We Definitely Know What Attention is

Mengting Zhang [1 *]
Monash Centre for Consciousness and Contemplative Studies, Department of Philosophy, Monash University, 29 Ancora Imparo Way, Melbourne, VIC 3800, Australia
* Correspondence to [mengting.zhang@monash.edu]

Abstract

There is a tendency to criticize traditional attention research and argue for abandoning the term attention. This paper identifies the main critique as reification, circularity, false dichotomy, and one–many problem, and offers a systematic response. Circularity refers to treating attention as both explanans and explanandum. As argued, the underlying assumption of the reification critique concerns the specification of neural processes of attention. Drawing on Marr's distinction between levels of analysis, I argue that attention is a phenomenon at the computational level, generated by neural processes, which helps address these two critiques. To address false dichotomy and one–many problem, I adopt the active inference approach and offer a definition of attention as the rule-governed filtering of inputs. This definition justifies the conceptual distinction between intention and attention, attention and preattention, and reshapes the endogenous–exogenous distinction. It also explains how various types of attention fall under a single definition.

Keywords: attention; rule-governed; filter; active inference

Introduction

Attention is part of everyday life and plays an important role in many cognitive and emotional processes. However, there is still no agreement on what attention is. Classical views describe attention as selection (Broadbent, 2013), as a filter (Broadbent, 2013), or as a spotlight (Posner et al., 1980). More recent accounts define attention as selection for task performance (Wu, 2011, 2024), as prioritizing information (Watzl, 2023), as organization of sensory inputs in accordance with the subject's interest (Jennings, 2020), or as a manner of performance (Mole, 2024). Despite these differences, these accounts have faced strong criticism over the past few decades. As a result, many researchers now suggest dropping the term attention and focusing instead on specific behavioral or neural processes (Anderson, 2011; Di Lollo, 2018; Hommel et al., 2019). Extreme claims such as "no one knows what attention is" (Hommel et al., 2019), "there is no such thing as attention" (Anderson, 2011), and "attention is a sterile concept" (Di Lollo, 2018) reflect this growing trend.

Here, I argue that rejecting the term attention is not well justified. This paper aims to examine these critiques and respond to them. While authors use slightly different labels, I identify critiques as four main types: reification, circularity, the false dichotomy, and the one-many problem.

The paper is organized as follows. Section 2 discusses the reification critique and reveals its assumption that the neural processes of attention must be specified for the legitimacy of the term attention. I reject this assumption and argue that attention is a phenomenon at the computational level, different from neural processes at the other level. Section 3 addresses the circularity critique. I argue that attention is the explanandum of neural processes, while within the computational level it can serve as both explanans and explanandum. Section 4 introduces a new definition of attention based on the active inference framework, namely, attention as the rule-governed filtering of inputs. This definition helps resolve the false dichotomy critique, which questions distinctions such as those between pre-attention and attention, intention and attention, and endogenous and exogenous attention. It addresses the one–many problem by showing how different categories of attention can be unified within a single definition.

Can we specify attention even if it is not a neural process?

Critics have noted a tendency to reify attention in existing attention research, that is, to treat attention as a thing. They argue that such reification is problematic because it encourages the view that attention is a cause responsible for further effects, while existing definitions fail to specify this cause (Anderson, 2011; Di Lollo, 2018). I argue below that what really constitutes the problem is not what the critics identify, namely treating attention as a thing. Rather, the problem lies in the difficulty of specifying attention as a neural process. In this light, compared with the label reification used by critics, a more appropriate label for the problem is specification. I then appeal to David Marr's distinction between levels to argue that specification does not in fact constitute a problem.

Actually, what critics mean by a thing is opposed to a process. According to the discussion offered by critics, I construct the critique as follows. If attention is conceived as a thing rather than a process, it is thereby a causal agent. Yet it is hard to identify the causal agent correspondent to attention. Therefore, conceiving attention as a thing is problematic (Anderson, 2011; Di Lollo, 2018).

This critique seems to rely on the assumption that if attention were instead conceived as a process rather than a thing, the problem would disappear. I am skeptical of this assumption. Take the influential view that attention is selection as an example. Since selection is an action, this view defines attention as a process. Defining attention as a selection process still implies that attention plays a causal role in subsequent processing, such that selection causes goal pursuit. In this light, it still requires identifying the process corresponding to attention. This is not easy, since attention is highly diverse, including perceptual and internal attention, endogenous and exogenous attention, and so on. This

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

difficulty makes definition that treats attention as a process rather than a thing still problematic.

In this light, what truly constitutes the problem is not that attention is treated as a thing and therefore as a causal agent. Rather, the problem is treating attention as a cause and requiring to specify attention. More precisely, I contend that what matters is the demand for specification of attention. Even if attention is conceived as an effect, the specification problem does not disappear. Hence, it is better to label the problem as the specification problem instead of the reification problem.

In the following discussion, I offer a response to the specification problem. The analysis shows that critics understand specifying attention as specifying the experimentally manipulable neural processes of attention. It is argued that current attention definitions fail to achieve specification and should therefore be dropped. My contention is that failing to specify empirically manipulable neural processes does not constitute a problem. As I argue below, attention is not essentially equated with neural processes. Thus, this failure does not undermine the legitimacy of attention, but instead highlights the essential difference between attention and neural processes.

David Marr distinguished three levels of analysis: the computational level, the representational and algorithmic level, and the implementational level (Marr, 2010). There are multiple interpretations of Marr's distinction. I adopt the interpretation advanced by Bechtel and Shagrir (2015)[1]. As emphasized by Bechtel and Shagrir, the computational level focuses on physical constraints from the environment and the embodied structure of the agent, and the representational and algorithmic level concerns patterns of connections among neurons, such as addition, multiplication, factorization, integration, and so on.

I apply this level distinction to situate attention and its relevant processes. Consider a task in which similar stimuli are presented at different positions on the screen and participants are asked to report the stimulus at the center of the screen. In this case, the environmental condition is that multiple similar stimuli are presented simultaneously, and the embodied constraint is that the scope of the retina is limited. These conditions constrain the agent's performance. The agent must deploy her limited retinal scope to filter stimuli at the center of the screen for further processing, while filtering out stimuli at the periphery. These psychological processes of attention locate at the computational level.

At the representational and algorithmic level, the representational and algorithmic rules to realize this task are upregulating the precision of inputs from the center and downregulating the precision of inputs from the periphery. At the implementational level, the relevant neural processes involve sustained increased activity in early visual cortex representing the center area and decreased activity in early visual cortex representing other areas (Silver et al., 2007).

Hence, in the case of attention, attention is located at the computational level, whereas neural processes are located at the implementational level. As emphasized by Bechtel and Shagrir, different levels stand in a mechanistic relationship (2015). This means that processes at the other two levels are mechanisms that produce the phenomena at the computational level. The phenomenon emerges from underlying mechanisms but is not fully reducible to them. Thus, attention emerges from neural processes but is not fully reducible to them. The inability to specify the neural processes of attention indicates that they are, by nature, different things. Since attention is not equated with any neural process, the failure to specify its neural processes does not constitute a problem. This rejects the critique that uses this inability to argue for dropping the term attention.

Can we treat attention as both explanans and explanandum?

Besides the specification problem, there is critique claiming that the term attention is used as both explanandum (which is to be explained) and explanans (which to explain) and that this dual use generates the conceptual ambiguity, since the same term is applied to different explanatory roles. Critics label this the circularity problem and use it to argue for abandoning the term attention (Di Lollo, 2018; Hommel et al., 2019).

I respond to the circularity problem as follows. Whether attention can function as both explanans and explanandum depends on the context. Attention stands in a reciprocal explanatory relationship with phenomena at the computational level, and treating it as both explanans and explanandum at this level does not constitute a problem. By contrast, attention is the explanandum of processes at other levels, and treating it as both explanans and explanandum at those levels is not legitimate. To reduce conceptual ambiguity, it is necessary to clarify the appropriate way to discuss explanatory relationships involving attention in different contexts.

Attention and other psychological phenomena stand in a reciprocal explanatory relationship. For example, in major depressive disorder, biased attention toward negatively valenced memories can contribute to depressive affect, while depressive affect can in turn strengthen this attentional bias (Beevers et al., 2015). Attention is both the explanandum and explanans emotional response. Likewise, a person's motivation to achieve a goal can explain why their attention is selectively directed toward goal relevant stimuli. In turn, sustained attention to these stimuli can enhance perceived task value or interest, thereby strengthening motivation. In this way, motivation and attention mutually explain and reinforce each other. By contrast, when it comes to the explanatory relation between attention and processes at other levels, attention should be treated merely as the explanandum. As argued in the previous section, attention

[1] Noteworthy is that Bechtel and Shagrir also developed Marr's idea on that levels are not concerned with different perspectives of the same happening, but concerned with the different happenings (2015).

stands in a mechanistic relationship to processes at other levels. It is the interactions of neural processes that give rise to the occurrence of attention but not the reverse. Hence, interactions of the neural processes explain attention but not the reverse. This indicates that whether it is legitimate to treat attention as both explanandum and explanans depends on the explanatory context.

Distinguishing contexts in this way has implications for appropriate forms of expression. Claims such as paying attention reinforces motivation, emotional responses drive biased attention, or attention arises from biased activity in early visual cortex are appropriate. By contrast, claims such as paying attention activates the prefrontal cortex or attention allocation modulates precision weighting in inputs are misleading, because they reverse the direction of explanation.

Can we maintain the unity of attention in the face of diversity?

In addition to the criticisms discussed above, two further critiques challenge the use of the term attention. One holds that various clear-cut distinctions commonly drawn in attention research are not well supported (Anderson, 2011; Di Lollo, 2018; Hommel et al., 2019). The other argues that attention refers to a wide and heterogeneous range of phenomena that are difficult to unify under a single concept (Di Lollo, 2018; Hommel et al., 2019). On this basis, critics suggest abandoning the term attention altogether and giving up the traditional strategy of drawing clear-cut conceptual distinctions (Anderson, 2011; Di Lollo, 2018; Hommel et al., 2019). I draw on the active inference framework to develop an account of attention that demonstrates its capacity to address both critiques.

Attention as rule-governed filtering of inputs

I define attention as rule-governed filtering of inputs. Filtering refers to the process of separating relevant information from irrelevant information, allowing the relevant to pass through while blocking the irrelevant. Attention characterizes the biased processing of information and thus aligns with the idea of filtering. Defining attention as filtering is not novel (Broadbent, 2013). My definition differs from earlier accounts in emphasizing that the filtering is rule-governed, an idea I adopt from discussions of attention in the active inference literature.

The active inference framework treats agents as possessing a model of the environment and interacting with it based on this model (Parr et al., 2022; Smith et al., 2022). To survive as long as possible, agents should keep their models as closely aligned with the environment as possible, minimizing the mismatch between model and environment. This mismatch is also referred to as the uncertainty of environment. Minimizing the mismatch is named as free energy minimization (Parr et al., 2022; Smith et al., 2022), a principle that governs various forms of decision making, including the filtering of inputs. Attention, therefore, is regarded as filtering governed by the rule of free energy minimization.

Distinction between attention and sub-personal filtering

Defining attention as free-energy-minimization-governed filtering helps distinguish it from sub-personal filtering. Philosophers have highlighted the importance of distinguishing attention from sub-personal processes that resemble attention (Watzl, 2023; Wu, 2024). For example, retina implements a filtering process by allowing only certain wavelengths to pass on for further processing, yet we would not say that retina is paying attention to particular wavelengths of light. This is because attention is properly attributed to an agent, whereas the retina is merely a component of an agent.

Defining attention as rule-governed filtering can meet this challenge. As discussed above, the rule that governs attention is the free energy minimization principle. The prerequisite of pursuing free energy minimization is entailing a model of the environment, as free energy is the mismatch between the model and the actual world. Since only agents, rather than their individual components, can be said to entail such a model, filtering governed by free energy minimization must necessarily be implemented at the level of the agent rather than at the level of its components.

As mentioned above, one critique is that the difficulty of establishing clear-cut distinctions in the attention field motivates abandoning the traditional attention research (Anderson, 2011; Hommel et al., 2019). Specifically, these distinctions include those between attention and intention (Anderson, 2011; Hommel et al., 2019), between endogenous attention (i.e., goal-directed attention) and exogenous attention (i.e., salience-driven attention), and between attention and pre-attention (Di Lollo, 2018; Anderson, 2011; Hommel et al., 2019). I propose that the account of attention within active inference also capable to address this critique.

Replacing the endogenous–exogenous distinction with the pragmatic-value-guided versus epistemic-value-guided distinction

Endogenous attention refers to attention guided by internal top-down goals, which is deliberate and slow whereas exogenous attention refers to attention guided by externally salient stimuli, which is automatic and quick (Carrasco, 2011). The binary distinction between endogenous and exogenous attention has been challenged by two lines of evidence. First, endogenous attention can sometimes operate automatically, and exogenous attention can depend on its relevance to the target, undermining a clear-cut distinction between the two (Folk et al., 1992; Hommel et al., 2019). Second, research has identified attention driven by learned rewards, which is neither guided by top-down goals nor triggered by salience. This type of attention cannot be classified as either exogenous or endogenous (Awh et al., 2012), further questioning the binary distinction.

I summarize these findings as including three cases: goal-guided automatic attention, target-relevant exogenous attention, and reward-based attention. As I concern, we can draw a distinction between pragmatic-value-guided attention and epistemic-value-guided attention within active inference framework and such a distinction can account for these cases. Goal-guided automatic attention and target-relevant exogenous attention are attention guided by the pragmatic value of inputs, whereas reward-based attention is guided by the epistemic value. This explanatory power demonstrates that the active inference framework's distinction is superior to the endogenous–exogenous attention distinction.

Pragmatic value refers to the value of fulfilling a specific goal (Parr et al., 2022; Smith et al., 2022). Attention guided by pragmatic value has a slightly broader scope than endogenous attention. Endogenous attention is typically defined as attention guided by a top-down goal that is consciously held. Within the active inference framework, however, a goal can guide attention even when it is not consciously represented (Parr et al., 2022; Smith et al., 2022). On this view, attention guided by pragmatic value also includes attention guided by unconsciously held goals. An example of unconsciously held goal is the goal of maintaining homeostasis, such as keeping body temperature within an appropriate range. Agents have the automatic response of dressing more on a cold day. It is guided by this homeostatic goal without the goal being consciously represented. This kind of goal can also guide attention, making it goal-directed yet automatic.

Regarding the target-relevant exogenous attention, it is drawn to an external stimulus not because it is salient, but because it is relevant to a current goal. For example, when searching for a friend in a crowd and expecting them to wear a red coat, a red object may automatically capture attention. The shift is exogenous in that it is stimulus-driven, yet target-relevant because the stimulus matches a goal-defining feature. I suggest that this kind of attention is also guided by pragmatic value, since inputs are filtered according to their relevance to the goal. In this way, cases that fall outside the account of endogenous attention can be accommodated within the pragmatic-value-guided attention.

Reward-based attention arises from learned associations between stimuli and rewards. For example, if you repeatedly receive positive feedback from seeing a thumbs-up icon on social media, the symbol may automatically capture your attention whenever it appears, even when you are not actively seeking positive feedback. Attention is drawn in this case not because the stimulus is salient or relevant to a current goal, but because it has a learned association with reward. I regard this as attention shaped by epistemic value.

Epistemic value contrasts with pragmatic value. Information is pragmatically valuable if it helps achieve a goal that is currently being pursued (e.g., information about the word name helps achieve the goal of reporting the word name in the Stroop task). Information is epistemically valuable if it is not relevant to a specific current goal but contributes to minimizing free energy in the long term (Parr et al., 2022; Smith et al., 2022). For instance, monetary rewards have been found to enhance participants' motivation to perform cognitive control tasks, such as resisting tempting food to maintain a diet, when they are exhausted (Muraven & Slessareva, 2003). In this case, the monetary reward is not pragmatically valuable, since it is not directly related to the immediate goal (e.g., maintaining the diet). However, it is epistemically valuable, as it contributes to the agent's long-term well-being. Learned rewards, though unrelated to a current goal, convey patterns or contingencies in the environment and hence contribute to long-term survival and well-being. Attention driven by such rewards is therefore guided by epistemic value.

In this light, even if the traditional endogenous–exogenous distinction is untenable, a clear-cut distinction of attention based on pragmatic and epistemic value remains viable. Within the active inference framework, various factors such as learning history, past experience, goals, and salience contribute to the pragmatic and epistemic value of information. These factors are deeply intertwined and dynamically interact with one another, making them difficult to delineate precisely. Nevertheless, epistemic and pragmatic value can still be distinguished in a clear and principled way.

Distinction between attention and intention

As noted, traditional attention research has emphasized the distinction between attention and intention (Hommel et al., 2019). This distinction is significant because salience-driven attention (e.g., attention captured by a kangaroo suddenly running out) is clearly not intentional, implying that intention is not associated with attention in any case. However, critics have questioned this distinction, arguing that attention and intention are not clearly dissociable at the level of experimental manipulation. They draw on the evidence of the overlapping neural substrates underlying attentional and intentional activities. The posterior parietal cortex, typically regarded as a core neural substrate of selective attention, has also been found to be involved in tasks that implement intention (Mountcastle et al., 1975; Snyder et al., 1997).

The active inference framework addresses this critique by highlighting two different ways in which input filtering can be implemented. One way is through the attribution of higher precision to a particular information channel. The other is through entertaining a stronger expectation of a future state (Parr & Friston, 2017).

Active inference is grounded in the predictive processing view of the brain. On this view, the brain continuously generates predictions about the world and receives sensory inputs to update or correct those predictions. Predictions and inputs are hierarchically organized. Predictions originate at higher levels of the hierarchy and are transmitted downward, whereas sensory inputs arise at lower levels and are transmitted upward.

From this perspective, there are two ways to prioritize certain kinds of inputs, such as sounds. One way is to assign higher precision to the channel processing auditory information relative to channels processing other forms of input. Because this channel is treated as more precise, the agent regards information coming from it as more reliable and informative, thereby prioritizing inputs from that channel.

The second way is through the formation of higher level predictions. The brain can simultaneously entertain multiple predictions about the world, such as “I should be in the office at 8” or “I should take my time and enjoy breakfast,” with each prediction assigned a different probability value. When one prediction has a substantially higher probability than the others, the agent has a stronger expectation that the corresponding state will occur. This stronger expectation constrains and prescribes a sequence of lower level states. In other words, beliefs at higher levels generate predictions about states at lower levels. As a result, information consistent with these predictions, such as the bus timetable relevant to getting to the office, will be preferentially filtered and processed in order to fulfill the higher level expectation.

Intention is easily conflated with attention in the second case because input filtering is implemented through the maintenance of an expectation, which effectively functions as an intention. Nevertheless, the two phenomena remain distinguishable. Within an active inference model, intention consists in entertaining an expectation of a future state, formally represented by one state possessing a substantially higher probability than alternatives. Attention, by contrast, consists in the actual prioritization of certain inputs over others, which can be formalized as the selective uptake of information from lower levels of the hierarchy.

Importantly, entertaining an expectation of a future state at a higher level does not by itself guarantee selective filtering of inputs. For example, if the intention is overly general, such as “I should keep thinking,” it may fail to prescribe sufficiently specific sequences of lower level states. In that case, inputs from lower levels will not be selectively prioritized. The agent therefore possesses an intention without exhibiting attention.

Active inference framework addresses this critique by highlighting two different ways to implement input filtering. One way is to attribute higher precision in the information channel. The other way is to entail a higher expectation of a future state (Parr & Friston, 2017)[2]. Active inference framework is grounded in the predictive brain assumption. The brain is conceived as generating the predictions of the world and receive inputs to adjust its predictions. Predictions

Distinction between attention and preattention

The distinction between attention and preattention was proposed to explain the difference between feature recognition and object recognition in visual attention. Research has shown that the recognition of basic visual features, such as color or orientation, differs from the recognition of objects. The former is rapid and relatively independent of distracting information, whereas the latter is not (Treisman & Gelade, 1980). To account for this difference, the distinction between preattention and attention was introduced. Preattention is defined as the rapid and parallel processing of basic features, whereas attention is defined as the sequential binding of these features. The theory that adopts this distinction is known as feature integration theory (Treisman & Gelade, 1980).

Critics have questioned the distinction between attention and preattention (Anderson, 2011; Di Lollo, 2018). For instance, Anderson draws on empirical findings to challenge this distinction. Some findings suggest that performance in basic feature recognition is also affected by location information, that is, information processed at the attention stage (Treisman & Gelade, 1980). As I interpret, this finding poses a challenge because it indicates that attentional processing may already operate at the preattention stage, thereby undermining a clear-cut distinction between the two.

I argue that the active inference framework can both preserve the conceptual distinction between attention and preattention and explain their close interaction. As noted above, attention is defined as sequential processing and preattention is defined as parallel processing (Treisman & Gelade, 1980). Parallel processing can be understood as the simultaneous processing of different sources of information without clear prioritization, whereas sequential processing is the processing with explicit prioritization, such that some information is processed earlier and other information later. Active inference framework can formalize this distinction.

In active inference models, prioritizing inputs is formalized as assigning a higher precision (γ) value to a particular likelihood mapping (A). A higher precision value increases the influence of inputs from that source on the agent’s belief updating, which corresponds to the function of attention allocation. By contrast, non-prioritized processing of inputs, namely, preattention, is formalized as assigning similar precision (γ) value to different likelihood mappings (A). Similar precision value ensures that inputs from different

[2] This paper primarily draws on the work of Parr and Friston (2017). In that work, Parr and Friston employ the active inference framework to clarify what is meant by the term attention, thereby undertaking an important project of conceptual clarification. They identify two different ways of understanding attention, each of which can be formalized differently within the active inference framework.

By contrast, other active inference studies, such as Misic et al. (2025), Holmes et al. (2021), Mirza et al. (2019), and Brown et al. (2011), are primarily concerned with constructing specific active inference models of particular forms of attention, such as selective attention or visual attention, rather than with clarifying the concept of attention itself. Moreover, these studies formalize attention in terms of increased precision within a particular information channel, thereby overlooking the possibility that attention may also be implemented through entertaining a specific expectation.

sources contribute to decision-making to a similar degree. In more intuitive terms, attention implements the filtering of inputs, whereas preattention does not. Defining attention as the filtering of inputs therefore is able to distinguish it from preattention.

Notably, even though preattention does not implement filtering, the result of preattention is not necessarily an evenly distributed representation of information—that is, not all information has equal value, and some information can still stand out. This is because, even when inputs are processed in a non-prioritized manner, the inputs themselves may have different inherent values. Differences in input value lead to differences in value after pre-attentive processing. For example, imagine looking at a page full of black text with a single word in bright red. At the preattention stage, the color information of the page, including the redness, is processed without prioritization. Nevertheless, the red word still stands out. This occurs because the redness input is rare compared to the blackness inputs, and this rarity assigns it a higher value. Even though the color information is processed non-prioritively, the higher inherent value of redness ensures it stands out after preattention. This interpretation suggests that differences in the value of information can arise from multiple factors, and the manner of filtering inputs is only one such factor.

Regarding the apparent conflation of attention and preattention, this occurs because the two processes are constantly interacting. In the active inference framework, the agent is understood as a prediction machine that continuously predicts the world and updates its beliefs based on prediction errors (Smith et al., 2022). I apply this to the relation between attention and preattention. Attention transmits predictions downward to preattention, shaping how sensory information is processed at the pre-attentive stage. Preattention processes sensory information in parallel and transmits prediction errors upward to attention. These two functions constantly interact and impact each other, creating a deep connection and the apparent conflation.

Unify diverse forms of attention under a single account

As mentioned earlier, one major critique of existing attention research is that the term attention encompasses too broad a range of phenomena to admit of a unified account (Hommel et al., 2019), a challenge I refer to as the one-many problem. The active inference framework offers a way to address this issue.

On this view, all forms of attention share a common feature: the rule-governed filtering of inputs. This constitutes the unifying characteristic of attention. However, because such filtering can be shaped by multiple contributors, including external stimuli, internal motivations, and working memory, which together determine the epistemic and pragmatic value of information, attention can manifest in a variety of forms. Different types of attention may therefore be categorized according to the factors that contribute to the filtering process. Examples include external attention and internal attention. Moreover, filtering itself can be implemented in two distinct ways: either through modulating the precision of an input channel or through establishing a higher level expectation. Attention can therefore also be classified according to its mode of implementation. Examples include attention realized through eye saccadic movements and attention realized not through eye saccadic movements.

Concluding remarks

This paper has analyzed the main criticisms for abandoning the term attention and provided detailed responses to each. Drawing on the active inference framework, I proposed a definition of attention as the rule-governed filtering of inputs. This definition explains why various types of attention can be unified under a single term and distinguishes attention from related phenomena, including preattention, intention, and sub-personal filters. The traditional distinction between exogenous and endogenous attention is revised to a distinction between attention guided by pragmatic value and attention guided by epistemic value, offering a clearer and more accurate account. In the future, this definition has the potential to address other issues in the study of attention, such as the debate between early and late selection, the relationship between attention and consciousness, the relationship between attention and the sense of self, and so on.

References

Anderson, B. (2011). There is no Such Thing as Attention. *Frontiers in Psychology*, 2. https://doi.org/10.3389/fpsyg.2011.00246

Bechtel, W., & Shagrir, O. (2015). The Non-Redundant Contributions of Marr's Three Levels of Analysis for Explaining Information-Processing Mechanisms. *Topics in Cognitive Science*, 7(2), 312–322. https://doi.org/10.1111/tops.12141

Beevers, C. G., Clasen, P. C., Enock, P. M., & Schnyer, D. M. (2015). Attention bias modification for major depressive disorder: Effects on attention bias, resting state connectivity, and symptom change. *Journal of Abnormal Psychology*, 124(3), 463–475. https://doi.org/10.1037/abn0000049

Broadbent, D. E. (2013). *Perception and communication* (Original work published 1958). Pergamon Press.

Brown, H., Friston, K. J., & Bestmann, S. (2011). Active Inference, Attention, and Motor Preparation. *Frontiers in Psychology*, 2. https://doi.org/10.3389/fpsyg.2011.00218

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, Vision Research 50th Anniversary Issue: Part 2, 51(13), 1484–1525. https://doi.org/10.1016/j.visres.2011.04.012

Di Lollo, V. (2018). Attention is a sterile concept; iterative reentry is a fertile substitute. *Consciousness and Cognition*, Visual Experience and Guidance of Action: A Tribute to Bruce Bridgeman, 64, 45–49. https://doi.org/10.1016/j.concog.2018.02.005

Holmes, E., Parr, T., Griffiths, T. D., & Friston, K. J. (2021). Active inference, selective attention, and the cocktail party problem. *Neuroscience and Biobehavioral Reviews*, 131, 1288–1304. https://doi.org/10.1016/j.neubiorev.2021.09.038

Hommel, B., Chapman, C. S., Cisek, P., Neyedli, H. F., Song, J.-H., & Welsh, T. N. (2019). No one knows what attention is. *Attention, Perception & Psychophysics*, 81(7), 2288–2303. https://doi.org/10.3758/s13414-019-01846-w

Jennings, C. D. (2020). *The Attending Mind* (1st ed.). Cambridge University Press. https://doi.org/10.1017/9781108164238

Marr, D. (2010). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. The MIT Press. https://doi.org/10.7551/mitpress/9780262514620.001.0001

Mirza, M. B., Adams, R. A., Friston, K., & Parr, T. (2019). Introducing a Bayesian model of selective attention based on active inference. *Scientific Reports*, 9(1), 13915. https://doi.org/10.1038/s41598-019-50138-8

Mišić, T., Koledić, K., Bonsignorio, F., Petrović, I., & Marković, I. (2026). An Active Inference Model of Covert and Overt Visual Attention. In M. Albarracin, D. Benrimoh, C. L. Buckley, P. Lanillos, R. J. Pitliya, H. Shimazaki, I. P. Stoianov, T. Verbelen, & M. Wisse (Eds.), *Active Inference* (pp. 167–181). Springer Nature Switzerland. https://doi.org/10.1007/978-3-032-16955-6_10

Mole, C. (2024). Attention And Attentiveness: A Defence of The Argument for Adverbialism. *Australasian Journal of Philosophy*, 102(2), 465–480. https://doi.org/10.1080/00048402.2023.2289109

Mountcastle, V. B., Lynch, J. C., Georgopoulos, A., Sakata, H., & Acuna, C. (1975). Posterior parietal association cortex of the monkey: Command functions for operations within extrapersonal space. *Journal of Neurophysiology*, 38(4), 871–908. https://doi.org/10.1152/jn.1975.38.4.871

Muraven, M., & Slessareva, E. (2003). Mechanism of self-control failure: Motivation and limited resources. *Personality and Social Psychology Bulletin*, 29(7), 894–906. https://doi.org/10.1177/0146167203029007008

Parr, T., & Friston, K. J. (2017). Working memory, attention, and salience in active inference. *Scientific Reports*, 7(1), 14678. https://doi.org/10.1038/s41598-017-15249-0

Parr, T., Pezzulo, G., & Friston, K. J. (2022). *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. The MIT Press. https://doi.org/10.7551/mitpress/12441.001.0001

Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology*, 109(2), 160–174.

Silver, M. A., Ress, D., & Heeger, D. J. (2007). Neural correlates of sustained spatial attention in human early visual cortex. *Journal of Neurophysiology*, 97(1), 229–237. https://doi.org/10.1152/jn.00677.2006

Smith, R., Friston, K. J., & Whyte, C. J. (2022). A step-by-step tutorial on active inference and its application to empirical data. *Journal of Mathematical Psychology*, 107, 102632. https://doi.org/10.1016/j.jmp.2021.102632

Snyder, L. H., Batista, A. P., & Andersen, R. A. (1997). Coding of intention in the posterior parietal cortex. *Nature*, 386(6621), 167–170. https://doi.org/10.1038/386167a0

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(1), 97–136. https://doi.org/10.1016/0010-0285(80)90005-5

Watzl, S. (2023). What attention is. The priority structure account. *WIREs Cognitive Science*, 14(1), e1632. https://doi.org/10.1002/wcs.1632

Wu, W. (2011). Attention as Selection for Action. In C. Mole, D. Smithies, & W. Wu (Eds.), *Attention: Philosophical and Psychological Essays* (pp. 97–116). Oxford University Press. https://philarchive.org/rec/WUAAS

Wu, W. (2024). We know what attention is! *Trends in Cognitive Sciences*, 28(4), 304–318. https://doi.org/10.1016/j.tics.2023.11.007