ATTENTION AND AI
(with James Williams)
September 1, 2023

14:00:46 From Peter Schmidt to Everyone:
    Hey Matthew!
14:01:46 From Peter Schmidt to Everyone:
    Oi Julio, tudo bem?
14:08:18 From Júlio de Ló to Everyone:
    Replying to "Oi Julio, tudo bem?"
    Oi Peter, tudo ótimo e com você? tks
14:17:10 From Ana to Everyone:
    YAAYY!
14:18:16 From Peter Schmidt to Everyone:
    My email, for any question in relation to the Strother School's various projects:
14:18:35 From Peter Schmidt to Everyone:
    peter@sustainedattention.net
14:21:18 From Peter Schmidt to Everyone:
    So cool! I have not seen this…

Notes on Jac Mullen's "Exercise of the AI Generated Environment"

14:41:19 From Sal Randolph to Everyone:
    Notes: The light in here is weird: way too bright, way too vivid. The light is coming from everywhere: so many windows, so many skylights. How can it possibly be so dark and shady-feeling in here? The light seems uncanny and excessive, the number of windows ridiculous, the interior still dark.

    Prompt: Create the dim interior of a vacation cottage with the bight slanting sunlight of a summer afternoon entering from many windows and skylights. Revision: more light. Revision: more light.
14:42:01 From D. Graham Burnett to Everyone:
    I chose to focus, for my "glitch," on the way that these systems can sometimes fail to capture the way depth of field works — and the false depth generated via blur can fail to map onto form correctly. This made for a very strange experience of the real space I am actually in.
14:42:09 From Elena Radice to Everyone:

prompt: You're a four years old child who is setting a home-ambient responding to different purposes: a trip, a drawing, an adult job, a castle, a forest, and adult "adulting" in a corner.

14:42:16 From Jeff Dolven to Everyone:

Behind my head when I zoom are two photographs of trees hung with halved apples. That looks like a glitch. Actually it's art. Really it's art? Virtually it's art? Anyhow—look through that lens, and all you see is half of anything. In fact, half of anything is just about everything you ever see. "Show me half of everything in the room."

14:42:23 From Kyle Winston-Lindeboom to Everyone:

Prompt: "Wood counter and wood ceiling frame domestic kitchen scene filled of white stripes on the surface of the materials."

14:42:27 From Matthew Clemente to Everyone:

A man sitting at a desk in a vacant office on a college campus in Boston.

What first drew my attention were the windows of the office. They look out onto a brick building in front of which are several trees. Boston College has a scenic campus with a lot of gothic architecture. I'm on the top floor of the building I'm in looking out at a good portion of the campus. I imagined, however, that without very specific promptings, the AI would generate a generic image of downtown Boston out my window, leaving it inaccurate.

14:43:17 From Alexandru Balgiu to Everyone:

A spark of forgetfulness in the library, an unknown animal feeding on dust

14:43:28 From Sal Randolph to Everyone:

Reacted to "A spark of forgetful..." with ❤️

14:43:33 From Peter Schmidt to Everyone:

Notes: my experience quickly became auditory - there's a weird irregular clanking coming from a construction site nearby. It made me aware of how indefinite and vague are the day to day sounds that make up our auditory field. The prompt was: the outskirts of a major city — construction — mild traffic — breezy day.

14:44:34 From Jac Mullen to Everyone:

Reacted to "Notes: The light in ..." with ❤️

14:58:12 From Ana Cristina Coura (Tininha) to Everyone:

Notes: I keep dried branches that I colect in my walks with my books, in the bookshelves - it makes sense to see that all wood goes together. The prompt: branches and books blending together. When I relaxed my attention, the curves of the branches created a lot of motion and I could see the moving jungle, just like I experience when I'm taking a walk in the woods, specially when I'm in a tropical forest. Lots of motions with my books and branches coming "alive' around me. Funny fact: as I was curiously looking around my environment and getting into the generated (by my own prompt) reality, my dog, who rests next to me, started barking, intrigued. I wonder if he got a sense of the strange sense I was experiencing my self. ;)

14:58:35 From Katie Murphy to Everyone:

Glitches: Either: what is that enormous television doing in this room full of books, and why is only half of it reflecting the rest of the room, while the rest is a totally void black? OR: in this mostly very orderly room, with everything geometrically aligned, why is there a pile of knitting on the couch, and haphazard notebooks heaped on that shelf, and tangled headphones on the table? I realised that my attention had slipped from AI glitches to thinking about what actual decisions had produced contradictions in the room — largely related to accommodating two different's people's needs and desires — and how far those decisions were about what we want to see, and what we want to be seen, and how far just the accretion of living. (I'm the messy one.)

Post-meeting timestamp: From Melissa Galvez: At first, I was very confused what a "tell" might be in this environment; after all, it wasn't created by AI and thus there shouldn't be any tells.

But then, as I looked around the room, I noticed something - two pieces of mail, books encased in brown cardboard boxes, stacked on top of each other. Mail for an old resident of this house, who no longer lives here. Exactly alike except for the number printed on the side of the box, just one number off - 05J0038497 vs 05J0038498. That struck me as odd, like there didn't need to be two books in this scene, only one. It's like someone copied/ pasted an image and thought that 2 books made more sense than one but didn't really know how things work. Like the file number on the image had simply been increased.

Then as I started to look around the room I noticed an uncanny trend - there were MULTIPLE examples of two nearly identical objects, with just a small difference. Two receipts stacked on top of each other, for the same transaction, but one has the tip. Two earbuds, but in different orientations. 2 pieces of mail. I was looking around the room, really paying attention to each detail, when I realized that I hadn't looked at myself - I looked down at my hands and saw two identical hands, except one hand has a ring.

This experience was truly extraordinary - it began to feel as if the room really WAS a simulation where the creator/ creating AI just kept making copies of things.

But the prompt was unclear to me. I could bring up words that described the feel of the scene, the things in it, but I had no idea how the duplication tell fit in.

It wasn't until Phase 4 - when I really closed my eyes and let my mind wander, that it came to me. Here's the prompt: "Draw a middle aged woman* sitting at a desk or table in a room filled with things of every day life - papers, books, items from past meals. The room teeters on the edge of depression/ chaos, and functionality."

To me, the key to the duplication was the idea of a room teetering on the edge of something. There is this dueling energy in here right now - a whiff of chaos, disorder, dirt - and yet, it is a functional space, where people work, eat, sort mail, get things done.

The duplication, it seems to me, arose from AI trying to capture a space that is on the edge of two energies. Each item represents one version of events, sitting comfortably against another version. One small difference between the two. A space unsure which direction to fall. A space sitting in between.

Notes on the prompt from James Williams: resisting AI-driven "Informationalism"

15:13:46 From Peter Schmidt to Everyone:
        Sal: I was thinking on my parakeet as we discussed the machinic metaphor. Wondering: in what ways is he, and is he not, like a machine? What is the quality of his attention? What kinds of motivations are we building into the AI? What desires do they see in us? What forms of attention is AI automating? We know it will be good at

detecting cancers in radiological scans, or at guiding self-driving cars. How much are we going to outsource our own attention?

15:14:43 From Julian Chehirian to Everyone:

Jeff - that exercise made me think of an exercise that might be along the lines of what James is interested in. Barthes idea of image having a "punctum". That would be an interesting thing to ask of an AI image. AI seems to give likeliest of what you are asking for... punctum is what eludes what you are asking for. The image as an averaged informationality... can you still find something there that does to us what Barthes talks about?

Ana - I've been having conversations about Youtubers with my children - 8 and 9 years old - about how their brain works. That it is a choice for your attention to be or not be a product. And I think the only way to preserve the attentional frame is the other thing - the emotional - that cannot be processed. As much as AI is trying to do this.

15:14:44 From Julian Chehirian to Everyone:

Alex - importance of senses and perceptions in one's experience of the world. I've been doing some work on improvisation lately. Thinking about ad-hoc relationship to environment and responding to existing conditions. And how use intuition and deep listening as a way to perhaps avoid any kind of systematic conditioned response. An interesting contrast in relation to this question of information.

Jeff - there is something raw about information. It doesn't know what to do with itself. It is not curated. It isn't necessarily for us exactly. We have to decide what to do with it. Information and concept have a necessarily dialectic relationship. And AI would like to spare you an encounter with information. Because who needs it when you can get what you want right away? And not what you don't want.

15:14:48 From Julian Chehirian to Everyone:

Ana - a ceramic cup, compared to a factory molded cup. The prior preserves something more than information - an imprint of its making.

Julian - to Ana's comment, I think about how there is a fantasy / illusion of seamlessness, when it comes to AI generated information. And the seams, the traces of the contingency and "madeness" of knowledge or objects, is a repository for alternatives, for what isn't but could be.

15:14:50 From Sal Randolph to Everyone:

Peter: Thoughts at the intersection of attention and AI

Elena: Coming from a big discussion of teachers and professors about all the courses proposed on these topics in a very superficial way. There was a colleague yelling into the mic about how we are all being manipulated by AI. I'm more interested in what is inside the AI. What is the difference between summary and details. What we see when we pay attention is details. We see the archive. What the AI does is sums up and produces summaries and generalizations. The way cameras/phones interpret the "photographs" you take in order to offer something it seems you want.

15:14:59 From Sal Randolph to Everyone:

Peter: AI sees everything equally. The idea that AI images feed on themselves, that they can only draw on the existing archive of images on the internet. Soon AI will be consuming mainly AI generated images. Like the water cycle, a closed system. A weird self-contained ecology. I think of attention in terms of connection, that it can help us perceive how we are connected to each other. These two ecologies of connection.

Sal: AI as an embodied/machinic form of attention. What does it mean to outsource our attention - to AI scanning medical records, to self-driving cars, the way that AI replicates our race, gender, and class bias.

15:15:05 From Sal Randolph to Everyone:

Peter: How do you correct ethical biases as the AI makes use of all our past material?

Elena: Training AI on particular voices.

Peter: More discrimination about what AI takes in. What AI takes in is itself a form of attention.

15:15:17 From Jac Mullen to Everyone:

Carlos:

There are two questions here.

What is our agency in a fully informational society? This is the question carried over from cybernetics, active in business school etc. [It is, in other words, receiving a lot of 'weird' attention?] How does technology dissolve agency?

The other question is what kinds of metaphors are emerging from Chat-GPT? Are we creating this in our own image? Is it good? Bad? Are we exceeding the limits?

15:15:21 From Jac Mullen to Everyone:

Kyle:

Focusing on image generation: ascribing words like 'art,' 'beauty,' etc. to images made by these models will blur the boundaries between human-tool-made and AI-generated image. It is really important to using different terms for these different activities/forms of visual objects. New words will allow for more agency.

15:15:25 From Melissa Galvez (she/her/hers) to Everyone:

In our group there was some discussion of informationalism vs X. What other lens are important? Informationalism vs ritual, action, group forming, interpretation

15:15:26 From Jac Mullen to Everyone:

Katie:

AI learning vs. human learning. Is learning the correct word for this context? For both?

What is the relationship between attention and intelligence? Are these two totally different concepts? What do these two things bear on one another? Does it make sense to say AI is 'attending'?

Carlos:

Google transformers paper about attention.

15:15:31 From Sal Randolph to Everyone:

Elena, what was the name of the artist you mentioned?

15:15:32 From Jac Mullen to Everyone:

Jac:

Stuff about attention in transformers. Amplification of salience stimuli and inhibition of non-salient stimuli.

Maybe we return back to the old idea of human as microcosm—we do contain all things, in a basic way. The kinetic, the chemical, the computational. We tend to mistake ourselves for the process when we encounter it in our media or tech for the first time. But we contain all of these processes.

Also—striking when even some of the same algorithms might be active in transformers and in subsets of neurons in the brain (in terms of e.g. learning).

Also—predictive coding, perceptual fill-in similar to fill procedures in e.g. Adobe photoshop, etc.

Kyle: Old experiment for identifying blind spot.
15:16:41 From Brad Fox to Everyone:
David, Justin, Alyssa, Graham, Brad:
15:17:02 From Brad Fox to Everyone:
David is in the room!

Justin:
I'm too close to these issues to know how to contribute. I've been writing about it.

Alyssa:
I've been following your work since your blog posts

Graham:
a machinic master-metaphor for the human is almost always an ethically problematic project. AI's nascent human-like qualities advances the machinic metaphor explicitly as an information-processing terms.
It's as if it's reminding us the last forms of resistance of this master metaphor will collapse, leaving us no recourse.
How can we protect a concept of attention in the face of this?
Wayne Wu's intentionalism?
Are there forms of attention that are unassimilable?
15:17:14 From Brad Fox to Everyone:
Justin
There's a Mott and Bailey (?) motion for people who deploy machinic metaphors. If you press they will insist that they don't want to deal with ?s of phenomenal consciousness or qualia.
Everything we think is unique to that is just a sophisticated autofill.

Alyssa
Everything can be rendered information except… What is the dot dot dot…?

Brad
It reminds us to seek the non-transactional

Graham
Attention to the non-determinant
Like Alyssa's cinematic cut speaks to the infinite
That which will not be assimilated

David:

Humans make infinite ends out of finite means

Our experience is attentional not informational — how do we shore up against the status quo bias that everything is information? How to specify against these specific routes?

Alyssa
How to recall people to the attentional nature of experience?
How do we protect what makes our lives our own?

15:17:28 From Jeff Dolven to Everyone:

Julien was our rapporteur, but just one follow-up to the thought on information and AI—I wonder if AI might end up making us nostalgic for information? Insofar as there is too much information, it is boring, etc.—it is raw, whereas the products of AI tend to be perfectly cooked, to the general taste. Information and concept are in a dialectic for modernity, and with text and image etc. prediction, will they disappear together? (This is about our use of information, as opposed to the way in which AI is trained on information.)

15:17:41 From Sal Randolph to Everyone:
Artist that Elena mentioned re: voice and AI was: https://www.hollyherndon.com

15:17:57 From Elena Radice to Everyone:
Reacted to "Artist that Elena me..." with ❤️

15:27:13 From D. Graham Burnett to Everyone:
The artist/architect Marshall Brown has a lovely essay on "SEAMFULNESS" — as a form of resistance to the seamlessness of digital systems…

15:28:00 From Julian Chehirian to Everyone:
Touch re-enactors

15:31:09 From Sal Randolph to Everyone:
Beauty as flaws

15:33:55 From Julian Chehirian to Everyone:
Period eye futurism

15:36:27 From Sal Randolph to Everyone:
A future consisting entirely of the regurgitated past.

15:37:38 From Sal Randolph to Everyone:
"this diorama of materiality"

15:38:57 From Jac Mullen to Everyone:
J. J. Gibson