



The attention schema theory in a neural network agent: Controlling visuospatial attention using a descriptive model of attention

Andrew I. Wilterson^a and Michael S. A. Graziano^{a,b,1}

^aDepartment of Psychology, Princeton University, Princeton, NJ 08544; and ^bPrinceton Neuroscience Institute, Princeton University, Princeton, NJ 08544

Edited by Gregg H. Recanzone, University of California, Davis, CA, and accepted by Editorial Board Member Michael S. Gazzaniga July 7, 2021 (received for review February 5, 2021)

In the attention schema theory (AST), the brain constructs a model of attention, the attention schema, to aid in the endogenous control of attention. Growing behavioral evidence appears to support the presence of a model of attention. However, a central question remains: does a controller of attention actually benefit by having access to an attention schema? We constructed an artificial deep Q-learning neural network agent that was trained to control a simple form of visuospatial attention, tracking a stimulus with an attention spotlight in order to solve a catch task. The agent was tested with and without access to an attention schema. In both conditions, the agent received sufficient information such that it should, theoretically, be able to learn the task. We found that with an attention schema present, the agent learned to control its attention spotlight and learned the catch task. Once the agent learned, if the attention schema was then disabled, the agent's performance was greatly reduced. If the attention schema was removed before learning began, the agent was impaired at learning. The results show how the presence of even a simple attention schema can provide a profound benefit to a controller of attention. We interpret these results as supporting the central argument of AST: the brain contains an attention schema because of its practical benefit in the endogenous control of attention.

attention | internal model | machine learning | deep learning | awareness

For at least 100 y, the study of how the brain controls movement has been heavily influenced by the principle that the brain constructs a model of the body (1–10). Sometimes the model is conceptualized as a description (a representation of the shape and jointed structure of the body, including how it is currently positioned and moving). Sometimes it is conceptualized as a prediction engine (generating predictions about the body's likely states in the immediate future and how motor commands are likely to manifest as limb movements). Both the descriptive and predictive components are important and together form the brain's complex, multicomponent model of the body, sometimes called the body schema. The body schema is probably instantiated in a distributed manner across the entire somatosensory and motor system, including high-order somatosensory areas such as cortical area 5 and frontal areas such as premotor and primary motor cortex (2, 11–14). Without a correctly functioning body schema, the control of movement is virtually impossible. Even beyond movement control, the realization that any good controller requires a model of the thing it controls has become a general principle in engineering (15–17). Yet the importance of a model for good control has been mainly absent, for more than 100 y, from the study of attention—the study of how the brain controls its own focus of processing, directing it strategically among external stimuli and internal events.

Selective attention is most often studied as a phenomenon of the cerebral cortex (although it is not limited to the cortex). Sensory events, memories, thoughts, and other items are processed in the cortex, and among them, a select few win a competition for signal strength and dominate larger cortical networks (18–23). The process

allows the brain to strategically focus its limited resources, deeply processing a few items and coordinating complex responses to them, rather than superficially processing everything available. In that sense, attention is arguably one of the most fundamental tricks that nervous systems use to achieve intelligent behavior. Attention has most often been studied in the visual domain, such as in the case of spatial attention, in which visual stimuli within a spatial “spotlight” of attention are processed with an enhanced signal relative to noise (23, 24). That spotlight of attention is not necessarily always at the fovea, in central vision, but can shift around the visual field to enhance peripheral locations. Classically, that spotlight of attention can be drawn to a stimulus exogenously (such as by a sudden change that automatically attracts attention) or can be directed endogenously (such as when a person chooses to direct attention from item to item).

The attention schema theory (AST), first proposed in 2011 (25–28), posits that the brain controls its own attention partly by constructing a descriptive and predictive model of attention. The proposed “attention schema,” analogous to the body schema, is a constantly updating set of information that represents the basic functional properties of attention, represents its current state, and makes predictions such as how attention is likely to transition from state to state or to affect other cognitive processes. According to AST, this model of attention evolved because it provides a robust advantage to the endogenous control of attention, and in cases in which the model of attention is disrupted or makes errors, then the endogenous control of attention should be impaired (28–30).

Significance

Attention, the deep processing of select items, is one of the most important cognitive operations in the brain. But how does the brain control its attention? One proposed part of the mechanism is that the brain builds a model, or attention schema, that helps monitor and predict the changing state of attention. Here, we show that an artificial neural network agent can be trained to control visual attention when it is given an attention schema, but its performance is greatly reduced when the schema is not available. We suggest that the brain may have evolved a model of attention because of the profound practical benefit for the control of attention.

Author contributions: A.I.W. and M.S.A.G. designed research; A.I.W. performed research; A.I.W. and M.S.A.G. analyzed data; and A.I.W. and M.S.A.G. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. G.H.R. is a guest editor invited by the Editorial Board.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹To whom correspondence may be addressed. Email: graziano@princeton.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2102421118/-DCSupplemental>.

Published August 12, 2021.

An especially direct way to test the central proposal of AST would be to build an artificial agent, provide it with a form of attention, attempt to train it to control its attention endogenously, and then test how much that control of attention does or does not benefit from the addition of a model of attention. Such a study would put to direct test the hypothesis that a controller of attention benefits from an attention schema. At least one previous study of an artificial agent suggested a benefit of an attention schema for the control of attention (31).

Here, we report on an artificial agent that learns to interact with a visual environment to solve a simple task (catching a falling ball with a paddle). The agent was inspired by Mnih et al. (32), who demonstrated that deep-learning neural networks could achieve human or even superhuman performance in classic video games. We chose a deep-learning approach because our goal was not to engineer every detail of a system that we knew would work but to allow a naive network to learn from experience and to determine whether providing it with an attention schema might make its learning and performance better or worse. The agent that we constructed uses deep Q-learning (32–34) on a simple neural network with three connected layers of 200 neurons each. It is equipped with a spotlight of visual attention, analogous to (though much simpler than) human visual spatial attention. Many artificial attention agents have been constructed in the past, including some that use deep-learning networks (31, 35–41). Here, we constructed an agent that we hoped would be as simple as possible while capturing certain basic properties of human visual spatial attention. These properties include a moving spotlight within which the visual signal is enhanced relative to background noise, the initial exogenous attraction of the spotlight of attention to the onset of a visual stimulus, and the endogenous control of the spotlight to maintain it on the stimulus. The agent was introduced to its task with no a priori information (the neural network was initialized with random weights) and was given only pixels and reward signals as input. Using this artificial agent, we tested the hypothesis, central to AST, that the endogenous control of attention is greatly enhanced when the agent has access to a model of attention.

AST also proposes a specific relationship between the attention schema and subjective awareness, at least in the human brain. In the final section of this article, we discuss the possible relationship between the artificial agent reported here and subjective awareness.

Methods

Deep Q-learning is a method of applying reinforcement learning to artificial neural networks that have hidden layers. It was developed by researchers at Google DeepMind (32) who found that such networks could effectively learn to play video games at human and even superhuman performance levels. It is now widely used, and a toolbox, the TF-Agents library for Python, is available (33, 34). Because of its standardization and public availability, we do not provide a detailed description of the neural network methods here. Instead, we provide only the choices specific to our agent that, in combination with the resources cited here, will allow an expert to replicate our study. In the following sections, we first describe the artificial environment (inputs and action space) within which the agent operates. Then, we describe the game that the agent plays within that environment. Then, we describe the specific parameter choices that define the internal details of the agent itself. Finally, we describe the steps by which the agent learns.

The Agent's Environment and Action Space. Fig. 1 shows the environment of the agent, including two main components: a visual field and an attention schema. The top panel shows the visual field: a 10×10 array of pixels where black indicates a pixel that is “on” and gray indicates a pixel that is “off.” When a “ball” is presented in the visual field, its location is indicated by a black pixel. If the agent had perfect, noise-free vision, the ball would display as an isolated black pixel. Over a series of timesteps, a sequence of pixels would turn on across the display, tracking the motion of the ball. However, we gave the agent noisy vision. Where the ball is present, the pixel is black, and where the ball is not present, each pixel, on each timestep, has 50% probability of being black or gray, thus rendering the array a poor source of visual information. In

one region of the array, in a 3×3 square of pixels that we termed the “attention spotlight” (indicated by the red outline in Fig. 1), all noise was removed and the agent was given perfect vision. Within that attention spotlight, if no ball is present, all nine pixels are gray, and if the ball is present, the corresponding pixel is black while the remaining eight pixels are gray. In effect, outside the attention spotlight, vision is so contaminated by noise that only marginal, statistical information about ball position is available, whereas inside the attention spotlight, the signal-to-noise ratio is perfect and information about ball position can be conveyed. In this manner, the attention spotlight mimics, in simple form, human spatial attention, a movable spotlight of enhanced signal-to-noise within the larger visual field.

In Fig. 1, within the upper 10×10 visual array, the bottom row of pixels contains no visual noise. Instead, it doubles as a representation of the horizontal space in which the paddle (represented by three adjacent black pixels) can move to the left or right. The paddle representation is, in some ways, a simple analogy to the descriptive component of the body schema. It provides an updating representation of the state of the agent's movable limb.

The bottom field in Fig. 1 shows a 10×10 array of pixels that represents the state of attention. The 3×3 square of black pixels on the otherwise gray array represents the current location of the attention spotlight in the visual field. This array serves as the agent's simple attention schema. Note that the attention schema does not represent the object that is being attended (the ball). It does not convey the message “attention is on the ball” and provides no information about the ball. Instead, it represents the state of attention itself. As the attention spotlight moves around the upper visual array, the representation of attention moves correspondingly around the lower attention schema array. Note also that it is only a descriptive model of attention. It represents the current spatial location of attention (and, with the memory buffer described in *The Agent*, recent past states of attention as well) but does not represent predictive relationships between the state of attention on adjacent time steps or predictive relationships between attention and other aspects of the game. It contains no information indicative of the improved signal-to-noise that attention confers on vision—i.e., a prediction about how attention affects task performance. Any predictive component of an attention schema must be learned by the agent and would be represented implicitly in the agent's neural network.

During each game that the agent plays, 10 timesteps occur. On each timestep, the agent selects actions (movements of the attention spotlight and of the paddle) that are implemented on the next timestep. The agent can move its attention spotlight in any of eight possible ways: one step to the left, right, up, or down or one step in each diagonal (e.g., up-and-right motion entails moving one step up and one step right). If the edge of the spotlight touches the edge of the visual array, further movement in that direction is not allowed. For example, when the spotlight touches the top edge of the array, if the agent then selects an upward movement, the spotlight does not move, and if the agent selects an up-and-to-the-right movement, the spotlight moves only to the right.

The agent can also move its paddle in either of two directions: one step to the left or one step to the right. If the paddle touches the edge of the available space, no further movement in that direction is possible. For example, if the paddle touches the left edge of the array, and the agent then chooses a leftward movement, the paddle will remain at the same location for the subsequent timestep.

The Game and the Reward. In this section we describe the game, and in the following sections we describe our choices for constructing the agent itself such that it can learn to play the game.

The agent is trained on a game of catch, illustrated in Fig. 2. At the beginning of each game, a ball appears at the top of the 10×10 visual array, moves downward one row of pixels on each timestep, and reaches the bottom row of the visual array in 10 timesteps. The trajectory is not fully predictable. For each game, the ball's start position (at one of 10 possible locations in the top row) and end position (at one of 10 possible locations in the bottom row) are chosen randomly. Thus, the start position is not predictive of the end position. To pass from start to end, the ball moves on a diagonal trajectory until it is directly vertically above its end point. In the remaining timesteps, the ball moves vertically downward to reach its final position. The effect of the bent trajectory is to introduce unpredictability. When the ball is first presented in timestep 1, though its initial position is now revealed, its direction is unpredictable. When the ball begins to move, though its direction is now revealed, the timestep at which the trajectory changes from diagonal to vertical, and thus the ball's final position, is unpredictable. This unpredictability prevents the agent from learning to solve the game in a trivial manner by relying on information gleaned in the first few timesteps.

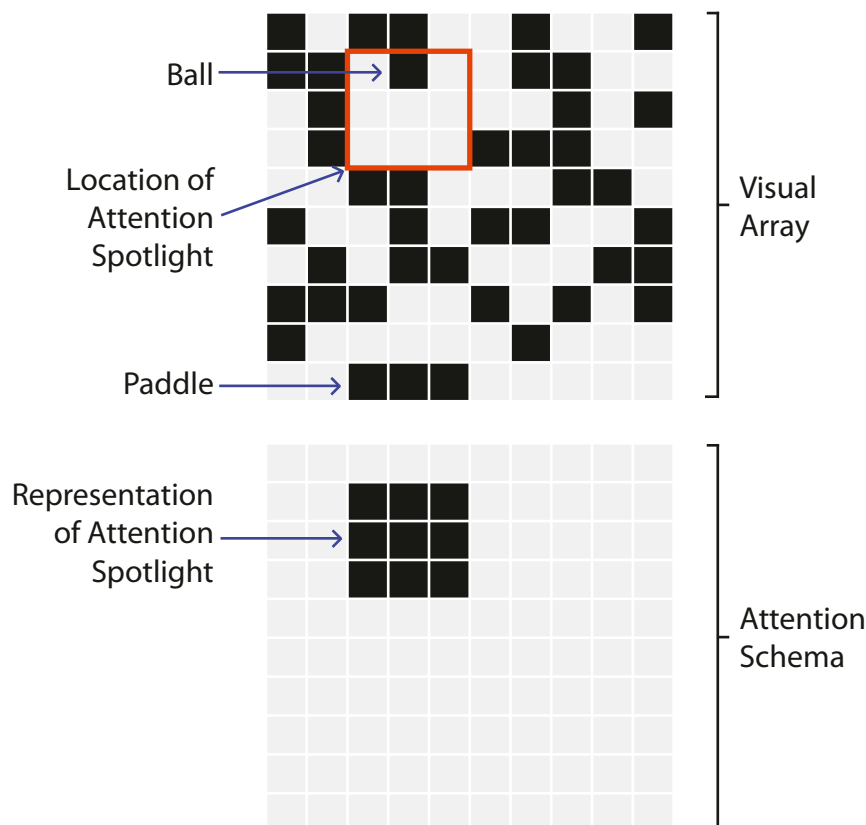


Fig. 1. Diagram of the agent's environment. (*Top array*) A 10×10 pixel visual field. Black pixels are "on," gray pixels are "off." The visual field contains noise (each pixel with 50% chance of turning black or gray on each timestep). The 3×3 pixel attention spotlight (orange outline) is shown here centered at visual field coordinates $X = 4$, $Y = 3$. Within the attention spotlight, all visual noise is removed. In this example, the ball is located within the spotlight, at visual coordinates $X = 4$, $Y = 2$. The lowest row of the visual array contains no visual noise and represents the space in which the paddle (three adjacent black pixels) moves. (*Bottom array*) A 10×10 pixel attention schema. The 3×3 area of black pixels represents the current state of the agent's attention spotlight.

To succeed at the game, the agent must learn to move its attention spotlight to track the ball's trajectory since reliable visual information is not present outside the spotlight. At the start of each game, the spotlight is initialized to a location at the top of the visual array, overlapping the ball. If the ball is initialized in an upper corner (left or right), the attention spotlight is initialized in the same upper corner (left or right). If the ball is initialized in a middle pixel of the upper array, the attention spotlight is initialized at the upper edge such that it is vertically centered on the ball. Thus, on the first timestep, when the ball appears, the agent's attention is automatically placed on the ball.

In analogy to the human case, one could say that the attention spotlight is exogenously attracted to the ball, similar to when the sudden appearance of a stimulus captures a person's spatial attention in an automatic, bottom-up manner. After the first timestep, the agent must learn to engage the equivalent of endogenous attention; it must learn to control its attention internally, keeping it aligned with the ball. Thus, the game involves an initial exogenous movement of attention to a stimulus and a subsequent endogenous control of attention with respect to that stimulus. The agent is rewarded by gaining 0.5 points on each timestep that the attention spotlight includes the ball and is punished by losing 0.5 points on each timestep that the attention spotlight does not include the ball. The positive and negative rewards that the agent can earn are balanced for better learning (32).

At the start of each game, the paddle is initialized to a position near the center of the paddle array, with three blank spaces to the left and four blank spaces to the right. By timestep 10, when the ball reaches its final position, if the paddle is aligned with the ball (if the ball location overlaps any of the three pixels that define the paddle), then the agent is rewarded by gaining 2 points for "catching" the ball. If the paddle is not aligned with the ball, then the agent is punished by losing 2 points. The largest number of points possible for the agent to earn in one game is therefore 7, including 5 for perfect attentional tracking of the ball over 10 timesteps and 2 for catching the ball with the paddle.

What the Agent Must Learn. As noted previously, the agent receives rewards for accomplishing two goals. The first goal is to track the ball with the attention spotlight. The second is to catch the ball with the paddle. If the agent fails at the first goal, the second goal becomes more difficult because, outside the attention spotlight, visual information about the ball is masked by noise.

To accomplish these goals, the agent presumably must learn basic properties of the environment. First, although Fig. 1 displays the pixels in a spatial arrangement, with a top and bottom, right and left, adjacency relationships between pixels, and neat division between two different functional fields, the agent begins with no such information. Pixels in its input array simply turn on and off. It has no a priori information that pixels can have meaningful correlational relationships, that the visual array has a spatial organization, or that the visual array is distinct from the attention schema array. It must learn every aspect of the environment needed to solve the task.

Second, the agent must learn, in essence, that it has a body part, a paddle, under its control. It must learn to move the paddle to the left or right in a manner correlated with the trajectory of the ball to earn rewards.

Third, to perform the task, the agent must learn to move its attention spotlight to track the ball along a partially unpredictable trajectory.

The Agent. We used a deep Q-learning agent implemented through the TF-Agents library for Python (33), which allows for computationally efficient training of reinforcement learning models using TensorFlow (34). The agent was built on top of a standard Q-Network as provided in the TF-Agents library. The agent was defined as a deep Q-learning network using the Adam optimizer with a learning rate of 0.001. We employed an "epsilon-greedy" policy ($\epsilon = 0.2$), to ensure exploration of the action space. We chose a network composed of three fully connected layers of 200 neurons each.

Because of the nonpredictable behavior of the ball, the agent's task was non-Markov, meaning task performance depended on information about the trajectory of the ball, which could only be inferred by observing a sequence of timesteps (not all the information needed to make an appropriate response was included in each timestep). Therefore, we equipped the model




Fig. 2. Trajectory of the ball across the visual field in a typical game. The start position (at one of 10 possible locations in the top row) and end position (at one of 10 possible locations in the bottom row) are chosen randomly. The start position is therefore not predictive of the end position. From the start position, on each timestep, the ball moves to an adjacent pixel on a diagonal trajectory until directly vertically above its end point. In the remaining timesteps, the ball moves vertically downward to reach its final position.

with a short-term memory. At the beginning of each game, the agent had an empty 200×10 pixel short-term memory buffer (corresponding to the 20×10 input array shown in Fig. 1 over 10 timesteps). On each timestep, the state of the environment was placed into the short-term memory buffer (with each observation containing 20×10 pixels of information). Subsequent timesteps in the same game were appended to previously observed timesteps, such that, by the end of the game, the short-term memory buffer contained every observation made during the game. Since the short-term buffer contained information not just from the current timestep, but from previous timesteps, the trajectory of the ball could be inferred and the task was rendered solvable. Note that the agent operated within an environment of two 10×10 input arrays (the visual array and the attention schema array, as shown in Fig. 1); but because the agent was given a 10-timestep memory, it made decisions on the basis of information from ten times as many inputs.

The Training Procedure. The agent was trained over the course of 1,500 iterations. In each iteration, the agent played 100 simulated games of catch, storing the results in a replay buffer. After completing its 100 plays, the agent was trained on the contents of the replay buffer in batches of 100 pairs of temporally adjacent timesteps, following standard training procedure for deep Q-learning (32–34).

After the 100 training games, the performance of the network was measured on 200 test games, to collect data on the progress of learning. The results of the test games were not stored in the replay buffer and therefore did not contribute to training. Thus, the test games sampled the model's ability to perform without contaminating the ongoing learning process. After the 200 test games, the ability to learn was reinstated for another iteration of 100 training games. This sequence (100 training games followed by 200 test games) was repeated until all 1,500 iterations of training were complete. Thus, in total, the agent was trained on 150,000 games and tested on 300,000 games.

After training, the agent was reinitialized and run through the training procedure five times, to generate data from five different artificial subjects. Since the initial synaptic weights were random, each subject was initialized differently, and therefore testing five subjects provided a way to check the reliability of the learning results. Some researchers use only one artificial subject, because of high reliability in the results (32), but testing several subjects is a useful, common way to assess the variance in the result (39), including testing five artificial subjects (41). Unlike human subjects, for which large numbers are required due to large statistical variation, artificial

agents can have more reliability. As described in subsequent sections, the small SE among the five artificial subjects shows that the present results have good reliability.

Results

Experiment 1: Baseline Performance. The purpose of experiment 1 was to determine whether the agent was capable of learning the task, and, if so, whether the agent was leveraging its spotlight of visual attention in order to do so. We therefore trained the agent with a functioning attention window, and then tested it both with and without a functioning attention window. As noted in the previous section, the agent was trained on 1,500 iterations. After each training iteration, the agent played 200 test games. The test games were divided into four categories of 50 games, each category testing a different aspect of the agent's performance. In the following, we describe each of the four tests. See Table 1 for a summary of training and testing conditions.

First, we tested the agent on its ability to track the ball with its attention spotlight. In this category of test games, the agent received points only for attention tracking, not for ball catching. Note, again, that these test games did not contribute to the agent's learning; the agent was never trained on this reduced reward contingency; learning only occurred during training games for which the full reward structure was implemented. Instead, the points earned by the agent in this category of test game served solely as a measure of how well the agent performed at tracking the ball with its attention. The agent received 0.5 points for successfully keeping the ball within the spotlight on each timestep, and lost 0.5 points for failing to do so. It could earn a maximum of 5 points across all 10 timesteps on a perfectly played game. Negative scores indicate that the agent lost more points than it gained during a game.

The blue line in Fig. 3A shows the learning curve for this first test condition (Attention Score, Normal Agent). At the start of the 1,500 training iterations, the agent scored below zero, indicating that it lost more points than it gained (it made more errors than it made correct choices). Note that chance performance is, by definition, the performance of the agent when its network weights are randomized and its actions are random, at the start of learning. During training, the agent increased the points it was able to earn for attention tracking, passing zero (making more correct choices than errors), and approaching a performance of ~ 3.5 points. The curve shows the mean and SE among the five artificial subjects (see *SI Appendix* for curve fits and quantification of learning rates). Note that the performance at the end (by 1,500 iterations) was far above the chance performance levels at the start of learning, and that the difference is well outside the extremely small SE.

In a study with highly variable human behavioral data, one would typically perform statistical tests to determine if the observed differences between means can be assigned a P value below the benchmark of 0.05. With the highly reliable performance of artificial agents, such statistical tests are often not relevant (32). Here, the difference between initial, chance level of performance and final, trained level is so far outside the measured SE that the corresponding P is vanishingly small ($P < 1 \times 10^{-10}$). Because the effects described here are so much larger than the measured SE, we will not present the statistical tests typical of human behavioral experiments but will instead adhere to the convention of showing the SE.

The orange line in Fig. 3A shows the results for the second category of test games (Attention Score, Attention-Disabled). For these test games, the agent's attention spotlight was altered. We adjusted the rules such that visual noise was applied both inside and outside the attention spotlight. All other aspects of the agent remained the same. The agent still had an attention spotlight that could move about the visual field, and a functioning attention schema array that signaled the location of the

Table 1. Training and testing configurations for all three experiments

		Training	Test 1	Test 2	Test 3	Test 4
Experiment 1	Attention window	Intact	Intact	Disabled	Intact	Disabled
	Attention schema	Intact	Intact	Intact	Intact	Intact
Experiment 2	Attention window	Intact	Intact	Intact	Intact	Intact
	Attention schema	Intact	Intact	Disabled	Intact	Disabled
Experiment 3	Attention window	Intact	Intact	Disabled	Intact	Disabled
	Attention schema	Disabled	Disabled	Disabled	Disabled	Disabled

For each experiment, the attention window and the attention schema could be intact or disabled during the training trials and during the testing trials. The testing trials were grouped into four tests. Test 1 and 2 involved measuring the agent's performance on the attention tracking component of the task. Tests 3 and 4 involved measuring the agent's performance on the ball-catching component of the task. For example, in experiment 1, the attention window was intact during training and during test 1 and test 3 and was disabled during test 2 and test 4.

attention spotlight. However, the attention spotlight itself was rendered inoperable, filled with the same level of visual noise as the surrounding visual space. The only change, therefore, was to the level of noise present within the spotlight. Note that the agent was still trained with functioning attention—only during the test games was the attention window disabled. In principle, an agent could perform the task with its attention spotlight disabled. The visual array, though filled with noise, does contain some visual information. In the region where the ball is located, the likelihood of a black pixel increases marginally above the background noise level of 50%. Likewise, the agent could in principle leverage other information. For example, on every timestep, for good performance, the agent should lower its attention window by one row. It should never raise its attention window. On these grounds, we expect the agent's performance to improve with training. As shown by the orange line in Fig. 3A, over the course of the 1,500 iterations, the agent's score did rise above the initial, chance level. It even rose above 0, indicating that the agent was able to gain more points than it lost on the attention tracking aspect of the task. However, performance was much worse without functioning attention than with functioning attention (orange line versus blue line in Fig. 3A). Note that the difference between the curves is much greater than the SE within each curve. These results demonstrate that the agent learned the task, and also that it used its attention window to help perform

the task, since when its attention was disabled, it could no longer perform as well.

The blue line in Fig. 3B shows the results for the third category of test games (Catch Score, Normal Agent). Here we tested the agent's ability to catch the ball with the paddle. Success earned 2 points per game and failure earned -2 points. The agent began by scoring below zero (failing more often than succeeding). Over the course of the 1,500 training iterations, the agent's score rose, approaching ~ 0.8 points (of the maximum possible 2 points). In comparison, the orange line in Fig. 3B shows the results for the final, fourth category of test games (Catch Score, Attention-Disabled). Here, we tested the agent's ability to catch the ball with the paddle while the attention spotlight was disabled (visual noise was applied both inside and outside the spotlight). The agent's performance on the catch task was worse without a functioning attention spotlight than with one (orange line versus blue line, Fig. 3B). The difference was, again, well outside the range of the SE.

An alternative perspective on the efficacy of the attention window may be possible. Consider Fig. 3A. The data show improved performance in the blue line (representing the agent with all noise removed from within the attention window) over the orange line (representing the agent without a functioning attention window). Noise was created by randomly turning on 50% of the pixels across the visual array. On average, within the 9-pixel area subtended by the attention window, 4 to 5 pixels should be

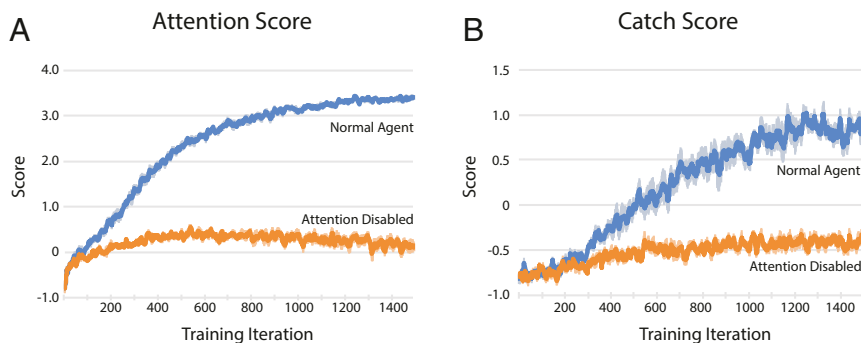


Fig. 3. Results of experiment 1. Training occurred over 1,500 iterations. Each iteration included 100 training games (during which the agent learned), followed by 200 test games to evaluate performance (during which the agent did not learn). The 200 test games were divided into four categories of 50 games each. Each data point shows the mean score among the 50 test games within a test category, and among 5 simulated subjects, spatially smoothed with a 10-iteration moving average. Lighter margin around darker line reflects SE of intersubject variability. (A) Performance on attention tracking task (how well the attention spotlight remained aligned on the ball through the game), for which the agent could earn a maximum of 5 points per game. The blue line denotes performance under normal conditions, identical to the conditions experienced in training. The orange line denotes performance under attention-disabled conditions, wherein the attention spotlight did not reduce visual noise in the subtended area. (B) Performance on ball-catching task (whether the paddle was aligned with the ball by the end of the game), for which the agent could earn a maximum of 2 points per game. The blue line denotes performance under normal conditions, identical to the conditions experienced in training. The orange line denotes performance under attention-disabled conditions, wherein the attention spotlight did not reduce visual noise in the subtended area.

turned on, but that noise was removed to make attention functional. Consider a manipulation in which, instead of removing those noise pixels from the attention window, we remove 5 noise pixels randomly from the entire visual array. With a marginal reduction in noise spread everywhere, instead of a total removal of noise within the restricted attention window, would the agent perform as well as in the blue line in Fig. 3A? The answer is no. In a separate test, we found that under these conditions the agent performed similarly to the orange line in Fig. 3A, with a peak score of about 0.5. Thus, reducing the background noise by an incremental amount everywhere does not provide an equivalent performance advantage as does clearing the noise within a controllable spotlight of attention.

The results of experiment 1 confirm that the agent can improve on the game through training, and that it does so partly by relying on the visual information obtainable in its high signal-to-noise attention spotlight, since removing the functionality of the spotlight causes a substantial reduction in performance. These results enable us to proceed to the next experiment, testing how well the agent can control its attention spotlight when its attention schema fails.

Experiment 2: Necessity of the Attention Schema for Performance. In experiment 2, we performed similar tests as in experiment 1. However, instead of testing the effect of disabling the attention spotlight, we tested the effect of disabling the attention schema. To disable the attention schema, we fixed every pixel in the attention schema array (lower array in Fig. 1) to an inactive state, such that the array no longer provided any information. The agent was trained with a functioning attention schema (trained under the same conditions as in experiment 1) but was tested both with and without the attention schema. Once again, we used four categories of test game, enumerated in the following. See Table 1 for the training and testing conditions.

The blue line in Fig. 4A shows the results from the first test category (Attention Score, Normal Agent). In this category, we measured the points that the agent received for tracking the ball with its attention spotlight. This test is a replication of the first test of experiment 1 (blue line in Fig. 3A). Just as in experiment 1, at the start of the 1,500 training iterations, with randomized weights, the agent scored at a chance level that was below zero, indicating that it lost more points than it gained. Over training,

the agent showed a rising learning curve, approaching ~3.5 points (of the maximum possible 5 points).

In comparison, the orange line in Fig. 4A shows the results from the second test category (Attention Score, Attention Schema Disabled). For these test games, the agent still had an attention spotlight that afforded noise-free vision within it, and the agent could still select movements of that attention spotlight. However, the attention schema (lower array in Fig. 1) was turned off (filled entirely with gray pixels). No signal was presented within it. Note that the agent was still trained with a functioning attention schema, but its performance was probed with its attention schema temporarily turned off. Performance was much worse without a functioning attention schema than with one (orange line versus blue line in Fig. 4A).

The blue line in Fig. 4B shows the results for the third test category (Catch Score, Normal Agent). Here, we tested the agent's ability to catch the ball with the paddle. Replicating experiment 1, the agent showed a rising learning curve, approaching ~0.8 points (of the maximum possible 2 points). In comparison, the orange line in Fig. 4B shows the results for the fourth test category (Catch Score, Attention Schema Disabled). Here, we tested the agent's ability to catch the ball with the paddle, while, once again, the attention schema was disabled. Performance was worse without a functioning attention schema than with one (orange line versus blue line in Fig. 4B).

The results of experiment 2 confirm that when the agent is trained on the game with an attention schema available to it, it can learn the task. The results also show that the agent's performance relies on the attention schema, since deactivating the attention schema on test trials results in a greatly reduced performance.

Experiment 3: Necessity of the Attention Schema for Learning. In experiment 2, described previously, we showed that as the agent learns to control its attention, if its attention schema is temporarily disabled during a test game, the agent's performance is reduced. However, suppose the agent were trained, from the start, without an attention schema, such that it never learned to rely on that source of information. Would the agent be incapacitated, or would it learn to play the game through an alternate strategy?

In principle, one might think the agent could learn the tracking task and the catch task without an attention schema informing it about the state of its attention. The reason is that the visual array itself already contains information about the attention spotlight.

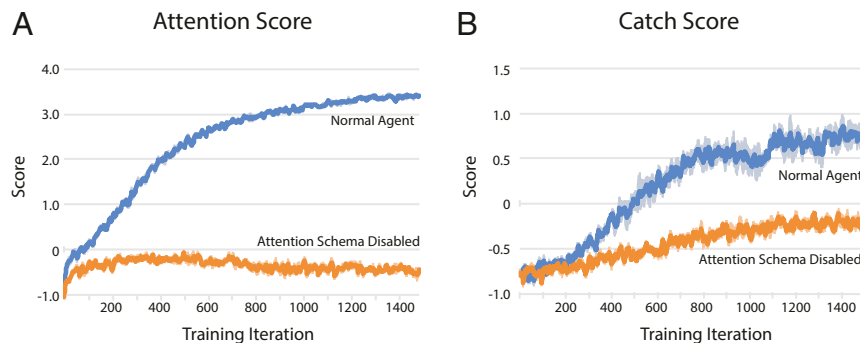


Fig. 4. Results of experiment 2. Training occurred over 1,500 iterations. Each iteration included 100 training games (during which the agent learned), followed by 200 test games to evaluate performance (during which the agent did not learn). The 200 test games were divided into four categories of 50 games each. Each data point shows the mean score among the 50 test games within a test category, and among 5 simulated subjects, spatially smoothed with a 10-iteration moving average. Lighter margin around darker line reflects SE of intersubject variability. (A) Performance on attention tracking task (how well the attention spotlight remained aligned on the ball through the game), for which the agent could earn a maximum of 5 points per game. The blue line denotes performance under normal conditions, identical to the conditions experienced in training. The orange line denotes performance under attention-schema-disabled conditions, wherein all pixels in the attention schema array were fixed to the off state. (B) Performance on ball-catching task (whether the paddle was aligned with the ball by the end of the game), for which the agent could earn a maximum of 2 points per game. The blue line denotes performance under normal conditions, identical to the conditions experienced in training. The orange line denotes performance under attention-schema-disabled conditions, wherein all pixels in the attention schema array were fixed to the off state.

Consider the visual array in Fig. 1 (upper 10×10 array). The array is filled with noise except in one 3×3 window, revealing, at least statistically, where the attention spotlight is likely to be. All necessary information to know the state of its attention and to perform the task is therefore already present in the visual array. Thus, it is not a priori obvious that the added attention schema array is necessary for the agent to learn the task. Moreover, the added attention schema could, in principle, even harm performance, because adding extra inputs to a limited neural network might slow the rate of learning. Thus, it is an empirical question whether supplying the agent with an added attention schema will help, harm, or make no significant difference to learning. The purpose of experiment 3 was to answer that question. It is the central experiment of the study. When the agent must learn the task without an attention schema present from the outset, is learning helped, harmed, or unaffected, as compared to learning with an attention schema present?

In experiment 3, in all games, whether training games or test games, the attention schema was disabled. To disable it, we turned off all the pixels in the attention schema array. The agent therefore never received any information from the attention schema array about the state of its attention. The blue line in Fig. 5A shows the results from the first test category (Attention Score, Attention Normal). This test measured the points received for tracking the ball with the attention spotlight. The agent had a normal attention spotlight (no visual noise within the spotlight), and retained the ability to move its attention spotlight, but lacked a functioning attention schema. Over the 1,500 training iterations, the agent showed some degree of learning, improving from its initial negative score to a final level of about 0.4 points. Compared, however, to the agent in experiment 1 that was trained with a functioning attention schema (Fig. 3A, blue line), the agent in experiment 3 (Fig. 5A, blue line) was severely impaired. The difference in performance is large compared to the SE.

The orange line in Fig. 5A shows the results from the second category of test games (Attention Score, Attention-Disabled). Here, not only was the attention schema disabled, but the attention spotlight itself was also disabled. All pixels in the visual array, whether in or out of the spotlight, were subject to the same visual noise. Again, the agent showed some learning, improving over its initial negative score. Comparing the orange and blue lines in Fig. 5A shows that the agent must have relied partly on the

visual information available in its attention spotlight, since without that spotlight, it performed consistently worse. However, performance was poor in both conditions compared to the performance of the fully intact agent. Without its attention schema, the agent was impaired.

Fig. 5B shows the results from the third and fourth category of test games. Here we measured the points that the agent received for catching the ball with its paddle when its attention schema was disabled. The blue line shows the results when the attention spotlight was normal, and the orange line shows the results when the attention spotlight was disabled. Again, comparing the orange and blue lines shows that the agent must have relied partly on the visual information available in its attention spotlight, since without that spotlight, it performed worse. However, the key finding here is that performance was poor in both conditions compared to the performance of the fully intact agent.

Experiment 3 shows that when the agent is trained from the outset without an attention schema, it does not learn well compared to its performance when trained with an attention schema. Perhaps because the agent learns to use the available information in the visual input array, performance does improve over the 1,500 training iterations. The results show that some of the agent's minimal improvement in performance must depend on the proper functioning of its attention spotlight, since learning was reduced when the spotlight was disabled. However, in all conditions tested in experiment 3, the agent was severely impaired at learning the task—either tracking the ball with its attention spotlight or catching the ball at the end of each game—as compared to the intact agent. Without an attention schema, the agent was compromised.

Discussion

The Usefulness of an Attention Schema. The present study demonstrates the advantage of adding a model of attention to an attention controller. An artificial agent, with a simple version of a moving spotlight of visual attention, benefitted from having an updating representation of its attention. The difference was drastic. With an attention schema, the agent learned to perform. Without an attention schema, the machine was comparatively incapacitated. When the machine was trained to control its attention while it had an attention schema, and then tested while its attention schema was disabled (experiment 2), performance was

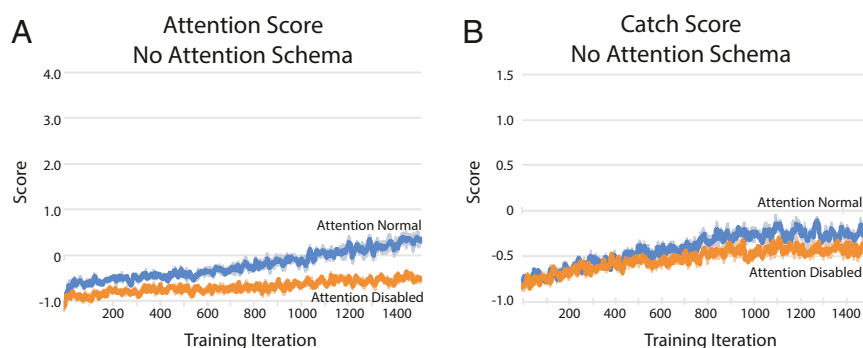


Fig. 5. Results of experiment 3. All games, both training and test games, were run under attention-schema-disabled conditions, wherein all pixels in the attention schema array were fixed to the off state. Training occurred over 1,500 iterations. Each iteration included 100 training games (during which the agent learned), followed by 200 test games to evaluate performance (during which the agent did not learn). The 200 test games were divided into four categories of 50 games each. Each data point shows the mean score among the 50 test games within a test category, and among 5 simulated subjects, spatially smoothed with a 10-iteration moving average. Lighter margin around darker line reflects SE of intersubject variability. (A) Performance on attention tracking task (how well the attention spotlight remained aligned on the ball through the game), for which the agent could earn a maximum of 5 points per game. The blue line denotes performance under conditions in which the agent had a functioning attention spotlight. The orange line denotes performance under attention-disabled conditions, wherein the attention spotlight did not reduce visual noise in the subtended area. (B) Performance on ball-catching task (whether the paddle was aligned with the ball by the end of the game), for which the agent could earn a maximum of 2 points per game. The blue line denotes performance under conditions in which the agent had a functioning attention spotlight. The orange line denotes performance under attention-disabled conditions, wherein the attention spotlight did not reduce visual noise in the subtended area.

drastically reduced. When the machine was trained from the outset without an attention schema available (experiment 3), learning was drastically reduced. These findings lend computational support to the central proposal of AST: the endogenous control of attention is aided by an internal model of attention (25–30). We argue that evolution arrived at that solution in the brain because it provides a drastic benefit to the control system.

One possible criticism of the present findings might be that, when the artificial agent was robbed of its attention schema, it did not have sufficient information to solve the game, and therefore its failure to perform was inevitable. Thus, the results are no surprise. This argument is not correct. In principle, an agent could solve the game without being given an attention schema. The reason is that complete information about the location of the attention window is present, though implicit, in the visual array. The patch of clear pixels in an otherwise noisy array provides information about where the attention spotlight is located. Therefore, in principle, the agent is being supplied with the relevant information. Moreover, adding the extra 100 input signals needed for an attention schema could, arguably, harm learning and performance by adding more inputs for the limited neural network to learn about. Based on a priori arguments like these, one could hypothesize that the attention schema might help, might do little, or might even harm performance. The question is therefore an empirical one. The results confirmed the prediction of AST. The addition of an attention schema dramatically boosted the ability of the agent to control its attention.

When the agent was not given an attention schema, could it have learned to perform well with more training? It is always possible. We can only report on the 150,000 training games that we tested, and do not know what might have happened with much more training. Moreover, a more powerful agent, with more neurons or layers, might have been able to learn to control its attention without being given an explicit attention schema. We do not argue that because our particular agent was crippled at the task when it was given no attention schema, therefore all agents would necessarily be so as well. We do suggest, however, that adding an attention schema makes the control of attention drastically easier, and the lack of an attention schema is such a liability that it greatly reduced the performance of our simple network.

Descriptive versus Predictive Models of Attention. When the agent contained all its components (a functioning attention spotlight and an attention schema), in order to perform, it must have learned some of the conditional relationships between the current state of its attention, past and future states, its own action options, the trajectory of the ball, and the possibility of reward. Therefore, arguably, the agent's attention schema had two broad components: a simple descriptive component (depicting the current state of attention) that was supplied to it in the form of the attention schema array, and a more complex, subtle, and predictive component that was learned and encoded in an implicit manner in the agent's synaptic weights. Without the descriptive component, the agent could not effectively learn the predictive component. Essentially, the present study shows the usefulness of giving the agent an explicit, descriptive attention schema as a foundation on which the agent can learn a predictive attention schema.

The type of attention that we built into the artificial agent is extremely simple: a square window that removes all visual noise within it, and that can move in single-increment steps. The descriptive attention schema that we supplied to the agent is, therefore, correspondingly simple: it informs the agent about the location of the attention window. Real human attention is obviously more complex (18–23). Human spatial attention is graded at the edges, the intensity can be actively enhanced or suppressed, the shape and size of the spotlight can change, it may be divisible into multiple areas of focus, and it can move around the visual

field in more complex ways than in stepwise increments. Beyond spatial attention, one can attend to features such as color or shape, to other sensory modalities such as audition or touch, or to internal signals such as thoughts, memories, decisions, or emotions. Moreover, human attention has predictable influences on far more than just a simple catch task; it influences memory, emotion, planning, decision-making, and action. Human attention is obviously of much higher dimensionality and complexity than a square spotlight that impacts one task. If humans have an attention schema as we propose, it must be correspondingly more complex, representing the great variety of ways that attention can vary and impact other events. We acknowledge the simplicity of the agent tested here, but we suggest that similar underlying principles may be at work in the artificial and biological cases.

The Relationship between the Attention Schema and Subjective Awareness. We previously proposed that the human attention schema has a relationship to subjective awareness (25–28). To understand how, first consider the analogy to the body schema. You can close your eyes, thus blocking vision of your arm, and still “know” about your arm, because higher cognition has access to the automatically constructed, lower-order body schema. Even people who lack an arm due to amputation can still have a lingering body-schema representation of it, which they report as a phantom limb (42). The ghostly phantom limb phenomenon illustrates a property of the body schema: the schema contains a detail-poor, simplified description of the body, a “shell” model, lacking any implementation details about bones, muscles, or tendons.

In the same manner, according to AST, we “know” about how our cortex is allocating resources to select items, because higher cognition has access to some of the information in the automatically constructed, lower-order attention schema. In AST, because the attention schema is a simplified model of attention, a kind of shell model that lacks the implementation details of neurons and synapses and pathways in the brain, it provides us with the unrealistic intuition that we have a nonphysical essence of experience, a feeling, a kind of mental vividness, that can shift around and seize on or take possession of items. In AST, the information in the attention schema is the source of our certainty that we have a nonphysical essence of experience.

Logically, everything that we think is true about ourselves, no matter how certain we are of it, must derive from information in the brain, or we would not be able to have the thought or make the claim. AST proposes that the attention schema is the specific set of information in the brain from which derives our common belief and claim to subjective awareness. In AST, when we say that we have a subjective awareness of an item, it is because the brain has modeled the state of attending to that item, the model is not perfectly physically detailed or accurate, and based on information contained in the model, higher cognition reports a physically incoherent state. In this approach, visual awareness of an apple, for example, depends on at least two models. First, a model of the apple constructed in the visual system supplies information about the specifics of shape and color. Second, a model of attention supplies the information underlying the additional claim that an ineffable, subjective experience is present.

If AST is correct—if subjective awareness in people indicates the presence of a model of attention—then the control of attention should depend on awareness. There ought to be situations in which the human attention schema makes a mistake—in which attention and awareness become dissociated—and in that case, the control of attention should break down. A growing set of behavioral data supports this prediction. Subjective awareness and objective attention usually covary, awareness encompassing the same items that attention encompasses. The two can become dissociated, however. When a person has no subjective awareness of a visual stimulus, attention can still be exogenously attracted to the stimulus, and the degree of attention is not necessarily

diminished (29, 30, 43–53). In that case, the person can no longer effectively control that focus of attention. Without subjective awareness of a stimulus, people are unable to strategically enhance or sustain attention on the stimulus, even if it is task-relevant (29); people cannot suppress attention on the stimulus, even if it is a distractor (54); and people cannot adjust attention to one or the other side of the stimulus, even if that controlled shift of attention is of benefit to the task at hand (30). Without awareness of a visual stimulus, the endogenous control of attention with respect to that stimulus is severely impaired. These findings lend support to the hypothesis that awareness is indicative of a functioning control model for attention. When the control model is compromised, the control of attention breaks down.

Many processes in the brain function outside of awareness. Blindsight patients have (very limited) visual ability, responding and even pointing to objects of which they are not subjectively aware (55). Patients suffering hemispatial neglect show high-level cognitive and emotional reactions to stimuli that they do not consciously see, such as recognizing that a flaming house is undesirable even though they are not aware of the flames (56). In healthy subjects, unconscious vision routinely influences decision-making and behavior, sometimes in striking ways (57). So many complex mental processes have been shown to occur without awareness of the relevant items, that sometimes it seems as if awareness must do nothing, or, at least, cognitive processes that require it must be the exception rather than the rule. The ability to endogenously control attention may be one of the few internal processes that are known to categorically break down without awareness of the relevant stimuli.

Suppose you are planning and executing the sequence of getting a knife out of a drawer, an apple out of the refrigerator, and peeling the apple. If you are not conscious of the drawer, knife, refrigerator, or apple, your exogenous attention might automatically flicker to one or another of those items, and you might even be able to reach for one of them (as blindsight patients do), but you will probably not coordinate or execute the whole plan. We suggest that creating a cognitive plan and executing it requires a strategic, sequential movement of attention from one item to the next. According to AST, that ability to control attention collapses without awareness of the items, and it is for this reason that our complex planning and behavior depends on awareness.

Is the Artificial Agent Conscious? Yes and No. If consciousness is linked to an attention schema, and if the artificial agent described here has an attention schema, then is the agent conscious? We suggest the answer is both yes and no.

Suppose one could extract and query all the information that the agent contains about attention. As noted previously, that information set has two components: a descriptive component (the lower array of inputs that was supplied to the agent, that describes the state of attention) and a predictive component (the information that the agent learns about the relationship between attention, vision, action, and reward—this predictive information is represented in an implicit manner in the connectivity within the neural network). Now suppose we ask the following: “There is a ball that moves. There is a paddle that moves. What else is present?” The machine presumably contains information about attention that, if extracted, systematized, and put into words, could be articulated as: “Another component is present. Like the ball and the paddle, it can move. It can be

directed. When it encompasses the ball, rewarded performance with respect to the ball can occur. When it does not encompass the ball, the ball is effectively absent from vision, and any rewarded performance is unlikely.”

Now consider human consciousness, which could be described something like this: “It is not the objects in the external world; it is not my physical body that interacts with the world. It is something else. It shifts and moves; it can be directed; it encompasses items. When it does so, its signature property is that those items become vivid, understood, and choices and actions toward those items are enabled.”

Given all of that discussion, does the agent have consciousness or not? Yes and no. It has a simple version of some of the information that, in humans, may lead to our belief that we have subjective consciousness. In that sense, one could say the agent has a simple form of consciousness, without cognitive complexity or the ability to verbally report. Our hope here is not to make an inflated claim that our artificial agent is conscious, but to deflate the mystique of consciousness to the status of information and cognitive operations.

Many other theories of consciousness have been proposed (e.g., refs. 58–70). We argue that most belong to a particular approach. They propose a condition under which the feeling of consciousness emerges. They posit a transition between the physical world of events in the brain and a nonphysical feeling. From this typical approach stems a common reluctance of scientists to study consciousness, perceiving it to be something outside the realm of the physical sciences. In contrast, AST is an engineering theory that operates only within the physical world. It proposes that the brain constructs a specific, functionally useful set of information, the attention schema, a rough depiction of attention, from which stems at least two properties: first, the brain gains the ability to effectively control attention; and second, we gain a distorted self-caricature. That caricature leads us to believe that we have a mysterious, magical, subjective consciousness.

Theories like AST are sometimes called “illusionist” theories, in which conscious experience is said to be an illusion (71, 72). While we agree with the concept behind that perspective, we argue that consciousness is probably better described as a caricature than an illusion. A caricature implies the presence of something real that is being caricatured. In AST, human subjective awareness is a useful caricature of attention. We cannot intellectually choose to turn the caricature on or off, because it is constructed by the brain automatically. Moreover, according to AST, the caricature is of essential importance; the model of attention helps to control attention. While not all aspects of AST have been tested yet, the present study supports at least one key part of the theory. A descriptive model of attention does appear to benefit the control of attention.

Data Availability. All data and code for this study are publicly available and have been deposited in Princeton DataSpace (<http://arks.princeton.edu/ark:/88435/dsp01kp78gk430>).

ACKNOWLEDGMENTS. Supported by the Princeton Neuroscience Institute Innovation Fund. We thank Isaac Christian for his thoughtful help and Argos Wilterson for his enthusiastic encouragement.

1. S. Gallagher, “Body schema and intentionality” in *The Body and The Self*, J. L. Bermudez, A. Marcel, N. Glan, Eds. (MIT Press, Cambridge, MA, 1988), pp. 225–244.
2. M. S. A. Graziano, M. M. Botvinick, “How the brain represents the body: Insights from neurophysiology and psychology” in *Common Mechanisms in Perception and Action: Attention and Performance XIX*, W. Prinz, B. Hommel, Eds. (Oxford University Press, Oxford, UK, 2002), pp. 136–157.
3. H. Head, *Studies in Neurology* (Oxford University Press, London, 1920), vol. 2.
4. H. Head, H. G. Holmes, Sensory disturbances from cerebral lesions. *Brain* **34**, 102–254 (1911).
5. N. P. Holmes, C. Spence, The body schema and the multisensory representation(s) of peripersonal space. *Cogn. Process.* **5**, 94–105 (2004).

6. M. Kawato, Internal models for motor control and trajectory planning. *Curr. Opin. Neurobiol.* **9**, 718–727 (1999).
7. M. J. Pickering, A. Clark, Getting ahead: Forward models and their place in cognitive architecture. *Trends Cogn. Sci.* **18**, 451–456 (2014).
8. R. Shadmehr, F. A. Mussa-Ivaldi, Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* **14**, 3208–3224 (1994).
9. J. A. Taylor, J. W. Krakauer, R. B. Ivry, Explicit and implicit contributions to learning in a sensorimotor adaptation task. *J. Neurosci.* **34**, 3023–3032 (2014).
10. D. M. Wolpert, Z. Ghahramani, M. I. Jordan, An internal model for sensorimotor integration. *Science* **269**, 1880–1882 (1995).

11. M. Desmurget, A. Sirigu, A parietal-premotor network for movement intention and motor awareness. *Trends Cogn. Sci.* **13**, 411–419 (2009).
12. M. S. A. Graziano, D. F. Cooke, C. S. R. Taylor, Coding the location of the arm by sight. *Science* **290**, 1782–1786 (2000).
13. P. Haggard, D. Wolpert, "Disorders of body schema" in *High-Order Motor Disorders: From Neuroanatomy and Neurobiology to Clinical Neurology*, J. Freund, M. Jeannerod, M. Hallett, R. Leiguarda, Eds. (Oxford University Press, Oxford, UK, 2005), pp. 261–271.
14. P. Morasso, M. Casadio, V. Mohan, F. Rea, J. Zenzeri, Revisiting the body-schema concept in the context of whole-body postural-focal dynamics. *Front. Hum. Neurosci.* **9**, 83. (2015).
15. E. F. Camacho, C. Bordonas Alba, *Model Predictive Control* (Springer, New York, 2004).
16. R. C. Conant, W. R. Ashby, Every good regulator of a system must be a model of that system. *Int. J. Syst. Sci.* **1**, 89–97 (1970).
17. B. A. Francis, W. M. Wonham, The internal model principle of control theory. *Automatica* **12**, 457–465 (1976).
18. D. M. Beck, S. Kastner, Top-down and bottom-up mechanisms in biasing competition in the human brain. *Vision Res.* **49**, 1154–1165 (2009).
19. M. M. Chun, J. D. Golomb, N. B. Turk-Browne, A taxonomy of external and internal attention. *Annu. Rev. Psychol.* **62**, 73–101 (2011).
20. M. Corbetta, G. L. Shulman, Control of goal-directed and stimulus-driven attention in the brain. *Nat. Rev. Neurosci.* **3**, 201–215 (2002).
21. R. Desimone, J. Duncan, Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* **18**, 193–222 (1995).
22. T. Moore, M. Zirnsak, Neural mechanisms of selective visual attention. *Annu. Rev. Psychol.* **68**, 47–72 (2017).
23. M. I. Posner, Orienting of attention. *Q. J. Exp. Psychol.* **32**, 3–25 (1980).
24. C. W. Eriksen, J. D. St James, Visual attention within and around the field of focal attention: A zoom lens model. *Percept. Psychophys.* **40**, 225–240 (1986).
25. M. S. A. Graziano, *Consciousness and the Social Brain* (Oxford University Press, New York, 2013).
26. M. S. A. Graziano, *Rethinking Consciousness: A Scientific Theory of Subjective Experience* (W. W. Norton, New York, 2019).
27. M. S. A. Graziano, S. Kastner, Human consciousness and its relationship to social neuroscience: A novel hypothesis. *Cogn. Neurosci.* **2**, 98–113 (2011).
28. M. S. Graziano, T. W. Webb, The attention schema theory: A mechanistic account of subjective awareness. *Front. Psychol.* **6**, 500 10.3389/fpsyg.2015.00500. (2015).
29. T. W. Webb, H. H. Kean, M. S. A. Graziano, Effects of awareness on the control of attention. *J. Cogn. Neurosci.* **28**, 842–851 (2016).
30. A. I. Wilterson et al., Attention control and the attention schema theory of consciousness. *Prog. Neurobiol.* **195**, 101844. (2020).
31. J. van den Boogaard, J. Treur, M. Turpin, "A neurologically inspired network model for Graziano's attention schema theory for consciousness" in *International Work-Conference on the Interplay Between Natural and Artificial Computation: Natural and Artificial Computation for Biomedicine and Neuroscience*, J. Ferrandez Vincente, J. Álvarez-Sánchez, F. de la Paz López, J. Toledo Moreo, H. Adeli, Eds. (Springer, Cham, Switzerland, 2017), pp. 10–21.
32. V. Mnih et al., Human-level control through deep reinforcement learning. *Nature* **518**, 529–533 (2015).
33. S. Guadarrama et al., TF-Agents: A Reliable, Scalable and Easy to Use TensorFlow Library for Contextual Bandits and Reinforcement Learning. <https://github.com/tensorflow/agents>. Accessed 15 January 2020.
34. M. Abadi et al., TensorFlow: Large-scale machine learning on heterogeneous distributed systems. arXiv [Preprint] (2013). <https://arxiv.org/abs/1603.04467> (Accessed 15 January 2020).
35. V. Mnih, N. Heess, A. Graves, K. Kavukcuoglu, Recurrent models of visual attention. *Adv. Neural Inf. Process. Syst.* **27**, 2204–2212 (2014).
36. A. Borji, L. Itti, State-of-the-art in visual attention modeling. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**, 185–207 (2013).
37. G. Deco, E. T. Rolls, A neurodynamical cortical model of visual attention and invariant object recognition. *Vision Res.* **44**, 621–642 (2004).
38. J. H. Reynolds, D. J. Heeger, The normalization model of attention. *Neuron* **61**, 168–185 (2009).
39. O. Le Meur, P. Le Callet, D. Barba, D. Thoreau, A coherent computational approach to model bottom-up visual attention. *IEEE Trans. Pattern Anal. Mach. Intell.* **28**, 802–817 (2006).
40. M. A. Schwemmer, S. F. Feng, P. J. Holmes, J. Gottlieb, J. D. Cohen, A multi-area stochastic model for a covert visual search task. *PLoS One* **10**, e0136097 (2015).
41. A. Vaswani et al., Attention is all you need. arXiv [Preprint] (2017). <https://arxiv.org/abs/1706.03762> (Accessed 3 December 2020).
42. V. S. Ramachandran, W. Hirstein, The perception of phantom limbs. The D. O. Hebb lecture. *Brain* **121**, 1603–1630 (1998).
43. U. Ansorge, M. Heumann, Shifts of visuospatial attention to invisible (metacontrast-masked) singletons: Clues from reaction times and event-related potentials. *Adv. Cogn. Psychol.* **2**, 61–76 (2006).
44. P. Hsieh, J. T. Colas, N. Kanwisher, Unconscious pop-out: Attentional capture by unseen feature singletons only when top-down attention is available. *Psychol. Sci.* **22**, 1220–1226 (2011).
45. J. Ivanoff, R. M. Klein, Orienting of attention without awareness is affected by measurement-induced attentional control settings. *J. Vis.* **3**, 32–40 (2003).
46. Y. Jiang, P. Costello, F. Fang, M. Huang, S. He, A gender- and sexual orientation-dependent spatial attentional effect of invisible images. *Proc. Natl. Acad. Sci. U.S.A.* **103**, 17048–17052 (2006).
47. R. W. Kentridge, T. C. Nijboer, C. A. Heywood, Attended but unseen: Visual attention is not sufficient for visual awareness. *Neuropsychologia* **46**, 864–869 (2008).
48. C. Koch, N. Tsuchiya, Attention and consciousness: Two distinct brain processes. *Trends Cogn. Sci.* **11**, 16–22 (2007).
49. A. Lambert, N. Naikar, K. McLachlan, V. Aitken, A new component of visual orienting: Implicit effects of peripheral information and subthreshold cues on covert attention. *J. Exp. Psychol. Hum. Percept. Perform.* **25**, 321–340 (1999).
50. V. A. Lamme, Why visual attention and awareness are different. *Trends Cogn. Sci.* **7**, 12–18 (2003).
51. P. A. McCormick, Orienting attention without awareness. *J. Exp. Psychol. Hum. Percept. Perform.* **23**, 168–180 (1997).
52. L. J. Norman, C. A. Heywood, R. W. Kentridge, Object-based attention without awareness. *Psychol. Sci.* **24**, 836–843 (2013).
53. G. F. Woodman, S. J. Luck, Dissociations among attention, perception, and awareness during object-substitution masking. *Psychol. Sci.* **14**, 605–611 (2003).
54. Y. Tsuchiya, Y. Sasaki, T. Watanabe, Greater disruption due to failure of inhibitory control on an ambiguous distractor. *Science* **314**, 1786–1788 (2006).
55. P. Stoerig, A. Cowey, Blindsight in man and monkey. *Brain* **120**, 535–559 (1997).
56. J. C. Marshall, P. W. Halligan, Blindsight and insight in visuo-spatial neglect. *Nature* **336**, 766–767 (1988).
57. U. Ansorge, W. Kunde, M. Kiefer, Unconscious vision and executive control: How unconscious processing and conscious action control interact. *Conscious. Cogn.* **27**, 268–287 (2014).
58. B. J. Baars, *A Cognitive Theory of Consciousness* (Cambridge University Press, New York, 1988).
59. G. Carruthers, A metacognitive model of the sense of agency over thoughts. *Cogn. Neuropsychiatry* **17**, 291–314 (2012).
60. D. Chalmers, Facing up to the problem of consciousness. *J. Conscious. Stud.* **2**, 200–219 (1995).
61. A. Damasio, *Self Comes to Mind: Constructing the Conscious Brain* (Pantheon, New York, 2015).
62. S. Dehaene, *Consciousness and the Brain* (Viking Press, New York, 2014).
63. D. C. Dennett, *Consciousness Explained* (Little-Brown, New York, 1991).
64. S. Grossberg, The link between brain learning, attention, and consciousness. *Conscious. Cogn.* **8**, 1–44 (1999).
65. V. A. Lamme, Towards a true neural stance on consciousness. *Trends Cogn. Sci.* **10**, 494–501 (2006).
66. H. Lau, D. Rosenthal, Empirical support for higher-order theories of conscious awareness. *Trends Cogn. Sci.* **15**, 365–373 (2011).
67. T. Metzinger, *Being No One: The Self-Model Theory of Subjectivity* (MIT Press, Cambridge, MA, 2003).
68. M. I. Posner, Attention: The mechanisms of consciousness. *Proc. Natl. Acad. Sci. U.S.A.* **91**, 7398–7403 (1994).
69. D. Rosenthal, *Consciousness and Mind* (Oxford University Press, New York, 2005).
70. G. Tononi, M. Boly, M. Massimini, C. Koch, Integrated information theory: From consciousness to its physical substrate. *Nat. Rev. Neurosci.* **17**, 450–461 (2016).
71. D. Dennett, Illusionism as the obvious default theory of consciousness. *J. Conscious. Stud.* **23**, 65–72 (2016).
72. K. Frankish, Illusionism as a theory of consciousness. *J. Conscious. Stud.* **23**, 1–39 (2016).