

A Virtue of Attention

Wayne Wu
Italian Academy of Advanced Studies
Columbia University

It is a remarkable tendency among scientists later in their careers to try to tackle the big questions, typically spelled out in books produced for public consumption. I mention this not to mock but to concede feeling this tendency myself. I've been toying with writing a book that engages with the public on attention. Yet serious theorists write professional articles and academic treatises! When one's colleague produces a popular book, there's quiet condescension. "Sure, they're entitled to try to cash in on their work, but it's all very soft and squishy isn't it?" The philosopher has, I think, a different problem that provides a counterpoint. Academic philosophy, at least the so-called Analytic tradition, is often far removed from engaging the desire for philosophical insight felt by the layperson. I think philosophers *ought* to do more to engage publicly. Yet there's the soft-and-squishy concern. I'm going to try to sketch how certain academic issues about attention might productively engage public concern. This will provide a backdrop to some more detailed things I will present on Wednesday.

Attention is something we all need to think more deeply about. Recent public discourse regarding attention is caught up with the rhetoric of the *attention economy*. Attention is a commodity that can be exchanged. It can be sold or stolen. This is what tech companies are charged with doing, buying our attention on the cheap or even illicitly making off with it.

Consider two ways of understanding the commodity idea, one that treats it as a resource like energy and another that links it to labor (here, I would be interested in an economist's take on the attention economy). Attention as resource? Attention in us is a biological phenomenon. Let's begin our reflection with a genuine biological resource: blood. One can donate blood, sell it, and perhaps even have it stolen. Blood transactions can be quantified. In one's mental and physical activity, one requires more blood flow to bring needed oxygen, another resource, so the amount of blood available constrains performance. Attention seems to be like that. If you pay less attention, then you are apt to make mistakes or miss things. The solution is paying more attention (what is the equivalent phrase in Italian...is it similarly transactional?).

That there can be more or less attention is apparently common sense. It does not follow, however, that payment talk reflects psychological reality. An interesting question I shall not pursue is whether any psychological capacities are literally resources in a quantifiable sense. One might imagine a future world where we can sell space in our neural memory buffers to companies, given a pressing social need to store more and more data. The harvesting of mental buffers! I'll leave these dystopian possibilities to science fiction novelists.

I'm not going to argue against the resource idea in this paper though I think it's wrong. Rather, I will assert that attention is not a resource doled out like energy or blood. If there is an economics of attention, it relates to a connection to labor, or better, to *laboring*. Here, I'm going to appeal to William James whose description of attention rings true to me:

Everyone knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others.

To extract the core idea, as a philosopher, I'll inelegantly but with logical perspicuity put the Jamesian idea this way (three variables, *S* for subjects, *M* for mode and *T* for targets):

S M-attends to T when S mentally selects in mode M target T in order to deal with.

My apologies to you and James! In a slogan, attention is *selection for action*. Notice that the values of *S* are sentient creatures, so you and I attend (our brains do not). The values of *T* are verified by empirical investigation for creatures of a given kind, but for us include space, features, objects, time, internal states, thoughts, images and so on. Science has helped us understand the targets of attention in great detail. Finally, the mode of attention, *M*, can be sensory, emotional, cognitive and so on, and indeed, mode descriptions can be quite complicated.


James' insight, which he teases out of common sense, is that attention is a critical part of one's doing things. Consider bona fide cases of your own attention in action: reaching for an object or for a memory, working out a problem by reasoning things through, imagining walking through a favorite trail or neighborhood, having a difficult conversation or a conversation in a language

you have not yet mastered and so on. Attention is embedded in our doings through our taking mental possession of targets in different modes (visually, mnemonically etc.), often ignoring other targets. Psychologists call this *top-down* attention and contrast it with *bottom-up* attention, say one's attention being captured by a loud sound or an embarrassing memory. Playing with prepositions, the Jamesian gloss is that top-down attention is exemplified in the taking possession *of* something *by* the mind while the bottom-up case is the taking possession *of* the mind *by* something. In one case, the mind is active, in the other, passive. We'll come back to the bottom-up case, more broadly the case of *automatic* attention.

Attention matters not because we have limited quantities of a mental good but because it is central to our agency, to our doing things with body or mind (philosophers speak of bodily and mental actions). The issue of an attention economy if we wish to speak in those terms is not that we have things to sell but that we have things to do. Our doings are what matters. The challenge of attention is not theft but the modification of agency.

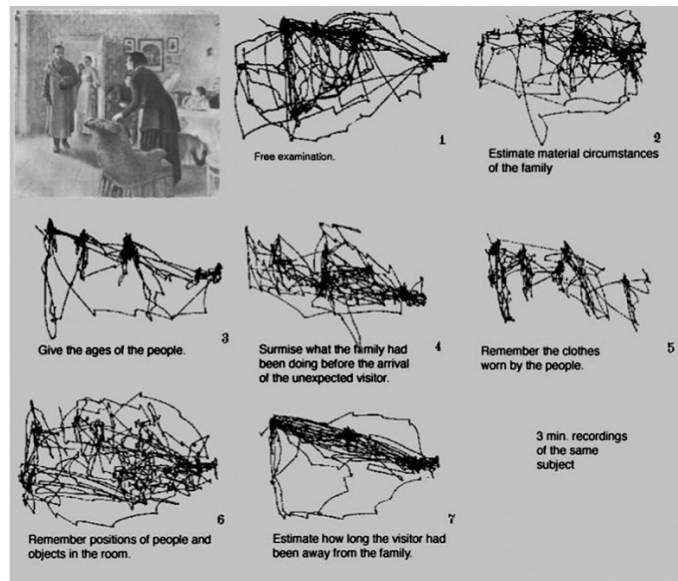
There is nothing wrong with modification per se. After all, we have a will, and in exerting it, we modify our own attention directing it as we wish (usually). Some of us still maintain ideals of social and political discourse where we try to convince each other through reasoning that certain things should be done. Many theorists emphasize rational processes where people make informed decisions through correct reasoning, individually or jointly.

Setting the will and concerted activity aside, what is of interest to me is the *unwitting* modulation of agency through the shaping of attention in ways of which we are largely unaware. This leads to what I prefer to call the automaticity of attention. Orienting towards a loud sound is but one form of automatic attention. There are many others. Let me give a picture of a common case, the automaticity of eye movement, drawing on pioneering work from the Soviet scientist, Alfred Yarbus, who developed ways of tracking these. Consider saccadic movements, ballistic ocular movement that happens between one to four times a second:




The “task” the subject is executing is *free viewing*. The subject need do nothing more than have a look at the painting. cursory observation reveals that there is a *bias* towards faces. The eye lands more often on faces than other targets though the subject need not be aware of this fact. What is striking about eye movements is their automaticity. Where and on what the eye lands is typically not explicitly planned.

Yarbus also showed that one way to strongly bias saccades is to give agents concrete tasks. Instructions, when followed, engage the will. The agent intends to do the instructed task. With changes in intention come changes in action, strikingly revealed in many of the observed patterns Yarbus recorded. For many cases, one can seemingly read off the intended action from the trace.



The experiment beautifully reveals an interplay between control and automaticity. Roughly, the agent's control leaves as trace that can be deciphered from the pattern, say gathering information about clothes by looking at the figures. That movements serve the action intended. Automaticity is revealed in the fine-grained movements that constitute the overall intelligible pattern. An agent would not be able to reconstruct the specific movements, say accurately report the last two or three saccades made (you might ask yourself what the last two landing points of your eye were as you read this text). This is the joy of automaticity, the freeing up of thought.

Consider another example which reveals the interplay of automaticity and control in learning. As in Yarbus's cases, there is a task. Radiologists look at images for medical assessment. What is striking in this series is given the same task to locate any anomalies in the x-ray, the eye movements reflect level of medical training. Learning shapes how we attend as part of acquiring skill and expertise. I'll talk more about this case in my talk but until then, you might speculate on what is going on (what is given is a heatmap where the "hot zones" are amount of time fixating that region; the flip side of that is that the eye moves often to that position or lingers there):



Let me return to the idea of biases on attention and summarize the main point. Given that attention matters to what we do, its influence on our actions reflects a myriad of biases. Intentions are one bias, but the vast majority of biases are subterranean in evading our awareness of them. Here we come to the crux of my present interests: what are these biases and how are they acquired, changed, and sometimes manipulated? This is where we come, I think, to matters of public interest.

One way bias matters is in the attention economy (I'd prefer to reframe the issue in terms of autonomy and agentic skill, but since everyone is talking about economies, I'll use the idea in this paper). The challenge of the attention economy is not in commodification and exchange of goods, wittingly or not, but in the shaping of our agency, specifically through tapping into the automaticity of bias. How would our thinking about the attention economy change by reframing things in this way? That's something I would be keen to discuss with you.

The other reason it matters, turning to academic philosophy, is that it opens up an unexplored area of research, namely our sensitivity to norms that make demands on what we *ought* to do. One way philosophers explicate the normative domain is to speak of reasons for actions, considerations that speak for or against a course of action. Investigating how players respond to reasons is part of explaining the rationality of agents. To build on an earlier example, consider a young resident who notices an anomaly in the x-ray, a smudge to untrained eyes. It is probably nothing serious, at least that is their gut reaction. Still, it is a radiographical anomaly. It doesn't have the obvious signs of a tumor but as a new resident, one is inclined to be conservative. Yet given a medical system of limited resources, e.g. the cost of imaging like an MRI or CT scan, what is the proper medical action? Is there reason to medically intervene? Deliberation of this sort has been a central topic in understanding rational agents in philosophy and other disciplines.

I want to step back, however, and note that in reflecting on the anomaly, the resident has already engaged attention. But look back at the eye movements in the x-rays and notice that there is an important prior step that philosophers at least have not engaged with, one's sensitivity to reasons. After all, it's easy to miss important things, to fail to notice. What would it be to have an appropriate sensitivity to things that matter, to reasons for action? What would it be to be appropriately attuned?

Here's an excerpt from a paper I'm currently trying to finish on *attunement*, the agent's attentional potential, that is propensity to attend to things in a given context and the agent's internal state at that time:

Excerpt from the introduction of "Attunement" (Draft)

To act well, one must "see correctly". One walks to work. Ahead just on the other side of the street is a person in distress. That person's condition imposes a demand: we should aid them. Perhaps the person is a tourist who is lost or perhaps the need is medical, the person has collapsed. Different needs make different demands, and the reasons exemplified in visible need provide material for practical reason and rational action.

Reasons and reasoning are the starting point for philosophical investigation of an agent's engagement with the normative, say with moral or epistemic reasons. Philosophers have focused on providing theories of practical or theoretical reasoning and the ensuing intentional action. Yet this theoretical orientation assumes the satisfaction of something the agent cannot take for granted: *noticing* relevant reasons. It is one thing to begin with relevant reasons,

another to notice them in the first place. Noticing is *attending*. To notice something is to pay attention to it, and to fail to notice is to fail to pay.¹ Nothing guarantees that attention will take possession of available reasons, and if it does not, there is nothing to reason with. Ordinary experience involves surprising examples of the perplexing failure of attention apparently dulled. “How did you miss *that*?” One might own up, “I wasn’t paying attention.” Empirical work shows that this is a statistically normal state of affairs. If so, philosophical reflection that tees up materials for the agent passes by a sizeable realm of the agent’s engagement with the space of reasons. Attention and the conditions for appropriate attention are philosophically significant topics over which many have passed in silence.

On a familiar view of consciousness, that a person in need is in my view means that I see them. Yet seeing said person is not enough to engage my response to them if I do not notice. Empirical work reveals that failure to notice is pervasive. Indeed, the associated rhetoric suggests that where we fail to notice the person, they impose no demands on us. The unnoticed person in need is effectively behind a wall. We are *inattentionally blind* to them. How can we be obligated to what we literally cannot see?

Walls are not necessarily barriers to obligation. A father who leaves his toddler alone by the pool while he goes into the house to (just so very) quickly check the game on television is not absolved of an obligation just because a wall separates him from his child who has just fallen into the pool. We cannot block normative demands by shutting doors. But what of the

¹ In many contexts where we wish to say that one *noticed X* or failed to do so, one can speak equivalently of *paying attention* to *X* or failing to pay attention. The problem in the opening case is that one failed to pay attention to a relevant reason. That said, I acknowledge that some might treat talk of noticing as referring to *one type* of attention, distinguished by being more cognitively loaded, say involving conceptualization and recognition. This seems to me a more “technical” notion. In this essay, I am treating the normative assessments phrased in terms of noticing to be equivalent to a claim about attention. If the invocation of noticing causes problems, readers can recast relevant claims in terms of attention paid or not.

mother who is coming home from work and walks along the outside wall of the garden which separates her from her child whom she cannot perceive. Is she at that point obligated to save her child of whom she is literally unaware? Empirical work suggests that where we fail to notice reasons for action, our situation is more akin to the mother than the father. Yet failing to notice is commonplace because inattention is the flipside of attention. Are we then pervasively absolved of normative demands just because of pervasive inattention?

Philosophical reflection papers this over by placing reasons on a platter for the imagined agents that for the targets of philosophical inquiry. The challenge of attention is answered by fiat, and a field of inquiry is thereby closed. In this essay, I argue for the need to consider the agent's propensity to attend, an attentional potential I shall call the agent's *attunements*.

Attunement is the basis for how the agent attends in a context, the subject's being inclined to notice aspects of the world, internal or external. Unless we are attuned to relevant factors in the world so that we are more likely to attend to them, reasoning is often stymied or veers off course. Attunement shapes our engagement with the world. It bridges merely seeing without noticing, to noticing and attending, necessary conditions for action. Attunement is a critical capacity in agents who can respond to reason, and in epistemic and ethical contexts, virtuous agents are properly attuned.

My stated project is to tell a useful and informative story about bias on attention across different domains that draws substantively on different levels of analysis including the philosophical, psychological, computational and neural. I would like to do this in a way that engages theoreticians as well as the general public. What would it be to fully understand

attention as something that is biased in different ways to yield virtuous and vicious behavior of theoretical and social concern? In my presentation, I'll discuss a possible narrative within the context I've presented here.