	$G - G$
	$A - A$	$A - I$	$\bar{N} - I$
<b>Range</b>	(0.43,0.15)	(0.44,0.19)	(0.47,0.27)

Table 5: The top distinctive features for each genre.

Each feature is an (expected,observed) pair e.g.,  $I - O$  in means that we observe Olfactory language when we expected Interoceptive language. This is the most representative feature for Novels and Lyrics as it has the lowest standard deviation. We include the range of variances of the top most distinctive and representative features.

expected olfactory language, 19% and 11% of the cases where we expected gustatory and haptic language, respectively. This usage of visual language as a replacement for lower senses was observed across all 3 genres. This replacement or *cross-modal compensation* might be because the lower senses do not have a strong relation with the perceptual system and consequently individuals might be relying on visual language as a semantic scaffold to compensate for the weaker perceptual system of the lower senses (Speed and Majid, 2020).

We also observe (in all three genres) that in more than 90% of instances where we expected to see non-sensorial language we instead observed interoception. This might also be because interoception dominates in our data and is consistent with observations about higher senses in the literature (Majid et al., 2018).

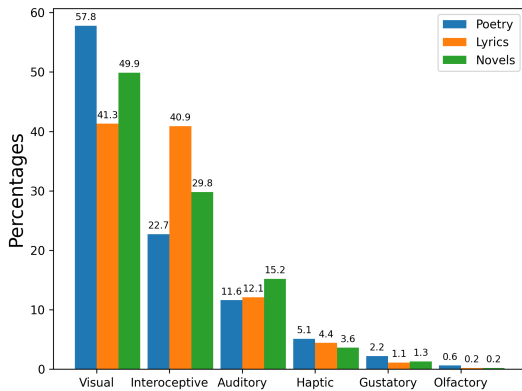


Figure 2: Distribution of observed modalities.

## 5.2 Results for research questions

**RQ1: Is sensorial style a product of random chance?** We investigate whether sensorial style is motivated by the individual choices or whether it is a product of randomness. Table 3 provides the results. If sensorial style is a non random phenomenon and a product of individual choice and intent, we expect the sensorial style vectors to be

distinct from random vectors. That is, we expect them to have a lower average similarity as compared to random vectors (generated by our random model). Methodological details are in Section 3.1).

Considering all individuals with more than 10 sense-focused sentences<sup>7</sup>, more than 90% of the individual sensorial style vectors in the novels and lyrics datasets are non-random. However, in the poetry dataset only around 23% of the individual vectors in the poetry dataset were non-random. A possible reason for the difference in poetry is likely data sparsity — fewer sensorial expressions/ author (see Table 2). Exploring this intuition further we find that the non-random vectors in poetry have on average 159 sense-focused sentences while the remaining vectors that looked random had on average 24 sense-focused sentences. Similarly, 8 out of the 10 most prolific individuals had non-random vectors. However, none of the 10 least prolific individuals had non-random vectors. These support our intuition regarding data sparsity being the cause of the difference in poetry.

**RQ2: How much data do we need to get a stable representation of sensorial style?** We evaluate the average similarity for each individual with progressively larger samples sizes,  $k$ , of their sense-focused sentences. We chose a range of values of  $k$  from  $k = 1$  to  $k = 10$  with a granularity of 1, from  $k = 10$  to  $k = 100$ , and  $k = 100$  to  $k = 1,000$  with granularities of 10 and 100. In Figure 3 we summarize these results with the median of average similarity across all individuals in each genre. We say the sensorial style vector has converged at a  $k$  value when the graph becomes more or less horizontal from that point onward. From the figure we see that as  $k$  increases similarity increases (within the  $m$  samples of size  $k$ ). We note that lyrics reaches a median average similarity of 0.90

<sup>7</sup>Because of the volume of lyrics, we limit the analysis to individuals with  $> 500$  sense-focused sentences.

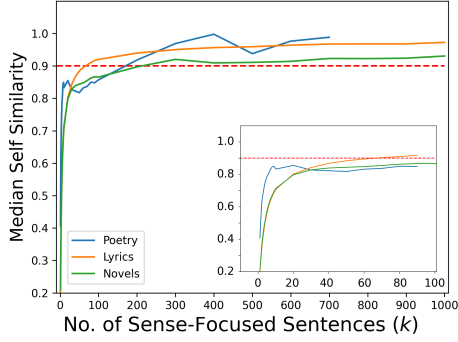


Figure 3: Convergence of style vectors as a function of  $k$ , the number of sense-focused sentences sampled.

with a sample of less than 100 sense focused sentences. In contrast, we need between 200 and 300 sentences to get the same 0.9 median average similarity for novels. Compared to novels and lyrics the plot for poetry has some fluctuations at  $k \geq 400$ , perhaps because there are only 7 poets with more than 100 sentences.

### RQ 3: Does sensorial style vary over time?

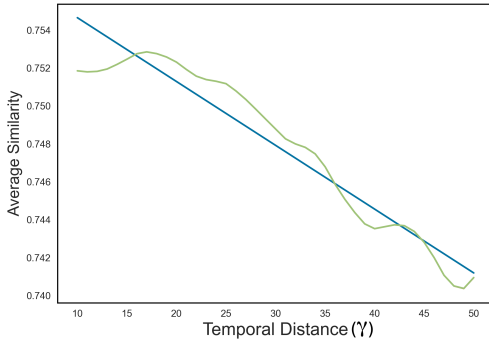


Figure 4: Average similarity between music lyrics as a function of their temporal distance. The blue line represents the linear approximation of the relationship. As temporal distance increases similarity decreases slightly.

We now explore whether sensorial style has changed over time in lyrics. As observed earlier, the poetry dataset is sparse so we do not include it. And with novels, it is sometimes challenging to pinpoint a single publication year, since some texts are written over multiple years<sup>8</sup>. Using the method described in Section 3.3, we measure average similarity in style of lyrics between pairs of individual sensorial style vectors that are a temporal distance  $\gamma$  away from each other. We use window length  $\delta = 1.5$  years. Figure 4 illustrates our findings.

We observe that as temporal distance between works increases the similarity of sensorial style decreases. The augmented Dickey-Fuller test had a

<sup>8</sup>As an example, Louisa May Alcott’s *Little Women* was published in two volumes spanning two years, 1868 and 1869.

$p$ -value = 0.96 meaning we cannot reject  $H_0$ , that the distribution is not-stationary (Cheung and Lai, 1995). However, the decrease is very slight; the 50 year drop in similarity is only 0.01. Approximating the relationship between the average similarity and temporal distance as a linear process, we note that the average similarity decreases very slowly at a rate of  $3.37 \times 10^{-4}$  per year.

### RQ 4: Which features are representative and distinctive of the individuals within each genre?

Table 4, shows the top-6 features that are representative of the members of each genre. We observe that these are all synaesthetic. Additionally the use of olfactory language in non-olfactory contexts is a representative feature in a majority of table cells (10/18). The standard deviations of top representative features is relatively low (between 0.00 to 0.02). This would indicate that the level of consistency for these top features is generally consistent for the three genres. One takeaway from these observations can be that in synaesthetic contexts, individuals are more prone to using lower senses (like olfaction) in a more consistent manner.

Table 5, shows the top-6 features that had the highest standard deviation between the members of each genre and were the most distinctive. The distinctive features were predominantly non-synaesthetic. In the cases where the distinctive feature was synaesthetic, the observed modality was either interoceptive or visual. This would indicate that there is greater diversity in expressions that rely on higher senses as a semantic scaffold.

For each genre, we rank all the sensorial style features by the standard deviation and compare them using Pearson’s correlation. We observe that, the features are highly correlated. Lyrics had a high correlation with both novels (0.75) and poetry (0.81), while poetry and novels had a slightly lower correlation of 0.48.

### Can sensorial style be used for prediction tasks?

We investigate whether sensorial style features can be used to identify genre. We compare against other standard style representations: LIWC (Tausczik and Pennebaker, 2010) and content-free words (CFW) vectors (Hughes et al., 2012). We use standard 5 fold cross validation for each experiment to train and test a random forest classifier. We consider the most prolific 50 authors/genre.

Table 6 shows the results. We observe that sensorial style predicts genre with a high level of accuracy ( $> 90\%$ ). While the other representations achieve close to perfect accuracies, key to note is

Method	Baseline	Sensorial Style	LIWC	CFW
Features	—	42	73	307
Accuracy	0.33	0.91	0.99	0.99

Table 6: Prediction accuracy of the different features.

that our goal is less about beating baselines and more about understanding the kinds of signals conveyed by sensorial style.

## 6 Case Study: Sensorial style in Lyrics

As a small illustration, we explore how sensorial style varies across different songs composed by the same artist. We consider all 962 artists who had at least 5 songs in the Hot 100 and extract a sensorial vector for each song. We then measure the average pair-wise cosine similarity (self similarity) amongst the songs of each artist. Almost 80% of the artists had an average self similarity  $\geq 0.70$ . Only two artists had a self similarity  $< 0.50$ .

The rapper *NF* was the most consistent artist with an average self similarity of 0.93. Conversely, the least consistent artist was the rock musician *Tommy James* with an average similarity of 0.42. For example, in the song “When I Grow Up”, *NF* used auditory language non-synaesthetically in 85.7% of the cases and for the song “NO NAME” this happened in 76.9% of the cases. The similarity between these two songs was over 95%.

In contrast, “Nothing to Hide” and “Ball and Chain” by *Tommy James* had a similarity of 0.32. In “Ball and Chain”, the artist uses visual language in all the instances where it was expected. However in “Nothing to Hide”, he uses interoceptive language synaesthetically instead of visual in about 57% of the cases. This case study demonstrates a method for exploring sensorial style and their variations across writings at the individual level.

## 7 Related Work

There are no directly comparable studies examining sensorial style for large numbers of individuals that consider interoception. Instead we briefly review closely allied topics.

**Sensorial language:** Sensorial language is not uniformly distributed across the six sensory modalities as reflected in sensorial lexicons, such as the one we use (Lynott et al., 2020). This is also observed in large text collections. In their analysis of 8 million words from around 7,000 English texts (Koblet and Purves, 2020) found over 28,000 visual descriptions and only 78 referring to the olfactory modality. Similar findings for multiple corpora are

observed in (Winter et al., 2018). Our results are consistent with these prior works.

**Sensorial language, the brain & emotion:** The salience of sensorial words is known to be highly correlated with the volumes of cortical activation in the brain (Reilly et al., 2020). Lievers (2015) show that there is directionality to how senses are substituted for each other. Winter (2016) found that gustatory and olfactory words (e.g., ‘stinky’, ‘delicious’) are on average more emotionally valenced than visual and auditory words and these also appear in more emotionally valenced sentences.

Bubl et al. (2010), show that alterations in mental states have a direct effect on perception; specifically, that depression directly impacts how the color blue was perceived. Kernot et al. (2016), found a decrease in the novelist Iris Murdoch’s use of olfactory language following her diagnosis of depression and Alzheimer’s. We credit this study for providing us with the hint that sensorial language may lead to a sensorial style. We take their sensorial style analysis forward with larger collections of authors, several genres and a more informative representation of sensorial style.

## 8 Limitations and Conclusions

We have shown that individuals have sensorial language style and that this sensorial style is a non random phenomenon for novelists and musicians and therefore is likely developed intentionally. Interestingly, we also found that it takes just a few hundred sentences to extract stable sensorial style representations. We also show that sensorial style in lyrics largely stable over time; the average similarity decreases at a rate of  $3.37 \times 10^{-4}$  per year.

Additionally, we show that sensorial style vectors seem to perform well at genre identification. The performance was high ( $> 0.90$ ), however, it was not close to perfect as with other style representations. The question about how sensorial style representations can be improved to increase performance requires further investigation.

Our study is limited in that our method relies heavily on the underlying Lynott et al. (2020) lexicon, and as with similar studies, is only as good as the lexicon. Additionally, we assume that each term is associated with a single sensorial modality. However, as research in psychology and neurology has shown, sensorial language is cross-modal. We leave this analysis to future work. In summary, we take a first step towards showing that sensorial style has a legitimate role in stylometrics research.



## References

- Marc Brysbaert, Michaël Stevens, Paweł Mandera, and Emmanuel Keuleers. 2016. How many words do we know? practical estimates of vocabulary size dependent on word definition, the degree of language input and the participant's age. *Frontiers in psychology*, 7:1116.
- Emanuel Bubl, Elena Kern, Dieter Ebert, Michael Bach, and Ludger Tebartz Van Elst. 2010. Seeing gray when feeling blue? depression can be measured in the eye of the diseased. *Biological psychiatry*, 68(2):205–208.
- Yin-Wong Cheung and Kon S Lai. 1995. Lag order and critical values of the augmented dickey–fuller test. *Journal of Business & Economic Statistics*, 13(3):277–280.
- Louise Connell, Dermot Lynott, and Briony Banks. 2018. Interoception: the forgotten modality in perceptual grounding of abstract and concrete concepts. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1752):20170143.
- Arthur D Craig. 2002. How do you feel? interoception: the sense of the physiological condition of the body. *Nature reviews neuroscience*, 3(8):655–666.
- Ilja Croijmans, Laura J Speed, Artin Arshamian, and Asifa Majid. 2019. Measuring multisensory imagery of wine: The vividness of wine imagery questionnaire. *Multisensory research*, 32(3):179–195.
- Stephen De Ullmann. 1945. Romanticism and synaesthesia: A comparative study of sense transfer in Keats and Byron. *PMLA*, 60(3):811–827.
- David I Holmes. 1998. The evolution of stylometry in humanities scholarship. *Literary and linguistic computing*, 13(3):111–117.
- David Howes. 2010. *Sensual relations: Engaging the senses in culture and social theory*. University of Michigan Press.
- James M Hughes, Nicholas J Foti, David C Krakauer, and Daniel N Rockmore. 2012. Quantitative patterns of stylistic influence in the evolution of literature. *Proceedings of the National Academy of Sciences*, 109(20):7682–7686.
- David Kernot, Terry Bossomaier, and Roger Bradbury. 2016. The impact of depression and apathy on sensory language. *Open Journal of Modern Linguistics*, 7(1):8–32.
- Olga Koblet and Ross S Purves. 2020. From online texts to landscape character assessment: Collecting and analysing first-person landscape perception computationally. *Landscape and Urban Planning*, 197:103757.
- Francesca Strik Lievers. 2015. Synaesthesia: A corpus-based study of cross-modal directionality. *Functions of language*, 22(1):69–95.
- Francesca Strik Lievers and Chu-Ren Huang. 2016a. A lexicon of perception for the identification of synaesthetic metaphors in corpora. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 2270–2277, Portorož, Slovenia. European Language Resources Association (ELRA).
- Francesca Strik Lievers and Chu-Ren Huang. 2016b. A lexicon of perception for the identification of synaesthetic metaphors in corpora. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 2270–2277.
- Francesca Strik Lievers and Bodo Winter. 2018. Sensory language across lexical categories. *Lingua*, 204:45–61.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Andres Lou, Diana Inkpen, and Chris Tanasescu. 2015. Multilabel subject-based classification of poetry. In *The Twenty-Eighth International Flairs Conference*.
- Dermot Lynott and Louise Connell. 2013. Modality exclusivity norms for 400 nouns: The relationship between perceptual experience and surface word form. *Behavior research methods*, 45(2):516–526.
- Dermot Lynott, Louise Connell, Marc Brysbaert, James Brand, and James Carney. 2020. The Lancaster sensorimotor norms: multidimensional measures of perceptual and action strength for 40,000 English words. *Behavior Research Methods*, 52(3):1271–1291.
- Asifa Majid, Seán G Roberts, Ludy Cilissen, Karen Emmorey, Brenda Nicodemus, Lucinda O'Grady, Benjie Woll, Barbara LeLan, Hilário De Sousa, Brian L Cansler, et al. 2018. Differential coding of perception in the world's languages. *Proceedings of the National Academy of Sciences*, 115(45):11369–11376.
- Marius Mosbach, Stefania Degaetano-Ortlieb, Marie-Pauline Krielke, Badr M Abdullah, and Dietrich Klakow. 2020. A closer look at linguistic knowledge in masked language models: The case of relative clauses in American English. *arXiv preprint arXiv:2011.00960*.
- Claire Murphy. 2019. Olfactory and other sensory impairments in Alzheimer disease. *Nature Reviews Neurology*, 15(1):11–24.
- Carita Paradis. 2003. Is the notion of linguistic competence relevant in cognitive linguistics. *Annual Review of Cognitive Linguistics*, pages 247–271.
- Jamie Reilly, Maurice Flurie, and Jonathan E Peelle. 2020. The English lexicon mirrors functional brain activation for a sensory hierarchy dominated by vision and audition: Point-counterpoint. *Journal of Neurolinguistics*, 55:100895.

Jacques Savoy. 2020. *Machine Learning Methods for Stylometry*. Springer.

Koustuv Sinha, Prasanna Parthasarathi, Joelle Pineau, and Adina Williams. 2020. Unnatural language inference. *arXiv preprint arXiv:2101.00010*.

Richard Sorabji. 1971. Aristotle on demarcating the five senses. *The Philosophical Review*, 80(1):55–79.

Laura J Speed and Asifa Majid. 2020. Grounding language in the neglected senses of touch, taste, and smell. *Cognitive neuropsychology*, 37(5-6):363–392.

Yla R Tausczik and James W Pennebaker. 2010. The psychological meaning of words: Liwc and computerized text analysis methods. *Journal of language and social psychology*, 29(1):24–54.

Joel Whitburn. 2010. *The Billboard Book of Top 40 Hits*. Billboard Book of Top 40 Hits.

Bodo Winter. 2016. Taste and smell words form an affectively loaded and emotionally flexible part of the english lexicon. *Language, Cognition and Neuroscience*, 31(8):975–988.

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

Bodo Winter, Marcus Perlman, and Asifa Majid. Vision dominates in perceptual language: English sensory vocabulary is geared towards usage.

Bodo Winter, Marcus Perlman, and Asifa Majid. 2018. Vision dominates in perceptual language: English sensory vocabulary is optimized for usage. *Cognition*, 179:213–220.

## A Distribution of Modalities

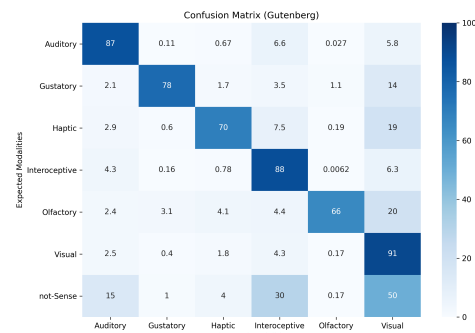


Figure 5: Distribution of expected-observed modalities in the Novels Dataset.

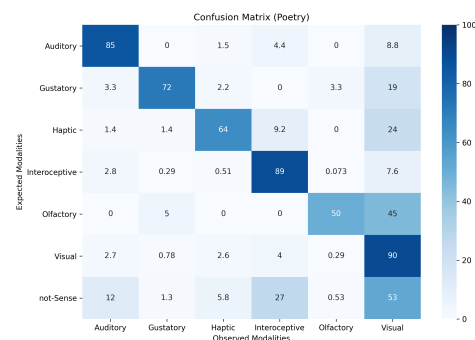


Figure 6: Distribution of expected-observed modalities in the Poetry Dataset.