

Evolutionary Adaptation to Social Information Use without Learning

James M. Borg and Alastair Channon

School of Computing and Mathematics, Keele University, UK
{j.borg,a.d.channon}@keele.ac.uk

Abstract. Social information can provide information about the presence, state and intentions of other agents; therefore it follows that the use of social information may be of some adaptive benefit. As with all information, social information must be interpretable and relatively accurate given the situation in which it is derived. In both nature and robotics, agents learn which social information is relevant and under which circumstances it may be relied upon to provide useful information about the current environmental state. However, it is not clear to what extent social information alone is beneficial when decoupled from a within-lifetime learning process, leaving evolution to determine whether social information provides any long term adaptive benefits. In this work we assess this question of the adaptive value of social information when it is not accompanied by a within-lifetime learning process. The aim here is to begin to understand when social information, here expressed as a form of public information, is adaptive; the rationale being that any social information that is adaptive without learning will be a good base to allow the learning processes associated with social information to evolve and develop later. Here we show, using grounded neuroevolutionary artificial life simulations incorporating simulated agents, that social information can in certain circumstances provide an adaptive advantage to agents, and that social information that more accurately indicates success confers more reliable information to agents leading to improved success over less reliable sources of social information.

Keywords: Social Information, Public Information, Evolutionary Adaptation, Neuroevolution, Artificial Life

1 Social Information, Learning and Evolution

Social information can broadly be thought of as information derived from the behaviours, actions, cues or signals of other agents [1]. As social information necessarily involves the direct or indirect broadcasting of information in to the public domain, it is sometimes known as (or conflated with) public information [2]. Here we assess whether the use of social information in populations of simulated neuroevolutionary agents is adaptive when decoupled from within-lifetime

learning processes. Within-lifetime learning processes confer significant adaptive advantages to agents employing them, be it through the development of a set of robust and flexible behaviours, the rapid adaptation to new environments or circumstances, the quick incorporation of new information, or the guiding of the evolutionary process itself [3]. The adaptive advantages of learning are particularly potent when social information is incorporated alongside innovation and individual learning [4], resulting in social learning and potentially even cultural evolution [5]. However, as beneficial as within-lifetime social learning processes are, it is unclear to what extent social or public information has an adaptive benefit when decoupled from these learning processes and evolution is left to determine the value of social information. Is the incorporation of social information alone enough to gain an adaptive advantage over non-social agents? Or are learning processes necessary to allow social information to confer any benefits? These are the questions that we address in this paper.

Social learning is seen widely in nature [6] in a range of species as diverse as humans and nine-spined stickleback fish [7]. The mechanisms and processes that underpin social learning are themselves broad, ranging from teaching, imitation and emulation to stimulus enhancement and exposure [8], with any of these mechanisms potentially leading to the formation of traditions and cultures [5,9]. However, within each social learning category there is some dependence on who information is obtained from, be it a teacher or which agent is unintentionally (or intentionally) exposing an individual to new information. As social learning is necessarily conformist, a poor social information model may result in the discovery and propagation of sub-optimal behaviours [10]. Despite the potential pitfalls of over-conformist social learning, including sub-optimal behaviour development [10] and even population collapse [11,12], social learning, and therefore social information transfer, can be of great benefit to agents, thus explaining why even simple forms of social information transfer are seen so widely in nature [13,8,14] and have been shown to produce complex behaviours that are easily attributed to more complex social learning mechanisms like imitation [15]. At the heart of the problem being addressed here are three core arguments. (1) Information is a fitness enhancing resource [16], even when information suppression is seen to be adaptive [17] or when information is encoded or interpreted incorrectly [16] - any new information about the world enables populations of agents to better adapt to the world they are in, even if this means disregarding or suppressing information. (2) Incremental evolution is not a process of unguided random variations, but a process that itself can adapt in a way that is analogous to the kind of learning seen in cognitive organism [18], leading to complex and robust adaptive traits in nature, autonomous robots [19] and simulated agents [20] in the same way that learning can lead to complex and adaptive behaviours (though on a different time scale). (3) Inadvertently expressed public information and simple mechanisms of social information transfer can lead to behaviours that are sufficiently complex to enable cultural evolutionary processes [21,5]. These three core arguments give us good reason to believe that social information without within-lifetime learning processes should still be adaptive, and therefore lead to

evolution adopting the use of social information to the benefit of social agents over non-social agents. Though we must still be mindful that social information may be at odds with personal beliefs [21] or lead to population-level conformism to sub-optimal behaviours [10], thus leading to a trade-off between the accommodation of social information and the evolution of robust evolved behaviours. This leads us to the hypothesis that agents making use of social information should outperform non-social agents: any additional information, that is not just noise, that provides more information about the environment should lead agents to an improved “performance” in the environment over agents without access to such information. However, social information may only be useful when it accurately indicates success or indirectly leads to success, and therefore may provide little or no concrete benefit in complicated or less predictable environments - in these more challenging environments learning may be necessary to allow temporarily useful social information to be quickly adopted and then rejected when it is no longer relevant. This hypothesis will be tested by modelling populations of agents who have no social information available to them and populations of agents with various forms of social information available to them. Each social information strategy will be tested against the non-social strategy, starting initially with the most basic strategy available: presence, with the null hypothesis in each case being that the social population does not show an improved ability to solve the task at hand compared to non-social agents. The social information strategies used here are: presence, action, health and age. Presence social information simply enables agents to detect the presence of other agents (non-social agents are essentially blind to other agents); action enables agents to see what other agents are currently doing; health enables agents to see the current energy or battery state of others; and age information enables agents to see how long others have lived for.

2 EnVar and Environmental Set-up

The task world used here is known as EnVar. EnVar is a bounded (non-toroidal) 2D environment containing a variety of consumable resources known as plants. Plants are recognised by agents simply as an RGB value. Plants are divided into a number of species, each with a base RGB value and a radius in RGB space. Plants are generated within these RGB regions and identified as belonging to the nearest species according to euclidean distance in RGB space. Each plant species is assigned an energy value, which is transferred to agents if the plant of that species is consumed; energy values may be positive or negative. Notionally the EnVar world is broken up in to cells, though here each cell represents a pixel and therefore the world can be considered to be continuous. Plants in the world take up a number of cells, forming a block, with each block only being able to be eaten a certain number of times before being exhausted (here set to be 200 eating events). Once a plant block has been exhausted it is no longer consumable and therefore removed from the world to be replaced by a new block from a random plant species somewhere else in the world - this maintains

a constant number of food blocks in the world at any time. Agents are permitted to share space with a plant resource but cannot overlap with each other, thus removing the possibility of agents piling up on top of one another on valuable food resources. In this work EnVar is set up to create a 700×700 pixel sized cell world, containing five hundred 10×10 pixel blocks of plants. In order to test our hypothesis we test populations of social and no social agents in a set of increasingly difficult environments. Environmental difficulty is dictated by the ratio of positive food resources to negative food resources. The simplest world used here has an equal (1 : 1) ratio of positive food species to negative food species. Tests get progressively harder by increasing the number of negative food species, whilst maintaining only one positive food species, resulting in the most difficult world used here having a 1 : 9 ratio of positive food species to negative food species. As each plant species has a equal chance of appearing in the world, and covers approximately the same portion of RGB space, agents in the most difficult environment are nine times more likely to experience a negative plant resources than a positive plant resource. In the results section below environment 1 relates to a 1 : 1 ratio environment, with environment 9 relating to a 1 : 9 ratio environment. For all tests here negative food species come with an energy value $E_{neg} = -10.0$, with positive food species contributing an energy value of $E_{pos} = 1.0$ when consumed. This provides a strong evolutionary pressure to avoid eating negative food species.

3 Neuroevolutionary Model

Agents in the EnVar simulation world are grounded 2D simulated agents, controlled by a hybrid neural network architecture known as the *Shunting Model*¹. The shunting model uses two interacting neural networks to determine agent behaviours, here represented as a discrete set of agent actions. The two interacting networks are known as the *Decision Network* and the *Shunting Network*. The decision network is simply a feed-forward neural network comprising of an input layer, one hidden layer and an output layer. Outputs from the decision network are used to produce a locally-connected, topologically-organised network of neurons known as the shunting network, which simply places and organises agent preferences for environmental features and states in such a way to allow the agent to hill climb in a shunting space (known as the activity landscape) that directly maps on to their immediate neighbourhood. The shunting network weights are fixed for all agents, whereas the decision network is genetically encoded and is subject to change via evolution.¹

3.1 The Shunting Network

The shunting network is a locally-connected, topologically-organised network of neurons that was originally used for collision free motion planning in robots [22]

¹ For a detailed overview of the Shunting Model architecture please see [22,23,24]

and has been subsequently applied in a number of 2D and 3D artificial life models [23,4,24,10]. Here the shunting network’s topology is simply superimposed on to the environment, with each cell in the network topology directly relating to a pixel within an agent’s visual field. Using the shunting equation (see equation 1) values for each cell (which can be interpreted as representing an environmental feature or state, and are initially set by the Iota output I obtained from the decision network) are propagated across the neurons/cells of the network, producing an activity landscape with peaks and valleys representing desirable and undesirable features in the environment. The result is a landscape which allows the agent to follow a route determined by the higher Iota values while avoiding undesirable valleys. A mock-up example of an activity landscape with a snapshot of the visual field it represents can be seen in figure 1.

$$\frac{dx_i}{dt} = -Ax_i + \sum_{j \in N_i} w_{ij} [x_j]^+ + I_i \quad (1)$$

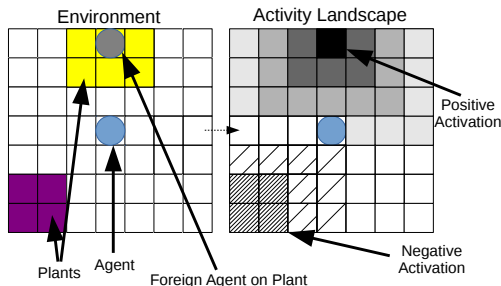


Fig. 1. Mock-up transition from agent visual field to shunting network activity landscape: The left-hand grid shows the agent’s visual field with two plant objects and one other agent occupying the same space as a plant. The right-hand grid shows an example activity landscape for the visual field. The agent determines that an agent on a plant is an interesting feature and therefore assigns it a strong positive Iota value (I), whereas the purple plant is seen negatively and is therefore assigned a strong negative Iota value. These Iota values propagate over the activity landscape using equations 1 and 2. The central agent then chooses to move within its immediate Moore neighbourhood to the cell with highest activity value.

In equation 1 each node in the shunting network corresponds to one pixel within an agent’s visual field; x_i is the activation of neuron i ; A is the passive decay rate; N_i in the receptive field of i ; w_{ij} is the connection strength from neuron j to i , specified to be set by a monotonically decreasing function of the Euclidean distance between cells i and j ; the function $[x]^+$ is $\max(0, x)$; and I_i is the

external input to neuron i (known as the Iota value). The shunting network is advantageous as it exhibits computational efficiency by not explicitly searching over all possible paths. In line with the work of Stanton and Channon [24], we use a simplified, stable solution for equation 1 as seen in equation 2. Here constant $x_i^{new} = x_i$ for all i . The maximum Iota value is $maxI = 15$, with the resulting value for x_i^{new} also being capped at a minimum Iota value $minI = -15$. This stops Iota values growing out of control, whilst providing a large enough maximum value (and a small enough minimum value) to ensure activity propagation across the network. In order to allow propagation to occur within a time-step, the shunting equation must be run a number of times, we take this number of iterations to be equal to the diameter of the visual field.

$$x_i^{new} = \min \left(\frac{1}{8} \sum_{j \in N_i} [x_j]^+ + I_i, maxI \right) \quad (2)$$

The shunting model implemented here differs in a number of significant ways from previous Artificial Life implementations [23,4,24,10]. In these previous implementations agents see their entire environment, have a set number of discrete environmental features and states to set Iota values for, and are in the environment alone to complete a predetermined task. Here agents have a limited view of the world, have the possibility of needing to set an Iota value for a plant of any given RGB value, and exist as a population within the environment (leading to possible input states where an agent can be seen on a particular plant). In order to accommodate these differences the shunting model here is run independently for each pixel in an agent’s visual field, which is set here to have a radius of 30 pixels from center of the agent, with information about that pixel being included as part of the agent’s decision network input layer. In this way an Iota value is calculated for each unique environmental state within an agent’s visual field (in previous models, each discrete environmental state was included as an output, with only an agent’s internal state or current cell’s state being accommodated in the input layer of the decision network). This change does not change the resulting behaviour of the shunting model or activity landscape, just the way in which information is passed to the shunting network from the decision network. In order to minimise the amount of processing time required to populate and create the activity landscape, Iota values are only collected for unique states experienced by an agent - for a state to be unique it must be a newly experienced set of decision network inputs (discussed below). To further optimise processing time, an agent will only produce an activity landscape if its outputs determine that it should move in the current time step; agents that are not moving do not need an activity landscape.

3.2 The Decision Network, Neuroevolution and Reproduction

Evolution in the model is applied only to the decision network. The decision network here is a feed-forward neural network comprised of seven standard in-

put nodes, and an additional social input node in social information tests, eight hidden units, and two output nodes, resulting in 112 - 128 weights. Each network layer is fully connected, with floating point weights in the range [-1:1] being directly encoded from an agent's genotype. A standard sigmoid activation function is used at each hidden and output node, though outputs processed for deriving agent actions are then scaled to be within the range [0:1] and the Iota output is scaled to be within the range [$minI$: $maxI$]. As the agent is expected to produce an Iota value to feed in to the shunting network for each unique environmental feature or state within its visual field, inputs into the decision network must accommodate both the internal state of the agent, the state of their current environment, and the state of the environmental feature they are assessing; this leads to there being two sets of input nodes. The first set of input nodes are simply plant RGB inputs - if the agent is viewing empty space these inputs are set to -1, else they are set to be the normalised RGB of the plant being viewed. Following these inputs are a series of generic inputs, which are dependent on the agent's internal state and the current environmental state. These inputs are the agent's current battery level in the normalised range [0:1], a moving average of the agent's battery level over the previous 100 time steps, the agent's current external environmental state and a moving average environmental state, which are both set to be +1 and do not change in the tests presented here (the model is set-up to accommodate external environmental change which is not used here). In social information tests agents have an additional input based on the agent they of viewing.

The genotype, which is essentially an array of weights, is subjected to both mutation and crossover should a reproduction event take place. The crossover mechanism used here is single point crossover, with per locus mutation occurring with probability $p_{mut} = 1/L$, where L is the length of the genotype. Mutation is achieved by way of Gaussian random noise, with a value taken from a normal distribution with $\mu = 0$, $\sigma = 0.01$ being either subtracted or added to the floating point value at the loci to be mutated. All weight values are bounded in the range [-1:1]. Reproduction events take place only in response to a death event. Agents can die if they run out of energy, or if they are in the lowest 10% of agents ranked by energy at the end of an epoch. The first method for removing agents from the population ensures that agents cannot remain in the population with no energy, the second method ensures space is made for new agents to be created even if the population as a whole is successful at maintaining above zero energy levels, thus maintaining a selection pressure for task improvement. Both methods of death are not directly related to task ability as it is possible for a good agent to be unlucky and never, or rarely, experience a positive food resource whereas less able agents may have the fortune to be born near an abundance of food resources or be born relatively close to the end of an epoch. This method of reproduction maintains a constant population size of 200 agents. The new agent, or child, created to replace the removed agent is the progeny of two agents, one of whom is selected in a tournament, the other of which is selected randomly from the remaining population. The tournament selection

mechanism applied here takes two agent from the population, compares their current energy levels, and selects the agent with the higher energy level as a parent. Like in nature this isn't a perfect measure of fitness as it is possible the agent is young and therefore has not yet had time to loose significant amounts of energy, or the agent could have simply been lucky or unlucky with available food sources. However, in general agents with more effective behaviours will on average find themselves with a better energy levels than agents with less effective behaviours, thus driving evolution toward behaviours that are more suited to the task or environment at hand. The second parent is selected randomly to ensure the population doesn't become dominated by the progeny of a small sub-set of the population, thus maintaining a level of exploration in the genotypic search space. New agents are placed in the world within the visual field of one of their parents.

3.3 Agent Actions and Action Energy Costs

The agents in the model have a set of simple, discrete, actions available to them, through the output layer of their decision networks: wait, eat or move. The decision network has two outputs, an Iota output to be fed into the shunting network and an eat/wait output. The agent first considers its current input state at its current position - if the agent produces an Iota value above the threshold $\theta_a = 0.5$ it indicates the agent is happy with its current state and position and therefore does not move (an activity landscape is therefore not calculated as it not needed). The agent's eat/wait output is then considered; if the output produces a value above the threshold $\theta_b = 0.5$ the agent attempts to eat whatever may be at its current position; agents are welcome to try and eat at locations where no plant is present, but no benefit for this action is conferred, and the eat action is considered to be an unsuccessful eating attempt rather than a wait action. If an agent decided to eat at a location containing a plant, the plant's energy is transferred to the agent, this does not necessarily lead to the exhaustion of the plant resource, as plants are considered as a mass. The Iota output is in the range [-1:1], which is then scaled to be within the range $[minI : maxI]$ for use in the shunting network, whereas the eat/wait output is limited to the range [0:1]. If the eat/wait output gives an output below the expected threshold the agent simply waits at its current location. Waiting and eating both reduce an agents energy by 0.1 energy units (though eating may result in a net energy gain), with moving using up 0.2 energy units per time step. Agents will only move if their Iota output for their current location is below threshold θ_a , in this case an activity landscape is created based on the Iota outputs for all visible environmental features. Agents are born with, and are able to achieve, a maximum energy level of 100 units. As epochs here constitute 1000 time steps, an agent would be able to survive for a maximum of one epoch, or one thousand time steps by remaining inactive. In order to avoid moving agents moving around in circles, or moving backwards and forwards, in neutral space (where there is no activity gradient from the activity landscape) consecutive neutral move actions maintain the same direction of travel with probability $p_{dir} = 0.9$.

Measurements are taken to determine whether an eat event was successful or unsuccessful. Any eat action that does not result in a non-negative energy providing food source being consumed is considered to be unsuccessful, so only eating non-energy reducing plants is a successful eating action. In order to measure a population of agents' success in a given environment, the difference between successful and unsuccessful eating actions is measured. This difference measure is useful as it is possible for agents to spend an equal amount of time eating successfully and unsuccessfully, which would demonstrate a strong performance on measure of successful eating, but a weak performance on a measure of unsuccessful eating - the difference instead demonstrates a neutral performance, so a population that spends very little time eating, but all of that time eating successfully (so a picky eating strategy) would be a better performing population than a locust like population that eats everything in sight.

3.4 Social Information Strategies

Populations of agents using social information differs only very slightly from non-social populations; social information populations have an additional input unit for social information, thus non-social agents are rendered blind to other agents in the world. The social information strategies explored here, including the no social strategy are discussed below:

NO SOCIAL: No input node is available to the agent to enable social information to be used by the agent's decision network. Agents proceed with no information about other agents.

PRESENCE: The social information input node receives an input of +1 if another agent is present within the visual field. No other information about the agent being viewed is used. This strategy is not dissimilar to the Inadvertent Information strategy used by agents in the work by Mitri et al. [17], though the agents explored in the work presented here do not have a choice about whether they express social information or not (this is the case for all social information strategies presented here).

ACTION: An input representing the current action state of the agent being viewed. The wait action is input as a value of 0, eat is input as 0.5 and move is represented as 1.

HEALTH: The current energy levels of the agent being viewed are normalised to be within the range [0:1] and input to the viewing agent's decision network.

AGE: The age (in time steps) of agent being viewed is normalised using a hyperbolic tangent function of the logarithm of the age, which is then normalised to be within the range [0:1]. See formula (3) where a represents agent age in time steps.

$$input_a = (\tanh(\log(a)) + 1) / 2 \quad (3)$$

4 Results and Discussion

Forty populations of each social information strategy (including no social) were tested on each environment (1 \rightarrow 9). Each population was permitted to evolve in the environment for 100 epochs of 1000 time steps. Reproduction and death events occurred both within and at epoch, meaning all populations were a mix of young and older agents at all stages of evaluation, with agents having no maximum age limit. Population data was accumulated for each epoch, and collected at the end of each epoch. As we are primarily interested here in the final test performance achieved by a population, not the pathway toward this achievement, average metrics were taken for each population, for each environment, for the last 25 epochs of a test, by which point performance had stabilised across measures. The results presented here are the median values of the 40 populations' average last 25 epochs of data - as this data was rarely normally distributed the medians were considered to be of more use than means. In order to derive the statistical significance between population data for each social information strategy a Mann-Whitney U test was used, with p values being derived from the resulting Z-scores. Figure 2 presents Z-score values on an inverted secondary y-axis, with p-value being represented by highlighting over Z-score data points. In order to test our hypothesis, that populations of agents making use of social information should outperform non-social agents, we measure the difference between how often agents successfully and unsuccessfully apply their eat actions, thus allowing us to measure the effectiveness of the eating behaviour within populations. Only comparisons for each social information strategy against the no social strategy are undertaken to see if any statistically significant differences arise. We go on to further analyse a wider array of metrics, including successful and unsuccessful eating actions in isolation, agent turnover, and average agent age.

4.1 Eat Action Performance

In figure 2 we can see the difference between successful and unsuccessful eating actions for each social information strategy compared to results for non-social populations. Looking first at populations with no social information (black line on all graphs in figure 2) we see that the median difference crosses zero, and therefore indicates the eat action is being applied unsuccessfully more often than successfully, at environment 3 (a 1:3 positive to negative food ratio). All social information strategies manage to maintain the eat action in favour of successful eating until a more difficult environment - this is most notable for both the Health and Age social information strategies where eat actions do not begin to favour unsuccessful eating until environment 5, with the Health strategy re-crossing zero briefly, and the Age strategy maintaining an almost neutral profile for all environments after environment 5. This suggests there is a benefit to social information in that social information may allow populations to maintain successful behaviours in more challenging and difficult environments. However, if we look more closely at the resulting Z-scores and p-values we see that

both the Presence and Action strategies rarely demonstrate a significantly better difference in eat actions over populations of no social agents, and even when significant differences are seen they are with relatively weak and therefore lead us to the conclusion that we cannot say with any certainty that either the Presence or Action social information strategy provides a significant improvement over having no social information at all. Despite the poor performance seen for all strategies in later environments, all strategies were capable of enabling at least one population to achieve a positive eat profile in all environments. It is also worth noting the inconsistent results observed with regard to the No Social strategy in environments 7 and 8. Despite the median result fluctuating in a way that suggest environment 8 was less challenging than environment 7, there was no statistically significant difference between the distribution of results for these environments.

Despite Presence and Action social information being of dubious value, it does seem that both Health and Age social information provide a more convincing benefit. We can see in figure 2(c) that populations using social information about the health of other agents demonstrate a significantly better difference in eating actions until environment 7, this performance difference is most noticeable in less difficult environments (environments 1 and 2) where we see a p value <0.01 . Populations using social information about age (as seen in figure 2(d)) also demonstrate a significantly better difference in eating action in less difficult environments, though the statistical significance over environments is less consistent. However, the two most significant Z-scores seen relate to no social information vs. age social information on environments 1 and 2, which suggest that social information about age is particularly useful in these less difficult, but still challenging environments. From this data we can begin to see the potential advantages of certain types of social information.

4.2 Social Information Performance in Less Difficult Environments

In figure 2 we see that environment 1, where there is a 1:1 ratio of positive to negative plant resources, gives rise to a significant difference in eating performance when social information populations are compared to populations with no access to social information, with this result being extended to environment 2 (a 1:2 ratio) for both Health and Age social information populations. This suggests a particular benefit to using social information in less difficult environments. It is worth noting here that whilst environment 1 and 2 are less difficult than later environments used here, they are still themselves reasonably challenging given that we could have tested in environments with positive plant resources in abundance. Having a 1:1 or 1:2 ratio of positive to negative plant resources provides a reasonable challenge, so much so that in environment 3 we see that non-social populations, relying on evolution alone and having no access to social information, now begin to struggle at the task. Figure 3 shows a wider range of performance metrics for environment 1, including the breakdown of successful and unsuccessful eating actions in isolation. Here we see that the success of social

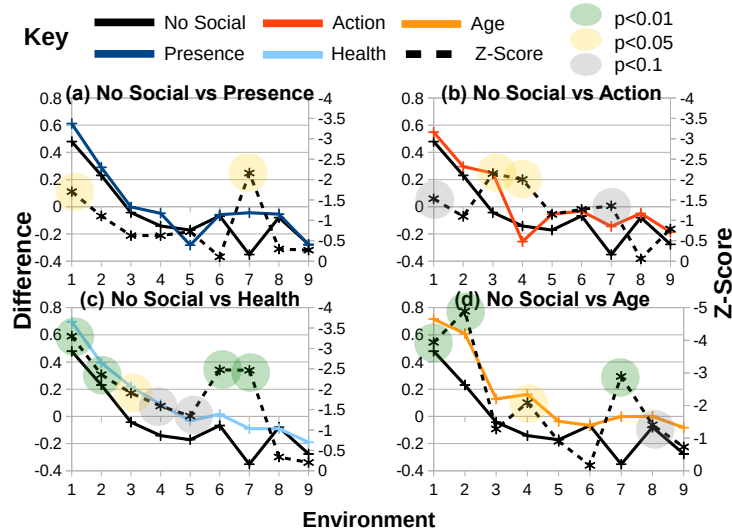


Fig. 2. Difference between successful and unsuccessful eat actions: Graphs showing the difference in % of actions that qualified as successful eating actions and % of actions that qualified as unsuccessful eating actions in each environment, for each social information strategy compared against no social information. All graphs show the Z-score from an Mann-Whitney U test on the secondary y axis, with highlighting included to indicate statistical significance. Each data point represents the median of the average results for forty populations.

information populations is as a result of both higher eat success rates and lower eat failure rates, though it is interesting to note that Age, Presence and Action social information populations are capable of demonstrating very low levels of eat success, even when compared to No Social populations, when the full data range is considered. The main driving force behind the success of social populations, especially Health and Age, seems to be consistently low eat fail rates across populations - the upper quartile ranges for both of these strategies not exceeding 0.02 (2% of actions). This suggests that social information is often being used to help agents avoid or not consume negative plant resources. Age and Health information may be particularly useful for this purpose as it would allow agents to avoid or ignore young or unhealthy agents whilst developing a preference for healthy and older agents. Whilst Presence or Action information may also be useful for the purposes of discrimination (move towards areas of high agent presence, or follow moving agents for example), they are both potentially riskier sources of information compared to Health or Age which both provide information about agent success. Figure 4, which shows performance metrics for environment 2, also shows that for Age social information this abil-

ity to maintain consistently low rates of unsuccessful eating alongside a strong eating success performance is maintained in slightly harder environments. We can also see that for unsuccessful eating actions, the upper quartile range for social information strategies is comparable to the median for no social populations.

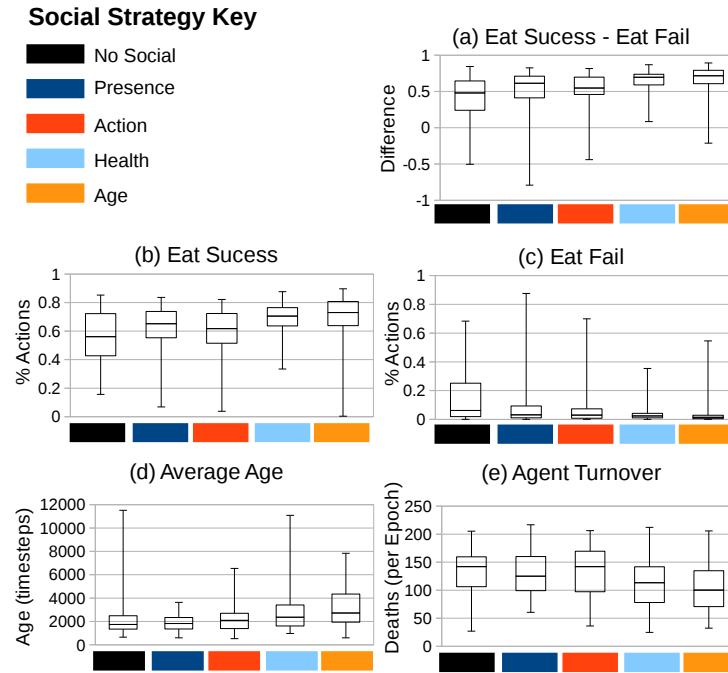


Fig. 3. Environment 1 box plots: Box plots for the eat action and other population metrics, including a breakdown of successful and unsuccessful eating actions, average agent age, and agent turnover, in environment 1.

Alongside information about eating, both figures 3 and 4 also give information on average agent age and agent turnover. For both environments 1 and 2 we see both Age and Health social information enabling populations to accomplish a high average agent age with an accompanying reduction in agent turnover (fewer agents dying within an epoch due to running out of energy), though the median agent turnover for Health social information is comparable to the no social tests. Both Presence and Action populations fail to distinguish themselves from No Social populations, suggesting the improvements in eating performance seen most notably in environment 1 do not necessarily translate directly to im-

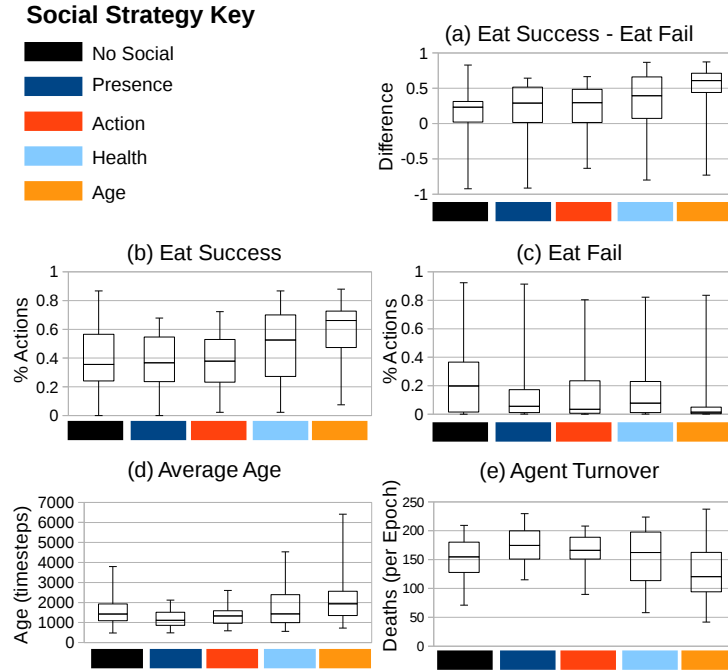


Fig. 4. Environment 2 box plots: Box plots for the eat action and other population metrics, including a breakdown of successful and unsuccessful eating actions, average agent age, and agent turnover, in environment 2.

proved survival, this suggests there must be other underlying behaviours that are causing these populations to use more energy thus resulting in lower average ages and a higher agent turnover when compared to the Age and Health social information populations. We suspect the indeterminate quality of both Presence and Action information causes agents using this information to be less discerning about which agents and plant resources they move towards, resulting in less informed movement and therefore less efficient energy expenditure - though further analysis will have to be done to confirm these suspicions.

5 Conclusions and Further Work

The work presented here, alongside results from of Mitri et al. [17], contribute to the discussion on the adaptive value of social information for evolved simulated agents by demonstrating that social information can provide an adaptive benefit to a neuroevolutionary process when decoupled from a within-lifetime learning

process. However, we do see that social information is only of an consistent adaptive benefit in less difficult environments, and when the social information itself is informative. This work also demonstrates the potential adaptive benefits of simple social and public information strategies such as social influence, social facilitation, stimulus enhancement, and local enhancement [25,13,8,14], adding further weight to the work by Noble and Todd [15] in which it was argued that simple social learning mechanisms are capable of producing complex adaptive behaviours that may easily be confused for the resulting behaviours of more complex social learning mechanisms. The social information strategies implemented here could be argued to be mechanisms of stimulus and local enhancement as the social information inadvertently expressed here by agents could be used by others as an attractor to unfamiliar plant resources or a promoter of eating (or other) behaviours. However, we also see evidence of social information potentially being used to ignore locations or being used to suppress eating (or other) behaviours, which may indicate some level of information suppression [17]. Moving forward we intend to undertake a greater analysis of the behaviours being expressed by agents in this work. It would be of interest to see how often non eating actions are utilised by agents and whether the amount of movement undertaken by agents is promoted by having social information, with further investigations being necessary to ascertain whether this movement results in greater or lesser agent grouping, as evidence from other fields suggests that social information should result in larger groups, thus promoting more informed individual decision making based on the larger quantities of social information made available as a result of having a larger social group [1].

References

1. King, A.J., Cowlshaw, G.: When to use social information: the advantage of large group size in individual decision making. *Biology Letters* **3**(2) (2007) 137–139
2. Bonnie, K.E., Earley, R.L.: Expanding the scope for social information use. *Animal Behaviour* **74**(2) (2007) 171–181
3. Nolfi, S., Floreano, D.: Learning and evolution. *Autonomous robots* **7**(1) (1999) 89–113
4. Borg, J.M., Channon, A., Day, C.: Discovering and maintaining behaviours inaccessible to incremental genetic evolution through transcription errors and cultural transmission. In: *ECAL 2011: Proceedings of the Eleventh European Conference on the Synthesis and Simulation of Living Systems*, MIT Press (2011) 101–108
5. Whiten, A., Van Schaik, C.P.: The evolution of animal cultures and social intelligence. *Philosophical Transactions of the Royal Society B: Biological Sciences* **362**(1480) (2007) 603–620
6. Reader, S.M., Biro, D.: Experimental identification of social learning in wild animals. *Learning & Behavior* **38**(3) (2010) 265–283
7. Kendal, J.R., Rendell, L., Pike, T.W., Laland, K.N.: Nine-spined sticklebacks deploy a hill-climbing social learning strategy. *Behavioral Ecology* **20**(2) (2009) 238–244
8. Galef, B.G.: Social learning and traditions in animals: evidence, definitions, and relationship to human culture. *Wiley Interdisciplinary Reviews: Cognitive Science* **3**(6) (2012) 581–592

9. Zwirner, E., Thornton, A.: Cognitive requirements of cumulative culture: teaching is useful but not essential. *Scientific Reports* **5** (2015)
10. Jolley, B.P., Borg, J.M., Channon, A.: Analysis of social learning strategies when discovering and maintaining behaviours inaccessible to incremental genetic evolution. In: *International Conference on Simulation of Adaptive Behavior*. (2016) 293–304
11. Whitehead, H., Richerson, P.J.: The evolution of conformist social learning can cause population collapse in realistically variable environments. *Evolution and Human Behavior* **30**(4) (2009) 261–273
12. Borg, J.M., Channon, A.: Testing the variability selection hypothesis - the adoption of social learning in increasingly variable environments. In: *ALIFE 13: The Thirteenth Conference on the Synthesis and Simulation of Living Systems*, MIT Press (2012) 317–324
13. Rendell, L., Fogarty, L., Hoppitt, W.J., Morgan, T.J., Webster, M.M., Laland, K.N.: Cognitive culture: theoretical and empirical insights into social learning strategies. *Trends in cognitive sciences* **15**(2) (2011) 68–76
14. van der Post, D.J., Franz, M., Laland, K.N.: Skill learning and the evolution of social learning mechanisms. *BMC Evolutionary Biology* **16**(1) (2016) 166
15. Noble, J., Todd, P.M.: Imitation or something simpler? modeling simple mechanisms for social information processing. In Nehaniv, C.L., Dautenhahn, K., eds.: *Imitation in animals and artifacts*. MIT Press (2002) 423–439
16. McNamara, J.M., Dall, S.R.: Information is a fitness enhancing resource. *Oikos* **119**(2) (2010) 231–236
17. Mitri, S., Floreano, D., Keller, L.: The evolution of information suppression in communicating robots with conflicting interests. *Proceedings of the National Academy of Sciences* **106**(37) (2009) 15786–15790
18. Watson, R.A., Szathmáry, E.: How can evolution learn? *Trends in ecology & evolution* **31**(2) (2016) 147–157
19. Floreano, D., Keller, L.: Evolution of adaptive behaviour in robots by means of darwinian selection. *PLoS Biol* **8**(1) (2010) e1000292
20. Channon, A.D., Damper, R.: The evolutionary emergence of socially intelligent agents. In Edmonds, B., Dautenhahn, K., eds.: *Socially Situated Intelligence: a workshop held at SAB'98*, University of Zurich Technical Report. (1998) 41–49
21. Danchin, É., Giraldeau, L.A., Valone, T.J., Wagner, R.H.: Public information: from nosy neighbors to cultural evolution. *Science* **305**(5683) (2004) 487–491
22. Yang, S.X., Meng, M.: An efficient neural network approach to dynamic robot motion planning. *Neural Networks* **13**(2) (2000) 143–148
23. Robinson, E., Ellis, T., Channon, A.: Neuroevolution of agents capable of reactive and deliberative behaviours in novel and dynamic environments. In: *Advances in Artificial Life*. Springer (2007) 345–354
24. Stanton, A., Channon, A.: Incremental neuroevolution of reactive and deliberative 3d agents. In: *ECAL 2015: Proceedings of the Thirteenth European Conference on the Synthesis and Simulation of Living Systems*, MIT Press (2015) 341–348
25. Acerbi, A., Marocco, D., Nolfi, S.: Social facilitation on the development of foraging behaviors in a population of autonomous robots. In: *European Conference on Artificial Life*. (2007) 625–634