MIT Press • Open Encyclopedia of Cognitive Science

Personal/Subpersonal Distinction

Zoe Drayson[1]

[1]University of California, Davis

MIT Press

Published on: Jul 24, 2024

DOI: https://doi.org/10.21428/e2759450.8593aeda

When we say that someone recognizes a famous painting, prefers Mexican food, or judges the winner of a competition, we are attributing cognitive capacities to the person. These *personal-level* attributions of cognition can be contrasted with the *subpersonal-level* attributions made by cognitive scientists when they claim that the fusiform gyrus in the brain recognizes faces, that pyloric neurons prefer a certain frequency, or that the early visual system judges depth from retinal disparity. In each of these latter cases, the cognitive capacity (e.g., recognition, preference, judgment) is attributed not to the person but to some part of their cognitive system. This distinction between personal-level and subpersonal-level attributions of cognition raises interesting questions, including about the relationship between personal-level and subpersonal-level attributions of cognition, and whether the personal/subpersonal distinction picks out two different kinds of cognitive processes or merely reflects two different kinds of explanatory projects we might have.

History

Personal-level attributions of cognitive capacities are familiar from our everyday explanations of others' behavior in terms of their abilities to reason, imagine, remember, plan, and decide. Subpersonal-level attributions of cognitive states, however, are a relatively recent scientific development. These arose during the *cognitive revolution*, when psychologists such as Miller (1956) and Broadbent (1958) started to model the brain as unconsciously storing, retrieving, and transforming information in the manner of a physical computer. A well-known example is Baddeley and Hitch (1974)'s model of working memory, on which a person's capacity to engage in mental arithmetic or follow directions, for example, is broken down into the different capacities of separate components or subsystems. One component, the phonological loop, stores and rehearses speech-like sounds; a second component, the visuo-spatial sketchpad, processes and manipulates visual information and spatial movement; and a third component, the central executive, coordinates the first two systems and discriminates between relevant and irrelevant information. The separate subsystems are described cognitively (e.g., as rehearsing, manipulating, discriminating) but the cognitive capacities in question are ascribed to subpersonal components rather than directly to the person.

Philosopher Daniel Dennett first introduced the personal/subpersonal distinction in *Content and Consciousness* in 1969, noting that in addition to our everyday accounts of people's cognitive capacities, there are also scientific accounts which focus on "the sub-personal level of brains and events in the nervous system" (Dennett, 1969, p. 93).

The so-called Pittsburgh School of philosophers (including Hornsby, 2000; McDowell, 1994) interpreted Dennett to be drawing a clear distinction between two very different kinds of explanations: normative explanations [see Normativity] about how we ought to think and act as rational beings (the personal level) and scientific explanations that describe how the mechanisms in our brains actually work (the subpersonal level). In doing so, the Pittsburgh School took the personal/subpersonal distinction to be a version of Wilfred Sellars (1956)'s distinction between a "space of reasons" and a "space of causes" (Drayson, 2012). This led them to

question whether there could even be a science of cognition that reduced human rationality to computational mechanisms in the brain.

Subpersonal-level attributions of cognitive capacities faced several other theoretical objections. Some philosophers argued that it is simply wrong to describe the brain (rather than the person) as evaluating or discriminating, drawing on Wittgenstein's claim that "only of a living human being and what resembles (behaves like) a living human being can one say: it has sensations; it sees; is blind; hears; is deaf; is conscious or unconscious" (Wittgenstein, 1953, section 281). Cognitive psychology was also accused by some philosophers of committing the *mereological fallacy* by assuming that a term ascribed to a whole entity is therefore ascribable to one of its parts (Bennet & Hacker, 2003).

Against these criticisms, Dennett argued that the explanatory work done by subpersonal attributions of cognition is what allows psychologists to ascribe cognitive capacities below the level of the whole person without thereby committing some fallacy or error:

> It is an empirical fact, and a surprising one, that our brains — more particularly, parts of our brains — engage in processes that are strikingly like guessing, deciding, believing, jumping to conclusions, etc. And it is enough like these personal level behaviors to warrant stretching ordinary usage to cover it. (Dennett, 2007, p. 86)

While the practice of attributing subpersonal cognitive capacities is now commonplace, questions remain about how to interpret the personal/subpersonal distinction (Drayson, 2014). When neuroscientists describe receptive-field neurons as "preferring" certain spatial orientations, for example, should we take their claims literally as evidence that cells can have preferences (Drayson, 2020; Figdor, 2018)? Do Libet et al. (1983)'s measurements of readiness potential in the brain suggest that there is no coherent self to be found at the personal level (Rupert, 2018)?

Core concepts

Explanatory levels

Since the cognitive revolution in psychology and the birth of cognitive science, it has been widely acknowledged that there are many different explanatory levels at which we might try to understand the mind: in terms of first-person experience, information-processing functions, or cellular mechanisms in the brain, for example (Colombo & Knauff, 2020). The standard view of the personal/subpersonal distinction takes it to be a distinction between two levels of explanation.

Examples of explanations at the personal level include our everyday explanations of people's behavior in terms of their cognitive capacities (e.g., to perceive depth, to comprehend language) and their propositional attitudes (e.g., the belief that the clock is slow, their hope that the meeting starts on time). Much of contemporary

cognitive science follows Dennett in assuming that we can apply a similar kind of explanation at the subpersonal level to parts of our cognitive system: Cognitive scientists talk about the *goals* of the visual system or the *calculations* performed by motor processes, for example, and describe neural populations as *representing*.

The existence of cognitive science owes much to these sorts of subpersonal-level explanations: they act as a bridge between our everyday psychological explanations of people and our neurophysiological explanations of their brains. By distinguishing between the personal and subpersonal explanatory levels, we can make clear when we are ascribing cognitive terms (such as recognizing, preferring, and judging) to the individual or person, and when we are ascribing these terms to information-processing subsystems or neural mechanisms.

The interface problem

One of the biggest challenges raised by distinguishing personal-level and subpersonal-level explanations is to say how they relate to each other. This is known as the *interface problem* (Bermudez, 2005): How do our everyday personal-level explanations of people's behavior interface with the subpersonal-level explanations of cognition found in the mind-brain sciences?

One way to address the interface problem is through the sort of functional decomposition found in Baddeley and Hitch (1974)'s model of working memory, in which the cognitive capacity of working memory is decomposed into separate functions including the phonological loop and the visuospatial sketchpad. The phonological loop can be further decomposed into the acoustic store and the articulatory loop, and the visuospatial sketchpad can be further decomposed into a visual cache and an inner scribe. On this account, the cognitive capacity at the personal level is explained by the combination of multiple subcapacities at the subpersonal level. If each capacity or function is decomposed into progressively less-complex capacities or functions, worries about explanatory regress can be avoided because the process will eventually bottom-out in neural structures.

An alternative way to address the interface problem is through the use of symbolic computation. The sorts of inferential thought processes that appear in explanations at the personal level, such as planning and decision-making, can be characterized as algorithmic transitions between symbols (Fodor, 1975). These symbols can themselves be physically implemented by neural properties, just as the 0 and 1 symbols of binary are physically implemented by voltages in computer hardware. On this approach, the relationship between personal-level explanations and subpersonal-level explanations is one of mirroring rather than decomposition.

We might, however, attempt to address the interface problem by explaining a person's cognitive capacities directly in terms of neural structures and processes. On this approach, personal-level explanations could be seen as characterizing the higher-level complex patterns that emerge from the more basic interactions between lower-level neurophysiological states. This approach attempts to account for our mental lives directly in

neurophysiological terms, rather than appealing to abstract intermediaries such functions and symbols, and so omits a separate cognitive level of explanation.

Each of these approaches attempts to solve the interface problem by offering a scientific account of how personal-level phenomena arise from subpersonal-level processes. Some philosophers, such as those who take the approach of the Pittsburgh School, deny that we should expect to find a scientific solution to the interface problem. They propose that the distinctively reason-giving explanations that are characteristic of the personal level cannot be adequately analyzed or comprehended in terms of scientific laws or mechanisms.

Entities

When we make attributions of cognition, either personal or subpersonal, there is still a further question about what sorts of cognitive entities we are talking about. When we use subpersonal-level explanations to posit neural representations, for example, we might be talking about the firing response of an individual neuron, the chemical connections between populations of neurons, or the abstract mathematical properties of a neural manifold. Different theories and research programs can posit different kinds of cognitive entities, and it is important to realize that these are not necessarily observable or concrete objects. Cognitive subsystems, for example, are distinguished by their functional roles and need not be identified with localized neural structures; and some theories of cognition posit abstract mathematical entities such as multidimensional state-spaces and response thresholds.

Some cognitive scientists talk of the cognitive entities themselves, rather than the explanations or theories, as being personal or subpersonal. Notice, however, that personal and subpersonal explanations can pick out the same cognitive entities in two different ways—just as the same systems can be described by both statistical mechanics and classical mechanics.

Many cognitive scientists are realists about cognitive entities: They think that the success of our theories commits us to the existence of the entities they posit, whether concrete or abstract, observable or unobservable. Realists propose that the function of a cognitive subsystem is something that we discover rather than assign and that there is a fact of the matter about whether a certain neural structure is a representation and what it represents. The use of personal and subpersonal explanations, however, is compatible with a more instrumentalist or pragmatic stance on the cognitive entities they posit: Perhaps the theories are just useful tools or frameworks, or perhaps their theoretical commitments are relative to individual research programs or explanatory goals.

Questions, controversies, and new developments

The scope and limits of the personal level

Some philosophers who talk about personal-level attributions of cognition have a robust notion of personhood in mind, in which a person is a rational and reflective agent who can be held responsible for their actions. On such a view, personal-level explanations seem not to be applicable to infants, nonhuman animals, or artificial systems. Most cognitive scientists, however, tend to understand a person as simply an individual or cognitive system, and are prepared to attribute personal-level capacities like language competence to an artificial intelligence or personality traits to a rat. Many of the disagreements between theorists about personal-level attributions of cognition turn on what concept of personhood it requires.

Even if we agree on a notion of personhood, there is still a question of how we decide whether a cognitive capacity is attributable to the person or to some subpersonal component or system (Westfall, 2022). On one view, this is a matter to be settled by our linguistic practice: which terms we do in fact tend to ascribe to persons. Alternatively, and more commonly, we can think of our linguistic practice as tracking some feature of the world that makes our attributions correct: Perhaps cognition is correctly attributable to the person if they are consciously aware of the cognition in question, for example. There are good reasons to challenge the conscious awareness criterion: We might attribute the belief that oranges are not green to a person even if they have never reflected upon or articulated this view, for example. And we might think that most English speakers have the cognitive capacity to order adjectives correctly (e.g., “cute old purple house” rather than “purple old cute house”) despite them being unaware of this.

Instead of consciousness, we might appeal to some form of cognitive integration as the criterion for correctly attributing cognition to the person: Cognitive capacities are generally attributed at the personal level when they are appropriately integrated into the person’s behavior. We can attribute the belief that oranges are not green to a person if they do not select oranges when asked to select only green fruit, for example, and we can attribute the cognitive capacity to order adjectives correctly to a person when they consistently do so in speech. There are further questions, however, on how much integration is required and what sort of behavior is relevant: Is voluntary action required, for example, or does reflexive behavior count?

The scope and limits of the subpersonal level

Prior to the cognitive revolution, there were few subpersonal-level *psychological* explanations. Mental states were generally understood to be things that people could self-attribute by introspecting. When Freud and other psychologists began talking about nonintrospectable beliefs and desires, these were still attributed to the person (although unconsciously) rather than to some computational subsystem or neural mechanism. The sorts of subpersonal explanations that kickstarted the cognitive revolution, on the other hand, make attributions of cognition to functional subsystems, computational symbols, and neuronal structures.

Even those philosophers who question whether it is appropriate or accurate to attribute subpersonal *cognition* can agree that there are neural explanations that are subpersonal in the sense of not being personal-level:

neurophysiological explanations that do not employ cognitive terminology but talk only of neural firing rates and neurotransmitters, for example. Some people further extend the idea of subpersonal explanations beyond the brain, describing explanations of respiration in terms of chemical reactions in the mitochondria of our cells as subpersonal, for example (Naeem, 2023).

The mind-body problem

The interface problem should be distinguished from the mind-body problem in philosophy. The interface problem concerns the explanatory relationship between the personal level and the subpersonal level: How does an appeal to subpersonal-level theories help us to account for personal-level phenomena? It is a question about the relationship between different kinds of scientific theories or explanations. The mind-body problem, on the other hand, is traditionally understood as concerning the metaphysical relationship between the mental and physical realms: Are minds necessarily physical things, or are they the sorts of thing that could (in principle, even if not in actuality) exist independently of any physical substrate?

Some philosophers have developed positions on the mind-body problem that are motivated by particular solutions to the interface problem, most notably homuncular functionalism (Lycan, 1988) and the computational theory of mind (Fodor, 1975). These positions claim that personal-level mental states are not merely explained by subpersonal phenomena but *identical* to them.

Broader connections

Unconscious perception

The personal/subpersonal distinction plays a role in recent debates about unconscious perception. The existence of unconscious perceptual processing in the visual system is uncontroversial, but perception attributed to the person is normally in the form of conscious perceptual experience. Some experimental data, however, appears to suggest that our behavior can be given a personal-level perceptual explanation even where conscious perceptual experience is lacking. Subjects presented with nude images to one eye while given rapidly-changing stimuli to the other eye report no conscious experience of nude images, but they seem to attend more closely to male or female nudes in ways that align with their gender preferences (Jiang et al., 2006). One way to interpret this data is to say that the attentional phenomena are best explained by the person having perceived the nude images despite reporting no such experience. An alternative interpretation is to attribute attentional capacities subpersonally and claim that a cognitive subsystem of the person is attending, which can in turn be explained by sensory processing of the stimuli by perceptual subsystems rather than personal-level perception. We can characterize this debate in terms of the criteria being used for personal-level attribution: If we think that only high-level agential phenomena are properly attributable to the person, then we might deny that the attentional and perceptual capacities being exercised are genuinely personal-level on the grounds that they are more reflex-like (Phillips, 2019; Phillips & Block, 2016). What seems to be a debate

about the nature of perception can be interpreted instead as a disagreement about how to characterize the domain of personal-level attributions.

Dual-process theory

Dual-process theory proposes that human cognitive processing comes in two different varieties: Some cognitive processes (*Type 1*) are fast, unconscious, and intuitive, whereas others (*Type 2*) are slow, conscious, and reflective (Kahneman, 2011). Some theorists propose to identify this distinction with the personal/subpersonal distinction, such that only Type 2 cognitive processes are attributed to the person, whereas Type 1 processes are attributed subpersonally (Frankish, 2009). Such a move, however, is in tension with the standard approach to dual-process theory, on which the distinction between Type 1 and Type 2 cognition is a distinction between two different subpersonal-level explanations. Type 1 processes are usually characterized as nonsymbolic parallel information-processing taking place in evolutionarily older areas of the brain, whereas Type 2 processes are usually characterized as symbolic sequential information-processing taking place in evolutionarily newer areas of the brain. Dual-process theory is thus best understood as the claim that some of the cognitive states attributed to a person can be given Type 1 subpersonal-level explanations, whereas other cognitive states attributed to the person can be given Type 2 subpersonal-level explanations.

Implicit bias

Sometimes people sincerely voice one opinion while behaving in a way which suggests they endorse a contradictory opinion. Consider, for example, data from the Implicit Association Test that appears to suggest that someone who sincerely articulates and defends anti-racist beliefs will still categorize stimuli in accordance with negative racial stereotypes (Greenwald et al., 1998). This categorization data is often explained by attributing an *implicit* cognitive state to the person: an implicit racial bias or an implicitly racist attitude. Implicit cognition, however, can be attributed at either the personal or subpersonal level, depending on how we understand the criteria for correct attribution at the personal level (Drayson, 2022). If a correct attribution of a personal-level cognition requires that such cognition can be introspectively accessed and endorsed by the person, then an implicit racial bias will not be correctly attributable at the personal level. But given a weaker construal of cognitive integration, on which personal-level cognition merely needs to be behavior-guiding, an implicit racial bias can be correctly attributable at the personal level. In the latter case, we might still appeal to subpersonal phenomena to account for the implicit bias. We might explain why someone lacks introspective awareness of their implicit bias by appealing to the nature of the computational processes involved or the format in which information is neurally encoded (Drayson, in press).

Further reading

- Bermudez, J. L. (2005). *Philosophy of psychology: A contemporary introduction*. Routledge. https://doi.org/10.4324/9780203642405
- Dennett, D. C. (1969). *Content and consciousness*. Routledge.

- Drayson, Z. (2014). The personal/subpersonal distinction. *Philosophy Compass*, *9*(5), 338–346. https://doi.org/10.1111/phc3.12124

References

- Baddeley, A. D., & Hitch, G. J. (1974). Working memory. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 8, pp. 47–89). Academic Press. https://doi.org/10.1016/S0079-7421(08)60452-1 ↩
- Bennett, M. R., & Hacker, P. M. S. (2003). *Philosophical foundations of neuroscience*. Wiley-Blackwell. ↩
- Bermudez, J. L. (2005). *Philosophy of psychology: A contemporary introduction*. Routledge. https://doi.org/10.4324/9780203642405 ↩
- Broadbent, D. E. (1958). *Perception and communication.* Pergamon Press. https://doi.org/10.1037/10037-000 ↩
- Colombo, M., & Knauff, M. (2020). Editors' review and introduction: Levels of explanation in cognitive science: From molecules to culture. *Topics in Cognitive Science*, *12*(4), 1224–1240. https://doi.org/10.1111/tops.12503 ↩
- Dennett, D. C. (1969). *Content and consciousness*. Routledge. ↩
- Dennett, D. C. (2007). Philosophy as naive anthropology: comment on Bennett and Hacker. In M. Bennett, D. C. Dennett, P. M. S. Hacker, & J. R. Searle (Eds.), *Neuroscience and philosophy: Brain, mind, and language* (pp. 73–96). Columbia University Press. ↩
- Drayson, Z. (2012). The uses and abuses of the personal/subpersonal distinction. *Philosophical Perspectives*, *26*(1), 1–18. https://doi.org/10.1111/phpe.12014 ↩
- Drayson, Z. (2014). The Personal/Subpersonal Distinction. *Philosophy Compass*, *9*(5), 338–346. https://doi.org/10.1111/phc3.12124 ↩
- Drayson, Z. (2020). Why I am not a literalist. *Mind & Language*, *35*(5), 661–670.

↩

- Drayson, Z. (2022). The fragmented mind: Personal and subpersonal approaches to implicit mental states. In J. R. Thompson (Ed.), *The Routledge handbook of philosophy and implicit cognition* (pp. 79-89). Routledge.

↩

- Drayson, Z. (in press). The psychology of implicit knowledge. In K. Sylvan, E. Sosa, J. Dancy, & M. Steup (Eds.), *The Blackwell companion to epistemology* (3rd ed.). Wiley Blackwell.

↩

- Figdor, C. (2018). *Pieces of mind: The proper domain of psychological predicates*. Oxford University Press. https://doi.org/10.1093/oso/9780198809524.001.0001

↩

- Fodor, J. A. (1975). *The language of thought*. Harvard University Press.

↩

- Frankish, K. (2009). Systems and levels: Dual-system theories and the personal—subpersonal distinction. In J. Evans & K. Frankish (Eds.), *In two minds: Dual processes and beyond* (pp. 89–108). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199230167.003.0004

↩

- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology*, *74*(6), 1464–1480. https://doi.org/10.1037/0022-3514.74.6.1464

↩

- Hornsby, J. (2000). Personal and sub-personal; A defence of Dennett's early distinction. *Philosophical Explorations*, *3*(1), 6–24. https://doi.org/10.1080/13869790008520978

↩

- Jiang, Y., Costello, P., Fang, F., Huang, M., & He, S. (2006). A gender- and sexual orientation-dependent spatial attentional effect of invisible images. *Proceedings of the National Academy of Sciences*, *103*(45), 17048–17052. https://doi.org/10.1073/pnas.0605678103

↩

- Kahneman, D. (2011). *Thinking, fast and slow.* Farrar, Straus and Giroux.

↩

- Libet, B., Gleason, C. A., Wright, E. W., & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). The unconscious initiation of a freely voluntary act. *Brain*, *106*(3), 623–642. https://doi.org/10.1093/brain/106.3.623

↩

- Lycan, W. G. (1988). *Judgement and justification*. Cambridge University Press.

↩

- McDowell, J. (1994). The content of perceptual experience. *The Philosophical Quarterly*, 44(175), 190–205. https://doi.org/10.2307/2219740

↩

- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, *63*(2), 81–97. https://doi.org/10.1037/h0043158

↩

- Naeem, H. (2023). Is a subpersonal virtue epistemology possible? *Philosophical Explorations*, *26*(3), 350–367. https://doi.org/10.1080/13869795.2023.2183240

↩

- Phillips, I. (2019). Object files and unconscious perception: A reply to Quilty-Dunn. *Analysis*, 80(2), 293–301. https://doi.org/10.1093/analys/anz046

↩

- Phillips, I., & Block, N. (2016). Debate on unconscious perception. In B. Nanay (Ed.), *Current controversies in philosophy of perception* (pp. 165–192). Routledge.

↩

- Rupert, R. D. (2018). The self in the age of cognitive science: Decoupling the self from the personal level. *Philosophic Exchange*, *47*(1), Article 2.

↩

- Sellars, W. (1956). Empiricism and the philosophy of mind. In H. Feigl & M. Scriven (Eds.), *Foundations of science & the concepts of psychology and psychoanalysis* (Vol. 1, pp. 253–329). University of Minnesota Press.

↩

- Westfall, M. (2022). Constructing persons: On the personal–subpersonal distinction. *Philosophical Psychology*, *37*(4), 831–860. https://doi.org/10.1080/09515089.2022.2096431

↩

- Wittgenstein, L. (1953). *Philosophical investigations* (G. E. M. Anscombe & R. Rhees, Eds., G. E.M. Anscombe, Trans.). Blackwell.

↩