

Scents and Sensibility: Evaluating Information Olfactation

Andrea Batch,¹ Biswaksen Patnaik,² Moses Akazue,³ and Niklas Elmqvist^{1,2}

¹College of Information Studies, University of Maryland, College Park, MD, USA

²Department of Computer Science, University of Maryland, College Park, MD, USA

³School of Computing Science, University of Glasgow, Glasgow, Scotland, United Kingdom
ajulca@umd.edu, bpatnaik@umd.edu, mooses.akazue@glasgow.ac.uk, elm@umd.edu

ABSTRACT

Olfaction—the sense of smell—is one of the least explored of the human senses for conveying abstract information. In this paper, we conduct a comprehensive perceptual experiment on *information olfactation*: the use of olfactory and cross-modal sensory marks and channels to convey data. More specifically, following the example from graphical perception studies, we design an experiment that studies the perceptual accuracy of four cross-modal sensory channels—scent type, scent intensity, airflow, and temperature—for conveying three different types of data—nominal, ordinal, and quantitative. We also present details of a 24-scent multi-sensory display and its software framework that we designed in order to run this experiment. Our results yield a ranking of olfactory and cross-modal sensory channels that follows similar principles as classic rankings for visual channels.

Author Keywords

Olfactory perception, information olfactation, olfactory displays, scents, smell, evaluation.

CCS Concepts

•Human-centered computing → Visualization; Empirical studies in visualization; Visualization design and evaluation methods;

INTRODUCTION

It is a truth universally acknowledged, that a new research topic, such as *information olfactation* [58], in possession of a good theoretical framework, must be in want of empirical validation. Following practice in graphical perception [11, 15, 54], in this paper we report on a controlled perceptual experiment designed to elicit internal rankings of four *sensory channels* (analogous to *visual channels* [57] or *visual variables* [11]) for three different forms of data: nominal, ordinal, and quantitative [66]. The channels included scent type, amount of scent, speed of the airflow, and air temperature.

Strictly speaking, only the former two are actually olfactory channels; wind speed and temperature are tactile stimuli sensed by the skin, and not the nose’s olfactory receptors. However, we include them here because they are easily generated and manipulated by an olfactory display. Furthermore, we opted not to include a stereoscopic channel due to the complexity of the implementation as well as its perceptual imprecision, where people tend to turn their head rather than relying on differential sensing from two nostrils [60].

Conducting this evaluation required fabricating an information olfactation display capable of supporting all of these sensory channels within the necessary data ranges. Thus, a secondary contribution of this paper is our olfactory display consisting of 24 essential oil bottles controlled using ultrasonic diffusers (Figure 1). The display is controlled using a software API interfacing with an Arduino based device. Beyond the essential oil containers, which typically is configured to emit six different smells at four different intensities each, the display can also control the temperature using thermoelectric heating and cooling (a separate chamber fitted with heating coils and Peltier modules) as well as wind speed using controllable fans.

Not surprisingly, our results mostly follow analogous results from graphical perception. In particular, based on accuracy perception for different stimuli, we found that quantitative data is best represented by temperature and wind speed, and nominal data is best represented by temperature and scent type. However, we were surprised that ordinal data was best conveyed using scent type, which typically has no encoded ordering. That scent intensity was outperformed by all other channels was also unexpected, although we explain this with the fact that an intensity of the same scent is difficult. This framework of olfactory channel effectiveness provides a reference for designers and opens the design space of encoding information through smells for interactive immersive displays, ubiquitous analytics, and general analytical environments.

Our contributions are threefold: First, we propose a methodological standard for evaluating olfactory display techniques based on existing graphical perception methods and general olfactory studies. Second, we present a olfactory display system capable of harnessing four sensory channels—scent type, scent intensity, wind speed, and air temperature—to repre-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CHI '20, April 25–30, 2020, Honolulu, HI, USA.

© 2020 Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-6708-0/20/04 ...\$15.00.
<http://dx.doi.org/10.1145/3313831.3376733>

sent data. Finally, we use the results from our experiment to establish a ranking of sensory channel effectiveness.

RELATED WORK

Our work spans perception, novel devices for visualization, aspects of multimodality and accessibility, and olfactory displays. Here we review relevant literature in these areas.

Perceptual Psychology

Perceptual psychology [34], a part of cognitive psychology, is concerned with the human sensory system, which in turn can be seen as the preconscious aspects of human cognition [33, 65]. In general, the interpretation of any external stimulus by our sensory system—e.g., sight, sound, touch, smell, or taste—falls under the umbrella of perception [34]. Thus, perceptual psychology is of interest to the data visualization community not because of the stimulus itself—typically visual—or even the characteristics of the sensory systems, but rather in terms of the information-carrying capacity of the stimulus as well as our bodies’ ability to interpret this information [72].

For data visualization, therefore, the primary instrument for understanding perception is the *graphical perception* experiment, where human participants are asked to interpret visual stimulus in a controlled laboratory setting. Some of the early work in this vein dates back to paper-based statistical graphics, such as results by Eells et al. [28] from 1927, by Croxton et al. [20] in 1927 on pie charts, bar charts, and circle diagrams, and by Croxton et al. [19] in 1932 on shapes for comparison. The work of Peterson and Schramm [59] from 1954 is seminal in that it systemically studied eight different statistical graphs and derived resulting guidelines.

Modern graphical perception for visualization can be said to start from such holistic surveys that not merely measure accuracy for individual chart types or shapes, but attempt to study and rank multiple ones. Already in 1967, Jacques Bertin, a cartographer by training, assembled a ranking of so-called *visual variables* (also known as *visual channels* [57]) from his personal expertise and current practice in cartography [11]. Cleveland and McGill [15] assembled results from many perception studies to provide a similar ranking of visual variables backed by empirical data; remarkably, the rankings are more or less identical. Mackinlay [54] later extended Cleveland and McGill’s ranking into a more fine-grained model suitable for automatic visualization design. While Mackinlay never empirically verified his model, his ranking is foundational in that it introduced variations depending on whether the data to visualize is nominal, ordinal, or quantitative [66].

Novel Devices for Visualization

The mouse, keyboard, and monitor have long been the reigning input and output hardware for visualization [52], but this is slowly changing as device technology progresses beyond the traditional personal computer. Recent years have thus seen advances for visualization on large displays [2, 7], digital tabletops [44, 71], multi-user environments [3, 25], mobile devices [10, 63], and even smartwatches [14, 41].

Ubiquitous analytics [29] harnesses this menagerie of available computing devices in a user’s immediate surroundings for

anytime and anywhere data analysis. Common among all these efforts are middleware infrastructures for managing such ecosystems [4, 5, 6]. Another focus is on the efficient use of available display space; in particular, Horak et al. [42] propose heuristics and algorithms for managing semantic layouts of visualization views across multiple surfaces.

Another trend is on the use of immersive technologies, such as virtual, mixed, and augmented reality, for data visualization. Typically captured by the shared moniker “immersive analytics” [55], this family of tools and techniques rely on providing fluid experiences for visualization [30] using the full 3D environment of the user—virtual or real—as a canvas for data analysis. Examples include the Glyphmaker system for visualizing highly correlated multidimensional data in VR [62], the multidimensional analysis system ImAxes [17], and a collaborative graph visualization tool for HMDs and CAVEs [18]. Such immersive tools are particularly well-suited for multimodal embellishments that go beyond mere vision, including the sense of smell. We discuss these topics in detail below.

Multimodality and Accessibility

While visualization is traditionally based on transforming symbolic data into geometric representations [12], there is a rich plethora of other media that could be used for this purpose. The motivation for such “multimodal visualization”—i.e., that goes beyond the visual medium—is typically to either (a) augment stimulus, (b) provide complementary stimulus for when the user’s eyes are busy, or (c) replace the stimulus altogether for visually impaired users. In assistive technologies, the latter is known as *sensory substitution* [16]: replacing input from one sensory modality with another modality, such as converting written text into spoken language.

Regardless of motivation, the most common multimodal visualization is *sonification*, where data is represented using audio. Zhao et al. [80] combine sound and speech to allow visually impaired users explore linked maps and tables. Goncu et al. [35] automatically translate floor plans into accessible ones. Finally, Ferres et al. [31] propose a natural language interface that uses speech to describe line graphs.

Touch is another powerful method for representing data beyond the visual medium. For example, embossed touch maps [22] convey data in a 2D area using shape, electrovibration can be used to generate tactile feedback of 2D data [78], and multimodal VR with force feedback is superior to traditional printed tactile media [79]. The natural extension to tactile visualization is *data physicalization* [46], where physical artifacts are used to convey data. Examples include 3D-printed tangible maps [48, 40], physical bar charts [69, 70], and wheeled micro-robots for representing data [51].

Olfactory Displays

A specialized form of data physicalization uses smell to represent data. If *olfaction* is the sense of smell, then an *olfactory display* is a programmable device that is capable of creating an olfactory stimulus by (a) emitting odorous molecules (chemostimulation) [64], or (b) directly activating odor receptors in the nose (electro-stimulation) [38]. The former category—creating olfactory stimulus by emitting odor—can be further

organized based on its mode of distribution: ultrasonic atomization, atomization through Venturi effect, and evaporative diffusion. Patnaik et al. [58] surveys these mechanisms.

The most straightforward usage of olfactory displays is for increasing presence in immersive applications, such as Virtual Reality training and recreation. In fact, *Sensorama* [39], the very first VR implementation and patented in 1962, included both “at least one” scent channel, as well as a fan to generate a breeze on the user’s face.

However, our interest in this paper is more narrow in that we focus on olfactory displays used for *information olfaction* [58]: using scent to convey abstract data—such as stock market price over time, node types in a social network, or the distribution of data in a histogram—rather than a realistic phenomenon—such as the smell of a damp cave in a dungeon crawler, the tang of gunpowder in a combat simulation, or the heavy aroma of motor oil in an airplane mechanic training application. Washburn and Jones [73] were among the first to suggest this practice, listing several existing olfactory devices that could be used for data visualization. However, most existing displays are typically used for a small number of notifications, such as Dobbstein et al.’s “scentifications” using the inScent pendant [24], Dmitrenko et al.’s use of odor for driving-related messages [23], and Grace and Steward’s peppermint scent to alert drowsy drivers to prevent them from falling asleep at the wheel [36]. Similarly, in his master’s dissertation, Kaye [49] talks about smell icons—*smicons*—and proposes a “symbolic” olfactory display that, for example, uses the scent of mint for a rising stock market, and lemon for a falling one. Most recently, Patnaik et al. [58] proposed two six-scent displays—a tabletop and a mobile device worn around the user’s neck—as well as a theoretical framework for information olfaction, along with several examples.

In the face of all these prior efforts, we have yet to come across any work that studies the efficiency of different aspects of scent—or, *olfactory channels*—for different forms of data. To the best of our knowledge, therefore, we believe this is the first study that does for olfactory displays what Bertin [11], Cleveland and McGill [15], and Mackinlay [54] did for visual displays: it ranks olfactory channels for quantitative, ordinal, and nominal data based on empirical findings.

Thermal Feedback

There has been extensive research exploring the utility of thermal stimulation as a feedback channel both for HCI and Virtual Reality (VR) [47, 76]. This ranges from the role users’ perception of thermal cues plays in the identification of objects or in the creation of a more realistic image of an object, to the suitability of thermal cues for encoding information in non-visual situations. Thermal stimuli consists of parameters such as direction of change (warming/cooling), amount of change (intensity), and rate of change, which all contribute to how it is perceived by users [76]. Beyond the ‘yes-no’ detection of thermal stimuli, different combinations of thermal parameters have been shown to be perceptually different and suitable for communicating with users both when presented in isolation, and when combined with other modalities (e.g., visual or auditory stimuli) [1, 76]. They have been demonstrated to be

suitable for communicating both discrete and continuous information. For example, Wettach et al.’s. [74] thermal navigation app guides users to their destinations using different levels of thermal stimuli (the hotter the stimulus, the closer the user is to their destination). Similarly, Wilson et al. [75] used thermal icons as non-visual notifications for text messages. Direction of change (warm and cool) was mapped to message source (‘personal’ and ‘work’), and intensity of change (moderate and strong) was mapped to message importance (‘standard’ and ‘important’) of a text message. Thermal icons were identified with an accuracy of 83%, thus showing the suitability of thermal feedback for communicating discrete data.

In terms of how users perceive and interpret thermal stimulation, in normal atmospheric conditions (approximately 20–40°C), warm, large or fast thermal stimuli changing from the skin’s resting temperature are generally perceived to be stronger and less comfortable than cool, smaller or slower changes [47, 76]. Also the higher the thermal stimuli intensity, the more arousing (level of activation/excitement) and dominant (level of control) users perceive them to be.

Similar to the natural associations between temperature sensations and subjective experiences in language and cognition (e.g., cold and distant or warmth and physical closeness) [43], temperature also has an inherent association with smell, such as the warmth in the smell of cinnamon and the coolness in the smell of mint. This makes thermal stimuli a natural choice for augmenting olfactory feedback by adding an extra thermal parameter to olfactory displays.

OLFACTORY DISPLAY: IMPLEMENTATION

ViScent 2.0 (Figure 1) is an olfactory display for information olfaction [58]. The system is organized into components that control unique sensory channels for encoding information.




Figure 1: The viScent 2.0 system.

System Overview

ViScent 2.0 is a tabletop olfactory display capable of producing a range of olfactory stimuli for information olfaction. More specifically, the viScent olfactory channels include

- **Scent type:** fragrance (e.g., leather, peppermint, coffee);
- **Scent intensity:** perceived air concentration of an odorant (essential oil in solution with water as inert solvent);

- **Airflow rate:** speed of the air carrying the scent; and
- **Air temperature:** temperature of the conveying air.

Each of these channels is managed by a subsystem, four in total, as well as a control system; we describe them below. All components are enclosed in two custom-designed physical modules: the *olfactory display unit* and the *control tower*. The build is a MakerBeam (anodised aluminium beams) framework covered with acrylic plexiglass panels, laser cut into shape. The olfactory display unit houses all component subsystems.

The control tower houses the control system and power supplies. It is responsible for controlling the functioning of all four subsystems. The tower has digital readouts (Figure 2) that display relevant information such as temperature (room, coolant, heating, and cooling core) and power (power and voltage drawn by the heating, cooling, and control systems).



Figure 2: The viScent 2.0 control tower.

Control System

The control system is housed in the uppermost section of the control tower, and modulates the entire system. It is run by an ATmega2560-based microcontroller. The microcontroller interfaces through a USB cable with a computer running Unity-based software controlling the entire system.

Channel: Scent Types

We define scent classes as discrete fragrances as a means of encoding information. We use essential oils diluted with water as the source of the scent. The oil-water mixture is atomized with an ultrasonic transducer controlled by the ATmega2560 microcontroller. The transducer sits on a cork fitted to a glass bottle containing the oil-water mixture. A cotton bud fitted underneath the transducer acts as a channel carrying the scent from the bottle up to the transducer. All bottles are fitted on the display with a custom designed 3D printed housing. We use 6 distinct fragrances for scent type (5 in the experiment).

Channel: Scent Intensity

We define scent intensity as the intensity of a certain fragrance. To define scent intensity quantitatively, we shall use the concept of *volume fraction* to represent the concentration of odorant presented to the user; for a more nuanced discussion of the conversion from odorant concentration to scent intensity than is in scope for the purposes of this paper, a rigorous comparison is presented by Wu et al. [77]. Volume fraction is defined

as volume of a constituent divided by the volume of all the constituents of the mixture prior to mixing. Volume fraction is a dimensionless quantity. For our experiment, we dilute essential oils by mixing them with water. Our system uses 5 levels of dilution producing 5 intensity levels of a certain fragrance (see supplemental material for details). The volume fractions of our essential oil mixtures (for scent intensity) are available in the supplemental material to this paper. These mixtures are then placed in the ultrasonic atomizer pods for atomization.

Channel: Airflow Rate

We define the airflow rate as an olfactory channel which relates to the speed of the air carrying the smell. In our implementation, we measure this as a function of fan speed as determined by voltage delivered to the device's fans; however, we must note that the preferred measure of airflow rate in olfactory literature is volumetric flow rate, which is the product of flow velocity and cross-sectional area [21]. Here, the idea is to evaluate if an increasing air flow rate carrying a certain scent relates to an increasing quantity in a dataset or vice versa. We use 12V brushless DC fans to diffuse the scent vapours towards the user. AirFlow rate is controlled by controlling the fan speed. We use L298N driver to control the fan speed with Pulse-Width Modulation (PWM). The L298N driver is controlled by the ATmega2560 microcontroller. The specific voltages for the fan speeds used in our study are available in the supplemental material.

Channel: Air Temperature

We define the temperature of the air carrying the scent as a channel where the temperature of the air carrying the scent is associated with a quantity from the data. Designing a thermal system to control air temperature is critical as achieving rapid temperature changes can be complicated. Here we present a detailed description of the thermal system design and challenges associated with it.

To achieve a programmable air temperature control, we segment and design the thermal interface in two parts, a heating and a cooling system. This helps us instantly switch between heating and cooling without delays.

Air Heating

We use resistive heating to maintain a heated core over which air flows. A blower produces an air stream through the heated core: drawing-in air from the surrounding, pushing it through the heated core and out through a vent that opens up to the user. A MOSFET (AOD4184A, N Channel) controls the current flowing through the resistive heating core, thereby controlling the heating core temperature. The microcontroller (ATmega2560) interfaces with this MOSFET with PWM. We also attach a temperature sensor adjacent to the heating core to monitor the temperature. A L298N driver controls the blower fan, thereby controlling airflow rate. We optimize the airflow rate and the resistive heating to obtain optimal warm air.

Air Cooling

The air cooling system is one of the most complex systems employed in this prototype. We use thermoelectric cooling to maintain a cooling core at subzero temperature. Mirroring the

heating system, a blower produces an air stream through the cooled core: drawing-in air from the surrounding, pushing it through the cooled core and out through a vent that opens up to the user. We use four thermoelectric modules (TEC1-12706) attached underneath an aluminum heat exchanger that acts as the cooling core. These modules sit on top of an aluminum liquid cooling block. The modules are sandwiched between the aluminum heat exchanger (cooling core) and an aluminum liquid cooling block with thermal adhesive. On supplying power, the thermoelectric modules act as heat pumps pulling heat from the heat exchangers/cooling core to the other side interfacing with the aluminum liquid cooling block. This rapidly cools down the cooling core thereby heating up the aluminum liquid cooling block. We circulate a coolant (XSPC EC6: a high performance, high thermal conductivity coolant) through the aluminium liquid cooling block and an aluminium heat exchanger that sits outside the olfactory display on the control tower. This coolant transfers the heat produced by the thermoelectric modules to a heat exchangers. Three 12V DC cooling fans create a steady stream of air flow through the heat exchangers to bring about efficient heat transfer. The coolant is stored in a coolant reservoir connected to a pump that does the circulation. Two temperature sensors are connected to this cooling system, one placed on the cooling core to monitor the cooling core temperature and the other dipped inside the coolant in the coolant reservoir tube to monitor the temperature of the coolant. The blower fan is controlled by a L298N driver interfaced with the ATmega2560 microcontroller to control cool air flow. The thermoelectric modules, the cooling core and the aluminum liquid cooling block are covered with thermal insulation to have maximum efficiency.

METHOD

We conducted a perceptual experiment evaluating the utility of scent for conveying abstract information. In doing so, we followed the analogy of past empirical work on graphical perception such as that catalogued by Cleveland and McGill [15], Mackinlay [54], and Bertin [11]. Similar to these studies, the ultimate purpose of our study was to determine an internal ranking between sensory channels and different types of data: quantitative, ordinal, and nominal [66].

We conducted pilot tests to find noticeable differences for all conditions (see below). Our study, including methods, design, and predictions, was preregistered¹ prior to collecting data.

Apparatus

We conducted our study using the viScent 2.0 device as the olfactory display (Section 3). The device was connected to a laptop computer running Microsoft Windows 10. The laptop ran the Unity-based viScent control system, as well as an automated testing framework implemented using the viScent API. Instead of the laptop display, we used a 55-inch display with a resolution of 1920×1080 pixels.

The study was conducted in an isolated laboratory space. The viScent tabletop display was arranged between the participant and the display in a position so that it would not obstruct the screen, yet was still at a comfortably distance from the

¹<https://osf.io/grdk7/>

user's face. Participants wore noise-canceling ear protection during the experiment to minimize confounds from ambient noise or sound from the olfactory display. Box and stand fans were used during experiments to maintain air circulation. Furthermore, the space was thoroughly aired out between sessions to eliminate vestigial scents.

The scent configuration was designed specifically for the experiment. For scent intensity, we used five bottles of different intensities of mango (see Section 3.4). For scent type, we used five scents based on cognitive psychology empirical work by Castro et al. [13] as it relates to the framework introduced by Patnaik et al [58]: leather, orange, peppermint, coffee, and pear. Each scent was represented in three different intensities: the low, mid, and high volume fractions from above. The remaining 4 bottles were not used during the experiment.

Participants

We recruited 20 paid participants (12 identified as male, 8 as female) for our experiment. Participant ages ranged from 22 to 30 years. All participants were university students and had a basic knowledge of data and statistics. No participant reported olfactory dysfunction, and we screened participants against allergies to any essential oil used in the experiment both during recruitment as well as during informed consent prior to the experiment. No participant reported discomfort.

Experimental Factors

We involved the following two factors in our experiment:

- **Sensory Channel (SC):** The sensory aspects used to convey data. We studied the following four channels:
 - SCENT TYPE (*S*): Using one of five scents to convey data (leather, orange, peppermint, coffee, and pear).
 - SCENT INTENSITY (*I*): The concentration of mango (five fractions, see Section 3.4) used to convey data.
 - AIRFLOW RATE (*A*): The speed of the air (i.e., wind) delivering the scent (conveyed using fan voltages).
 - TEMPERATURE (*T*): Temperature of air delivering the scent (one cooling, one neutral, and three heating).
- **Data Type (DT):** The specific type of data being conveyed using scent [66]. Informed by Mackinlay's three-part ranking [54], we study three data types:
 - QUANTITATIVE (*Q*): Numbered items that support all arithmetic operations (a combination of interval-scale and ratio-scale levels [66]).
 - ORDINAL (*O*): Labeled items that support rank order, but not relative degree of difference between items.
 - NOMINAL (*N*): Categorical items that differentiate only on their names or identifiers.

Tasks and Stimuli

The experiment involved a single task—identifying a data item conveyed using scent—with different instantiations depending on the data type *DT*. For all tasks, the screen showed a visual representation of the data type (Figure 3):

- **Quantitative sensing task (T_Q):** Recover a number encoded using the sensory channel. *Display:* A slider with a continuous color scale.




Figure 3: Examples of each task (one per data type DT). Participants select the value conveyed using scent and then click the “Submit” button. This is followed by a dialog asking for the participant’s confidence on a 5-level Likert scale.

- **Ordinal sensing task (T_O):** Recover an ordered data item encoded using the sensory channel. *Display:* A slider with a five-segment color scale.
- **Nominal sensing task (T_N):** Recover a nominal data item encoded using the sensory channel. *Display:* An unordered list of checkboxes.

Similarly, the mapping from data values to scent differed depending on which sensory channel SC was used (nominal data was assigned in random order):

- **Scent Type:** Items were assigned to scents depending on their position in the range of possible values.
- **Scent Intensity:** ascending items (if ordered) assigned to ascending scent intensities (essential oil saturations).
- **Airflow Rate:** ascending items (if ordered) assigned to ascending fan voltages (yielding increasing airflow).
- **Temperature:** ascending items (if ordered) assigned to increasing temperature (cooling and heating).

For ordinal and nominal data types, the data range for tasks was five distinct values for each sensory channel, which translated to five different scents for S , five different scent intensities for I (Section 3.4; we did not use a 0% intensity as the absence of smell is not a reliable signal), five different airflow rates for A (five distinct voltage values to the fans), and five different temperatures for T (one cooling, one neutral, and three heating settings with increasing voltage).

For quantitative data, we used continuous voltage for the airflow rate R . However, since scent type S and scent intensity I rely on discrete bottles where the diffuser can only be turned on or off, we had to blend bottles to generate additional smells to carry more than five values. Blending scents is a non-linear process [58], so more research is needed here.

With this caveat in mind, we generated additional scent types S by blending the three different intensities of the five scents used so that any value between a scent type S_A and scent type S_B was subdivided into three regions using a blend of scent intensities (H, M, L for high, medium, low) as follows: $[0, 0.17) \rightarrow (H \times S_A, 0)$, $[0.17, 0.5) \rightarrow (M \times S_A, L \times S_B)$, $[0.5, 0.83) \rightarrow (L \times S_A, M \times S_B)$, and $[0.83, 1) \rightarrow (0, H \times S_B)$. This yielded a total of 13 unique scent blends.

Five bottles yielded $2^5 = 32$ distinct intensities. To represent a value, the value was normalized to the range $[0, 32)$, converted to binary, and used to activate the corresponding bottles 0–4.

Each experimental condition ($SC \times DT$) was repeated three times. Prior to each block of three repetitions, participants were given a tutorial where they were given the “sensory legend” that corresponded to the visual display. For example, for scent type S , the participant would get to smell each scent as its associated value on screen was highlighted, e.g., that a lemon scent corresponded to “Volvo.” A visual label persisted on the screen showing this scent-to-data mapping throughout the block of repetitions, but the sensory legend was not repeated again. In the example above, the olfactory label “lemon” would be placed under the data label “Volvo.”

Experimental Design

We used a within-participants factorial design where each participant was exposed to trials for all sensory channels and data types. This yielded the below design, the order of each experimental condition $SC \times DT$ randomized to counterbalance systemic effects of practice:

	4	Sensory Channels SC (S, I, A, T)
×	3	Data Types DT (Q, O, N)
×	3	repetitions
	36	trials per participant

For each trial, we collected the accuracy (both whether the answer was correct, and for ordinal and quantitative data, the normalized distance from the correct answer), the completion time, and the Likert-scale confidence rating. The completion time was measured from the beginning of a trial until the end of the 9-second habituation period (see below) or when an answer was submitted, whichever was shorter.

Procedure

Upon arriving at a session, participants were first given informed consent in an antechamber outside the laboratory space. The purpose here was to screen for allergies to essential oils prior to entering an area that could be potentially hazardous to a person with allergies.

After giving consent, the participant was allowed to enter the laboratory space and was given a brief explanation of the purpose of the study. The experimenter demonstrated the

olfactory display and the testing framework. The participant was allowed to train on several example trials using different sensory channels. Timed trials began once the participant indicated they felt comfortable to proceed.

Each block of experimental conditions $SC \times DT$ began with the above tutorial, during which the sensory legend was displayed. This was followed by the three repetitions, each with a new random data value to sense. The same visual legend persisted during the entire block of three trials. During a trial, sensory stimulus was active for a total of 9 seconds. This corresponds to the typical sensory habituation period of the human olfactory system [58]. After this period, a visual feedback indicated that the stimulus was no longer active. Participants were not able to repeat the stimulus.

After submitting a data value corresponding to the sensory stimulus, the software would pop up a dialog box polling the participant about their level of confidence in their answer on a 5-point Likert scale. This was followed by a blank screen during which a participant could rest between trials.

Once all trials had been completed, the participant was given an exit survey. They were then compensated \$10 for their participation. A typical session lasted between 50 to 60 minutes; no single session lasted more than one hour.

Predictions

We formulate the following basic predictions and our corresponding motivation about our experiment. We want to emphasize, however, that the goal of this paper is not to accept or reject hypotheses, but rather to derive rankings of sensory channels for different data types.

- P1 *Participants will be significantly **more accurate** when sensing nominal (N) data using scent type (S) than all other sensory channels.* Distinct scents lend themselves to distinguishing between discrete sets of data items.
- P2 *Participants will be significantly **more accurate** when sensing ordinal (O) and quantitative (Q) data using scent intensity (I) than all other sensory channels.* Our sensory systems are sensitive to intensity, and its increasing nature fits ordered data types.
- P3 *Participants will be significantly **less accurate** when sensing ordinal (O) and quantitative (Q) data using scent type (S) than all other sensory channels.* Distinct smells are ill-suited to representing ordered data (cf. P1).

RESULTS

Here we review the results from our study, starting with an overview and then organizing findings into the three data types: nominal, ordinal, and quantitative. For each data type, we will discuss accuracy/error and completion time. The reason we slice our results by data type first is that we are not primarily trying to compare different types, but rather to derive an internal ranking within each type (similar to Mackinlay's ranking [54]). We also report on participants' subjective feedback.

Following our preregistration, we use bootstrapping ($N = 1,000$ iterations) on repetitions aggregated by simple mean

for each participant by olfactory channel and data type to calculate 95% confidence intervals, and then analyze the results graphically [26].

Overview

Figure 4 shows a summary of the three dependent measures arranged by data type DT . As stated above, ranking between different data types is of no real consequence to our study, so we will not discuss this data further other than to say that there appears to be little support for claiming that any of the data types N , O , and Q exhibits divergent performance compared to any other data type. The only exception may be that the small overlap between CIs (Figure 4c) suggests that participants were faster for trials with nominal than with ordinal data.

We also study the perceived confidence rating given by participants after each trial in Figure 5. The fact that all confidence intervals are well above the neutral is another indication that our experiment is a success, at least as perceived by the participants themselves. Results per data type in Figure 5a seem very similar, and there is little evidence to suggest that participants expressed different confidence ratings for different data types. The confidence ratings in Figure 5b are more divergent based on the sensory channel. In general, there is some evidence to suggest that temperature T was the channel that participants felt most confident about. Certainly, participants appear to rate their confidence for temperature T as stronger than for intensity I and airflow rate A . Furthermore, the confidence intervals for scent type S and scent intensity I (both olfactory) are larger than the other two (tactile) channels, which may indicate that tactile sensing is more accurate than olfactory sensing. Other pairwise comparisons are more difficult to assess.

Figure 6 summarizes error (distance from the correct value), correctness (ratio of participant's answers that were exactly correct), and completion time. We refer to this figure below.

Nominal Data

The top row of plots in Figure 6 represent participant performance for nominal data N . Note that the error metric is not applicable to nominal data, as there is no distance property for nominal data (it only supports equality). Correctness, however, is defined as whether the participant's response exactly matched the stimulus, and taken in aggregate it represents the ratio of trials that were exactly correct. The plot in Figure 6b depicts 95% confidence intervals that are rather large, indicating that this was a challenging experimental condition. Since there is significant overlap between CIs, there is only a moderately strong ordering between sensory channels SC .

For completion times (Figure 6c), which may be less important for our ranking except to provide context, the spread is smaller. Temperature T , intensity I , and scent type S appear to yield similar completion times, with a slight advantage for scent type, but they are all outperformed by airflow A .

Ordinal Data

Data collected for trials involving the ordinal data type O are shown in the middle row of Figure 6. The error rate (Figure 6d) here exhibited relatively high spread, with trials that used intensity I seemingly resulting in higher error than




Figure 4: Results (error, correctness, and completion time) for each data type. Error bars show 95% confidence intervals, dots show means. Nominal data N has no error since the data type only supports identity, not distance. Quantitative data Q has no correctness as it represents a continuous input range, and thus providing the exactly correct answer is not a relevant measure.




Figure 5: Perceived confidence on a Likert scale ranging from -2 (uncertain), through 0 (somewhat), to 2 (certain).

the other three channels. Scent type S and airflow A are close in performance, but the correctness plot (Figure 6e), where scent has a clear advantage, shows that scent wins out. Intensity, however, appears to be outperformed by the other channels even in this plot. Airflow A performs slightly better than temperature T for both error and correctness.

As for completion time (Figure 6f), the data suggests that airflow rate A required less time for trial completion than for scent S , temperature T , and scent intensity I .

Quantitative Data

The bottom row in Figure 6 shows data for the quantitative data type Q . In this case, we do not plot correctness, as the quantitative data trials asked participants to answer using a continuous data scale. It is rather unlikely that participants would be able to answer the exact correct value being conveyed using the sensory stimulus, so instead we rely on the distance from the correct value (e.g., the error) as the accuracy metric. Studying this error metric (Figure 6g), there is ample evidence that temperature T was the most accurate sensory channel for perceiving quantitative data. The data also suggests that the airflow rate A is moderately more accurate than the scent type S , and that they both are more accurate than scent intensity I .

Completion times in Figure 6i seem to indicate that airflow rate A is faster than scent type S and temperature T , and possibly even scent intensity I . Scent intensity also appears to exhibit shorter completion times than both scent and temperature.

Subjective Feedback

None of the participants reported ever having used an olfactory display in the past; in fact, many were intrigued by the concept and volunteered for the study mainly to experience it. Several

expressed curiosity about applications of our work; “*I look forward to see how you will implement this in real life.*”

In practice, participants spent approximately 45 minutes on each session. While all participants who begun the experiment also completed it, several noted that they felt saturated at the end, their ability to smell diminished. However, we saw no indication of this in our analysis. Nevertheless, participants expressed some surprise in the level of difficulty in the trials; said one participant, “*this was a lot harder than I thought.*” This may have arisen from the high granularity expected of participants, where some noted that they were easily able to discern the “big picture,” but not minute details.

We also conducted a Likert questionnaire of participants after the experiment to evaluate their experience of the display system (Figure 7). Participants were asked if they agreed with statements of positive sentiment regarding ease of adoption, enjoyment, ease of use, learning curve, and efficiency.

DISCUSSION

Our results on the error, correctness, and completion time for sensing information-carrying stimulus begin to show the utility of information olfaction. Below we summarize these results after discussing the benefits of encoding data with smell.

The Perks of Smelling a Wallflower

While prior work by Patnaik et al. [58] provides a more detailed examination of contexts in which scent may be preferable to vision for evaluating information, it is worth briefly enumerating why this is a good idea in the first place. Put succinctly, olfactory encoding is useful in situations when a person cannot *look* (i.e., their eyes are busy elsewhere, such as when driving, or for ambient displays in an office), or cannot *see* (i.e., data analysis by people with low vision or blindness). It also offers benefits in representing information relevant to the gradual buildup of findings as part of the sensemaking process by encoding “slow-moving” data as an ambient feature of the environment [32, 49].

The virtues of ambient information displays have been amply discussed elsewhere [27, 56, 61]; for example, we see data olfaction as a useful addition to an ambient information environment such as the ambientROOM [45]. We note that the latter situation—*accessible data visualization*—is receiving increasing interest in the visualization community [53] We also find it worth highlighting that, in our own results, we find that the self-reported user experience of information olfaction

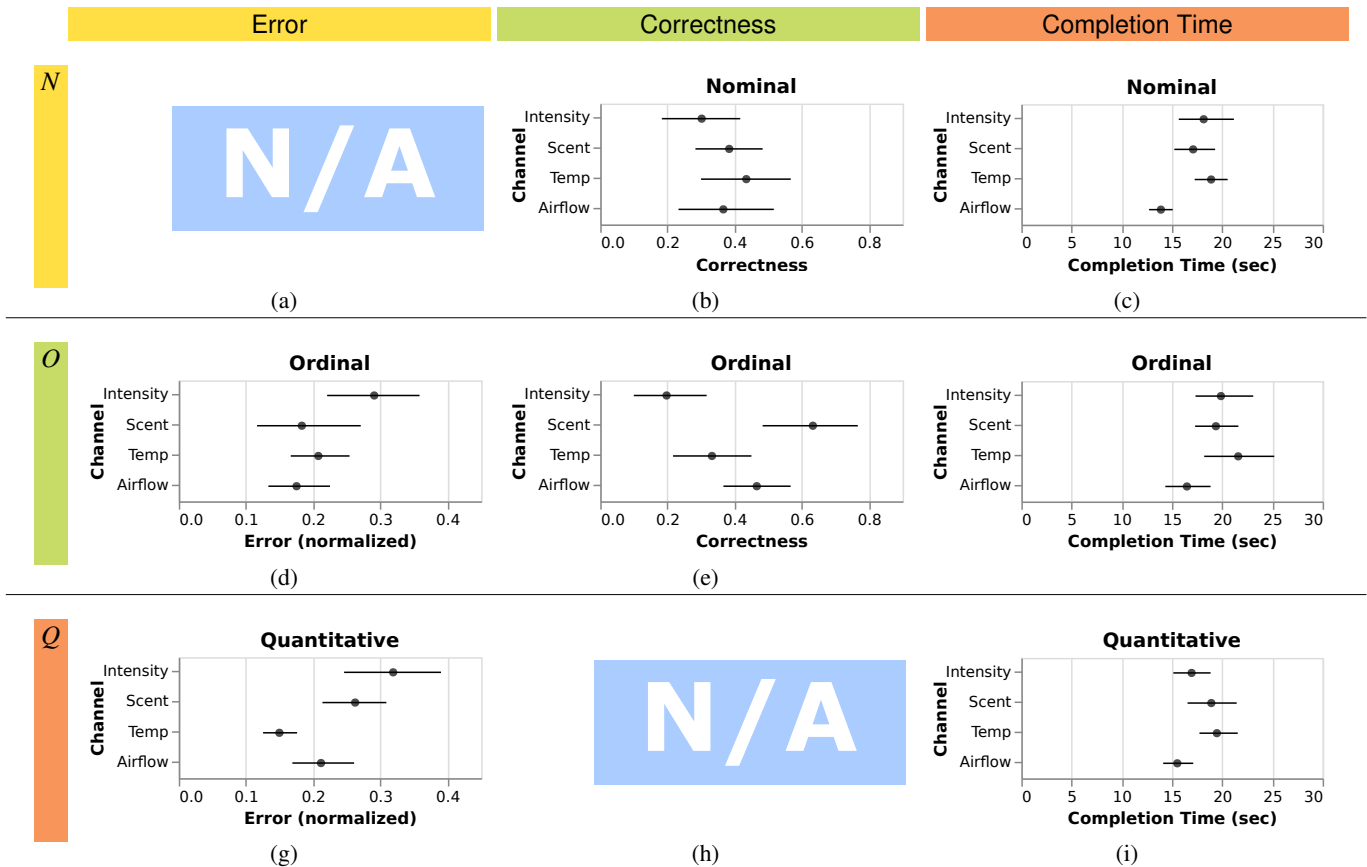


Figure 6: Error (top row), correctness (middle), and completion times (bottom) organized by data type DT and sensory channel SC . Error bars show 95% confidence intervals, and dots show means. Plots marked N/A are those where the metric is undefined.




Figure 7: Perceived confidence (Likert scale) from “Strongly Disagree”, through “Neutral”, to “Strongly Agree”.

display systems are generally quite positive (Figure 7). This work is an important entry into that conversation. Blind or not, all of us navigate a spatial world and need to access data.

Quantitative Data

Synthesizing the above findings, we now derive a ranking of sensory channels for quantitative data. Just like Mackinlay’s ranking [54], the below ranking is not entirely based on empirical findings; rather, it is *supported* by empirical data. For each channel, we give a brief motivation for the ranking:

Q1. TEMPERATURE (T), on account of being most accurate;

Q2. AIRFLOW RATE (A), on account of being slightly more accurate and faster than scent type;

Q3. SCENT TYPE (S), on account of being more accurate than intensity; and

Q4. SCENT INTENSITY (I), on account of being the least accurate sensory channel for quantitative data.

Ordinal Data

Scent type outperformed all other channels in error and correctness for ordinal data, indicating that the sharp distinction between scents aided in solidifying the users’ perception of the boundaries between values along a discrete scale.

O1. SCENT TYPE (S), on account of being the most accurate (comparable in error, better in correctness);

O2. AIRFLOW RATE (A), on account of being slightly more accurate and faster than temperature;

O3. TEMPERATURE (T), on account of being more accurate (both in error and correctness) than intensity; and

O4. SCENT INTENSITY (I), on account of being the least accurate sensory channel for ordinal data.

Nominal Data

The results for nominal data are inconclusive for **P1**: there is no clear evidence that scent type is better than temperature for

nominal data (in fact, it may be worse). However, scent type exhibited a shorter completion time, thus yielding our ranking:

- N1. SCENT TYPE (S), on account of tied best accuracy while being the faster;
- N2. TEMPERATURE (T), on account of tied best accuracy;
- N3. AIRFLOW RATE (A), on account of being the fastest;
- N4. SCENT INTENSITY (I), on account of being the least accurate sensory channel for nominal data.

Smelling Least → Smelling Best

One of the more surprising findings from our study was that scent intensity was outperformed by basically all other sensory channels for all data types (the plots in Figure 6 give the detailed results). This clearly disproved our prediction **P2**, which was the exact opposite. However, this phenomenon has an explanation in the literature.

In psychophysics, the quantitative study of physical stimuli and the sensations and perceptions they produce, the concept of *sensory scaling* in assigning perceived numbers to sensory experiences is well-known [50]. Sensory experiences are subjective, and building a personalized scale for specific senses is a time-consuming process based on past experience and exposure. What one person ranks as strong stimulus—say, a 9 on a scale of 1 to 10 commonly referred to as the Labeled Magnitude Scale (LMS) [37]—may merely rate as moderate for someone else, e.g., a 6. Furthermore, some people merely have a higher sensory range than others; for example, so-called “supertasters” [8] experience taste with far greater intensity than others. Also, intensities are modified by their context; for example, a word such as “large” or “small” all depend on the noun it describes. This is why Stevens [68] can give the following example without ambiguity: “*Mice may be called large or small, and so may elephants, and it is quite understandable when someone says it was a large mouse that ran up the trunk of the small elephant.*”

To address this, psychophysics researchers have introduced the so-called “general” Labeled Magnitude Scale (gLMS) [9] where instead of labeling the rungs on the scale using the same specific sense—e.g., “10 is the most intense smell you have ever experienced”—the scale is labeled using the strongest imaginable sensation of any kind, i.e., not restricted to the specific sensory channel. This begins to address the personalized concern, but arguably still makes for a subjective scale.

Nevertheless, mitigating the scaling problem takes time, and scent intensity is typically untrained for most people. Since our goal was to empirically understand information olfaction with participants representative of the general population, we did not provide any extended training in the intensity tutorial. Furthermore, the nature of our experiment precluded us from leveraging the gLMS scale since all trial blocks were preceded by an “sensory legend.” However, we did base our scent intensity on the power law of psychophysics [67] (Section 3.4).

Limitations

Perhaps the most significant limitation to our study is the possibility of olfactory contamination. To combat this, we

designed our olfactory display to physically separate scents to avoid direct contamination. We also conducted informal tests of our system prior to conducting our user study in order to verify that the stimuli were being presented to the participant as expected. However, many of the adjustments made during this period were based on the subjective sensory perception of olfactory stimuli by the authors. A more robust validation step would involve either capturing instrument readings from positions relative to the olfactory display corresponding to the position of the participants’ noses, or inserting validation trials in the user study in which participants are asked explicitly to distinguish between scents.

Another limitation, implicit in both our study design and Patnaik et al.’s [58] theoretical model, is that we assume, based on existing empirical work, that olfaction is cross-modal. We do not evaluate how effective temperature or airflow rate are without an odorant. While we supervised the participants during the study to ensure that the thermal and airflow channel modifications were centered on their faces, there is still room for improving the impact these features have on olfaction.

As noted earlier, our study involves both olfactory (scent and intensity) and tactile (airflow and temperature) channels rather than olfactory channels alone. However, this is also consistent with our pragmatic philosophy to information olfaction; we are not merely interested in the information-carrying capacity of smell alone, but rather what information can be conveyed using a typical olfactory display. This philosophy is also consistent with the human-computer interaction audience.

Finally, our study was a laboratory study, which limits the pool of potential participants. As a result, our study included 20 participants; not a large number of participants by any account. With that said, to quote Dragicevic [26], “there is no magic number of participants.” We are confident in the number of participants and the validity of our results.

CONCLUSION AND FUTURE WORK

We put the theory of information olfaction to practical test by empirically evaluating the olfactory perception of information. While our efforts mirror seminal work by Mackinlay [54] and Cleveland and McGill [15], it is the first study of its kind.

The disappointing results for odor intensity for all data types, as well as for scent type in encoding quantitative and nominal data, warrants further exploration. While we believe our hardware implementation accurately conveyed these signals to the user, it is still an approximation simulating the desired stimulus. The possibility remains that our participants may not be perceiving the scent intensity with as granular a level of detail as is required for the task at hand. Refining the granularity of detail in the level of intensity in scent for presenting users with abstract information is an open area of research.

Acknowledgments

This work was partially supported by the U.S. National Science Foundation award IIS-1539534. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the funding agency. We thank Steve Franconeri for inspiring our title.

REFERENCES

- [1] Moses Akazue, Martin Halvey, Lynne Baillie, and Stephen Brewster. 2016. The Effect of Thermal Stimuli on the Emotional Perception of Images. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 4401–4412. DOI : <http://dx.doi.org/10.1145/2858036.2858307>
- [2] Christopher Andrews, Alex Endert, Beth Yost, and Chris North. 2011. Information visualization on large, high-resolution displays: Issues, challenges, and opportunities. *Information Visualization* 10, 4 (2011), 341–355. DOI : <http://dx.doi.org/10.1177/1473871611415997>
- [3] Sriram Karthik Badam, Fereshteh Amini, Niklas Elmqvist, and Pourang Irani. 2016. Supporting Visual Exploration for Multiple Users in Large Display Environments. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*. IEEE, Piscataway, NJ, USA, 1–10. DOI : <http://dx.doi.org/10.1109/vast.2016.7883506>
- [4] Sriram Karthik Badam and Niklas Elmqvist. 2014. PolyChrome: A Cross-Device Framework for Collaborative Web Visualization. In *Proceedings of the ACM Conference on Interactive Tabletops and Surfaces*. ACM, New York, NY, USA, 109–118. DOI : <http://dx.doi.org/10.1145/2669485.2669518>
- [5] Sriram Karthik Badam, Eli Fisher, and Niklas Elmqvist. 2015. Munin: A Peer-to-Peer Middleware for Ubiquitous Analytics and Visualization Spaces. *IEEE Transactions on Visualization and Computer Graphics* 21, 2 (Feb 2015), 215–228. DOI : <http://dx.doi.org/10.1109/TVCG.2014.2337337>
- [6] Sriram Karthik Badam, Andreas Mathisen, Roman Rädle, Clemens N. Klokrose, and Niklas Elmqvist. 2019. Vistrates: A Component Model for Ubiquitous Analytics. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 586–596. DOI : <http://dx.doi.org/10.1109/TVCG.2018.2865144>
- [7] Robert Ball and Chris North. 2007. Realizing embodied interaction for visual analytics through large displays. *Computers & Graphics* 31, 3 (2007), 380–400. DOI : <http://dx.doi.org/10.1016/j.cag.2007.01.029>
- [8] Linda M. Bartoshuk. 1991. Sweetness: history, preference, and genetic variability. *Food Technology* 45, 11 (1991), 108–113.
- [9] Linda M. Bartoshuk, Valerie B. Duffy, Barry G. Green, Howard J. Hoffman, Cia-Wen Ko, Laurie A. Lucchina, Lawrence E. Marks, Derek J. Snyder, and James M. Weiffenbach. 2004. Valid across-group comparisons with labeled scales: the gLMS versus magnitude matching. *Physiology & Behavior* 82, 1 (2004), 109–114. DOI : <http://dx.doi.org/10.1016/j.physbeh.2004.02.033>
- [10] Dominikus Baur, Bongshin Lee, and Sheelagh Cappendale. 2012. TouchWave: kinetic multi-touch manipulation for hierarchical stacked graphs. In *Proceedings of the ACM Conference on Interactive Tabletops and Surfaces*. ACM, New York, NY, USA, 255–264. DOI : <http://dx.doi.org/10.1145/2396636.2396675>
- [11] Jacques Bertin. 1967. *Sémiologie graphique*. Mouton/Gauthier-Villars, Paris, France.
- [12] Stuart K. Card, Jock D. Mackinlay, and Ben Shneiderman. 1999. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers, San Francisco, CA, USA.
- [13] Jason B. Castro, Arvind Ramanathan, and Chakra S. Chennubhotla. 2013. Categorical Dimensions of Human Odor Descriptor Space Revealed by Non-Negative Matrix Factorization. *PLOS ONE* 8, 9 (Sep 2013), 1–16. DOI : <http://dx.doi.org/10.1371/journal.pone.0073289>
- [14] Yang Chen. 2017. Visualizing Large Time-series Data on Very Small Screens. In *Short Paper Proceedings of the IEEE VGTC/Eurographics Conference on Visualization*. Eurographics Association, Geneva, Switzerland, 37–41. DOI : <http://dx.doi.org/10.2312/eurovisshort.20171130>
- [15] William S. Cleveland and Robert McGill. 1984. Graphical Perception: Theory, Experimentation and Application to the Development of Graphical Methods. *J. Amer. Statist. Assoc.* 79, 387 (Sept. 1984), 531–554. DOI : <http://dx.doi.org/10.2307/2288400>
- [16] Albert M. Cook and Janice Miller Polgar. 2015. *Assistive Technologies: Principles and Practice* (4th ed.). Mosby, St Louis, MO, USA.
- [17] Maxime Cordeil, Andrew Cunningham, Tim Dwyer, Bruce H. Thomas, and Kim Marriott. 2017a. ImAxes: Immersive Axes as Embodied Affordances for Interactive Multivariate Data Visualisation. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, USA, 71–83. DOI : <http://dx.doi.org/10.1145/3126594.3126613>
- [18] Maxime Cordeil, Tim Dwyer, Karsten Klein, Bireswar Laha, Kim Marriott, and Bruce H. Thomas. 2017b. Immersive collaborative analysis of network connectivity: CAVE-style or head-mounted display? *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 441–450. DOI : <http://dx.doi.org/10.1109/TVCG.2016.2599107>
- [19] Frederick E. Croxton and Harold Stein. 1932. Graphic Comparisons by Bars, Squares, Circles, and Cubes. *J. Amer. Statist. Assoc.* 27, 177 (1932), 54–60. DOI : <http://dx.doi.org/10.2307/2277880>
- [20] Frederick E. Croxton and Roy E. Stryker. 1927. Bar charts versus circle diagrams. *J. Amer. Statist. Assoc.* 22, 160 (1927), 473–482. DOI : <http://dx.doi.org/10.2307/2276829>

- [21] Fabrizio Davide, Martin Holmberg, and Ingemar Lundström. 2001. Virtual olfactory interfaces: Electronic noses and olfactory displays. In *Communications Through Virtual Technology: Identity, Community, and Technology in the Internet Age*, G. Riva and Fabrizio Davide (Eds.). IOS Press, Amsterdam, The Netherlands, 193–220.
- [22] Regina Araujo de Almeida (Vasconcellos) and Bruce Tsuji. 2005. Interactive mapping for people who are blind or visually impaired. In *Cybercartography*, D. R. Fraser Taylor (Ed.). Modern Cartography Series, Vol. 4. Academic Press, Amsterdam, The Netherlands, 411–431. DOI: [http://dx.doi.org/10.1016/S1363-0814\(05\)80021-8](http://dx.doi.org/10.1016/S1363-0814(05)80021-8)
- [23] Dmitrijs Dmitrenko, Emanuela Maggioni, and Marianna Obrist. 2017. OSpace: Towards a Systematic Exploration of Olfactory Interaction Spaces. In *Proceedings of the ACM Conference on Interactive Surfaces and Spaces*. ACM, New York, NY, USA, 171–180. DOI: <http://dx.doi.org/10.1145/3132272.3134121>
- [24] David Dobbstein, Steffen Herrdum, and Enrico Rukzio. 2017. inScent: A Wearable Olfactory Display as an Amplification for Mobile Notifications. In *Proceedings of the ACM International Symposium on Wearable Computers*. ACM, New York, NY, USA, 130–137. DOI: <http://dx.doi.org/10.1145/3123021.3123035>
- [25] Jakub Dostal, Uta Hinrichs, Per Ola Kristensson, and Aaron J. Quigley. 2014. SpiderEyes: designing attention- and proximity-aware collaborative interfaces for wall-sized displays. In *Proceedings of the ACM Conference on Intelligent User Interfaces*. ACM, New York, NY, USA, 143–152. DOI: <http://dx.doi.org/10.1145/2557500>
- [26] Pierre Dragicevic. 2016. Fair Statistical Communication in HCI. In *Modern Statistical Methods for HCI*, Judy Robertson and Maurits Kaptein (Eds.). Springer, Boston, MA, USA, 291–330. DOI: http://dx.doi.org/10.1007/978-3-319-26633-6_13
- [27] James R. Eagan and John T. Stasko. 2008. The Buzz: supporting user tailorability in awareness applications. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1729–1738. DOI: <http://dx.doi.org/10.1145/1357054.1357324>
- [28] Walter C. Eells. 1926. The relative merits of circles and bars for representing component parts. *J. Amer. Statist. Assoc.* 21, 154 (1926), 119–132. DOI: <http://dx.doi.org/10.2307/2277140>
- [29] Niklas Elmqvist and Pourang Irani. 2013. Ubiquitous Analytics: Interacting with Big Data Anywhere, Anytime. *IEEE Computer* 46, 4 (2013), 86–89. DOI: <http://dx.doi.org/10.1109/mc.2013.147>
- [30] Niklas Elmqvist, Andrew Vande Moere, Hans-Christian Jetter, Daniel Cernea, Harald Reiterer, and TJ Jankun-Kelly. 2011. Fluid interaction for information visualization. *Information Visualization* 10, 4 (2011), 327–340. DOI: <http://dx.doi.org/10.1177/1473871611413180>
- [31] Leo Ferres, Gitte Lindgaard, Livia Sumegi, and Bruce Tsuji. 2013. Evaluating a tool for improving accessibility to charts and graphs. *ACM Transactions on Computer-Human Interaction* 20, 5 (2013), 28. DOI: <http://dx.doi.org/10.1145/2533682.2533683>
- [32] Avery N. Gilbert, Robyn Martin, and Sarah E. Kemp. 1996. Cross-Modal Correspondence between Vision and Olfaction: The Color of Smells. *The American Journal of Psychology* 109, 3 (1996), 335–351. <http://www.jstor.org/stable/1423010>
- [33] E. Bruce Goldstein. 2015. *Cognitive Psychology: Connecting Mind, Research and Everyday Experience* (4th ed.). Cengage Learning, Stamford, CT, USA.
- [34] E. Bruce Goldstein and James R. Brockmole. 2017. *Sensation and Perception* (10th ed.). Cengage Learning, Stamford, CT, USA.
- [35] Cagatay Goncu, Anuradha Madugalla, Simone Marinai, and Kim Marriott. 2015. Accessible on-line floor plans. In *Proceedings of the International Conference on the World Wide Web*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 388–398. DOI: <http://dx.doi.org/10.1145/2736277.2741660>
- [36] Richard Grace and Sonya Steward. 2001. Drowsy driver monitor and warning system. In *Proceedings of the International Driving Symposium on Human Factors in Driver Assessment, Training and Vehicle Design*. Public Policy Center of Iowa, Iowa City, IA, USA, 64–69. DOI: <http://dx.doi.org/10.17077/drivingassessment.1010>
- [37] Barry G. Green, Gregory S. Shaffer, and Magdalena M. Gilmore. 1993. Derivation and evaluation of a semantic scale of oral sensation magnitude with apparent ratio properties. *Chemical Senses* 18, 6 (12 1993), 683–702. DOI: <http://dx.doi.org/10.1093/chemse/18.6.683>
- [38] Surina Hariri, Nur Ain Mustafa, Kasun Karunanayaka, and Adrian David Cheok. 2016. Electrical Stimulation of Olfactory Receptors for Digitizing Smell. In *Proceedings of the ACM Workshop on Multimodal Virtual and Augmented Reality*. ACM, New York, NY, USA, 4:1–4:4. DOI: <http://dx.doi.org/10.1145/3001959.3001964>
- [39] Morton L. Heilig. 1962. Sensorama simulator. (Aug 1962). US Patent 3050870.
- [40] Leona Holloway, Kim Marriott, and Matthew Butler. 2018. Accessible Maps for the Blind: Comparing 3D Printed Models with Tactile Graphics. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 198:1–198:13. DOI: <http://dx.doi.org/10.1145/3173574.3173772>

- [41] Tom Horak, Sriram Karthik Badam, Niklas Elmqvist, and Raimund Dachsel. 2018. When David Meets Goliath: Combining Smartwatches with a Large Vertical Display for Visual Data Exploration. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 19:1–19:13. DOI: <http://dx.doi.org/10.1145/3173574.3173593>
- [42] Tom Horak, Andreas Mathisen, Clemens Nylandsted Klokmose, Raimund Dachsel, and Niklas Elmqvist. 2019. Vistribute: Distributing Interactive Visualizations in Dynamic Multi-Device Setups. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 19:1–19:13. DOI: <http://dx.doi.org/10.1145/3173574.3173593>
- [43] Hans IJzerman and Gün R. Semin. 2009. The Thermometer of Social Relations: Mapping Social Proximity on Temperature. *Psychological Science* 20, 10 (2009), 1214–1220. DOI: <http://dx.doi.org/10.1111/j.1467-9280.2009.02434.x>
- [44] Petra Isenberg and Sheelagh Carpendale. 2007. Interactive Tree Comparison for Co-located Collaborative Information Visualization. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1232–1239. DOI: <http://dx.doi.org/10.1109/TVCG.2007.70568>
- [45] Hiroshi Ishii, Craig Wisneski, Scott Brave, Andrew Dahley, Matt Gorbet, Brygg Ullmer, and Paul Yarin. 1998. ambientROOM: Integrating Ambient Media with Architectural Space. In *Conference Summary on ACM Human Factors in Computing Systems*. ACM, New York, NY, USA, 173–174. DOI: <http://dx.doi.org/10.1145/286498.286652>
- [46] Yvonne Jansen, Pierre Dragicevic, Petra Isenberg, Jason Alexander, Abhijit Karnik, Johan Kildal, Sriram Subramanian, and Kasper Hornbæk. 2015. Opportunities and Challenges for Data Physicalization. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 3227–3236. DOI: <http://dx.doi.org/10.1145/2702123.2702180>
- [47] Lynette A. Jones and Michal Berris. 2002. The psychophysics of temperature perception and thermal-interface design. In *Proceedings of the Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. IEEE, Piscataway, NJ, USA, 137–142. DOI: <http://dx.doi.org/10.1109/HAPTIC.2002.998951>
- [48] Shaun K. Kane and Jeffrey P. Bigham. 2014. Tracking @stemxcomet: teaching programming to blind students via 3D printing, crisis management, and Twitter. In *Proceedings of the ACM Symposium on Computer Science Education*. ACM, New York, NY, USA, 247–252. DOI: <http://dx.doi.org/10.1145/2538862.2538975>
- [49] Joseph Nathaniel Kaye. 2001. *Symbolic Olfactory Display*. Ph.D. Dissertation. Massachusetts Institute of Technology, Cambridge, MA, USA.
- [50] Harry T. Lawless and Hildegard Heymann. 1999. *Scaling*. Springer, Boston, MA, USA, 208–264. DOI: http://dx.doi.org/10.1007/978-1-4419-7452-5_7
- [51] Mathieu Le Goc, Charles Perin, Sean Follmer, Jean-Daniel Fekete, and Pierre Dragicevic. 2019. Dynamic Composite Data Physicalization Using Wheeled Micro-Robots. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 737–747. DOI: <http://dx.doi.org/10.1109/TVCG.2018.2865159>
- [52] Bongshin Lee, Petra Isenberg, Nathalie Henry Riche, and Sheelagh Carpendale. 2012. Beyond Mouse and Keyboard: Expanding Design Considerations for Information Visualization Interactions. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2689–2698. DOI: <http://dx.doi.org/10.1109/TVCG.2012.204>
- [53] Alan Lundgard, Crystal Lee, and Arvind Satyanarayan. 2019. Sociotechnical Considerations for Accessible Visualization Design. In *Proceedings of the IEEE Conference on Information Visualization*. IEEE, Piscataway, NJ, USA, 16–20. DOI: <http://dx.doi.org/10.1109/VISUAL.2019.8933762>
- [54] Jock D. Mackinlay. 1986. Automating the Design of Graphical Presentations of Relational Information. *ACM Transactions on Graphics* 5, 2 (1986), 110–141. DOI: <http://dx.doi.org/10.1145/22949.22950>
- [55] Kim Marriott, Falk Schreiber, Tim Dwyer, Karsten Klein, Nathalie Henry Riche, Takayuki Itoh, Wolfgang Stuerzlinger, and Bruce H. Thomas. 2018. *Immersive Analytics*. Lecture Notes in Computer Science, Vol. 11190. Springer International Publishing, New York, NY, USA. DOI: <http://dx.doi.org/10.1007/978-3-030-01388-2>
- [56] Todd Miller and John T. Stasko. 2002. Artistically conveying peripheral information with the InfoCanvas. In *Proceedings of the ACM Conference on Advanced Visual Interfaces*. ACM, New York, NY, USA, 43–50. DOI: <http://dx.doi.org/10.1145/1556262.1556268>
- [57] Tamara Munzner. 2014. *Visualization Analysis and Design*. CRC Press, Boca Raton, FL, USA.
- [58] Biswaksen Patnaik, Andrea Batch, and Niklas Elmqvist. 2018. Information Olfaction: Harnessing Scent to Convey Data. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 726–736. DOI: <http://dx.doi.org/10.1109/TVCG.2018.2865237>
- [59] Lewis V. Peterson and Wilbur Schramm. 1954. How accurately are different kinds of graphs read? *Educational Technology Research and Development* 2, 3 (June 1954), 178–189. DOI: <http://dx.doi.org/10.1007/BF02713334>
- [60] Jess Porter, Brent Craven, Rehan M. Khan, Shao-Ju Chang, Irene Kang, Benjamin Judkewitz, Jason Volpe, Gary Settles, and Noam Sobel. 2007. Mechanisms of

- scent-tracking in humans. *Nature Neuroscience* 10, 1 (2007), 27. DOI:<http://dx.doi.org/10.1038/nn1819>
- [61] Zachary Pousman and John T. Stasko. 2006. A taxonomy of ambient information systems: four patterns of design. In *Proceedings of the ACM Conference on Advanced Visual Interfaces*. ACM, New York, NY, USA, 67–74. DOI:<http://dx.doi.org/10.1145/1133265.1133277>
- [62] William Ribarsky, Jay Bolter, Augusto Op den Bosch, and Ron Van Teylingen. 1994. Visualization and analysis using Virtual Reality. *IEEE Computer Graphics and Applications* 14, 1 (Jan. 1994), 10–12. DOI:<http://dx.doi.org/10.1109/38.250911>
- [63] Ramik Sadana and John T. Stasko. 2014. Designing and implementing an interactive scatterplot visualization for a tablet computer. In *Proceedings of the ACM Conference on Advanced Visual Interfaces*. ACM, New York, NY, USA, 265–272. DOI:<http://dx.doi.org/10.1145/2598153>
- [64] Charles Spence, Marianna Obrist, Carlos Velasco, and Nimesha Ranasinghe. 2017. Digitizing the chemical senses: Possibilities & pitfalls. *International Journal of Human-Computer Studies* 107 (2017), 62–74. DOI:<http://dx.doi.org/10.1016/j.ijhcs.2017.06.003>
- [65] Robert J. Sternberg and Karin Sternberg. 2012. *Cognitive Psychology* (6th ed.). Cengage Learning, Stamford, CT, USA.
- [66] Stanley Smith Stevens. 1946. On the Theory of Scales of Measurement. *Science* 103, 2684 (1946), 677–680. DOI:<http://dx.doi.org/10.1126/science.103.2684.677>
- [67] Stanley Smith Stevens. 1957. On the psychophysical law. *Psychological Review* 64, 3 (1957), 153–181. DOI:<http://dx.doi.org/10.1037/h0046162>
- [68] Stanley Smith Stevens. 1958. Adaptation-Level vs. the Relativity of Judgment. *The American Journal of Psychology* 71, 4 (Dec. 1958), 633–646. DOI:<http://dx.doi.org/10.2307/1420322>
- [69] Faisal Taher, John Hardy, Abhijit Karnik, Christian Weichel, Yvonne Jansen, Kasper Hornbæk, and Jason Alexander. 2015. Exploring Interactions with Physically Dynamic Bar Charts. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 3237–3246. DOI:<http://dx.doi.org/10.1145/2702123.2702604>
- [70] Faisal Taher, Yvonne Jansen, Jonathan Woodruff, John Hardy, Kasper Hornbæk, and Jason Alexander. 2017. Investigating the Use of a Dynamic Physical Bar Chart for Data Exploration and Presentation. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 451–460. DOI:<http://dx.doi.org/10.1109/TVCG.2016.2598498>
- [71] Matthew Tobiasz, Petra Isenberg, and Sheelagh Carpendale. 2009. Lark: Coordinating Co-located Collaboration with Information Visualization. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (2009), 1065–1072. DOI:<http://dx.doi.org/10.1109/tvcg.2009.162>
- [72] Colin Ware. 2012. *Information Visualization: Perception for Design* (3rd ed.). Morgan Kaufmann Publishers, San Francisco, CA, USA.
- [73] Donald A. Washburn and Lauriann M. Jones. 2004. Could olfactory displays improve data visualization? *Computing in Science Engineering* 6, 6 (Nov 2004), 80–83. DOI:<http://dx.doi.org/10.1109/MCSE.2004.66>
- [74] Reto Wettach, Christian Behrens, Adam Danielsson, and Thomas Ness. 2007. A Thermal Information Display for Mobile Applications. In *Proceedings of the International Conference on Human Computer Interaction with Mobile Devices and Services*. ACM, New York, NY, USA, 182–185. DOI:<http://dx.doi.org/10.1145/1377999.1378004>
- [75] Graham Wilson, Stephen Brewster, Martin Halvey, and Stephen Hughes. 2012. Thermal Icons: Evaluating Structured Thermal Feedback for Mobile Interaction. In *Proceedings of the International Conference on Human-Computer Interaction with Mobile Devices and Services*. ACM, New York, NY, USA, 309–312. DOI:<http://dx.doi.org/10.1145/2371574.2371621>
- [76] Graham Wilson, Martin Halvey, Stephen A. Brewster, and Stephen A. Hughes. 2011. Some Like It Hot: Thermal Feedback for Mobile Devices. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 2555–2564. DOI:<http://dx.doi.org/10.1145/1978942.1979316>
- [77] Chuandong Wu, Jiemin Liu, Peng Zhao, Martin Piringer, and Günther Schaubberger. 2016. Conversion of the chemical concentration of odorous mixtures into odour concentration and odour intensity: A comparison of methods. *Atmospheric Environment* 127 (2016), 283–292. DOI:<http://dx.doi.org/10.1016/j.atmosenv.2015.12.051>
- [78] Cheng Xu, Ali Israr, Ivan Poupyrev, Olivier Bau, and Chris Harrison. 2011. Tactile display for the visually impaired using TeslaTouch. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 317–322. DOI:<http://dx.doi.org/10.1145/1979742.1979705>
- [79] Wai Yu and Stephen Brewster. 2002. Multimodal Virtual Reality versus printed medium in visualization for blind people. In *Proceedings of the ACM Conference on Assistive Technologies*. ACM, New York, NY, USA, 57–64. DOI:<http://dx.doi.org/10.1145/638249.638261>
- [80] Haixia Zhao, Catherine Plaisant, Ben Shneiderman, and Jonathan Lazar. 2008. Data Sonification for Users with Visual Impairment: A Case Study with Georeferenced Data. *ACM Transactions on Computer-Human Interactions* 15, 1, Article 4 (May 2008), 28 pages. DOI:<http://dx.doi.org/10.1145/1352782.1352786>