

Article

# The Evolution of Neuroplasticity and the Effect on Integrated Information

Leigh Sheneman <sup>1,2,†</sup>, Jory Schossau <sup>2,3,†</sup>  and Arend Hintze <sup>1,2,3,\*</sup> 

<sup>1</sup> Department of Computer Science and Engineering, Michigan State University, East Lansing, MI 48823, USA; leighasheneman@gmail.com

<sup>2</sup> BEACON-Center for the Study of Evolution in Action, Michigan State University, East Lansing, MI 48823, USA; jory@msu.edu

<sup>3</sup> Department of Integrative Biology Michigan State University, East Lansing, MI 48823, USA

\* Correspondence: hintze@msu.edu

† These authors contributed equally to this work.

Received: 12 March 2019; Accepted: 21 May 2019; Published: 24 May 2019



**Abstract:** Information integration theory has been developed to quantify consciousness. Since conscious thought requires the integration of information, the degree of this integration can be used as a neural correlate ( $\Phi$ ) with the intent to measure degree of consciousness. Previous research has shown that the ability to integrate information can be improved by Darwinian evolution. The value  $\Phi$  can change over many generations, and complex tasks require systems with at least a minimum  $\Phi$ . This work was done using simple animats that were able to remember previous sensory inputs, but were incapable of fundamental change during their lifetime: actions were predetermined or instinctual. Here, we are interested in changes to  $\Phi$  due to lifetime learning (also known as neuroplasticity). During lifetime learning, the system adapts to perform a task and necessitates a functional change, which in turn could change  $\Phi$ . One can find arguments to expect one of three possible outcomes:  $\Phi$  might remain constant, increase, or decrease due to learning. To resolve this, we need to observe systems that learn, but also improve their ability to learn over the many generations that Darwinian evolution requires. Quantifying  $\Phi$  over the course of evolution, and over the course of their lifetimes, allows us to investigate how the ability to integrate information changes. To measure  $\Phi$ , the internal states of the system must be experimentally observable. However, these states are notoriously difficult to observe in a natural system. Therefore, we use a computational model that not only evolves virtual agents (animats), but evolves animats to learn during their lifetime. We use this approach to show that a system that improves its performance due to feedback learning increases its ability to integrate information. In addition, we show that a system's ability to increase  $\Phi$  correlates with its ability to increase in performance. This suggests that systems that are very plastic regarding  $\Phi$  learn better than those that are not.

**Keywords:** information integration theory; neuroevolution; autonomous learning

## 1. Introduction

Information Integration Theory [1,2] defines the neural correlate  $\Phi$  that seeks to quantify how conscious a system is. It does this by measuring the degree to which the components of a system compute more information than the sum of its parts [3]. At the same time, the ability to integrate information is probably beneficial for cognitive systems in general, as it allows these systems to make decisions based on information obtained from multiple environmental sensors and integrate this with knowledge of past outcomes. Imagine a child playing soccer for the first time trying to score a goal. To kick the ball, she must visually locate the ball and determine how to make contact with her foot

to move the ball in the desired direction. The cognitive structure she has as a child is drastically different from that of the professional soccer player she might become. We assume that the amount of information integration is different for the child and the professional because the neural structures are different [4], but we are unsure if the amount of information integration increases or decreases with neural change. The child's cognitive system might need to process a great amount of information while the professional's cognitive system is already streamlined by training to the point that kicking goals becomes a near unconscious reflex. Alternatively, the professional, being much more accurate and with more experience, might integrate more information than the child.

Information integration has been proposed to be a neural correlate for conscious experience [5] under the assumption that consciousness requires information integration. The neural correlate  $\Phi$  has undergone many iterations and refinements, with each new iteration also slightly redefining the measurement. Therefore, we will use  $\Phi$  to refer to the concept of information integration in general, and the special nomenclature when referring to specific versions of  $\Phi$ , such as:  $\Phi^{Max}$ ,  $EI$ ,  $\Phi_{atomic}$ , or  $\Phi_{maxH}$  (see below for their definitions). Also, attention is a different process than consciousness [6] and thus  $\Phi$  should not necessarily correlate with it. Obviously, the task of kicking the ball could also be interpreted in the light of attention. Since we do not have a neural correlate for attention, and our focus is on information integration, we disregard attention for now.

Unfortunately, we do not yet have the ability to thoroughly measure  $\Phi$  in humans, and thus we must use computational models instead. A lot of previous research [7–11] used animats, which are virtual organisms that are controlled by a Markov Brain [12] (a form of neural networks, MB in short). Since the introduction of information integration theory [4,13–15], it has been shown that the amount of information integrated by a system increases over the course of Darwinian evolution [7–9]. We have established that task complexity defines a lower bound of  $\Phi$  [8] for the task to be solved and that more complex environments require more  $\Phi$  [9]. The use of animats in these investigations has shown that they are technically easy to use, have a small computational footprint, and can be perfectly introspected, making them the ideal surrogate in the absence of natural organisms. However, the interaction between  $\Phi$  and lifetime learning has not been explored. Here, we evolve animats to perform a learning task, and measure  $\Phi$  at various stages during their life to answer the question "How does  $\Phi$  change in a learning organism?"

### 1.1. Markov Brains

MBs are networks of computational logic units (which from now on we will refer to as gates) that are genetically encoded and subject to evolution [7–9,11,16–20] (For a detailed description of Markov Brains, see [12]). Conventionally, the MBs contain deterministic and probabilistic gates, each of which can be thought of as a  $2^i \times 2^o$  state transition table,  $P$ , where  $i$  is the number of inputs and  $o$  is the number of outputs. Probabilistic gates are coded so each input is associated with a vector of probabilities determining the likelihood for each possible output to occur. Deterministic logic gates work the same way, except their vector contains a single probability of 1.0 and the rest are 0.0. Each MB is represented by a sequence of numbers that serves as the genome and contains stretches encoding computational units similar to how genes encode proteins. These genes determine the number, type, and the probability tables if necessary. In addition, each gene contains information about how gates are connected to each other and how they interface with sensors or actuators.

Markov Brains built from deterministic and probabilistic logic gates can perform complex computations and can control animats in virtual environments. They can use past experiences to inform future decisions, and form representations about their environment [16]. However, they do not change their composition, connections, or logic, and instead remain static throughout their lives. Changes only occur as a result of mutation, which happens between generations. If we want to study learning and how learning changes ability to integrate information, then the animats need past experiences to change the computation over time. Similar to long-term potentiation (LTP), which strengthens some neural connections and weakens others, here we allow probabilities of logic gates to

change based on internal feedback. We pioneered this method by adding so-called *feedback gates* to the repertoire of gates from which Markov Brains can be built. For a detailed description, see Sheneman and Hintze 2017 [21].

In essence, feedback gates have two additional inputs for positive or negative reinforcement. At each update, a feedback gate records which output was mapped to the current input. For positive feedback, the probability that leads to this mapping becomes increased; in the case of negative feedback, this probability becomes decreased. After that, the probability table is normalized so that each row of the probability table sums to 1.0. The strength of the feedback as well as the number of time points the feedback gates record is evolvable as well, and encoded in the genes describing each feedback gate. Observe that the feedback itself is not supplied externally from a supervisor, but the MB itself must evolve a mechanism to detect which actions have a positive or negative outcome. Based on this, the MB then applies feedback to their feedback gates internally.

Unfortunately, feedback gates and probabilistic logic gates violate a critical requirement of information integration theory 3.0 [2]. To accurately compute  $\Phi$ , components must be decomposable [22]. Decomposable gates receive inputs and have more than one output. However, the outputs of these gates need to be probabilistically independent from one another as they are not allowed to share information. For this reason, we created probabilistically independent (called *decomposable*) versions of all the gates, including feedback gates (see the Methods for a detailed description of decomposable gates). In addition to facilitating  $\Phi$  analysis, they provide a small benefit for evolutionary adaptation. For example, in [21] it took 500,000 generations for animats to evolve, but we show with decomposable gates the same performance after 200,000 generations.

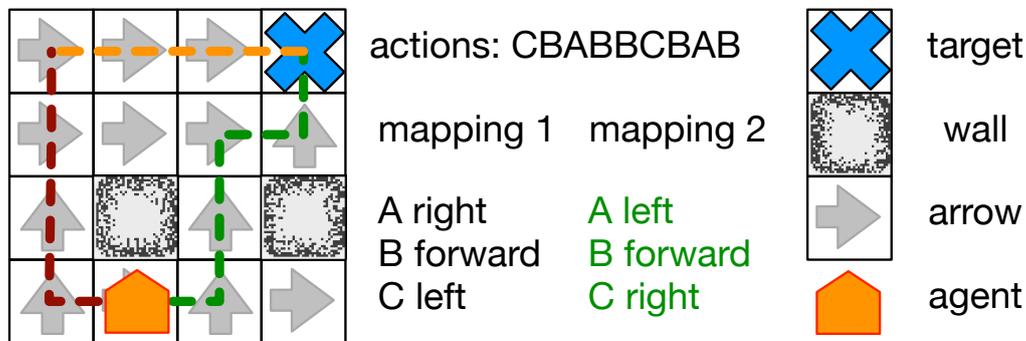
A key part of MB evolution is that their interaction with their environment requires them to be embodied [23]. These learning animats evolve to use an internal feedback mechanism to change the function of their logic gates. In this way, we can compare animats that are naïve about their task with those that have experienced the task before. These experiences can structurally change the neural substrate, and consequently  $\Phi$  may change.

### 1.2. Learning Environment

In previous work, we evolved animats who learn to solve a navigation task (see Figure 1). We now use that framework to investigate the effect of learning on  $\Phi$ . To evolve these animats, we placed them in a 2D lattice (64 cells wide and 64 cells long) environment with walls as exterior-facing cells and a few of the inner cells. From the remaining cells, we selected a random one as the goal position and computed Dijkstra paths [24] to the goal from all other cells. Each of the cells on the path is assigned a direction indicating the next cell on the path to guide the animat's navigation. Simply put, the animat only needs to follow this series of directions from cell to cell to reach the goal. The animat was randomly placed 32 steps from the goal and oriented in a random direction. The animat was then allowed to reach the goal as many times as it could within 512 timesteps. At each step, the animat could either move forward, turn 90°, or wait. If the animat reached the goal, then the animat was replaced at a random cell 32 steps away from the goal.

This presents a very easy task and would not require animats to remember or learn anything. To improve the complexity of the task and to require animats to remember and learn, the animats were not allowed to directly choose to move or turn but instead had to choose from a set actions associated with results. An animat would choose one of four actions (A, B, C, or D), which were each associated with a different one of four results: turning 90° left, turning 90° right, doing nothing, or moving forward. This indirection of action to result allowed different mappings between those actions and their associated results, preventing the agents from evolving for any one particular mapping. There are 24 possible combinations of mappings from actions to results, so every time an animat was placed in a new variation of the environment the next combination was used. For example, in one mapping the action A may result in a left turn, B a right turn, C nothing, and D a movement forward. In the next map, the selection of A moves the animat forward, B does nothing, C leads to a left turn, and D leads to

a right turn. After each animat in the population experienced all 24 combinations, tournament selection was used to determine which animats reproduced. Over 200,000, generations animats evolved the ability to properly navigate their environment given unknown action-result mappings. To do so, they needed to learn by trial and error how actions mapped to movement and rotation.



**Figure 1.** Overview of the navigation task. An animat (orange) must navigate towards a goal (blue cross). Each cell contains an arrow (gray arrow) indicating the shortest path towards the goal, circumventing obstacles and walls (black/gray blocks). However, the animat has no definite control over its body and can only choose actions which are mapped to movements. An example sequence of action, as seen in the figure (CBABBCBAB) might lead towards the goal (mapping 2 green path) while given a different mapping (mapping 1 dark red path) it might lead the animat astray. If the animat can integrate the information it collected while navigating erroneously (dark red path), it could learn the correct mapping and proceed to the goal (orange path).

### 1.3. Different Ways of Measuring $\Phi$

Using animats evolved in this environment, we can now measure the amount of  $\Phi$ . However, the method used to measure  $\Phi$  has changed over the history of information integration theory. The latest version of  $\Phi^{Max}$  [9] seems to be the most appropriate method to measure the amount of information integration, because this version can measure  $\Phi$  for a particular moment in time. Unfortunately, computing  $\Phi^{Max}$  is computationally costly, and current implementations take about 24 h for one animat at one point in time just to find the main complex (nodes that are essential for the integration of information). Here, we need to measure  $\Phi^{Max}$  for hundreds of data points in hundreds of replicates over hundreds of generations, precluding us from using  $\Phi^{Max}$ .

Fortunately,  $\Phi_{atomic}$  (see Equation (1)) is a computationally much simpler measure. While the states of all components at each time point still need to be known, the computation of  $\Phi_{atomic}$  does not require an iteration over all parts of all partitions—a calculation with a complexity of  $O(B_n N^2)$ , where  $B_n$  is the Bell number of  $n$  (nodes), which scales super exponentially. Instead, using  $\Phi_{atomic}$  reduces the problem to a complexity of  $O(n^2)$  (for a more detailed description of  $\Phi_{atomic}$ , see [7]).

$$\Phi_{atomic} = SI_{atomic} - \mathcal{I} \tag{1}$$

$\Phi_{atomic}$  depends on the definition of information that is processed by the entire system over time. To do this, we use joint random variables  $X = X^{(1)}x^{(2)} \dots X^{(n)}$ , where  $X^{(i)}$  represents the system nodes changing the system state over time ( $t$ ). Each  $X_t$  is defined by the probability  $p(x_t)$  to observe  $X_t$  in state  $x_t$ . Time changes each node  $i$  as  $X_0^{(1)} \rightarrow X_1^{(i)}$  and each  $X_t^{(i)}$  is described by probability distribution  $p(x_j^i)$ .

Therefore, for any time step  $t$  to  $t + 1$ , we can calculate the information processed by the system as

$$I(X_t : X_{t+1}) = \sum_{x_t, x_{t+1}} p(x_t, x_{t+1}) \log \frac{p(x_t, x_{t+1})}{p(x_t)p(x_{t+1})} \tag{2}$$

Thus,  $\Phi_{atomic}$  is a quantification of the information processed by the system that is more than the sum of the information processed by the individual units:

$$\Phi_{atomic} = I(X_t : X_{t+1}) - \sum_{i=0}^n I(X_t^{(i)} : X_{t+1}^{(i)}) + \mathcal{I}. \quad (3)$$

where  $I(x_t^{(i)} : X_{t+1}^{(i)})$  represents the information processed by node  $i$ , and  $\mathcal{I}$  is the nonindependence between the network variables [25–29]:

$$\mathcal{I} = \sum_{i=0}^n H(X_t^{(i)}) - H(X_t). \quad (4)$$

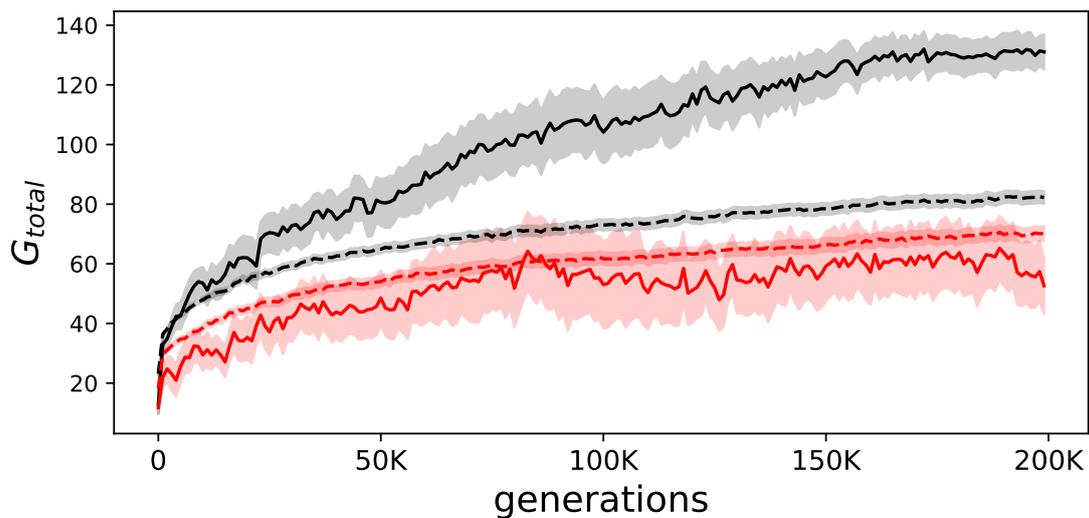
With this “more than the sum of its parts” system information, and spatial integration of nonindependence,  $\Phi_{atomic}$  captures both temporal and spatial network information characteristics. These terms together simplify to

$$\Phi_{atomic} = \sum_{i=0}^n H(X_t^{(i)} | X_{t+1}^{(i)}) - H(X_t | X_{t+1}) \quad (5)$$

It has also been shown that  $\Phi_{atomic}$  is a very good correlate with more complex versions of  $\Phi$  [7,8], making it an ideal substitute, with the exception of no previous method to use on individual time points. However, here we discovered a way to calculate  $\Phi_{atomic}$  for individual time points. In previous experiments [7,8],  $\Phi_{atomic}$  was calculated by recording the animat’s brain states ( $X$ ) at every time point ( $t$ ), and one needs all of these  $X_t$  to  $X_{t+1}$  transitions in order to compute  $\Phi_{atomic}$ . However, for each time point  $t$  we can compute all possible state-to-state transitions the animat’s brain could theoretically experience. This is done by computing the state-to-state transition probability matrix (TPM) at time point  $t$  assuming  $H_t$  to be a uniform random variable encompassing all possible brain states, and  $H_{t+1}$  to be defined by the TPM. Here, we derive the TPM for each time point by stochastic sampling. Observe that a machine or brain that would not change during its lifetime will have the same TPM for each time point. Only if the probabilities defining the state-to-state transitions change, for example as it happens here due to learning, a machine would be able to have different TPMs for each time point. We define the  $\Phi_{atomic}$  measurement performed in this manner as  $\Phi_{maxH}$ . In this way, we can quantify an animat’s ability to integrate information before, during, and after learning has affected change.

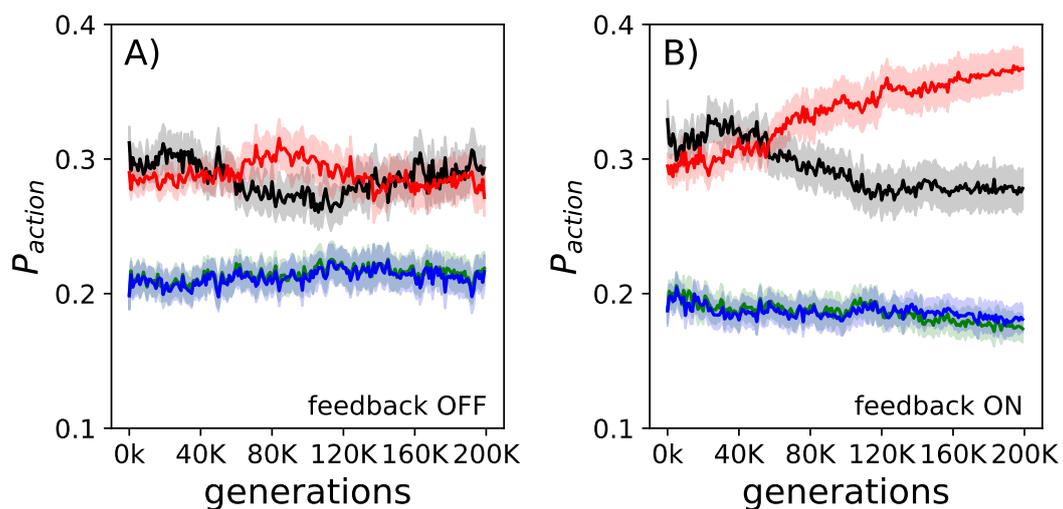
## 2. Results

Animats were evolved to perform a navigation task for which they had to use environmental clues to determine the correct motor outputs, and thus the correct behavior. Given the difficulty of the task, only 13 of all the 200 replicate experiments resulted in animats that performed well at the end of evolution. Specifically, all animats were tested on all possible 24 variations of the environment, one with learning (allowing them to take advantage of their evolved internal feedback mechanisms), and one without learning. After 200,000 generations, only those animats that reached the goal two-fold more with learning than without were considered “elite.” From these 13 elite performers, all animats on the line of decent (LoD, see Methods for details) were evaluated in the same environment to analyze  $\Phi$ . As expected, animats evolved the ability to reach the goal, and used their internal feedback mechanism to do so. The difference between performance with and without internal feedback between the elite animats and all animats shows this (see Figure 2). We further find that elite animats perform on average worse when their internal feedback mechanism is disabled when compared to all other animats. This could be due to the different behavioral strategies. Elite animats obviously have evolved the ability to learn from their environment, and are thus probably more dependent on their feedback mechanisms as opposed to the not so well evolved other animats.



**Figure 2.** Evolution of goal-reaching. The solid lines show the average performance of the elite animats, while the dashed lines show the average performance of all animats. In black, the summation of goals reached ( $G_{total}$ ) over all 24 environmental variations, in red the same but when measured in animats that were prevented from using their evolved internal feedback mechanism. The difference between the solid (elite) and dashed black line (all animats) illustrates how much better elite animats perform. The shadows indicate the standard error of the means.

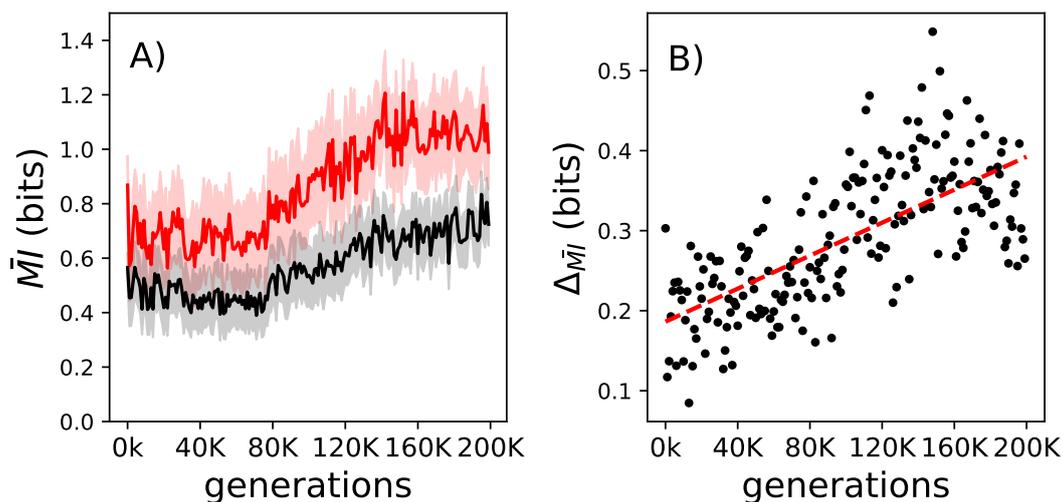
Over the course of evolution, animats optimized their ability to navigate the environment, and to learn how the possible choices map to their movements. As expected, animats learn to move more deliberately through the environment, by reducing the number of unnecessary turns, and idle behavior, while moving forward more often (see Figure 3). These changes in action frequencies are slightly different to those observed before [21] for agents that used non-decomposable feedback gates. We assume that these differences are due to the small functional differences necessary to ensure decomposability. However, they are neglectable.



**Figure 3.** Evolution of goal-oriented action selection. The fraction of actions used (black idle, red forward, green left turns, blue right turns) over the course of evolution averaged over all elite performers. In panel (A), measured when the internal feedback mechanism was turned off, we find almost no change in action selection preference. In panel (B), measured when the internal feedback mechanism was turned on, we find that the number of turns and idles are reduced, while the number of forward movements becomes increased. The shadows represent the standard error of the means for each action.

Decomposable feedback gates must change their probability matrix to allow for lifetime learning. The mutual information that these gates convey between their inputs and outputs should increase over the lifetime, and we expect that the delta of this change also increases over evolutionary time, as observed before [21] for non-decomposable logic gates. As expected, the mutual information is increasing over the lifetime of an animate and due to evolution (see Figure 4).

To obtain an additional impression of how animats behave and how their ability to learn changes over evolutionary times and over their lifetime, see the Supplementary Movie.

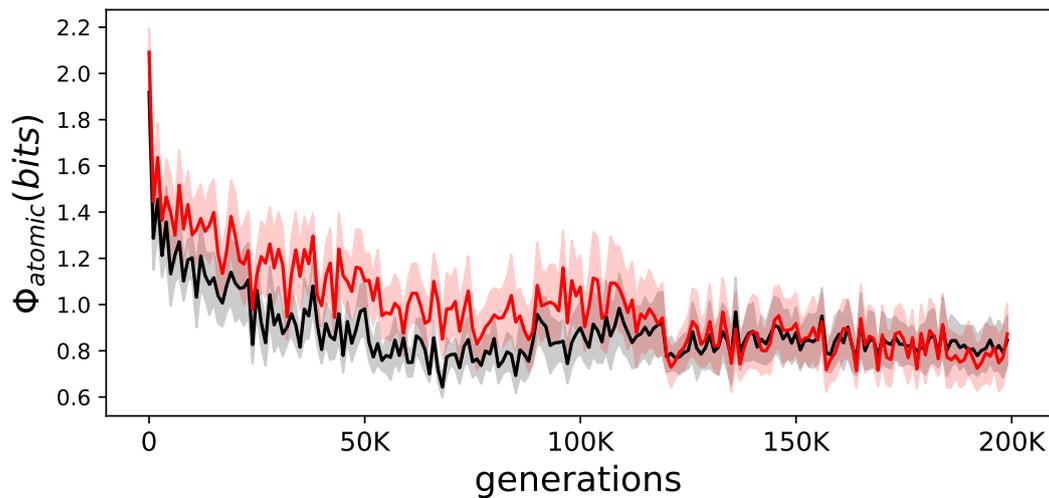


**Figure 4.** The change to mutual information conveyed by decomposable logic gates over their lifetime and over evolutionary time. Panel (A): In black is the average mutual information that the decomposable logic gates convey for all elite performers at the beginning of their lives. In red is the same but measured at the end of the animats' lives. The shadows indicate the standard error of the means for those measurements. Panel (B): The lifetime increase of the mutual information that decomposable logic gates convey between birth and the end of life ( $\Delta \bar{M}I = \bar{M}I_{death} - \bar{M}I_{birth}$  measured for all elite performers on the line of descent. The red dashed line shows the linear regression trend ( $r^2 = 0.697$  and a  $p$ -value of  $2.034 \times 10^{-30}$ ).

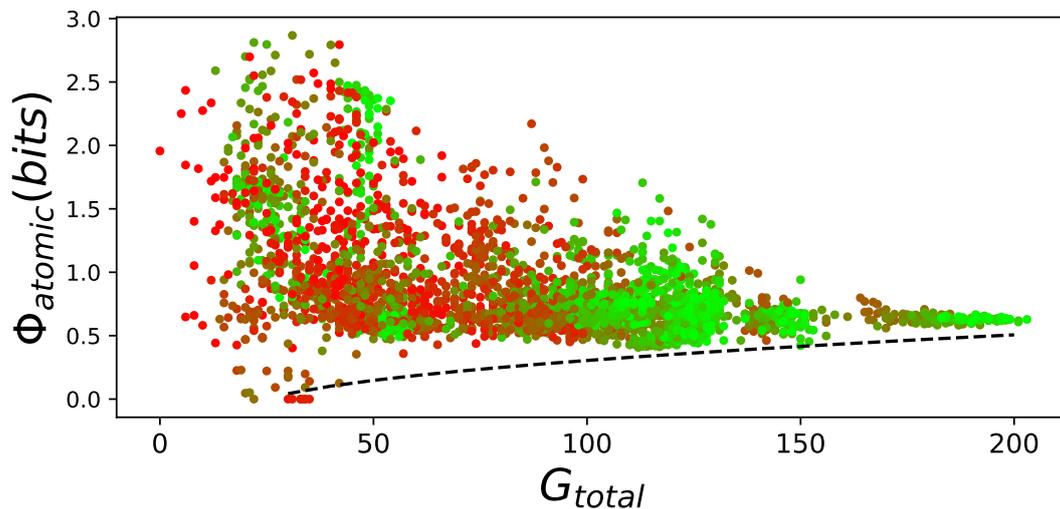
### 2.1. Effects of Evolution on $\Phi$

Earlier experiments with  $\Phi_{atomic}$  showed increases over evolutionary time [7–9,11]. Here, it seems that on average,  $\Phi_{atomic}$  becomes lower, and animats that are not allowed to learn have a slightly higher value (see Figure 5).

However, this result can be explained. Random animats can have high values of  $\Phi_{atomic}$  regardless of their ability to perform. Only well-performing animats require a minimal amount of  $\Phi_{atomic}$ . This has been shown before [7,8] and we observe the same phenomena here (see Figure 6). We find the same lower bound, where animats require a minimal amount of  $\Phi_{atomic}$  to perform. In addition, we observe that at the end of evolution with learning, variance of  $\Phi_{atomic}$  over all animats converges towards zero, whereas without learning evolved animats greatly varied in their ability to integrate information.



**Figure 5.** Evolution of  $\Phi_{atomic}$  for elite animats. The average  $\Phi_{atomic}$  for elite animats on the LoD. In black are the animats that were allowed to learn; in red are the same animats but prevented from learning during their lifetime. The shadows indicate the standard error or the means.

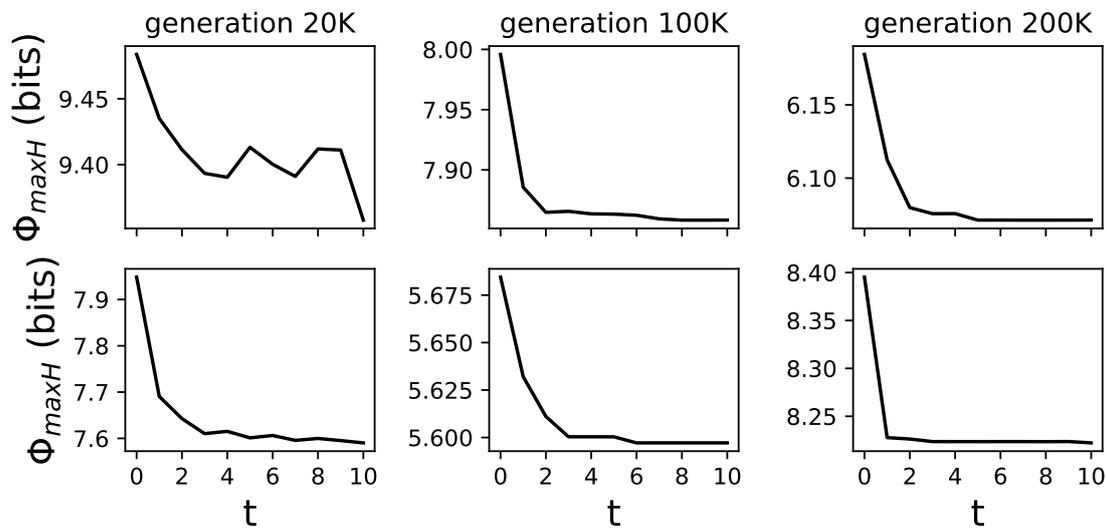


**Figure 6.**  $\Phi_{atomic}$  versus performance on the LoD. For all elite performers on the LoD,  $\Phi_{atomic}$  was plotted against their performance (total number of goals reached, pooled over all possible 24 action mappings  $G_{total}$ ). The LoD generation time point is indicated by a color gradient from red (beginning of evolution) to green (end of adaptation). The lower bound relation of  $\Phi_{atomic}$  to performance  $G_{total}$  is indicated by a dashed black line (hand drawn, since there is no proper way to fit a lower bound). This figure visualizes how  $\Phi_{atomic}$  changes due to evolution for all elite performers. The better animats become, the more goals they reach, and the further right their data points move on the X-axis. Similarly, the less  $\Phi_{atomic}$  they have, the lower they move on the Y-axis. To incorporate the generations (which is roughly similar to the animats' performance), each data point is color coded, starting with red at the beginning of evolution, and becoming green at the end of evolution. This behavior is very similar to that found before [7,8].

## 2.2. Effects of Learning on $\Phi$

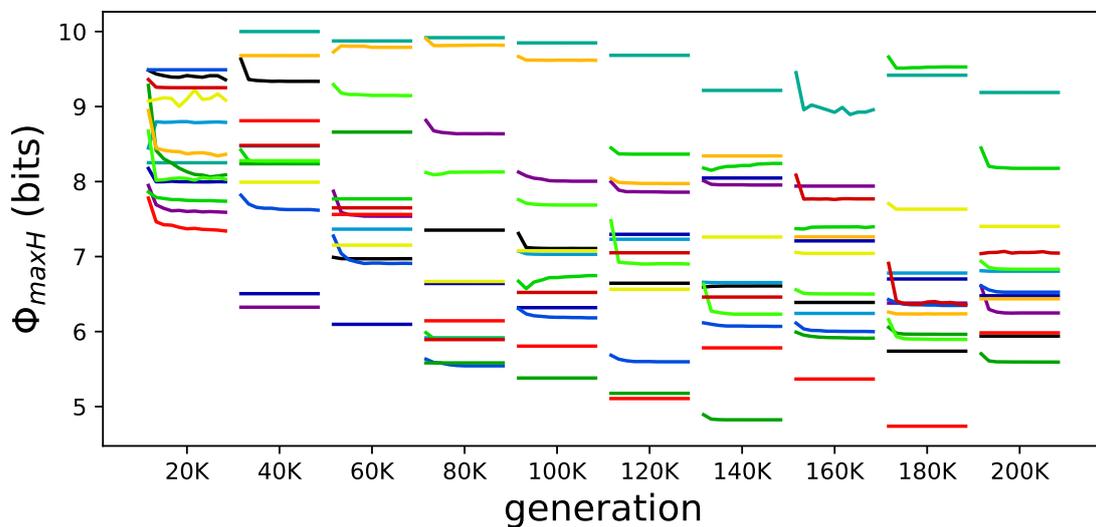
The addition of feedback gates to MBs allows these brains to change from generation to generation and during their lifetime. This allows us to test how these lifetime changes (learning) affect  $\Phi$ . We use  $\Phi_{maxH}$  to assess an animat's ability to integrate information during a specific time point in its life. This is possible because  $\Phi_{maxH}$  only requires a TPM which can be derived at any time point and does not need an observation of the animat in its environment, which might alter its experience through learning. Animats were given 512 time points in the environment, and we calculated  $\Phi_{maxH}$  every

50 time points. As an observation, we find animats to reduce the amount of  $\Phi_{maxH}$  over their lifetime (see Figure 7 for examples of differently evolved animats and how  $\Phi_{maxH}$  changes over their lifetime).



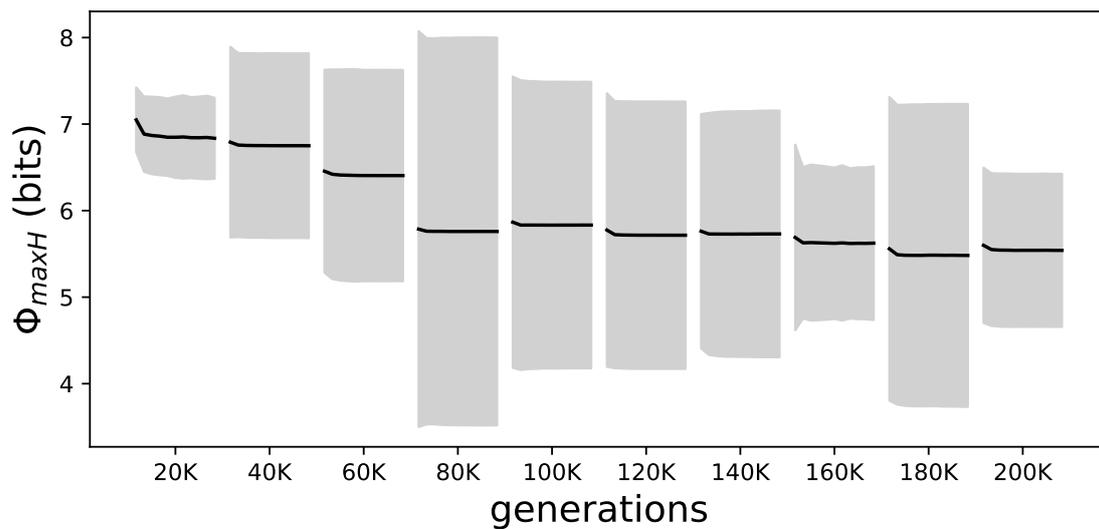
**Figure 7.** Examples of different animats and how  $\Phi_{maxH}$  changes over their lifetime. Six randomly chosen animats and their lifetime changes in  $\Phi_{maxH}$ , averaged over all 24 actions-effect mappings. Animats were chosen from the LoD at generations 20,000, 100,000, and 200,000. These are representative samples, to illustrate how  $\Phi_{maxH}$  changes. All examples can be found in Figure 8.

Beyond this apparent reduction of  $\Phi_{maxH}$  over the animats’ lifetime, we also see two additional trends. The first is a general reduction of  $\Phi_{max}$  over evolutionary time, similar to the reduction of  $\Phi_{atomic}$ . The second observation is an additional but total smaller reduction of  $\Phi_{maxH}$  over the lifetime of elite animats (see Figure 8).



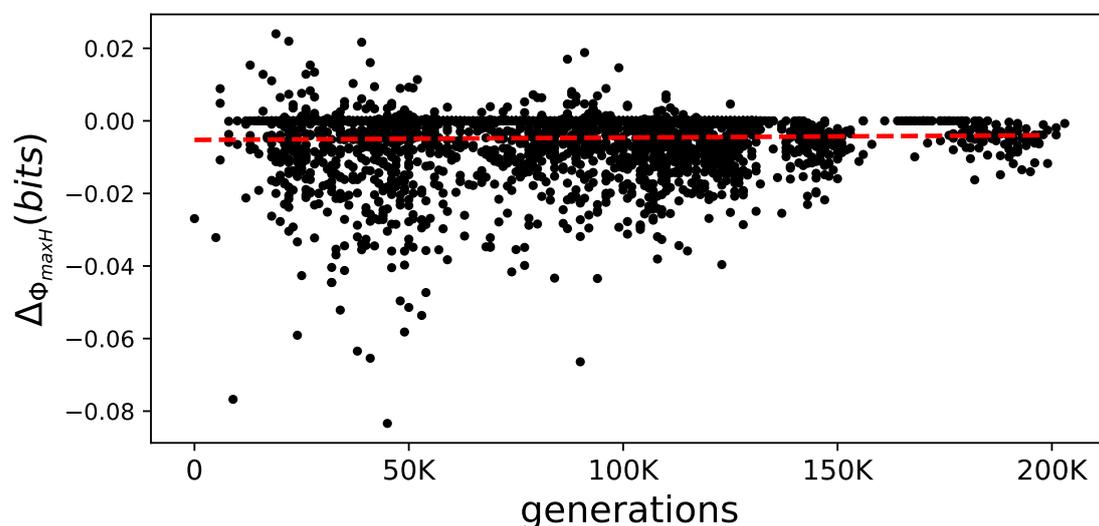
**Figure 8.** Samples of all elite animats and how  $\Phi_{maxH}$  changes over their lifetime for various stages of evolution. Each line represents the change in  $\Phi_{maxH}$  experienced by each animat at different times over their evolutionary history (generations are indicated on the X-axis). Each color represents an independent experiment; the Y-axis is kept to scale so that the general reduction of  $\Phi_{maxH}$  can be observed.

Specifically, we find that the amount of  $\Phi_{maxH}$  at the beginning and end of animats’ lifetimes decreases for all elite performers during evolution (see Figure 9).



**Figure 9.** Evolutionary and lifetime changes of  $\Phi_{maxH}$  for all elite performers. Samples of the elite performers from the LoD every 20 K generations. Each interval shows the average  $\Phi_{maxH}$  from all elite performers over their lifetime as a black line; the gray shadow marks the standard deviation. This figure presents the averages of the values shown in Figure 8 per time point on the LoD. The reduction of  $\Phi_{maxH}$  over evolutionary time as well as within each lifetime can be seen.

In addition, we characterize the reduction of  $\Phi_{maxH}$  by fitting a linear function and using the slope of that fit as a proxy for the reduction. This slope correlates with evolutionary time (generations) over all elite performers (see Figure 10). Both types of analysis (Figures 9 and 10) show that the amount of information integration due to learning is reduced and that the reduction itself becomes smaller.



**Figure 10.** The lifetime change in  $\Phi_{maxH}$  over generations. From all elite performers, the change in  $\Phi_{maxH}$ , measured as the difference in  $\Phi_{maxH}$  at the end of their lifetime and the beginning ( $\Delta_{G_{total}}$ ) over generations. The correlation coefficient of 0.224 is indicated as a linear fit as a red dashed line.

### 3. Discussion

In this work, we measured how lifetime learning influences an animat's ability to integrate information. We had no intuition of how information integration would respond to learning: increase, decrease, or remain. Because such measurement is currently prohibitively difficult with natural

organisms—let alone human beings—we use evolved animats as the next best substitute. Markov Brains controlled these animats, because they were a proven study system used for similar work and thus allow for direct comparisons. The biggest constraint on this type of assessment for any continuous or probabilistic computational system is the necessity to be computationally decomposable. To date, only decomposable probabilistic logic gates [22] have had this feature, and here we introduced decomposable feedback gates. We know of no other system that has this property, is evolvable, and has sufficient neural plasticity to adapt during its lifetime. This precludes other systems for identical studies until their decomposability has been shown (artificial neural networks [30], for example). Regardless, in future work this analysis should be repeated with other computational systems, and eventually on natural organisms. Until then, we conjecture that these results can be generalized to those systems.

Instead of measuring  $\Phi^{Max}$  from IIT 3.0 or  $EI$  from IIT 2.0, here we used the more computationally tractable  $\Phi_{atomic}$  to quantify an animat's ability to integrate information during its lifetime, and  $\Phi_{maxH}$  to measure changes in an animat's ability to integrate information. We assume that these results should generalize to  $\Phi^{Max}$  and hope to use that measure or any successor in the future when computational simplifications can be found.

#### 4. Methods

To study how information integration changes over an organism's lifetime, we created a computational model in which animats are required to navigate in a changing environment. Animats needed to learn how to properly move within this environment. The brain of each animate was a probabilistic finite state machine that consisted of sensor nodes (input), memory nodes (hidden), and action nodes (output). This was implemented by a Markov Brain (MB) [7,12]. The animates were evolved using a genetic algorithm (GA) that selected for individuals that successfully decoded a variety of symbols.

##### 4.1. Environment

Populations of 100 animats were evolved in 200 replicates to perform a spatial navigation task [21] described above. Populations could evolve for 200,000 generations testing their performance on all 24 possible variations of that environment at every generation. Upon evolutionary completion, the LoD for every replicate was constructed. These animats were then again tested in the same environments but their internal feedback mechanism was disabled. The difference in performance between using feedback and not being allowed to use feedback was used as an indicator to identify elite performers. Only animats that on average reached each of the navigation goals two times more often with feedback than without were considered elite. This resulted in 13 independent evolutionary experiments.

##### 4.2. Animat

Animats' brains consisted of 16 states: 4 input sensors, 2 output motors, and 10 hidden states. Each input sensor was assigned a direction to detect: sensor 0 fired when the animat was told to stay put, 1 indicated turn right, 2 go forward, and 3 turn left. As described above, the environment was composed of tiles with arrows that point towards the shortest path to the goal. Sensors could read the arrow of the tile that the animat was currently on. When the sensor detects its given direction, the state was set to 1, otherwise the state was 0. For example, if the map indicated the animat's next closest step was to the left, its first four states would then be: 0, 0, 0, 1.

##### 4.3. Line of Descent

The LoD for an experiment was determined by choosing a random animat in the final generation and tracing its lineage to the first generation [31]. Through this method, the most recent common ancestor (MRCA) and all mutations that reached fixation in the population were recorded.

#### 4.4. Feedback Gates

The key feature of feedback gates is that every input  $i$ , which results in output  $o$  is defined by a probability that has the potential to change with each update due to the feedback received. Each time these numbers change, so does the amount of predictive information shared between nodes, resulting in a change to  $\Phi_{atomic}$  [7]. It is important to note that the typology, or the connections, between nodes does not change during the lifetime of an animat; only the information contained in the connection changes. Topological changes, or changes to the start condition of the probabilities, however, can occur due to mutations from one generation to the other.

#### 4.5. Decomposable Feedback Gates

Decomposable feedback gates superset includes features from feedback gates, decomposable gates, and probabilistic gates. A probabilistic gate relies on a probabilistic mapping between input and output patterns. However, each input pattern is associated with a separate set of probabilities over all output patterns. Decomposable gates are based on the design of probabilistic gates. To ensure decomposability of outputs for a single input pattern, the output patterns must have probabilities such that the probabilities of single outputs are independent. In other words, if an input pattern could produce output 0 0 with probability 0.36, 0 1 with 0.16, 1 0 with 0.16, and 1 1 with 0.04, it is possible to compute the independent probabilities for each of the two bits  $a$  and  $b$  of 0.4 and 0.4. 0 0 (both bits off) =  $(1 - a) * (1 - b)$ , 0 1 =  $(1 - a) * b$ , 1 0 =  $a * (1 - b)$ , and 1 1 =  $a * b$ . However, if the output pattern probabilities related to a single input pattern were 0.02, 0.48, 0.48, and 0.02, then it would be impossible to determine the independent probabilities for each bit because their likelihoods are dependent on the state of one another. This decomposability of independence is important when we use the information theoretic analysis of  $\Phi$  so as not to invalidate the fundamental assumptions of independence.

The probabilistic gates for an animat are constructed upon the *birth* (construction) of the organism, and derive their probabilities directly from the genome. For a decomposable gate, the genome provides the *factors* for each input pattern that are then multiplied together to produce a full row of output pattern probabilities. In this way, each input pattern will probabilistically produce any of the possible output patterns such that the probability of individual output-bit activations within each pattern is independent across all patterns associated with that input.

To allow learning via feedback over a lifetime, we add to the design of each gate a positive feedback connection to one of the states (input, hidden, output), and a negative feedback connection to one of the states. When the input bit for positive feedback is activated, the probability for the previous input–output activation is increased by a small amount, and one of the other output pattern probabilities randomly chosen for that input is decreased by the same delta. Negative feedback happens in a similar way when the input bit wired for negative feedback is activated. This configuration allows feedback to time  $t$  of the time point  $t - 1$ . Additionally, we allow feedback on further time points into the past  $t - 2$ ,  $t - 3$ , and so on, each with their own feedback deltas. Evolution controls how any of these feedback mechanisms are used. For every gate, the genome encodes the following: which two states are used for positive and negative feedback, the feedback history length between 0 and 3 time points, and the positive and negative feedback deltas between 0 and max probability. Since each individual probability is modified separately due to feedback, and then used to reconstruct the probability table of the gate decomposable feedback, gates remain decomposable even after feedback is applied.

## 5. Conclusions

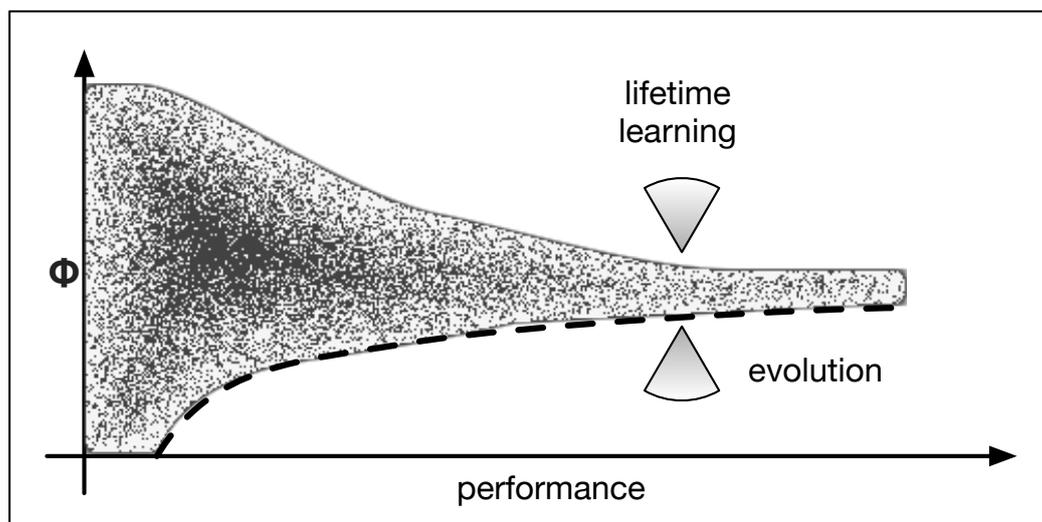
From earlier experiments, we understood how  $\Phi$  changes over evolutionary time: there is a lower bound such that performance only increases together with  $\Phi$ , and that increasingly complex environments also require an increase in  $\Phi$ . While we knew that learning increased the animats'

performance, we did not know how lifetime learning would affect  $\Phi$ . Furthermore, it was unclear if the increase in performance was due to streamlining the computation, thus reducing integration, or if better performance required more integration of information? To determine how the integration of information was affected by learning, we tested how  $\Phi_{atomic}$  behaves over evolutionary times as well as how it is affected by lifetime learning. The tests performed here confirm all prior observations about  $\Phi$ .

Over the course of the animat's life,  $\Phi_{maxH}$  is reduced, suggesting that at birth the animat had to integrate more information to streamline the computation. As a result, after learning, the animat integrates less information. Applying this idea to the soccer player analogy implies that the well-trained player integrates less information than she did at the beginning of her career.

The second interesting observation is that the reduction of information integration becomes further minimized over evolutionary time. The more evolved an animat becomes, the less it must reduce its information integration due to learning. This suggests that the learning process itself becomes optimized to reduce the change to information integration. This seems surprising. Because learning reduces  $\Phi_{maxH}$ , one might assume that better learning correlates with more reduction. However, our results contradict this notion and instead show that the better an animat learns, the less it reduces  $\Phi_{maxH}$ . This phenomenon might result from faster learning being more advantageous than slower learning in this particular environment. The larger the probability table change necessary, the slower learning is. Consequently, there is an evolutionary advantage to change the system as little as possible, which might explain why animats adapt  $\Phi_{maxH}$  minimally at the end of evolution. This phenomenon could also explain the observation that  $\Phi_{atomic}$  converges to the smallest possible value (see Figure 5) instead of being more distributed towards the end of evolution.

Therefore, we hypothesize that the behavior of  $\Phi$  depends on two "forces" (influences). Evolution optimizes the function of the animat which is only limited by its ability to integrate information. Consequently, evolution increases  $\Phi$  over time (see Figure 11 lower arrow). On the other hand, lifetime learning streamlines the computation and reduces the amount of integration. The degree to which this needs to be done is further minimized to improve performance and to accelerate learning. This creates a second force which minimizes  $\Phi$  (see Figure 11 upper arrow). These two influences now force  $\Phi$  to converge.



**Figure 11.** Illustration of two forces. Over the course of evolution, we find  $\Phi$  to necessarily increase to allow for higher levels of performance. This is illustrated by the lower bound to  $\Phi$  (dashed line) and suggests that evolution seeks to increase  $\Phi$  (lower of the two arrows). At the same time, lifetime adaptations are time-consuming in this environment and should thus be minimized. Consequently, we hypothesize that the second influence due to learning reduces  $\Phi$  (upper arrow).

It is entirely possible that the phenomenon we observe is dependent on the specific task, type of evolutionary optimization, and the life time learning mechanism implemented here. Other forms of learning for other tasks, particularly those where time is not of the essence, may behave differently. Investigating the changes to  $\Phi$  due to learning (or other forms, such as back propagation or reinforcement [30]) presents an interesting opportunity for future research.

**Supplementary Materials:** The following are available at <http://www.mdpi.com/1099-4300/21/5/524/s1>.

**Author Contributions:** conceptualization, A.H. and L.S.; methodology, A.H., J.S. and L.S.; software, L.S. and J.S.; validation, J.S. and A.H.; formal analysis, A.H.; writing—original draft preparation, L.S.; writing—review and editing, A.H. and J.S.; visualization, L.S. and A.H.; supervision, A.H.; project administration, A.H.; funding acquisition, A.H.

**Funding:** This work was supported in part by Michigan State University through computational resources provided by the Institute for Cyber-Enabled Research, and is-based in part upon work supported by the National Science Foundation under Cooperative Agreement No. DBI-0939454.

**Acknowledgments:** We would like to thank Larissa Albantakis, Giulio Tononi, and Christoph Adami for the many useful discussions on  $\Phi$  and their valuable input on this work.

**Conflicts of Interest:** The authors declare no conflict of interest.

## References

- Balduzzi, D.; Tononi, G. Integrated information in discrete dynamical systems: Motivation and theoretical framework. *PLoS Comput. Biol.* **2008**, *4*, e1000091. [[CrossRef](#)] [[PubMed](#)]
- Oizumi, M.; Albantakis, L.; Tononi, G. From the Phenomenology to the Mechanisms of Consciousness: Integrated Information Theory 3.0. *PLoS Comput. Biol.* **2014**, *10*, e1003588. [[CrossRef](#)] [[PubMed](#)]
- Tononi, G. *Phi: A Voyage from the Brain to the Soul*; Pantheon Books: New York, NY, USA, 2012.
- Tononi, G. Integrated information theory of consciousness: An updated account. *Arch. Ital. de Biol.* **2012**, *150*, 56–90.
- Tononi, G.; Koch, C. The neural correlates of consciousness. *Ann. N. Y. Acad. Sci.* **2008**, *1124*, 239–261. [[CrossRef](#)] [[PubMed](#)]
- Koch, C.; Tsuchiya, N. Attention and consciousness: Two distinct brain processes. *Trends Cogn. Sci.* **2007**, *11*, 16–22. [[CrossRef](#)] [[PubMed](#)]
- Edlund, J.A.; Chaumont, N.; Hintze, A.; Koch, C.; Tononi, G.; Adami, C. Integrated information increases with fitness in the evolution of animats. *PLoS Comput. Biol.* **2011**, *7*, e1002236. [[CrossRef](#)] [[PubMed](#)]
- Joshi, N.J.; Tononi, G.; Koch, C. The minimal complexity of adapting agents increases with fitness. *PLoS Comput. Biol.* **2013**, *9*, e1003111. [[CrossRef](#)]
- Albantakis, L.; Hintze, A.; Koch, C.; Adami, C.; Tononi, G. Evolution of Integrated Causal Structures in Animats Exposed to Environments of Increasing Complexity. *PLoS Comput. Biol.* **2014**, *10*, e1003966–19. [[CrossRef](#)]
- Albantakis, L.; Tononi, G. Fitness and neural complexity of animats exposed to environmental change. *BMC Neurosci.* **2015**, *16*, P262. [[CrossRef](#)]
- Schossau, J.; Adami, C.; Hintze, A. Information-Theoretic Neuro-Correlates Boost Evolution of Cognitive Systems. *Entropy* **2016**, *18*, 6. [[CrossRef](#)]
- Hintze, A.; Edlund, J.A.; Olson, R.S.; Knoester, D.B.; Schossau, J.; Albantakis, L.; Tehrani-Saleh, A.; Kvam, P.; Sheneman, L.; Goldsby, H.; et al. Markov Brains: A Technical Introduction. *arXiv* **2017**, arXiv:1709.05601.
- Ay, N.; Wennekers, T. Temporal infomax leads to almost deterministic dynamical systems. *Neurocomputing* **2003**, *52*, 461–466. [[CrossRef](#)]
- Ay, N.; Bertschinger, N.; Der, R.; Güttler, F.; Olbrich, E. Predictive information and explorative behavior of autonomous robots. *Eur. Phys. J. B* **2008**, *63*, 329–339. [[CrossRef](#)]
- Ay, N. Information geometry on complexity and stochastic interaction. *Entropy* **2015**, *17*, 2432–2458. [[CrossRef](#)]
- Marsteller, L.; Hintze, A.; Adami, C. The Evolution of Representation in Simple Cognitive Networks. *Neural Comput.* **2013**, *25*, 2079–2107. [[CrossRef](#)]

17. Chapman, S.; Knoester, D.B.; Hintze, A.; Adami, C. Evolution of an artificial visual cortex for image recognition. In Proceedings of the ECAL 2013: The Twelfth European Conference on Artificial Life, Sicily, Italy, 2–6 September 2013; pp. 1067–1074.
18. Hintze, A.; Mirmomeni, M. Evolution of Autonomous Hierarchy Formation and Maintenance. In *Artificial Life 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*; The MIT Press: Cambridge, MA, USA, 2014; pp. 366–367.
19. Kvam, P.; Cesario, J.; Schossau, J.; Eisthen, H.; Hintze, A. Computational evolution of decision-making strategies. *arXiv* **2015**, arXiv:1509.05646.
20. Olson, R.S.; Hintze, A.; Dyer, F.C.; Knoester, D.B.; Adami, C. Predator confusion is sufficient to evolve swarming behaviour. *J. R. Soc. Interface* **2013**, *10*, 20130305. [[CrossRef](#)] [[PubMed](#)]
21. Sheneman, L.; Hintze, A. Evolving autonomous learning in cognitive networks. *Sci. Rep.* **2017**, *7*, 16712. [[CrossRef](#)] [[PubMed](#)]
22. Schossau, J.; Albantakis, L.; Hintze, A. The role of conditional independence in the evolution of intelligent systems. In Proceedings of the Genetic and Evolutionary Computation Conference Companion, Berlin, Germany, 15–19 July 2017; pp. 139–140.
23. Clark, A. *Being There: Putting Brain, Body, and World Together Again*; MIT Press: Cambridge, MA, USA, 1998.
24. Dijkstra, E.W. A note on two problems in connexion with graphs. *Numer. Math.* **1959**, *1*, 269–271. [[CrossRef](#)]
25. Schneidman, E.; Still, S.; Berry, M.J.; Bialek, W. Network information and connected correlations. *Phys. Rev. Lett.* **2003**, *91*, 238701. [[CrossRef](#)]
26. McGill, W.J. Multivariate information transmission. *Psychometrika* **1954**, *19*, 97–116. [[CrossRef](#)]
27. Tononi, G.; Sporns, O.; Edelman, G.M. A measure for brain complexity: Relating functional segregation and integration in the nervous system. *Nat. Acad. Sci.* **1994**, *91*, 5033–5037. [[CrossRef](#)]
28. Lungarella, M.; Pegors, T.; Bulwinkle, D.; Sporns, O. Methods for quantifying the informational structure of sensory and motor data. *Neuroinformatics* **2005**, *3*, 243–262. [[CrossRef](#)]
29. Lungarella, M.; Sporns, O. Mapping information flow in sensorimotor networks. *PLoS Comput. Biol.* **2006**, *2*, e144. [[CrossRef](#)] [[PubMed](#)]
30. Russell, S.J.; Norvig, P. *Artificial Intelligence: A Modern Approach*; Pearson Education Limited: London, UK, 2016.
31. Lenski, R.E.; Ofria, C.; Pennock, R.T.; Adami, C. The evolutionary origin of complex features. *Nature* **2003**, *423*, 139–144. [[CrossRef](#)] [[PubMed](#)]



© 2019 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).