Toward an Imagination Science

Neuroscience Imagination Retreat

December 2-4, 2016

Philadelphia, PA

Participants:

Jessica Andrews-Hanna, Assistant Professor in the Department of Psychology and the Cognitive Science Program at the University of Arizona

Drake Baer, Senior writer at *Thrive Global*

Roger Beaty, Postdoctoral Fellow in the Department of Psychology at Harvard University

Randy Buckner, Professor of Psychology and Neuroscience at Harvard University

Kalina Christoff, Professor in the Psychology Department and the Brain Research Centre at the University of British Columbia

Elizabeth Hyde, Research Specialist, Imagination Institute, University of Pennsylvania

Scott Barry Kaufman, Scientific Director, Imagination Institute, University of Pennsylvania

Martin Seligman, Director, Positive Psychology Center and the Zellerbach Family Professor of Psychology at the University of Pennsylvania

Chandra Sripada, Associate Professor in Psychiatry and Philosophy at the University of Michigan

Diana Tamir, Assistant Professor of Psychology at Princeton University

Report prepared by: Roger E. Beaty[1] (lead author), Scott Barry Kaufman[2], and Elizabeth Hyde[2]

[1]Department of Psychology and Center for Brain Science, Harvard University, USA

[2]The Imagination Institute and the University of Pennsylvania

Toward an Imagination Science

The past decade has seen an explosion of research into the psychology and neuroscience of imagination, with rapidly evolving literatures on topics ranging from mind-wandering, daydreaming, mental simulation, theory of mind, and creative problem solving. Despite considerable progress, however, several fundamental questions remain: What is imagination, and how do we measure it? Is imagination a fixed ability, or can it be enhanced through targeted intervention? These questions inspired a recent gathering of leading experts on the neuroscience of imagination in Philadelphia for an "Imagination Retreat," featuring neuroscientists Jessica Andrews-Hanna, Randy Buckner, Kalina Christoff, Chandra Sripada, and Diana Tamir. The meeting was organized by Scott Barry Kaufman, Martin Seligman, and Elizabeth Hyde from the Imagination Institute, and was funded by a grant from the John Templeton Foundation, and it consisted of a series of discussions that took place over the course of a three-day period.

This report highlights some of the common themes from the retreat, with a focus on the brain network that is thought to give rise to imagination: the so-called *default network*. The retreat began by discussing the history of the default network and its serendipitous discovery in the late 1990s by Marcus Raichle and other neuroscientists, including one of the Imagination retreat participants, Randy Buckner. Participants explored the various cognitive functions that have since been ascribed to the default network—from mind-wandering to mental simulation to creative cognition—and their experiences as scientists who study these complex mental phenomena. The retreat discussions also focused on the measurement and cultivation of imagination, which touched on practical applications in educational contexts. The report concludes with a central theme that emerged from the retreat: the need for a new field of *imagination science*.

Discovering the Default Network

A primary goal of cognitive neuroscience is to understand how the brain gives rise to our thoughts and behaviors. To study the brain in action, neuroscientists track blood flow in the brain using functional magnetic resonance imaging (fMRI) during some cognitive process of interest. In the early days of cognitive neuroscience, researchers often employed a passive resting condition to compare neural activity during some active task (e.g., memory retrieval). The consensus at the time was that neural activity during such rest periods, where research participants simply relaxed in the brain scanner without a task to complete, reflected a baseline level of unpredictable "noise" in the brain signal—not a stable pattern of brain activity or cognitive process. Although resting conditions were widely used in cognitive neuroscience for some time, a debate as to their appropriateness lead some researchers to question whether some meaningful signal—and cognitive process—existed in this seemingly random noise.

One observation that supported the potential importance of resting brain activity was the consistent pattern of activation that emerged when contrasting activity during cognitive tasks. Using the so-called "subtraction" method—a common approach in cognitive neuroscience where neural activity measured during one task is subtracted from activity during another task of interest—researchers documented the recurring involvement of a set of brain regions, including medial prefrontal cortex (mPFC), posterior cingulate cortex (PCC), and lateral inferior parietal lobes (IPL). This pattern of activity that emerged when the brain was thought to be "resting" was first referred to as the "default mode network" by Marcus Raichle in 2001, based on the idea that it reflected a special metabolic state that the brain returned to when not engaged with the external world. Notably, however, the empirical focus at the time was centered on the metabolic *state* and not the *network* per se: Raichle and colleagues initially characterized the resting-state as a special “mode” wherein the oxygen extraction fraction (i.e., proportion of oxygen consumed) was uniform across the cortex. Buckner, a neuroscientist who played a key role in the discovery of the default network, described Raichle’s seminal paper and the history behind the term “default mode network”:

> “…This landmark paper that is titled Default Mode, is actually an argument that this is a special state, because the oxygen extraction fraction across the brain is uniform…At a metabolic level, it was about the mode, not about the network…It's interesting to see. There's a discussion about the network. There's a lot of meat in that paper, but it's interesting to see how the name got plucked, and got applied to the network, when that actually isn't what the original paper was using the term default mode for. I find that quite interesting as a sort of story of science, how that landmark paper sort of gets used in many ways by the field to crystallize the topic even beyond, in some sense, what is written in the paper itself.”

Although many foundational papers were published in the early-2000s based on the work of Raichle and colleagues, another cognitive neuroscientist at the time, Nancy Andreasen, was among the first to postulate that resting-state brain activity within the default network reflected meaningful cognitive processes. In a series of brain imaging studies, Andreasen discovered that patterns of neural activity identified during cognitive tasks involving episodic and semantic memory retrieval were strikingly similar to the brain regions engaged at rest. Andreasen postulated that such passive resting-states reflected active internal processing, which she referred to as *random episodic and semantic thinking* (REST)—what many researchers now typically refer to as “mind-wandering.” Buckner noted Andreasen’s keen insight into the importance of the network as a source of meaningful cognition:

> “Nancy Andreasen was one of these people who wrote a beautiful paper in 1995 about this network: its presence during rest, and how it relates to mind wandering and spontaneous thought. From start to finish, from the anatomy to the current ideas, she captured a lot of this in 1995. She had written letters to Marc Raichle and myself talking about how, in some sense, it was unfortunate that we were using these conditions as a reference, because they involved all this spontaneous thought and memory. I talked to her about this, and if I recall correctly, I still didn't get it at all. I just dismissed this, and set it aside, and it wasn't until many years later, where I started to think deeply about what she was saying.”

The discovery of the default network lead to an explosion of research investigating its role in cognition. At the same time, researchers also sought to identify other networks of the brain, and to understand how the default network relates to these networks “at rest” and during performance on cognitive tasks. The advent of seed-based functional connectivity, pioneered by Biswall and colleagues, allowed researchers to explore correlations in the BOLD signal between a given brain region and the rest of the brain. These methods were central to the discovery of several functional brain networks beyond the default network, including the frontoparietal control, dorsal attention, and salience networks, among others. Kalina Christoff pointed out during the retreat how this network approach led to a paradigm shift in cognitive neuroscience:

"I think that is a huge shift in thinking yet again, because now we're…talking not just about this one network, but about interactions."

Building on early research showing "deactivation" of the default network during externally-focused cognitive tasks, Michael Fox and colleagues provided evidence that the brain is organized into distinct networks, and that the default network and dorsal attention network tend to work in opposition. A key feature of this competitive framework was that, when engaged with a challenging cognitive task during fMRI, the default network tended to show reduced activity while the dorsal attention and other cognitive control networks showed increased activation. Reduced default activity during cognitive control tasks has long been considered to reflect a reduction of internally-directed thought processes, such as mind-wandering and daydreaming, which can interfere with ongoing task performance. Jessica Andrews-Hanna elaborated on the common finding of default network deactivation in the context of how neuroscience experiments were typically designed:

> "Often those papers don't show activity in the default network and in fact, they show deactivation in the default network. But, again, going to this issue of baseline. What is a baseline? A working memory test is going to drastically deactivate the default network compared to rest. But…since those tasks were not designed…to determine whether or not the default network plays any role, I wonder if it's worth going back and reviewing the literature to see whether, if there was a different comparison condition, control condition that was used, maybe activity in the default network would be present for some of these more basic processes."

The Default Network and Imagination

Building on early theories and evidence pointing to its role in mind-wandering, researchers have since broadened the repertoire of mental activities linked to activation of default regions, including remembering personal past experiences, imagining possible future experiences, constructing mental images of scenes and places, considering the thoughts and emotions of other people, and coming up with novel solutions to creative problems. The seemingly complex and diverse cognitive process supported by the default network raises questions as to whether a common factor underlies their engagement.

Some researchers have pointed to the role of the self, since many of the cognitive processes that activate the default network involve self-referential thinking. Others have focused on the largely stimulus-independent activity of the network, using terms such as “self-generated” or “spontaneous” thought to describe default-related cognition. Although these terms have been beneficial in conceptualizing the default network and its role in imagination, like all constructs, there are important limitations to consider. For example, the term “self-generated” implies that the brain initiates or produces something—in this case, thoughts. However, one could extend the term “self-generated” to include essentially all activity of the brain, from perception to executive functions, because all mental activity originates in the brain. Andrews-Hanna reinforced this issue of terminology, noting that “internally-directed, self-generated certainly encompasses much of cognition.”

Likewise, the term “spontaneous thought” seemingly captures some of what the default network does, especially in the context of mind-wandering—a mental state often characterized by spontaneity of thought content. Christoff, who studies spontaneous thought, also commented on this term during the retreat in the context of constraints on imagination:

> "I see spontaneous thought is in the superordinate category of different phenomena in the mind and that, essentially, spontaneous thought is a free movement of the mind. So, when we change from one mental state to another, and we do so in a relatively free way, free of constraints—and the constraints could be external such as a task, or they could be internal such as emotional concerns, or anything that draws our attention. So, this is a state of essential freedom of mind, where the mind is free to move to anywhere it chooses, and the more free it is, the more spontaneous it is, so something like dreams are probably the most free, healthy thinking that could happen where the mind moves anywhere it wants."

Although the default network is commonly associated with spontaneous thought, recent research has also linked the network to goal-directed cognitive processes, which are typically considered less spontaneous in nature, such as working memory tasks. These tasks tend to recruit regions of the frontoparietal control network—a set of brain regions involved in cognitive control and executive functioning. Neuroimaging evidence suggests that the default network can cooperate with the frontoparietal network during working memory tasks that require people to draw on internal representations. Chandra Sripada discussed this recent work at the retreat:

> "…in general, what we find is that there's concurrent activation of intentional goal directed subserving systems with individual components of the default network, the most common being MTL [medial temporal lobe]. Which makes sense that sometimes goal-directed intentional systems would have to rely on internally-generated episodic representations. It makes sense that what is fundamentally an intentional goal-directed

> process could activate components of the default network, but it remains to be seen whether there are really clear cases of frontoparietal network cooperating with the entirety of the default network."

This observation is consistent with recent theories that emphasize a more nuanced approach to studying the default network, focusing on subsystems within the greater network. These models build on early research demonstrating a fractionation of the larger-scale network into component systems linked to specific cognitive functions. Although many of the cognitive processes linked to each subsystem appear to tap aspects of imagination in general, each subnetwork may also make domain-specific contributions. Moreover, a subnetwork approach acknowledges the complex anatomical organization of the larger default network which, although functionally correlated at rest, is comprised of a massively expansive collection of distributed neural circuits. Buckner elaborated on this point:

> "…what we're talking really about is these higher order association circuits that are distributed throughout the brain, that are tremendously expanded in humans, have properties that allow you to detach from your current environment and think and explore, but they're not one thing."

Sripada also spoke to the issue of "lumping" vs. "splitting" when it comes to understanding the default network:

> "I think it sounds more conservative to be a splitter and not try to create the more abstract lumping hypothesis. But, they're on equal footing. I mean, if you fail to lump when lumping is the right thing to do, you've made an empirical mistake. If you've lumped incorrectly, you've made an empirical mistake. So, I'm not afraid to lump and I think that there are data points that have to be captured. So, Diana [Tamir] made a really good observation that the default network, like MTL has a concrete detail-oriented function whereas, the dorsomedial system involves itself in conceptual representations and abstraction. That may make you hesitant to say that they're part of the unified network. Alternatively, there are views that can make sense of why you would want a self-referential system to interface with an episodic memory system, which in turn connects with the conceptual representation system. That would be a view that says that spontaneous generation of self-referential episodes helps enhance conceptual knowledge about the world. And that's an attractive view and it could be pursued. So, it remains to be seen whether, in lumping, there are good generalizations to be had that unifies some of this or not. Fair enough, or not. It could go either way. At this point, I'm not willing to say that the lumping, at all, looks like the door is closed to it. It may very well be that there are powerful lumping hypotheses that might win out."

Another important function of the default network is mental simulation. Diana Tamir, a neuroscientist who studies social cognition and mental simulation, described a deep learning, hidden layer hypothesis that may explain how the brain carries out complex mental simulations:

> "Think of a computer science neural network metaphor, where there are hidden layers within the networks and you're trying to go from input to output. The episodic simulation that is really rich and vivid is this, closest to the observable layer of inputs. The neat thing that we can do with our minds is run these simulations, you know? We can start off in one point, have some stimuli in our environment that have led to some information that's more highly accessible to them and then we run a simulation. We start there and there's some transitional probabilities that bring us through this abstract space of cool dreams or daydreams or imaginations. If you think about this layer embedded within multiple layers, some of which are hidden, then you can make predictions. The function of a simulation is to understand something about your environment or to be able to function better in the future. You can run a lot of simulations and that's almost a generation of information for yourself, if you have this hidden layer and you can model transitions between the more observable layer and this hidden layer. The hidden layer acts as a method of abstracting upon this information that you've created for yourself."

As mentioned earlier, another central component of default network engagement is self-relevance. Several studies have reported default network activity during fMRI tasks that require people to process self-relevant information. Further evidence for the role of the default network in self-related processes comes from studies of mindfulness. Studies with expert meditators, for example, have reported a downregulation of default regions, both during the resting-state and tasks that require mindful awareness of one's thoughts. Because mindfulness training often involves suppressing mind-wandering (e.g., focused meditation), one possibility is that this decreased default activity may reflect a simultaneous decrease in spontaneous mind-wandering.

Christoff described her recent work on the neural basis of mindfulness during the retreat, and highlighted the importance of other brain systems beyond the default network:

> “We've done meta-analysis to see what kind of systematic functional differences happen with different meditation practices. With open monitoring, also mindfulness, you get a lot of executive regions that become altered after meditation. There are a lot of issues, like what does it mean for that to become altered? Are you altering the resting state, or are you altering some functional mechanisms?”

Neuroimaging studies with expert meditators have also reported expertise-related increases in functional connectivity between default network regions and brain areas involved in focused external attention. Increased coupling between default and control network regions raises interesting questions about underlying cognitive processes in expert meditators. For example, such coupling may reflect “open monitoring” processes that allow meditators to observe their internal trains of thought with a sense of detachment. This interpretation is consistent with a seminal fMRI study by Christoff and colleagues of mind-wandering showing co-activation of default and cognitive control regions during periods of “meta awareness,” or moments when participants reported being aware of their mind-wandering episodes during the fMRI task.

The ability to monitor and control spontaneous thoughts could be beneficial for creativity and imagination as it may allow people to direct and manipulate internal trains of thought toward a specific creative aim. To date, however, the extent to which meditation benefits imagination remains somewhat unclear. Although a recent meta-analysis of behavioral studies found a modest association between mindfulness and performance on creativity tasks, specific forms of meditation (e.g., focused mediation) that encourage the suppression of spontaneous thoughts may not be conducive to creativity, and appear somewhat antithetical to imagination. Scott Barry Kaufman, the Scientific Director of the Imagination Institute, shared his thoughts on mindfulness and creativity, based on his recent experiences with a mindfulness training program:

> "It increases clarity of mind, I've experienced that first hand, absolutely, so I see the benefits of it. I think I feel less creative though, after one of the sessions. I still feel like there's a tension somewhere there that has never been fully resolved…I feel like when I have the greatest clarity of mind, I feel like everything's organized in my brain and I'm calm and the mind is quiet, or the default network is quiet."

Another brain system that plays an important role in imagination is the salience network. The salience network is involved in several complex cognitive and affective functions, from cognitive control to emotion processing. But perhaps its most central function is *salience detection*—orienting attention to behaviorally-relevant external and internal stimuli in a bottom-up fashion—a function described by Christoff:

> "Understanding salience is really interesting to also try to apply the lower level of processing…because a lot of salience assessment happens very low level. For me, my thinking on that is influenced by the literature on emotional salience. There is psychology literature on that independent of neuroscience. The idea that something like the amygdala is not assessing, is not detecting threat. It's not detecting fear. It's not detecting disgust or

> even happiness. It’s detecting emotional salience and I think a lot of regions, including the dorsal cingulate and the anterior cingulate are processing a form of salience with slightly different flavors. That's why pain is very salient and it's something that we evolved to alert us to salience.”

The dorsal salience network has also been found to be critical for facilitating interactions between the default and cognitive control networks during creative thinking tasks. But how does the salience network “know” when to make this switch? Current frameworks conceptualize the salience network as a multi-level gating system that acts to “sift through the gold” of internally-generated thoughts and flag salient stimuli for further processing in a hierarchical fashion, consistent with neuroimaging evidence showing a graded pattern of salience processing stemming in limbic regions and moving to cortical salience hubs such as the ACC. Christoff described her views on the salience network, and how it relates to other brain systems:

> “I think what happens it is that process of sifting through the gold. Like, what is gold and what is not, happens on multiple levels. It happens through the salience network, for a lack of a better term right now. Let's call them salience mechanisms…I think, for me, the salience mechanisms need to remain a flexibility of being instantaneous or having a phase shift. There's a benefit in the system and I think that's maybe why we evolved to have this variability by default in our thoughts because you want to be able to shift, detect something salient. You want to shift from detecting something that's not salient or having an easygoing moving thought with nothing in particular interesting to something that's

interesting. You want to have that contrast and I think that's where the salience mechanism comes in there. Just essentially there to flag the salient things.
I think that starts from the level of limbic mechanisms, including the amygdala and probably nucleus accumbens and all these striatal mechanisms going to the level of cingulate cortex. Anterior cingulate cortex, probably not just the dorsal but the rostral and to regions such as anterior insula, which is a big component of it. They kind of interface. They tell us: ‘I've been internally-oriented, but suddenly something important has come up,’ or ‘I've been watching the environment, it's been boring and now suddenly some really interesting internal thought has come up.’ They kind of act as this switch sometimes in terms of states.”

In terms of imagination, when one is engaged with an internal train of thought (e.g., searching for a creative solution to a given problem) the salience network may function to identify novel and interesting ideas for further examination and processing by higher-order cognitive systems associated with idea evaluation. Although this notion of the salience network makes some sense conceptually, empirical support for the role of the salience network in imagination remains minimal. One of the key questions going forward will be to address how the salience network “knows” that spontaneously generated ideas are worthy of further cognitive processing. This process of discernment may be a fruitful direction for future research on the role of brain systems and imagination.

Imagination Interventions

One question that frequently came up during the retreat concerned the potential adaptive functions of imagination. For example, if mind-wandering and daydreaming can lead to detrimental outcomes, why do they occupy nearly half of our waking life? From an evolutionary perspective, why would the brain dedicate so much time and resources to something that's not conducive to survival? Retreat participants explored the potential benefits of imagination in the context of education in a classroom setting. For example, the tendency to "go inward" during instruction could be adaptive by allowing the mind to reflect on and integrate recently encoded information, which could have downstream benefits for learning and memory. If so, educational interventions could benefit from harnessing natural fluctuations of attention that allow students to switch between an internal and external focus, a point addressed by Martin Seligman:

> "I noticed as an undergraduate that in a lecture I would listen for about 60 to 90 seconds and take notes, then I'd stop listening and I'd take notes on what I was thinking about. I'd go inward, the last thing I heard and for about 60 seconds take notes on what I was thinking about what I had heard…It feels to me that my wellspring of creativity is that I leave the table every 90 seconds for about 60 seconds to contemplate, free associate to what I just heard externally and then of course I miss the 60 seconds. Thirty years ago, David [Meyer] gave a colloquium in which he said that the most basic, well replicated fact in all of human memory and learning is spaced versus mass trials, that we learn better when things are spaced than massed…Of course, this just follows directly from the notion that when you have space between material, then you can go inward…and when you have massed, the time is excluded."

Because children are inherently curious and exploratory, intervention programs that aim to enhance imagination by capitalizing on this natural cognitive and behavioral flexibility could be particularly fruitful. For example, educational research suggests that conventional classroom settings might constrain imagination and creativity. Interventions that target the classroom context by modifying the structure and daily routine of students may therefore encourage more imaginative behavior in children. In general, such interventions might be beneficial to the extent that they maximize opportunities for internally-directed exploration and minimize externally-oriented activities. At the neural level, one would expect these approaches to promote default network engagement and corresponding imaginative processes, while decreasing the frequency and intensity of the dorsal attention network—a brain system that is routinely engaged to complete the many external demands in a classroom. Scott Kaufman raised this issue during the retreat:

> “We could actively be robbing children of the opportunities for developing their imaginative capacities by the kind of way we force their attention. It is true that there are some anti-correlated brain regions. Well, there is some deep implications there for pedagogy, right? The dorsal attention network…is highly active whenever we are being forced to focus on an external stimulus, like a teacher in front of us. Isn't that actively suppressing our ability to develop these key capacities?”

Developmental creativity research suggests that young children perform well on tasks that require the generation of original ideas (e.g., the alternate uses task). Yet although their ideas tend to be original, they may not always be practical or useful in solving creative problems. This relative lack of idea evaluation abilities may correspond to the protracted development of frontal brain networks that support executive functions. Interventions that encourage evaluative skills in children may be beneficial, then, to the extent that such skills can be harnessed in the absence of a fully developed cognitive control system. Seligman raised this question at the retreat:

> "We often think that children have on their mind theory of mind, another mind, so the question is in some ways: What can you do if it's sense of audience? What can you do to enhance a child's theory of mind?"

Current imagination interventions with children typically involve some form of play, with the goal of promoting idea generation, imagination, and originality. Although play has largely been used to enhance originality in children, it can also be used to facilitate higher cognitive processes such as executive control and theory of mind. In this context, future intervention programs that employ play-based methods to increase both idea generation and evaluation could be particularly fruitful, because idea generation and evaluation are critical components of creative thought.

In addition to education-based interventions, researchers are beginning to develop methods to better assess components of imagination. Such novel assessments, in turn, may inform intervention programs that can significantly increase imaginative abilities across the lifespan. The Imagination Institute, supported by a grant from the John Templeton Foundation, has recently accelerated discovery in the science of imagination by funding several research projects aimed at developing imagination-based assessments and interventions. Using behavioral and brain-based techniques, these projects have the potential to dramatically increase our understanding of the critical cognitive and neural mechanisms underlying imagination across different domains, from daydreaming to mentalizing to creative problem solving and beyond.

This work, along with ongoing research in the fields of mind-wandering, mental simulation, and creativity, may ultimately contribute to a comprehensive and unified science of imagination.