A Puzzle about Discrimination

Eugen Fischer

It is no surprise that empirical psychology refutes, again and again, assumptions of uneducated common sense. But some puzzlement tends to arise when scientific results appear to call into question the very conceptual framework of the mental to which we have become accustomed. This paper shall examine a case in point: Experiments on colour-discrimination have recently been taken to refute an assumption of first-person authority that appears to be constitutive of our ordinary notion of perceptual experience. The paper is to show that those experiments do not refute this assumption, and will suggest that the impression to the contrary is, ultimately, due to mistaken translation from the technical idiom of empirical psychology into the plain English we use every day. This is to take the mystery out of what we shall see to constitute a pretty puzzle; and it is to remind us just how careful we need to be when drawing conclusions from results of scientific psychology.

I.

The puzzle arises from the results of discrimination-experiments on colour-vision involving the

> ***Forced-choice Task T:*** Two stimuli are repeatedly presented to a subject, switched around in a random fashion. The subject is to indicate, each time, which of the two is the previously indicated 'test-stimulus'. When he feels unable to tell, the subject just is to guess.

Such experiments have established

> ***Result E:*** Even stimuli they sincerely and carefully judge to look the same to them are discriminable for non-pathological human subjects,

where the technical term is (partially) defined by means of this sufficient condition:

> ***Discrimination****:* Two stimuli are *discriminable* for a subject if he is able to identify the previously indicated test-stimulus at a significantly higher than random rate, over a series of trials, in the context of forced-choice task T,

where what is to count as a 'significantly higher than random rate' remains, up to a point, a matter of stipulation, viz., of the desired level of confidence (Clark 1993, 45-50).

That this result conflicts with a central assumption of 'folk psychology' is suggested by this

> ***Common Intuition:*** When two stimuli are discriminable for a subject, the subject 'picks up' some difference between them, when they are presented to him. As he picks it up visually, those stimuli must actually look different to him, even if he does not realise. So the subject may be wrong not only about how things *actually look* but also about how they currently *appear to him*.

As substantiated by the particular Result E, this conclusion goes against our ordinary beliefs: Common-sense is of course more cautious than some armchair philosophy, and credits us with the last word on how things look to us only if some further conditions are fulfilled. The subject must master the concepts he employs, and be able to apply them properly in the given case; he must have enough time to make a careful choice, and must exercise enough care in making it. But all these conditions are fulfilled in the present case: Subjects are not asked to describe the colours they see, creating room for mistakes as to whether, say, 'magenta' is quite the right word to describe the particular shade of colour the test-stimulus looks to them. Rather, subjects are here only asked whether one patch of coloured light and another look the same to them. Nor can Result (E) be taken to reveal mistakes due to carelessness: Subjects find even stimuli discriminable that they previously had all the time in the world to consider and declared to look the same to them even upon the most careful scrutiny.

Interpreted in the light of the Common Intuition, Result (E) thus seems to refute a genuinely common-sense assumption of first-person authority (and not merely a piece of armchair philosophy). What is more, this common-sense assumption has often been taken to be 'constitutive' of 'our ordinary notions of perceptual experience'. So the present result seems to refute not only an assumption of common-sense but to reveal the emptiness of a central subset of our mental vocabulary, by establishing

> ***Conclusion C***: First-person authority fails to hold even for basic features of our current perceptual experiences such as, e.g., their qualitative identity, and fails even under the favourable conditions under which common sense takes it to obtain.

This conclusion, then, is rather puzzling, viz., the moment one puts oneself into the position of the experiments' subjects and asks: 'How can I be wrong not only about whether two patches of colour *actually* look the same, but about whether they now look the same *to me*, if, stare at them as I might, I just cannot *see*, cannot get myself to *grow aware* of, any difference between them?'

Lately, some philosophers have been explicitly arguing to precisely this puzzling conclusion, viz., from results of discrimination experiments on blindsight patients: Such patients are able to discriminate light stimuli they sincerely claim not to see at all (Stoerig & Cowey, 1992). As the discriminations mirror those of ordinary subjects and are in line with the prevailing opponent-process model of colour-vision, such patients, it is argued against what would be the common-sense view, *experience* colours they are *not aware* of seeing (e.g., Nelkin, 1993, 1995). The argument, it is conceded (Nelkin, 1995), is not particularly strong. So why is it put forward in spite of combining this short-coming with a most highly counter-intuitive conclusion? Presumably, it is only because the Common Intuition forcefully suggests the move from the discrimination results to this conclusion, already prior to explicit theorising, that these philosophers are disposed to countenance the admittedly weak arguments they then explicitly adduce.

If this is right, their puzzling position is best attacked by addressing that tacit motive, the powerful intuition that prepares the ground for the more involved empirical argument without being explicitly stated itself. The present paper shall therefore drag this quietly influential intuition out into the open, and refute and remove it, so as to pave the way for a more sober appreciation of the pertinent empirical work.

II.

As arising in the context of the experiments on non-pathological subjects mentioned first, the Common Intuition, when spelled out more fully, translates into this persuasive ***Reasoning R:***

(A1) When two stimuli are discriminable for a subject, that subject discriminates them when they are tested for discriminability.
(A2) When tested for discriminability, the stimuli are presented to the subject's sight; he looks at them and does not perceive them by exercising any other sense. Hence:
(C1) Subjects discriminate the stimuli presented while looking at them, and not by exercising any other sense.

This first conclusion almost compels us to make two inferences:

(A3) If that is so, the subjects must discriminate those stimuli *by* exercising their sense of sight, by *looking* at them. But:
(A4) If *that* is how they discriminate them, these stimuli must *look* different to the subjects.

Two compelling inferences, apparently to the best explanation, thus yield:

(C2) When two stimuli are discriminable for a subject, they look different to him. But:
(E) Subjects can discriminate stimuli they sincerely and carefully report to look the same to them. Hence:
(C3) Some stimuli they sincerely and carefully judge to look the same to them actually look different to subjects.

I.e.: Even the sincere and careful reports of linguistically competent subjects like ourselves are not, as we ordinarily think, the last word on whether things look the same to them. Instead, such reports are at best the last word on whether we *believe* that things look the same to us. And such beliefs may be wrong. Whence **Conclusion (C)**.

This line of reasoning makes an otherwise obvious slip look pardonable: Subjects who can discriminate two stimuli are often (e.g., Clark 1993, 48) loosely described as being able to 'distinguish' the stimuli or to 'tell' them 'apart'. But of course this step is not logically valid: Suppose that, confronted with forced-choice Task T, I can discriminate stimuli I sincerely and carefully judge to look the same to me. Reduced to a series of (unbeknownst to me: mostly correct) guesses, I cannot be said to be able to 'distinguish' these stimuli, as the word is ordinarily used: 'He cannot distinguish them; he is merely guessing which is which.' (The same goes for 'to tell apart'.) So we cannot translate the technical 'discriminate' with the plain 'distinguish', and would be wrong if we reasoned: 'Subjects can discriminate even stimuli they carefully judge to look the same to them, hence they can distinguish some such stimuli, so some such stimuli must look different to them.' Of course Reasoning R is not quite as quick-and-dirty. But, as we shall see, it owes its *prima facie* plausibility to a mistake of the same kind.

I now want to show that this reasoning, innocent though it looks, does not license the puzzling conclusion (C); and I want to remove the impression of mystery this fallacious conclusion may leave us with. To refute the reasoning, I shall establish two points: First, the crucial inferences (A3) and (A4) construe the discrimination of stimuli x and y as an achievement; and both inferences are acceptable, viz., as inferences to the best explanation, only if this achievement is attributable to the subject (section 3). Second, this requirement is not met to the extent required to infer conclusion (C2) without an important qualification; and this qualification makes it impossible to infer the mystifying conclusion (C3) from the by itself straightforward finding (E) (section 4). Eventually, I will spell out the conclusions we are left with instead, in such a way as to dispel whatever sense of mystery their predecessor might enshroud our experimental finding (E) in (section 5).

III.

To set the stage for the first point, note that it is indeed plausible to construe discriminating as an achievement. By definition, two stimuli x and y are

discriminable for a subject iff, in the context of some such forced-choice task like T, he is able to identify the target stimulus x at a significantly higher than random rate, over a series of trials. Apparently, he can thus be said to *discriminate* x and y iff he manages to identify x at a significantly higher than random rate, over a series of trials, when engaged in such a forced-choice task (or, on an alternative construal: iff he would manage ... when..., and is actually presented with the stimuli). The only sensible choice then seems to be that between achievement and activity, and brief reflection reveals discrimination to flout some necessary conditions for the latter, due to White (1964, 23-5): First, while testing for discriminability of two stimuli is something that *may take some time*, discriminating them is not. The stimuli may be presented, again and again, for five minutes, and Mr.X may take one second, each time, to make up his mind which is which. But to say that he discriminates them is to make an assertion about his tracking record, viz., that he manages (or would manage) to identify the target at a significantly higher than random rate over a series of trials. Accordingly, second, the verb *has* neither *a present continuous* nor *a past habitual tense.* As different pairs of stimuli are being tested, the subject may discriminate stimuli continually but not continuously. While looking for differences between stimuli is something he may be engaged in, discriminating stimuli is not. Third, this is not *done in a certain manner*: To ask how the subject discriminates the stimuli is not to ask whether he did it, say, eagerly or in an orderly fashion, but to ask, e.g., for an explanation of the subject's sensitivity (or that of his visual system) to the difference between the stimuli. And, fourth, it is not the sort of thing *one can have reasons for:* Mr.X may have reasons of his own for peering down the eye-piece, but not for discriminating the stimuli presented. But if discriminating thus is no activity, it seems plausible to assume it an achievement.

Inferences (A3) and (A4) keep in line with this conclusion and construe the discrimination of stimuli x and y as a (potentially unconscious) achievement of the subject, akin to the (conscious) achievement of distinguishing x and y: They address the question of how the subject achieves this, of how he works the trick, so to speak. (A3) invokes as the most plausible possibility that he achieves this by looking at the stimuli. And (A4) adduces what looks like the most plausible explanation of why the subject can accomplish the feat in this way: He can discriminate the stimuli by looking at them, because they look different to him.

Indeed, for our two inferences to work, the achievement of discrimination *has* to be attributable to the subject, rather than to 'his visual system'. Only if *he* is the one to discriminate the stimuli presented does it make sense to say, as (A3) does, that he does this 'by exercising his sense of sight', 'by looking at the stimuli': This 'by' requires that the verbs preceding and following it share the same subject. And only Mr. X, but not his visual system, can be said to '*exercise* his sense of sight' or to '*look at*' the stimuli pre-

sented. Of course we can reformulate our assumption if, after all, only 'the visual system' can be credited with the achievement of discrimination. Replacing the uncongenial 'by' in (A3), we obtain

(i) 'The visual system discriminates the stimuli presented only if the subject looks at them' or, at best,
(ii) 'The visual system discriminates the stimuli presented due to the fact that the subject looks at them.'

But even if true, neither statement tells us *how* it is that the stimuli presented to the subject are being discriminated. I.e.: If the subject himself cannot be credited with this achievement, (A3) has to be reformulated in a way that deprives the step of its alleged status as an inference to an explanation, let alone: an inference to the best explanation.

Indeed, suitable reformulation of (A3) also deprives (A4) of this status: Once the former is reformulated in one of the ways envisaged, the latter is no longer called upon to explain why the subject can discriminate stimuli by looking at them. Rather, it then is to explain why (i) the visual system discriminates them only if the subject looks at them. Note that here, in (i) and (ii), we are employing 'to look' in a new and stretched sense: Of course my visual system responds to stimuli I am not looking at. It responds to, and I see, both the pretty girl and her dog, but I would be lying if I answered anything but 'At the girl' when asked, 'What are you looking at?' If (i) or (ii) are to be true, we hence need to interpret 'to look at x' as covering all cases of (roughly) 'so directing one's gaze as to have x in one's visual field.' But the explanation of why (i) the visual system discriminates stimuli only if the subject 'looks at' them in *this* sense is straightforward enough: Otherwise, there is no relevant causal interaction between the stimuli and the system's receptor cells, as no light from the former reaches the latter. And therefore one can also say that, in a sense, (ii) the achievement of the system is 'due to' the fact that the subject 'looked at' the stimuli. But then there is no need to invoke, as in (A4), the assumption that the stimuli presented to the subject look different to him. The pertinent explanation can clearly do without anything like (A4). In other words: If at all, also (A4) is acceptable as an inference to the best explanation only if the subject himself can be credited with the achievement of discriminating the stimuli presented to him.

IV.

So to which extent, if any, can such an achievement be attributed to the subject? The obvious idea is that 'Mr.X discriminated x and y' marks a complex achievement the man himself can be credited with if he managed a series of more modest achievements, viz., if he identified the target stimulus correctly, repeatedly. For this idea to work out, the pertinent use of '*to iden-*

tify' needs to mark an achievement of the subject. It is useful to compare 'identifying' to such paradigmatic achievement-terms like 'finding', 'winning', 'arriving', 'discovering', 'concealing', or 'detecting'. The comparison yields a truism that provides us with some necessary conditions:

> An achievement is something one may try for or resolve to try for and manage or fail to obtain; it is something which may be reached by various means and methods or by luck, by one's own efforts or with help. (White 1964, 25)

Identifying seems to fit this bill: Each time the stimuli are presented to me, I can *try to* identify the target stimulus no less than the inspector can try to establish the identity of the corpse; either of us may *manage or fail to obtain* the correct identification; I may identify the stimulus *by* carefully attending to, say, its hue (while ignoring brightness and intensity); and I may thus identify it by an own effort or with help (somebody suggests I consider the hue). So it seems that if 'to identify' also takes some further hurdles, the present idea will work out, and 'Mr.X discriminated x and y' will credit the man with an achievement. Presumably, this is why Reasoning R looks so innocent.

But an achievement-term need not always be used to mark an achievement. To bring the relevant difficulty into view, consider first a paradigm then a limiting case. Winning a research prize is a paradigmatic achievement: Among other things, I can try to win by focusing on fashionable topics and employing an approach favoured by an influential referee; I can win by fair means (devising a clever counter-example to an influential assumption) or foul (networking), or by luck (a better entry is lost in the mail), and can win by my own effort or with help (the example being either due to a colleague or promoted by a referee). But unless the lottery is a fraud, it is only in a very impoverished sense (if at all) that I can be said to 'try to win' the grand lottery prize: One *might* say this if I do not merely buy a ticket in the pious hope that it will win but systematically purchase all tickets still on the market. ('He tried to win by buying up all unsold tickets.') Also, I cannot win 'by' any 'means or method', but *only* by luck: I can increase my chances, but cannot win, *by* systematic purchase of as many tickets as possible. Nor can I, strictly speaking, *win* 'by' any 'effort' or 'with' somebody's 'help': I can only buy and bankroll the tickets, but cannot win, 'by own effort' or 'with help'. Winning the lottery prize thus is, at best, a limiting case of achievement. So let's not jump to conclusions from the fact that psychologists use the achievement-term 'to identify', but consider more closely whether this verb (and, hence, 'discriminate') is used to mark achievements of subjects, in the particular experimental context in question.

When describing their performance in forced-choice experiments, the experimenters we considered describe their subjects' (say) button-pushing responses as 'identifications' of the target-stimulus, regardless of whether the

response expresses a careful choice of the subject or the sort of mere guess the rules of the experiment compel him to make whenever he finds himself unable to tell. This stretches the use of the term a bit: If the inspector concludes a 'blind' bet as to the identity of the corpse prior to looking at any evidence whatsoever, he would ordinarily not be said to have then 'identified', or even 'tried to identify' the corpse, even if eventually he wins his bet. And this linguistic innovation matters: Suppose you cannot tell any difference between the two stimuli the experimenter keeps switching around in random fashion, as you peer down the eye-piece. The rules of the experiment do not allow you to ask for *help* or to try any tricks; they oblige you to simply guess, if you cannot tell. But then you are not in a position to even '*try* to identify' the test-stimulus: You can only make a random choice and *hope* that it is right. (Here you are even worse off than in the lottery: In the experiment, you cannot buy more tickets or do anything to increase your chances.) And you cannot 'identify' the test-stimulus *by any means or method*: Reduced to a mere guess, you cannot defend yourself against the charge of having 'merely guessed' by adducing any method you followed or any means you employed. Hence, if it happens to be right, your choice will count as a lucky guess but will not qualify as an 'achievement'. Accordingly, if your guesses track the target stimulus with a significantly higher than random rate of success, your discrimination will not be an achievement. I.e.: In those cases in which the subject of the experiment is reduced to guessing, he cannot be credited with an achievement, even if the two stimuli in question prove discriminable for him.

These are the cases in which he would sincerely and carefully report both stimuli to look the same to him, at each presentation. So discrimination is the subject's achievement at best in those cases in which it would not conflict with his sincere and careful reports of matching. But both (A3) and (A4) require that the discriminations in question be achievements of the subject. So all we are, at best, entitled to conclude by Reasoning R is: (C2') When two stimuli *he does not sincerely and carefully judge to match* are discriminable to the subject, they look different to him. But then the experimental finding (E), that subjects can discriminate also stimuli they sincerely and carefully judge to match, no longer warrants the conclusion (C3) that some stimuli we sincerely and carefully judge to match look different to us.

Note that this argument is immune to play with words (due to a member of the Bielefeld audience): One might wish to insist that I can be said to 'try to win' the lottery prize even if I buy a single lottery ticket in the pious hope that it will win; and that even if I am reduced to hazarding a guess, I can be said to 'try to identify' the test-stimulus. Of course I can then identify it *only* 'by luck'. But achievements need not be meritorious, so the lucky identification is an achievement, for all that. But note that to explain *such* a (non-meritorious) 'achievement' there is no need to invoke, as in (A4), the as-

sumption that the stimuli *look* different to the subjects: This assumption does not provide any explanation of why they guess right significantly more often than not – indeed, if the stimuli did look different to them, they would not have to guess. So loose verbiage does not get us around the present conclusion: Reasoning R does not license the move from the straightforward experimental finding (E) to the rather mystifying conclusions (C3) and (C).

In passing, we already noted why this fallacious reasoning seemed plausible: Activity and achievement were the only two candidates to come to our minds, and brief reflection revealed discrimination to flout some necessary conditions for the former. We then thought this achievement to be attributable to the subject because we thought it to consist in a number of minor achievements of identification, and because 'to identify' is usually used to mark a meritorious achievement attributable to subjects. But we found that psychologists stretch the use of the verb in describing subjects as 'identifying' stimuli in forced-choice experiments: Their technical usage does not require that the 'identification' be their subjects' (meritorious) achievement. Ultimately, the plausibility of Reasoning R thus rests on a mistaken, homophonic, translation of a technical term of psychology with a verb of everyday English that is used differently – a mistake of the same sort as the conflation of the technical notion of *discrimination* with the everyday one of *distinguishing*.

The general moral is hence this: Psychologists use technical terms (e.g., 'discriminable') for which there is no short-hand in ordinary English, and use words borrowed from ordinary English (e.g., 'identify') in a regimented, deviant, or wider sense. This makes it necessary to attend carefully to the rules according to which the man in the lab and on the Clapham omnibus, respectively, uses his words, if we are to relate the results of the former (such as 'Subjects can discriminate stimuli they sincerely and carefully report to match') to the rules and convictions the latter goes by (like 'My sincere and careful report is the last word on whether things look the same to me').

IV.

So which conclusions do we now find ourselves left with, if we heed this advice? Granting subjects the last word on whether two patches of coloured light now look the same to them, we can move from Result (E) on to: (a) When pressed to indicate which of two stimuli is the previously presented test-stimulus, subjects sometimes guess right at a significantly higher than random rate, even when the two stimuli look the same to them. This strongly suggests that our visual system is capable of differential response to some stimuli that look the same to us. Indeed, psychologists since Cornsweet (1970, 217-19) argue that discriminable stimuli must have different effects at the level of retinal processing. If this is right, we may add a

second conclusion: (b) Some stimuli that have different effects at the level of retinal processing look the same to us.

(b) is pleasantly pedestrian: It is not at all levels of neural processing that differences of stimuli effects translate into differences in how things look. To use a helpful picture: Not all such differences break through to the phenomenological surface. But (a) might, at first sight, still seem perplexing. For how can our behaviour be sensitive to differences between stimuli of light, if the stimuli don't look different to us? Perplexity vanishes upon a closer look at what exactly the sensitivity established by the experiments under consideration comes down to: When asked to pick out the target from among discriminable stimuli they cannot distinguish, subjects typically feel uncomfortable about the task and find the insistence of the experimenter ludicrous, and often declare themselves unable to oblige or adopt idiosyncratic 'guessing-strategies' (cp. Marcel, 1983, 204). If our life depended on making the right choice, we would panic, as we stare at two lights that look exactly the same colour to us, and feel that we put our life in the hand of the Lord as we blindly guess which of the two is the target stimulus. If we knew whether we can discriminate a given pair of stimuli that look the same to us, we would develop some confidence in our guesses, and maintain it as long as we were assured that we continue to be exposed to the same stimuli. But of course we cannot tell whether we can, or cannot, discriminate any such pair of stimuli. So, unless the experimenter tells us, we continue to regard our 'choices' as entirely arbitrary – even though they actually are right significantly more often than not.

Once we realise that *this* sort of sensitivity to the physical differences between stimuli is all that is established by the experiments we considered, conclusion (a) seems no more mystifying than (b): *Such* behavioural sensitivity is obviously consistent with the truth of the subject's report, that both stimuli look exactly the same to him. I.e.: *Such* sensitivity does not challenge our common conviction that under certain conditions first-person authority holds for the phenomenological features of our current perceptual experience. The failure of Reasoning R thus allows us to replace a mystifying conclusion, (C3), by two others, (a) and (b), that strike us as pedestrian, at the latest once we are clear on their content. And this takes the mystery out of the experimental finding (E) that initially puzzled us: that subjects can discriminate stimuli they sincerely and carefully report to look the same to them.

Eventually, (a) and (b) suggest a more general point: They might jointly make us inclined to say that, even though we don't 'consciously perceive' differences between such stimuli, we 'perceive' them 'unconsciously'. What we then see is that *such* 'unconscious perception' makes no difference to the way things look to us. And this might well also hold for blindsight and further phenomena we are inclined to label instances of 'unconscious perception'. Indeed, the analogous treatment of discrimination in blindsight is now

obvious. What the discrimination experiments on blindsight patients reveal is: (a') When pressed to indicate which of two stimuli presented in the blind areas of their visual fields is the previously presented test-stimulus, such patients (are reduced to guessing, as they cannot see either of the stimuli presented, and) sometimes guess right at a significantly higher than random rate. This strongly suggests that: (b') Also stimuli blindsight patients cannot see have (differential) effects at the level of retinal processing. (I.e.: The relevant pathology is not located at this level.) Again, the 'unconscious perception' in question makes no difference to the way the world looks to subjects. To take up the previous helpful picture: Unconscious perception need not make a difference to the phenomenological surface.[1]

Notes

[1] The present material is included in the more comprehensive paper 'Discrimination: A Challenge to First-Person Authority?', in: *Philosophical Investigations*, vol. 25, January 2002. I am grateful to the editor and the publisher for their kind permission to reproduce the material here. © Basil Blackwell Ltd.

References

Clark, A. (1993): *Sensory Qualities*, Oxford: Oxford University Press

Cornsweet, T. (1970): *Visual Perception*, London: Academic Press

Marcel, A.J. (1983): Conscious and Unconscious Perception, *Cognitive Psychology 15*, 197-237 (Part 1 - Experiments on Visual Masking and Word Recognition) + 238-300 (Part 2 - An Approach to the Relations between Phenomenal Experience and Perceptual Processes)

Nelkin, N. (1993): What Is Consciousness?, *Philosophy of Science 60*, 419-34

Nelkin, N. (1995): The Dissociation of Phenomenal States from Apperception. In: Metzinger, T. (ed.), *Conscious Experience*, Paderborn: Schöningh

Stoerig, P. & Cowey, A. (1992): Wavelength Discrimination in Blindsight, *Brain 115*, 425-44

White, A.R. (1964): *Attention*, Oxford: Blackwell