

Jolie Abadir

Professor Kramer

COGS/PHIL UN3655

May 17, 2026

Priority Map Dissociation in Autism and PTSD: An Architectural Constraint on Selection-for-Action

1. Introduction

The category “attention disorders” groups together autism, PTSD, ADHD, and several other conditions on the basis that each involves some apparent dysfunction of attention. This grouping is a clinical convenience. The philosophical question it raises is whether the category names a single kind of attentional breakdown or collapses mechanistically distinct ways the priority map can be tuned into a single diagnostic bucket. If the latter, the category does more than describe; it imposes a normative baseline against which any departure counts as a deficit.

This paper argues that autism and PTSD dissociate two components of priority map construction, that this dissociation gives Wu’s selection-for-action framework a constraint it must accept in order to remain explanatory rather than descriptive, and that once that constraint is in place, the deficit framing of these conditions becomes visible as a normative imposition the architecture does not license.

The argument proceeds as follows. Section 2 sets out selection-for-action as the framework the paper defends and addresses Watzl’s priority-structure account as the main competitor. Section 3 introduces the priority map and grants the complication Awh, Belopolsky, and Theeuwes raise against the clean bottom-up/top-down dichotomy, retreating to a weaker architectural claim sufficient for the paper’s purposes. Sections 4 and 5 develop autism and

PTSD as the two empirical cases. Section 6 states the architectural payoff and the Wu constraint. Section 7 makes the difference-not-deficit move. Section 8 concludes by noting open questions about other conditions and what the cases reveal about attention itself.

2. Selection-for-action as the right framework

The framework this paper defends is Wu's. On selection-for-action, attention solves a problem cognitive systems face continuously: at any moment many possible inputs are available and many possible actions are open, and something has to determine which inputs will drive which actions. Attention is that determination. There is no separate faculty that surveys options and then picks one; the picking is what attention does. Visual attention, motor attention, working memory, and the attentional aspects of cognitive tasks are the same structure operating across different domains.

The main competitor is Watzl's priority-structure account. Watzl agrees that priority maps exist but disagrees with Wu about what attention fundamentally is. For Watzl, attention is the process by which some mental contents become central in consciousness while others recede. The view has clear phenomenological support. Attended contents are foregrounded in awareness in a way that seems part of what attending is rather than a side effect, and Watzl's account makes that centrality constitutive. The view also handles cases that are difficult for selection-for-action: receptive attention, meditative attention, aesthetic attention. In each, consciousness is clearly structured without any obvious action being selected.

Selection-for-action can absorb these cases. Receptive attention involves action in a non-motor sense. Attending to a painting consists in sustaining a particular relation to it, and that sustained relation is what selection is selecting. Watzl's phenomenal structure, on this reading, is not a rival account of attention but the experiential aspect of selection-for-action when the

coupling is in place. The framework does not require that the action be motor or serve any external goal; it requires that the system be doing one thing with its inputs rather than another.

The autism and PTSD cases will test selection-for-action's explanatory commitments. If attention is what Wu's framework says it is, conditions that alter attention should alter it in ways the framework can specify. Whether the framework can specify them is the question Section 6 returns to.

3. The priority map and Awh's complication

The priority map is the standard computational model of attentional selection. On the account developed by Fecteau and Munoz, each location in the visual scene is assigned a priority value that integrates two kinds of input: bottom-up salience, which depends on the physical properties of the stimulus, and top-down relevance, which depends on the goals of the agent. Attention is allocated to whichever location wins the priority competition. Bisley and Mirpour have identified the lateral intraparietal area as the most likely neural site where these inputs are integrated, with single-neuron activity in that region tracking salience, relevance, and other factors simultaneously. Carrasco's review of two decades of psychophysics provides the broader empirical context: attention does not merely select among percepts but modulates perception at early stages, sharpening contrast, accelerating processing, and altering the appearance of attended stimuli.

The textbook version of the priority map carves its inputs cleanly into bottom-up and top-down. Awh, Belopolsky, and Theeuwes argue that this dichotomy fails. Selection history and reward-based effects produce attentional biases that fit neither category. When a target has repeatedly appeared in a certain color or location, that feature acquires priority even after it ceases to be the goal and even when it is not physically salient. The bottom-up/top-down framing

forces these effects to be classified as one or the other; on Awh's view, that classification misdescribes the phenomenon. Priority is shaped by multiple sources, and the dichotomy is, as the title of their paper has it, a failed theoretical division.

The argument here grants Awh. The clean dichotomy does fail. But a weaker and more defensible claim survives: the priority map's construction involves at least two functionally distinct stages of computation, an early salience-driven stage and a later goal-directed stage, and those stages can be independently parameterized. "At least two" rather than "exactly two." Whatever the full inventory of priority sources turns out to be, those sources are integrated at distinct stages of map construction, and the autism/PTSD asymmetry is what that distinction looks like when stress-tested empirically. The architectural commitment the rest of the paper uses is this weaker, post-Awh version.

4. Autism: early-stage parameterization

Mottron, Dawson, Soulières, Hubert, and Burack propose the Enhanced Perceptual Functioning model of autistic perception. Their central claim is that autistic perception is not impaired at the low-level perceptual stage; it is differently weighted. Low-level features such as edges, textures, pitch, and fine spatial detail are processed with greater autonomy from higher-level cognition than in non-autistic perception, and autistic individuals show measurable advantages on tasks that depend on local feature detection. The model is supported by a converging body of evidence across visual and auditory modalities, though it remains in dispute with the Weak Central Coherence account, which interprets the same data as a deficit in global integration. For present purposes, both accounts agree that autistic attention weighs early perceptual information differently from neurotypical attention, which is what the dissociation requires.

Keehn, Müller, and Townsend develop the architectural picture further. Drawing on Posner's three-network framework, they argue that the primary attentional impairment in autism localizes to the orienting network, and specifically to its disengagement function. Autistic individuals can construct priority maps but have difficulty reconfiguring them. Once content has been prioritized, releasing it so that attention can be reallocated is impaired. The pattern appears as early as infancy and is one of the earliest neurobehavioral markers of later ASD diagnosis.

Shafritz, Dichter, Baranek, and Belger provide the neural-level evidence. Using fMRI, they show that during response-shifting tasks, individuals with autism exhibit hypoactivation in dorsolateral prefrontal cortex, anterior cingulate cortex, intraparietal sulcus, and basal ganglia. The result that matters most for the present argument is the correlation they report between activation in ACC and left IPS and the severity of restricted, repetitive behaviors: subjects with higher RRB scores showed smaller activations in these regions. The architectural finding maps onto the clinical picture in a direct way.

Taken together, these results give the autism side of the dissociation a specific shape. Autism is a case in which the early, salience-driven stage of priority map construction is parameterized differently from the neurotypical case, while the later goal-directed stage remains relatively intact. The local-feature weighting, the disengagement profile, and the RRB-correlated hypoactivation are consistent with one another at the level of architectural commitment.

5. PTSD: late-stage parameterization

Hayes, VanElzaker, and Shin synthesize neuroimaging evidence on the interaction between emotion and cognition in PTSD. Their review identifies a consistent neural pattern: hyperactivity of the amygdala, which tags stimuli for emotional significance, combined with hypoactivity of the ventromedial prefrontal cortex, which normally regulates the amygdala's

tagging signal. The result is that threat-relevant content is weighted heavily in the goal-and-affect-modulated stage of priority map construction. The regulatory machinery that would normally down-weight threat-irrelevant content in non-threat contexts is impaired, and threat-related representations capture attention disproportionately.

Three features of this neural pattern matter for the architectural argument. First, the dysregulation is located in a circuit (amygdala-vmPFC) whose function is regulatory rather than perceptual; the vmPFC does not modify the input image, it modulates which representations get treated as motivationally salient. Second, the alteration is direction-specific: it amplifies content already tagged by threat history, leaving content lacking that tag unaffected. Third, the impairment is documented in studies using stimuli with normal physical salience (faces, neutral words, scenes), which means the threat-reweighting effect cannot be attributed to altered low-level processing of the stimuli themselves. Each of these features locates the parameterization difference at the same point: the stage at which the priority map is regulated by goal-and-affect-tagged content, downstream of low-level salience computation.

This is not a salience-level change. People with PTSD do not see threat more vividly than others. The early perceptual machinery responsible for processing physical stimulus features is intact, and there is no evidence of altered low-level perception in PTSD comparable to the differences documented in autism. What is altered is the later, goal-directed stage at which the priority map's regulation by current goals and affective tagging takes place. The priority map's input-side is operating normally; the regulatory side is not.

This gives the PTSD side of the dissociation a shape that is the inverse of autism. The early stage of priority map construction is intact while the later goal-directed stage is reweighted by threat history. The two conditions hit different stages of the same dynamic process.

6. The architectural payoff and the Wu constraint

The dissociation between autism and PTSD does philosophical work. Autism's parameterization difference localizes to the early, salience-driven stage of priority map construction; the local-feature weighting and impaired disengagement that Mottron and Keehn describe are differences in low-level perceptual prioritization. PTSD's parameterization difference localizes to the later, goal-directed stage; the amygdala-vmPFC pattern Hayes and colleagues describe is regulatory rather than perceptual, and what gets reweighted is not what the system sees but what the system treats as relevant given its history. Autism and PTSD are the same kind of attentional alteration—priority map parameterization—happening at different stages of the same dynamic process.

This asymmetry is evidence for an architectural claim weaker than the dichotomy Awh, Belopolsky, and Theeuwes argued against, but sufficient for the philosophical work the paper aims to do: the priority map's construction involves at least two functionally distinct stages, and those stages can be independently parameterized. The claim is not that the map has only two stages, or that the bottom-up/top-down distinction is theoretically clean. The claim is that whatever the full inventory of priority sources turns out to be, those sources are integrated at distinct stages of construction, and the autism/PTSD asymmetry is what that distinction looks like when stress-tested empirically.

The strongest objection comes from Todd and Manaligod, who propose treating priority as a high-dimensional state space rather than a sequence of stages. On their view, autism and PTSD might both be shifts along the selection-history dimension, and the apparent asymmetry would dissolve into one dimension behaving two different ways. The reply is that the state-space framing and the staged framing are not in competition. Even within a multidimensional state

space, the shifts in autism and PTSD happen at different points in the computational pipeline: autism's shift acts before goal-directed regulation enters, PTSD's shift acts within it. The asymmetry tracks where in the priority map's construction the parameterization difference applies, not which dimension of an abstract space it occupies.

This is where selection-for-action needs to accept the architectural distinction. Wu's framework holds that attention is the structure resolving the many-many problem at the input-output interface, and its explanatory ambition is to unify the attentional phenomena under that structure. For selection-for-action to predict that autism and PTSD will produce systematically different attentional profiles, rather than to accommodate whichever profile turns out to obtain, the priority map underwriting selection must have components that can be parameterized independently. Without the architectural commitment, the framework can describe any pattern of dissociation post-hoc but cannot say in advance which patterns are possible.

One reply available to Wu is that selection-for-action is an account of what attention is rather than a predictive model of how attentional profiles vary across populations, and that the parallel demand on a metaphysical account of water as H_2O would not require it to predict regional variation in water hardness. The reply is fair as far as it goes, but the explanatory ambition Wu defends in "We Know What Attention Is!" goes further. The 2024 paper argues that selection-for-action unifies the attentional phenomena under a single mechanism and that the field should treat the unification as explanatory rather than merely classificatory. A framework with that ambition cannot remain silent about the architecture that underwrites the unification. If selection-for-action unifies attentional phenomena, it must be willing to say something about why those phenomena come apart in the ways they do, and the autism/PTSD dissociation is exactly the kind of case where the framework is asked to deliver.

The commitment is not that selection-for-action must adopt the early/late distinction as its central structure; it is that the framework must accept at least one architectural distinction of this kind in the priority map it depends on. Otherwise selection-for-action becomes compatible with empirical findings only by remaining silent on them, and the autism/PTSD dissociation makes the silence costly. Selection-for-action remains the right account of what attention is. The architectural constraint is what selection-for-action needs in order to be the account it claims to be.

7. Difference, not deficit

Once the architectural commitment is in place, the deficit framing of autism and PTSD becomes visible as a normative move rather than a neutral medical fact. The medical framing treats neurotypical priority map parameterization as the standard against which any departure counts as dysfunction. The architecture does not license that treatment. If the priority map has stages that can be independently parameterized, then a non-neurotypical parameterization is not a broken version of the standard map; it is a different parameterization optimized for a different action-environment.

Munton makes the relevant structural point in a different context. On her account, prejudice consists in salience being misattributed: features become attention-worthy on grounds that do not track the property they purport to track. The same structural critique applies to the deficit framing of attention conditions. The property the framing purports to track is dysfunction; what it actually tracks is parameterization difference. Whiteley extends the point by arguing that some salience perspectives are harmful not because they are descriptively wrong but because of what they do to subjects whose maps differ from the foregrounded standard. The deficit framing

of autism and PTSD is one such perspective. It treats neurotypical parameterization as the property worth tracking and pathologizes the rest.

What the cases reveal in a non-disorder register is that the priority map is parameterized for the action-environments a cognitive system has historically inhabited. Autistic parameterization reflects an action-environment in which fine perceptual detail and local feature processing matter more than rapid global integration. PTSD parameterization reflects an action-environment in which threat-relevant stimuli have had unusually high prior probability. Neither parameterization is broken; each is what the system has become responsive to. The architecture leaves open the question of which parameterization counts as standard, and the medical framing's answer to that question is a substantive normative claim, not a discovery about how attention works.

8. Conclusion

The autism/PTSD dissociation provides selection-for-action with the architectural constraint it needs to be explanatory rather than descriptive. Once the framework accepts that priority map construction involves at least two functionally distinct stages that can be independently parameterized, it can predict the shape of the dissociation rather than merely accommodate it. The improvement is not a refutation of Wu's view but the specification of what his view must commit to in order to do the unifying work it claims.

This leaves open the question of whether other conditions, such as ADHD, schizophrenia, and depression, fall along the two axes the autism/PTSD dissociation identifies or expose further dimensions of priority map construction. The framework, with the architectural commitment added, can absorb new cases without treating any one parameterization as the norm, which is what allows the philosophical machinery to do work the medical framing cannot do.

What the dissociation shows, in the end, is that attention is not one thing happening uniformly across cognitive systems but a structured process whose stages can be tuned to different action-environments. Selection-for-action remains the right account of what attention is; the autism and PTSD cases show what selection-for-action must commit to in order to be that account. The deficit framing of attention conditions is what happens when a normative parameterization is mistaken for the architecture itself. The architecture is silent about which tuning is correct, and that silence is where the philosophical work begins.

Works Cited

- Awh, Edward, Artem V. Belopolsky, and Jan Theeuwes. "Top-Down versus Bottom-Up Attentional Control: A Failed Theoretical Dichotomy." *Trends in Cognitive Sciences* 16, no. 8 (2012): 437–443.
- Bisley, James W., and Koorosh Mirpour. "The Neural Instantiation of a Priority Map." *Current Opinion in Psychology* 29 (2019): 108–112.
- Carrasco, Marisa. "Visual Attention: The Past 25 Years." *Vision Research* 51, no. 13 (2011): 1484–1525.
- Fecteau, Jillian H., and Douglas P. Munoz. "Salience, Relevance, and Firing: A Priority Map for Target Selection." *Trends in Cognitive Sciences* 10, no. 8 (2006): 382–390.
- Hayes, Jasmeet P., Michael B. VanElzakker, and Lisa M. Shin. "Emotion and Cognition Interactions in PTSD: A Review of Neurocognitive and Neuroimaging Studies." *Frontiers in Integrative Neuroscience* 6 (2012): 89.

- Keehn, Brandon, Ralph-Axel Müller, and Jeanne Townsend. "Atypical Attentional Networks and the Emergence of Autism." *Neuroscience and Biobehavioral Reviews* 37, no. 2 (2013): 164–183.
- Mottron, Laurent, Michelle Dawson, Isabelle Soulières, Bénédicte Hubert, and Jake Burack. "Enhanced Perceptual Functioning in Autism: An Update, and Eight Principles of Autistic Perception." *Journal of Autism and Developmental Disorders* 36, no. 1 (2006): 27–43.
- Munton, Jessie. "Prejudice as the Misattribution of Salience." *Analytic Philosophy* 64, no. 1 (2023): 1–19.
- Shafritz, Keith M., Gabriel S. Dichter, Grace T. Baranek, and Aysenil Belger. "The Neural Circuitry Mediating Shifts in Behavioral Response and Cognitive Set in Autism." *Biological Psychiatry* 63, no. 10 (2008): 974–980.
- Todd, Rebecca M., and Mana R. Manaligod. "Implicit Guidance of Attention: The Priority State Space Framework." *Cortex* 102 (2018): 121–138.
- Watzl, Sebastian. *Structuring Mind: The Nature of Attention and How It Shapes Consciousness*. Oxford: Oxford University Press, 2017.
- Watzl, Sebastian. "What Attention Is. The Priority Structure Account." *WIREs Cognitive Science* 14, no. 1 (2023): e1632.
- Whiteley, Ella. "Harmful Salience Perspectives." In *Salience: A Philosophical Inquiry*, edited by Sophie Archer. London: Routledge, 2022.
- Wu, Wayne. *Movements of the Mind: A Theory of Attention, Intention and Action*. Oxford: Oxford University Press, 2023.

Wu, Wayne. “We Know What Attention Is!” *Trends in Cognitive Sciences* 28, no. 4 (2024): 304–318.

Wu, Wayne. *Attention*. 2nd ed. Abingdon: Routledge, 2025.