

## Comments on François Recanati's *Mental Files*: Doubts about Indexicality

**David Papineau**  
King's College London

DOI: 10.2478/disp-2013-0015

BIBLID [0873-626X (2013) 36; pp. 159-175]

François Recanati has written a valuable and timely book. The suggestion that thought involves the deployment of 'mental files' has been around for some decades, and is becoming increasingly popular, but until now there has been no sustained examination of the idea. Recanati has developed a detailed theory of mental files, and future treatments will take his book as their starting point.

I am very much in favour of the general lines of Recanati's approach. In particular, I fully support his policy of dealing with Frege cases, not by introducing some extra semantic level additional to referential value, but simply by appealing to the possibility that different *vehicles* of thought can be used to refer to the same entity. Recanati shows convincingly that a multiplication of descriptive senses or extra 'intensions' is quite unnecessary to deal with Frege cases. All we need is the idea that distinct mental files can refer to a single entity.

I also fully support Recanati's decision to focus on the individual rather than the community. Recanati's mental files are possessed by individuals, and he offers no general account of when different individuals might be said to 'grasp the same concept'. In this respect his book contrasts with Mark Sainsbury and Michael Tye's recent *Seven Puzzles of Thought* (Oxford University Press 2012), which is in many respects consistent with Recanati's approach, but which seeks primarily to articulate a notion of a public concept rather than that of an individual's mental file. I myself am very doubtful that any one notion can do justice both to the public and individual dimensions of thought; indeed I remain to be convinced that there is any real work for the idea of a public concept, once we have a good account of individual mental files and the use of words to communicate them. So

I think that Recanati does well to focus on the individual rather than the social level.

However, as is only to be expected, there are aspects of Recanati's book with which I disagree. In particular, I think that his adoption of an 'indexical model' for mental files leads him astray in various ways. In these comments I shall focus on his use of this model. In the first two sections below I shall point out some ways in which this model can be misleading. After that I shall argue that, when it comes to 'demonstrative files', the model is not only misleading but positively erroneous.

## 1 The indexical model

Recanati models the workings of mental files on the way that indexical *words* function in language.

Linguistic indexicality is familiar enough. Indexical word types — such as 'here', 'now', 'I', 'you' — do not have referents to call their own. There is no particular time referred to by the word type 'now', nor any particular person referred to by the word type 'you'. In this respect, indexical word types contrast with proper names types — 'Timbuktu' — or natural kind term types — 'gold' — which have standing referents, so to speak.<sup>1</sup>

It is only *tokens* (specific dated uses) of indexical words types that have referents. The way this works is familiar. The indexical word types have a 'character' that specifies how the context of utterance of any *token* of that type will determine a referent for that token. Thus the character of the type 'now' specifies that any token use of this word type will refer to the time at which the relevant utterance is being made; similarly the character of the type 'you' specifies that any token use of this word type will refer to the person to whom the relevant utterance is being addressed; and so on.

Recanati's thought is that mental files work like indexical words. Token mental files fall into types, depending on the kind of 'epis-

<sup>1</sup> Perhaps proper names are not so non-indexical as they initially appear. More than one person is called 'David Papineau'. One account of how such proper names manage to refer is that 'David Papineau' is a type the tokens of which get attached indexically to specific people on occasion of use (Pelczar 2001).

temically rewarding relation' that the token bears towards its referent. For example, tokens of the perceptually demonstrative mental file type *that man* will bear a potentially fruitful epistemic relation to the man to whom the relevant thinker is currently attending. This species of 'epistemically rewarding relation', in conjunction with the relevant thinker's context, then determines a referent for any token of *that man*, analogously to the way that the character of an indexical word type plus a context determines a referent for any indexical word token. Similarly, tokens of the mental file type *here* will refer to the place where the thinker is currently located; and tokens of the mental file type *I* will refer to the person who is using that token in thought.

Recanati applies this indexical model to mental files in general. He argues that even his 'encyclopaedic' files conform to the model. Encyclopaedic files are distinguished by the fact their existence does not depend on any specific epistemic relation to their referent. My *Barack Obama* file can survive the loss of any specific epistemic rewarding relationship to him (I might forget his name, or alternatively forget what he looks like and cease to be able to identify him visually) as long as some (any) such rewarding relationships remain. For Recanati, this is enough to bring encyclopaedic files under the indexical model. The crucial point, as he sees it, is that their referents are still contextually determined. The referent of my *Obama* file is determined as that item in my context to which the file bears some (any) epistemically rewarding relations.

## 2 There is less mental than linguistic indexicality

I have no objection to the idea that the referents of mental files are generally contextually determined, and Recanati is of course free to use the term 'indexicality' to express this idea if he wishes. But there is a danger that this usage will obscure the fact that there is a lot less indexicality in the mental realm, so to speak, than Recanati's analogy between mental and linguistic 'indexicality' might lead one to expect.

Consider first Recanati's '*I*' files. These are mental files which each subject possesses, distinguished by the special epistemic relationship of self-knowledge that each subject bears to itself. I have my

*I* file, you have yours, and so on. In the normal case, each of us uses our *I* file throughout our life as a repository in which to accumulate information about ourselves. These are the tokens of the *I* file type: that type, to be specific, whose tokens acquire their reference in virtue of bearing the self-knowledge relationship to particular thinking subjects.

Now consider the English word 'I'. The character of this word type specifies that any of its tokens will refer to the user of that token. So on many occasions English speakers will use the type word 'I', and on each of those occasions the token so uttered will refer to whomever is speaking. In this case, and by contrast with mental *I* files, the normal understanding is that even the different uses of 'I' by a given individual on different occasions will each comprise different *tokens* of the type. If I use 'I' this evening when talking to my wife on Monday, and then again on Thursday when talking to a student, these are naturally taken to be two different tokens of 'I', each separately assigned me as referent in virtue of the principle that all such uses refer to the speaker.

It is understandable enough that we should so cut things up differently in the mental and linguistic cases. Mental *I* files normally function as repositories for all the information that individuals acquire about themselves in the course of their life. Viewed in this way, it is essential that they persist for as long as their possessors survive. They need to be able to accumulate information over lifetimes.

'I' words, by contrast, have no such corresponding function. There is no obvious sense in which bodies of information get attached to tokens of the 'I' word type, and a fortiori no sense in which an accumulating body of information is attached to all of a given person's 'I' utterances. So there is no reason to lump all of a given person's 'I' utterances into different manifestations of a single 'I' token, in the way that Recanati lumps together all a given person's rehearsals of their mental *I*.<sup>2</sup>

<sup>2</sup> Perhaps this is a bit quick. If I use the type word 'I' repeatedly in the course of a *single* conversation, then shouldn't we count these as different *uses* of the *same* token, on the grounds that in general intra-conversational uses of indexical words are most naturally read, given the conventions of anaphora, as repetitions of the same token term (and so as 'de jure' co-referential)? I am happy to grant this. But my more general point remains. For these considerations provide no argument

The upshot is that mental *I* files display a lot less indexicality than linguistic ‘I’ words. Whereas I must have deployed millions of linguistic tokens of the type word ‘I’ in the course of my life, I have only ever deployed one mental token of the *I* file type.

Now consider encyclopaedic files like my *Obama* file. For Recanati, this is an ‘indexical’ file in virtue of having its referent fixed contextually (as that item to which the file bears some epistemically rewarding relations). Again, I worry that this terminology may be misleading about the structure of encyclopaedic files.

Let us ask how encyclopaedic files are supposed to conform to the type-token structure displayed by other indexical constructions. My personal *Obama* file is presumably a token of the relevant type. This token will persist as long as I remain able to think of Obama, and will serve as a repository for all the information that I accumulate about him in this time. It is this token whose reference is contextually fixed as Obama himself.

But what now is the type of which this personal file is a token? A first thought might be that it is the category which contains all the other token personal encyclopaedic files possessed by people who can think about Obama, whether by recognizing him, or by knowing his name, or any mix of these and other epistemically rewarding relations. But this seems wrong. After all, those tokens all have the *same* referent, namely Barack Obama. And given this there seems nothing to stop us saying that the type itself has this standing referent. But this is in tension with the idea that indexical types have no referent to call their own.

As far as I can see, the only good way to fit the personal encyclopaedic *Obama* files into the standard indexical type-token structure would be to view them as tokens of the type *encyclopaedic file* (or perhaps *encyclopaedic person* file). This type would have no reference of its own, and each of its tokens (such as an *Obama* file, or a *my first teacher* file, or a *that-woman-down-the-road* file...) would then have its referent fixed as that thing (person) in the relevant thinker’s environment to

---

for identifying tokens of ‘I’ *across* conversations. If I use ‘I’ on Monday in one conversation, and then on Thursday in another, these are surely *two* tokens each separately assigned a reference in context. So we still have a contrast with mental *I* files, where the *same* token of the *I* type needs to stay with me from Monday to Thursday, and indeed for the rest of life.

which the relevant file bears some epistemically rewarding relation.

This would work all right, but it would be strange, and certainly wouldn't line up with any indexical constructions present in natural languages. We certainly don't have some type *word* the tokens of which refer variously to Obama, my first teacher, that woman down the road, ..., depending on the context in which those tokens are uttered.

This is not yet a substantial criticism of Recanati's account of encyclopaedic files. I have no objection to his central idea that they have their referents fixed contextually. Still, it does seem unhelpful to use the term 'indexical' to express this idea. If such paradigmatically permanent files as my *Obama* file are classed together with 'here' and 'now' files as similarly 'indexical', I cannot help feeling that some important distinctions are being lost.

### 3 *You* thoughts

So far I have argued only that Recanati's 'indexical' model for mental files needs to be treated with care, given the extent to which some of his mental files display rather less indexicality than this terminology might initially lead readers to expect.

However, I worry that there is a more substantial danger in Recanati's emphasis on indexicality: the thought that mental files share the structure of linguistic indexicals can encourage us to multiply mental files beyond necessity. In this section I shall illustrate this danger by considering the possibility of 'you' files. In the next section I shall apply the morals I draw to 'perceptual demonstrative' files. (While Recanati does not himself posit 'you' files, 'perceptual demonstrative' files play a central role in his project.)

Suppose I am thinking that John Colleague gave a good talk yesterday, and that I must tell him this next time I see him. At that moment, lo and behold, he comes into the common room. 'You gave a good talk yesterday', I say to him.

Now, my utterance unquestionably contains an indexical type *word*, 'you', tokens of which are conventionally understood as referring the person to whom the utterance is addressed. But should we think of me as expressing a thought involving a correspondingly indexical *you* file? When John looms into view, do I form a token

thought file of the *you* type, a distinctively 'second-person' file whose reference is fixed via the relation it bears to the addressee of the thinker's current utterance?

This would seem an unattractive move (and indeed Recanati does not make it). I don't need to form any extra *you* file when I am about to address John, beyond any files about him I already have. So there is no reason to suppose that my utterance 'You gave a good talk yesterday' expresses some corresponding indexical *you* thought. Rather it is just the linguistic means that I use to express a pre-existing non-second-person thought.

Recanati agrees, in the course of discussing how thoughts get communicated in speech. His view (222-3) is that when a speaker's utterance *u* contains 'you', the information that the hearer 'is the addressee of *u*' will be in the file that the speaker is using to think about the hearer, and also in the hearer's first-person I file. This information will 'stand for' the relevant files, and 'trigger their activation', but it will not contribute to the content of what is being communicated.

This seems quite right to me. We don't need *you* files to explain the content of what speakers say about their hearers.

Still, perhaps the point bears a bit more examination. Note that *something* inside the speaker needs to figure out that, in the context, 'you' is the appropriate word with which to give public verbal expression to the thought that, say, *John gave a good talk yesterday*.

Presumably, the production of the relevant utterance is informed by the speaker's information that (a) John is the (potential) addressee, and (b) 'you' is the right word to communicate to a current addressee any thought that refers to that same addressee.

But if this is right, then it seems that there must be re-usable mental term *current addressee* in the speaker's mental economy, with which to formulate the information in (a) and (b). And this then argues that we should recognize that in thought there is a *type* of referring mental term (*current addressee*), *tokens* of which are used on particular occasions to mediate inferences about when to use the word 'you' to express some thought.

So it does look as if there is indeed a referring mental type (*current addressee*) which works very like the linguistic type indexical 'you', in that tokens of both will refer to whoever is currently (potentially)

being addressed.

Still, even though we do need to recognize these re-usable mental types, there is no reason to think that they should be dignified as *files*, nor that they play any significant role in *thought*.

As I see it, the mental type in question operates mainly in the sub-personal speech production system. Once this speech-production system ‘knows’ that *John is the current addressee*, it will set itself to express all John-referring thoughts using the word ‘you’. In the normal case, I don’t need to think about how to express my thoughts. My selection of words is generated automatically, courtesy of an automatic and unconscious system that figures out what grammatical string of words will best serve to express my thought in the current context.

True, there will need to be some interaction between the sub-personal speech production system and personal level conceptual thought, in order to derive the crucial interfacing information that, say, *John is the current addressee*. We can’t eliminate conceptual thought in deriving this conclusion, for after all there is no limit to the kind of conceptual information about John that might be relevant (suppose you know that John often goes around in disguise, but can’t resist custard cream biscuits ... and you use this information to figure that the strange bearded man in the common room must be John.)

Still, this doesn’t mean that the term *current addressee* itself needs to function as an information-accumulating file in our person-level cognitive economy. Once a judgement like *John is the current addressee* has been arrived at, it will be handed over to the speech production system, and this token of *current addressee* will cease to play any role in person-level thought.<sup>3</sup> Moreover the speech production system itself certainly won’t treat this token as some file whose function is to gather and preserve any further information about its referent. Its only use for this token is to register that John is the current addressee and thence direct that *John* thoughts should be expressed with ‘you’.

<sup>3</sup> This argues against Recanati’s suggestion that some such item as ‘is the addressee of *u*’ will generally get entered into speakers’ files for the person in question. Since this information is only needed by the unconscious mechanisms that put thoughts into words, there is no obvious reason to keep a record of it. Perhaps this is part of the explanation of why I can often remember conversations in some detail without being able to remember whom I had them with.



So, all in all, even though we need to recognize that there are referring mental terms *current addressee* with an indexical structure like ‘you’, these terms will not function as mental *files*.

#### 4 Demonstrative files and their problems

Let me now turn to demonstrative files. This is where I think that the indexical analogue is genuinely damaging. It encourages the view that there are token mental files corresponding to token linguistic demonstratives, when in truth there is nothing corresponding in our actual cognitive structure.

Recanati has a distinct species of demonstrative files (*that thing, that woman*) which are opened when a thinker is in perceptual rapport with some item, and which survive and accumulate information as long as that rapport is maintained.

An immediate query about these files concerns cases where we think about some previously perceived item even though we have ceased to be in continuous perceptual contact with it. So for example I might recall the woman I saw this morning, along with the information I then acquired about her. Or I might recognize that woman when I see her again this afternoon, again remembering the information I acquired earlier.

A natural first thought is that in such cases we reactivate the demonstrative file that we opened when we first saw the woman this morning. This would explain the current availability of any information we acquired in that earlier encounter. However, Recanati cannot say this. The demonstrative file that I originally formed disappears along with the termination of the epistemically rewarding perceptual contact on which it is based.

Recanati’s response to this query is to multiply files. As well as the original perceptual demonstrative file, I will also have a ‘memory demonstrative file’ (62), and a ‘recognitional file’ (71). Memory demonstrative files exist in virtue of an epistemically rewarding *memory* relation that the thinker bears to the relevant item; as long as one can remember the earlier encounter, one can think of the woman in question as *that woman [whom I saw]*. Recognitional files exist in virtue of an epistemically rewarding relation of *familiarity*; as long as one is capable of recognizing the item in question, one can think of it via a

recognitional file.

So, on Recanati's view, when I later recall the woman that I saw this morning, or later recognize her, I am not reactivating my original perceptual file, but rather activating new and different files, a memory demonstrative file, or a recognitional file.

However, now Recanati faces a different query. When I later think about the woman, either via memory or via recognition, I will presumably have available all the information about her that I acquired from my earlier perceptual encounter this morning. But Recanati cannot take this as given. That earlier-acquired information was originally deposited in my perceptual demonstrative file, a file that is distinct from memory file and recognitional file, and indeed no longer exists. So there is no immediate guarantee that the information it contained will be available elsewhere.

Recanati's response is that the relevant information from the original demonstrative file will be *transferred* to the memory file and recognitional file. These latter files will inherit the information originally deposited in the perceptual demonstrative file, and so will be able to activate it in thought, and augment it via further encounters with the referent.

## 5 A simpler view

Well, all this adds up to a cogent story, but it strikes me as gratuitously complicated. Why have so many files when one would do? Here is an alternative picture. When I first encounter some item perceptually, I open a potentially permanent file in which to accumulate information about that item. That file outlasts the original encounter, and the same file is reactivated when I remember the relevant item or re-encounter it. The information earlier acquired is thus automatically available on those later occasions, and can be added to when new facts are acquired, without any need for any multiplication of files. (See Papineau 2006.)

On this view, the files that we open on first perceptual encounters, and in general on coming into any contact with any new item of thought, are *name-like*. They are designed to be permanent repositories of information about the item in question, and are not dependent on any particular sources of information about that object. In

this respect they are akin to Recanati's 'encyclopaedic' files, whose function is to gather information about some referent from whatever sources offer themselves.

Of course, there will be occasional cases where we open two such files for what is in fact the same referent. I encounter a woman (or a tree, or a chair, ...) and then later on I encounter it again without realizing that it is the same one. Or I already have a well-developed file for *John Perry*, say, and then don't realize that the man I am talking to at the party is him. But nothing in this requires us to multiply *types* of file. In such cases we will simply have two name-like files containing different bodies of information that we don't yet recognize are co-referential. And, if we do later realize that the two files refer to the same thing, then we can merge them, or more cautiously link them, and thereby bring the two bodies of information together.

Many of the name-like files that we open in this way will prove temporary. Not every perceived tree that we have occasion to think about — or chair, or coffee cup, or indeed person — will prove worthy of a lasting entry in our mental filing system (or turn out to be the same as something for which we already have a lasting entry). And in some such cases the files we have opened for these things will no doubt atrophy away and cease to be available for forming thoughts. But again this doesn't require us to multiply *types* of file. There is no need to view the files that get closed down as special demonstrative files which by their nature cease to exist once their defining epistemic relation is lost. From my perspective, there is no constitutive feature of the closed-down files that prevents them persisting indefinitely as repositories of information about their referents. The reason they get closed down is not that they cannot survive the loss of some epistemic relation, but simply that they have faded away from disuse.

## 6 More complications

Recanati's multiplication of files generates even more difficulties than those I have drawn attention to so far. Suppose I remember that *Paul Churchland is tall*. Then I recognize Paul at a conference *and* note that *Paul Churchland has a beard*. I conclude, quite logically, that *someone is tall and has a beard*. However, if my memory demonstrative

file and my recognitional file are different files, as Recanati's story has it, then this inference is invalid, for nothing yet guarantees that *Paul-remembered* is the same person as *Paul-recognized*. We would seem to need an additional premise, to the effect that *Paul-remembered* = *Paul-recognized*. (As Recanati puts it, arguments that use the *same* file throughout are de jure entitled to presumptions of co-reference. But when different files are in play we need extra information to establish co-reference de facto.)

However, as Recanati himself agrees, it is highly implausible that my argument that *someone is tall and has a beard* stands in need of any extra premise that *Paul-remembered* = *Paul-recognized*. When I see Paul at the conference, I surely recognize him *as* the Paul that I can remember, and don't need explicitly to judge that the Paul I am recognizing to have a beard is the same as the one I remember to be tall.

Recanati's response to this extra problem is to introduce a yet further file, a *recognitional-demonstrative* file. This file is activated when you re-encounter someone whom you could previously remember, and is distinguished by the fact it now enjoys *two* epistemically rewarding relationships with its referent, the memory relationship and the current perceptual relationship. As a result, *this* file will acquire both the information that *Paul Churchland is tall* and that *Paul Churchland has a beard*, and since this information is now housed in a single file we will have the desired de jure presumption of co-reference to draw the desired conclusion, without any need of an extra identity premise.

But once more this multiplication of files seems quite unnecessary, forced on us only by Recanati's commitment to the idea that different files are constituted by different epistemically rewarding relations to their referents. If we drop this idea, then there is nothing to stop us supposing that I have always had one name-like encyclopaedic Paul Churchland file that I originally formed when I first read *Scientific Realism and the Plasticity of Mind* in the early 1980s, and into which I have since placed all my Paul Churchland information acquired from whatever sources, including the perceptual sources I came to be able to use once I met him.

## 7 Why not be simple?

I can see no advantages in multiplying files in the way that Recanati does. When Recanati first defends the idea that demonstrative files die off, to be succeeded by memory files and recognitional files, with consequent transfers of information, he appeals, quoting Frege, to the fact that certain *linguistic expressions* need to be replaced in a corresponding way.

‘If someone wants to say today what he expressed yesterday using the word ‘today’, he will replace this word with ‘yesterday.’ (81)

Well, this is of course true. But it does not bear on the point at issue. We can all agree that we often *express* our thoughts using type *words* whose reference depends on the context of utterance, and thus need to shift type words to keep referring to the same entity when the context of utterance changes. But it does not follow that the mental files constituting the thought expressed need to be shifted similarly. (Indeed the rest of the quotation Recanati takes from Frege makes just this point: ‘Although the thought is the same its verbal expression must be different...’)

We often use indexical words when it is clear that there is nothing correspondingly indexical in our thoughts. We have already had the example of ‘you’ expressions of non-second-person thoughts. There are plenty of other similar cases. You don’t know Jane’s name, though I do; I know that you can see her, though I can’t. ‘*That woman* is an eminent barrister’ I tell you, invoking our mutual knowledge that you are looking at her. The expression may be indexical, but the thought I am expressing involves my permanent encyclopaedic file for Jane, not some temporary demonstrative file (remember that I myself can’t see her). There is nothing indexical about my thought. Even though I express my thought using the phrase ‘that woman’, my thought itself involves my permanent file for Jane, which does not shift reference with context in the same way as the phrase ‘that woman’ does.

Examples could be multiplied. There is a surprisingly widespread tendency to infer, from the use of indexical words to express some thought, that the thought expressed must be similarly indexically structured. But it does not take much reflection on cases to show

that this inference is generally invalid.

Is there any other reason to suppose that we have short-lived ‘demonstrative’ mental file tokens of the kind that Recanati posits? A strongly verificationist account of concepts might hold that every distinct criterion of application demands a distinct concept, in order to rule out irresolvable disagreements. But even the logical empiricists rejected this extreme ‘operationalism’ on the grounds that it would require an absurd proliferation of concepts, and were happy to recognize concepts with multiple criteria of application. Similarly, in the present context, there seems no rationale for requiring that every ‘epistemically rewarding relation’ generates its own mental file. Why not simply allow that many such relations can become attached to stable encyclopaedic files?

## 8 Perception and action

A rather different thought would be that we need to recognize demonstrative files, not because they have distinguished epistemic inputs, but because they have special behavioural outputs.

Suppose that I want to grab, or parry, or move away from or towards some specific physical item. To guide my behaviour, won’t I need to think of it as *that thing there* [*that I can see/feel/hear*]?

I think that this is indeed right, and that it does argue in favour of a re-usable type of mental term (*that thing there*), tokens of which are used on particular occasions to guide behaviour. But just as with the *current addressee* terms discussed earlier, there is no reason to view these terms as associated with any *files*, nor to suppose that they play any significant role in *thought*.

I take it that the direct control of fine-tuned motor behaviour is managed by an automatic sub-personal system, analogous to the system that determines which words we use to voice our thoughts. This motor control system will respond to directives like *grab that thing there* [*that I can see/feel/hear*], and to this extent will indeed deploy tokens of a type mental term *that thing there*. But this automatic motor control system is not in the business of storing information about the things it refers to, and so will not have any information-accumulating files associated with its tokens of *that thing there*.

Just as with the speech production system, the motor control system will need to interface with conceptual thought along the way to arriving at instructions like *grab that thing there*. Suppose I want to look up a passage in *Naming and Necessity*. I will need physically to get hold of the book and leaf through it. In order to do this I need to arrive at a judgement that *Naming and Necessity is that thing there*, which I can then put together with my desire to *grab Naming and Necessity* to generate the motor instruction *grab that thing there*.<sup>4</sup>

As before, there is no limit to the kind of conceptual information that might help generate the judgement *Naming and Necessity is that thing there*. In a simple case, I might simply see that *Naming and Necessity* is in its normal place in my bookshelf. But I might also note that a blue hardback is on the common room table, and remember that John Colleague had told me that he hadn't been able to find his copy of *Naming and Necessity* since he'd taken it to the common room to show the new lecturer that he didn't understand Kripke...

Still, once I have generated the information that *Naming and Necessity is that thing there*, by whatever means, and thence generated the instruction *grab that thing there*, I can simply hand matters over to the behaviour-control system. We can think of this system as representing the book as a grabbable item standing in a specific relation to my body, limbs and possible behaviour. To this extent the automatic system will be representing the book in the same way as it would represent any similarly sized and shaped item that is similarly grabbable. That is, it will be using a temporary token of a type representation, a type that may well be re-used on other occasions when a similar object is to be grabbed. But this automatic system won't treat this token as a file in which to accumulate information about its referent. After all, the behaviour-control system doesn't want to know anything about this referent, beyond its egocentric location and grabbability. That is why it is perfectly adequate for its purposes simply to represent it using a temporary token of the type *that thing there*.

<sup>4</sup> It is tempting to view conscious perception as the medium of this interface. It seems plausible that conscious perception simultaneously represents entities *conceptually*, as re-identifiable items about which we have stored information, and *egocentrically*, as items to which we bear such-and-such a current spatial relation.

I have already argued that personal-level conceptual thought has no good use for temporary files corresponding to such temporary tokens of *that thing there*. As I pointed out in section 6 above, at the personal level it makes far more sense to coin non-temporary name-like files for the things that we perceive, files which can outlast our perceptual contact with those items and in which we can preserve any information we glean about them. At the personal level, temporary files which do not outlast perceptual contact would simply generate extra cognitive work to no good advantage.

So, to sum up, reflection on the nature of behavioural guidance does indeed point to the existence of ‘perceptual demonstratives’ that come in types and whose tokens do not survive the loss of perceptual contact. But there is no reason why these terms should function as mental *files* of the kind Recanati is interested in. The function of these terms is to direct the motor control system to perform certain types of behaviour, not to accumulate information about their referents.

## 9 Conclusion

Recanati is greatly to be thanked for developing a detailed theory of mental files. His book will bring shape to the debate on this topic and define the agenda for future discussions. In my view, however, his emphasis on the ‘indexicality’ of mental files is misplaced. At best this emphasis is misleading, and at worst it leads to the postulation of far more mental files than are needed.<sup>5</sup>

David Papineau  
Room 509, Philosophy Building  
King’s College London  
Strand  
London, WC2R 2LS  
United Kingdom  
david.papineau@kcl.ac.uk

<sup>5</sup> I would like to thank Mark Textor for helpful comments on an earlier version of this paper.



*References*

- Papineau, David. 2007. Phenomenal and Perceptual Concepts. In *Phenomenal Concepts and Phenomenal Knowledge*, ed. by Alter, Torin and Walter, Sven. Oxford University Press.
- Pelczar, Michael W. 2001. Names as Tokens and Names as Tools. *Synthese* 128