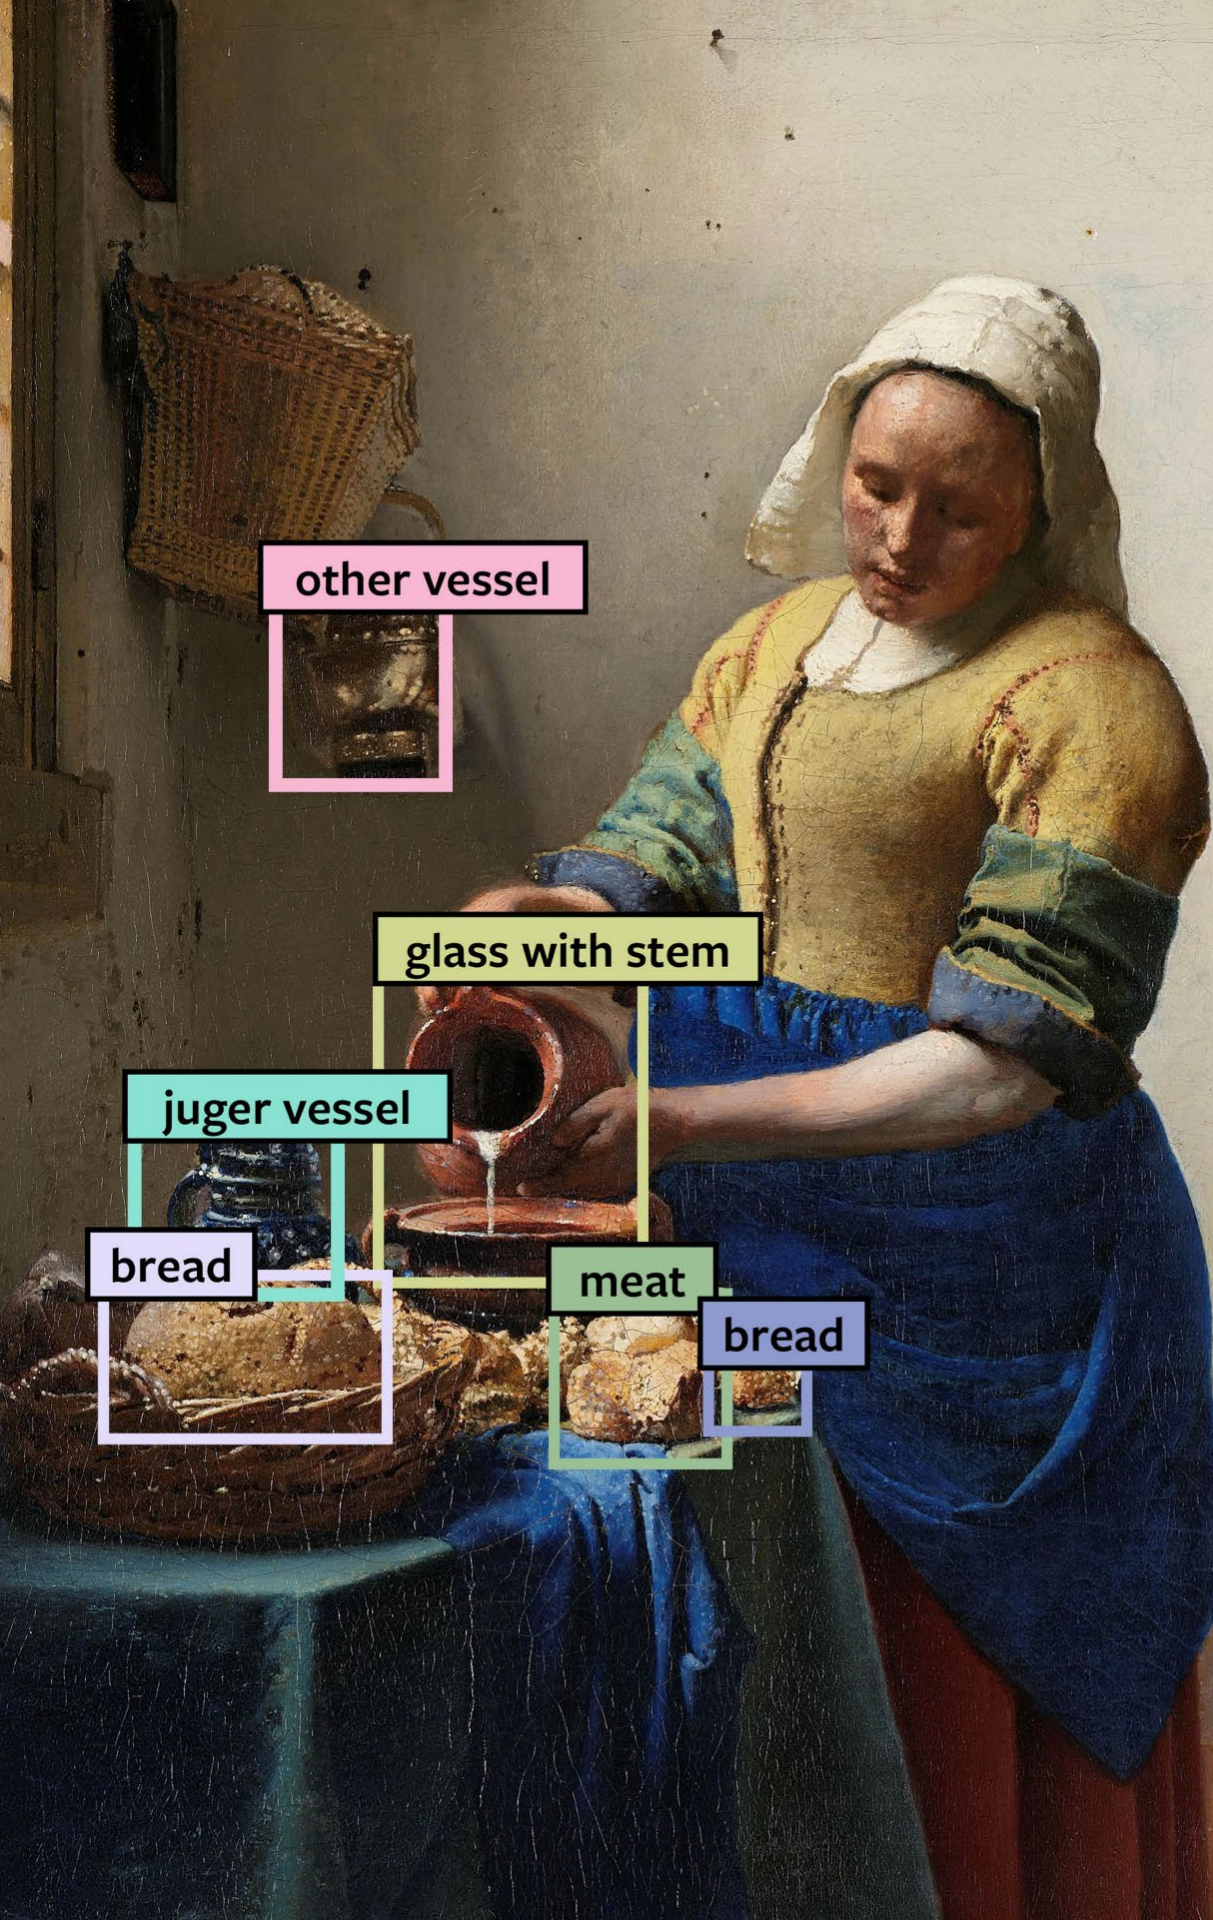


Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

Marieke van Erp, William Tullett, Vincent Christlein, Thibault Ehrhart, Ali Hürriyetoğlu, Inger Leemans, Pasquale Lisena, Stefano Menini, Daniel Schwabe, Sara Tonelli, Raphaël Troncy, and Mathias Zinnen

More than the Name of the Rose

How to Make Computers Read, See, and Organize Smells

We are all familiar with the famous Shakespearian statement that a rose by any other name would smell as sweet. Yet throughout history the scent of the rose has absorbed, as if by a kind of magnetic attraction, a multitude of meanings. It has often been the flower of power. Contained in expensive casting bottles, the scent of rose oil (commonly known as rose attar) was distributed to the mistresses of Henry VIII. The potent scent evoked an intimacy with the body of the monarch and was a representation of Tudor rule. In the eighteenth and nineteenth centuries, European visitors to the Ottoman Empire noted the use of rose-water, sprinkled on the hands, to welcome guests in an expression of hospitality. The corpses of medieval men and women who emitted the odor of sanctity instead of pungent putrefaction were often said to smell of roses. The scent of the rose has also been central to Islam: the first rose supposedly bloomed from the tear of the Prophet, and the indispensability of rose water in Islamic culture has left its impact in material culture in the form of the *gulab*, a rounded bottle with an elongated spout for pouring the perfume. The rose has never merely smelled sweet; it has evoked a whole range of meanings.[1]

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

The same might be said of the rose's material history. The variety of roses (gallica, moschata, damascena), the different parts of the plant (petals, thorns, leaves), and the way it has been processed (raw flowers, distillations, enfleurage, and more) all affect the type of scent one encounters. In fact, in nineteenth-century Europe, with a new fashion for the beautiful colors and shapes of Chinese tea roses, the fragrance of popular rose varieties changed markedly: from a heavy, rich, intense, floral scent to a hay-like, green, faintly earthy and woody odor.[2] Scholars have suggested that this change marks a broader shift in nineteenth-century culture. As sight eclipsed olfaction in the West's hierarchy of the senses, the smell of the rose became less important than its visual beauty.[3] In early modern medicine, the cooling properties of rose, according to the Galenic system, meant that it was frequently employed in medical recipes in the form of buds, waters, oils, dried petals, and powder to treat a variety of illnesses.[4] Rose continues to be valued by the cosmetics industry today. In Iran, where rose water has a long history stretching back to at least the seventh century, rose is an integral flavor in food and an ingredient in face creams (to name just two of its many uses). The scent of the rose is thus not singular but plural: it can be linked to a whole series of objects, practices, and meanings.

But how do we locate these objects, practices, and meanings over time and cultural realms? Here, the rich digitized text and image collections currently available to scholars offer many opportunities and challenges. The combination of vast digitized archives and the development of advanced digital tools has also opened up new avenues of digital or computationally driven research in the field of history. Digital humanities projects such as Living with Machines now combine data science and machine learning on large amounts of digitized data, for example census records, maps, and newspapers. The recent foundation of the *Journal of Digital History* also marks a shift in the methodological orientation of history scholars.[5]

The Odeuropa project can bank on this expanding field of interdisciplinary collaborations in the digital humanities. The aim of the project is to develop a "computer nose": design AI strategies to capture references to smells in the past from digital heritage collections. In this article, we will provide a step-by-step explanation on how Odeuropa employs language and computer vision technology to extract smell-related information from texts and images, and how these extracted bits of information are then encoded in a so-called knowledge graph, a graph-like database built on Semantic Web technologies (the European Olfactory Knowledge Graph, or EOKG). From here, we can extract storylines and follow smell sources (such as roses), noses, and olfactory practices over time and space.

The potential that such an approach holds for sensory history is huge. One criticism often leveled at history—especially cultural history—is that it tends to get fascinated with exciting discursive flotsam and

1 Holly Dugan, *The Ephemeral History of Perfume: Scent and Sense in Early Modern England* (Baltimore, MD, 2011), 42–69; Beyza Uzun and Nina Macaraig, "Scenting the Imperial Residence: Objects from the Topkapı Palace Museum Collections," *The Senses and Society* 17, no. 1 (2022): 68–89; Katelynn Robinson, "The Anchoress and the Heart's Nose: The Importance of Smell to Medieval Women Religious," *Magistra* 19, no. 2 (2013): 41–64.

2 Harold McGee, *Nose Dive: A Field Guide to the World's Smells* (London, 2020), 238–41.

3 Constance Classen, *Worlds of Sense: Exploring the Senses in History and across Cultures* (London, 1993).

4 Rebecca Laroche, "Roses in Winter: Recipe Ecologies and Shakespeare's Sonnets," in *Ecological Approaches to Early Modern English Texts*, ed. Jennifer Munroe and Edward J. Geisweidt (London, 2016), 51–61.

5 David M. Berry, "The Computational Turn: Thinking about the Digital Humanities," *Culture Machine* 12 (2011): 1–12. Living with Machines investigates how technology impacted labour, social, and cultural transformations during the first Industrial Revolution: https://livingwithmachines.ac.uk/

Frontis: Detail from Johannes Vermeer, Het melkmeisje (The Milkmaid), ca. 1660, oil on canvas, Rijksmuseum Rijksstudio.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

jetsam without truly exploring the resonance of the attitudes that have been identified.[6] Finding a particular idea, thought, or feeling in the past is not in and of itself particularly interesting or difficult. What is trickier is attempting to track how widespread a particular attitude was or how representative it might have been of a community or segment of society.[7] Here, computer methods can be extremely helpful.

Let us give an example: At an early stage in the Odeuropa project, we used n-grams to look for groupings of words that were frequently connected to a set of olfactory "seed" terms (e.g., smell, odor, stink, stench). An n-gram is a sequence of words of a given number: a bigram would be "smell this," a trigram "please smell this," etc. This was quite revealing. Performing the analysis on the English Google Books corpus, it became clear that many of the most culturally resonant references to smell were not discussions of contemporaneous odors but invocations of past scents.[8] For example, eighteenth- and nineteenth-century books regularly refer to the scent of Lebanon and a place called Bucklersbury. It turned out that this was because writers frequently quoted the biblical Song of Songs (which evokes the smell of Lebanon) and because discussions of perfumery and medicine often talked about the smell of a street in London called Bucklersbury, which in the sixteenth and seventeenth centuries had played host to a potent mix of scents emanating from druggists, apothecaries, and grocers. What the n-gram analysis revealed was that Europe was already developing a keen sense of its own olfactory heritage and history before sensory historians started their own excavations of the sensate past.

Working with big data is also a cautionary tale. Computationally driven or computationally supported historical research requires complementary skills to "traditional" historical research. In the Odeuropa project, we take an approach that relies on an interdisciplinary team of experts in both computational methods and (art) history who closely work together to join deep domain expertise in their respective fields to tailor complementary research methods to a unified approach. The benefits for both fields are not to be underestimated. For historians, it is an opportunity to work with and adapt cutting-edge technology to their research questions. For computer scientists, it is an opportunity to field test their approaches in the challenging environment that humanities data brings.[9] For instance: Digital analyses can be based on too many data points to check manually, and the workings of tools can be opaque. Making techniques more transparent, explainable, and reproducible and combating bias in both data and tools is a hot topic in the artificial intelligence community, where algorithms have been likened to "black boxes."[10]

Collaborations between humanists and computer scientists can help overcome these challenges. However, how to productively work together, so that knowledge is gained and tools improved, is a challenge in itself.

In Odeuropa, the groups worked together on so-called "competency questions" to scope and steer the computer science process.[11] In

6 Peter Mandler, "The Problem with Cultural History," *Cultural History* 1, no. 1 (2004): 94–117.

7 Dror Wahrman, "Change and the Corporeal in Seventeenth- and Eighteenth-Century Gender History: Or, Can Cultural History Be Rigorous?," *Gender and History* 20, no. 3 (2008): 584–602.

8 J. B. Michel, Yuan Kui Shen, Aviva P. Aiden, Adrian Veres, Matthew K. Gray, The Google Books Team, Joseph P. Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, Steven Pinker, Martin A. Nowak, and Erez Lieberman Aiden, "Quantitative Analysis of Culture Using Millions of Digitized Books," *Science* 331, no. 6014 (2011): 176–82.

9 We are not the first to see the benefits of this. See for instance: Marieke van Erp et al., "The Future of DH Research: Alone You May Go Faster, but Together You'll Get Further," *Debates in the Digital Humanities* (forthcoming).

10 Max Kemman, Martijn Kleppe, and Stef Scagliola, "Just Google It: Digital Research Practices of Humanities Scholars," *Proceedings of the Digital Humanities Conference 2012* (2014). See also the recent call for papers on reproducibility and explainability in digital humanities in the *International Journal of Digital Humanities* (2023), https://www.springer.com/journal/42803/updates/23255012.

11 M. Grüninger and M. S. Fox, "The Role of Competency Questions in Enterprise Engineering," in *Benchmarking: Theory and Practice*, ed. Asbjørn Rolstadås (Boston, 2011), 22–31.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

Table 1. Example of Odeuropa competency questions by theme.

Themes	Competency Questions
Smells Core properties of smell sources and carriers	*What are the most frequent smell sources in Paris in the 18th century?* *Which smells were perceived during spring?*
Noses and Gestures Involving the actors perceiving the odors	*Which professions are more present in smelling experience descriptions?* *Which smelling gestures are described frequently by tea-merchants?*
Identities Meanings of smells and their capability of being representative of individuals and groups	*Which smells did people associate with femininity in England?* *What are the odors most associated with Ashkenazi Jewish practices?*
Emotions Interactions between olfaction and feelings	*What odors disgusted upper-class Europeans most?* *Which smell triggers memories of childhood?*
Practices Smell-producing practices; olfactory practices significant to certain groups	*What types of cooking produce a bad smell?* *What smells are connected with religious practices?*
Sites and Contexts The presence of odors in particular places and moments in time	*Which smells are associated with ships?* *Which smells could be perceived during the Crimean War?*
Texts and Images How smells are represented in texts and images	*What are the adjectives used for orange aroma in the 15th century?* *Which smells can be found in paintings in the Rijksmuseum's collections?*

software engineering, competency questions are a core method used to promote user-centered design. In Odeuropa, the competency questions were centered around different olfaction-related core themes (Table 1).

In designing the computational tools, care was taken that they would contribute to answering the competency questions. This meant taking concrete, empirical questions and trying to figure out what the key factors would be; these factors would end up as nodes in the knowledge graph. For example, the first two questions in our list might be reframed as:

- What are the most frequent smell sources in [place: city; neighborhood; street] in the [time: century; decade; year]?
- Which smells were perceived during [time; season]?

What do we want to know about the scent of the rose? We can formulate the following competency questions:

- What adjectives are used to describe the smell of roses in historical texts?
- Do these qualifications change over time?
- Are the qualification words different in different languages?
- In what particular places were rose scents likely to be found, in Europe, between 1600 and the 1920s?

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

- What smells might have co-occurred with the scent of roses?
- What was the smell of roses used for? Was it deployed in specific ways to specific ends?

With our competency questions arranged, we can move on to some of the digital humanities tools we use, starting with text.

Reading for Smell Experiences

Natural language processing (NLP) is a discipline that allows for the processing of textual resources, performing tasks such as the extraction or classification of linguistic information. Taking advantage of NLP allows us to process large collections of documents, which are now digitized, to gain new insights into their content at scale. In Odeuropa we focus on what we call "olfactory events": explicit mentions of smells. So we do not capture all roses (e.g., a rose mentioned in a word list or recipe), but only the roses that appear around explicit smell words (e.g. *smell*, *sniff*, *stink*, *perfume*). To capture these events, we extract text snippets around smell words and smell sources, and we identify the qualities used to describe them. We also try to analyze when and where the smell event took place. This, along with metadata of the documents such as the genre, the author, and the date and place of publication, can provide an extended representation of how smells and their perception changed over the course of time. By further extending the extraction to multiple languages, we can potentially study which and how smells are perceived in different countries.

The approach used to extract olfactory information is supervised machine learning. In this approach, the system learns how olfactory events are described in texts by starting from example sentences that were manually labeled by human annotators in a so-called *training step*. During training, the supervised algorithm creates a model able to infer which parts of a text are related to smell. In the examples provided to the algorithm, each word is associated with a label indicating (in our task) whether it is a smell word, a smell source, a quality, or none of the above. The algorithm analyzes the training data (representing words as vectors) to define a function to generalize the examples used for training in order to assign labels to new, unstudied texts. When applied to those texts, the model can indeed predict whether olfactory events are present based on what was learned during training.

Let us retrace the steps in this approach. As a first step toward a system for olfactory information extraction, in Odeuropa we created a benchmark (a set of exemplary texts used as a reference) to be used as training data covering seven languages (English, Italian, French, German, Dutch, Latin, and Slovene). For each language we selected around one hundred texts to be manually annotated by native speakers.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

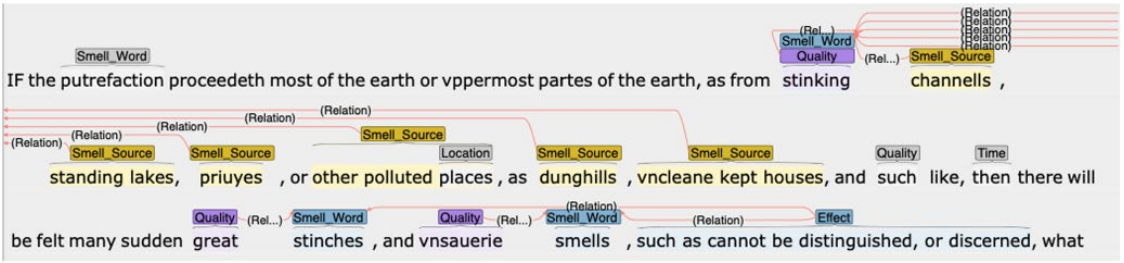


Figure 1. An annotated excerpt from James Manning, *Complexions Castle* (London, 1604) in the INCEpTION tool.

The texts were selected by taking into consideration publication date and genre. Thus we could guarantee representation in the data of different time periods and domains of interest, such as medicine, religion, travel, theater, science, and law.

To carry out annotation of olfactory information in texts, we rely on INCEpTION,[12] a web-based text-annotation environment. INCEpTION provides ample functionalities allowing continuous quality control. It, for instance, supports multiple languages and multiple simultaneous annotators. This is important, given the complexity of the task. In INCEpTION, annotators were asked to mark in the texts the presence and the extent of three elements:

Smell words: Nouns, verbs, adverbs, and adjectives evoking an olfactory event (e.g., *stink*, *scent*, *smell*, *odor*, *reek*, *fetor*, *fragrance*, *sniff*, *stinking*, *aromatic*, etc.).

Smell source: The person, object, or place that has a specific smell. It can also refer to any human or nonhuman source or object that produces an odor (e.g., plant, animal, perfume, human). The smell source, indicated [*between brackets*], is the entity or phenomenon that the perceiver experiences through his or her senses:

He carefully smelt [*the rosewater*] for any trace of poison.
The place was strongly scented [*with attar of roses*].

Quality: The quality associated with a smell and used to describe it. This is typically expressed by qualitative adjectives (e.g., *rancid*, *fresh*, etc.), and often preceded by an intensifier such as "very" or "really." The intensifier has to be annotated with the related adjective as part of the same span. Qualities include *intensity* (not perceptible, weak, distinct, strong), *volume / reach* (far-reaching), *duration* (lasting, permanent), *state* (old, deteriorated), *character* (humid, dry, garlicky, fruity, woody), and *hedonic characteristics* (malodorous, aromatic, healthy):

The rose had an [*earthy*] smell.
The waiter smelled something [*foul*] in the kitchen.
It was filled with an odor that was [*hard to conceive*].

12 INCEpTION can be accessed at: https://inception-project.github.io/.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

Table 2. Overview of annotated instances in the benchmark for each language.

	Dutch	English	French	German	Latin	Italian	Slovenian
Smell Word	1,788	1,530	845	2,659	1,228	1,254	1,973
Smell Source	1,922	1,313	710	2,297	772	952	1,638
Quality	1,071	1,084	450	1,730	552	707	936

The annotation effort involved different teams of annotators for each language. The final number of instances of smell words, smell sources, and qualities present in the benchmark is shown in Table 2.[13] Each benchmark has been used to obtain a model for olfactory information extraction in the specific language.

After this first phase, each model was applied to different large historical corpora available in digital format. For this article, we will focus on the results of the English-language analysis. For English, the model trained on the data created for the English benchmark has been used to extract olfactory information from four document collections so far (2022):

1. **EEBO**: works represented in the Early English Books Online collections known as EEBO TCP Phase I and II.
2. The **Royal Society Corpus**: scientific periodicals issued between 1665 and 1869.
3. A preprocessed version of the larger **Old Bailey Corpus** dataset, containing court proceedings published between 1720 and 1913.
4. **Project Gutenberg**: a digital library of over sixty thousand books.[14]

Although systems for information extraction are far from perfect, they are able to automatically extract data from massive amounts of documents, allowing *distant reading*: a view on document content that is driven by quantitative analysis and that tries to put in connection different pieces of information coming from different sources.[15] While an in-depth analysis of sources is possible only when humans read single documents, an information extraction system can provide an alternative and complementary view on data, and a small fraction of classification mistakes is generally considered acceptable in this respect. In the next Odeuropa phase, more sources will be processed in order to both increase the number of English texts and extend the extraction to all the seven languages for which we have training data available.

By manually checking the output of the information extraction system, some common types of mistakes can be identified. The main difficulty in extracting olfactory elements from texts stems from the specific domain we are targeting. Terms used to describe olfactory information are often ambiguous and relate to other senses as well (e.g., sweet, fresh, unpleasant). Linguists state that smell is the sense showing the least specialized vocabulary in European languages.[16] Thus, a system for olfactory information extraction should be able to disambiguate terms that are being used to refer to smells and not to other, more dominant senses. Furthermore, almost any object can be a smell

13 The benchmark can be downloaded at: https://github.com/Odeuropa/benchmarks and corpora.

14 Early English Books Online, University of Michigan Library, https://quod.lib.umich.edu/e/eebogroup/; Royal Society Corpus (RSC) 4.0, http://fedora.clarin-d.uni-saarland.de/rsc_v4/; see also Hannah Kermes, Stefania Degaetano-Ortlieb, Ashraf Khamis, Jörg Knappen, and Elke Teich, "The Royal Society Corpus: From Uncharted Data to Corpus," *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC '16)* (2016): 1928–31. LREC 2016 was held in Portorož, Slovenia; the Old Bailey Corpus can be found at https://fedora.clarin-d.uni-saarland.de/oldbailey/; the Project Gutenberg library can be found at https://www.gutenberg.org/.

15 Ted Underwood, "A Genealogy of Distant Reading," *DHQ: Digital Humanities Quarterly* 11, no. 2 (2017).

16 Bodo Winter, *Sensory Linguistics: Language, Perception and Metaphor* (Amsterdam, 2019).

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

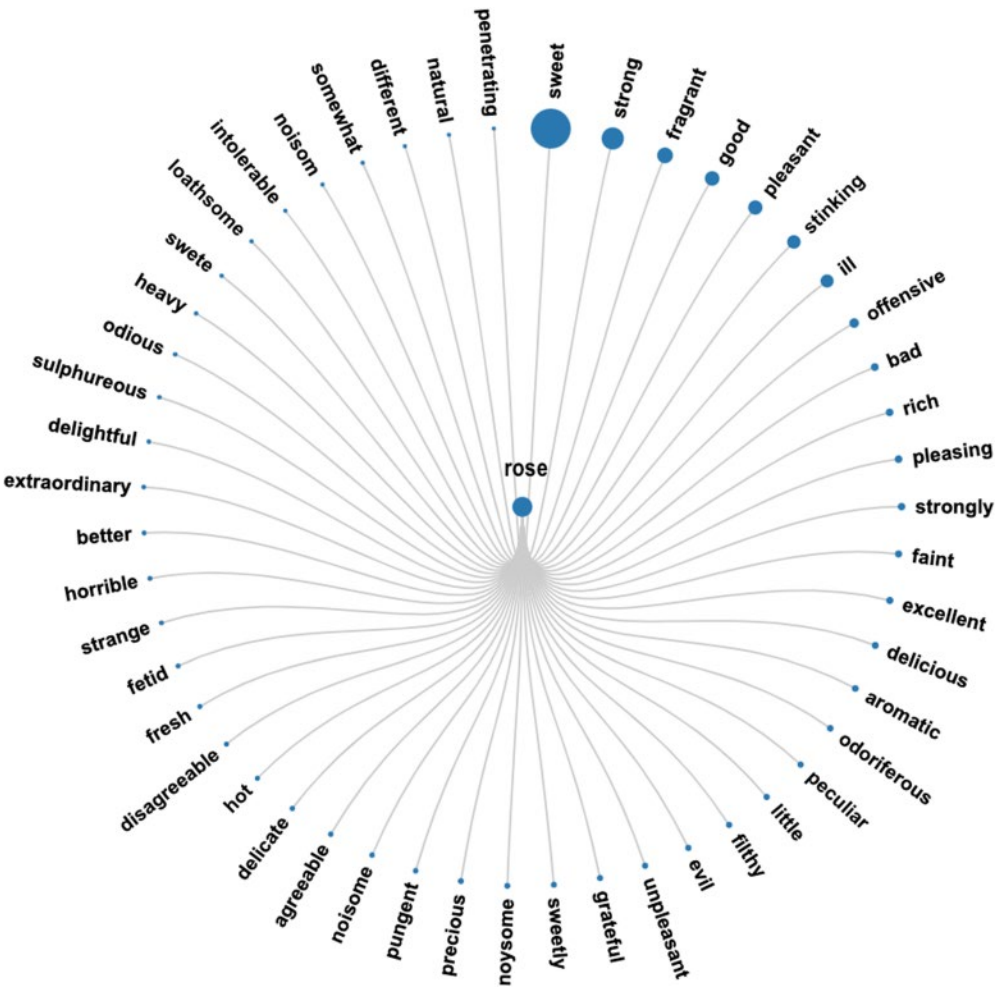


Figure 2. Qualities (adjectives) most frequently associated with rose, based on data in the EOKG as of August 2022 (EEBO, Old Bailey Papers, Royal Society Corpus, Gutenberg).

source. For example, "air," "new shoes," or "rose" can appear in a variety of situations where most of them are unrelated to the smell. Disambiguating the presence of these elements in olfactory situations based on the context is still a challenging task. To minimize the issue, we designed the algorithm to be aware of contextual information in order to discern whether a generic word is used as part of an olfactory event.

Figure 2 shows one way that we can visualize the results of this approach. Here we have the most common terms used to describe the smell of the term "rose," arranged in a circle in order of the frequency with which they appear, from most regular (*sweet*) to least regular (*penetrating*). We do indeed find that a rose has often smelled as sweet as William Shakespeare once suggested, but the results also demonstrate the broader diversity of the vocabulary used in describing rose odors. The smell of rose can be of different strengths (*strong*, *faint*, *little*, *delicate*) and qualities (*hot*, *fresh*, *foetid*), it can produce quite different

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

reactions from noses (*delicious*, *pleasing*, *stinking*, *offensive*), and it can evoke different valuations (*excellent*, *peculiar*, *precious*, *odoriferous*). So there is more in the name, but what is in the image? How can we *see smells*?

The Odeuropa project also focuses on extracting emotions from text references. To do this, it employs the same approach for text processing as already described, simply adding a specific emotions vocabulary going through the same annotation process and subsequent development of a system for automated recognition of references to emotions in texts.

Seeing Smells

For research on the visual representation of olfaction in digital image collections, we can turn to the field of artificial intelligence known as *computer vision*: the application of machine-learning techniques to the analysis of visual data.[17] Common tasks of computer vision algorithms are classification, object detection, and segmentation. These tasks differ in the scope they take into account. In *classification*, the algorithms look at the whole image and assign it to one of a number of previously defined categories based on what is being imaged. *Object detection* works on smaller regions of the image and entails the localization of predefined objects and their classification. A special case of object recognition is *pose estimation*, in which the posture of the persons depicted is determined on the basis of a number of so-called *keypoints*, such as elbow, wrist, etc. One example is the work of Peter Bell and Leonardo Impett, who apply pose estimation to analyze large amounts of Annunciation and Baptism scenes in historical paintings, finding patterns in the body poses of the depicted characters.[18] Another example considering a broader context would be the analysis of image composition based on body posture, gaze estimation, and foreground-background separation, which has been conducted by Madhu et al.[19] *Segmentation*, on the other hand, considers each image pixel individually and assigns it to a category. These classical computer vision tasks can be applied to the fields of digital heritage and digital humanities to provide a quantified perspective that complements or offers new questions for more qualitative approaches.

In the Odeuropa project, we use a combination of image-level classification techniques, pose estimation, and object detection to find traces of past smells in artworks and other historical imagery, such as book illustrations or cartoons. This requires the identification of ways in which odors can be visually represented or referenced. Qualitative approaches, to a certain extent, can rely on a commonsense understanding of smell references to enable the recognition of olfactory dimensions of an artwork. However, if we want to automate parts of the analysis to

17 An overview of the field: Richard Szeleski, *Computer Vision: Algorithms and Applications* (Redmond, WA, 2010).

18 Peter Bell and Leonardo Impett, "Ikonographie und Interaktion: Computergestützte Analyse von Posen in Bildern der Heilsgeschichte," *Das Mittelalter* 24, no. 1 (2019): 31–53.

19 Prathmesh Madhu, Tilman Marquart, Ronak Kosti, Peter Bell, Andreas Maier, and Vincent Christlein, "Understanding Compositional Structures in Art Historical Images Using Pose and Gaze Priors," *Proceedings of the European Conference on Computer Vision* (2020): 109–25.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

take larger amounts of visual data into perspective, we need a more formal definition of these references.

Complicating matters is that odors can hardly be directly referenced in a visual medium, unlike in texts.[20] Apart from specific exceptions such as smoke, the substrate of smell usually remains invisible. Instead, we have to rely on visual proxies that hint toward the presence of smells in an artwork. Noses busily sniffing, objects held to the nose, allegorical stand-ins for smelling, and the much broader list of visual elements in any one period or genre that, thanks to our historical knowledge, we know to have particular olfactory importance. The most straightforward way of finding such proxies is by localizing objects that either emit smells themselves, such as pipes, incense burners, or pomanders, or indirectly indicate the presence of smells, such as sniffing dogs or flies buzzing around fruit.[21]

When applied to modern photographs, computer vision algorithms designed to detect objects in images exhibit an impressive performance.[22] However, applying these algorithms to historical imagery entails two major challenges:

> The artistic representation of objects can differ significantly, not only from their photographic representation, but also between different periods, styles, or artists.
>
> The classes of objects that are of interest for the Odeuropa project are usually not represented in modern object-detection datasets. Some of these objects might also have changed their appearance over the course of the centuries, which complicates their recognition.

We address these problems by designing detection algorithms that are robust toward variations in style and capable of detecting a large variety of objects. Therefore, we are experimenting with multiple techniques that enable the transfer of knowledge from modern large-scale photographic datasets to the domain of historical artworks, where there is much less training data available.[23] Additionally, we continuously annotate artworks from multiple digital museum collections to create more training data for fine-tuning computer vision algorithms on artworks. The resulting Odeuropa Dataset of Smell-Related Objects (ODOR) is freely accessible on Zenodo for others to download and use.[24]

How well this adaptation to the domain of historical artworks works depends heavily on the kind of objects we are trying to localize. For example, pomanders (scented jewelry worn around the waist or the neck) or censers (used to distribute the scent of burning resins, often in religious contexts) can be detected quite reliably. Such objects might be interesting for our story about roses. Recipes for pomanders often included rose petals or rose water and so, in identifying pomander containers in images, we may be identifying objects that held rose-like scents.[25]

Other types of smell-relevant phenomena can be much more difficult to automatically recognize. Smoke, for example, has a high olfactory significance since it directly conveys the presence of vapors in the air that might

20 Elizabeth S. Marx, *Visualising, Perceiving, and Interpreting Smell in Seventeenth-Century Dutch Art* (PhD thesis, Cambridge, 2022), 18–32.

21 Due to their strong sense of smell, dogs were often depicted as metaphors for the sense of smell in general. Flies can be considered as a hint toward decay and decomposition, hence their common co-occurrence in still-life paintings with rotten fruit or withering flowers; see the entries in James Hall, *Dictionary of Subjects and Symbols in Art* (London, 1974).

22 Li Liu, Wanli Ouyang, Xiaogang Wang, Paul Fieguth, Jie Chen, Xinwang Liu, and Matti Pietikäinen, "Deep Learning for Generic Object Detection: A Survey," *International Journal of Computer Vision* 128, no. 2 (2020): 261–318.

23 Mathias Zinnen, Prathmesh Madhu, Peter Bell, Andreas Maier, and Vincent Christlein, "Transfer Learning for Olfactory Object Detection" (paper presentation, Digital Humanities 2022, Tokyo, Japan, July 25–29, 2022).

24 Mathias Zinnen, Prathmesh Madhu, Ronak Kosti, Peter Bell, Andreas Maier, and Vincent Christlein, Odeuropa Dataset of Smell-Related Objects (2.0.0) [dataset], Zenodo (2022): https://doi.org/10.5281/zenodo.7125961

25 *A Closet for Ladies and Gentlewomen* (London, 1608), 57; William Bullein, *Bulleins Bulwarke of Defence against all Sicknesse, Soarenesse, and Woundes that doe Dayly Assaulte Mankinde* (London, 1579), 18.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024



Figure 3. Pieter Lastman, *David in the Temple*, 1618, oil on panel, Herzog Anton Ulrich-Museum Braunschweig, GG 208. RKDImages (214918).

have a scent. However, smoke's lack of clear boundaries and its tendency to blend into the image background makes it extremely difficult to recognize it in artworks and localize it correctly. This does not mean it cannot be done, as demonstrated by the images we have included in this article.

Another challenge with some objects is the particularity of their categorization. In the ODOR dataset, we distinguish between more than twenty different flower species based on their different odors. However, the visual differences between different flower species can be subtle, especially given the various degrees of artistic abstraction in the dataset. Correctly classifying a specific flower species thus becomes an extremely challenging task.

These challenges for categorization arose early on in our work, during the preparation of the training data. As we annotated images, we noticed that annotating the correct flower species is time-consuming and requires a good amount of botanical knowledge that our annotators did not have at that time. We thus annotated all flowers as generic flowers in the first iteration and postponed the assignment of the correct species to the second round of annotations. Due to the time-consuming nature of this task, we decided to crowdsource it and specified Amazon Mechanical Turk jobs, where crowdsourced workers were given batches of the annotated flower instances and asked to assign the correct flower species.[26] To ensure a sufficient annotation quality, we implemented

26 Amazon Mechanical Turk (website), accessed January 2023, https://www.mturk.com/. The Odeuropa image recognition team has chosen to use Mechanical Turk with knowledge of the ethical issues that the platform poses (such as poor labor conditions). This is one of many issues that digital humanities scholars should consider. Other issues include the environmental costs of using server-intensive methods and cloud-based data storage that involve the expenditure of large amounts of energy.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

multiple measures of quality control, including inter-annotator agreement and a final consistency check by our expert annotators.

Up to now, this process resulted in the manual annotation of around one thousand samples of roses, most of which could be found in rich flower still-life paintings, such as those of Ambrosious Bosschaert.

Even without doing predictions on larger image corpora, we can already use this data to trace the smell of roses. A natural way to leverage this data is by using it as a query tool to support qualitative analysis. Tools that require little to no programming experience, such as curated Jupyter notebooks or graphic dataset exploration tools, enable scholars to browse through the annotated roses in the database.[27] This, in turn, facilitates gaining an overview of contexts in which the flower occurs and might already result in identifying unexpected connections.

However, the query results in a corpus of manual annotations necessarily reflect the biases and preselections underlying the creation of the dataset. Although this limitation is not exclusive to computer-assisted data analysis, we are well aware of it and acknowledge the need to reveal the biases underlying the dataset creation process as explicitly as possible.

To some extent, this problem can be mitigated when going from manual annotations to automatic predictions. On the one hand, we can simply consider more data when we do not have to rely on manual annotations. On the other hand, since predictions can, in principle, be made on all kinds of image corpora, we can adapt the data domain to the topic under investigation (e.g., analyze only images of perfume advertisements when writing a history of roses and perfumery).[28]

This usage of manual annotations and computer-generated object predictions as a query tool can support classical approaches to art history or cultural history. In addition, providing data in a machine-readable format also enables new perspectives and approaches that would not be possible with classical approaches. Deriving statistical information from large amounts of visual data, scholars can take a step back from singular works of art and take a perspective of *distant viewing*, where, similar to *distant reading* for texts, larger trends and structures can be revealed that remain invisible to the detailed analysis of singular sources.[29]

How can we apply this approach to the smell of roses and their representation in the visual arts? One obvious question we could try answering using our dataset is how the number of depicted roses has changed throughout the specific time period investigated by the Odeuropa project. Of course, using these trends alone does not allow us to conclude anything about historical perspectives on the smell of roses. It might serve, however, as one piece of evidence about the changing importance of the smell over time. To this end, we can simply count the number of annotated roses in our dataset and relate them to the creation year (Figure 6). Looking at the chart, we notice that the number of depicted roses has a peak between 1600 and 1650 and then decreases until 1920, which marks the end of Odeuropa's period of interest.



Figure 4 Detail from David Teniers (II), Village in the Snow (1625-1690). Oil on panel. Private collection. RKDImages (290572).

27 Project Jupyter can be found at https://jupyter.org/.

28 Note that the artistic style of the analyzed image data can still be a large restriction since performance of computer vision algorithms can decrease severely when applied on data that is different from their training data.

29 Taylor Arnold and Lauren Tilton, "Distant Viewing: Analysing Large Visual Corpora," *Digital Scholarship in the Humanities* 34 (2019): i13–i16. Ted Underwood, "A Genealogy of Distant Reading," *Digital Humanities Quarterly* 11, no. 2 (2017): 1–12.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024



Figure 5. An annotated example of an image in the Odeuropa Smell Explorer. George Dunlop Leslie, Rosenzeit (1880). Hamburg Kunsthalle, HK-2393. Wikicommons.

Does this correlation indicate the decreasing importance of roses after the second half of the seventeenth century? We have to be careful and consider the dataset's biases. Comparing the trend to the overall temporal distribution of the current annotations in our dataset (Figure 6), it appears that it is more likely to be explained by our sampling of artworks to annotate than by an evolution of the cultural impact of roses.

On the contrary, when we jointly plot the distributions of rose annotations and of overall annotations (Figure 7), we notice that the proportion of roses increases over time—a trend that directly contradicts the naïve conclusion obtained from looking at the absolute numbers. This correlation becomes even more apparent when we plot the percentage of rose annotations among all annotations (Figure 8).

Plotting the rose annotations in relation to all annotations still does not allow us to draw conclusions about the significance of roses over time. Instead, it could just as well be explained by an increasing amount of still-life paintings focusing on vases of flowers in our dataset over the second half of the eighteenth century. We can try to correct for this possible bias by comparing the number of rose annotations to all annotations of flowers over time (Figure 10).

Doing that comparison reveals that the increasing share of rose annotation still holds when compared to all flower annotations. This gives us a strong indication of an increasing cultural significance of flowers in visual culture, which, in turn, might provide us information about the changing significance of the smell of roses.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

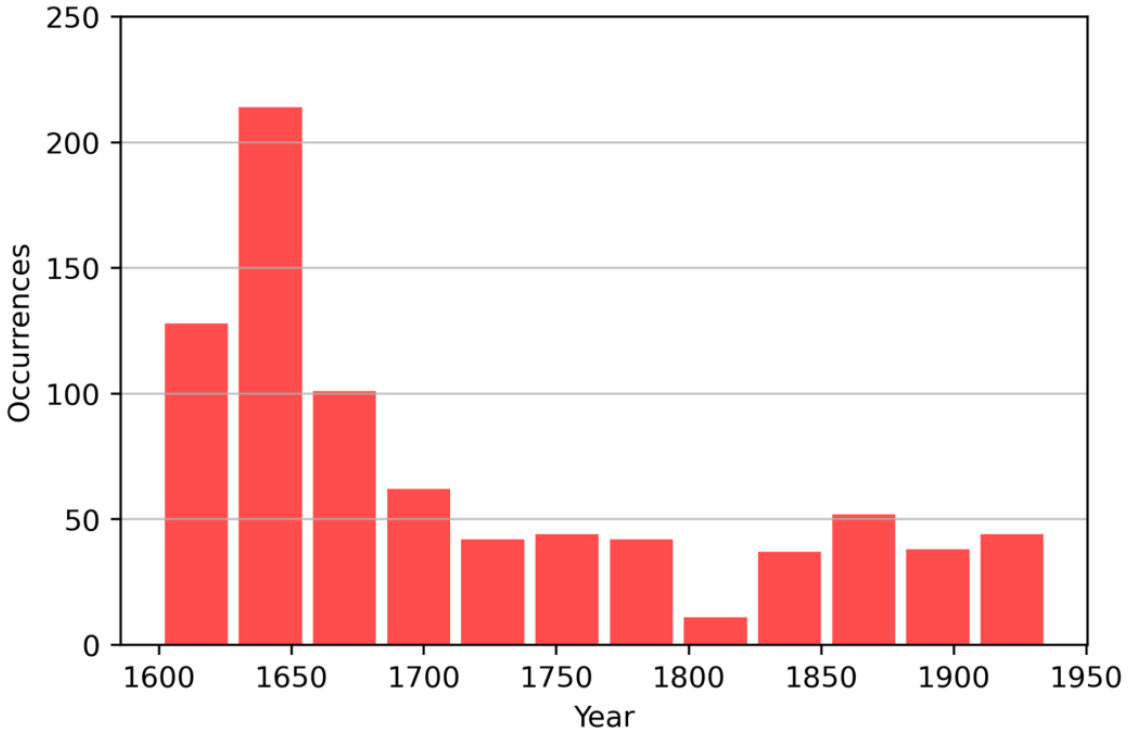


Figure 6. Distribution of roses in paintings over time in the ODOR dataset version 2.1 (2022).

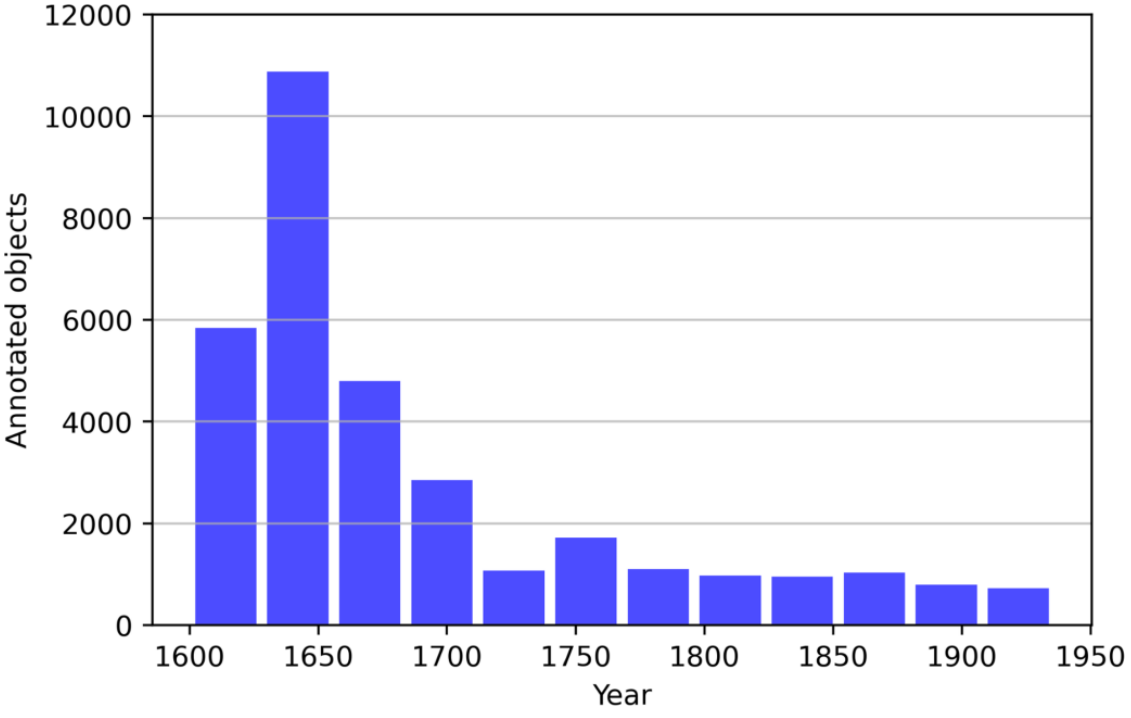


Figure 7. Distribution of all annotations in paintings over time in the ODOR dataset version 2.1 (2022).

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

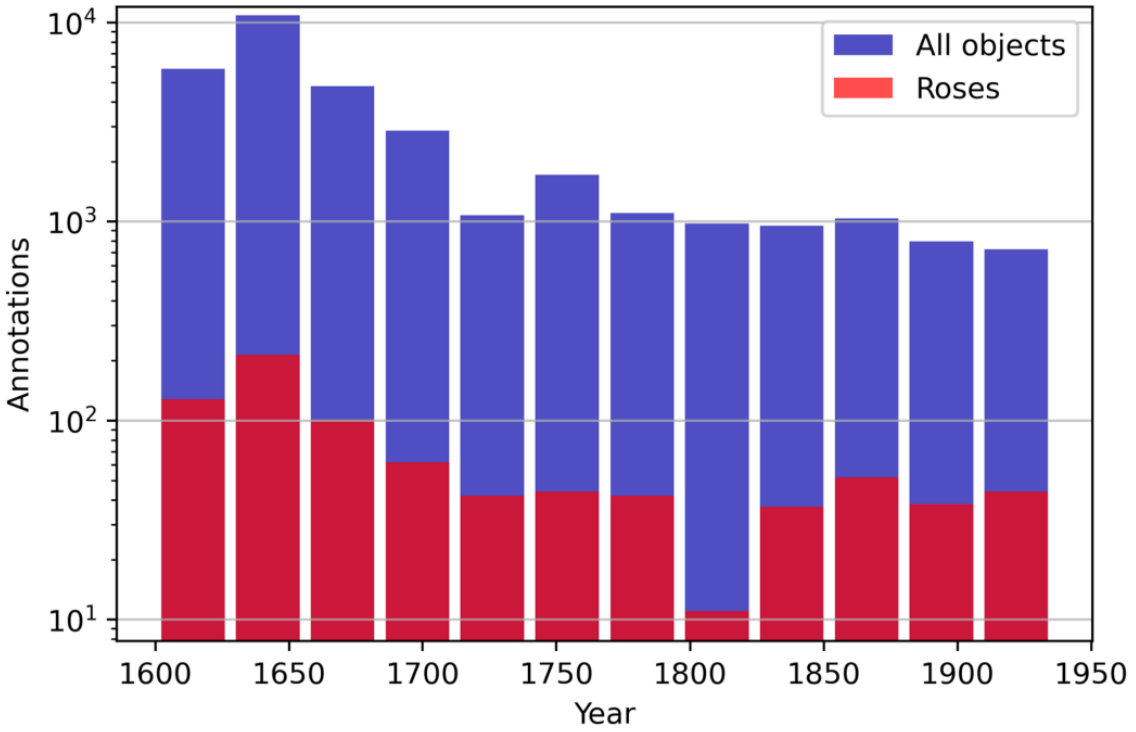


Figure 8. Joint distribution of roses (red) and all objects (blue) in paintings over time in the ODOR dataset version 2.1 (2022). The number of annotations is plotted in log scale to increase readability.

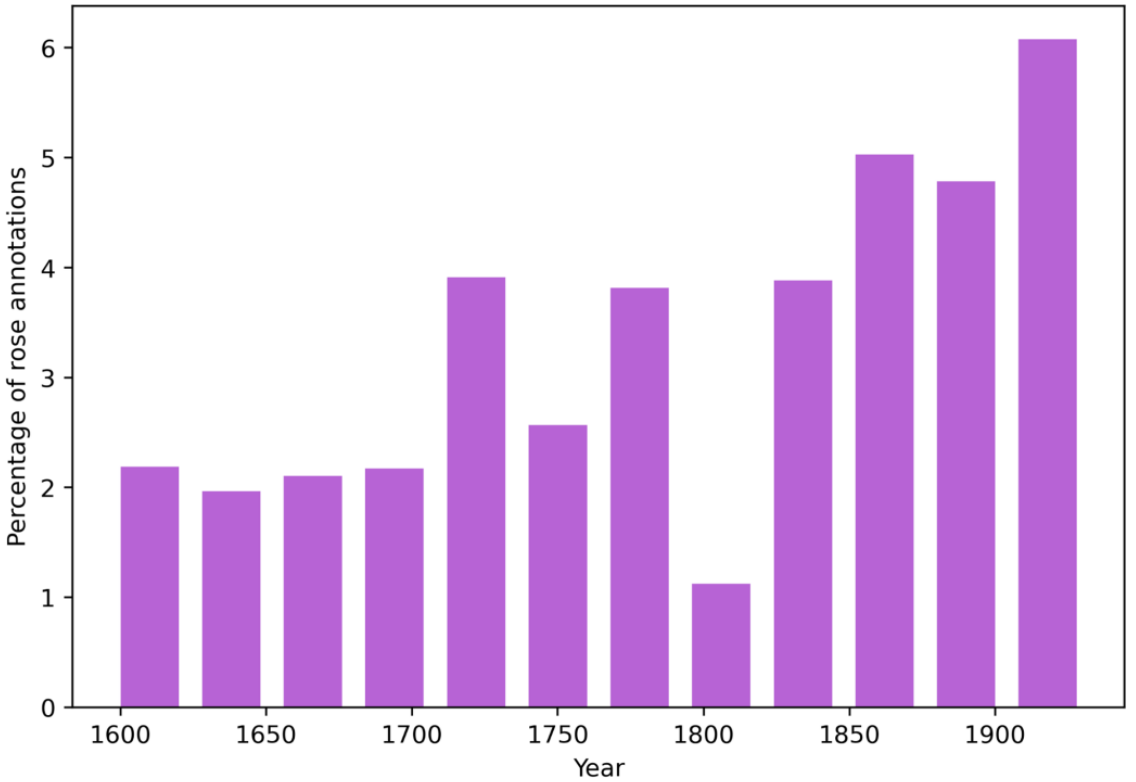


Figure 9. Percentage of rose annotation in relation to all objects in paintings over time in the ODOR dataset version 2.1 (2022).

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

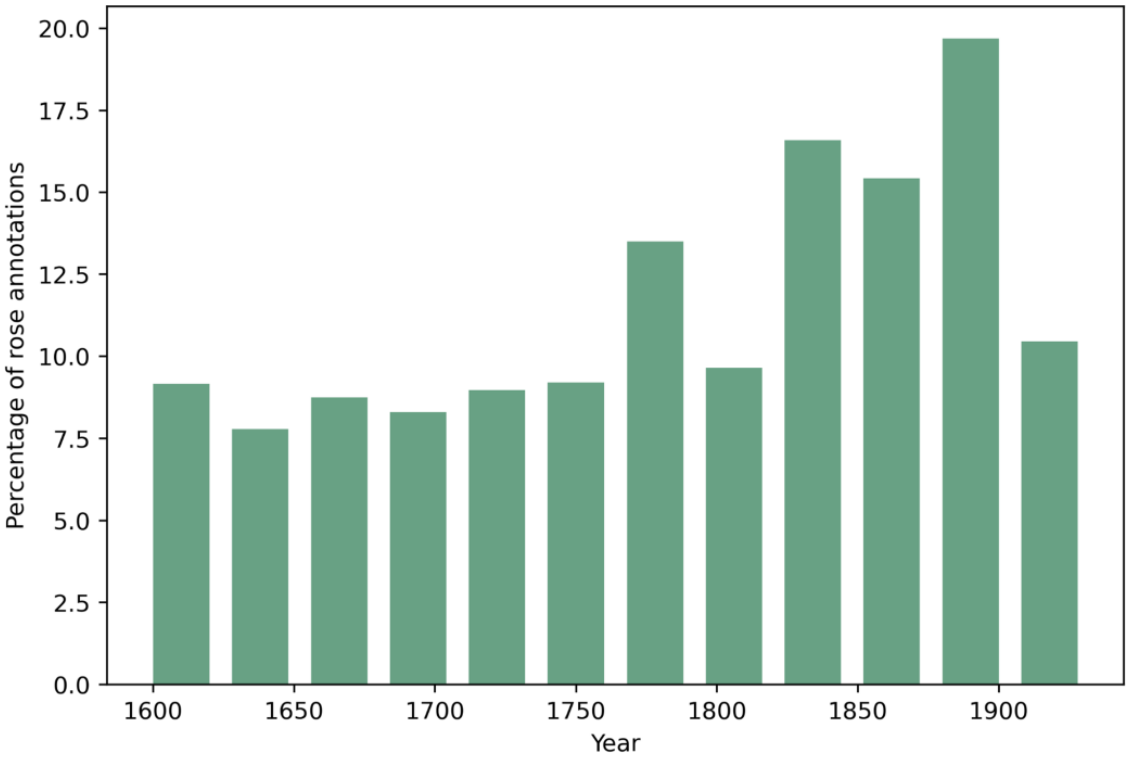


Figure 10. Percentage of rose annotations in relation to all flower annotations in artworks over time in the ODOR dataset version 2.1 (2022).

But what do we make of this visual information in relation to smell? Consider the argument we floated at the start of this piece, put forward originally by Constance Classen, that there was a shift away from the lower senses to vision over the course of the eighteenth and nineteenth centuries, which was in turn linked to a shift in how roses were valued and cultivated—from an interest in fragrance to a focus on visual beauty. The data here might support or contradict this argument. On the one hand, the increased proportion of roses in the dataset over time might suggest that roses were increasingly valued for their visual beauty and so were depicted more often. To test this hypothesis, we can shift our focus away from *olfactory objects* toward *smell gestures* and investigate whether we notice a development in the number of roses that have been depicted in conjunction with smelling gestures. In contrast to isolated depictions of roses, we can infer that the rose depictions have a connection to olfaction with higher certainty.

One of the gestures we have annotated is "sniffing." How can we automatically find out which objects are being sniffed in these instances?

We want to find out which flower is relevant for the sniffing gesture; we are only interested in the single flower that the depicted woman brings to her nose. The other flowers in the painting should not be taken into consideration. Of all objects annotated in the painting, we identify only those objects as relevant to the sniffing gesture that have

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

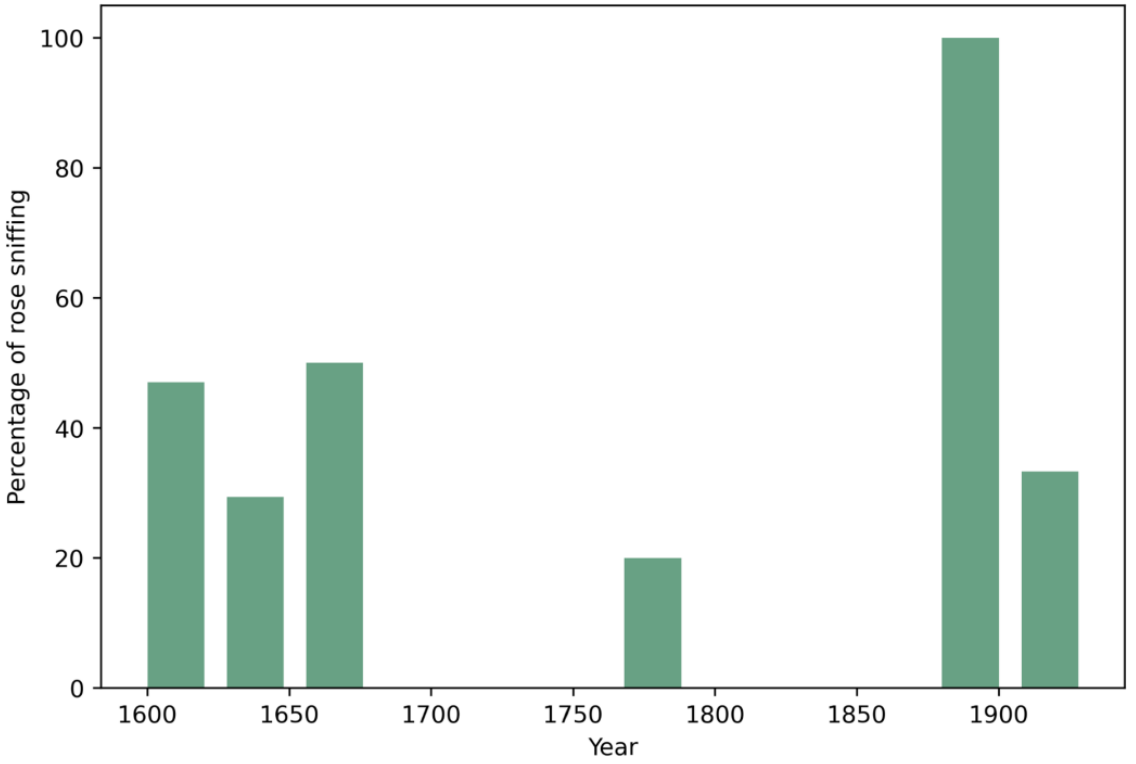


Figure 11. Percentage of rose annotations in relation to all objects involved in sniffing gestures over time in the ODOR dataset version 2.1 (2022).

an overlap with the bounding box defining the smell gesture. If there are no such objects, we only consider one object that is closest to the box and below a certain distance threshold.

Using this criterion for smell objects, we find that 34 percent of the objects involved in a sniffing gesture are roses. Figure 11 plots this percentage over time. Given the high value around 1900, this seems to contradict the theory of the deodorized roses. Note, however, that currently we can only base our analysis on forty-seven depictions of smelling gestures, which prevents us from making broader conclusions. The limited amount of data also explains the gaps in some time spans of the plot. We simply did not yet have any images in that period to evaluate. In the next phase of the project, we hope to fill this gap.

Although, at this moment, we are still lacking a sufficient number of gesture annotations, it should have become evident that information about smell gestures can be extremely valuable in the visual analysis of past smells. They can tell us a lot about historical conceptualizations of olfaction as well as reveal the emotions that were associated with smells. At the same time, they provide relatively compelling evidence that artworks were indeed perceived to have olfactory elements. Their connection to smell is more immediate: whereas, ultimately, every object could be considered a smell source, smell gestures can hardly be depicted without considering their connection to olfaction.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

Table 3. Number of the ten most common object co-occurrences with roses annotated in the ODOR dataset version 2.1 (2022).

Object	Number of Co-occurrences
Flower	1957
Carnation	970
Anemone	481
Jasmine	277
Tulip	262
Iris	236
Grapes	225
Heliotrope	224
Peach	202
Cornflower	196

In order to use them effectively for quantitative approaches to sensory history, however, we need to increase the amount of smell gestures in our dataset. One obvious way of doing this is to find more gesture depictions and annotate them manually. Indeed, we are continuously working on finding more annotations. However, as smell-related metadata in digital collections is sparse, finding a sufficient number is challenging.[30]

Another approach is to automatically query for smell gestures in large visual corpora and only manually correct the resulting annotations. Unfortunately, smell gestures are also much more difficult to recognize than olfactory objects. Most objects share some features that are common to their various representations, even though they present some stylistic and historical variations. Recognizing a smell gesture, on the other hand, requires an understanding of the visual content that goes beyond visual features of specific objects: relations between persons and objects, body posture, or the semantic content of the image context all take part in their interpretation as a specific smell gesture. We try to solve the issue by jointly analyzing the poses of the depicted persons and the relationship between their body parts and specific detected objects. This is a complex task, but also an innovative approach, which would for computer vision moves further beyond the state of the art.

Another way of approaching the representation of the smell of roses in the visual arts could be to analyze which other objects accompany the depiction of roses. Table 3 lists the top-ten objects next to the number of times they have been annotated on the same artwork as roses. Investigating those co-occurrences reveals that roses are mostly depicted in conjunction with other flowers and, to a lesser degree, with fruits. This corroborates the assertion that the vast majority of rose depictions occur in still-life paintings, which often contain large bouquets of various flowers and fruits. As before, however, we have to consider the biases of our data collection and annotation process. Are flowers really overwhelmingly represented in these images rather than in other images? Considering the bad image quality and low resolution of many digitized artworks from online collections, flower depictions in other

30 Sofia Collette Ehrich, Caro Verbeek, Mathias Zinnen, Lizzie Marx, Cecilia Bembibre, and Inger Leemans, "Nose-First: Towards an Olfactory Gaze for Digital Art History" In proceedings *Workshops and Tutorials held at LDK 2021*, Zaragoza, Spain, September 1–4, 2021, https://ceur-ws.org/Vol-3064/mdk6.pdf).

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

genres (such as in landscape paintings) might just as well be simply too small to be identifiable as such.

Investigating this hypothesis in detail could be done through an analysis of image-level metadata, specific co-occurrences, or temporal evolution, as it has been exemplified for the number of rose annotations over time. Such a detailed analysis is out of scope of this article. Here, we have aimed for something more modest: to demonstrate how smell-related visual data can be used for quantitative approaches to sensory history. At the same time, we have been keen to demonstrate how inevitable biases of data selection and annotation can be mitigated by a deeper analysis of confounding factors.

Organizing Smells

Now that we have all this historical smell data—extracted from texts and located in images—we have to make it usable. This is where Odeuropa's European Olfactory Knowledge Graph comes in. In the last few decades, so-called knowledge graphs (KGs) have been introduced as a data structure for representing information. The KG structure, at the same time flexible and meaningful, has already been successfully used in organizing information on social networks and in search engines. Researchers have now started to use KGs to harmonize and interconnect heterogeneous data from cultural archives and libraries, offering new possibilities for searching, interacting with, and using the data contained within them.[31] Europeana and SILKNOW are two successful examples of KGs and semantic technologies applied to cultural heritage.[32]

In a knowledge graph, the information is organized into nodes—representing real-world concepts, such as people, places, etc., and even including even abstract concepts—and connecting arcs known as *properties*. These properties have clear meanings as relationships that can connect (i.e., relate) a node with another one (e.g., "France *has capital* Paris" or "Mary *lives in* Paris") or a node with a literal property (e.g., "Mary *has age* 29"). Often, KGs follow the Semantic Web principle, so that every node and property is identified with a uniform resource indicator (URI). So, "France *has capital* Paris" becomes "<http://dbpedia.org/resource/France> <https://dbpedia.org/ontology/capital> <http://dbpedia.org/resource/Paris>" or, using the common convention of using prefixes for namespaces, "<dbr:France dbo:capital dbr:Paris>". Apart from providing an identifier, the URI can be interpreted as an address that allows the users (or the machine) to "follow the link" (navigate) and visit the target resource, which is in this case shown as a web page on their browser. The set of available kinds of properties usually comes from a so-called *ontology*, which describes the rules involving each property as well as the possible *types* or *classes* the entities can

31 Eero Hyvönen, "Using the Semantic Web in Digital Humanities: Shift from Data Publishing to Data-analysis and Serendipitous Knowledge Discovery," *Semantic Web Journal* 11, no. 1 (2020): 187–93, https://www.semantic-web-journal.net/system/files/swj2310.pdf.

32 Antoine Isaac and Bernard Haslhofer, "Europeana Linked Open Data—data.europeana.eu," *Semantic Web Journal* 4, no .3 (2013): 291–97, http://semantic-web-journal.org/sites/default/files/swj297.pdf; Thomas Schleider, Raphaël Troncy, Thibault Ehrhart, Mareike Dorozynski, Franz Rottensteiner, Jorge Sebastián Lozano, Georgia Lo Cicero, "Searching Silk Fabrics by Images Leveraging on Knowledge Graph and Domain Expert Rules," *Proceedings of the 3rd Workshop on Structuring and Understanding of Multimedia heritAge Contents (SUMAC'21)* (2021): 41–49, https://doi.org/10.1145/3475720.3484445.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

belong to (expressions like "dbr:France a dbo:Country" or "dbr:Paris a dbo:City" are examples).

The combination of KG and Semantic Web brings a number of advantages:

- The capability of dereferencing entities and properties with the browser, helping human consumption of the graph content. Dereferencing is basically taking a reference and retrieving the information to which it refers: in the same way that a library call number is a reference, and dereferencing is actually going and taking the volume to which it refers off the shelf.
- Methods and standards for integrating different data sources, promoting the reuse of resources through shared identifiers (URIs) and the interlinking of resources with other datasets in the Web of Data. This possibility allows for greatly enhancing the information in the KG with other existing knowledge sources at little additional cost.
- The possibility to use the data using AI techniques, exploiting the large collection for building automatic systems for different purposes: detecting the topic of a document, guessing some missing metadata, recommending the next element to read with a recommender system, etc.

The first step for building a knowledge graph of smells is to design the data model, which is the description of the data structures to put in place: which type of nodes (classes) and through which properties they are linked. In practice, a data model can be implemented as an ontology or, more frequently, as a set of combined ontologies.

The challenges are multiple. How to represent an entity such as smell, which has characteristics such as ephemerality, indescribability, and subjectivity?[33] How to decide which are the attributes to take into account? A possible solution is to look at expert-driven competency questions like the ones we described above. It can be readily observed that the sources used actually refer to smell experiences. From these, we can realize that time and space are quite important dimensions for the study of olfactory heritage and that these are at the center of the olfactory event, whether that is the generation of a smell (smell emission) or its perception by an agent (olfactory experience).

We decided thus to create the Odeuropa Data Model by extending a well-used ontology in the heritage field: CIDOC CRM.[34] This ontology is widely adopted by libraries and cultural archives and includes ways of connecting the information with cultural objects such as paintings and textual resources. Moreover, the structures of CIDOC CRM help in organizing the information around events. These events can describe the generation or creation of something, or of experiencing something (like in the CRMsci extension for scientific observation). We added to CIDOC CRM a number of classes and properties that are specifically oriented toward smell and can be extended to include other senses such as taste or touch.[35]

33 Pasquale Lisena, Marieke van Erp, Cecilia Bembibre, and Inger Leemans, "Data Mining and Knowledge Graphs as a Backbone for Advanced Olfactory Experiences" (paper presentation, STT 2021: Smell, Taste, and Temperature Interfaces Workshop at CHI 2021, Yokohama, Japan, May 7, 2021), https://stt21.plopes.org/wp-content/uploads/2021/05/STT2021_Odeuropa.pdf.

34 Pasquale Lisena, Daniel Schwabe, Marieke van Erp, Raphaël Troncy, William Tullett, Inger Leemans, and Lizzie Marx, "Capturing the Semantics of Smell: The Odeuropa Data Model for Olfactory Heritage Information" (paper presentation, The Semantic Web: 19th International Conference, ESWC 2022, Hersonissos, Greece, May 29–June 2, 2022): https://doi.org/10.1007/978-3-031-06981-9_23; Martin Doerr, "The CIDOC Conceptual Reference Module: An Ontological Approach to Semantic Interoperability of Metadata," *AI Magazine* 24, no. 3 (2003): 75–92.

35 The full list of these properties can be observed at: https://data.odeuropa.eu/ontology/.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

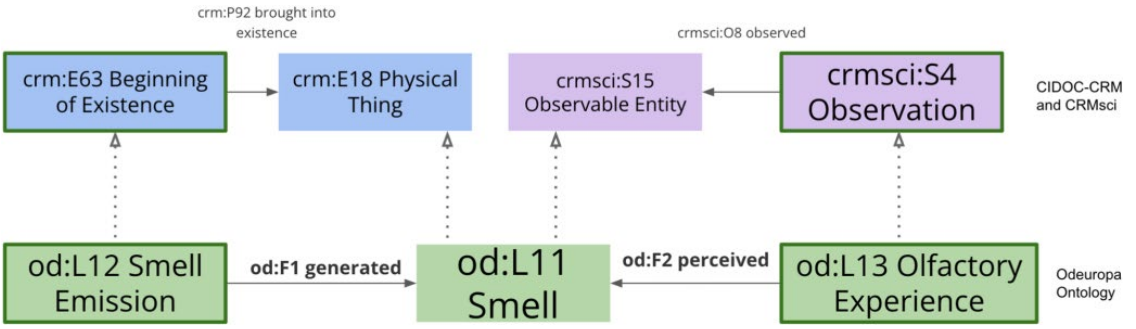


Figure 12. The core elements of the Odeuropa ontology (depicted in green): a *smell*, the event tthat generated it (*smell emission*), and the perception event (*olfactory experience*). These three extend existing elements in CIDOC CRM and CRMsci (depicted respectively in blue and purple). Events are shown with thick borders, and relations as labeled arcs.

The Odeuropa Data Model has been created using a process of iteration, critique, and reiteration involving both history and heritage experts, chemists, and computer scientists. The result is organized around a core structure describing three key elements: the *smell emission* event, the resulting *smell*, and one or more *olfactory experiences* involving the people who perceive the smell.

These elements can be further detailed by relating them (adding an arc) to several aspects: it is possible to link the smell source and the carrier to the smell emission; it is possible to link the perceiver, the emotion, the gestures, and the words used to describe it (attributes) to the olfactory experience; and it is possible to link time and space information to both.

Consider an example coming from *Blagrave's supplement or enlargement to Mr. Nich. Culpeper's English Physitian [...]* (1677), included in the Early English Books Online collection:

> The **Eglantine,** as well as the manured **Roses**, is planted in **Gardens**; if it be set against a wall, under a Window it will cast a **most pleasant** smell into the **room**, [...] in the **Summer Months**.

This example can be represented through the Odeuropa data model using the following graph representation (Figure 13).

Here, the elements in green are the core elements mentioned before. The smell emission happens in the garden, having the two flowers as a source, and producing the smell that is described as "most pleasant" by our unknown perceiver in the room. All of this happens in the summer.

The smell of rose can be present in textual passages in multiple different languages (and Odeuropa works across seven European languages), so how do we make sure that results are harmonized among languages and among synonyms that we may find in the text? In the context of Odeuropa we produced a set of vocabularies listing olfactory objects, fragrant places, gestures, and so on, organized hierarchically and including several labels in different languages.[36] So, the *rose* will contain its

36 The set of vocabularies is browsable at: https://vocab.odeuropa.eu.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

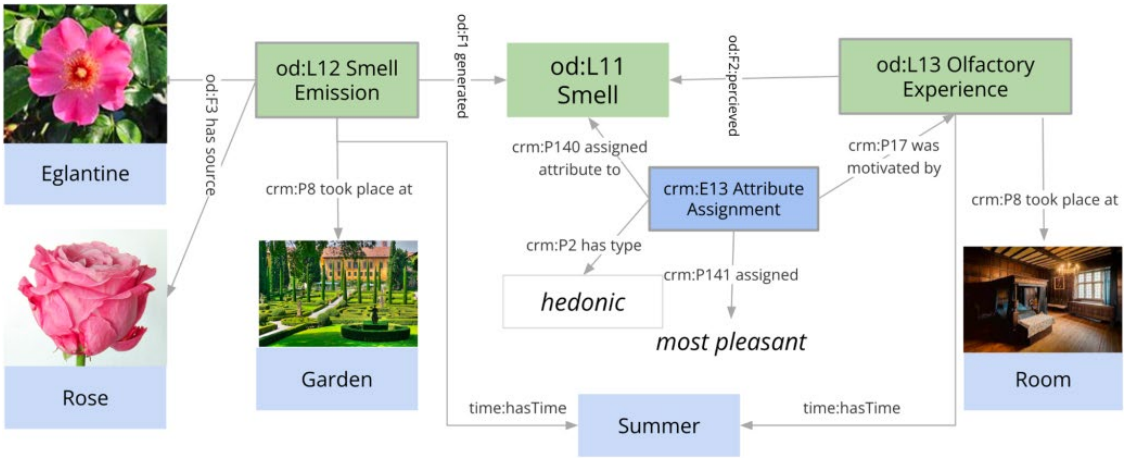


Figure 13. The representation as a graph of a text excerpt using the Odeuropa ontology.

label in the different languages ("rose," "rosa," "roos," etc.) and is represented as part of *flowers*, which is in turn part of *plants*.[37] This allows for grouping results at different abstraction or specificity levels, so results about roses will also appear in queries about flowers in general. Moreover, synonyms are included in order to retrieve results about "tussie mussie" even when you are searching for "posy," both being similar ways of arranging flowers in small bouquets.[38]

After having decided the structure to adopt, the next step is KG *population*, which involves organizing the information extracted from text (section 1) and images (section 2) around this schema and loading it into the knowledge graph. This process consists of the following actions: parsing the output of the text and image annotations, assigning each element to the proper place in the graph; identifying the elements that may include terms in the vocabularies (e.g., the smell source) and matching them with the right standardized vocabulary entry (e.g., "rose"); identifying the part of the graph that can be interlinked with (related to) external resources in the Web of Data and perform the entity matching. Those resources include, for example, encyclopedic resources (WikiData, DBpedia), geographical data (GeoNames), and art archives (Getty ATT).

This process has led to the creation of the European Olfactory Knowledge Graph (EOKG), which has of today more than 220,000 smell references extracted from text and images, represented by over 7 million triples in the format <subject - predicate - object>. This knowledge graph is publicly available and it is possible to interrogate it using SPARQL, a semantic query language.[39] This language uses the same triple structure described above by replacing some positions with placeholders, identified by prepending the "?" sign. These placeholders are the values that we do not know and that we are asking the KG to return, by matching the structure against existing triples in the KG. As an example, the query below, extracts all smells of roses:[40]

37 The "rose" concept in the Odeuropa vocabulary of olfactory objects: http://data.odeuropa.eu/vocabulary/olfactory-objects/138

38 The "tussie mussie" concept in the Odeuropa vocabulary of olfactory objects: http://data.odeuropa.eu/vocabulary/olfactory-objects/162

39 The database's SPARQL interface can be accessed at: https://data.odeuropa.eu/.

40 The query for retrieving all smells emitted by a rose can be directly executed with this link: https://data.odeuropa.eu/sparql?savedQueryName=Smells%20of%20roses&owner=admin. It renders 3,051 results (accessed November 6, 2022).

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

```
PREFIX od: <http://data.odeuropa.eu/ontology/>
PREFIX crm: <http://erlangen-crm.org/current/>

SELECT * WHERE {
        ?emission od:F1_generated ?smell .
        ?emission od:F3_had_source ?smell_source .
        ?smell_source  crm:P137_exemplifies  <http://data.
odeuropa.eu/vocabulary/olfactory-objects/138> .
}
```

After defining the prefixes, we ask the database all triples that match the following pattern (placeholders are in italic): a *smell emission* generates a *smell*; the *smell emission* has a certain *smell source*; and this same *smell source* is exemplified by the concept "rose," here identified by its specific URI.

Whereas reading the query is possible even to people without any background in computer science, creating and manipulating the data in the KG using SPARQL is not trivial and can be an important obstacle to access by the broader community. For this reason, all the EOKG data are also accessible via an online, user-friendly web application called the Odeuropa Smell Explorer.[41] The tool does not require any programming by the user, and it has two main features:

First, it offers an advanced search function for smell references, with the possibility to select textual or visual references (or both), perform a full-text search, or select one of the terms of the multilingual vocabularies for filtering the relevant results in a multilingual fashion.

Second, it offers the possibility to look at the list of smell sources and fragrant spaces and see, for each of them, aggregated results such as *timelines* (to see the evolution of a smell in time), *maps* (to identify areas in which a smell is particularly present), and *word clouds* (to understand how this smell has been perceived and described by people).

In addition, it is also possible to log into the application and create a named list of elements. In this way, while using the tools, it would be possible to save the visualized results in one of these lists, creating personal collections of olfactory references like a scrapbook, with the possibility to keep them private or to share them with other people. Be aware that the EOKG and its interface are still in development. More data is coming in regularly.

Is the Technology Any Good?

To know whether the tools we are developing are any good (by the standards of computer scientists), we need some way to evaluate them. In computer science, there are three different dimensions to

41 The Odeuropa Smell Explorer can be accessed at: https://explorer.odeuropa.eu/.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

evaluations: *evaluation task* (intrinsic or extrinsic), *assessment type* (quantitative or qualitative), and *evaluation metrics* (the indicators used to assess the performance). In the following paragraphs, each dimension is discussed. While these assessments are quite different from those practiced by historians, if interdisciplinary projects are to be successful then we must be aware of what success means in every participating discipline: we have to know what makes our respective collaborators tick.

In an intrinsic evaluation, the expected output of a system is rated, often down to a detailed task or even subtask. In an extrinsic evaluation, the tool is rated as part of a larger task. An extrinsic evaluation of the text-extraction efforts could be to have users assess the usefulness of the extracted text snippets while they are using the explorer. An informative evaluation furthermore consists of a quantitative and a qualitative part. The quantitative part is driven by metrics and provides overall scores, much like a test score in a university exam. The qualitative part is driven by an investigation of what (types of) errors the system makes or what it performs well on, much akin to giving a student feedback on what they still need to work on.

A core resource in an intrinsic evaluation is a benchmark dataset, such as the ones we have devised for the text and image recognition tasks described in sections 1 and 2. A benchmark dataset can be considered a "test" for a computer tool and serve as a reference against which to compare different approaches being evaluated. As we develop tools to recognize references to smells in hitherto unknown texts and images, we ensure that we have a dataset in which we have marked those smell references—or "the correct answers." We then give the tool the texts or images from that benchmark dataset and ask it to recognize the correct smell elements. We then compare the system output to the "answers" in the benchmark dataset. Table 2 in section 1, for example, states that the English section of our benchmark dataset contains 1,530 smell words, 1,313 smell sources, and 1,084 qualities. This means that we can test our language technology tool on potentially 1,530 smell words, 1,313 smell sources, and 1,084 qualities.

As the test is set up a little bit differently from a regular test as we may know it from our student days, a set of metrics is reported that is more elaborate than simply accuracy, or "*x* out of *y* correct." Say we have a sentence containing nine words:

The place was strongly scented [*with attar of roses*].

In this sentence, "scented" is the smell word and "attar of roses" is the smell source. In such linguistic detection tasks, it is customary to evaluate the system on the word level; this means that we aim for the system to predict that there is one smell word, four smell source words, and four words that do not denote any smell-related references.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

There are several types of mistakes that a system can make. A system can "miss" a reference (i.e., a false negative), or it can erroneously mark a word as a smell reference (i.e., a false positive). Say our system recognizes "scented" as a smell word, that it erroneously marks "attar" as a smell word, and that it correctly marks "roses" as a smell source. If we were to report accuracy for this system, it would get a score of six out of nine, or 66.7 percent. Per category of things to predict for the system, it got a 100 percent score on the non-smell reference words, 100 percent on the smell word, and 25 percent on the smell source reference words.

When reading language technology papers, one will often come across the terms *precision*, *recall*, and *F-score*. Precision and recall were developed in 1955 to assess the performance of an information retrieval system.[42] Precision is a measure of quality—in a set of results, it is the fraction of correct results over all returned ones. In our example sentence, the precision score for the smell words would be one out of two (50 percent), as it returned two smell words but only one is correct. The precision score for the smell sources would be one out of one (100 percent), as one term was correctly recognized, and the score for the non-smell words is four out of six (66.7 percent). Recall is a measure of quantity, or sensitivity. In a set of results, it measures how many of the correct results were returned by the system, among all possible correct ones. In our case, for the smell words, one out of one was returned (100 percent); one out of the four smell source words was returned (25 percent); and four out of four of the non-smell word references were returned (100 percent).

To provide a summary of the scoring, the F-score measure is used as a way to combine the different aspects measured by precision and recall. This metric was defined in 1979 by van Rijsbergen as the *effectiveness function*.[43] Later this was renamed the F-score or Fβ score. The Fβ score allows for giving either precision or recall more weight, but generally the harmonic mean (F1), giving each an equal weight, is reported. Thus, our example system achieves a score of 66.7 percent for the smell words, 40.0 percent for the smell sources, and 80.0 percent for the non-smell words. For our example sentence, the overall weighted average precision comes to 79.6 percent, the precision to 66.7 percent, and the F1-score to 60.7 percent.

It is also possible to devise a stricter evaluation of the recognition of sequences of words. In a phrase-based evaluation, we would evaluate "with attar of roses" as a whole. A system would only "get a point," so to speak, if it recognized this entire phrase correctly. Since our example system only recognized "roses," the score would thus be zero for the recognition of smell sources. While in a sense this type of evaluation can provide a better indication of the performance of a system (since in many recognition problems, detecting fragments of the phrase is not that useful to end-users), there are issues with this type of evaluation.

42 Allen Kent, Madeline M. Berry, Fred U. Luehrs Jr., J. W. Perry, "Machine Literature Searching VIII: Operational Criteria for Designing Information Retrieval Systems," *American Documentation* 6, no. 2 (1955): 93–101.

43 C. J. van Rijsbergen, *Information Retrieval* (London, 1975).

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

One core issue is that all tokens in the sequence are equally important, whereas in our example, not recognizing the word "with" is less problematic than not recognizing "attar" or "roses." In fact, a system recognizing "attar of roses" still yields a useful result to the end-user. Overall, the current Odeuropa system for olfactory information extraction in English yields an F_1score of 88.5 percent on smell words, 52.1 percent on smell sources, and 70.6 percent on qualities based on this stricter evaluation approach. However, as part of our error analysis, we are also looking into the types of words the system is forgetting and whether this is harmful with respect to the end goal.

Evaluating computer vision systems comes with its own challenges. Whereas in an English, Dutch, French, or German text there are sentences and words that can denote possible segment endings, in a painting a computer vision tool needs to figure out where an object's boundaries are. To evaluate this, metrics need to take into account the overlap between an object bounding box in the benchmark dataset to compare it with an object bounding box in the system prediction. This is measured in terms of intersection over union (IoU)—i.e., the ratio of the area of the predicted and the so-called ground truth (reference) bounding box. Then, the intersection of those bounding boxes is computed as well as their union. The closer the IoU ratio is to 1, the more precise the prediction. For evaluation purposes, cutoff points have to be chosen. For example, a IoU ratio of 0.95 would state that all results that have an IoU larger than 0.95 are considered as "positive" predictions, while all predictions lower would be denoted as "negative." In practice, multiple thresholds are often considered and afterward averaged. The commonly used COCO evaluation protocol uses the average over all ten thresholds (0.5, 0.55, ..., 0.95) to compute the mean average precision (AP).[44] The final metric is then computed as the mean of the AP values for each class to be detected.

In 2022, we organized the ODOR challenge.[45] Organizing a challenge is a common method to advance the state of the art in the computer science community. Basically, a research group or consortium devises a task and an accompanying dataset and invites the community to develop a tool or system to solve this task within a particular time frame. Often the datasets will remain available to the research community afterward, and they can become core benchmark datasets against which to test systems. The main objective of this challenge was to develop a system for object detection of smell-related objects. The mean average precision (mAP) of the best competitor of the ODOR challenge achieved a mAP of 11.49 percent. This is very promising given the challenging nature of the dataset. The best current object detection model for photographs achieves 64.5 percent mAP on the COCO dataset. However, the COCO dataset was trained on a much larger corpus. Moreover, the image bounding boxes in our ODOR dataset are much smaller in comparison to the COCO dataset. Our current best performing model achieves an mAP of 14 percent, which we aim to further improve in the course of the project. Nevertheless, the current state is already good enough to

44 Tsung- Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, Piotr Dollár, "Microsoft COCO: Common Objects in Context," in *Computer Vision: European Conference on Computer Vision 2014*, vol. 8693, ed. David Fleet et al. (2014), 740–55, https://doi.org/10.1007/978-3-319-10602-1_48.

45 You can find more information for the ODOR challenge at: https://odor-challenge.odeuropa.eu/.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

use the automatic recognition system as a query tool to quickly explore unseen data (Figure 14).

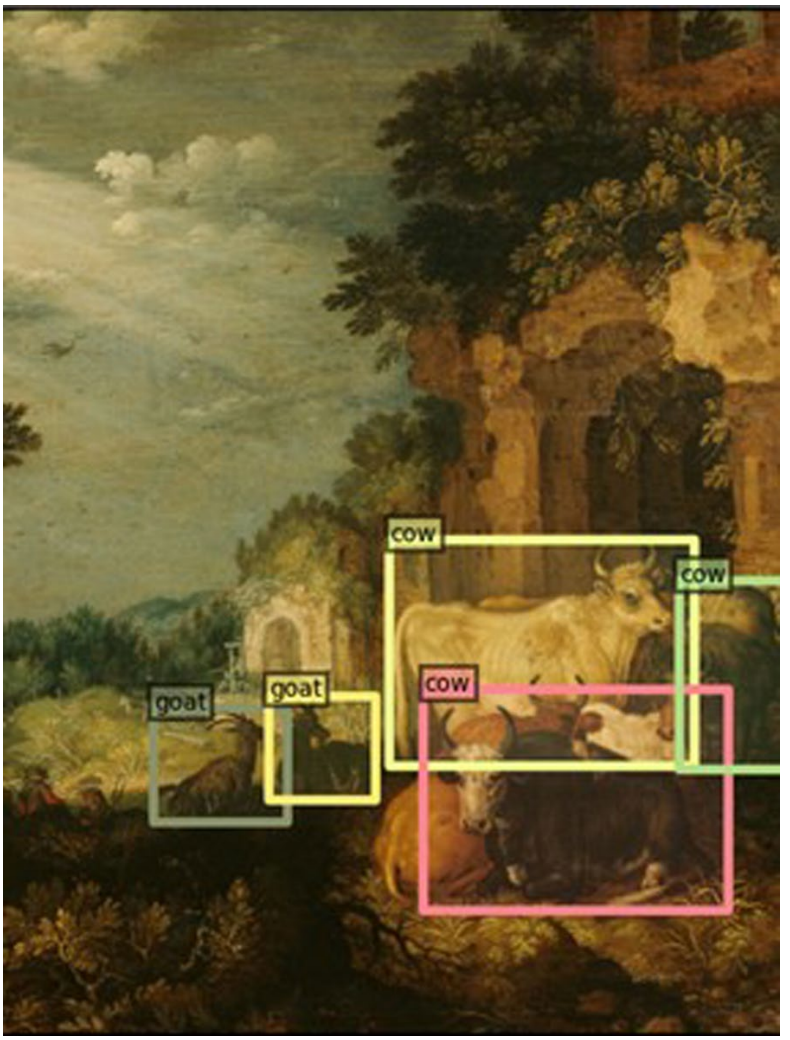


Figure 14. Detail from Roelant Savery, *Landschap met ruïnes, vee en herten* [Landscape with ruins, cattle and deer], ca. 1614–20, oil on panel, Rijksmuseum Rijksstudio.

At the time of writing, we are organizing the MUSTI challenge as a task in the Multimedia Evaluation (MediaEval) benchmarking initiative.[46] This challenge is focused on both language and computer vision technology, as the goal is to recognize references to smells in texts and images and to connect them—i.e. match a text snippet mentioning a smell to the correct artwork.

46 The MediaEval 2022 benchmarking tool can be found at: https://multimediaeval.github.io/editions/2022/.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

Table 4. Number of competency questions per category and number of answerable questions with the Odeuropa data model (the ontology), in combination with AI techniques (AI), with the addition of external data (ExtData), and with a further extension of the model (Extension).

Category	Ontology	External			Total
		AI	Data	Extension	
A. Smells	10	0	0	1	11
B. Noses and gestures	6	0	0	0	6
C. Identities	6	0	0	0	6
D. Emotions	6	0	0	0	6
E. Practices	8	5	0	0	13
F. Sites and contexts	9	0	2	0	11
G. Texts and images	19	0	2	0	21
TOTAL	64	5	4	1	74

Because the text and image extraction has generated millions of statements to ingest into the knowledge graph, an intrinsic evaluation of all elements is not possible. For the knowledge graph, we therefore focused on evaluating the data model against the competency questions; we then carried out further extrinsic evaluations, with users performing tasks using the knowledge graph. In total, seventy-four competency questions were collected from humanities and heritage science scholars. The competency questions served both as a guide in the development of the model and as a means to test it. Our test on the final model yielded four different results (see Table 4 for a summary of results):[47]

The vast majority of questions can be answered with a SPARQL query.[48] Example: *What are the most frequent smell sources in Paris in the 18th century?*

A few questions cannot be answered by simple SPARQL queries but instead require additional AI methods to find a proper solution. Example: *Was muck perceived as more disgusting than smoke?*

Four questions are answerable with SPARQL but require external information that is available only in sources that are outside the scope of the Odeuropa model (e.g., with the addition of knowledge bases such as WikiData). Example: *Which smell could be perceived during a war?*

One question requires an extension of the model. Given that the challenging element of this question is not directly related to the olfactory/heritage information but rather directly related to time representation, we decided to keep this issue open for future work. Example: *Which smells were perceived during the morning?*

Future work will involve investigating whether the EOKG is eventually returning correct results to corresponding queries for these questions, establishing its usefulness for the intended research purposes.

47 Pasquale Lisena, Daniel Schwabe, Marieke van Erp, Raphaël Troncy, William Tullett, Inger Leemans, and Lizzie Marx, "Capturing the Semantics of Smell: The Odeuropa Data Model for Olfactory Heritage Information" (paper presentation, The Semantic Web: 19th International Conference, ESWC 2022, Hersonissos, Greece, May 29–June 2, 2022):

48 SPARQL is the query language of the Semantic Web as specified in Eric Prud'hommeaux and Andy Seaborne, "SPARQL Query Language for RDF: W3C Recommendation 15 January 2008," https://www.w3.org/TR/rdf-sparql-query/.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024



Figure 15. Automatically generated word-cloud of corresponding adjectives from the Smell Source entry for "Rose" in Odeuropa Explorer.

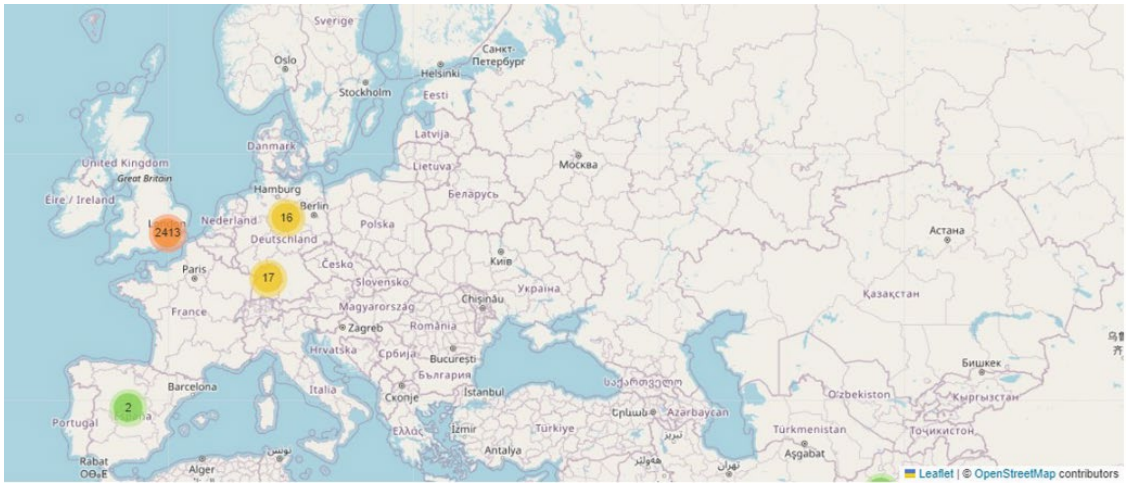


Figure 16. Map of results with locatable places from the Smell Source entry for "Rose" in Odeuropa Explorer.

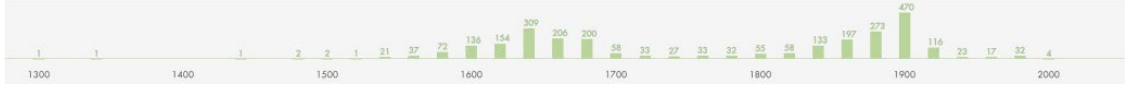


Figure 17. The number of occurences over time from the Smell Source entry for "Rose" in the Odeuropa Explorer.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024



Figure 18. Examples of textual and image occurrences of rose from the Smell Source entry for "Rose" in the Odeuropa Explorer.

Of course, the question of whether the technology is any good will vary for different scholars. For historians, the question will probably be how well the technology can help us answer historical questions. To think about this, we can turn to the Odeuropa Smell Explorer, where all of the processes we have discussed so far come together. Let us return to our rose-related "competency questions" and see what we can get through the Smell Explorer tool. Throughout the article we have already begun to answer some these questions: in our discussion of text processing, we addressed the words used to describe the scent of the rose; and in the section on images, we began to explore the co-occurrences of roses with other scents and the ability of roses to stand in as a visual signifier for smelling in general. Now we can turn to some of the visualizations included as part of the entry for "rose" in the Smell Explorer: word clouds, maps, and timelines. As we will see, the Smell Explorer tool does not provide easy answers. Instead, it prompts new questions that lead us, by the nose, down different tracks into the past: What adjectives are used

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

to describe the smell of roses in historical texts? Does this change over time? Are the words different in different languages?

The first group of questions centered on attitudes to smells, as evidenced by the changing words used to describe smells over time. Consider, for instance, dynamic word clouds generated by the entire set of processed historical data from the two decades with the most results for roses in the Smell Explorer (the 1640s and the 1900s). By comparing them, it becomes evident that the adjectives used to describe the smell of the rose are remarkably consistent: *fragrant*, *natural*, *faint*, *indescribable*, *rich*, and *sweet smelling* occur in the mid-seventeenth century and in the early twentieth century in several languages (principally Dutch and English). One explanation for this adjectival consistency could be that descriptions of rose odors are recycled and reused, becoming tropological in character. The durability of descriptors for rose scents therefore turns not on changing olfactory perceptions of the rose but on literary tradition and intertextual reference.

There are, however, subtle shifts in relative frequency. In the 1900s the terms "indescribable" and "faint" appear larger in our dynamic word clouds and therefore more common (frequent) in the data. Could this reflect the softer and weaker scent of the rose varieties that became popular in the nineteenth century, when horticulturalists increasingly bred their plants for visual beauty and when the lighter, hay-like scent of tea roses reigned supreme over flora's European domain? This is naturally where our close reading will need to come in and follow up our new queries: Are there particular places where rose scents were likely to be found in Europe in the period from the 1600s to the 1920s?

At this moment, this question is hard to answer. We are still training the computer to recognize places and spaces in the smell-event quotes (for instance, a rose being smelled in a *kitchen* [space] in *Amsterdam* [place]). For all smell events we cannot locate, the EOKG now takes location information from the metadata of the text or image: the place of publication of a book, or the place of the gallery that holds the image. Hopefully in 2023, the results will provide better insights for this question. Right now, this shows us the biases in our datasets.

There are additional questions as well: What smells might have co-occurred with the scent of roses? What was the smell of roses used for? Was it deployed in specific ways to specific ends?

On the question of co-occurrence, the image data is pretty clear: roses tend to appear with other flowers in historical images that include—among others—carnations, irises, and jasmine. However, as we have discussed, roses tend to be easiest to identify in images where they are front and center, and these are more likely to be found in still-life paintings of flowers in vases. However, a number of other interesting relationships also emerge from the images if we cycle through those collected in the Smell Explorer. For example, wilting roses are often portrayed alongside tobacco pipes, carcasses, and skulls, with

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

the implication that the fading scent of the rose was one element in the composition of memento mori images that meditated on the inevitability of decay, death, and impermanence. An open question, though, is how the scent of roses was used and in which contexts might this scent be deployed. Over the coming year, as we begin to apply other frames to our text annotation, we will begin to excavate the particular *practices* that relate to the smells we find in our scent extracts.

Computationally Supported Humanities Research

The Odeuropa consortium strives to facilitate computationally *supported* (historical) research. We do not aim for the computer to answer the questions; we designed it to support scholars and other interested parties in finding the right information to answer their questions. In order to assure the quality of the data and the tools, we organize various "hackathons" where different audiences are invited to play around with the data and tools, evaluating the assets and taking them—if wanted—into their own hands and research.

In August 2022, during the PastScent Hackathon, the Odeuropa Smell Explorer was first put to the test by a group of six historians, five art historians, four heritage scholars and professionals, two literature scholars, and five fragrance experts. For two days, the participants immersed themselves in Odeuropa's processed text and image data via the Smell Explorer interface to develop olfactory storylines.[49] The overall participants' responses to the Odeuropa tools were positive, and they provided the team with helpful comments and suggestions to further improve the tooling. Furthermore, the workshop highlighted the flexibility of working with a tool based on a Semantic Web model that allows for different entry points to the underlying data. The different teams worked with the same data, but each had a different focus and came to a different storyline based on a different perspective: a single scent through time and space; the shifting smellscapes of a landscape over time; an exploration of the lost scents of paradise and the lost smells of ecological collapse; an olfactory journey through the decks of a historical ship; and a narrative following the smells of pilgrimages across the European continent.

It is important to note that during the workshop, Odeuropa team members were on hand to explain peculiarities of the system and the underlying technology. As scholars start working with data derived via (semi-)automatic means, they need to be made aware of the methodological and technological background of the data and tools. This means, for example, being aware of the data selection, which affects any quantitative observations made on the data—simply because we can access *more* data through automated processing, it does not mean we can access *all* data. Limitations here lie in availability (both licensing

49 A full overview of the workshop can be found in William Tullett, "Deliverable D5.2: Academic Conference on Olfactory History and Digital Heritage," PastScent (website), forthcoming. This and other Odeuropa deliverables, once approved, are available at: https://odeuropa.eu/deliverables/.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

and the extent of a collection's digitization) as well as in the quality of the data. In particular for objects that were digitized longer ago, the photograph or scan and potential ensuing steps undertaken to make it machine-interpretable could be of too low a quality to perform automatic analyses.[50] These are of course also issues in traditional archival research, but the sheer size of digitized collections might sometimes make it more difficult to assess the gaps in the data.

The other main issue to be aware of is that the tools are not perfect, because they are extensions of human intentions and choices. In Odeuropa, we aim to be transparent about the strengths and limitations of the developed tools, and where possible we provide confidence scores that can indicate how certain the system is of a particular decision (choice) as well as link to the original text and image sources. When encountering errors in automatically derived data, it may be easy to immediately dismiss the entire system as useless, but without proper feedback and without working together on improving the tools, we will never be able to develop good enough tools. It also ignores the fact that humans make errors too. Ultimately, both computers and humans get things wrong: sometimes they make the same errors (and humans build errors *into* technologies), and sometimes they make different ones.

With these caveats in mind, we hope you will take a look at the exciting developments in computationally supported history research and try out the tools for yourself, be it the Odeuropa Smell Explorer or another project's interface—because only through transdisciplinary collaborations can we start to make sense of all the great digitized data that is currently available to us.

50 For texts, the digitization process often involves performing optical character recognition (OCR) or handwritten text recognition (HTR) to turn a photo or scan of a text into a machine-readable text. Fraktur, smudged ink, folds in the paper, strikethroughs, or sloppy handwriting can hamper these processes.

Marieke van Erp leads the Digital Humanities Research Lab (DHLab) at KNAW Humanities Cluster. She holds a PhD (2010) in language technology from Tilburg University in the Netherlands and specializes in combining natural language processing and semantic web technologies to further humanities research. She is both project manager and researcher with the Odeuropa project.

William Tullett is Associate Professor of sensory history at Anglia Ruskin University in Cambridge. His first book was *Smell in Eighteenth-Century England* (Oxford University Press, 2019) and his second book, *Smell and the Past: Noses, Archives, Narratives* will be published by Bloomsbury in May 2023. As part of the Odeuropa project, he is developing an online Encyclopedia of Smell History and Heritage.

Vincent Christlein received his PhD (Dr.-Ing.) in computer science in 2018 from the Friedrich-Alexander University of Erlangen-Nürnberg (FAU) in Germany. He heads the Computer Vision Group at the Pattern Recognition Lab (FAU), which covers projects such as glacier segmentation and document and art analysis.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

Thibault Ehrhart is a research engineer in the data science department at EURECOM. He has experience in software engineering, notably in web and mobile development. He currently works on projects related to semantic web technologies, information extraction, and machine learning. With the Odeuropa project, he contributed to the exploration of the Knowledge Graph through the development of the Smell Explorer and the mobile app City Sniffers.

Ali Hürriyetoğlu works on historical multilingual text processing with the Odeuropa project at the Royal Netherlands Academy of Arts and Sciences (KNAW) Humanities Cluster DHLab. He has worked in academic, industrial, and governmental settings to process news and social media text in various domains for information extraction tasks. His recent research focus is on the robustness and the generalizability of text processing systems across contexts.

Inger Leemans is Professor of cultural history at Vrije Universiteit Amsterdam. She is PI of both the Odeuropa project and the NL-Lab, a research group on Dutch culture and identity at the Humanities Cluster within the KNAW. Her research focuses on early modern cultural history, the history of emotions and the senses, cultural economy, and digital humanities, and she has published about the history of pornography, the (radical) Enlightenment, stock markets, and financial crises.

Pasquale Lisena is a researcher in the data science department at EURECOM, working on knowledge modeling and information extraction in digital humanities. He received his PhD in computer science from the Sorbonne in 2019. He works on population and exploitation of knowledge graphs, such as the European Olfactory Knowledge Graph.

Stefano Menini is a researcher in the digital humanities research group at Fondazione Bruno Kessler in Italy with an interdisciplinary background of computational linguistics and philosophy. The focus of his work is on the application of natural language processing techniques to explore the content and extract different typologies of information from texts. With the Odeuropa project, he has been developing systems to extract olfactory descriptions and the perception of smells from multilingual historical texts.

Daniel Schwabe is a visiting researcher at the AI Laboratory at the Joszef Stefan Institute in Italy. He is a retired Full Professor in informatics at the Pontifical Catholic University of Rio de Janeiro, and collaborating researcher with the Knowledge Graph Center at the University of Southern California Information Sciences Institute. His research interests focus on knowledge modeling and knowledge graphs.

Sara Tonelli is the head of the digital humanities research group at Fondazione Bruno Kessler. She received a PhD in language sciences from Università Ca' Foscari, Venice in 2010. She has been involved in several national and European projects dealing with historical

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024

archives, event and temporal information extraction and political stance detection.

Raphaël Troncy has been an associate professor of data science at EURECOM since 2009. He obtained a PhD in 2004 from Grenoble University, supported by a fellowship at the National Audio-Visual Institute (INA) of Paris. He is an expert in combining semantic technologies, natural language processing and recommender systems in particular for the creative industries and the galleries, libraries, archives, and museums sector.

Mathias Zinnen studied philosophy, history and computer ccience in Mainz, Berlin, and Erlangen. Currently, he is pursuing his PhD in computer science at the Pattern Recognition Lab of Friedrich-Alexander-Universität Erlangen-Nürnberg (Germany). With the Odeuropa project, he is working on the automatic recognition of smell references in historical artworks using methods of computer vision. His research interests include domain adaptation techniques for computer vision in the visual arts such as transfer learning, contrastive pretraining, and style transfer.

Author contributions (by author initials) are listed according to the Contributor Roles Taxonomy (CRediT). The use of this taxonomy makes the various contributions of the authors explicit to give authors proper recognition. Information on the CRediT can be found at: https://credit.niso.org/.

Conceptualization: MvE, WT, IL
Data curation: MvE, MZ, SM, PL, AH
Formal Analysis: MZ, WT, SM, PL, DS
Funding acquisition: MvE, WT, IL, ST, VC
Investigation: MZ, WT, SM, ST, PL, DS
Methodology: WT, IL, PL, DS
Project administration: MvE
Resources: SM, ST, PL
Software: MZ, SM, PL, TE
Supervision: VC
Visualization: MZ, SM, PL
Writing (original draft): MvE, WT, MZ, SM, ST, PL, VC
Writing (review and editing): MvE, WT, IL, VC, DS

This work is supported by the European Union's Horizon 2020 research and innovation program within the Odeuropa project (grant agreement no. 101004469). The authors would like to thank the rest of the Odeuropa team, the data annotation team, and the PastScent workshop participants who test drove the first version of the Smell Explorer and contributed to the first version of the storylines.

Downloaded from https://academic.oup.com/ahr/article/128/1/335/7098185 by guest on 18 November 2024