



Contextual variation and objectivity in olfactory perception

Giulia Martina¹ 

Received: 6 May 2020 / Accepted: 16 July 2021 / Published online: 31 July 2021
© The Author(s) 2021

Abstract

According to Smell Objectivism, the smells we perceive in olfactory experience are objective and independent of perceivers, their experiences, and their perceptual systems. Variations in how things smell to different perceivers or in different contexts raise a challenge to this view. In this paper, I offer an objectivist account of non-illusory contextual variation: cases where the same thing smells different in different contexts of perception and there is no good reason to appeal to misperception. My central example is that of dihydromyrcenol, a substance that can smell both woody and citrusy depending on what other odourants one has recently been exposed to. I first argue that the subjects' apparently conflicting reports about the way dihydromyrcenol smells are best understood as comparative characterisations of a smell. Given this understanding, different reports can be correctly made in response to perceiving the very same smell. I then argue that the phenomenal difference between the experiences subjects have across contexts can be explained compatibly with Smell Objectivism. On the account proposed, subjects perceive the very same smell but different qualities, notes, or aspects of it are salient to them, depending on the context of perception. I then consider how the proposed defence of Smell Objectivism can be adapted to other cases where the same thing is reported as smelling different in different contexts.

Keywords Perception · Olfaction · Smell · Perceptual variation · Appearance reports

✉ Giulia Martina
giulia.martina@unito.it

¹ Department of Philosophy and Education Sciences, University of Turin, Turin, Italy

1 Introduction

Olfactory experience, one may think, presents us with objective aspects of the world around us. When we take in the smell of freshly ground coffee, sniff to check whether the cake in the oven smells burnt, or are suddenly hit by the pungent smell of gasoline, we seem to be aware of entities in our environment that we can discover, track, or try to avoid.¹ Smell Objectivism articulates this thought: the smells we perceive in olfactory experience are objective and do not depend for their existence or nature on individual perceivers, their experiences, or their perceptual systems.² In this paper, I focus on smells understood as the qualities or properties—or composites of these—that account for the qualitative identities, similarities and differences among olfactorily perceivable entities. For instance, I will talk of the smell of this cup of coffee as the quality or complex of qualities shared across different cups of coffee that are olfactorily identical or that, in other words, smell exactly the same.

Philosophers disagree over what sort of entities we can olfactorily perceive. Candidates include ordinary objects such as cups of coffee, the stuffs that these objects are made of, special olfactory objects such as clouds of volatile molecules—or the stuff that these are made of—, or properties only.³ In this paper, I remain neutral on this issue: the smell of coffee may be a property—or composite of properties—of the cup, of the stuff in it, of a cloud of volatile molecules emanating from it, or of different sorts of entities in different cases. For simplicity, I will talk of ordinary objects as having or giving off smells, and will formulate Smell Objectivism accordingly. On Smell Objectivism, for instance, when I lift my cup of coffee and inhale I am perceiving a smell—a coffee-y, or more precisely, say, a hazelnut-chocolaty coffee-y smell—which does not depend on me, my experience of it, or my perceptual system—in short, as I will say, an objective and perceiver-independent smell.

In the philosophical literature, the general claim that smells are objective and independent of perceivers has been defended by Batty (2009).⁴ More specific versions of Smell Objectivism include views on which olfactory qualities are, or are determined by, chemical properties of stuffs or odour clouds (e.g. Carvalho, 2014;

¹ In this paper, I am concerned with orthonasal olfactory perception, i.e. olfactory perception that occurs when we breathe air in from the external environment.

² Two notes on this definition. First, I focus on claims about smells we perceive, as opposed to merely seeming to perceive. The main reason is that, on my proposed account, the phenomena discussed in this paper do not involve illusions or hallucinations—I consider an alternative account in Sect. 4. Second, the definition leaves it open whether smells are ‘anthropocentric’ in the sense that which of the properties in our environment count as smells depends in part on our perceptual systems. This anthropocentrism is compatible with the properties themselves being fully independent of our perceptual systems: what the perceptual systems do is ‘select’ some properties but not others to be the olfactory properties, and thus determine which properties are olfactorily available—see Hilbert (1987) and Kalderon (2007, pp. 590–594) for discussion of this issue in the colour domain.

³ See e.g. Batty (2010), Carvalho (2014), Lycan (2014), Mizrahi (2014), Young (2016), Richardson (2018), Barwich (2019). Pluralist accounts on which different sorts of entities can be olfactory perceived are also defended (e.g. Aasen, 2019).

⁴ Batty (2009) focuses on the thesis that smells are real and non-relational, but she seems to defend Smell Objectivism as I formulate it. For instance, she defends the claim that ‘the lilac smell is a perceiver-, or mind-, independent property of an odor’ (Batty, 2009, pp. 327).

Lycan, 2014; Young, 2016) and views on which olfactory qualities are objective but irreducibly qualitative, although possibly supervenient on chemical properties (e.g. Richardson, 2018). In this paper, I will only be concerned with Smell Objectivism as the general thesis formulated above.

As we know from the extensive literature on colour and colour perception, one of the main challenges faced by objectivist views of sensible properties is the phenomenon of perceptual variation, and in particular non-illusory perceptual variation. A coloured object can look different to different perceivers (e.g. due to differences in their visual systems) and in different conditions (e.g. under different illuminants or against different backgrounds) without any illusion or misperception being involved. The phenomenon can be easily and satisfactorily explained if we suppose that different colours are being perceived by different perceivers—or by the same perceiver in different conditions. But if we allow that the object has different colours at one and the same time, so the argument goes, we need to construe colours as perceiver-dependent properties.⁵

While in the colour domain similar arguments have been widely debated, discussion of the implications of perceptual variation for the metaphysics of smells has so far been very limited—Batty (2009) is a notable exception. But as a growing body of research is revealing, perceptual variation in the olfactory domain is widespread. The evidence suggests that the way an object or odourant smells to one can be affected by diverse factors such as the conditions in which one perceives, what other smells are present, one's degree of sensitivity to certain odourants, one's familiarity with certain smells, and even the way the thing one smells is labelled.⁶ In some cases it may be plausible to argue that an illusion or misperception is involved—for instance, if a perceiver suffers from the olfactory dysfunction *parosmia*, which causes them to experience familiar objects as having intense, unpleasant smells. But in many other cases we seem to have non-illusory variation.

Consider a case of intersubjective variation where the same object smells different to different perceivers, but we have no reason to take one perceiver but not the other to be suffering an illusion or misperception. How can both perceivers be 'getting things right'? How can they both genuinely perceive the smell of the object?⁷ An opponent of Smell Objectivism who takes smells to be perceiver-dependent has a good explanation.⁸ The thing really has two different smells, which is why it smells different to the two perceivers. And these smells are compatible, i.e. the object can have both smells at one and the same time, because the smells depend in part on the perceiver, their experience, or their perceptual system. For instance, smells may be construed as relations to perceivers or as dispositions to produce certain kinds of experiences.

⁵ See e.g. Cohen (2009) and Kalderon (2007) for formulations of this kind of argument.

⁶ For an overview of the research see e.g. Barwich (2019) and Wilson and Stevenson (2006, Ch. 6).

⁷ In using the expression 'getting things right', I intend to leave open what exactly allowing that neither perceiver is having an illusion amounts to. In what follows, I focus on the claim that both subjects are genuinely perceiving the smell of the object at stake.

⁸ In this paper, I will only engage with an opponent of Smell Objectivism who holds that there exist smells. See Batty (2009) for a discussion of eliminativism about smells.

We now have material for an argument against Smell Objectivism. For this explanation of perceptual variation is not available to Smell Objectivism, and it is not clear what alternative explanation there is. On the one hand, if the smelly object stays the same across the two perceiving subjects, and in particular if it does not differ in its objective and perceiver-independent properties, then Smell Objectivism seems to predict that the thing's smell does not differ across the two perceivers. Whether olfactory qualities are identical with, determined by, or supervenient on the objective and perceiver-independent properties of sources, odourants, or clouds of molecules these give off, if the latter properties do not differ, the former properties do not differ either. On the other hand, there are no good reasons to claim that one perceiver but not the other is having an illusion or misperception, so—unless we take both perceivers to be misperceiving—there is pressure to concede that both are genuinely perceiving the smell of the object they are presented with. But it seems that Smell Objectivism cannot concede this. This is because the object smells different to them. And how can the object smell different to them if it has—as Smell Objectivism predicts—the same, constant smell, so that there is only one smell there to be perceived, and if neither perceiver is suffering an illusion? The smell objectivist cannot, unlike their opponents, explain the difference in how things smell in terms of a difference in the smells perceived. They thus seem to be at a loss when it comes to accounting for this kind of case.

In this paper, I focus on a subset of cases of non-illusory perceptual variation where the way an object or odourant smells is affected by contextual factors of various kinds—cases of contextual variation. These are cases where the same odourant smells different to different perceivers—or to the same perceiver at different times—in different contexts of perception, although the perceivers' olfactory capacities are functioning normally, nothing in the environment is impairing their perception, and there is no reason to think that the perceivers differ in their olfactory sensitivity to the odourant involved. I argue that in cases of contextual variation the defender of Smell Objectivism can allow, contrary to what the above argument suggests, that both perceivers—or the same perceiver on both occasions—are 'getting things right' smell-wise. More precisely, the objectivist can allow that both contexts of perception are non-illusory, and so that both perceivers—or the same perceiver on both occasions—genuinely perceive the smell of the object they are presented with. This is so even if the object smells different across the two contexts of perception, i.e. if sniffing it results in phenomenally different olfactory experiences across the two contexts.

My argument proceeds in two steps. At each step, I discuss a certain source of evidence for the anti-objectivist claim that the best explanation of olfactory contextual variation is that the different perceivers—or the same perceiver at different times—perceive different smells. I argue that this explanation is not the only one available, and that the objectivist can offer an alternative explanation of the evidence that is at least as satisfactory. Some of the evidence for the anti-objectivist explanation is linguistic: in cases of contextual variation, perceivers make different and apparently conflicting reports about the way things smell. The first step in my argument is to outline an independently plausible understanding of the perceivers' reports as comparative characterisations of a smell. Given this understanding, the

different and apparently conflicting reports can be true of the same smelly object; we can thus understand how perceivers may make those reports in response to perceiving the same smell.

Some of the evidence for the anti-objectivist explanation of contextual variation is phenomenological: the same smelly object smells different, i.e. the perceivers undergo phenomenally different olfactory experiences in sniffing that object. The second step in my argument is to argue that differences in olfactory experience are not always best explained by differences in what smells are perceived. Here the discussion of the subjects' reports as comparatives will give us a positive suggestion for an alternative explanation. On my proposed account, while the perceivers perceive the same smell—the one, constant, objective smell that is present—different qualities, notes, or aspects of the smell are salient to them depending on the context of perception. As a result, different similarities that the smell bears to other smells become noticeable.

Here is a plan for the paper. In Sect. 2, I discuss a case of olfactory contextual variation: the chemical dihydromyrcenol smells citrusy in certain contexts and woody in others, depending on what other odourants the perceiver has just been exposed to. I articulate how the two steps in my argument apply to this example (Sects 2.2 and 2.3). In Sect. 3, I discuss how my defence of Smell Objectivism can be adapted to account for other cases that seem to involve contextual variation: one where labels affect the perceivers' responses to an odourant (Sect. 3.1), one where repeated exposure to a mixture of odourants affects how a single component of the mixture is experienced (Sect. 3.2). In Sect. 4, I compare my account with an alternative on which the cases of perceptual variation I discuss involve illusory or hallucinatory experiences.

2 A case of contextual variation

2.1 The challenge

Not every case where the way something smells differs across perceivers or conditions of perception is a case of perceptual variation. For instance, coffee smells different after one adds milk to it. Here the coffee smells different because it gives off a different smell, i.e. its olfactory qualities have changed. This is plausible because the odourant or source of the smell has changed—we added to coffee another substance with a smell of its own. Among cases of perceptual variation, i.e. cases where the source or odourant does not change in its objective and perceiver-independent properties and yet smells different, some are relatively easy for the defender of Smell Objectivism to explain. For instance, the cup of coffee smells different to a perceiver with normal olfactory capacities and to a perceiver suffering from anosmia who has lost their capacity to smell. Batty has argued that this kind of case is not a problem for Smell Objectivism: the difference between the normal and the anosmic perceiver amounts to the difference between perceiving certain smells and failing to perceive them (2009, pp. 331–332). Those smells may well be objective, perceiver-independent, and potentially available to appropriately equipped perceiving subjects. Similar

considerations apply to cases where one cannot perceive the smells in one's environment due to a temporary impairment of one's olfactory capacities—e.g. a bad cold—or the presence of intervening media—e.g. a face mask.

However, there are cases where these explanations are not available to the defender of Smell Objectivism. An example is that of dihydromyrcenol. This substance is usually described as having a citrusy, fresh, and lavender-like smell. However, the context in which dihydromyrcenol is presented, and in particular which other odourants are presented immediately before it, can make a difference to how it smells: whether citrusy and lime-like or woody and pine-like. Subjects consistently rate dihydromyrcenol as smelling very citrusy when it is presented after woody-smelling odourants such as pine, sandalwood and cedar oil, and as smelling more woody and significantly less citrusy when it is presented after citrusy-smelling odourants such as lemon, orange, and mandarin oil (Lawless et al., 1991; Lawless, 1991a, b).

We plausibly have a case of perceptual variation. The perceivers' ratings change across different contexts (a 'woody' context and a 'citrusy' context) in which they are presented with the same odourant, and it seems natural to think that the ratings change in response to a change in how the odourant smells to the perceivers or, in other words, in response to a change in the phenomenology of their olfactory experience. Moreover, this case poses a new challenge to Smell Objectivism, because Batty's explanation does not seem to apply: the difference between the two ways dihydromyrcenol smells cannot be explained by a difference between perceiving and failing to perceive the smell of dihydromyrcenol. The evidence reported by Lawless and colleagues concerns subjects without any known olfactory impairments, who can detect the presence of dihydromyrcenol by smelling in both the woody and the citrusy context. Furthermore, there is no reason to think that the difference in how dihydromyrcenol smells across the two conditions is due to variation in the subjects' degree of olfactory sensitivity to dihydromyrcenol: what matters is which odourants are presented before. Both the citrusy and the woody way of smelling are in principle available to every subject, as long as the subject is in the 'right' context.

On the face of it, this case of contextual variation does not involve an illusion or misperception. There is no obvious reason to think that the perceivers presented with dihydromyrcenol in the woody context, but not those presented with dihydromyrcenol in the citrusy context—or vice versa—are suffering an illusion as of some smells that are not really there. If so, the opponent of Smell Objectivism can apply the argument from non-illusory perceptual variation outlined in Sect. 1.

Consider two perceivers, Sam and Ana, participating in Lawless' experiment: they are both presented with the very same sample of dihydromyrcenol, but Sam has been presented with woody odourants just beforehand while Ana has been presented with citrusy odourants. As a result, Sam rates the smell of dihydromyrcenol as 'very citrusy' while Ana rates it as 'woody'. Since neither Sam nor Ana is suffering an illusion, the right verdict seems to be that they are both genuinely perceiving the smell of dihydromyrcenol—they are both 'getting things right'. If we take smells to be partly dependent on perceivers, their experiences, or their perceptual systems, we can respect this verdict. We can say that dihydromyrcenol has two compatible smells simultaneously: a very citrusy smell relative to Sam and a rather woody smell

relative to Ana. A good explanation of the fact that dihydromyrcenol smells different ways to our perceivers is that Sam perceives the former smell while Ana perceives the latter smell. Depending on my opponent's preferred metaphysics, the proposal would be spelled out somewhat differently, for instance as the claim that there are two different olfactory objects with different olfactory qualities, both having the same source, or simply as the claim that there are different olfactory qualities or composites of these.

This explanation does not seem to be available to the objectivist. Not only does Smell Objectivism predict that the olfactory qualities of the dihydromyrcenol sample stay the same across the two perceivers—as there is no change in the objective and perceiver-independent properties of the sample. But it seems that it would also predict that there are no other differences in what smells are present across the two conditions where Sam and Ana, respectively, are presented with the sample. This is because the woody and citrusy odourants are not presented together with dihydromyrcenol, but before. In fact, the contextual variation in ratings remains in a version of the experiment where subjects smell the woody or citrusy odourants in one room and then move to a different room to smell dihydromyrcenol (Lawless, 1991b). So this case of perceptual variation is unlike a case where something comes to smell different because new odourants have been introduced in the same environment, giving rise to a new 'mixed' smell—as would occur, for instance, if one was smelling coffee while perfume is sprayed in the room.⁹ Structurally, the case of dihydromyrcenol is more similar to the visual phenomenon of successive colour contrast, where, for instance, a yellow circle looks a cooler, more greenish yellow when seen after perceiving a red circle and a warmer yellow when seen after a green circle. In the dihydromyrcenol case, the objective and perceiver-independent properties that are present when dihydromyrcenol is sniffed by Sam and Ana are the same; Smell Objectivism would thus predict that there is no difference in the smells that are present.

Given this, it seems that Smell Objectivism cannot accommodate the non-illusory nature of the variation. More precisely, it seems difficult to reconcile the plausible claim that both Sam and Ana are genuinely perceiving the smell of the dihydromyrcenol sample with the fact that the sample smells different to them, if we suppose that there is only one, constant smell there to be perceived. In the remainder of this section, I will argue that this conclusion can be resisted. The defender of Smell Objectivism can explain the contextual variation in how dihydromyrcenol smells by holding that both Sam and Ana genuinely perceive the smell of dihydromyrcenol—i.e. the one objective, perceiver-independent smell that is present—even though dihydromyrcenol smells different to them.

The anti-objectivist argument from contextual variation relies on two key claims. First, a good explanation for the variation in how dihydromyrcenol smells across two perceivers is that the perceivers are perceiving different smells, i.e. different

⁹ This kind of explanation is implausible for our case also because dihydromyrcenol does not smell more woody in the woody context, as one would expect if the smell of dihydromyrcenol and the smells of the woody odourants had mixed.

olfactory qualities. Second, if we hold that the smell that is present, and thus the smell that is perceived, does not differ across the two perceivers, we cannot provide an explanation of the contextual variation that is consistent with its non-illusory nature—or at least our explanation will not be as good as the anti-objectivist one. I will not take issue directly with the first claim. Instead, I will dispute the second claim, considering the linguistic evidence based on the perceivers' ratings and reports in Sect. 2.2 and the phenomenological evidence, i.e. the fact that subjects have phenomenally different olfactory experiences, in Sect. 2.3.

2.2 Smell reports

The initial evidence for taking the case of dihydromyrcenol to be an example of perceptual variation is linguistic and concerns the perceivers' ratings and reports. Lawless originally observed the contextual variation in how dihydromyrcenol smells when a panel of trainees in food evaluation described this odourant as 'almost wholly citrus in character' in a session following typically woody odourants, but judged it to be too woody or pine-like to be a good example of a citrus odourant when evaluating it with a battery of citrus-smelling odourants (Lawless, 1991a, p. 49). This finding was confirmed in a series of experiments in which subjects rated dihydromyrcenol along a certain qualitative dimension—on a bipolar scale from 'very citrusy' to 'very woody' (Lawless et al., 1991; Lawless, 1991b). For simplicity, I will assume that these ratings can be translated into linguistic reports about the way things smell such as 'this smells woody', 'this smells of wood', or 'this smells like wood'. This would be either because expressions like 'woody', in the context we are considering, are best interpreted as contracted versions of smell reports such as 'this smells woody', or because in categorising an odourant as woody—by applying the descriptor 'woody' to it—subjects implicitly accept similar smell reports as true. So, Sam sniffs dihydromyrcenol after being exposed to woody smells and reports:

(1) This smells citrusy.¹⁰

Ana, presented with the same sample of dihydromyrcenol after being exposed to citrusy smells, reports:

(2) This smells woody.

On the face of it, this evidence from ratings and reports provides support for the anti-objectivist account of the contextual variation. Reports (1) and (2) seem to conflict: one characterises the smell as citrusy, the other as woody; and citrusy-ness and woodiness are very different olfactory qualities. But then, one may think, the reports

¹⁰ For simplicity, I am considering this report instead of a graded report like 'this smells very citrusy' even though on average the subjects in Lawless' and colleagues' experiments rate dihydromyrcenol as very citrusy. I do not think this significantly affects the ensuing discussion.

can both be true or correct only if they are characterisations of different smells: a very, or mainly, citrusy smell, and a rather woody smell. Since Sam and Ana plausibly make reports (1) and (2) each on the basis of their own olfactory experience, i.e. the olfactory experience they have when sniffing the sample of dihydromyrcenol, they can both be making correct reports only if they are reporting on, and perceiving, different smells. And just like we do not seem to have good reasons to take one perceiver but not the other to be suffering an illusion, we do not seem to have good reasons to take one perceiver but not the other to be making a false or incorrect smell report by wrongly characterising the smell of dihydromyrcenol with report (1) rather than (2)—or vice versa. The evidence from ratings and reports thus seems to favour an account on which dihydromyrcenol has two different smells at one and the same time, and Sam and Ana each perceive a different smell. As we have seen in Sec. 2.1, this explanation is arguably incompatible with Smell Objectivism.

In response, the defender of Smell Objectivism can argue that there is an independently plausible way of understanding smell reports such as (1) and (2) on which the linguistic evidence does not support the above explanation. It is not obvious how we should understand reports about the way things smell of the form ‘*o* smells *F*’, ‘*o* smells like an *F*’, or ‘*o* smells of *F*’. I propose that smell reports like these primarily characterise the smells of the smelly objects they are about: they characterise odourants or sources of smells by characterising their smells. This is especially plausible for reports (1) and (2), made by Sam and Ana in the context of Lawless’ experiment. These reports are primarily concerned with the smell of the sample rather than with its non-olfactory properties—they are concerned with the way things smell independently of how they otherwise are. So, for instance, an object may smell citrusy without being a citrus fruit (and even without containing any citrus fruit), and a report such as (1) may be true of that object regardless.

What, then, does it take for smell reports such as (1) and (2) to be true? I argue that these reports can be understood as having comparative import: they characterise the smell they are about comparatively with reference to the smells of certain kinds of odourants or sources of smells.¹¹ In particular, report 1) can be understood, approximately, as having this kind of structure:

1*) This has a smell which is similar to the characteristic smell of citrus fruit.

For the purposes of this discussion, we can take the characteristic smell of citrus fruit to be the class of smells typically or normally given off by citrus fruit—for instance, smells typically given off by lemons, oranges, and mandarins.¹² These are smells that, generally, we would recognise as likely having a certain kind of source (i.e. a citrus fruit). Given this comparative understanding of the reports, when Sam rates dihydromyrcenol as ‘very citrusy’, Sam is implicitly characterising the smell of dihydromyrcenol in terms of its similarity to the smells typical of citrus fruit.

¹¹ The proposal I sketch here is inspired by Martin’s (2010) comparative account of reports about the way things look. However, since there are significant differences between looks talk and smells talk, a fully-developed comparative account of smell reports will likely look somewhat different from Martin’s.

¹² In general, what counts as the characteristic smell of *F*s in part depends on the context: the conversational focus and communicative goals, as well as the speaker’s experience and knowledge may affect what class of smells is being referred to.

Analogously, in implicitly making the smell report (2), Ana is characterising the smell of dihydromyrcenol in terms of its similarity to the smells typical of wood.

Understanding smell reports such as (1) and (2) as comparatives is independently plausible as it fits well with the way we ordinarily talk about smells.¹³ English—like many other languages—has very few, if any, specifically olfactory predicates, i.e. terms that exclusively or primarily apply to smells in the way that colour predicates like ‘red’ apply to colours.¹⁴ Evidently, this feature of the English language does not prevent us from talking about smells. In fact, we can characterise smells in very subtle detail if needed, as shown by professional perfume or wine reviews. However, we usually characterise smells using terms that apply—or at least primarily apply—to entities other than smells (e.g. Majid & Burenhult, 2014; Majid et al., 2018). These include kinds of sources (e.g. banana, coffee), properties of these sources (e.g. burnt, ripe), the stuff these sources are made of (e.g. woody, metallic), taste properties (e.g. sweet, sour), and tactile properties (e.g. cool, pungent). For instance, we say that the peaches smell ripe and sweet, that a cake smells burnt, that a house smells of old wood, or that a green tea smells grassy.

If smell reports are understood as having comparative import, we have a good explanation of how we can successfully convey information about the way things smell and even characterise smells in detail in the absence of specifically olfactory predicates that pick out olfactory qualities directly. The smells of both familiar and unfamiliar objects and stuffs are characterised in terms of their similarity to the characteristic smells of things we are familiar with.

Given this comparative understanding of smell reports, we can argue that the linguistic evidence does not support the account of the dihydromyrcenol case that the opponent of Smell Objectivism proposes. If smell reports (1) and (2) are comparatives, then they are not conflicting or incompatible characterisations of a smell. Report (1) does not directly attribute a certain quality to a smell—say, citrusyness—such that report (2) would attribute an incompatible quality—woodiness. So we should not conclude that the reports can both be correct or true only if they characterise two different smells—a very citrusy one, and a woody one—which Sam and Ana, respectively, would be perceiving and reporting on. Instead, (1) and (2) may well be true of the very same smell, because a certain smell may bear a similarity to different kinds of smells at one and the same time. In particular, the smell of the sample of dihydromyrcenol that both Sam and Ana are presented with and reporting on may well be both similar to typical smells of citrus fruit and to typical smells of wood insofar as it may be similar to the first class of smells in a respect and similar to the second class of smells in another respect.

¹³ I think that a comparative account is plausible not only for (1) and (2) but also for many, if not all, smell reports. For the purposes of my argument, however, it is sufficient to accept the more restricted claim.

¹⁴ As Majid and colleagues show, this is not true of all languages, as hunter-gatherer societies tend to have a richer olfactory vocabulary (e.g. Majid & Burenhult, 2014; Majid et al., 2018). In this paper, I will not consider whether and how my discussion of the linguistic evidence would apply to languages other than English.

Plausibly, comparatives involve an implicit specification of a respect in which the thing characterised is similar to some other things. Accordingly, smell reports specify a respect in which the smell they are about and the smells in the comparison class are similar. Consider a more familiar example involving flavour. The flavour of milk chocolate is similar to the flavour of dairy products in some respects—sweetness, milkiness, creaminess—and similar to the flavour of dark chocolate in others—cocoa notes, nuttiness. We can thus report the flavour of milk chocolate as milky and sweet, but also as cocoa-y or nutty. In making the report (1), then, Sam may be picking out the smell of dihydromyrcenol as similar in a respect to certain smells, while Ana, in making the report (2), may be picking the same smell out as similar in another respect to some other smells. The citrusy and the woody contexts of perception may make a difference to which respects of similarity subjects focus on and deem relevant.

The claim that our smell reports are sensitive to respects of similarity, and that this affects our choice of smell report in a context fits well with some features of olfactory categorisation more generally. As Jraissati and Deroy (2021) argue, due to the complexity of smells, there are always many different aspects along which smells can vary and be compared. While the qualities of the stimuli categorised impose constraints, different speakers as well as the same speaker in different contexts may choose different categories for the same stimulus as they adopt different criteria for assessing similarity (Jraissati & Deroy, 2021, p. 5–7). Which comparison class is considered, i.e. which set of smells a certain smell is being compared to, also affects categorisation. Lawless and colleagues point out that this may play a role in the dihydromyrcenol case. Having been exposed to a smell which is a good example of a category (e.g. the smell of lemon oil for the category of citrusy smells), one's criteria for accepting a less prototypical smell (such as the smell of dihydromyrcenol) as a member of the same category become more stringent, so that one would be less inclined to label the smell as citrusy (Lawless et al., 1991, pp. 359–360).

Moreover, given this understanding of Sam's and Ana's smell reports, it is not only theoretically possible, but plausible that their reports are true of the very same smell. As Lawless puts it, dihydromyrcenol seems to be either an 'ambiguous' or 'multistable' smell or a 'two-note' smell, in that it has both a citrusy and a woody aspect (e.g. Lawless, 1991a, 1991b pp. 317, p. 52). But if both reports (1) and (2) can be correct characterisations of the very same smell, then we can understand how Sam and Ana may make those different smell reports in response to perceiving the very same smell, where this smell may well be objective and fully independent of the two perceivers, their experiences, and their perceptual systems. Since we have been supposing that those reports adequately capture Sam's and Ana's ratings, then those ratings do not support an account of the contextual variation that appeals to different smells being perceived by Sam and Ana.

2.3 Similarities and salience

Understanding smell reports as comparatives allows the defender of Smell Objectivism to explain how different and apparently conflicting ratings and reports can be

made in response to perceiving the very same smell. However, the objectivist still faces the challenge of explaining why two perceivers would make different reports, rate an odourant differently or more generally categorise it differently, if there is no difference in the smells they perceive. In the case of dihydromyrcenol, addressing this challenge requires moving beyond the level of language, because it is plausible that the difference between Sam's and Ana's ratings of the odourant is at least partly explained by the fact that dihydromyrcenol smells different to Sam and to Ana, i.e. that Sam and Ana are having different olfactory experiences in response to sniffing dihydromyrcenol.

To see what the challenge is, it is useful to compare the case of dihydromyrcenol with other cases where different subjects, or the same subject in different conditions, use different smell reports when presented with the same odourant. For instance, different subjects may characterise the same solution of benzaldehyde as smelling of almond, of marzipan, or of maraschino cherry (Lawless, 1991a). As these are all typical sources of smells with the distinctive almond-y and fruity character that a solution of benzaldehyde has, the different reports do not give us reason to think that benzaldehyde smells different to different subjects, i.e. that the subjects are having phenomenally different olfactory experiences.

The case of benzaldehyde, it is natural to think, is importantly different from that of dihydromyrcenol. All three characterisations of the smell of benzaldehyde (referring to almonds, cherries, and marzipan) can be appropriate in the very same context of perception. In fact, if a subject is familiar with all three kinds of sources, they can recognise the aptness of all three reports upon perceiving the sample of benzaldehyde. By contrast, reports (1) and (2), uttered in response to sniffing dihydromyrcenol, are used selectively. For instance, only when one is presented with dihydromyrcenol after having been exposed to citrusy odourants does one find report (2)—'this smells woody'—to be appropriate. Why does a subject find a certain comparative characterisation of dihydromyrcenol apt only in a certain context of perception? Here, unlike in the case of benzaldehyde, there is pressure to appeal to a difference in olfactory experience.¹⁵ According to the opponent of Smell Objectivism, moreover, this difference in olfactory experience is best explained by a difference in what smells are perceived by subjects who find different comparative reports apt. For instance, Ana finds report (2) apt because, due to a contextual effect after exposure to citrusy odourants, they perceive a woody smell; Sam, by contrast, perceives a very citrusy smell. But then it is plausible that these different smells will be perceiver-dependent.

The defender of Smell Objectivism can agree that there is a difference in the phenomenology of Sam's and Ana's experiences of dihydromyrcenol: the case of dihydromyrcenol is an example of perceptual variation. However, they can point out, it is a further claim that this evidence is best explained in terms of a difference in what smells Sam and Ana perceive, because a different explanation is

¹⁵ See Rudolph (2020) for a defence of the idea that various cases where subjects seem to disagree—on her view, faultlessly disagree—about the way something smells or tastes are best explained by supposing that subjects have phenomenally different perceptual experiences.

available. Our discussion of the linguistic evidence in the first step of my defence of Smell Objectivism (Sect. 2.2) gives a suggestion to pursue for this second step.

As we have seen, contextual factors can affect one's choice of comparative smell report and one's ability to appreciate a report as appropriate on an occasion. As Rudolph points out, it is in general common for subjects to use apparently conflicting comparative reports about the way things appear because 'they might be focusing on different aspects of the appearances, or they might differ in which aspects they think are most relevant' (Rudolph, 2020, pp. 211). Consider describing the flavour of milk chocolate when comparing it with different kinds of chocolate. One might focus on the milkiness and creaminess of milk chocolate if one wants to highlight the respect in which its flavour differs from that of dark chocolate, while focusing on the cocoa notes when one is comparing milk chocolate with white chocolate.

In the case of dihydromyrcenol there is a difference in perceptual context: Sam has just been exposed to woody odourants, Ana to citrusy ones. The perceptual context, I propose, affects Sam's and Ana's choice of report by making certain qualities, aspects or notes of the smell they are presented with more or less salient or noticeable to them. When one has just perceived some citrusy smells, for instance, the woody character of the smell of dihydromyrcenol tends to stand out and be noticed. In that context, one finds it apt to describe dihydromyrcenol as smelling woody. Given the perceptual salience of a certain aspect of the smell one perceives, one's characterisation of the smell focuses on a certain respect in which the smell is relevantly similar to certain classes of smells one is familiar with—here, the respect in which the smell is similar to the characteristic smell of wood. The salience of different qualities, notes, or aspects of the smell across the citrusy and the woody contexts makes a difference to the phenomenology of Sam's and Ana's olfactory experience. The analogy with the flavour of milk chocolate is again helpful: milk chocolate tastes very milky and sweet after one has eaten some dark chocolate, but its cocoa and nutty aspects stand out if one eats it after having some white chocolate.

There are different ways of spelling out this proposal, depending on one's commitments and preferences concerning the metaphysics of smells and of olfactory experience. I will outline three options here. On the first option, simply, Sam and Ana perceive the same olfactory object, but experience different olfactory properties or qualities of it (citrusy-ness vs. woodiness) depending on the context. This olfactory object really has both properties at one and the same time, but, contrary to what the anti-objectivist argument suggests, these can be construed as objective and perceiver-independent. To defend this thesis, one could appeal to a form of pluralism about olfactory qualities on the model of colour pluralism (Mizrahi, 2016; Kalderon, 2017). On this view, olfactory objects would have multiple but perceiver-independent olfactory qualities simultaneously, where different qualities can be perceived in different conditions and by different perceivers. Alternatively, one could hold that the olfactory object has multiple properties analogously to a multicoloured object having one colour on its front side and another on its back side. The different contexts, on this view, make one but not the other property perceivable by a subject at any one time, analogously to how different parts of the multicoloured object may be occluded from view in different viewing conditions.

On the second option, Sam and Ana perceive the same olfactory properties—whether or not they also perceive an object—but different properties are more or less salient depending on the context. Here the smell of dihydromyrcenol is construed as a composite of olfactory properties or qualities comprising both a citrusy component quality and a woody one, analogously to the surface of a striped shirt comprising large white stripes and thin blue stripes, perceived together as a pattern. In the visual case, while the shirt looks predominantly white, the blue stripes may become salient in certain conditions, whether because one deliberately attends to them or because they stand out—for instance, if one sees the shirt next to one with thin red stripes or next to a blue pair of trousers. As Lawless observes, it may be that dihydromyrcenol behaves somewhat like a smell mixture with distinguishable components, and that the salience of one component over the other is due to an adaptation effect.¹⁶ Having been exposed to citrusy smells, the perceiver adapts to the citrusy quality in the smell of dihydromyrcenol and the woody quality becomes more evident, similarly to how the cinnamon component of a mixture of vanillin and cinnamaldehyde ‘pops out’ after one is exposed to vanillin and adapts to the vanilla component of the smell.

On the third option, Sam and Ana perceive the same smell, but this is construed as one complex olfactory property with multiple notes, aspects or dimensions that may not be perceptually distinguishable from the overall smell. The citrusy-ness and the woodiness, on this picture, are not distinct olfactory properties, but rather aspects of one smell—those aspects responsible for the similarity the smell bears to typical smells of citrus and typical smells of wood, respectively. Across different contexts of perception, different such aspects or notes are more or less salient. For an approximate visual analogy, consider aspects of a surface such as glossiness and metallic quality, which we do not experience as visually distinct from colours but rather as dimensions of the appearance of coloured surfaces. The metallic-ness of a surface may become salient as different sources of light hit the surface or when the surface is seen next to one that has the same hue but is not metallic.

All three options are compatible with Smell Objectivism and all three allow one to satisfactorily address the challenge from perceptual variation by providing an account of the phenomenal difference in the subjects’ experiences of dihydromyrcenol. However, there is a significant difference between the first option, on the one hand, and the second and third options, on the other. On the first option there is after all a difference in what olfactory properties Sam and Ana perceive, even if there is a sense in which they perceive the same smell insofar as they perceive the same olfactory object and both olfactory properties are simultaneously instantiated by the object. By contrast, on the third and second option we are not appealing to a difference in the olfactory properties perceived in order to account for the phenomenological difference between Sam’s and Ana’s experiences. In general, differences in perceptual salience may not be accompanied by differences in what one perceives. Holding the objects and properties perceived fixed, there can be differences in what parts of a scene, objects, or properties a subject attends to, whether because they

¹⁶ See Lawless (1991a, b), and Lawless, Glatter and Hohn (1991) for discussion.

voluntarily direct their attention in a certain way, for instance in order to complete a task, or because certain aspects of their environment simply stand out to them. Extending this idea, I suggest that different components, notes, or aspects of the complex smells that both Sam and Ana perceive can stand out to their attention and be more noticeable on account of the different perceptual context. Importantly, these component qualities or notes are not perceiver-dependent: they are aspects of how the smell objectively is, and they are there to be experienced independently of perceivers.

Given this difference, which option a defender of Smell Objectivism chooses and develops will depend not only on their preferred metaphysics of smells and on whether they think that we can olfactorily experience objects, but also on their wider commitments concerning the nature of olfactory experience. In particular, the second and third option allow for the phenomenology of olfactory experience to differ—either for two perceivers or for the same perceiver at different times—even when what is perceived, i.e. the olfactory properties or complex property, stays the same. As a result, there does not seem to be co-variation between the phenomenology or qualitative character of an experience and what is experienced.¹⁷ Philosophers who want that relation to hold can still adopt a version of my proposal by pursuing the first option, on which the difference in phenomenology is a matter of which olfactory properties of the same smell are experienced.

I conclude this section by presenting two reasons why a defender of Smell Objectivism may prefer an account of the dihydromyrcenol case that appeals to either the second or third option over the first—their wider commitments notwithstanding. First, we do not have reasons to think that in one context of perception one simply fails to perceive a certain property or component of the smell of dihydromyrcenol that one would then come to perceive in the other context. Consider two cases where this explanation may be plausible. Suppose one is wearing a face mask and is presented with a perfume which is primarily lavender-y with a weak note of bergamot; while one perceives the smell of the perfume, one may only perceive the dominant lavender note; as one takes one's mask off, one comes to perceive the delicate bergamot note too. To take a visual example, think of seeing a large abstract painting through a crowd of museum visitors; the painting first looks uniformly blue; but when one gets closer, one discovers that the painting has a red patch which was initially occluded from view. The case of dihydromyrcenol is different, as there does not seem to be anything in the context—such as an intervening medium, an occluder, or another overpowering smell—that would impede one's perception. Given that we are supposing the subjects to be equally sensitive to dihydromyrcenol, we should thus conclude that both subjects are in a position to perceive whatever olfactory qualities there are.

A second consideration is that the second and third options—on which a complex smell is perceived in both contexts—fit better with the experimental evidence.

¹⁷ Depending on one's views, there may be ways to preserve the co-variation. For instance, one could hold that perceptual salience makes a difference to what one experiences by increasing the determinacy of the properties one experiences—see e.g. Nanay's (2010) account of visual attention.

While subjects rate dihydromyrcenol as smelling very citrusy in the woody context and as smelling more woody in the citrusy context, they nonetheless rate it as having both notes in both contexts, i.e. as smelling also a bit woody and, respectively, a bit citrusy (Lawless, 1991b, pp. 318). So subjects' ratings do not indicate that the smell of dihydromyrcenol appears like a purely citrusy smell in one context and like a purely woody smell in the other. It thus seems plausible that the olfactory experience one has in each context is different from the olfactory experience one would have if one was presented with a citrusy quality and, respectively, a woody quality in isolation. On the second and third way of spelling out my proposal, this is because in both contexts one is presented with, and perceives, either a composite of both olfactory qualities or a complex smell with both a citrusy and a woody aspect.

3 Extending the strategy

The account outlined for the case of dihydromyrcenol exemplifies a general strategy that the defender of Smell Objectivism can adopt in various cases of non-illusory contextual variation, i.e. cases where the same odourant smells different to different perceivers, or to the same perceiver at different times, even though the perceivers' olfactory capacities are functioning normally, nothing in the environment is impairing their perception, and there is no reason to think that the perceivers differ in their olfactory sensitivity to the odourant. On this general strategy, we first appeal to a comparative understanding of the perceivers' smell reports to argue that the different reports can be made in response to perceiving the very same smell; second, we offer an account of why perceivers make different reports while perceiving the same smell that is compatible with Smell Objectivism—in the case of dihydromyrcenol, we can appeal to the contextual salience of different qualities, components, or aspects of its smell.

If this objectivist strategy did not apply beyond the example of dihydromyrcenol, it would be very limited in scope and thus unlikely to be of much significance in the development of an objectivist approach to perceptual variation. In this section, I look at two other candidate cases of contextual variation and discuss how the general strategy can be adapted and extended to them.

3.1 Effects of labels

The first example is a case of variation in how an odourant is reported as smelling depending on the linguistic context in which it is perceived. In particular, presenting the same odourant with different labels can elicit very different hedonic reactions and associations (de Araujo et al., 2005; Djordjevic et al., 2008; Herz & von Clef, 2001). For instance, subjects respond very differently to a mixture of isovaleric and butyric acids with the label 'parmesan' and to the very same mixture with the label 'vomit': they give contrasting pleasantness ratings (e.g. pleasant vs. very unpleasant), they categorise or name the odourant differently (e.g. as cheese vs. as vomit),

and associate it with different behaviours, desires, and memories (e.g. grating over pasta vs. leaving the room) (Herz & von Clef, 2001, pp. 384–387).

According to Herz and von Clef, this evidence supports the conclusion that subjects have different and contrasting ‘perceptual responses’ depending on whether the mixture of acids is ‘experientially perceived as “vomit” or “parmesan cheese”’ (2001, pp. 235, 388, 390). This interpretation is also advanced by Barwich, who argues that the labels elicit a ‘markedly different phenomenological experience’ and ‘diverging perceptions of the same stimulus’ (2019, p. 7). We may thus have a case of cognitive influence on perception: the subjects’ conceptual identification of the smell as a smell of a certain kind (e.g. a parmesan smell) or their belief, prompted by the label, that a certain kind of source is present (e.g. parmesan cheese) affects their olfactory experience.¹⁸ If the subjects’ perceptual experiences differ across the two linguistic contexts, then, unless one of the experiences is illusory (or both are), the opponent of Smell Objectivism could argue that this difference is best explained by a difference in what the subjects perceive: two very different smells. On this proposal, smells are not only partly dependent on the perceivers’ perceptual systems or olfactory sensitivity, but also partly dependent on non-perceptual cognitive factors such as beliefs and expectations about what sort of sources of smells are present.

On the face of it, appealing to different smells being perceived has some intuitive plausibility. How could one give such contrasting pleasantness ratings and have such different associations in response to perceiving the very same smell? How can the same smell elicit both a desire to eat the odourant and an urge to avoid it? When we consider the experiment more carefully, however, we can see that this anti-objectivist explanation is not supported by the evidence. While the experiment shows that labels for sources of smells with either positive or negative connotations can make a difference to one’s hedonic and cognitive response to the same odourant, it does not show that different labels can elicit different perceptual, and more specifically olfactory, experiences. As authors of further studies using a similar paradigm with very similar results conclude, labels can affect ‘odors’ affective properties’, ‘perceived pleasantness’, subjects’ ‘emotional reactions to the odor’, and ‘affective responses to stimuli’ (Djordjevic et al., 2008, pp. 386, 393; de Araujo et al., 2005, pp. 676). If one’s hedonic, emotional, or affective response changes across the two linguistic contexts, this may well make an experiential difference, so that the overall phenomenology of one’s experience of the same odourant may differ across contexts. But these differences are compatible with one’s olfactory experience—or the olfactory component of one’s overall experience—staying the same. A fortiori, the experimental evidence does not motivate an explanation—which would be incompatible with Smell Objectivism—on which the subjects are experiencing different smells or olfactory qualities in those contexts.

¹⁸ Neither Herz and von Clef nor Barwich explicitly defend this account of the case. However, Barwich considers the case as part of various motivations for thinking that ‘olfactory processing [...] is fundamentally shaped by other sensory and higher-level cognitive processes (e.g., cross-modal cues, verbal descriptors, selective attention, or the mood of the perceiver)’ (2019, p. 7).

Instead, the mixture of isovaleric and butyric acids may smell qualitatively the same to the subjects even if the smell strikes them as pleasant and desirable in one context and as unpleasant and repulsive in another context. We are familiar with situations of this kind: for instance, the smell of chocolate cake strikes us as delicious when we are savouring the first slice, but as unappetising when we have had two slices and feel full. This is perfectly consistent with the chocolaty, rich, sweet smell of the cake staying the same across the two contexts; rather, something about us—our sensations, appetite, and desires—has changed. In this scenario, one might think, it is not our olfactory perception that changes—we olfactorily experience the very same chocolaty, rich, sweet qualities—but our hedonic response to that perception.¹⁹

Similarly to one's sensation of fullness, labels can make a difference to one's hedonic and cognitive response to the same smell. In the experiment, the 'parmesan' and 'vomit' labels are linguistic cues that prompt subjects to think about certain kinds of sources of smells among the ones that they could anyway associate with the smell they perceive. We know that labels can help us identify smells: while subjects generally struggle to spontaneously generate labels for the smells they perceive, they quickly recognise the appropriateness of a label when given a range of options (e.g. Jönsson & Stevenson, 2014).²⁰ In the case we are considering, it is plausible that the labels would generate very different hedonic reactions and thus pleasantness ratings because of the strong positive and negative connotations that parmesan cheese and vomit, respectively, have for most people. In the Herz and von Clef version of the experiment, moreover, subjects are encouraged to think of memories, desires, and behaviours in response to the smell; these may elicit further emotions and possibly mental imagery of the two smell sources and of circumstances one associates with them, contributing to the overall phenomenology of one's experience.

Herz and von Clef argue that the subjects' responses to the same odourant were so dramatically different that subjects 'perceived the odor as being different' and 'thought they were smelling a different odor' across the two contexts (2001, pp. 384–385). The motivation offered for these claims is based on the linguistic evidence: independent judges who did not know about the experimental design compared each subject's pleasantness ratings and reports across the two contexts, assessing the 'perceptual interpretation' as 'different' if the ratings and reports were overall different.²¹ Given this criterion, it is unsurprising that the judges would take

¹⁹ We can of course use the verb 'smell' with hedonic and evaluative predicates: for instance, we can say that the cake smells delicious. But this fact alone does not show that when we employ a contrasting hedonic or evaluative predicate to describe the same cake our olfactory experience has changed—just like going from judging that an all-black outfit looks boring to judging that it looks cool as our taste in clothing changes does not per se show that our colour experience has changed.

²⁰ This is true for speakers of many languages, but it is not universally true. Speakers of languages that—unlike English—have a dedicated olfactory lexicon, such as Jahai speakers, spontaneously generate consistent labels for smells with ease, much like English speakers do for colours (e.g. Majid & Burenhult, 2014).

²¹ Herz and von Clef offer the following example of a combination of ratings and reports which would be assessed as 'different perceptual interpretation': the subject recalls food-related memories, says that the odourant would go well with pasta, that they would eat it, and that they would call it 'cheese' in one

the subjects in the experiment to believe that what they were smelling in the two contexts differed: because of the different pleasantness ratings and associations, the subjects' reports were very different across contexts. Moreover, since the 'parmesan label' condition and the 'vomit label' condition were not in immediate succession, it is unsurprising that the subjects did not recognise the odourant as the same.²² That they did not recognise it as the same in those conditions, however, does not show that they had different olfactory experiences. Imagine presenting the subjects with both the 'vomit'-label jar and the 'parmesan'-label jar, containing the same mixture of isovaleric and butyric acids. If we asked the subjects to sniff again and judge whether the jars smell the same or different, intuitively, we would expect at least some of the subjects to come to recognise—by smelling—that the jars contain two samples of the same odourant; that is because, it is natural to think, the jars smell exactly the same. Supposing that the jars come to smell different in the experimental setting just because subjects cannot compare them side-by-side seems implausible, and it is in any case unnecessary to explain the ratings and reports.

We have seen some reasons why appealing to different smells being perceived across the two linguistic contexts is not necessary in order to account for the effects of labels, which therefore do not threaten Smell Objectivism. Applying the comparative account of smell reports outlined in Sect. 2.2 to the case under discussion will help us understand why it is not only possible, but very plausible to argue that the same smell is genuinely perceived across the two contexts. In addition to rating pleasantness and reporting their memories, subjects in the experiment are asked what they would call the odourant (Herz & von Clef, 2001, pp. 384). A subject replying 'parmesan' can be taken to implicitly be making a report such as 'this smells of parmesan' or 'this smells like parmesan'. On the comparative account of smell reports, this report characterises the smell comparatively in terms of its similarity with smells typically given off by parmesan cheese. But then the report 'this smells of parmesan' and the report 'this smells of vomit'—said in response to smelling the same mixture of acids in a different context—may be true of the very same smell. The same smell may well be both similar in a respect to the characteristic smell of parmesan cheese, and similar in a respect to the characteristic smell of vomit. In principle, the smell may be similar to those two comparison classes in the same respect.²³

Moreover, we have reason to think that the smell of the mixture of acids is indeed similar to the typical or characteristic smells of both kinds of sources. The effects

Footnote 21 (continued)

context; but they indicate that they would call it 'vomit' and that they would like to run out of the room in the other context (2001, pp. 385).

²² In the Herz and von Clef study, the two conditions were one week apart. In the Djordjevic et al. study, the same odourant was presented with different labels within the same session, but interspersed with many other odourants and labels—in this case, some subjects did realise that the same odourant had been presented multiple times (2008, pp. 389).

²³ Of course there are respects in which typical smells of parmesan cheese are not similar to the smell given off by the mixture of acids, as we should expect given that there are respects in which typical smells of parmesan cheese are not similar to typical smells of vomit.

of labels on subjects' pleasantness ratings are found with certain odourants but not with others. In the experiment by Herz and von Clef, for instance, subjects gave different pleasantness ratings and reports also when violet leaf was presented with the labels 'fresh cucumber' and 'mildew'; by contrast, menthol did not give rise to a comparable effect when presented with labels 'chest medicine' and 'breath mint'. As Herz and von Clef point out, this suggests that the presence and magnitude of the effect depends not only on whether the labels refer to sources that have strong negative and positive connotations, but also on 'the degree of anchoring the odor had with the verbal labels' (2001, pp. 388–389). I take this to mean that the presence and magnitude of the effect depends in part on the degree to which the smell of the odourant can be associated by a subject with both kinds of sources (e.g. parmesan and vomit, cucumber and mildew).

In the case of the mixture of acids, the association is strong because both parmesan cheese and vomit are typical and familiar sources of smells similar in quality to that given off by the mixture.²⁴ As Barwich observes, this is because of the chemical properties of those sources: 'butyric acid can be perceived as parmesan as well as vomit *because it is an element of both* parmesan and vomit aroma' (2019, p. 7). But then it is plausible that the same smell is present and is perceived in both linguistic contexts: a smell which is similar to the typical smells of both parmesan and vomit — and recognisably so. It is in virtue of perceiving that smell that subjects come to have different hedonic responses, cognitive associations, and possibly imagery related to those sources.

3.2 Odour-odour learning

A second example of contextual variation that my general strategy may be applied to is the phenomenon of odour-odour learning.²⁵ In a series of experiments, Stevenson has shown that repeatedly exposing subjects to mixtures of two odourants affects their ratings of the qualities of the component odourants and of how similar these odourants are to one another (e.g. Stevenson, 2001a, b, c). For instance, subjects who have been exposed to a mixture of mushroom-smelling champignon and lemon-smelling citral rate champignon as lemony—and as more lemony than a control odourant that had not been paired with citral—and rate champignon and citral as more similar to each other than odours that they have not experienced together in a mixture (Stevenson, 2001b, pp. 165–173). Moreover, similar effects were obtained when champignon was experienced in a mixture with cherry odour: in this case,

²⁴ Interestingly, there is at least anecdotal evidence that people sometimes find certain types of cheese as smelling like vomit in everyday contexts where they know that what they are smelling is cheese. See e.g. discussions on whether cheese smells like vomit: <https://www.quora.com/Why-does-the-Parmesan-cheese-in-some-restaurants-smell-more-like-vomit-than-like-cheese>; https://www.reddit.com/r/Cheese/comments/4gbfc3/cheese_tastessmells_like_vomit/

²⁵ Another promising example is the phenomenon of odour-taste learning (e.g. Stevenson et al., 1995): exposure to combinations of an odourant (e.g. vanilla) and a sweet-tasting food causes the odourant to be reported as smelling sweet when presented alone. This case is more complex because it plausibly requires discussing cross-modal associations, retronasal olfaction, and flavour perception.

subjects rated champignol as more cherry-like and sweet than a control odourant that had not been paired with cherry. Another experiment compared subjects' ratings of minty-smelling L-carvone before and after exposure to a mixture of L-carvone and musty-smelling p-anisaldehyde, and found that L-carvone was rated as mustier post-exposure (Stevenson, 2001a).

Stevenson argues that odour-odour learning involves a change in one's olfactory experience: champignol comes to smell more lemony or, respectively, more cherry-like; L-carvone comes to smell mustier. Stevenson goes on to suggest that the reason why, for instance, champignol is rated as lemony and as similar to citral after exposure to the mixture is that champignol has acquired a quality of citral—a lemony quality (e.g. Stevenson, 2001a, b, pp. 567, 574, pp. 156). On a possible reading of this suggestion, a subject comes to perceive, in smelling champignol, an olfactory quality (lemony-ness, say) that they would not have perceived when presented with champignol before the exposure phase. But then the opponent of Smell Objectivism would argue that the lemony olfactory quality that the subject perceives after the exposure phase is partly dependent on the subject, and in particular on their previous experiences of the champignol-citral mixture. This account of the case is suggested by Nudds, who comments that Stevenson's experiments show that 'past experience determines the qualities or character of the smells we experience, and that the qualities of a smell can change as a consequence of past experience'; as a result, Nudds observes, subjects might perceive different different smells even though they are exposed to the same chemicals (2017, pp. 98–99).

The defender of Smell Objectivism can agree that odour-odour learning makes a difference to the olfactory phenomenology of one's experience—not only to one's ratings, reports, hedonic reactions, or cognitive associations. If it did not make a difference to the olfactory phenomenology of one's experience, it is natural to think, why would one report that champignol smells lemony or, respectively, cherry-like and that minty L-carvone smells mustier? Nonetheless, the objectivist can argue that this difference does not have to be explained in terms of a new olfactory quality (e.g. a lemony quality) being perceived, as a result of one's past experience, in the presence of an odourant that is not normally judged as having that quality (e.g. champignol, which does not normally smell lemony). To be sure, the explanations given for the other cases of contextual variation discussed so far do not seem promising. When one reports that champignol smells lemony, one might convey that one is having an experience that is somewhat similar to the experience one normally has when smelling lemons.²⁶ But it seems implausible to hold that champignol really is similar to typical smells of lemon, where exposure to the mixture simply makes that similarity noticeable to one. A reason to doubt this hypothesis is that champignol can equally come to smell sweet and cherry-like to someone who has been exposed to

²⁶ In this context, one may be using the report 'this smells lemony' to convey how the smell of champignol affects one or what sort of psychological response it elicits. As Martin argues for the visual domain, a comparative looks report can be used to express how the speaker finds things to be similar, characterising the appearance of an object with reference to one's psychological response to perceiving the object (2010, pp. 215–217).

a mixture of champignol and cherry: the effect really seems to depend just on one's past experience of the mixture. However, Stevenson himself suggests an alternative explanation that the defender of Smell Objectivism can take on board.

Stevenson argues that odour-odour learning may be understood as 'an implicit memory effect on perception' (2001a, pp. 576). His hypothesis is that perceiving a component of the mixture one has been exposed to triggers the olfactory memory of the whole mixture, i.e. a memory of the mixed smell that has been encoded 'in a relatively unanalysed configural form' (e.g. 2001a, pp. 574). For instance, smelling champignol triggers the olfactory memory of the champignol-citral mixture, which is a memory of a mixed mushroomy-lemomy smell. Given this hypothesis, a simple explanation of the effect of one's past experience would appeal to non-perceptual olfactory phenomenology: memory-based olfactory mental imagery of the smell of the mixture, including the absent lemony component, accompanies one's perceptual experience of the (mushroomy) smell of champignol.²⁷ Such olfactory mental imagery can make a difference to both the phenomenology of one's experience and one's reports: champignol is reported as smelling a bit lemony because its smell reminds subjects of the mushroomy-lemomy smell of the mixture that they have recently and repeatedly perceived.

This explanation of the case is in line with Stevenson's own hypothesis and is compatible with Smell Objectivism. Stevenson sometimes writes that, as a result of the configural memory being triggered, the smell being perceived acquires 'something of the quality of its absent mixture partner' (2001a, p. 574). But this further claim can be resisted: an explanation appealing to the contribution of memory-based mental imagery does not commit us to the claim that, as a result of one's past experience, one comes to perceive a new lemony olfactory quality, which will have to be construed as dependent on one's past experience. When one is presented with champignol, whether before or after being exposed to the champignol-citral mixture, one perceives the same mushroomy smell, which can thus be fully objective and perceiver-independent.

In this section, I argued that the objectivist strategy can be adapted to account for two candidate cases of olfactory contextual variation that are significantly different from the dihydromyrcenol case. This illustrates not only that the scope of the strategy is wide, but that the strategy is flexible and can appeal to different explanatory resources in different cases — a feature that we should regard as advantageous.

4 An alternative: misperception

Throughout this paper, I have been assuming that in all the cases discussed subjects genuinely perceive the smells that are present. In this section, I consider an alternative account that does not share this assumption, and argue that my account has at least as much explanatory power as the alternative while being less committal.

²⁷ For an empirically-informed discussion of the similarities between olfactory mental imagery and olfactory perception, see Young (2020).

Some of the authors of the empirical studies presented in this paper hold that all or some of the phenomena discussed involve an illusion or misperception. Herz and von Clef (2001) claim that the effects of labels are a case of olfactory illusion, and Stevenson presents these effects as well as the case of dihydromyrcenol as illusions where ‘the same stimulus is reportedly perceived in a different manner’, whether ‘qualitatively (what it smells like)’ or ‘hedonically’ (2011, pp. 1889–1890). Odour-odour learning is cited as an example of illusion involving ‘different stimuli, same percept’—both champignol and citral can give rise to an experience of something smelling lemony (Stevenson, 2011, pp. 1892). The plausibility of these claims, I suggest, depends on what notion of illusion or misperception we appeal to.

Herz and von Clef rely on the assumption that ‘an illusion is observed when a stimulus is invariant but context alters its perception’ (2001, pp. 383, 388, 390). If all it takes to have an olfactory illusion is a change in one’s perceptual experience while the odourant does not change, then it is difficult to deny that at least some of the cases we have been discussing are illusions—except perhaps the very case Herz and von Clef focus on, if, as I have suggested, it does not involve a change in perceptual phenomenology. But this definition is so weak that it is arguably too liberal. Given this definition, many ordinary olfactory experiences would count as illusory, including familiar cases where an object smells more or less intensely depending on whether we are closer or further away from it or on whether or not there is an intervening medium such as a face mask or a closed window. In effect, adopting this definition amounts to denying that there can be non-illusory contextual variation in how things smell.

By contrast, Stevenson’s definition of an illusion or misperception is stronger. He observes that the notion of a misperception relies on the idea that in perceptual experience there can be ‘a veridical state, in which the mind accurately reflects some objective state of the environment’ (Stevenson, 2011, pp. 1893). The cases of contextual variation we discussed, then, are such that there is ‘a disparity between some objective state of the world and one’s perception of it’, and thus satisfy what Stevenson presents as a ‘consensus definition’ of illusion in the psychological literature (Stevenson, 2011, pp. 1888–1889). In her discussion of Stevenson’s proposal, Batty (2014) notes that this definition is similar to the traditional philosophical notion of an illusion or misperception as involving a contrast between the way things appear and the way things are. The case of dihydromyrcenol, for instance, can be described as involving a change in ‘how a subject perceives the stimulus to be’ or in ‘the odourant’s apparent properties’ (Batty, 2014, p. 5). This change is not meant to be merely a change in what the subject takes the odourant or the smell to be like, or in the properties they would attribute in their judgements or reports. Rather, the subject’s olfactory experience presents or represents the environment as being a way it is not, it presents some smells or olfactory qualities as present even though they are not. Within Batty’s framework, olfactory misperceptions are hallucinations as of smells or olfactory qualities that are not instantiated in the subject’s surroundings.²⁸

²⁸ On Batty’s view we do not olfactorily experience objects, and so we cannot olfactorily experience objects as having properties they do not have—which is the standard philosophical definition of an illusion. Misperceptions are thus hallucinations. I here do not distinguish this version of the misperception account from one appealing to illusions.

If one thinks that we olfactorily experience objects, then some misperceptions may be illusions as of those objects as instantiating olfactory qualities that they do not instantiate.

The misperception account—even given this stronger notion of misperception—is, *per se*, compatible with Smell Objectivism. Suppose that one’s characterisation of dihydromyrcenol as smelling woody is made in response to an illusory experience of a non-woody smell as being woody, or even in response to a hallucination as of a woody smell. These smells could still be construed as objective and independent of perceivers, their experiences, and their perceptual systems. In general, philosophers who defend objectivist views of sensible qualities are often happy to understand certain cases of perceptual variation as involving misperception—for instance, when strong coloured light makes a white object look blue and one is misled by one’s visual experience.

However, appealing to the misperception account to explain olfactory contextual variation is unjustified. Most obviously, as we have seen in Sect. 2.3, the empirical evidence fits well with taking the smell of dihydromyrcenol to genuinely include both woody and citrusy qualities, notes or aspects. Moreover, there is no obvious reason to think that the conditions in which dihydromyrcenol is perceived, *i.e.* after exposure to woody or to citrusy odourants, would give rise to misperception. Those conditions are quite ‘normal’, as in ordinary circumstances we constantly perceive different smells both at the same time and in close succession and we do not take this either to prevent us from genuinely perceiving the smells that are present or to cause illusions or hallucinations.

The misperception account also does not suit the effects of linguistic labels discussed in Sect. 3.1. While the labels ‘parmesan’ and ‘vomit’ influence one’s reports, cognitive associations, and hedonic reactions to the mixture of isovaleric and butyric acids, we do not have good reasons to think that they result in different olfactory experiences. Moreover, there is nothing illusory or misleading about the case: the smell of the mixture is indeed similar both to the smells typically given off by parmesan and to those typically given off by vomit, and is thus a good indicator of the presence of both kinds of sources.

Taking the phenomenon of odour-odour learning to be illusory may seem more promising. According to Stevenson, when the currently presented mushroomy-smelling champignon evokes the olfactory memory of the mixed mushroomy-lemmony smell, this memory ‘forms an integral part of the current percept’ and champignon comes to smell a bit lemony as a result (2001a, pp. 576). If so, Stevenson could argue that we have an illusion: if we rule out that the smell of champignon really has a lemony note, a component of the subject’s olfactory experience — their experience as of a lemony quality — does not ‘match’ the objective state of the world. In response, I would argue that the objectivist proposal appealing to memory-based olfactory imagery outlined in Sect. 3.2 offers an equally satisfactory as well as less committal explanation. First, it is plausible that the mushroomy smell that is in fact present is genuinely perceived—it is not the case that champignon does not smell mushroomy anymore after exposure to the champignon-citral mixture. Second, following Stevenson’s own hypothesis, we can appeal to the contribution of olfactory mental imagery to explain the effect of past experience on one’s current experience.

It thus seems unnecessary to hold that, while genuinely perceiving the mushroomy quality of the smell that is in fact present, one has an illusion or hallucination as of an un-instantiated lemony olfactory quality. Finally, my proposal allows that we consider the experience subjects have when champignol is reported as smelling lemony as misleading. For instance, due to the accompanying mental imagery, one's experience may mislead as to the character of the smell that is there, and thus not be a good guide to the way things are in one's environment—for instance, the experience may not be good grounds for judging that lemons are present. In light of these considerations, I propose to account for both misleading and non-misleading cases of contextual variation without appealing to an apparent awareness of, or misrepresentation of, an un-instantiated olfactory quality or smell.

5 Conclusion

I argued that Smell Objectivism can account for a range of cases where, at least on the face of it, the same thing smells different to different perceivers, or to the same perceiver at different times, due to contextual factors. In these cases there are no good reasons to think that one subject but not the other is either simply failing to perceive the smell of the thing or having an illusory or hallucinatory experience as of a different smell that is not really there. The objectivist, just like their opponent, can account for the plausible verdict that both perceivers genuinely perceive the smell of the thing.

To show this, I first argued that the perceivers' reports about the way the thing smells are plausibly understood as comparative characterisations of the smell they perceive as similar in some respect to the smells characteristic of certain kinds of things. Given this comparative reading, different and apparently conflicting reports can be correctly used to characterise the same smell, and so can be made in response to perceiving the same smell. The reports differ because different similarities the smell bears to other smells are more salient in different contexts.

In the dihydromyrcenol case, the perceivers' different ratings and reports are accompanied by phenomenally different olfactory experiences. To account for these phenomenal differences compatibly with Smell Objectivism, I proposed that different perceivers can perceive the same smell—whether this is understood as an olfactory object, as a composite of olfactory properties or as a complex olfactory property—even though different qualities, notes, or aspects of the smell are salient to them in different contexts of perception. As a result, different similarities that the smell bears to other smells can be more or less noticeable.

The same general argumentative strategy I proposed for the case of dihydromyrcenol can be adapted to other phenomena that, at least on the face of it, are cases of contextual variation in how things smell. In this paper, I considered two further examples: one where the variation in reports does not seem to be accompanied by different olfactory experiences, but rather by different hedonic and cognitive responses, and one where the experiential difference may be explained in terms of memory-based mental imagery elicited by the perceptual experience of the smell.

The paper also highlighted some questions for future research concerning the metaphysics of smells. I have outlined different ways to develop the idea that smells are complex entities having multiple qualities, components, or aspects that can be differently manifest in experience depending on contextual factors. For an objectivist who wants to adopt the argumentative strategy proposed here for cases of contextual variation, the question is which of these options best suits their specific version of Smell Objectivism and how exactly it should be developed.

Acknowledgements Material that resulted in this paper has been presented at workshops and conferences at the University of Bochum, the University of Eastern Piedmont, the University of Antwerp, and the University of Durham; I thank the participants for their helpful questions. For their comments and suggestions on various versions of this material, I am grateful to Tom Crowther, Mark Kalderon, Hemdat Lerman, Matt Nudds, and Simon Wimmer. I am also grateful to the anonymous referees for this journal for their constructive comments.

Funding Open access funding provided by Università degli Studi di Torino within the CRUI-CARE Agreement. OeAD (Ernst Mach Grant); PRINproject The Mark of the Mental (2017P9E9NF).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aasen, S. (2019). Spatial aspects of olfactory experience. *Canadian Journal of Philosophy*, 49(8), 1041–1061.
- Barwich, A.-S. (2019). A critique of olfactory objects. *Frontiers in Psychology*. <https://doi.org/10.3389/fpsyg.2019.01337>
- Batty, C. (2009). What's that smell? *Southern Journal of Philosophy*, 47, 321–348.
- Batty, C. (2010). Olfactory experience: Objects and properties. *Philosophy Compass*, 5(12), 1147–1156.
- Batty, C. (2011). Smelling lessons. *Philosophical Studies*, 153, 161–174.
- Batty, C. (2014). The illusion confusion. *Frontiers in Psychology*, 5, 231.
- Carvalho, F. (2014). Olfactory objects. *Disputatio*, 6(38), 45–66.
- Cohen, J. (2009). *The Red and the Real: An Essay on Color Ontology*. Oxford University Press.
- de Araujo, I. E., Rolls, E. T., Velazco, M. I., Margot, C., & Cayeux, I. (2005). Cognitive modulation of olfactory processing. *Neuron*, 46, 671–679.
- Djordjevic, J., Lundström, J. N., Clément, F., Boyle, J. A., Pouliot, S., & Jones-Gotman, M. (2008). A rose by any other name: Would it smell as sweet? *Journal of Neurophysiology*, 99, 386–393.
- Herz, R. S., & von Clef, J. (2001). The influence of verbal labeling on the perception of odors: Evidence for olfactory illusions? *Perception*, 30, 381–391.
- Hilbert, D. (1987). *Color and Color Perception: A Study in Anthropocentric Realism*. Center for the Study of Language and Information, Stanford, CA.
- Jraissati, Y., & Deroy, O. (2021). Categorizing smells: A localist approach. *Cognitive Science*. <https://doi.org/10.1111/cogs.12930>
- Jönsson, F., Stevenson, R. (2014). Odor knowledge, odor naming, and the 'tip-of-the-nose' experience. In B. Schwartz, A. Brown (Eds), *Tip-of-the-tongue states and related phenomena* (305-326). Cambridge: Cambridge University Press.

- Kalderon, M. E. (2007). Color pluralism. *The Philosophical Review*, 116(4), 563–601.
- Lawless, H. (1991a). The sense of smell in food quality and sensory evaluation. *Journal of Food Quality*, 14(1), 33–60.
- Lawless, H. (1991b). A sequential contrast effect in door perception. *Bulletin of the Psychonomic Society*, 29(4), 317–319.
- Lawless, H., Glatter, S., & Hohn, C. (1991). Context-dependent changes in the perception of odor quality. *Chemical Senses*, 16(4), 349–360.
- Lycan, W. G. (2014). The intentionality of smell. *Frontiers in Psychology*, 5, 436.
- Majid, A., & Burenhult, N. (2014). Odors are expressible in language, as long as you speak the right language. *Cognition*, 130(2), 266–270.
- Majid A., Burenhult, N., Stensmyr, M., de Valk, J., Hansson, B.S. (2018) Olfactory language and abstraction across cultures. *Philosophical Transactions of the Royal Society B* 373 (1752).
- Martin, M.G.F. (2010). What's in a look? In B. Nanay (Ed.), *Perceiving the world*. Oxford: Oxford University Press.
- Millar, B. (2019). Smelling objects. *Synthese*, 192, 4279–4303.
- Mizrahi, V. (2006). Color objectivism and color pluralism. *Dialectica*, 60(3), 321–336.
- Mizrahi, V. (2014). Sniff, smell, and stuff. *Philosophical Studies*, 171(2), 233–250.
- Nanay, B. (2010). Attention and perceptual content. *Analysis*, 70, 263–270.
- Nudds, M. (2017). The non-visual senses: The distinctive role of sounds and smells. In B. Nanay (Ed.), *Current controversies in philosophy of perception*. New York: Routledge.
- Richardson, L. (2018). Odours as Olfactibilia. In C. Mac Cumhail, T. Crowther (Eds.), *Perceptual ephemera*. Oxford: Oxford University Press.
- Rudolph, R. E. (2020). Talking about appearances: The roles of evaluation and experience in disagreement'. *Philosophical Studies*, 177(1), 197–217.
- Stevenson, R. J. (2001a). The acquisition of odour qualities. *Quarterly Journal of Experimental Psychology*, 54A, 561–577.
- Stevenson, R. J. (2001b). Associative learning and odor quality perception: How sniffing an odor mixture can alter the smell of its parts. *Learning and Motivation*, 32, 154–177.
- Stevenson, R. J. (2001c). Perceptual learning with odors: Implications for psychological accounts of odor quality perception. *Psychonomic Bulletin & Review*, 8, 708–712.
- Stevenson, R. J. (2011). Olfactory illusions: Where are they? *Consciousness & Cognition*, 20, 1887–1898.
- Stevenson, R. J., Prescott, J., & Boakes, R. A. (1995). The acquisition of taste properties by odors. *Learning & Motivation*, 26, 433–455.
- Wilson, D., & Stevenson, R. J. (2006). *Learning to smell: Olfactory perception from neurobiology to behavior*. The Johns Hopkins University Press.
- Young, B. D. (2016). Smelling matter. *Philosophical Psychology*, 29(4), 520–534.
- Young, B. D. (2020). Olfactory imagery: Is exactly what it smells like. *Philosophical Studies*, 177, 3303–3327.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.