THE TRANSMITTER | BRAIN INSPIRED

Alison Preston explains how our brains form mental frameworks for interpreting the world

Preston discusses her research examining differences in how children, teenagers and adults integrate new information into their memories.

12 MARCH 2025 | by PAUL MIDDLEBROOKS

This transcript has been lightly edited for clarity; it may contain errors due to the transcription process.

Alison Preston
There's always a challenge linking psychological concepts to actually brain computation, right? Because the brains don't think in terms of these schemas. That's not what the brain is really doing. What our brains, especially the mature brain, because I studied development too, which we can talk about that, what mature brains are doing is they're trying to get things right 80% of the time. They're building heuristic models of the world that get us adaptive actions most of the time.

This is where the large language models are interesting is it allows us to actually make these predictions. It allows us to understand the moments like where they're going to choose to encode something with high dimension versus a lower dimensional representation later. The large language models are useful in terms of an information theoretic point of view and helping us understand what does a stimulus look like and what quantitatively can we link to human memory for that stimulus?

[music]

Paul Middlebrooks
This is "Brain Inspired," powered by *The Transmitter*. Schemas, you know them, you love them, but do you know them? Do any of us know them? Are schemas in our brains? How do schemas relate to our ongoing cognition? What is a schema, anyway? All right. The concept of a schema goes back at least to the philosopher Immanuel Kant in the 1700s, who used the term to refer to a built-in a priori mental framework to organize our sensory experiences.

It was the psychologist Frederic Bartlett in the 1930s who used the term schema in a psychological sense to explain how our memories are organized and how new information gets integrated into our memory. Fast forward, another hundred years to today and we have ourselves a podcast episode, a "Brain Inspired" podcast episode with my guest today, Alison Preston, who runs the Preston Lab at the University of Texas at Austin.

On this episode, we discuss Ali's neuroscience research, explaining how our brains might carry out the processing that fits with our modern conception of schemas and how our brains do that in different ways throughout our development from childhood to adulthood. You may have just noticed I said, "Our modern conception of schemas," or something like that. Like everything else, there isn't a complete consensus among scientists these days or probably ever exactly how to define schema. Ali has her own definition. She shares that and how it differs from other definitions commonly used.

I like Ali's version and think that it should be adopted, and you'll hear it in just a minute. I think it should be adopted in part because it helps distinguish schemas from a related term, cognitive maps, which we've discussed a plenty on previous "Brain Inspired" episodes. Primarily, the issue is that the term schema and the term cognitive map can be used interchangeably. We discuss how to think about schemas versus cognitive maps versus concepts versus semantic information. Manifolds come up as they often do on this podcast.

On the last episode with Ciara Greene, Ciara and I talked a little bit about schemas and how they underlie our memories, our learning, and our predictions and how they can lead to inaccurate memories and predictions. Today, Ali and I really talk more about how circuits in the brain might adaptively underlie this process as we develop and we talk about how to go about measuring it in the first place. All right. Are schemas in our brains? Is schema a psychological construct? I'm going to let Ali sort that all out in a moment here.

I link to some of the papers that we discuss and to Ali's lab in the show notes of the episode. I am Paul Middlebrooks. You are awesome for being curious and being here. Thank you for listening or watching. All right. Here is Alison.

[transition]

When I was an undergraduate at the University of Texas at Austin back in the 20th century, they didn't even have a neuroscience major. I don't believe. Not that I would have.

Alison Preston
No. It started as the Department of Zoology actually.

Paul Middlebrooks
Is that where it came out of?

Alison Preston
Yes. It was well before my time, but it was in the zoology building and it emerged from the Department of Zoology. It was under this director, the Directorate of Biological Sciences, and actually the neuroscience department at UT is only about a decade old, right? Interestingly.

Paul Middlebrooks
Is that right?

Alison Preston
Yes. It used to be like a division, and now, it's full department and our major is actually quite new relative to things. It's been an evolution since you've been there.

Paul Middlebrooks
You started at 2007, right?

Alison Preston
Yes. This all happened while I was here. I was recruited through the Center for Learning and Memory, and I came in through a research center. I was appointed in psychology, and then there was a division of neuroscience under this Directorate of Biological Sciences, but then they decided to make it a full department. That happened during my time here at UT.

Paul Middlebrooks
That's great. Ali, it's nice to finally have you on because I was looking back and you've been on my list. I don't remember when I first emailed you.

Alison Preston
It's been years, yes.

Paul Middlebrooks
It's been years. About five years ago, I interviewed Brad Love. It was through reading about his work. We talked about concept learning and his sustained computational model, which you've used in your own research. That's how I believe I came to know your work and then reached out to you. For whatever reason, it's gotten kicked down the road multiple times. It's nice to finally have you on.

Alison Preston
I'm glad to be here.

Paul Middlebrooks
All right. My main goal in our discussion here is to better understand schemas. To do that, we'll talk about a host of your own research, but there are these different abstract levels of conceptual tools that we use to understand minds and brains. Schemas is one of them, but it's one of those concepts. The more I read about it, the less I think maybe I understand about it. Sometimes you talk about schemas, sometimes you talk about cognitive maps, and now, I don't know the difference between those, perhaps. I want you to strangle this out for me.

Alison Preston
Yes. My own thought on this has evolved over time because I think the word schemas obviously comes from Bartlett's work in the '20s and '30s. Certainly, how we use it now is not probably how he would have used it. I think it's important for us to think about what operational definitions are, right? My own definition for schemas is, and I've pushed myself to think about this and what they mean, especially when you're a neuroscientist, because obviously, there's always a challenge linking psychological concepts to actually brain computation. Because the brains don't think in terms of these schemas, right? That's not what the brain is really doing.

Paul Middlebrooks
Is it not? Okay.

Alison Preston
In terms of what the actual cells are signaling, certainly not, but what arises from it.

Paul Middlebrooks
If there's not a schema in the brain, it's still a psychological term.

Alison Preston
I think the way I think about schemas nowadays and to differentiate them from other forms of knowledge, like semantic knowledge or concepts, is that there's a couple of things that you have to think about. One is there not a moment in time, a schema is about a sequential set of actions. You can think, often we talk about schemas or narratives. We think about, going to the airport. You have a schema for what it's like to go to the airport. That involves, even things that happen days before you go to the airport. I have to arrange my ride or arrange my parking at the airport. I have to pack for that visit.

Then the morning of I'm going to put my things in the car or the vehicle, which whether that's the tube or a bus or a car to go to the airport. When you get to the airport, you tip the driver who helps you with your bags, then you go in, you check in, and you go to-- There's this routine, right? After you go through security, you get coffee at this specific place that you prefer in the airport. Then this is typically where you wait for the flights. When you prefer to queue up, if you're like one of those people who sits in your seat until your group is called, which is no one nowadays, or whether--

Paul Middlebrooks
I do that.

Alison Preston
I admire you. I'm too impatient to do that. If you're one of the people who herds the gate, and that goes-- I think of schemas as having this multiple set of sequential actions, that you expect. It's like a knowledge base. I think that's what makes it different from a concept, too. I think that's where we have to differentiate what a schema is from a concept because a concept doesn't necessarily have the sequential set of routines and actions for how something should unfold.

I think schemas are used right as predictive models, so that they help us understand, like given this set of principles, A, B, C, what is the next thing that's going to fall in line? What is the adaptive thing for me to do now? For me, schemas goes beyond concepts. You can think about a concept of a dog and it has relational features. It has multiple features, like dogs tend to have fur, they tend to have four legs, they most bark, though not all. There are these stereotypical features, they have relational structure, but they're not the same as a schema.

In some sense, they're different from, and that's where I think we have to, as psychologists in particular, think about how we're going to operationalize these different terms to think about how they might be differently instantiated by the brain. For me, I think schemas involves this sequential aspect, at least to some degree. This idea that action is part of it, like schemas have-- and that action can be overt or it can be like decision-making actions that they have us make predictions, make choices, and do behaviors as a function of what those schemas are.

I think that it may not be the perfect definition. My definition keeps evolving for this because I suppose there's other forms. This is where it gets a little trickier for us psychologists. You can also maybe have like a George Washington schema, the president of the United States.

Paul Middlebrooks
Wait a second though.

Alison Preston
That's the thing. That's the thing. I think this is where people differ. I would not call that a schema. I would probably call that a concept of George Washington because it's more stationary involves relational factors. I bring it up. Some people might not agree with me and I think that's okay. My own definition of schemas is different than a concept of George Washington because it's more dynamic. I think of schemas is more dynamic.

Paul Middlebrooks
I really like that because it helps me differentiate what a schema is and what it isn't. I was just talking with Ciara Greene who studies human episodic memory. I think she would disagree with your definition in the way that she uses it. Because the way she uses it, it's more like a concept or a cognitive map almost. Her thing is like all about when you're forming new memories, you have these relations between the things in your memories and you would agree with that, but--

Alison Preston
I totally agree with that.

Paul Middlebrooks
Yes. That's where schema, I think, ends conceptually in her conception of it. In your conception, you need those associative things to be all about some goal or some action moving forward in the world.

Alison Preston
Yes, exactly. I've pushed myself to think about this carefully because I also studied categories and concepts and cognitive maps. What is different--? My thought is that you can use cognitive, and we can talk about cognitive maps because that's another hairy definition, too. You can use cognitive

maps, you can use concepts and categories in a schema. You might have a schema, like they're embedded things. We talked about the airport example. Within that airport example, I talked about going to your favorite coffee shop.

You can have another schema or another concept of a coffee shop that dictates what I do there. The way I think of schemas, it's not that they're detached from cognitive maps or concepts. It's like hierarchical, it's a level above, like what makes it different from those things? Otherwise, you might as well use them synonymously as the same. Why do we have two words for what is essentially the same thing?

That's how I've chosen to think about a schema being different than a cognitive map where a cognitive map is like a set of interrelations that comes from not just direct experience, but inferred connections between things. It doesn't necessarily have to be static itself, but it gives you a global map of the relations between things, but it doesn't necessarily have this sequential action plan that goes along with it.

Paul Middlebrooks
The cognitive map, at least the term I think, and everyone points to Tolman who studied-

Alison Preston
Space.

Paul Middlebrooks
-rodents in mazes. You already have this dynamic environment. The idea is that these rodents who were allowed to sort of explore the maze before they ever had a task to do in the maze, would be able to make shortcuts, use shortcuts in the maze to solve the task. Therefore, hey thought Tolman suggested that they have this kind of model, this cognitive map in their head, which allows them to skip running through the entire maze and figure out the shortcuts and stuff. That's already a dynamic. A lot of the way, it studied is in navigation.

Alison Preston
Yes, heavily spatial.

Paul Middlebrooks
Only recently with your work and people like Tim Behrens have, has that concept been extended to logical relations between things, and between concepts. Only recently has that static aspect been explored more perhaps.

Alison Preston
That's what I'm saying. Maybe we can think about it as a continuum. There's maybe places where cognitive maps and schemas, but I think cognitive maps-- I think you can use them flexibly and it doesn't necessarily have a specific sequence of relationships you have to do. You can almost navigate it as needed, by the way the rats do, through these shortcuts, whereas to some degree schemas do have a routine associated with them.

Again, in my personal definition, how I've chosen to make some boundaries around these different psychological ways that we refer to what we're trying to understand what humans do or rodents do, and specifically, nerve brains do. We have to differentiate them if we're going to study them. This is where the field is at. Like you said, your previous guests, maybe they won't agree with me and I think that's okay because that gives us lots of interesting science to test, right?

Paul Middlebrooks
Yes. Right. Her whole thing, like yours also, is that-- well, one thing is the memory is reconstructive.

Alison Preston
Yes, it's constructed.

Paul Middlebrooks
We talked all about the sort of faults of our memories. One of the reasons why we have "faulty memories" is that when we're processing something and trying to store it, we're processing it within a schema. The whole nature of a schema is to generalize knowledge. This would be integration in your terms, perhaps.

Alison Preston
That's one of the good-- you do integration and then it can be generalized. Yes, exactly.

Paul Middlebrooks
Therefore, when you actually put something, put a piece of knowledge into your schema, you're actually losing some of the detail.

Alison Preston
I don't think you can do that with cognitive maps too, because the way I think about it is the individual experience that give rise to a cognitive map or schema, they may fade away in memory, but that doesn't mean there's not still specifics left in the schema. I think this is also where we have to push ourselves about thinking through these things. I agree in general, what our brains, especially the mature brain, because I studied development too, which we can talk about that. What mature brains are doing is they're trying to get things right 80% of the time.

They're building heuristic models of the world that get us adaptive actions most of the time because it's hard for brains to be perfect. Brains have limited capacity and resources. You have to have algorithms that allow you to get things right most of the time. Then the consequence of building these adaptive models of the world, let's call them schemas, is that sometimes they will be wrong, that you'll make a judgment where there's be some exception to that general pattern that you have in your brain.

That's one thing is that you'll make non-adaptive decisions when the context is somewhat varied, but sometimes as a function of having these routine schemas and your brain, you'll make assumptions about the world. That's one thing of false memory is, it's an assumption. Because it's adaptively built, like this is true, most of the time, I'm going to assume it's also true right now. That's, I think, where false memories come from. You're making a prediction about an associative network. That prediction turns out in that circumstance not to be true.

Paul Middlebrooks
It doesn't matter if it's true as long as it's useful, right?

Alison Preston
Exactly. I think that's true. Sometimes false memories can be adaptive, sometimes they're very not, sometimes they're detrimental to a certain degree. I think our brain is using its limited computational power to get things right most of the time. That's, I think, where schemas and cognitive maps help us. The flip side of the coin is sometimes those predictions will be wrong.

Paul Middlebrooks
I was going to bring this up later, but since you already mentioned, you were just discussing this in terms of the mature brain, but one of the things you study is the schema formation, and the brain processes and areas involved as we develop. Then I thought, wait, 80% of the adult brain. All right. Teens, you've shown actively suppressed details from the past, but children don't. Where's that percentage? Let's go ahead and talk about that.

Alison Preston
I wasn't saying 80% of the brain, but 80% of the things that you do are right. It's more like the adaptive action thing. Here's what I think is going on. In our mature work, we've shown that a part of this knowledge formation, whatever we want to call it, schemas, cognitive maps, it relies on interactions between the hippocampus, very well-known, important brain region for memory. That hippocampus interacts with other structures in the prefrontal cortex and the parietal cortex, which do executive function, high-level decision-making processes, et cetera.

Some of the pathways connecting the hippocampus to, for instance, the prefrontal cortex are among the last to develop. They don't develop until the third decade of life. That has to mean something for knowledge formation. I think we're in the early days of really understanding the developing brain in terms of how it forms and uses knowledge, because for decades, until maybe two decades ago, people thought the memory systems were done at the age of seven. It's only been really fairly recently we've chosen to change our view that memory systems continue to develop through adolescence.

There's a lot we don't know. I'm going to caveat my ideas with that saying, this is a very open field and we need to know more. My thought about the developing brain is because even the hippocampus isn't mature through adolescence. When you have a learning experience, so one of my work does is when you're learning something new, we know it's not disconnected to the past. Even when we're talking now, and I'm talking to you about schemas, you're bringing to mind what your previous conversation with your past guests.

That's influencing not only what you're attending to now, but what you're encoding and how you're encoding what I'm talking about now. We use our preexisting knowledge to guide how we learn about new experiences. Because of the lack of maturity of the hippocampus, even in middle childhood, those retrieval processes in children, 7 to 10, they're not as readily reactivating prior knowledge during new learning experiences. They're treating each learning experience more in isolation.

You have to do a lot to cue past memories for them to think about that this thing is connected to that other thing I experienced two weeks ago in class. That's partly the retrieval mechanisms in hippocampus aren't even mature in middle childhood to allow them to make connections between learning experiences separated in time. That requires a--

Paul Middlebrooks
Another way to say that in your terminology is their children are like super differentiators of knowledge.

Alison Preston
I'm going to come to that. They don't even reactivate. They end up with orthogonal representations for related experiences too because they use different ensembles to code them at different times. I think adolescents are super differentiators. We'll get to this. Adolescents, what's happening is the hippocampus is coming online. They'll start to think about past experiences. They'll think about this previous podcast that you did, because of the lack of maturity of the executive function regions of the frontal parietal cortex, they don't know what to do with that information. They do something different with adults.

Whereas you, as an adult, are going to link those experiences together to form like a schema of a schema for lack of better terms. I'm being a little cutesy here. What happens in adolescence is they choose to suppress it, overlapping content.

Paul Middlebrooks
Choose to suppress it.

Alison Preston
I'm anthropomorphizing the brain a little bit. The brain ends up suppressing that information. They end up with highly differentiated traces. It takes a top-down process. When you have overlapping memory experiences, there's a conflict there. Adults choose to resolve that contact by relinking them. Adolescents choose to resolve that contact by differentiated them and memory. They can still do things. A lot of what I study is inferential reasoning as well. When it comes to past, I later ask you to use knowledge in a way where you have to connect those learning events.

I teach you, A event goes with B event, and you later see B event go with C event. I ask you to infer these relationships between A and C. A child and an adolescent can do that. An adult does it much more easily because they've already formed a direct relationship between A and C in their memory. They have a link that they can use to infer that. They're fast and act at those decisions. Children and adolescents can still do it, but it's much more effortful because they have separate memory traces, whether they be like orthogonalized or truly differentiated.

They can go, well, A went with B, B went with C. They have to retrieve multiple memories, recombine them at the time of inference. They can still get it right, but it's a much more effortful and probably error-prone process. It's not that we don't think children don't reason and can't reason, it's that they're the way they form their knowledge is actually fundamentally different than adults. The way they're making these decisions is off of different underlying knowledge. This is something we're really pushing in our latest work where we're doing work now to really look at this.

We're doing longitudinal work to look within subjects during adolescence about this shift in knowledge formation within individuals across time to show when they start forming knowledge in different ways. We've made these inference decisions where there's no correct answer. The correct answer, we can objectively use it to model what kind of underlying knowledge structure you have. The choices you'll make will be determined based on your underlying structure. It won't make you wrong or right.

It just makes you make a different choice based on your underlying knowledge structure. These are things that I think we want to push is the idea that if I have different types of knowledge, I actually will make different choices at different ages because that underlying knowledge differs. We want to push that a little bit and understand that. I think maybe I've been talking-- I talked about a parent, but the most elegant idea about what's going on in adolescence, I think comes from some rodent fear conditioning work. I don't know if these studies.

It's a group who, in fear conditioning, you stick a rodent in a chamber and you have a foot shock. They learn to fear that particular spatial context. In a version of this, the experimenters put in young rodents, like kid rodents, adolescent rodents, and adult rodents. They all in the immediate situation. They put them all in the chamber and then they froze. What rodents do in foot shock, they'll freeze because that's their fear response.

Paul Middlebrooks
By foot shock, it's like shocking someone.

Alison Preston
It's like a mild shock that you would get from touching an outlet or something like that. It's not damage.

Paul Middlebrooks
An outlet by sticking your finger?

Alison Preston
Yes.

Paul Middlebrooks
It's not like sticking your finger in an outlet. It's like some static electricity that you feel.

Alison Preston
Yes, exactly. Anyway, they all shocked. What happens though, is you have a retrieval period where maybe 24, 48 hours later, they stick the rodents back in the context and they don't shock them, but they see what they do. They see if they immediately freeze when you put them back in the context. In this set of experiments, the kid rodents froze, the adult rodents froze, the adolescents did not.

Paul Middlebrooks
Oh, man.

Alison Preston
Okay. Here's where it gets super interesting. Then accidentally, they reuse these teenage rodents in a separate experiment when they were adults. The rodents matured. They happened to stick them back in the context and they immediately froze. It's not like that memory was not there. They chose not to use it during that adolescent time period. In a third set of experiments, they showed that it was related to memory suppression within the hippocampus. Those teenage rodents were actively suppressing hippocampal memories during that retrieval of the fear context.

The way I think about adolescence, and I'm not the only one who thinks that, is that adolescence is a time for exploration. You're trying to gather as much information about the world as possible. You don't care as much about how situations relate to one another. It's just trying to curate a body of what are all situations like for me. This is really interesting. Later, you're may be able to resolve all the relationships between those things. In adolescence, you're not necessarily, your brain is not necessarily exploiting what it already knows as much as exploring the environment as much as possible.

Paul Middlebrooks
The active suppression part is the really interesting part. When you're talking about those rodent studies, my mind went to complementary learning systems there. I thought, well, maybe what's happening is, which would be fascinating, is they encode the memory in their hippocampus as teenagers still. Then over time, the idea of complementary learning systems is that the hippocampus then is sending that information to the cortex, which takes a long time to sort of generalize and get encoded in the cortex. I thought, "Well, maybe that's just what's happening," but no--

Alison Preston
Yes.

Paul Middlebrooks
Go ahead.

Alison Preston
They didn't do that. They didn't measure that shift in the systems that in those studies. It's an interesting idea. It could be true that maybe because the reason they retrieved it later in adulthood is because it's not like a hippocampal memory anymore, but it's still got there. That could easily be true. I think that's true. At least in the moment, it was still in the hippocampus and they were actively suppressing it, whether it required a hippocampal retrieval when they stuck them back in the cages later. I don't know.

I don't think they did that particular manipulation to test that, but it's an interesting idea. I think there is probably some transformation that is happening that could be related to complementary learning systems. It's an interesting idea.

Paul Middlebrooks
The behavior of teenagers would suggest there is definitely active suppression going on. [chuckles]

Alison Preston
Yes, exactly. When we talk about this, especially if you remember your own teenage years or have teenagers as children--

Paul Middlebrooks
Your kids are what? 18? Your kids are 18?

Alison Preston
18 and 15. Yes. I see my daughter transitioning from this. I think it's intuitive. Anecdotally, we can see like, "Oh, yes, that is what my kids do." We know this intuitively, but yet we don't know why and why this is adaptive for the human brain, and what at time points it's transitioning. I think that's important for understanding how to educate children because the way you might teach a 21-year-old in college should be different than the way you teach a 15-year-old, should be different than you teach a 7-year-old.

Thinking about these things, to make learning environments the most adaptive, you have to understand what the brain's capacity and what the brain's tendencies are at this particular points of time. That's where we hope to go with this research. I also think it imbues a lot for mental health disorders and just looking at risk. A lot of these things about-- we're talking about are-- could explain why teenagers are more risky than others because they're so explorative. They're not going to exploit what they already know or what even they've seen other people do as much as have to experience itself.

I think there's a lot of interesting implications about how our memory systems work and change over development for everyday behaviors and childhood and adolescence that I think hopefully we'll be able to think about.

Paul Middlebrooks
In terms of schema formation, then one way to summarize it, and you should correct me here, is that when you're a child, you're just trying to get your foot in the door, trying to get a grasp on things. Then as a teenager, you start actively suppressing because your foot is in the door. You feel like you have a solid base. Now, it's time to explore.

Alison Preston
There's much information as possible in the system. Yes.

Paul Middlebrooks
Then many of us that dies down coming into adulthood when we're perfect.

Alison Preston
When you're perfect or you're learning to exploit what you know as much as possible to build these schemas that are adapt-- You learn what's correct 80% of the time, and you're okay with being wrong 20% of the time. You shift from this trying to have as much knowledge about specific instances to having the most generalizable knowledge base that you possibly can. That's how we talk about it. Psychological terms is this shift from having instance memory to having completely connective, adaptive relational schematic memory.

That's where what you learn in one context is easily applied in any other context in which you're involved in. I think that's what's shifting. I think it becomes a shift from exploring your environments as much as possible and then becoming this more adaptive system in adulthood. I think what's really interesting is some of the things we're thinking about is that a lot of things-- One of the questions that I think you wanted to ask me, so I'll skip ahead in our script is like how my own thought about neuroscience and the human brain has shifted.

Paul Middlebrooks
Yes.

Alison Preston
I've been influenced by computer science and artificial intelligence because we use a lot of those tools in our research, but I've become an information theorist [chuckles] in some ways.

Paul Middlebrooks
What does that mean?

Alison Preston
When you think about, what is the actual information coming into the brain, how many bits does it have, what does it look like, how much capacity does the system have to encode those bits, how does it have sampling algorithms to sample that information as efficiently as possible? You can think about mutual information theory, you can think about perplexity and surprise, and these models. I think it's important for us to understand the stimulus as much as what we do with it. It also relates to human attention, too.

If you think about information theory, and I only have limited capacity, so I have to attend the most important feature to the environment. What is important depends on my goals. It depends on what the task is and what the goals are. We've just talked about different goals at different ages at very high level. The goal of an adolescent is to explore the environment. The goal of an adult is to exploit it is what they know as much as possible. That actually has interesting implications for the dimensionality of memory representations at different ages, too.

This is yet to be shown. This is something we're very interested in testing. The idea that in adult, there's going to be a U-shaped curve in terms of dimensionality. In fact, what we've talked about is if adolescents are suppressing past knowledge when learning new experiences, informing all these differentiated traces, they're going to have very high dimensional representations. Whereas as you form a schema, there's a schemas inherently involved dimensionality reduction. We can circle back.

That's where the false memories and fuzziness maybe come in sometimes, is you reduce dimensions that are less goal relevant. When you have to fill in those dimensions later, when you are using your knowledge basis, you may fill in those compressed dimensions incorrectly.

Paul Middlebrooks
You have to sample from the exact right space of that low dimensional schema representation.

Alison Preston
Exactly. I think what there's also like, and this is where thinking-- Information theory is useful in thinking in terms of how the structure of knowledge is changing, allows us to quantify the structure of knowledge across time, whether that be within an individual or across development in ways that test these theories about how schematicized is human memory and how that's evolving. You can use dimensionality as a proxy for how right schematic. What I'm thinking, I'll point out, I alluded to this before is I don't think schemas are always fuzzy.

You maintain the details in a schema that are important to behavior. Sometimes those are highly resolved, and other parts, which tend to be less important, will be less resolved in that schema.

Paul Middlebrooks
I know that schemas are supposed to be adaptable and flexible because we're always updating them, but are some of the features of a given schema more crystallized and less adaptable are you saying?

Alison Preston
I don't know. I think you keep the things in the schema that are more resolved or those that help differentiate what to do in certain experiences. You can have a restaurant schema, but there's probably sub restaurant schemas. What I do in a fast casual restaurant is different than what I do in a formal sit-down restaurant. They share features, and they share lots of features than my goal in the restaurant, is to go in and get myself fed. How the routine of actions and predictions that you make in that restaurant differ by whether it's a fast casual or a sit-down restaurant.

It's like those points that differentiate different possible decision things that are going to be highly resolved. In this context, I should do A, but in this other context I should do B. That's where schemas are going to be more resolved. Whereas in the schema if the things about a fast casual restaurant, that are similar to like a formal restaurant, those will be compressed in the schema, where the action patterns are the same for both. I think where you have the resolution is where it actually leads to different choices.

You can use that to generalize in a new setting. I can't think of a new restaurant setting, but if I walk into a new setting and something is surprising, like I do an action, then it ends up being socially embarrassing or something, that's where you'd want to update your schema and say, "Oh, I need to have a new branch on this, because where I filled in that other-- my current action pattern is now wrong in this context."

Paul Middlebrooks
When my ex-girlfriend and I walked into a all-nude restaurant without knowing it's all nude.

Alison Preston
Exactly. It's like, "This is what I should do in all-nude restaurants. I need a new schema for that. I need a new schema for that."

Paul Middlebrooks
Yes. [chuckles]

Alison Preston
That's where I think it's like you have levels of schemas where differentiate different action patterns in different contexts. What I do in airports is different. My action pattern for what I do when I'm traveling alone is very different than when I'm traveling with my family. This is what I'm saying. It's not that schemas are super fuzzy. They're fuzzy where they need to be and resolve where they are--

Paul Middlebrooks
They're like the solution to everything. That's one problem with ideas like this. It's like you want to make them fit all of your problems. This is how they could address that. This is how they could address that. It's also important to think about what a schema is not.

Alison Preston
Exactly.

Paul Middlebrooks
I want to come back, though, to information theory in a sec. Remind me if I forget, but so what does a schema not? What can it not do?

Alison Preston
What does a schema not do? That's a really good question. I don't think I've ever thought through that question.

Paul Middlebrooks
Let me give you a side example here. We've just been talking about manifolds and low dimensionality. You used the term manifold? I can't remember.

Alison Preston
Yes, I do. Sometime, yes.

Paul Middlebrooks
I started learning about manifolds and a lot of research in neuroscience has found that a lot of-- when you record lots and lots of neurons, let's say you record 100 neurons, you could say that their spiking activity is in 100-dimensional space because every neuron can have its own firing activity. What's been shown over and over recently is that when we do things like reach, those 100 neurons form this lower dimensional manifold structure, which is a smooth, low dimensional thing. That the trajectory along that manifold dictates when and how we reach.

I started thinking about manifolds and everything looks like a manifold now. I think you could have two neurons and that forms a manifold. It almost explodes. Now, I'm at that point where like, "Oh, shit, everything's a manifold. That means nothing now."

Alison Preston
Yes, I use the word manifold in front of a systems neuroscientist, too. For the same reason, he went, "Oh, manifolds."

Paul Middlebrooks
Oh, really?

Alison Preston
[laughs] Yes. I think people are using these. I'm not saying manifolds are true, but I do think manifolds are a way to think about schemas, honestly. You can think about trajectories through a state space, which is the manifold.

Paul Middlebrooks
I was going to ask you where does manifold fit? It's at a lower abstraction because you define the manifold by measuring the neural activity. It's somewhat defined by the--

Alison Preston
Yes. We do that, too. We're using those tools to actually look at the geometry of the representations for these schemas. You can look at the distance and direction vectors between two points in between stimuli and a space, and use those to make predictions about whether you have a schema or not. We do things like that. In the simple task example I use here, and I'm not going to call this task a schematic task. Let me be very clear. If I learn stimulus A, let's say it's a basketball. I pair that with something like a coffee cup. Then I later see the coffee cup with a picture of a plant. I'm doing this.

Paul Middlebrooks
She's especially doing this with her hands.

Alison Preston
[crosstalk] Basketball and coffee cup there. Then I have a different triad where I see three other images. In these tasks, we teach these overlapping triads of images. You'll learn, and the ideas we'll ask you to infer, does the basketball share a relationship with the plant? We don't just teach you that one. We'll teach you 37 of them or whatever. What turns out early happens is that in, at least in the mature brain, we know the mature brain from my early work, forms a direct connection between the basketball and the plant.

When I ask you that question, it's easy to infer it. What we've also recently shown is that all the triads align with one another in the brain. The vector that predicts the distance and direction in neural space between the basketball and the plant will also allow you to predict how other A/C relationships and the other triads are doing. They all align in a task space. That's actually a schematic representation. If you can make predictions about how to traverse a neural space and those apply to other aspects of the task, that's a true schema. It's a second level inference in a way.

Paul Middlebrooks
Sorry, but just to take that a step further, when they all align, those are the good performers, but you've also shown that people who are less accurate at those tasks have less aligned geometries.

Alison Preston
Yes, exactly. That's right. The people who don't do that, they may be doing the first level or inference, but they don't align the things. They're less generalizable. There's different layers in which your brain can infer. I can keep A and B, I can keep the basketball and the cup representation separate from the cup and the plant representation. Some people will link the basketball and the plant, but they won't necessarily align it with everything else. There's different layers of maturity. We expect that. What makes those individuals different? We don't yet know.

I can't tell you, predict who's going to be a person who's really good at lying their neural geometries versus those who don't. We have some things, we know the white matter connectivity between hippocampus and the prefrontal cortex actually is predictive of these kinds of abilities in these tasks.

Paul Middlebrooks
The white matter are the axons.

Alison Preston
Yes. It's the way that the hippocampus sends signals into this region and receives them back. Those people who have more stronger connections between those regions are better at these tasks. We have a couple of things that we know are related to these things. We also know that, in development, the volume of the hippocampus, so how the size of it affects, like your ability in this task. We have inklings to what might drive these individual differences, but not necessarily a lot.

I think that's an interesting question is even in adults, you get a lot of variance in the way people do these tasks. The question is why? That's an open question, too. Is it a good thing? Is it a good thing that I have? I don't know that it's always good to have these aligned representations.

Paul Middlebrooks
Right. Because I was going to ask you maybe where does creativity fit into that story?

Alison Preston
I think that's a super fascinating question. We do have some data and others have data linking, like core memory functions to people's ability to do creative ideation. I think there is something to this, like the idea that you can make connections between unobserved experiences is in some ways the essence of creativity.

Paul Middlebrooks
To do that, is it better to be aligned or orthogonal?

Alison Preston
I don't know. I actually don't know. I think I can say this. I do think the reactivation part where you reactivate prior experiences and learning about new things, I think that is very important part of creativity because that's how you make connections between like, this is what's happened to me now versus being able to see like, "Well, that's happened to me later. I wonder about this." To some degree, I wonder about the breadth of the reactivation tuning curve, for lack of a better word, like how broadly you reactivate during a learning experience, probably relates to creativity.

Because to some extent, there's really far, those far connections are the most unusual and maybe most connected. I think that reactivation part of it, like how much your brain chooses to reactivate other things when learning experiences is the essence of creativity. Now, how high or low dimensional your representations have to be for creativity. I don't know. I don't know that I have a preferred hypothesis.

Paul Middlebrooks
It's definitely high.

Alison Preston
You think so? Yes. You think so?

Paul Middlebrooks
Yes.

Alison Preston
Yes. I don't know. There's some people out there. I've seen some work out there that suggests the opposite. The people hypothesis--

Paul Middlebrooks
You know how I know that? I'm a high dimensional thinker, Ali.

Alison Preston
[laughs] Okay. You're right. It's always high dimensional. That prediction would mean that adolescents are more creative than adults.

Paul Middlebrooks
There's evidence for this, right? It depends on how you define operationally creativity.

Alison Preston
Yes. That's a whole-- I study creativity peripherally. I've had trainees in my lab who are really interested in it. It's not been my core focus, but I do think they are really-- knowledge and creativity are intimately related. That's why we've gotten into this. There are overlaps

there I'm thinking about. Creativity is inherently relational. The same way we were talking about these relational memory experiences.

Paul Middlebrooks
There's the old bit of knowledge that mathematicians over 24 are done because they're not creative enough now to solve problems. It's supposed to be a story about creativity, right?

Alison Preston
Yes.

Paul Middlebrooks
I want to ask here because I realized. Let's talk about maturity. Much of your work focuses on the interactions between areas surrounding, and the hippocampus itself, and frontal cortex, medial prefrontal cortex. It's fairly well known or at least in my schema or my cognitive map of things that prefrontal cortex, as you're developing, as your child, you can think of it as all these millions of new connections being formed every day. It's just forming all the connections it can. Then there's a pruning process. That pruning happens in your late teens, right?

Alison Preston
It hits around puberty, right? It depends on whether you're male or female and what your age is. It starts from there.

Paul Middlebrooks
All right. I'm still pruning. Then we forgive teenagers their faults because they can't inhibit things that they shouldn't do, and they're exploratory. It's because the prefrontal cortex is not pruned yet. It's not honed yet, right?

Alison Preston
It's the whole brain that's undergoing pruning. It's not just the prefrontal cortex.

Paul Middlebrooks
I was going to ask you what that story is with the hippocampus.

Alison Preston
It's the whole brain. I mentioned we have these volume behavior correlations. The volume of your hippocampus tracks your performance. What's really interesting is that it's a smaller specific part of the hippocampus that predicts your ability to do this task, which you can relate to pruning.

In some ways, what happens during adolescence is you have fewer connections. That automatically on an MRI scan or relate to reduced volume and those structures. It's reduced volume in the hippocampus that predicts performance and better performance in this task, which is-- We interpret that as related to pruning, right? That you end up with these lower dimensional possibilities. When you're pruning, you're automatically reducing the potential dimensionality of the space.

Paul Middlebrooks
The entire capacity space is--

Alison Preston
Yes. Smaller. You don't have as much dimensionality possibilities. It's actually smaller hippocampi that actually predict performance of this task, which is--

Paul Middlebrooks
That's really cool. Okay. Getting back, because there's a thousand asides I wanted to go on. Getting back to your example of the basketball, and the plant, and relating these things, right?

Alison Preston
Yes.

Paul Middlebrooks
When you were talking about in terms of schema. I've potentially--

Alison Preston
Which I don't think that's a schema. I'll be very clear.

Paul Middlebrooks
That's my question. Is that a cognitive map? Is that what that is?

Alison Preston
That's the beginning of a cognitive map. The cognitive map becomes more the aligned part of it. We've always used that. That's called the associative inference task. This idea that you see the basketball and the cup and then the cup and the plant. It's the beginnings of how we form experiences that go beyond our direct observation, which is essential to both cognitive maps and schemas. It's first level inference that you need to do it. It is itself. You can define cognitive maps and schemas. They have to have a number of relations.

We've tried to use this very simple task to get to the mechanisms by which you can form a more complex cognitive map. I'll be very clear about. That said, this idea that I align all of these different triads, that begins to signal there is a cognitive map in there.

Paul Middlebrooks
That's right. Yes. But not a schema.

Alison Preston
But not a schema, right? Because it's not necessarily-- I think there's not a sequential set of action plans that go there too. That's my personal definition of schema because I'm trying to differentiate it from a cognitive map to make it different.

Paul Middlebrooks
There you go. I asked you, what is a schema not? You answered that upfront by saying what it is. A schema, because it necessarily, in your definition, involves sequentiality and time.

Alison Preston
Yes, and actions, meaning some choice and behavior. Yes, exactly.

Paul Middlebrooks
All right.

Alison Preston
That helps differentiate it from more-- I don't know if static is the right word, but again, I started forcing myself to do this is because how does schema is different from semantic knowledge? That's often the question that people ask too. I had to think about-- that's why I chose this sequential action is a way to differentiate it from it because semantic knowledge is-- does not have that same dynamism that--

Paul Middlebrooks
Semantic knowledge, an example of that would be George Washington was the first president of the United States. It's knowledge that--

Alison Preston
Exactly. Yes. He was a general in the United States Army.

Paul Middlebrooks
He has wooden teeth. That fits--

Alison Preston
He owned Mount Vernon. All of it.

Paul Middlebrooks
We have a cognitive map of George Washington, and that's not a schema.

Alison Preston
We have a concept of George Washington.

Paul Middlebrooks
A concept.

Alison Preston
Yes.

Paul Middlebrooks
How does semantic knowledge fit?

Alison Preston
I don't know. This is where I would like the field to really push and have better definitions. For people to align on definitions. I've chosen to push myself to think about these things, because a category and a concept that- even our definitions for those are a bit fuzzy at times. Let alone a cognitive map, let alone a schema. I think of them as highly related, but where we draw the boundaries between them and how you think about what behaviors and predictions I can make from a category. I can determine whether this stimulus should belong in category A or category B, but that doesn't tell me a sequence of actions. It doesn't have super interlinked knowledge.

Concepts are different. Concepts may be more abstract than categories. That's one possible definition. Where you think of-- I think of things like concept things. What is loyalty? You can probably even think of actions, and behaviors, and people that are loyal, but that's very different than what's a dog versus what's a cat. Even concepts, what's a noun and what's a verb. I don't know. That's very different than what we were talking about, cognitive maps. I think about cognitive maps and schemas is models of the world. Those are different than, at a broader level, than concepts and categories.

Paul Middlebrooks
They could be models of relations between concepts and categories.

Alison Preston
I think that's right. Exactly. They could be that too. I think that's embedded in them. That's why I said there, these are potentially-- I think of them hierarchically is levels of complexity and what's getting layered on at each level of relation. Again, my definitions might prove wrong, but I think we have to force ourselves to think carefully about how we want to define these things.

Paul Middlebrooks
It's an ongoing personal frustration. We have to use words to communicate. Those words are so low dimensional. They can hang on to lots of other things. One of the things that is good about science is that you have to operationally define something if you're going to study it. Do you think that is where-- I was going to ask you this later, but do you think that is where some of the criticisms of psychology from the neuroscience field is because these psychologists, they're using these fuzzy terms that could mean anything? Schema could mean anything you want it to mean, therefore it's not an operational thing to study.

Alison Preston
I do think that. There's always the tension in that. The way to think about this is one of the first right areas that people took psychological principles and try to apply them to the brain early in the advent of neuroimaging was this idea of recollection and familiarity. When you have a memory, you come across a person on the street and you know you know them but you don't know from where, that's a sense of familiarity. You can't maybe recall their name or where you're first met them or when the last time you saw them is. That's a sense of familiarity.

Whereas recollection, you see that person, and you're like, "Oh, the last time I saw them was in my barre class six months ago at this particular studio." You may not even remember their name, but you have more contextual information about where you know them from, why you know them, et cetera. That's just to find these things. For a long time, people were like, "Oh, the hippocampus does recollection. Some other parts of the brain does familiarity." That is both true and not true.

Paul Middlebrooks
It's not even wrong.

Alison Preston
It's not even wrong, but it's vague. The brain computes things, for lack of a better word, through synaptic plasticity, through the way that changes in small, individual cells, small networks, large networks change. They're computing something about that experience. That leads to these feelings of recollection, but they're not recollection themselves. This is where you have to use a low-dimensional word to describe an entire cascades of things that happen to get to that phenomena of recollection. That's how I think about it. It's not that they're wrong, but they're so low dimensional and so poor explanatory of things that happen to happen, the epigenetic, the molecular, the cellular, the systems level that lead to that experience of knowing that I saw that person at my bar class six months ago.

Paul Middlebrooks
We're not good at thinking in those system level ways, maybe. At least for me, I want to go as low dimensional as possible and say, "All right, well, the hippocampus does familiarity," because it's so nice.

Alison Preston
The brain's like shortcuts and heuristics for everything.

Paul Middlebrooks
That's right.

Alison Preston
That's what I'm telling you. We're building heuristics to get things right 80% of the time. We do that in our language as well.

Paul Middlebrooks
What I studied in my PhD dissertation is on metacognition. I don't know what that means.

Alison Preston
Wow. We can have a whole-- We can turn the tables and I can ask you questions about it.

Paul Middlebrooks
Of course, but I did it in monkeys. Actually, your familiarity versus recollection made me think of that because a lot of the metacognitive work is based on familiarity in a lot of the things.

Alison Preston
It's not that there isn't something wrong and that the essence of those principles. It is interesting to think about when I know something versus when I don't, but those are inadequate descriptors of the actual mechanisms that underlie those experiences.

Paul Middlebrooks
I returned to this over and over. I find so much of language is inadequate to express what I want to express or what I want to convey. Not because I'm super high dimensional thinker, just because it isn't sufficient, I feel, to--

Alison Preston
Sometimes. Yes. That's where, thinking about schemas, we have insufficient language to really think about what is schemas a lot of the times. That's maybe why my push to trying to think in terms of information theory. What information is actually coming in, but also what information does the brain store then as a result of that experience? Thinking about it in terms of these information theoretic terms, allows me to maybe quantify something. Then that might then shift my linguistic definition for that principle too. There's a natural give and take there to understand, "Well, my current definition of schema won't explain that information about what's being stored, and so how do I need to adjust it?"

Paul Middlebrooks
I'm glad you brought the information back up because I wanted to come back to this. I didn't know that we would be talking about information theory and Shannon information. One of the, I guess, criticisms-- information is sometimes hot in neuroscience in terms of quantifying things and it's all about information flow. Giulio Tononi's integrated information theory is really hot in consciousness and it's all about mutual information and what's passed on.

Even Shannon himself, Claude Shannon, who brought Shannon Information into existence, cautioned against using his information theory in things like biological systems. One of the criticisms is that to measure information, you have a sender and a receiver. The receiver already has to know the set of possible messages it can receive. What you were saying earlier that schemas are all about your goals and taking actions. Then you started talking about information theory. I thought, what if the set of possible messages are the goals? If your goals dictate the set of possible messages you can receive, if the messages you can receive are defined by the goals, then you could potentially have a viable measure of information. Does that make any sense to you?

Alison Preston
I think so because I think your goal set your attention filter in some ways and what are the potential messages I received? We're not using this in everything yet, but we are using it in-- My lab is starting to study more real-world stimuli, naturalistic narratives together with Alex who's also here at UT Austin. If you haven't talked to Alex, maybe he's--

Paul Middlebrooks
I've emailed him before.

Alison Preston
I'll give a plug for you to Alex.

Paul Middlebrooks
Please do. You Texas people, you take at least four years to get there.

Alison Preston
Sorry. But no, Alex and I are collaborating on some things. We have a joint graduate student. Alex is known for how we encode semantic knowledge. The human brain. He's been using a lot of large language models to study that. Now, we're collaborating because I'm interested in saying, "What do people really remember?" What he's studying is actually memory and knowledge.

We're actually looking at having people listening to stories and then having them recall them later. Then look at how the relationship between the information in the stimulus and how that's driving the moments of time that people are remembering, at what resolution of detail they recall later. There, we're using a lot of information theory too because we're using these large language models that actually can token to token, word to word, can make predictions. In that sense, the receiver and the communicator do-- because we're using these large language models, you do have a lot in the stimulus because these models have been trained on. How to update--

Paul Middlebrooks
Yes, you have millions of possible messages.

Alison Preston
Exactly. Those, we can do that. We're using that as a way to characterize, the moments in time that people actually remember or not and how they sample it. In terms of other more simplistic tasks, your goals do matter, right? Your goals set what you're going to pay attention to. You're going to ignore some facets in there. That's something we have to be careful as psychologists. I say psychologists in purpose because how we design experiments is ultimately in what we instruct, at least the human participants, or how you train a non-human primate or a rodent is going to set their filter for the signals that they expect to receive. That is automatically going to bias them in certain ways towards knowledge.

We have to be careful in thinking about our task creation, when thinking about how we're influencing these structures. That's really interesting to think about. There is work out there with getting back to schemas where people have tried to give children the strategy of connecting and linking experiences. Even if you instruct them to do so, they're less likely to do it. There is some fundamental limit in what they're able to do, but you can give them the goal set of what to expect.

This is work by a research associate in my lab. She worked with Patricia Bauer. Her name's Nicole Varga. She showed this and very similar paradigms to our basketball plant, but more realistic ones. You can learn a fact. These are in the domain of semantic knowledge. I can learn that all baby kangaroos are called joeys. Then they later teach the children blue flyers are a type of kangaroo. Then the inference question is what is a baby blue flyer called? The answer is a joey, right?

Adults are able to do that better than like a six-year-old. Even when you tell the six-year-old, you want to link these things in memory, they still don't do it as well as an adult does.

Paul Middlebrooks
[noise]

Alison Preston
Exactly. That's what I'm saying. You can give people goals and shift them around. Whereas probably in adults, there's going to be a naturalistic variance in how much people link blue flyers to joeys. If you give an adult that goal, you can push them around. The kids, there's still a fundamental limit in their system for their ability to do that. I do think in terms of information theory, the goals, the goals are set an important filter on what are the messages that the recipient is expecting to receive.

Paul Middlebrooks
Do you see bits everywhere now, is that what you're saying?

Alison Preston
No. I use information theory as a tool. It's limited like every tool. I figure, is this going to be useful in the task we have to-- When we intend to use it, we design the tasks in a certain way where we know what we're going to quantify about it.

Paul Middlebrooks
It helps you design a better task if you know what you're--

Alison Preston
Exactly. That's when we use it as we bring it in from the beginning and try to design the tasks to look at those things.

Paul Middlebrooks
We mentioned this a little bit earlier, but I want to get your thoughts on modern artificial intelligence, and what your work with schemas, and development, and how schemas are used over time. When you look at large language models or just modern artificial intelligence models in general, are they laughable? Of course they're impressive, but what principles from your work do you think should be paid attention to in terms of building better models? Do models need to develop like-- do you want your model to suppress certain information for a while?

Alison Preston
I don't think models are ever going to act like humans. They're not meant to do that in some ways. You can design a model to act more like a human, whether that's a laudable goal. I don't know.

Paul Middlebrooks
It doesn't matter. It's happening either way.

Alison Preston
Yes, it's happening.

Paul Middlebrooks
It doesn't have to be your goal.

Alison Preston
I'll talk about how I think about them. I'll talk about the way that we're using AI to answer neuroscience questions. Then we can talk about the reverse, how neuroscience influences AI, which I have less formulated ideas about.

Paul Middlebrooks
Great.

Alison Preston
Obviously, the large language models that they exist now are not human and do not-- the way they learn what they're coding is just-- they've been given stimuli that are very different than what a human experience is. Just by the training data set, we know that they're a little bit different than humans. They're bad at things human-- This is actually improving rapidly now. Two or three years ago, they were bad at things like inference and reasoning. That's been a big push in AI, to make them--

Paul Middlebrooks
That's a common theme in the AI world is like, "Oh, we're running up against a wall," and then the models get bigger and it just steps over the wall.

Alison Preston
Exactly. Yes. They don't fully understand why, what's actually allowing the model to make that step.

Paul Middlebrooks
That's cool.

Alison Preston
It's cool, but it's also scary at the same time.

Paul Middlebrooks
Yes. It's cool in a very scary way.

Alison Preston
Yes. The way that I use models is that we use them as a way to, again, characterize predictions. When I've trained something on a billion sets of inputs or more, I can make predictions about which word is going to follow, which word in a narrative with relatively strong ability because it's characterizing all of human language. That's useful in us thinking about when we're giving stimuli to people, especially these naturalistic

narratives, it allows you to characterize the fine details of the stimulus and then relate them to those of the-- when somebody produces the exact words back at you versus the recall versus when they give you something like blurrier.

Here's an example about this that might be an example of a schema. Imagine that you're listening to a story and a character decides to binge on a bunch of junk food. In the story, the character sits down in front of the TV and they eat a pint of ice cream. They have a bowl of popcorn. They eat a bag of potato chips. If you have a high dimensional representation, your recall for those three exact types of junk food, you're going to give them right back to me.

A lot of people, if you're maybe using a schema to recall, when you're encoding that story, it's like, "Oh, they binged and had a bunch of junk food." Somebody who's recalling it might say just what I said, "Oh, they binged and had a bunch of junk food," and that's as far as they go. Somebody might say, "They sat in front of the TV and they ate a pint of ice cream, ate a bowl of popcorn, and ate a bag of potato chips." That would be a very high dimensional representation. Versus somebody who-- here's where the schema becomes interesting, "Oh, they ate a bunch of junk food. They had a bunch of Twizzlers. They ate a piece of pizza. They give you junk food, but they fill in different types of junk food.

Paul Middlebrooks
Wait, this is the model or this is the--

Alison Preston
People do this. The model you have to tune. It'll either give you the specifics. You can tune the model to give you any of those outputs.

Paul Middlebrooks
Why do people do this? Why do they do both?

Alison Preston
Different people do different things, and this is fascinating, right? What matters is happening-- It's not just what's happening during recall, it's happening-- what's happening during the encoding matters. That's why I'm making this story. This is where schemas also influence what you remember because I can tune out and say, "Oh, they're just eating a bunch of junk food." I'm actually not going to choose to encode the specific items they eat because that's not important to the gist of the story. They're just eating junk food. When you ask the person to recall the story later and you ask them, "Give me as much detail as possible," they'll actually say-- use their schema for junk food to fill in details.

Paul Middlebrooks
As an entry port to get to the details.

Alison Preston
Exactly. That's where you get false memories from. It's like I'm using my schema to reproduce things that I didn't encode with sufficient high dimensionality. Now, you're asking me for it later.

Paul Middlebrooks
That's what you do with concepts also and words.

Alison Preston
Exactly.

Paul Middlebrooks
That's what a word is.

Alison Preston
Yes, exactly. This is where the large language models are interesting is it allows us to actually make these predictions. It allows us to understand the moments where they're going to choose to encode something with high dimension versus a lower dimensional representation later. The large language models are useful in terms of an information theoretic point of view and helping us understand what does a stimulus look like and what quantitatively can we link to human memory for that stimulus. That's where it's useful.

If we learn enough about that, you could actually then take the models and-- This is something that I'm talking about Alex for is you take the models and you build narratives in certain ways with certain properties and you can then give them to humans and say, "Which of these stories are the most memorable? In some ways, you can use the models to actually then create things that will induce a certain knowledge state in an individual in ways that you can manipulate and control.

Paul Middlebrooks
Oh, manipulation and control.

Alison Preston
Exactly. You can use these models to test what is really memorable for people in interesting ways.

Paul Middlebrooks
The interesting thing there is you are dictating how they form their schemas just by generating a story in a certain way.

Alison Preston
Exactly.

Paul Middlebrooks
Which is what authors do.

Alison Preston
Honestly, humans have been doing this for eons, right?

Paul Middlebrooks
That's what storytellers do.

Alison Preston
Exactly. Understanding what makes somebody a good storyteller is an interesting question. It gets back to this creativity thing. A lot of things, it's like we don't know-- There's lots of work on human narrative and human storytelling, but it's not been at this quantitative way. That's interesting to think about. How can neuroscience influence AI? That, I have--

Paul Middlebrooks
Do you see your work as being like, "Oh, why don't they-- They should take this on. This would be useful." Do you have those thoughts--

Alison Preston
Some of the work that we're starting to do, yes, because we're using the same tools. We're talking in the same language now. If we talk in this information theoretic way, it will come out. I've certainly reviewed a lot of papers that are now building neural nets that are meant to do inference that reference our work. They are doing that. They're starting to deploy these large language models on the same tasks that we use. They're trying to have the models solve tasks that we've been working with for decades. There is a natural crosstalk there. The field of AI is working on this. The large language models are about language and only language.

Paul Middlebrooks
People would disagree. Some people would disagree.

Alison Preston
I'll be more concrete. Let me say this. They don't take into account visual-- There are other AI models that work on visual stimuli in that way. We know the brain integrates multi-modal information into a single thread of unit. The AI models that combine-- This is a growing field and there are models that are being built to do this. This idea that the brain naturally integrates multiple sources of information and uses that simultaneously to build schemas of things like that. You don't have AI models that are good at that at all. How do I watch a movie? There's narrative content, there's emotional content, there's visual content. There's humans, there's objects.

Paul Middlebrooks
There's your junk food.

Alison Preston
Yes. There's lots of facets to that. The large language models don't capture the visual, spatial parts of those things because, often, when you watch a movie, there's no words going on. There's just actions. There's not even people on the screen sometimes. How we think about this multi-modal integration, the human brain does that naturally, and it uses those different sources of prediction and surprise to easily-- AI models don't do that yet.

Paul Middlebrooks
How would you build an AI model?

Alison Preston
I don't know. I'm not an AI thing. People are out there working on this. They're trying to build these large multi-modal models that can make predictions off of different sources of data. That's a big push in AI. That's where maybe AI can look to the human brain to understand how the human brain maybe integrates that information. That's just one idea. I want them to build those models because I need them to actually test what is truly naturalistic, episodic experience. I'm eager to do that.

Paul Middlebrooks
In general, you're all for using, for example, deep neural networks as proxies, as models for brain activity to then infer something about how our real brains are doing it. If we can predict the activity of our brains using these models, you're in favor.

Alison Preston
I wouldn't go as far as to say they're a full proxy. They're useful in the sense that they allow a quantification of a certain set of parameters. You can see whether the brain tracks those parameters or not. It probably isn't that they're exactly those parameters, necessarily, that the brain's computing. Again, I'd see it as an extension of math psychology where we were trying to mathematically quantify human behavior using much more simple models than we use now. The idea that they're-- knowing that those math models aren't exactly what's going on, but they're a tool to quantify these things.

I don't see models as trying to instantiate the human brain yet. They are useful tools to say, "Well, here's how they compute a specific computation that makes a prediction about that behavior. I can see if brain activity tracks that computation. That's where I see them being super useful.

Paul Middlebrooks
One of the beautiful things about mathematical psychology models, these really tiny models, is that they confer an enormous degree of understanding. If you buy into the premise that the model is supposed to be addressing, and it fits the behavior, then you have a story. Then you can understand it in terms of very understandable a few parameters. Then one of the problems with these larger models is that it's out the door.

Alison Preston
Yes. Exactly.

Paul Middlebrooks
If you're assessing a very particular question and the rest of it you don't need to interpret, then I suppose that that's--

Alison Preston
That's the set. There are limits to how much-- I want to make assumptions about what those models are doing. For some of the narrative work, we're tuning an attention layer to look at how the attention heads in the model influence encoding too. You can manipulate the models too and test predictions. Again, I'm not going to interpret what's going on in layer 9 of the model versus layer 12 of the model. What the heck is that? There are limits to my ability to understand what those models are, what they should do, and how they should align to the brain. Because they are fundamentally different than the brain.

Paul Middlebrooks
You're preaching to the choir here. It's interesting that there are swaths and swaths of people who believe the opposite. That if their model is doing it, that is the brain. You've just described the brain. One of the interesting things about that is that all of these models are built on a conception of a neuron that is decades and decades old. It's crazy to me.

Alison Preston
They require so much data to train. It's fundamentally different than the brain. Human brains don't take that much to learn.

Paul Middlebrooks
I was thinking in terms of your own children, but we can also think in terms of the students coming through your lab. Does your research alter or affect the way that you're giving this, your allowance for diverse behaviors?

Alison Preston
I'll talk about my children first because definitely being a neuroscientist who studies adolescent development has helped me at times.

Paul Middlebrooks
Has it? Because a lot of what I study, I think, "Oh, I don't even--" it's like a separate world, my home life. It's like I don't apply any of the things.

Alison Preston
I do because sometimes when my teenagers do something, I'll be like, "Oh, you're doing this because your prefrontal cortex isn't developed." It's in very broad way. I'll say this is not going to trigger me because your brain is doing what it's supposed to do. I literally--

Paul Middlebrooks
You told them that? "Your brain is not developed enough."

Alison Preston
Sometimes I'll say it out loud when they're in a mood to hear a joke about it. Honestly, the mantra runs through my head, I'm like, "I'm not going to let them trigger me right now because they are doing exactly what their brain is supposed to do." It makes me react to their teenage behavior, perhaps in a more sanguine way.

Paul Middlebrooks
Would they say the same thing, however?

Alison Preston
Maybe at times.

Paul Middlebrooks
You were telling me offline beforehand that you're amazed that your children haven't hated you, essentially, at any given point during their development.

Alison Preston
They like to spend time with me, yes.

Paul Middlebrooks
They like to spend time with-- That's all you could ask for as a parent.

Alison Preston
No, I feel like it was a parenting win. I don't know what I'd do. I accept them as individuals. Understanding their developmental process and that it's not about-- their reactions are not about me and what I'm doing. It's about what's going on inside of them is helpful.

Paul Middlebrooks
That is helpful

Alison Preston
The other thing about it-- I'll say this, my daughter was diagnosed with ADHD when she was in high school. I was able to pull out my slides. We actually had her genetically tested to understand which medications would work best for her because there's tools for that. We learned that she has-- her D2 receptors are too efficient. They're flushing dopamine out of the synaptic cleft. A little too quickly. That has implications. I pulled out my undergraduate slides for her. I showed her what was going on in her brain. That actually made her feel better about her ADHD to understand what the biology was behind it and made her feel there wasn't something necessarily wrong with her, but that her brain--

Paul Middlebrooks
Yes, depersonalizes it.

Alison Preston
In some ways, her brain was being more efficient in ways. This is the consequence. That has been helpful as a parent. My son, actually, he's 15 now. He's diagnosed with brain cancer when he was three. Certainly being a neuroscientist helped deal with those kinds of things. There's lots of ways that I use my knowledge of neuroscience to help parents. Not always, but sometimes it gives you perspective. In terms of mentorship, mentorship too, I don't know if I use it as much in lab. I use more human empathy to mentor.

Paul Middlebrooks
At that 18 plus age where the previous--

Alison Preston
We have a lot of undergraduates in my lab. They tend to work. I only worked with a small-- I only directly interact with this-- we have anywhere from 25 to 40 undergraduates in the lab at a time. I don't have time to work with all of them. I typically get to know the ones that are there for multiple years and doing an honors thesis. The ones that I'm directly mentoring are probably from the age of 23 and up, 23 into their 30s. They're at the end of the developmental perspective.

Paul Middlebrooks
Part of your hiring process is to scan their brains and see how well they differentiate and integrate, right?

Alison Preston
No, I don't do that. I don't know if I would be able to do-- I don't know if there's a point where I could do that and say, "Oh, you're the type of person I want who's going to be a creative thinker."

Paul Middlebrooks
Maybe one day.

Alison Preston
One day.

Paul Middlebrooks
Okay. Allie, my goal here, first of all, thank you for spending the time. Also, I wanted to say, yes, do, if you see Alex-

Alison Preston
I'll bug him.

Paul Middlebrooks
-tell him to look in his inbox because we had an exchange and I'll have to look it up myself. It was in his court probably two years ago now. Maybe he's not interested.

Alison Preston
I'll do you a solid.

Paul Middlebrooks
Okay. Because I only bother people-- my rule is twice. If they don't respond, I don't--

Alison Preston
No, I'll tell them this has been super fun. I will advocate.

Paul Middlebrooks
My goal here, as I stated at the beginning, was to better understand schemas. You did something right off the bat that helped me do that and talked about the necessity of sequentiality, of time series, essentially, and connecting that to the goals. Part of my whole thing was how do I differentiate this from a cognitive map? That solves that. I feel much better about schemas now.

Alison Preston
Right.

Paul Middlebrooks
It's still a tool-- or sorry, a psychological construct.

Alison Preston
It's still a low-dimensional description of probably what is actually going on, for sure.

Paul Middlebrooks
That's right. Thank you for this. Your work is cool and a lot of fun. Obviously, you're going to continue for many, many years doing it. Thanks for joining me here.

Alison Preston
Yes. Thanks, Paul. It's been great.

[music]

Paul Middlebrooks
"Brain Inspired" is powered by *The Transmitter*, an online publication that aims to deliver useful information, insights, and tools to build bridges across neuroscience and advanced research. Visit thetransmitter.org to explore the latest neuroscience news and perspectives written by journalists and scientists. If you value "Brain Inspired," support it through Patreon to access full-length episodes, join our discord community, and even influence who I invite to the podcast. Go to braininspired.co to learn more. The music you're hearing is "Little Wing" performed by Kyle Donovan. Thank you for your support. See you next time.

[music]
