Introduction. Human memory is imperfect. We encounter far more information each day than we have the capacity to remember. In our daily lives, failing to remember something is a nuisance; we stumble through social encounters with people whose names we can't retrieve; we lose points on academic exams when facts escape our recollection. But in military contexts, failing to remember information — how to operate a key piece of machinery, for example — can be much costlier.

Fortunately, the information we do remember is not random. Previous work has revealed that we are better at remembering information associated with rewarding experiences,[1] information associated with strong negative emotion,[2] and information we are explicitly told will be valuable to recall in the future.[3] These insights into the mechanisms of memory have already been applied to boost the efficacy of content delivery in certain educational settings.[4] However, the utility of these findings is limited because they focus on reward and value signals that are *present at the time of learning.*

In real-world environments, explicit signals of the value of information are often absent, in part because the utility of information *changes* across contexts. Contexts themselves are defined by the co-occurrence of the people, things, and scenarios we typically encounter within them. This means that *within a particular context,* we are likely to *re-encounter* common items again and again, such that the utility of remembering information associated with an item is *higher* for items that appear more frequently.[5] For example, if a person spends the first day of a week-long trip trekking through the jungle and encounters the same type of snake multiple times, remembering whether that snake is poisonous is likely to be useful since she will probably encounter that type of snake again. However, if the trekker decides to cut her trip short and fly home to New York City, the value of remembering information about the snake will decrease sharply. In other words, *since the utility of information dynamically shifts across contexts*, we cannot rely on static value signals to shape how we control memory encoding. Instead, we must use our *knowledge of our environments* to determine what to prioritize in memory. But little is known about how we rapidly acquire and apply knowledge about our environments to shape future learning. My research aims to address this gap by examining the neurocognitive mechanisms that influence how individuals use the structure of their environment to determine what information will be most useful in the future, and in turn, prioritize memory for that more important or useful information.

Relevance to the Department of Defense. This line of research is highly relevant to the Office of Naval Research's (ONR) focus on Human and Bioengineered Systems within their Warfighter Performance Program (Code 34), as well as the Army Research Laboratory's (ARL) Human Sciences campaign. Specifically, my proposed research will advance the ONR's agendas within its Cognitive Science of Learning and Cognitive Neuroscience of Executive Control programs, as well as the ARL's Human Behavior and Human Capability Enhancement initiatives. Within these umbrellas, the ONR has recently prioritized research on "coaching strategies for fast-moving, dynamically evolving military tasks" as well as "creating options for future (perhaps unanticipated) naval decisions," while the ARL aims to discover "methods that can accelerate learning" and "mechanisms governing the differences between individual humans." My proposed research connects to these programs by *probing the mechanisms that underlie individuals' ability to optimize both learning and executive control in novel and variable environments*.

Warfighters must be prepared to deploy rapidly to a wide range of changing environments and operate constantly advancing technologies. The number of different situations that soldiers may encounter is massive, and so *directly coaching trainees through every possible scenario is not feasible*. Additionally, it is impossible to narrow these scenarios to those that are most important because *the importance of information changes across contexts* as wartime environments and weaponry change and advance. Rather than learning every needed piece of information ahead of time, soldiers need to learn *strategies* to infer information utility from the environment around them and allocate memory resources accordingly. Additionally, past work suggests that the executive control systems that underlie value-gated learning continue to develop throughout late adolescence and early adulthood, and even then remain highly variable across people.[3] To develop optimal learning strategies for a wide range of individuals, my research aims to examine *how people to use the structure of their environment to estimate the utility of information and modulate memory encoding accordingly.*

Aim 1. Determine how environmental structure influences memory. I propose that two potential processes that rely on different neurocognitive mechanisms and emerge over different developmental time-courses may enable people to dynamically prioritize memory for different stimuli across contexts. If individuals first learn about what information is likely to be useful in their environment, they may rely on *top-down executive control systems* to enhance attention when they encounter high-value information. But if individuals learn about the structure of their environment *after* they initially encode to-be-remembered information, they might enhance memory for more useful information through *bottom-up, stimulus-driven reactivation*, prompted by encountering relevant cues. I will run two behavioral experiments to determine whether individuals can prioritize memory for important information *through either or both of these different processes.*

Experiment 1. First, people may be able to learn about the structure of their environment and use that knowledge to determine what information is likely to be useful. Then, when they encounter that information, they may be able to enhance attention during encoding to selectively boost memory. This process likely requires top-down executive control instantiated in the prefrontal cortex (PFC), and as such, may become more effective through late adolescence and early adulthood as the PFC continues to develop.[6] To test whether individuals use this strategy, I will recruit 60 participants between the ages of

16 and 25 and examine their ability to complete a goal-directed memory task. For the purpose of participant engagement, I will turn this task into a narrative-based game that requires the same cognitive processes that likely underlie strategic encoding across real-world environments. In the game, participants' overarching goal is to earn as many points as possible by feeding different animals with the correct foods. In the first task, *the frequency task,* participants will see pictures of 30 different *types* of animals in a zoo. The animals will vary in the frequency with which they are presented — some will be presented once, some 3 times, and some 6 times, such that participants will view 100 animals total. I will measure participants' ability to track the frequency of animals by asking them to press a button in response to repeated images and examining how their accuracy increases and reaction times decrease across repetitions. Next, participants will be told that they will have to feed all 100 animals in the zoo, earning one point for each animal they feed correctly. To feed the animals correctly, participants will attempt to learn the type of food each type of animal likes in the *paired-associates* task, in which they will briefly view a picture of each of the 30 types animals presented next to a picture of a different food. If participants remember the food associated with an animal that appeared 6 times, they will have the potential to earn 6 points. If they remember the pairing for an animal that appeared once, they will only have the potential to earn one point. Critically, regardless of how many of each type of animal live in the zoo, participants will only learn what type of food that animal likes *once*. In the final *memory test*, participants will have to select which food to feed each of the 100 animals in the zoo. My main analyses will examine how memory performance differs as a function of the frequency with which each animal was presented. Additionally, I will test for differences in frequency-modulated memory as a function of participant age, cognitive ability (as measured by an IQ test), and their ability to learn the structure of their environment. **I hypothesize that with increasing age and cognitive ability, individuals will demonstrate greater capacity to deploy top-down executive control to apply their knowledge of their environment to preferentially encode the more useful information.**

Experiment 2. The second route through which environmental structure may enhance memory for more useful information is via *reactivation* of memories associated with more frequent stimuli. To test this possibility, a new group of 60 participants will complete the same three tasks as in the first experiment, but in a different order. In this version, participants will first complete the *paired-associates* task and *then* complete the *frequency task* and the *memory test.* In other words, participants will first learn to-be-remembered information, and *then* they will learn about the structure of their environment, which will determine the relative importance of the paired associates they previously learned. Rather than modulating memory *encoding* based on the inferred importance of information, in this task ordering, participants may use the structure of their environment to preferentially *reactivate* learned information. Each time participants view an animal in the frequency task, they may bring to mind associated information. Thus, the information associated with more frequent animals may be reactivated *more often*, strengthening its representation in memory.[7] Previous work suggests that this type of reactivation is instantiated by neural activity in the hippocampus and ventral pathway of the visual stream.[8] As such, it may *not* rely on prefrontal executive control mechanisms and may better support learning in individuals across the entire age and IQ range. **I hypothesize that individuals, regardless of age and cognitive ability, will reactivate information associated with cues in their environment and preferentially strengthen memories for more useful information.**

Aim 2. Investigate the neural mechanisms that support context-driven memory prioritization. Two new groups of 60 participants each will complete modified versions of the experiments described above while undergoing functional magnetic resonance imaging (fMRI). With fMRI, I will probe the precise neural mechanisms that may underlie individual differences in people's ability to use frequency information to modulate memory. To ensure that the two categories of stimuli are distinguishable in the brain, I will use images of animal *faces* and *buildings* that the animals need to be brought to, rather than foods that they want to eat. Animal faces have been previously shown to be represented in the fusiform face area (FFA), whereas images of buildings have been shown to activate the parahippocampal place area (PPA).[9]

Experiment 1. As before, participants will first learn the *frequency* of items in their environments and then the *paired associates* that will be probed. Here, I am interested in the neural mechanisms underlying both people's ability to learn the statistics of their environment and to *use* the structure of their environment to modulate memory encoding. I expect frequency learning to be reflected in a *decrease* neural activity in the hippocampus and the FFA when repeated images are presented.[10,11] This *decrease* in neural activity would reflect *repetition suppression*, a widely observed phenomenon in which individuals exhibit *reduced* neural responses to stimuli that they have previously seen.

The extent to which learning in the *frequency task* affects memory performance will depend upon individuals' ability to subsequently *use* these learned statistics to modulate encoding processes in the subsequent *paired-associates task.* I hypothesize that the effect of *repetition suppression* on frequency-based modulations of memory will be mediated by the extent to which participants are able to exert top-down executive control over memory encoding. To test this hypothesis, I will conduct a psychophysiological interaction (PPI) analysis to examine differences in correlated co-activation, often referred to as functional connectivity, of the dorsolateral prefrontal cortex (dlPFC), which has been implicated in executive control over relational memory,[12] and the hippocampus, which plays a critical role in encoding information, for high vs. low-value animal-building pairs. **I hypothesize that participants will better at encoding high-value information to the extent that they a.) demonstrate sensitivity to the statistical structure of their environment,** as measured by repetition suppression in the hippocampus and mPFC during the frequency task **and b.) engage top-down control over memory encoding to implement an optimal policy based on these learned statistical regularities** as measured by increased functional connectivity between the dlPFC and hippocampus in the paired-associates task.

Experiment 2. In the second neuroimaging experiment, participants will learn the paired associates *prior* to learning which pairs are more useful in their environment. As such, I do not expect to see differential engagement of the dlPFC across pairs. Instead, I expect that when participants re-encounter stimuli in their environments, they will exhibit stimulus-evoked pattern completion and selectively reactivate the relevant paired associate. For example, if a participant learns that cows live in the grey shack, I expect that each time they see a cow in their environment, they will reinstate the neural representation of the shack in their PPA. In this way, despite learning each pair once, participants will selectively strengthen representations of the associates of more frequent stimuli, and *the extent to which participants reactivate the paired associate should predict subsequent memory*. To measure the extent to which participants reinstate the paired buildings when each animal is presented, I will first run a localizer task that will enable me to determine the relative activity of each voxel in the brain for animal faces vs. buildings. From these measurements, I will be able to define participant-specific "face" regions of interest (ROI) that I expect to overlap with the FFA and "building" ROIs that I expect to overlap with the PPA. Then, I will be able to examine whether there is increased activity in the "building" ROI during the *frequency task* relative to the average activity in those voxels during the "face" localizer task. Critically, in the *frequency task*, no buildings themselves will be presented on the screen — activity in the "building" ROI should reflect retrieval and reinstatement of the relevant paired associate. **I hypothesize that participants who show greater reinstatement of the paired associates during the frequency task will demonstrate a.) overall better memory on the final memory test and b.) larger frequency-based modulations in memory performance.**

Aim 3. Investigate the flexibility of learning mechanisms across contexts. My final experiment will probe the flexibility of these learning systems to test whether individuals can rapidly update representations of information value across different environments. Specifically, participants will learn the statistical structure of *two* environments — the frequency of animals in a zoo *and* the frequency of animals in a circus. *Then*, they will be told which set of animals they need to feed, and learn the food pairs for each animal. The value of encoding each pair will depend on what context the participant is in. For example, there might be 6 lions in the zoo but only 1 lion at the circus, meaning remembering the food the lion likes is much more valuable if the participant is in the zoo context. I will test whether participants are able to a.) keep track of the statistical structure of two environments, as measured by their accuracy in the frequency task, and b.) selectively control memory based on the *relevant* environment, as measured by their performance in the final memory test. **I hypothesize that as the context becomes more complex, age and general cognitive ability will influence memory for high-value information to a greater degree.**

Benefit to society. Prior research on learning has examined how explicit reward and value signals modulate memory performance across individuals, but little work to date has examined how people instantiate their own learning policies based on *inferring* the importance of information from their environments. Understanding whether and how people are able to do this is critical to understanding how learning operates in real-world settings in which the value of information dynamically changes across contexts. In particular, the studies described aim to determine what types of environments *require* executive control for optimal learning and what types of environments circumvent the need for such control. The studies described will inform potential coaching strategies to enhance these forms of learning in individuals across a broad age range by pinpointing what cognitive and neural mechanisms underlie these processes and where they might break down. Teaching people to pay attention to the statistical regularities of their environment, to practice proactively considering not just the immediate reward value but the long-term utility of information, and to use cues in their environment to remind themselves of information they have previously learned, may all be effective strategies for preparing people — or soldiers specifically — to enter unknown, high-stakes environments. In other words, providing people with the tools to direct their own learning would serve as an incredibly important complement to directly teaching people information itself.

Relevant qualifications and institutional support. I have spent the past five years researching the development of memory, attention, and executive control, and over that time have gained experience with experimental design and the analysis of both behavioral and neuroimaging data. With the support of the NDSEG, I will be able to conduct this novel line of work under the mentorship of Dr. Cate Hartley, who has expertise in the development of the cognitive and neural systems that support flexible learning and decision-making. Additionally, I will benefit from NYU's strong quantitative training program (through which I have already completed two advanced math courses), its Center for Brain Imaging, and myriad faculty with expertise in neuroimaging and statistical methods, which will equip me with the tools I need to carry out my planned research.

1. Patil et al. (2016). *Learn Mem*. 24(1):65-69. 2. Dunsmoor et al. (2015). *Nature*. 520(7547):345-348. 3. Castel et al. (2011). *Developmental Psychology*. 47(6):1553-1564. 4. Staus & Falk (2017). *Mind, Brain, and Education*. 11(2):45-53. 5. Anderson (1991). *Behav Brain Sci*. 14(3):471-485. 6. Gogtay et al. (2004) *PNAS*. 101(21):8174-8179. 7. Smith et al. (2013) *JEP: General*. 142(4):1298-1308. 8. Skinner et al. (2010). *Neuropsychologia*. 48(1):156-164. 9. Spiridon et al. (2002). *Neuron*. 35(6):1157-1165. 10. Barron et al. (2013). *Nature*. 16(10):1492-1498. 11. Pajani et al. (2017). *Sci Rep*. 7(1):160. 12. Murray et al. (2007). *Journal of Neuroscience*. 27(20):5515-5522.