THE AI GUIDE TO YOUR ADHD BRAIN

by Dr. Gillian Hayes

Who am I?

Hi, I'm Dr Gillian Hayes, coach for techy, scientific and creative adults with ADHD. And I spent 20 years doing AI at the world-renowned Department of AI at the University of Edinburgh, in Scotland.

I got really interested in how to make robots and software agents learn - from trial and error, from each other, and by imitating humans. Along the way, I figured out that I must have ADHD.

I had a lot in common with those software agents who couldn't seem to see the goal they were aiming at in their gridworld environment until they were right up against it.

Last-minute merchants.

And so, via a few twists and turns, I decided to become an ADHD coach, and use all those insights I'd gained from understanding how to make robots behave intelligently to help actual humans - you and me - act more as we'd wish to. And for us to realise that it's not "just us".

AI can help us to figure out our own mechanisms and ways of working, and then to stack the odds in our favour so that we're much more likely to get to where we want to go to.

What's in this guide?

You'll find a couple of ways of thinking about your ADHD brain, followed by five different things you can do with your ADHD brain, all inspired by artificial intelligence.

Table of contents

THE AI GUIDE TO YOUR ADHD BRAIN

Intro

AI. Artificial Intelligence, that is. What on earth does it have to do with your, and my, ADHD brain?

AI means simulating natural (human!) intelligence in robots or computer programs and – the other way round – figuring out how natural intelligence must work by seeing what it takes to get a robot to act intelligently.

So, AI helps us to figure out how our brains must be working. It gives us ways of understanding ourselves better and of finding better ways of working.

Ways of thinking about your ADHD brain

AI can help us to make models of brains and how they might be functioning when they're doing their processing - of what comes into them through our eyes, ears, touch, smell, balance, etc. And then how they go on to reason and come up with some responses -- moving to get out of the way of a divebombing seagull, answering a workmate's question, placing a chopped up onion in a pan without spilling it on the floor.

These models don't necessarily look exactly like our brains do, with neurons and synapses and glial cells, but they describe behavior and functionality. And they help us to understand what might be going on in our brains, how our behavior is caused, and how we interact with one another and our world. You can think of them as our digital twins.

Building our digital twins, running them in simulation, or running them on real physical robots can help us to understand ourselves better and explain ourselves to other people.

Let's start with a couple of AI models that help us to understand (1) why we find it hard to answer questions succinctly and sort through all the relevant material that magically pops up in our brains and (2) why we seem to mess around doing random things until we're way too close to the deadline.

1.Spreading activation networks

...or why our answers to questions are slow and come out in paragraphs.

What is it like being in an ADHD brain? Here's the one thing that so many people I've talked with about it say YESSSS, that's exactly what it is like. And no-one's described it like that before.

Here we go.

Imagine that, when I say something to you, there's a network in your brain, one node of which lights up. It recognises, matches with, the concept I've just expressed, the question I've just asked you. (A node might be a neuron or a cluster of neurons.)

Let's suppose that node is connected with about 5 other nodes - think branching spider's web - and each of those is connected with about another 5 nodes each. And activation spreads from the initial node first to its neighbours, then to their neighbours, and so on.

Eventually the activation fades out and no more neighbours are activated.   So you have a network of lit-up nodes. The nodes represent concepts that are connected to that first concept that I mentioned to you a paragraph ago. They're all relevant, but their relevance fades out the further away they get from the starting node.

Let's suppose you're thinking up the answer to my question. Well, that first node activated 5 of its neighbours. Each of those activated another 5 - so we're up to 25 in total. And each of those activated another 5, so now we're at 125. And so on. In a few hops, a few jumps from the initial node, you have 5x5x5x5x5.... (and so on, another 5 for each jump) nodes activated. And each of those nodes encapsulates another concept. So you have all this brain "lighting up" in no time at all after you've heard the question. That's a lot of information sticking up its hand and saying "I'm relevant".

In order to answer the question, you have to decide which of these concepts you're going to pick to put together a story that will pass muster as an answer. There are so many possibilities, some very outlandish, some very sensible-sounding, some that seem like what you want to say but not just now in public, etc.

It takes a fair amount of processing to get a story selected, put in order, and then piped out of your mouth in the form of an utterance. Now, supposing we ask the same question of someone who comes up with an answer very quickly, whose brain isn't wired in quite the same way - tentatively, neurotypically.

Let's say each node is only connected to 2 or 3 others. So for, let's say, six hops, that brain has to consider 2x2x2x2x2x2 = 64 nodes, or 3x3x3x3x3x3 = 729 nodes. Whereas 5x5x5x5x5x5 = 15625 nodes.

If your brain is sparsely wired, 64 nodes. If your brain is lushly wired, it's 15625. A factor of 15625/64 = 244.1406 difference. The lushly wired brain has to do 244 times as much work to come up with an answer. No wonder you have to sit and think for a bit!!!

I don't know if our brains are wired exactly as I've described. But this spreading activation, through an incredibly-well-connected network, with quite a few hops, in response to a question or concept, is how it feels to be inside an ADHD brain.

And so any explanation that neuroscientists come up with of how ADHD brains work has to be able to support this feeling of what it's like.

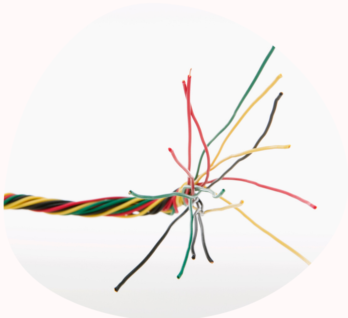

So, what can you do with it?

Well, at least know that you're doing a lot of work with that brain. And tell your questioner that you're thinking and would they please wait for your answer. Or check with them that they do actually want an answer (as opposed to them wanting you to reassure them that you're glad of their presence and you're taking notice of them).

In meetings, jot down a few key words to remind you of what you want to say, and use the hand-raised function in Zoom calls to get to say your piece (talk to the chair of the meeting about this first) so that you don't lose your place in the conversation.

And if someone suggests that your cognitive tempo is sluggish, maybe query the terminology. It sure doesn't feel sluggish. Unless we're talking a really speedy slug!

2. Reinforcement learning and getting things done when the payoff is far into the future - like next Friday.

It's one of the most well-known characteristics of people with ADHD: we find it hard to do tasks where the payoff is some distance into the future. Even when we know it's important.

It's as if the reward we get from achieving our goal just doesn't reach back from the goal to where we're standing at the start-line. If the goal and the reward we get from reaching it are a shining lighthouse, it's as if there's a thick fog obscuring our view. Whereas someone without ADHD is able to see that lighthouse shining far away in the distance and to gradually and deliberately approach it - by starting the first subtask and then the next and then the next, steadily, inexorably, and in good time.

In AI, this looks very like the area of reinforcement learning, or learning from interaction with the environment to get to a goal and receive a reward.

Young children learning to stack bricks on top of one another, you learning which route is best to take from home to the grocery store, a robot learning whether it should deliver another parcel or head for its charging station, or a computer games character learning the best way through the game level to reach the goal - these are all examples where actions have to be taken but the payoff isn't instant, it's at some time into the future.

We, or the robot or the game character, have to learn by trial and error, but once we've learnt, we can see the path to the goal. And then we just have to choose the actions that will get us to our goal as quickly as possible.

Let's imagine ourselves as a software agent controlling a game character that is heading through a grid-world, moving one square at a time, avoiding obstacles, pits and barriers, to get to the sack of treasure at the goal state. (You have to think of this as a simplified equivalent of us doing our homework assignment, where the goal state is to hand it in just before the deadline and the actions are answering each question in turn.) The goal square has a sack of treasure on it – 100% of the payoff. And the game agent can see, or sniff, the treasure from a distance, but the amount it can perceive drops off the further away it is. It might drop off 5% with each step. So one step away from the goal, it only sees a payoff of 95%. Two steps away it drops off another 5%, down to about 90%, and so on, the further away it is from the goal. Ten steps away from the goal the payoff it can see is down to about 60% which is still quite reasonable – it can still tell which direction to head in.

What does this have to do with learning in actual human brains? The passing back of the payoff down the chain of actions from the goal that's ahead of us to where we are now is – roughly speaking – implemented by dopamine in our brain.

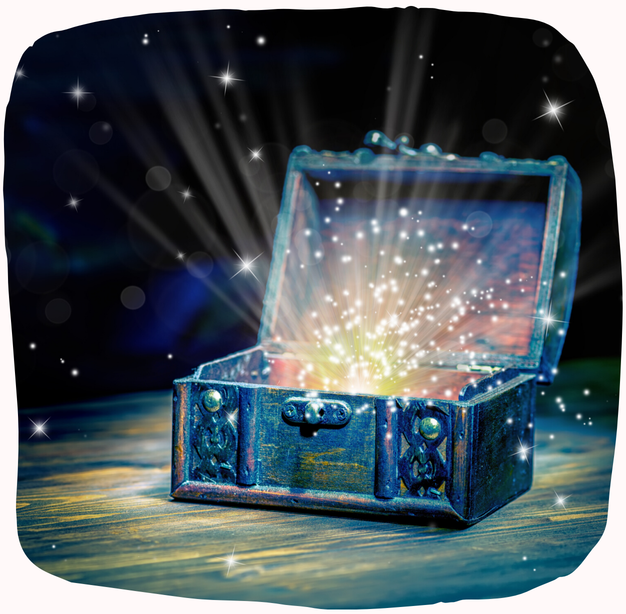

Here's the rub.

ADHD brains are insensitive to dopamine. So the payoff more likely drops off faster in ADHD brains – actions just don't look that rewarding even when we're only a few steps away from the goal. We can't tie more than a few actions together, not easily or naturally, that is.

Let's say the payoff drops off at 15% with each step (as it might do in a dopamine-insensitive ADHD brain). By the time the agent gets to 10 steps back from the goal the payoff it can see is down to a measly 19%.

So a brain that's not great at passing back this payoff from the goal really doesn't find actions very rewarding until it's rather close to the goal. Checking Insta or doing some discretionary housework looks like it'll give us (because we are that software agent!) a much bigger reward in the short term.

It's not that we're lazy or bad. It's that our brain isn't **able** to see doing the first question in our assignment that's due in 3 weeks' time as particularly rewarding.

So we do something else that **does** seem rewarding.

Whereas someone whose brain is very sensitive to dopamine – maybe neurotypical – can find that first question rewarding 3 weeks before the deadline. Because the dopamine doesn't drop off as quickly and that brain sees quite a good payoff from starting the question now.

So, what can you do with it?

Again, feel better! You're not morally decrepit!! You have a brain that's not very sensitive to dopamine and this is the logical consequence.

As for getting tasks completed before the deadline, it's a matter of splitting the tasks up into smaller pieces (as usual), and then finding ways of getting yourself a reward as you complete each piece. So that you get intermediate rewards along the way and don't have to wait for that payoff that's sooooo far into the future. Or find other ways of getting some (useful) dopamine. Like doing the assignment with someone else who's working at the same time (body double). Or competition - try to beat your classmate at getting a subtask finished. Or try to beat yourself by allowing yourself only a certain time (I bet I can get this finished inside an hour). Or promise someone else that they can read it through at the end of the day. And, of course, some ADHD medications lead to having more dopamine in useful parts of the brain.

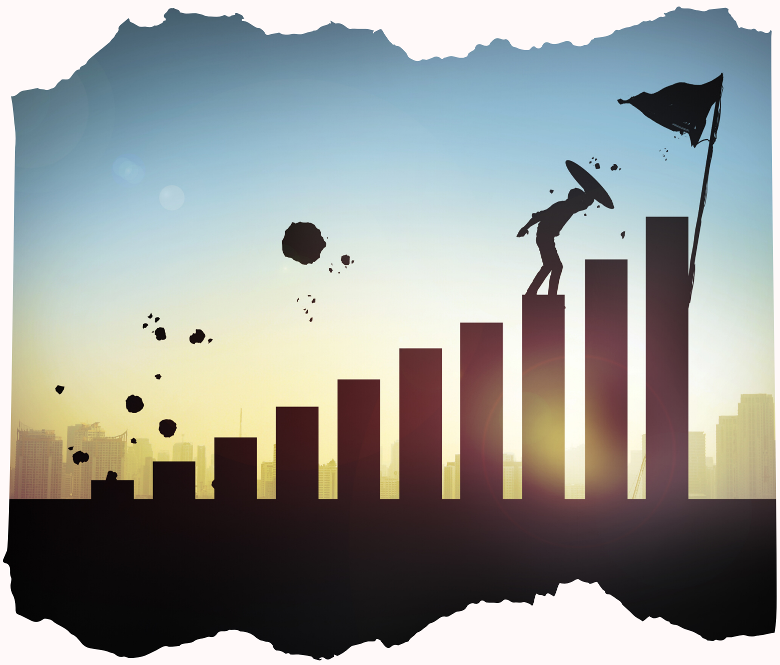

Things to do with your ADHD brain

Models give us a way of thinking about how our brains work. What comes next are a few useful things we can do.

They're inspired by some of the topics that go to make up the whole area of artificial intelligence: neuroscience, computer programming, a bit of practical philosophy, robotics, engineering, mathematics, game design, imagination, and so on.

3. Rehearsing it in your head

Mental practice makes it better, though probably not perfect.

One of the things we want our robots to be able to do is to pick up and manipulate objects - teacups or tomatoes - with their arms and grippers. And we'd like them to learn this for themselves, so that they can deal with all varieties of teacups and tomatoes placed all over the table top in front of them, not just a few specific items for which we can program the arm's movements explicitly. So they need to practise. But it might be expensive or time-consuming for the arm to practise doing movements using the actual hardware. So we give the robot a representation of itself - almost a "sense of self" - and let it "practise" in its mind/brain/computer.

When it comes to doing the grasp for real, it still might not be perfect but, if the robot's model of itself is true to "life", the grasp will be better than it was at the start and the arm will grab the cup or tomato.

Likewise, in humans, practising something in your "mind's eye" and feeling it in your body too can be helpful. It sets up some neural pathways devoted to the movement or task. (Like athletes at the start of the 100m, visualising and feeling every step of those hundred metres, before they get down into their blocks.) Of course, imagining and doing are different. Your brain can tell the difference. But they are close.

What can you do with it?

Use mental rehearsal when you have to do something that you really don't want to do. For example, at 3pm I am going to sit down to write my homework assignment. Or at 6.30pm I am going to go out for a run. Or at 10.30pm I will switch off the TV and get ready for bed. When the time comes, you won't feel like doing any of those.

Let's take the run. Imagine yourself sitting there as the clock hits 6.30pm. You don't want to go out. It's too hard. It's not fair. I should be fit effortlessly. I should be having my tea. I should have done this earlier today, at 6.30am. That's the time of day for virtuous activities. Boohoo. And then imagine yourself putting your running shoes on and going out the door.

Here's the thing. You don't imagine yourself feeling all Pollyanna-like and how wonderful it is to be going out for exercise. You DO imagine yourself with all that griping and complaining and inner two-year-old grousing, and -- AT THE SAME TIME -- you imagine yourself putting your shoes on, opening the door, and heading out. Your mental model has to be as true to life as you can make it. You're not trying to suppress the strife and argumentation. That'll happen for sure. So instead, ADD ON the virtuous activity you're committing to, and see yourself going out of the door and complaining loudly, simultaneously.

Do this many times - practise. And when the time does eventually come round for you to go for your run, you'll have primed yourself to take the actions (get your running shoes on) even though you're complaining and it'll be that bit easier. And you know that when you get outside the grumbling will drop away after about 30 seconds and you'll be off. Do this with any activity that you know you won't want to do when the time for it comes around.

4. External memory: it's easier to read it than to remember it

AI aims at replicating human reasoning capabilities in machines. It also aims to come up with intelligent-looking behaviour for robots and computer programs that is implemented in non-human-like ways. Intelligence-as-it-could-be, rather than simply mimicking human intelligence.

Robots have computers for brains. In common with human brains, they don't have infinite amounts of memory. But no-one says they have to be built the same as humans. We have very limited amounts of short-term memory - the memory we use to keep the information we're focussing on at any given moment. Whereas robots can actually have a lot of short-term memory, if that's how the robot builder decides to do it. So in respect of short-term memory, robots can do rather better than humans.

How can we get round this challenge of limited short-term memory? We don't have a solid-state drive that will automatically store everything. But we do have access to external memory - a lowish-tech version at that.

What can we do?

Write it down.

If we want to make a point in a conversation with our boss, we put it down on a notepad and take it into the meeting with us so that it doesn't disappear in the stress of the moment. If we've set our intention – our stated aim – for the next half-hour of work, we write it down so that it's there for us to come back to after a distracting chat with a co-worker. (Use Sticky Notes on a desktop.)

If we're testing a program with a bunch of different parameters or doing experiments with varying concentrations of reactants (or pick your own area of study), we write down what we do. (Lab notebook?)

Why? It's a lot easier to read it than to remember it!

For the same reason, if you're working at a computer, using a big primary screen along with a second screen is much more efficient than using a small tablet. You can spread your windows and apps and programs out and simply move your head to look at them, rather than having to remember what's hidden behind your top window. Dual monitors are good for productivity.

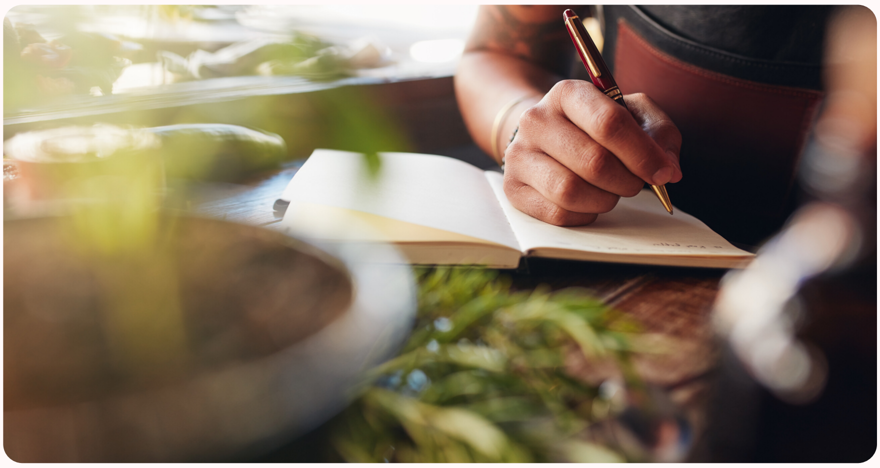

5. The world is its own best model, or, stuff is good!

Robot builders used to think that a robot should have a complete model of the world – its environment – in its "brain".

The robot would read its sensors, reason about what to do on the basis of the model, and then move. The trouble was, it was really hard to build a complete model of the world in all its gory detail. The only things that went into the model were those the robot builder thought to include. And they were approximations.

A door would appear as a perfect rectangle. A wall would be totally vertical, and at right-angles to the next wall and the floor. Whereas real walls and rooms have rounded edges, with chunks taken out of them, scrapes, bits of leftover paint, dust. A robot feeling its way on a perfect model of an imperfect world would soon get lost.

The roboticist Rod Brooks (who dreamt about robot vacuum cleaners – that became Roomba) coined the phrase: "the world is its own best model".

Let's suppose the robot wants to get across the room to the opposite corner to get to its charger. Instead of plotting an ideal – but theoretical – path on a map and following it, and hoping the room hadn't changed since the last time the map was updated, the robot might instead just move parallel to the walls keeping a certain distance away as measured by a laser ranger (the sort of thing you use to measure a room before you calculate how much wallpaper you need) and keep going around the edge of the room until it sees its charger.

Much easier, and no map needed. Just a way of moving along a wall and recognising a charger. If the room changes and the map is out of date, so what? No need to build a complicated model of the world. Just let the world be itself.

Why is this useful?

Stuff is good!!

Or, at least, it can be.

Controversial! To be enjoyed with caution!

If you need to take your gym kit to work with you tomorrow morning, put it on the floor in the doorway of your apartment. Position it so you can't leave the place without treading on it. You'll have a much better chance of taking it with you.

The alternative would be to make a list of everything you want to take with you and check that before you leave for work. But why? Your gym kit might change depending on what's clean. Why not just get everything together and put it by the front door.

Many people with ADHD like to use these "launching pads" by the front door for keys, change, ID lanyards. No need for a model - a list in this case. Let the world - your stuff - represent itself.

We also find it useful to stack the things we need to do on a big pile, or several spread-out piles. This can be good, if we are going to deal with the piles in a definite amount of time. By the end of Saturday, rather than "soon". In making the piles we've done some of the work of sorting and categorising and bringing order. There's nothing to say that you have to put everything away every evening if that doesn't serve you. But (let's be sensible) to avoid the piles becoming part of the furniture, do put an expiry date on them. Marker pen and a big post-it. If you've not cleared them by that date you have to put them back where they came from (which means you will likely forget them and therefore have to use another strategy for bringing them back to mind) or you give someone else in your household permission to get rid of them (which gives you external accountability along with your deadline - both useful in helping us to get things done).

The caution: I'm not discounting the fact that people with ADHD can have a big problem with too much stuff overwhelming them. But using its space-occupying properties can be helpful too.

6. Doing a long list of things when we can't get started...

...because we don't know exactly how to do every last thing on it in minute detail

Some of the useful tricks from AI just come from the type of computer programming languages you have to use if you're going to easily replicate human-type reasoning.

This one comes from the language Prolog - programming in logic. It works with lists.

How to do a long list of things:

- Do the first thing on the list.
- Run this method on the rest of the list

That's it.

When it runs, you would take your list of stuff. You do the first thing. And then the next. And then the next. And so on till you find your list has nothing left on it.

In practice, we often get overfaced with long lists of tasks, especially when we don't quite know how to do them. So it works more like this:

How to do a long list of things:

- Pick the first subtask off the list. This subtask can be as small as you like, as small as possible.
- Specify that subtask in more detail.
- Do it.
- Take the rest of the list and repeat.

Why is this useful?

It means that we don't have to know in advance how to do everything on the list before we even start.

We have to break off the first subtask only, chunk it down some more if we like, specify what it takes to do it, and carry out those actions. And we also have a method that we can follow to deal with the rest of the list. Knowing that often makes us feel much more comfortable doing the first thing, comfortable enough to start, comfortable enough to act even though we don't have complete knowledge of how to do the whole complete task.

Example: how to clean the living room

- Clean the window ledges
  - First subtask: take the plants off the ledges.
  - Specify & do: vacuum the ledges with a brush attachment.
  - Specify & do: put the plants back on again.
  - Specify & do: clean the rest of the living room (apply this method to the rest of the living room).

7. What time do you go to bed?

Aehm, when I feel tired? vs At 10:50pm

This next one is even simpler and extremely useful.

We turn nebulous concepts – like bedtime – into things. We reify. (Res is Latin for "thing".) For example, many of us go to bed when we feel tired. And our sleep patterns are all over the place. We don't know the answer to the question "What time do you go to bed?"

Whereas if we reify, we turn our bedtime into a specified thing, a time in this case. 10:50pm.

Why?

Because then we can compute with it. Do calculations with it, ask questions about it. For example, suppose it's 10pm. Now I can ask: Is there sufficient time to watch another episode of NCIS: Los Angeles before my bedtime? Answer: No. (So don't watch it!) Or, if it takes me 15 minutes to prepare and eat my supper, what time do I need to start it? Answer: 10:35pm. (Set a timer.)

Reifying is very useful for concepts to do with time, because we generally have a very imprecise notion of time and dealing with time by "feel" doesn't often turn out well.

Examples:

- Time-at-which-I'm-eating-today's-breakfast (because then you know that feeling tired 4 hours later means you need some more food, not a lie down) (check the clock while you're eating your scrambled eggs)
- Day-on-which-I-change-the-towels (Saturday. If it's Saturday change the towels, else don't.)
- Number-of-spare-printer-cartridges-I-keep-on-hand (two or more. So if replacing a printer cartridge leaves me with just one, buy some more straight away.)

So for anything for which the answer is "Aehm, when I feel like it", usually said in a not-quite-sure tone of voice, try reifying it.

Make it into a thing and assign a value to it.

Work with me

You have amazing ideas and wonderful talents. And you finally want to get them out of your head and implemented.

You want more time for the things you're really good at vs the things you should do.

You want the chance to figure out what you're really good at.

You want to stop dancing between being stuck and reacting to external circumstances.

You want to stop living life on hold and re-activate yourself.

Contact me for 1-to-1 personalised coaching and get your own unique intelligence expressed.

Contact me

Email: gillian@brightshinycoaching.com
Tel: +44 7767 143125
Website: www.brightshinycoaching.com

Follow me