

1 The effect of social information use without
2 learning on the evolution of social behavior

3

4

James M. Borg^{1*} Alastair Channon¹

5

¹School of Computing and Mathematics, Keele University, Staffordshire, ST5

6

5BG, UK

7

j.borg@keele.ac.uk, a.d.channon@keele.ac.uk

8

* corresponding author

9

Abstract

10 In a recent paper by Borg & Channon [6] it was shown that social
11 information alone, decoupled from any within-lifetime learning, can result
12 in improved performance on a food foraging task compared to when so-
13 cial information is unavailable. Here we assess whether access to social
14 information leads to significant behavioral differences both when access to
15 social information leads to improved performance on the task, and when it
16 does not; do any behaviors resulting from social information use, such as
17 movement and increased agent interaction, persist even when the ability to
18 discriminate between poisonous and non-poisonous food is no better than
19 when social information is unavailable? Using a neuroevolutionary arti-
20 ficial life simulation, here we show that social information use can lead to
21 the emergence of behaviors that differ from when social information is un-
22 available, and that these behaviors act as a promoter of agent interaction.

23 The results presented here suggest that the introduction of social infor-
24 mation is sufficient, even when decoupled from within-lifetime learning,
25 for the emergence of pro-social behaviors. We believe this work to be the
26 first use of an artificial evolutionary system to explore the behavioural
27 consequences of social information use in the absence of within-lifetime
28 learning.

29 **Keywords:** social information; social behavior; local enhancement;
30 agent interaction; behavioral persistence

32 1 Introduction

33 The idea that agents may be socially attracted to each other by way of actively
34 seeking each other out in order to benefit from the proximity of others, be it
35 to avoid predators, breed or cooperatively raise their young, or to discover new
36 resources or habitats is a well established one [1, 2]. However, it is difficult to
37 establish precisely why and when social information leads to increased social
38 interaction and pro-social behavior; social information here being defined as in-
39 formation derived from the behaviors, actions, cues or signals of other agents
40 [23]. As social information necessarily involves the direct or indirect broad-
41 casting of information in to the public domain, it is sometimes known as (or
42 conflated with) public information [5]. Here we will use term social information
43 when describing any information about an individual which is broadcast in to
44 the public domain.

45 1.1 General Hypotheses for the emergence of social inter- 46 action promoting behavior

47 In reviewing social information use, Valone [36] outlines three general hypotheses
48 to explain why individuals might prefer to settle near conspecifics (leading to
49 what may be described as habitat copying via local enhancement):

- 50 1. Individual fitness is enhanced via the Allee effect [1, 2, 32]; which is defined
51 by Stephens et al. [34] as “a positive relationship between any component
52 of individual fitness and either numbers or density of conspecifics”. Allee
53 observed that individuals were better able to survive and reproduce when
54 found in groups, concluding that there is a positive correlation between

55 population density or group size and individual fitness (known as the Allee
56 effect). If this effect holds true we would expect there to be selection pres-
57 sure in favor of agents being in close proximity to one another; increased
58 use of social information may therefore be as a result of increased social
59 interaction due to agent proximity.

60 2. Social information based resource discovery results in a reduction in search
61 costs, enabling a more efficient use of energy [32, 17]. As social information
62 may be used to reduce search costs, and increase the chance of experienc-
63 ing new resources which may have been otherwise overlooked, increased
64 agent interaction may result from a selective pressure to obtain social in-
65 formation rather than increased social information use being a secondary
66 consequence of increased agent interaction itself; the Allee effect resulting
67 as a consequence of this selective pressure to access social information.

68 3. Individuals use the presence of other (established) individuals as an indi-
69 cator of the high-quality of a habitat without necessarily requiring them to
70 rely on their own (possibly incomplete or poor) evaluation of the habitat
71 [35, 39]. Here social information not only reduces the search costs when
72 discovering resources, but also enables individuals to derive the quality of
73 a unfamiliar resource based on social information about the action, state
74 or presence of others. Again, agent interaction and the Allee effect result
75 as a consequence of selective pressures in favor of social information use,
76 rather than social information use resulting as a consequence of a selec-
77 tive pressure in favor of agent interaction. This hypothesis is similar to
78 hypothesis 2 (listed above), but differs subtly; hypothesis 2 is associated
79 with simply discovering resources, whereas this hypothesis is associated
80 with judging the quality of a resource once found. This hypothesis may
81 be a direct result of hypothesis 2.

82 1.2 Behavior in the presence of social information

83 Here we assess three questions regarding agent behavior in the presence of so-
84 cial information. Firstly, we assess whether the well established notion that
85 social information leads to behaviors that promote agent interaction is true in
86 simple artificial evolutionary systems such as the one used by Borg & Channon
87 [6]. Secondly, we assess whether agent private information reliability (or envi-
88 ronmental predictability) impacts on agent interaction and social information
89 use. Finally, we assess whether any observed social behaviors (i.e. behaviors
90 resulting from the use of social information) can be seen to persist even when
91 social information use does not lead to an improved task performance - that is
92 to say when agents with access to social information no longer perform better
93 than agents with no access to social information on a simple food foraging task,
94 where performance is measured by the proportion of eating activity dedicated
95 to consuming “positive” foods compared to “negative” foods.

96 The question of the persistence of what may be described as non-adaptive
97 social information use, or social learning, was addressed by Higgs [20] in his
98 meme-based simulation study of learning by imitation. One of the many things
99 Higgs concluded was that memes (discrete, replicating, units of “culture” [12, 13,
100 3, 4]) even when they provided a negative biological fitness, still led to imitation
101 evolving. In Higgs’ model individuals had both a biological and a cultural
102 fitness. Both of these fitness values were determined by the set of memes held
103 by an individual, with reproduction being determined by biological fitness, and
104 the chance of being imitated being determined by cultural fitness. In one of
105 Higgs’ test cases the biological fitness provided by a meme was the reverse of
106 the cultural fitness, resulting in all biologically fit memes being culturally unfit
107 and all culturally fit memes being biologically unfit - even in this test case
108 imitative learning evolved. This suggests that behavior which increases social

109 interactions may still be adaptive even when task performance is poor.

110 Higgs' [20] result is not necessarily that surprising, as it is more than reason-
111 able to expect to see agents with access to social information of any kind seek-
112 ing this information out regardless of the contribution this information makes
113 to fitness, provided some of the social information could provide an adaptive
114 advantage. Bullinaria [10] rationalises this expectation by stating that "If there
115 exists a set of memes with a range of positive and negative contributions to the
116 overall performance, then not imitating them will leave performance at some
117 baseline, while imitating them will result in a range of performance levels above
118 and below that baseline. Any selection on the basis of performance will then
119 favour those individuals that have imitated the good memes, and hence favour
120 higher imitation rates" - therefore we can see why agents may wish to collect
121 around sources of information; sometimes that information will be useful, so
122 gaining access to it is important. We would therefore expect to see agents at-
123 tempting to find sources of information even when obtaining that information
124 does not necessarily lead to an improved performance. Agent and social inter-
125 action for the purpose of habitat copying is also found to be adaptive in highly
126 variable environments [38], though with the potential pitfall of population col-
127 lapse during overly conformist social interaction [40, 8]. It has also been noted
128 by Rendell et al. [28] that strategies that rely heavily on social learning seem
129 to be remarkably successful, even when information obtained from non social
130 sources is no more costly than social information. We would therefore expect
131 behaviors that maximize access to social information to emerge.

132 In the model set-up developed by Borg & Channon [6], which forms the basis
133 for this work, there are a large number of possible food resources available to
134 agents, resulting in agents often being uncertain about whether any given food
135 resource will provide a positive or negative amount of energy. As environments

136 in the Borg & Channon [6] model become more difficult, a strategy whereby all
137 food is ignored may evolve, but this strategy would always be outperformed by
138 a strategy that sought to minimise uncertainty about available food resources
139 in order to discover a positive energy providing resources. Social information,
140 especially about the performance or fitness of an agent, may therefore be sought
141 in order to allow for decisions on whether to consume any given food resource
142 to be influenced by others, thus reducing uncertainty about the safety of a new
143 food resource. This kind of social information seeking behavior in order to seek
144 out information about new or novel food resources is often seen in Norway rats
145 [15, 27], though it is interesting to note that this social behavior is only used
146 to develop food preferences and not food aversions; this property of rat social
147 behavior has been suggested to be as a result of the high levels of lethality
148 associated with poor food choices in rat populations [27], thus resulting in very
149 little social information about negative food resources being available to the
150 population. We may see a similar scenario in the more difficult environments
151 presented here, providing a continued pressure for social behavior under extreme
152 environmental difficulty. van Bergen et al. [37] reports that when individually
153 learned information is less reliable, nine-spined stickleback fish tend to use social
154 rather than individually learned information, this could also be re-phrased as
155 social learning is more likely to take place when a task is difficult to individually
156 learn. Therefore, it is not unreasonable here to expect agents in populations who
157 have access to social information to seek this information out in order to reduce
158 the unreliability of their own internal models of the world; it is far easier to
159 evolve prestige based social strategies such as “trust older individuals” or “trust
160 successful individuals” [19, 25], or conformist social strategies such as “trust
161 the majority” [18, 25], than evolve a rule about each possible food resource or
162 situation one may experience, especially when it is likely that any given food

163 resource or situation is new to an agent and therefore is yet to be evaluated.

164 **1.3 Previous work: the EnVar model**

165 The work discussed here follows on from previous work by Borg & Channon
166 [6]. In the Borg & Channon [6] work an artificial life model, called EnVar,
167 was created to investigate the evolutionary adaptation to social information
168 use without learning. The question posed by Borg & Channon [6] was “does
169 the addition of social information enable agents to evolve to perform better
170 on a simple food foraging task than when social information is not available”.
171 The EnVar model places a population of agents in a 2D simulated environment
172 containing a large variety of food/plant resources. Food resources are recognised
173 by agents by their color (RGB values), with food grouped in to species of plants
174 based on their color. Some of the plant species provided positive energy when
175 consumed, other provided negative energy when consumed. The simplest task
176 tested involved two food species, with a 1 : 1 ratio of positive to negative
177 food species, the most difficult task involved ten food species with a 1 : 9
178 ratio of positive to negative food species. A series of different populations with
179 access to differing types of social information were tested, with performance
180 on the task being measured by the how much time agents spent consuming
181 positive food resources compared to the how much time agents spent consuming
182 negative food resources. All agents had a limited amount of energy which was
183 lost through eating negative food and re-gained through eating positive food.
184 Residual amounts of energy were also lost when agents simply did nothing or
185 when they were moving, with energy lost due to movement being greater than
186 energy lost due to waiting. Agents were replaced when they ran out of energy,
187 with the replacement agents being the progeny of two surviving agents from the
188 population.

189 Each population of agents was set-up to use one of five social information
190 strategies. One of these strategies involved no social information at all,
191 whereas the other four involved social information either about the activity or
192 state of other agents. The two social information strategies associated with
193 activity were “presence”, where agents could only see whether another agent
194 was present or not, and “action”, where agents could see whether another agent
195 was eating, waiting, or moving (the only three actions available to agents in the
196 model). The two social information strategies associated with agent state were
197 “health”, where agents could see the battery level of another agent, and “age”,
198 where agents could see how old others agents were.

199 As expected, absolute performance on the task dropped with environmental
200 difficulty in all cases. However, populations using social information did out-
201 perform non social populations on simpler environments, thus demonstrating an
202 evolutionary advantage to using social information. In some cases, social infor-
203 mation also enabled populations to maintain a positive task performance across
204 a wider range of environments; the best social information strategy observed
205 was able to perform well (eating more positive food than negative food) up to
206 an environment with five negative food species to one positive food species; non
207 social populations were only able to achieve a positive task performance until
208 an environment with an 1 : 3 ratio of positive to negative food species. Despite
209 populations making use of social information generally outperforming non social
210 populations, there was no significant difference between social and non social
211 population in the more difficult environments that were tested.

212 The model used in this work, including the social information strategies
213 implemented, exactly matches the Borg & Channon [6] model (more details on
214 this model can be found later in the paper). Therefore, all hypotheses should
215 be considered in the context of the Borg & Channon [6] model and results.

216 1.4 Hypotheses

217 The large amount of evidence to suggest the persistence of social information
218 promoting behaviors in unreliable and challenging environments, and evidence
219 from simulations that social learning mechanisms such as imitation provide a
220 selective advantage even when the information being obtained is not necessarily
221 fitness increasing, along with the well established principle that the desire to
222 obtain social information leads to agent interaction, leads us to postulate the
223 following hypotheses to be assessed to here.

224 1. Social Information should lead to behaviors that result in increased agent
225 interaction (i.e. movement to seek to social interactions): We will test this
226 hypothesis by comparing the amount of movement undertaken by agents
227 from social information using populations with non social agents. If we
228 see a significant difference in the amount of movement, we will then assess
229 how often agents from social populations spend around other agents. We
230 require a significantly larger number of movement actions combined with
231 agent interaction to demonstrate not only socially influenced interaction,
232 but also behaviors that promote social interactions. Sergio & Newton [31]
233 provides evidence that in some cases even simple information such as the
234 presence of other individuals (or occupancy) can be a suitable indicator of
235 resource quality and therefore enough to lead to agents coalescing around
236 a food source, therefore we would expect this hypothesis to hold true in
237 all social information strategies presented here; though when the presence
238 of another agent is used as a source of social information, some measure of
239 resource quality may still be required, as no information about the success
240 or state of the agent present on the resource is available to act as a proxy
241 for resource quality [36].

242 2. Social interaction between agents will be more likely when environments

243 are more unpredictable, and less likely when environments are more pre-
244 dictable. In the model environment used here it could be argued that
245 the more difficult environments are more predictable. The most difficult
246 environment tested here has a ratio of one positive food resource to every
247 nine negative food resources, therefore agents have a 90% chance of cor-
248 rectly guessing that a food resource will be dangerous. We may therefore
249 expect agent social interaction (should it be seen) to be at its highest in
250 lower difficulty environments, despite the possibility of non social agents
251 also performing well in these environments. From an artificial life and
252 evolutionary robotics perspective it would be useful to know under which
253 conditions pro-social behaviors, such as agent social interaction and co-
254 operative foraging, may emerge.

255 3. Behaviors resulting in increased agent interaction will persist (though at
256 reduced levels) even when task performance is poor, poor task perfor-
257 mance being characterised by agents spending more time eating negative
258 food than eating positive food: The adaptive value of social information,
259 even when potentially unreliable, should still be high enough to motivate
260 agents to seek others out more often than if social information was not
261 available. In the more difficult environments tested here we would ex-
262 pect social information to be relatively poor, due to the large quantities
263 of negative food resources populating the environment. However, it would
264 still be beneficial for agents to engage in movement sometimes in order to
265 provide potential access to any positive behaviors that may emerge in the
266 population. Therefore we would expect behaviors that encourage social
267 interaction, i.e. movement, to still appear more often in social populations
268 than in non-social ones, in all environments.

269 We will also go on to to assess whether social information leads to any significant

270 difference in the application of the other behaviors available to agents here
271 when compared to non-social populations, and whether task performance has
272 any implications for the application of behavior - we are especially interested to
273 assess whether a change in task performance from the predominantly successful
274 application of eat actions to the predominately unsuccessful application of eat
275 actions is accompanied by any notable transitions in behavior.

276 **2 Simulation Model and Experimentation**

277 The experimental set-up matches that used in Borg & Channon [6]. Summary
278 tables of the key parameters used can be found in the Appendix section at the
279 end of the paper.

280 Populations of neuroevolutionary agents (making use of the hybrid neural
281 network model known as the shunting model [43, 42, 30, 7, 33, 22]), each popula-
282 tion employing a different social information strategy, are tasked with surviving
283 in environments of differing difficulties. In order to test our hypotheses we test
284 populations of social and non-social agents in a set of increasingly difficult envi-
285 ronments; forty populations of each social information strategy being evaluated
286 per environment. Environmental difficulty is dictated by the ratio of positive
287 food resources to negative food resources. The simplest world used here has an
288 equal (1 : 1) ratio of positive food species to negative food species. Tests get
289 progressively harder by increasing the number of negative food species, whilst
290 maintaining only one positive food species, resulting in the most difficult world
291 used here having a 1 : 9 ratio of positive food species to negative food species All
292 data presented here relates to the final 25 epochs of evolution (of a total of 100
293 epochs) where population behavior and fitness had broadly stabilised (based on
294 the results of [6]). An epoch here is defined as 1000 time-steps, with a time-step
295 being defined as one full simulation loop.

296 The task world used here is known as EnVar. EnVar is a bounded (non-
297 toroidal) 2D environment containing a variety of consumable resources known
298 as plants. Plants are recognised by agents simply as an RGB value. Plants are
299 divided into a number of species, each with a randomly selected base RGB value.
300 Plants are generated within these RGB regions and identified as belonging to
301 the nearest species according to euclidean distance in RGB space to a species
302 base RGB value. The number of plant species is determined by the test being
303 conducted. In the tests conducted here, the number of species ranges from two
304 to ten. Each plant species is assigned an energy value, which is transferred to
305 agents if the plant of that species is consumed; energy values may be positive or
306 negative. Notionally the EnVar world is broken up in to cells, though here each
307 cell represents a pixel. Plants in the world take up a number of cells, set here to
308 100 pixels/cells, forming a 10x10 block, with each block only being able to be
309 eaten a certain number of times before being exhausted (here set to be 200 eating
310 events). Once a plant block has been exhausted it is no longer consumable and
311 therefore removed from the world to be replaced by a new block from a random
312 plant species somewhere else in the world - this maintains a constant number
313 of food blocks in the world at any time. Agents are permitted to share space
314 with a plant resource but cannot overlap with each other, thus removing the
315 possibility of agents piling up on top of one another on valuable food resources
316 - this can result in an agent's path to a food resource being blocked by agents
317 already on that resource, though agents cannot intentionally choose to block
318 other agents. For all tests here negative food species come with an energy
319 value $E_{neg} = -10.0$, with positive food species contributing an energy value of
320 $E_{pos} = 1.0$ when consumed. This provides a strong evolutionary pressure to
321 avoid eating negative food species. In this work EnVar is set up to create a
322 700×700 pixel sized cell world, containing five hundred 10×10 pixel blocks of

323 plants.

324 **2.1 Neuroevolutionary Model**

325 Agents in the EnVar simulation world are grounded 2D simulated agents, con-
326 trolled by a hybrid neural network architecture known as the *Shunting Model*
327 [43, 42]. The shunting model uses two interacting networks to determine agent
328 behaviors, here represented as a discrete set of agent actions. The two interact-
329 ing networks are known as the *Decision Network* and the *Shunting Network*. The
330 decision network is simply a feed-forward neural network comprised of an input
331 layer, one hidden layer and an output layer. Outputs from the decision network
332 (known as Iota values) are used to produce a locally-connected, topologically-
333 organised network of neurons known as the shunting network, which simply
334 places and organises agent preferences for environmental features and states in
335 such a way to allow the agent to hill climb in a shunting space (known as the ac-
336 tivity landscape) that directly maps on to their immediate neighborhood. The
337 shunting network weights are fixed for all agents, whereas the decision network
338 is genetically encoded and is subject to change via evolution.

339 **2.1.1 The Shunting Network**

340 The shunting network is a locally-connected, topologically-organised network
341 of neurons that was originally used for collision free motion planning in robots
342 [43, 42] and has been subsequently applied in a number of 2D and 3D artificial
343 life models [30, 7, 33, 22, 6]. Here the shunting network’s topology is simply
344 superimposed on to the environment, with each cell in the network topology
345 directly relating to a pixel within an agent’s visual field. Using a simplified and
346 stable version shunting equation developed by Stanton & Channon [33] (see
347 equation 1) values for each cell (which can be interpreted as representing an

348 environmental feature or state, and are initially set by the Iota output I ob-
 349 tained from the decision network) are propagated across the cells of the network,
 350 producing an activity landscape with peaks and valleys representing desirable
 351 and undesirable features in the environment. The result is a landscape which
 352 allows the agent to follow a route determined by the higher Iota values while
 353 avoiding undesirable valleys. A mock-up example of an activity landscape with
 354 a snapshot of the visual field it represents can be seen in Figure 1.

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

355 In equation 1 each node in the shunting network corresponds to one pixel
 356 within an agent’s visual field; x_i is the activation of neuron i ; N_i in the receptive
 357 field of i ; the function $[x]^+$ is $\max(0, x)$; and I_i is the external input to neuron
 358 i (the Iota value). The maximum Iota value is $\max I = 15$, with the resulting
 359 value for x_i^{new} also being capped at a minimum Iota value $\min I = -15$. This
 360 stops Iota values growing out of control, whilst providing a large enough maxi-
 361 mum value (and a small enough minimum value) to ensure activity propagation
 362 across the network. In order to allow propagation to occur within a time-step,
 363 the shunting equation must be run a number of times, we take this number of
 364 iterations to be equal to the diameter of the visual field.

365 The shunting model implemented here differs in a number of significant ways
 366 from previous artificial life implementations [30, 7, 33, 22]. In these previous
 367 implementations agents see their entire environment, have a set number of dis-
 368 crete environmental features and states to set Iota values for, and are in the
 369 environment alone to complete a predetermined task. Here agents have a lim-
 370 ited view of the world, have the possibility of needing to set an Iota value for
 371 a plant of any given RGB value, and exist as a population within the environ-
 372 ment (leading to possible input states where an agent can be seen on a particular

373 plant). In order to accommodate these differences the shunting model here is
374 run independently for each pixel in an agent’s visual field, which is set here to
375 have a radius of 30 pixels from the center of the agent, with information about
376 that pixel being included as part of the agent’s decision network input layer. In
377 this way an Iota value is calculated for each unique environmental state within
378 an agent’s visual field. This change does not change the resulting behavior of
379 the shunting model or activity landscape, just the way in which information is
380 passed to the shunting network from the decision network.

381 **2.1.2 The Decision Network, Neuroevolution and Reproduction**

382 Evolution in the model is applied only to the decision network. Here, the deci-
383 sion network is a feed-forward neural network comprised of seven input nodes,
384 and an additional social input node in social information tests, eight hidden
385 units, and two output nodes, resulting in 112 - 128 weights. Each network layer
386 is fully connected, with floating point weights in the range $[-1 : 1]$ being directly
387 encoded from an agent’s genotype. A standard sigmoid activation function is
388 used at each hidden and output node, though outputs processed for deriving
389 agent actions are then scaled to be within the range $[0 : 1]$ and the Iota out-
390 put is scaled linearly to be within the range $[minI : maxI]$. As the agent is
391 expected to produce an Iota value to feed in to the shunting network for each
392 unique environmental feature or state within its visual field, inputs into the
393 decision network must accommodate both the internal state of the agent, the
394 state of their current environment, and the state of the environmental feature
395 they are assessing; this leads to there being two sets of input nodes. The first
396 set of input nodes are simply plant RGB inputs - if the agent is viewing empty
397 space these inputs are set to -1, else they are set to be the normalised RGB of
398 the plant being viewed, with RGB values being normalised to be within the range
399 $[0 : 1]$ by way of linear normalisation. Following these inputs are a series of

400 generic inputs, which are dependent on the agent’s internal state and the cur-
401 rent environmental state. These inputs are the agent’s current battery level in
402 the normalised range $[0 : 1]$, a moving average of the agent’s battery level over
403 the previous 100 time steps, the agent’s current external environmental state
404 and a moving average environmental state, which are both set to be +1 and
405 do not change in the tests presented here (the model is set-up to accommodate
406 external environmental change which is not used here). In social information
407 tests agents have an additional input based on the agent being viewed.

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

427 The new agent, or child, created to replace the removed agent is the progeny
428 of two agents, one of whom is selected in a tournament, the other of which is
429 selected randomly from the remaining population. The tournament selection
430 mechanism applied here takes two agents from the population, compares their
431 current energy levels, and selects the fitter agent (i.e. the agent with the higher
432 energy level) as a parent. Like in nature, this isn't a perfect measure of fitness
433 as it is possible the agent is young and therefore has not yet had time to lose
434 significant amounts of energy, or the agent could have simply been lucky or
435 unlucky with available food sources. However, in general, agents with more
436 effective behaviors will on average find themselves with better energy levels than
437 agents with less effective behaviors, thus driving evolution toward behaviors
438 that are more suited to the task or environment at hand. The second parent
439 is selected randomly to ensure the population doesn't become dominated by
440 the progeny of a small sub-set of the population, thus maintaining a level of
441 exploration in the genotypic search space. New agents are placed in the world
442 within the visual field of one of their parents, selected at random - this does
443 place agents within close proximity of each other without the need for agents
444 to explore, providing a pressure against the evolution exploratory movement to
445 seek out other agents.

446 **2.2 Agent Actions and Action Energy Costs**

447 The agents in the model have a set of simple, discrete, actions available to
448 them, through the output layer of their decision networks: wait, eat or move.
449 The decision network has two outputs, an Iota output to be fed into the shunting
450 network and an eat/wait output. The agent first considers the input state at
451 its current position - if the agent produces an Iota value above the threshold
452 $\theta_a = 0.5$ it indicates the agent is happy with its current state and position and

453 therefore does not move. The agent's eat/wait output is then considered; if the
454 output produces a value above the threshold $\theta_b = 0.5$ the agent attempts to eat
455 whatever may be at its current position; agents are welcome to try and eat at
456 locations where no plant is present, but no benefit for this action is conferred,
457 and the eat action is considered to be an unsuccessful eating attempt rather
458 than a wait action. If an agent decided to eat at a location containing a plant,
459 the plant's energy is transferred to the agent, this does not necessarily lead to
460 the exhaustion of the plant resource. The Iota output is in the range $[-1 : 1]$,
461 any values in the range $[-\theta_a : \theta_a]$ are evaluated as neutral and resolve to 0.
462 The Iota output is then scaled to be within the range $[minI : maxI]$ for use
463 in the shunting network, whereas the eat/wait output is limited to the range
464 $[0 : 1]$. If the eat/wait output gives an output below the expected threshold the
465 agent simply waits at its current location. Waiting and eating both reduce an
466 agent's energy by 0.1 energy units (though eating may result in a net energy
467 gain), with moving using up 0.2 energy units per time step. Agents will only
468 move if their Iota output for their current location is below threshold θ_a . In
469 this case an activity landscape is created based on the Iota outputs for all
470 visible environmental features. Agents are born with, and are able to achieve, a
471 maximum energy level of 100 units. As epochs here constitute 1000 time steps,
472 an agent would be able to survive for a maximum of one epoch, or one thousand
473 time steps, by remaining inactive. In order to avoid agents moving around in
474 circles, or moving backwards and forwards, in neutral space where there is no
475 activity gradient from the activity landscape, consecutive neutral move actions
476 maintain the same direction of travel with probability $p_{dir} = 0.9$.

477 2.3 Social Information Strategies

478 The social information strategies explored here, including the no social strategy
479 are discussed below:

480 NO SOCIAL: No input node is available to the agent to enable social in-
481 formation to be used by the agent’s decision network. Agents proceed with
482 no information about other agents. There is very little evidence in nature for
483 agents being totally ignorant of the presence of other agents - this strategy was
484 simply to be used a baseline to compare the other social information strategies
485 against.

486 PRESENCE: The social information input node receives an input of +1 if
487 another agent is present within the visual field. No other information about the
488 agent being viewed is used. This strategy is not dissimilar to the “Inadvertent
489 Information” strategy used by agents in the work by Mitri et al. [26], though
490 the agents explored in the work presented here do not have a choice about
491 whether they express social information or not. In nature the presence of other
492 agents has been established as key motivator of where to eat or explore in a
493 number of vertebrates [16]. Social facilitation, defined as the mere presence of a
494 demonstrator affecting an observing agent’s behavior [21, 29], is an example of
495 a social learning strategy observed in nature arising the mere presence of other
496 agents.

497 ACTION: An input representing the current action state of the agent being
498 viewed. The wait action is input as a value of 0, eat is input as 0.5 and move
499 is represented as 1. Amalgamating these action inputs into one input rather
500 than two or three categorical inputs, whilst not ideal, was implemented in order
501 to ensure the input layer size for all social strategies was equal. Being able to
502 observe and interpret the activity or actions of other agents can lead to a variety
503 of social learning strategies seen in nature - these strategies include observational

504 conditioning, social enhancement, response facilitation and contextual imitation
505 [21, 29].

506 HEALTH: The current energy levels of the agent being viewed are normalised
507 to be within the range [0 : 1] and input to the viewing agent’s decision network.
508 Health information here is used a possible proxy for the success of agents, though
509 a noisy one as high energy levels could indicate that the agent is young (and yet
510 to expend any energy) or lucky, alongside indicating that an agent has evolved a
511 suite of adaptive behaviors that minimises energy use and maximises successful
512 eating events. The social learning strategy “copy successful individuals” is seen
513 regularly in nature [25], and is well established in theoretical modeling as viable
514 social learning strategy [9].

515 AGE: The age (in time steps) of agent being viewed is normalised using a
516 hyperbolic tangent function of the logarithm of the age, which is then normalised
517 to be within the range [0 : 1] (with 1 being asymptotic). Normalising age in this
518 way is necessary as agents may live for the entire duration of the simulation,
519 and are not selected against based upon their age. See formula (2) where a
520 represents agent age in time steps. Using information about the age of other
521 agents can result in a “copy older individual” social learning strategy [25], with
522 such strategies being observed in mate choice copying in fish [14, 24]. As avoiding
523 being removed from the population is also an indication of successful behavior,
524 copying individuals can also be seen as another form of the “copy successful
525 individuals” strategy.

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

526 It is worth noting that despite references to social learning, this work contains
527 no learning, therefore we would not expect complex social strategies such as
528 those seen in nature to emerge here. All references to social learning in nature

529 here are instead supposed to justify why the social information being used here
530 may be justified as forming the basis of more complex social learning strategies
531 seen in nature.

532 **3 Results**

533 **3.1 Action Profiles**

534 Figure 2 shows the median action profiles for each social information strategy
535 applied here, an action profile being the proportion of total actions each indi-
536 vidual action contributed. The most immediate difference between the social
537 information using populations and non social populations from Figure 2 is the
538 application of the move action. Whilst all populations show a reduction in move-
539 ment (as environmental difficulty increases), with an accompanied increase in
540 waiting, non social populations have extremely low levels of movement even in
541 environments of lower difficulty when compared to social information popula-
542 tions. In social populations movement is applied more frequently than waiting
543 in lower difficulty environments. This suggests that the increased performance
544 associated with populations that use social information in simpler environments
545 seen previously [6] is as a consequence of this greater willingness to move, either
546 to find new food resources or to find new sources of social information. As the
547 only difference between social and non social populations is the addition of social
548 inputs to agent neural networks, movement to seek new sources of information
549 is probably closer to the truth; as agents in all populations spend the majority
550 of their time in simpler environments eating, any movement motivated by the
551 desire to be around other agents would lead to a secondary consequence of being
552 around more food resources, enabling agents who are less able to distinguish be-
553 tween positive and negative food resources to defer some of their judgments on

554 the likely pay-off of a food resource, and instead rely on the social information
555 being provided by the agents they now find themselves around to make more
556 informed decisions. However it is not clear from Figure 2 whether or not this
557 difference in movement between non social and social populations is significant,
558 and whether this additional movement does lead to more opportunities for social
559 information use.

560 The immediate difference in movement behavior between non social and
561 social populations seen in Figure 2 is demonstrated to be significant by way
562 of Mann-Whitney U tests between the resulting application of move actions for
563 social populations compared to non social populations, this can be seen in Figure
564 3. The continued significance in the difference between social and non social
565 populations regarding movement is in contrast to the general lack of significance
566 in task performance difference between social and non social populations in
567 environments past environment 2 (as seen in [6]); these results indicate that the
568 introduction of social information leads to behavioral differences that persist
569 even when these behaviors do not result in improved task performance.

570 Regarding the other actions available to agents; eating (see Figure 4) and
571 waiting (see Figure 5), neither show any particular significant differences (where
572 $p < 0.01$) between social and non social population other than in environment
573 1 where waiting actions for all social populations are applied significantly less
574 than in non social populations ($p < 0.01$), and eating actions are applied sig-
575 nificantly less for social populations using the Presence and Action strategies
576 than in non social populations ($p < 0.01$). This broad lack of any significant
577 differences beyond environment 1, between non social and social populations for
578 eating and waiting, further demonstrates that movement is the primary driving
579 force in the improved task performance seen in earlier environments, especially
580 in environment 2 where only movement is significantly different despite previous

581 work [6] showing a significant difference in task performance; though it should
582 be noted that in environment 1 social information availability also leads to sig-
583 nificantly different eating and waiting behaviors, indicating that some adaptive
584 action profile across actions is available to drive improved task performance,
585 rather than just a reliance on movement behavior. The fact that in environ-
586 ment 1 differences in eat and wait actions result in less eating and waiting taking
587 place in social populations in favor of more movement, also indicates that social
588 agents are willing to risk higher energy expenditure, and are willing to spend
589 less time potentially obtaining energy via eating. This demonstrates that the
590 accommodation of social information leads to a more refined, and ultimately
591 more effective, eating strategy as a result of an increased willingness to move.
592 However, as we can see from the action profile box-plots in Figure 6, the ap-
593 plication of eating and waiting actions is drawn from quite a large range in all
594 populations, though the interquartile ranges for all actions do indicate some
595 level of consistency in the application of actions in environment 1.

596 The suggestion here is that the significant improvement in task performance
597 seen in social populations over non social populations in less difficult environ-
598 ments (as in [6]) is as a direct result of the behavior differences enabled by the
599 accommodation of social information. However, this does lead us to something
600 of a “Chicken and Egg” situation; did social information use follow as a result
601 of good foraging (with good foragers acting as useful sources of social informa-
602 tion), or did social information use result in the development of good foraging
603 strategies? As no information about plant resources are communicated by so-
604 cial agents, with only information about the agents themselves being expressed,
605 it would be sensible to assume that the improved task performance seen by
606 social populations in simpler environments is caused by agents developing be-
607 haviors that cause greater exposure to other agents (and therefore more sources

608 of social information), which then leads to an improve task performance as a
609 secondary outcome. The fact movement behavior remains significantly different
610 throughout all tests indicates that some behavioral differences persist despite
611 them providing no improvement in task performance.

612 **3.2 Reasons for Moving**

613 It is apparent from Figure 3 that movement behavior for populations permitted
614 to use social information differs significantly from non social populations - this
615 is in contrast to both eating actions (see Figure 4) and waiting actions (see
616 Figure 5) which only show significant differences between social and non social
617 populations in selected environments. Therefore some analysis on why social
618 agents move is necessary.

619 One possible indication that increased movement is a direct consequence
620 of an increased motivation for agents to interact, and arguably the clearest
621 demonstration of the Allee effect [1, 2], would be if agents were found to aggre-
622 gate/cluster (i.e. herd or shoal). Figure 7 shows the distributions of the size
623 of agent clusters for each social strategy compared to non social populations,
624 with cluster size simply being the number of other agents an agent has within
625 its visual field. Figure 7 demonstrates an increase in cluster size as environ-
626 mental difficulty increases, but no clear or significant difference in cluster size
627 between social and non social populations is observed. The increase in clus-
628 ter size as environmental difficulty increases is explainable as a consequence of
629 the increased waiting exhibited by all populations; agents move less in difficult
630 environments, resulting in new agents being less likely to move away from the
631 parent agents they are placed close to following a reproduction event. The lack
632 of significant difference in cluster size would likely be a result of moving agents
633 regularly encountering other moving agents due to the density of agents, and

634 agents clustering around good food resources. Therefore, this result is not to-
635 tally surprising and thus leads to the conclusion that the increased movement
636 seen in social populations does not lead to higher levels of aggregation or clus-
637 tering.

638 The fact that agents in social populations can actually view other agents
639 enables a second level of analysis regarding why social agents might be moti-
640 vated to move. When an agent decides to move, as opposed to wait or eat, they
641 evaluate their preference for each pixel/cell within their visual field. If a cell
642 contains another agent (either alone, or on a food resource) then a social agent
643 can register an *agent view*. Non social populations are blind to other agents, and
644 therefore unable to register agent views. Agent views can be positive (result-
645 ing in attraction), negative (resulting in repulsion), or neutral (ultimately not
646 affecting movement behavior). Should non neutral (positive or negative) agent
647 views be registered, we can conclude that social information is being actively
648 used by agents when moving.

649 Figure 8 shows the distributions of the number of neutral and non neutral
650 agent views accumulated by individuals in social populations. It is clear from
651 Figure 8 that for most social strategies, in most contexts, the social informa-
652 tion provided by the proximity other agents is considered to be of little use, and
653 therefore does not affect movement decisions. But it is key to note that for every
654 social strategy in all environments (barring environment 9 for Age populations),
655 some non neutral agent views are registered. Sometimes social information is
656 useful, and is therefore used to influence agent behavior. However, the distribu-
657 tion of agent views for populations using Age information (Figure 8(d)) stands
658 out; unlike the other social strategies, agent views are split relatively evenly
659 between neutral and non neutral activity. This indicates that some forms of
660 social information can often be useful, and thus worth seeking out. Consider-

661 ing how each social strategy was operationalized in the model, it is clear why
662 information about the age of other individuals might prove to be more useful
663 than other types of social information; age is the only unambiguous indicator of
664 success explored here. Presence can help agents in simpler environments decide
665 whether or not to move over to a food resource, but it is unlikely to promote
666 general exploratory movement. Action provides more information, but without
667 additional social information it is difficult for agents to determine whether an
668 action (especially eating) is being applied by a successful or reliable individual.
669 Health is better at indicating success, and therefore a more reliable source of
670 information, but is still noisy; young agents are born with full energy levels and
671 some agents can just be lucky when eating. Age is unambiguous; older agents
672 (especially those who have lived beyond a few epochs) can only have done so
673 by being successful at the task. We see that by the most difficult environment,
674 information about the age of others is not often accumulated, and is never used
675 - this is as a result of the environment being so challenging that agents rarely
676 live very long.

677 Figure 9 assesses whether the non neutral views accumulated by social pop-
678 ulations is perceived to be positive or negative. Whilst being measured on
679 drastically different scales across social strategies, we do see a shift from largely
680 positive agent views in Presence populations, through to largely negative agent
681 views in Health and Age populations, Action populations demonstrate little
682 preference either way. Whilst not analysed here, these result do suggest that
683 populations with access to more reliable social information (Health and Age)
684 are able to be more discerning about whether they wish to move toward another
685 agent, whereas populations with only the presence of other agents available to
686 them have very little cause to be repulsed (agents cannot directly interpret the
687 density of agents in an area, and therefore cannot be disinclined to move to-

688 wards over-saturated resources). But the fact that for social populations, social
689 information can not only be non neutral but can also be attractive/positive,
690 does provide an explanation for why social populations choose to move more
691 often than non social populations, even when this doesn't necessarily result in
692 improved task performance; movement may either be as a result of attractive
693 additional stimuli (other agents) or from a motivation to move to seek out other
694 other agents, as opposed to just waiting.

695 **3.3 Behavioral Transitions**

696 From Figure 10 we can see that that non social populations do not exhibit any
697 statistically significant transitions ($p < 0.01$) between environments in regard to
698 movement behavior. However, statistically significant transitions in movement
699 behavior between environments can be seen in all social populations. For popu-
700 lations using Presence information we see this statistically significant transition
701 happen between environments 2 and 3; the transition from primarily eating
702 positive food resources to primarily eating negative food resources also occurs
703 between environments 2 and 3. The association between a statistically signifi-
704 cant transition in movement behavior and the transition to primarily consuming
705 negative food resources is also apparent for populations using Action informa-
706 tion and populations using Health information - for Health populations it is
707 also interesting to note that statistically significant movement behavioral tran-
708 sitions occur on both occasions when positive food consumption drops below
709 zero. These results demonstrate that movement behavior in social populations
710 is strongly driven by agent task performance; when agents can no longer suc-
711 cessfully solve the task, social populations are less inclined to explore their
712 environment in order to seek out new food resources or new sources of social
713 information. In the case of populations using Age social information, the only

714 significant transition associated with movement behavior occurs before the tran-
715 sition to non-positive food consumption. The point at which this transition in
716 movement behavior occurs does correspond with a large drop in task perfor-
717 mance between environments 2 and 3, demonstrating that movement behavior
718 is still highly sensitive to task performance in Age social information popula-
719 tions.

720 From Figure 2, and Figures 4, 3 and 5 we can see that agent behavior
721 changes as environments become more difficult. These behavioral changes lead
722 to a reduction in movement and eating, and an increase in waiting. The primary
723 driving force behind the motivation to eat less, move less and wait more, inde-
724 pendent of social information strategy, is that food resources are increasingly
725 likely to be negative in their energy provision, and therefore it makes sense for
726 agents to spend more time conserving their energy waiting for a positive food
727 source to appear near to them or (in the case of social populations) for an agent
728 who's information suggests they can be trusted to move into their visual field.
729 However, in most cases the increase or decrease in actions as environments be-
730 come more difficult is not necessarily smooth, this being most apparent with
731 move actions (Figure 3) which for many social information strategies shows a
732 sudden reduction in action rather than a steady degradation. It is not clear
733 from earlier figures whether these changes between environments are statisti-
734 cally significant nor what is driving these sudden changes when they occur.

735 In Borg & Channon [6] it what shown that task performance (the ability
736 to eat positive food resources more frequently than negative food resources)
737 deteriorates as environments get more difficult - this difficulty being defined
738 by the ratio of positive food resources to negative food resources available in
739 the environment. The point at which task performance changes from success-
740 ful to unsuccessful (the point at which eating actions result in more negative

741 food resources being consumed than positive food resources) varies depending
742 on the social information strategy being tested, but occurs in all scenarios. For
743 No Social and Presence populations this transition (or zero crossing) occurs
744 between environments 2 and 3, Action populations experience this transition
745 between environments 3 and 4, and both Health and Age populations experi-
746 ence this transition to primarily negative food resource consumption between
747 environments 4 and 5 (though Health populations do not permanently cross into
748 negative task performance until after environment 6). Here we assess whether
749 any statistically significant changes to behavior, or behavior transitions, could
750 be associated with these zero crossing events for food type consumption.

751 When considering the total proportion of actions agents dedicate to eating,
752 as seen in Figure 11, we do not see any significant changes in eating behavior
753 that correspond to the point at which task performance transition from predom-
754 inantly successful application of the eat action to predominantly unsuccessful
755 application of the eat action. Instead, as seen in Figure 4, the median total eat
756 action degrades gradually with task performance. It is also worth noting the
757 extremely large data ranges seen with the total application of each action in the
758 box plot data in Figure 11. The large interquartile ranges especially show that
759 all populations, social and non social, are capable of exhibiting very high and
760 very low levels of eating activity. This is in stark contrast to movement, which
761 we can see from Figure 10 has reasonably small interquartile ranges for all popu-
762 lation types across all environments, and if anything becomes more consistent as
763 environmental difficulty increases, this being in contrast to the general increase
764 in the range of eat action data which generally increases as the environment
765 becomes more difficult. Increasingly large data ranges are also seen when we
766 consider the wait action (as seen in Figure 12). Any significant transitions seen
767 in waiting behavior, in all populations barring Health, do not seem to occur

768 in relation to the transition from positive to negative task performance. These
769 results further indicate that social agents are driven to seek out new sources
770 of social information, but with the caveat that social interactions are likely to
771 result in better task performance; though the fact that social populations move
772 more often than non social populations even when task performance is poor
773 suggests that social populations still persist in a residual amount of socially
774 motivated movement.

775 4 Discussion and Conclusion

776 In this work we attempted to address three questions. (1) Does social infor-
777 mation lead to increased agent interaction? (2) Is agent interaction, and by
778 extension social information use, dependent on the environmental predictabil-
779 ity? (3) Do social behaviors persist even when task performance is poor?

780 Social information transfer is highly prevalent in nature [41], and even the
781 simple presence of other agents have been demonstrated to encourage inter-
782 esting and novel behaviors in other agents [11], so it is not entirely surprising
783 that the results presented in this work provide strong evidence that social in-
784 formation can lead to interaction promoting behaviors, namely movement for
785 the purpose of increasing the probability of agent interaction. We also see so-
786 cial behaviors being favored in the simpler environments tested here. These
787 simpler environments did provide agents with a large variety of food resources
788 that could be either negative or positive with an equal probability, resulting in
789 a task which was reasonably easy to solve but also very difficult for individuals
790 to develop a complete set of categorisations for each food resource's edibility.
791 Social behaviors being favored here are likely to be as a result of social informa-
792 tion being more reliable than private information. As environments progressed
793 in difficulty, private information about the edibility of any given food resource

794 became more reliable, as it was increasingly likely that any given food resource
795 was energy reducing and therefore not worth consuming. Any social interac-
796 tion in later, more difficult, environments would still have yielded some benefits
797 though. In the presence of a food resource in any environment the presence,
798 actions, health or age of other local agents could potentially result in a novel
799 or new food resource being evaluated correctly. Despite private information
800 based on the likelihood of edibility encouraging a conservative policy on eating,
801 this new social information could sometimes yield positive results leading to an
802 adaptive advantage over agents who eschew social interaction. Here we see a
803 continued preference for movement in social information populations compared
804 to non social populations, even in more difficult environments where task per-
805 formance in both social and non social populations was similar. This continued
806 desire to move for the purpose of social interaction was less apparent in later
807 environments, with waiting actions being preferred due to the risk of unneces-
808 sary or unrewarding energy expenditure in more difficult environments, but still
809 significantly different from non social cases.

810 The results presented here add additional evidence to the idea that a pres-
811 sure for evolution to adapt to accommodate social information, be it via social
812 information transfer or imitation, is maintained even when social information
813 is either unreliable or risky [20], and therefore suggest that the introduction of
814 simple social information is sufficient, even when decoupled from any within-
815 lifetime learning processes, for the emergence of pro-social behaviors.

816 Following on from this work, and the work of Borg & Channon [6], a num-
817 ber of additional tests are required to fully establish how social information
818 is affecting agent behavior and to what extent agent behavior is affected by
819 parameters such as the cost of movement, cost of stationarity, population den-
820 sity, proportion of unfit agents replaced at the end of each epoch, and food

821 density, persistence, and energy. As it currently stands it is difficult to fully
822 establish whether agents are attracted to certain actions, older individuals, or
823 healthy individuals - the work here simply establishes that the availability of
824 social information can elicit changes of behavior, with these behaviors acting as
825 promoters of agent interaction.

826 5 Acknowledgements

827 Much of this work was completed thanks to PhD funding from the School of
828 Computing and Mathematics, Keele University, UK. We would also like to thank
829 the three anonymous reviewers for their fair and critical comments, and their
830 sound and valuable advice, on earlier versions of this work; this work is un-
831 doubtedly stronger as a result of their comments.

832 References

- 833 [1] Allee, W. (1931). *Animal aggregations: a study in general Sociology*. Uni-
834 versity of Chicago Press Chicago:.
- 835 [2] Allee, W. C. (1939). *The social life of animals*. W. Heinemann.
- 836 [3] Blackmore, S. (1998). Imitation and the definition of a meme. *Journal of*
837 *Memetics*, 2(11), 159–170.
- 838 [4] Blackmore, S. J. (1999). *The meme machine*. Oxford University Press.
- 839 [5] Bonnie, K. E., & Earley, R. L. (2007). Expanding the scope for social
840 information use. *Animal Behaviour*, 74(2), 171–181.
- 841 [6] Borg, J. M., & Channon, A. (2017). Evolutionary adaptation to social
842 information use without learning. In G. Squillero, & K. Sim (Eds.) *European*

- 843 *Conference on the Applications of Evolutionary Computation*, (pp. 837–
844 852). Springer.
- 845 [7] Borg, J. M., Channon, A., & Day, C. (2011). Discovering and maintain-
846 ing behaviours inaccessible to incremental genetic evolution through tran-
847 scription errors and cultural transmission. In T. Lenaerts, M. Giacobini,
848 H. Bersini, P. Bourguine, M. Dorigo, & R. Doursat (Eds.) *Proceedings of the*
849 *European Conference on Artificial Life 2011*, (pp. 102–109). MIT Press.
- 850 [8] Borg, J. M., & Channon, A. D. (2012). Testing the variability selection
851 hypothesis - the adoption of social learning in increasingly variable envi-
852 ronments. In C. Adami, D. M. Bryson, C. Ofria, & R. T. Pennock (Eds.)
853 *ALIFE 13: The Thirteenth International Conference on the Synthesis and*
854 *Simulation of Living Systems*, (pp. 317–324). MIT Press.
- 855 [9] Boyd, R., & Richerson, P. J. (1985). *Culture and the Evolutionary Process*.
856 University of Chicago Press.
- 857 [10] Bullinaria, J. A. (2017). Imitative and direct learning as interacting factors
858 in life history evolution. *Artificial Life*, 23.
- 859 [11] Channon, A. D., & Damper, R. (1998). The evolutionary emergence of
860 socially intelligent agents. In B. Edmonds, & K. Dautenhahn (Eds.) *Socially*
861 *Situated Intelligence: a workshop held at SAB'98, University of Zurich*
862 *Technical Report*, (pp. 41–49).
- 863 [12] Dawkins, R. (1989). *The Selfish Gene*. Oxford University Press, 2nd edition
864 ed.
- 865 [13] Dennett, D. C. (1995). *Darwin's Dangerous Idea: Evolution and the mean-*
866 *ings of life*. Simon and Schuster.

- 867 [14] Dugatkin, L. A., & Godin, J.-G. J. (1993). Female mate copying in the
868 guppy (*Poecilia reticulata*): age-dependent effects. *Behavioral Ecology*,
869 4(4), 289–292.
- 870 [15] Galef Jr, B. G. (1996). Social enhancement of food preferences in Norway
871 rats: a brief review. In C. M. Heyes, & B. G. Galef (Eds.) *Social learning*
872 *in animals: the roots of culture*, (pp. 49–64). Academic Press San Diego.
- 873 [16] Galef Jr, B. G., & Giraldeau, L.-A. G. (2001). Social influences on for-
874 aging in vertebrates: causal mechanisms and adaptive functions. *Animal*
875 *Behaviour*, 61, 3–15.
- 876 [17] Greene, C. M., & Stamps, J. A. (2001). Habitat selection at low population
877 densities. *Ecology*, 82(8), 2091–2100.
- 878 [18] Henrich, J., & Boyd, R. (1998). The evolution of conformist transmission
879 and the emergence of between-group differences. *Evolution and Human*
880 *Behavior*, 19, 215–241.
- 881 [19] Henrich, J., & Gil-White, F. J. (2001). The evolution of prestige: Freely
882 conferred deference as a mechanism for enhancing the benefits of cultural
883 transmission. *Evolution and Human Behavior*, 22(3), 165–196.
- 884 [20] Higgs, P. G. (2000). The mimetic transition: a simulation study of the
885 evolution of learning by imitation. *Proceedings of the Royal Society of*
886 *London B: Biological Sciences*, 267(1450), 1355–1361.
- 887 [21] Hoppitt, W., & Laland, K. N. (2008). Social processes influencing learning
888 in animals: a review of the evidence. *Advances in the Study of Behavior*,
889 38, 105–165.
- 890 [22] Jolley, B. P., Borg, J. M., & Channon, A. (2016). Analysis of social learn-
891 ing strategies when discovering and maintaining behaviours inaccessible to

- 892 incremental genetic evolution. In E. Tuci, A. Giagkos, M. Wilson, & J. Hal-
893 lan (Eds.) *From Animals to Animats 14: 14th International Conference on*
894 *Simulation of Adaptive Behavior*, (pp. 293–304). Springer.
- 895 [23] King, A. J., & Cowlshaw, G. (2007). When to use social information:
896 the advantage of large group size in individual decision making. *Biology*
897 *Letters*, 3(2), 137–139.
- 898 [24] Kirkpatrick, M., & Dugatkin, L. A. (1994). Sexual selection and the evolu-
899 tionary effects of copying mate choice. *Behavioral Ecology and Sociobiology*,
900 34(6), 443–449.
- 901 [25] Laland, K. N. (2004). Social learning strategies. *Animal Learning & Be-*
902 *havior*, 32(1), 4–14.
- 903 [26] Mitri, S., Floreano, D., & Keller, L. (2009). The evolution of information
904 suppression in communicating robots with conflicting interests. *Proceedings*
905 *of the National Academy of Sciences, USA*, 106(37), 15786–15790.
- 906 [27] Noble, J., Todd, P. M., & Tuci, E. (2001). Explaining social learning of
907 food preferences without aversions: an evolutionary simulation model of
908 Norway rats. *Proceedings of the Royal Society of London B: Biological*
909 *Sciences*, 268(1463), 141–149.
- 910 [28] Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman,
911 M. W., Fogarty, L., Ghirlanda, S., Lillicrap, T., & Laland, K. N. (2010).
912 Why copy others? Insights from the social learning strategies tournament.
913 *Science*, 328(5975), 208–213.
- 914 [29] Rendell, L., Fogarty, L., Hoppitt, W. J., Morgan, T. J., Webster, M. M., &
915 Laland, K. N. (2011). Cognitive culture: theoretical and empirical insights
916 into social learning strategies. *Trends in Cognitive Sciences*, 15(2), 68–76.

- 917 [30] Robinson, E., Ellis, T., & Channon, A. (2007). Neuroevolution of agents
918 capable of reactive and deliberative behaviours in novel and dynamic en-
919 vironments. In F. Almeida e Costa (Ed.) *Advances in Artificial Life: Pro-*
920 *ceedings of the 9th European Conference on Artificial Life*, (pp. 345–354).
921 Springer.
- 922 [31] Sergio, F., & Newton, I. (2003). Occupancy as a measure of territory
923 quality. *Journal of Animal Ecology*, 72(5), 857–865.
- 924 [32] Stamps, J. A. (1988). Conspecific attraction and aggregation in territorial
925 species. *The American Naturalist*, 131(3), 329–347.
- 926 [33] Stanton, A., & Channon, A. D. (2015). Incremental neuroevolution of reac-
927 tive and deliberative 3D agents. In *Proceedings of the European Conference*
928 *on Artificial Life 2015*, (pp. 341–348).
- 929 [34] Stephens, P. A., Sutherland, W. J., & Freckleton, R. P. (1999). What is
930 the Allee effect? *Oikos*, (pp. 185–190).
- 931 [35] Valone, T. J. (1989). Group foraging, public information, and patch esti-
932 mation. *Oikos*, (pp. 357–363).
- 933 [36] Valone, T. J. (2007). From eavesdropping on performance to copying the
934 behavior of others: a review of public information use. *Behavioral Ecology*
935 *and Sociobiology*, 62(1), 1–14.
- 936 [37] van Bergen, Y., Coolen, I., & Laland, K. N. (2004). Nine-spined stickle-
937 backs exploit the most reliable source when public and private information
938 conflict. *Proceedings of the Royal Society of London, Series B: Biological*
939 *Sciences*, 271(1542), 957–962.
- 940 [38] van der Post, D. (2008). *Learning What to Eat: emerging cultural phenom-*
941 *ena in group foragers*. Ph.D. thesis, Utecht University.

- 942 [39] Wagner, R. H., & Danchin, É. (2003). Conspecific copying: a general
943 mechanism of social aggregation. *Animal Behaviour*, *65*, 405–408.
- 944 [40] Whitehead, H., & Richerson, P. J. (2009). The evolution of conformist
945 social learning can cause population collapse in realistically variable envi-
946 ronments. *Evolution and Human Behavior*, *30*(4), 261–273.
- 947 [41] Whiten, A., & Van Schaik, C. P. (2007). The evolution of animal 'cultures'
948 and social intelligence. *Philosophical Transactions of the Royal Society of*
949 *London B: Biological Sciences*, *362*(1480), 603–620.
- 950 [42] Yang, S. M., & Meng, M. (2000). An efficient neural network method
951 for real-time motion planning with safety consideration. *Robotic and Au-*
952 *tonomous Systems*, *32*, 115–128.
- 953 [43] Yang, S. X., & Meng, M. (2000). An efficient neural network approach to
954 dynamic robot motion planning. *Neural Networks*, *13*(2), 143–148.

955 **Appendix: Tables of Parameters**

956 Here a series of tables providing an overview of the parameter settings used in
957 this work.

958 **A1: EnVar Parameters**

EnVar and Plant Parameters	
Parameter	Settings/Range
World Size (cells)	700×700
Number of Plants	500
959 Plant Size (cells)	10×10
Negative Plant Species Energy	$E_{neg} = -10$
Positive Plant Species Energy	$E_{pos} = 1.0$
Eating Events until Plant is exhausted	200

960 **A1 justifications and rationale**

961 Plant size and quantity was set in order for plants to take up approximately
 962 10% of the world area. During preliminary testing of the system this density
 963 of plants ensured plants were a frequently encountered feature of environment,
 964 without being densely packed; agents still often had to search for plants. Further
 965 exploration of the results presented here in regard to plant density would be
 966 worthwhile, as it would be expected to have a significant effect on movement
 967 behavior. The eating events until plant exhaustion parameter was set to equal
 968 population size - increasing or decreasing this variable would be expected to
 969 affect the proportion of time agents spent eating.

970 **A2: Simulation and Population Parameters**

Simulation and Population Parameters	
Parameter	Settings/Range
971 Simulation length (epochs)	100
Epoch length (timesteps)	1000
Population Size	200

972 **A2 justifications and rationale**

973 Epoch length was set to be approximately the number of timesteps required for
 974 an agent to move from one corner of the world to the other. Simulation length
 975 set based on preliminary testing, both task performance and agent behavior
 976 was seen to stabilise for all population types and environments by 100 epochs.
 977 Population size was set in regard to computational time to obtain results - larger
 978 population sizes had the negative consequence of both longer run times, and
 979 more densely packed environments. Varying population size would be expected
 980 to affect movement behavior, as a higher density of agents would reduce the
 981 need to search for other agents.

982 **A3: Agent and Evolution Parameters**

Agent and Evolution Parameter	
Parameter	Settings/Range
Agent size (radius, cells)	2
Visual field (radius, cells)	30
Max/initial battery	100
Stationary energy loss (per timesteps)	0.1
983 Movement energy loss (per timesteps)	0.2
% of population replaced at epoch	10%
Genotype length	$L = 112$ or $L = 128$
Mutation rate (per locus)	$p_{mut} = 1/L$
Gaussian random noise (mean)	$\mu = 0$
Gaussian random noise (standard deviation)	$\sigma = 0.01$
Crossover	Single point

984 **A3 justifications and rationale**

985 As agents could not inhabit the same cells as one another, agent size was kept
 986 small to minimise the need to recalculate agent movement choices. Visual field
 987 size was set to be as large as possible in regard to computational time to obtain
 988 results. The creation of activity landscapes was computationally expensive, and
 989 therefore limited the size of visual fields. Larger visual fields would be expected
 990 to encourage a higher proportion of movement behavior. Max battery was set
 991 in relation to stationary energy loss; the current configuration results in an
 992 agent losing all energy within one epoch should they remain static throughout.
 993 Movement energy loss was set to be double that of stationary energy loss to
 994 discourage movement behavior unless selected for. Energy loss could be further
 995 explored to better understand the dependency of agent behavior on the cost of
 996 behavior.

997 **A4: Neural Network (Shunting Model) Parameters**

Neural Network (Shunting Model) Parameters	
Parameter	Settings/Range
Decision network input units	$i = 7$ or $i = 8$
Decision network hidden units	$h = 8$
Decision network output units	$o = 2$
998 Maximum Iota value	$maxI = 15$
Minimum Iota value	$minI = -15$
Negative Iota output thresholds	$-\theta_a = -0.5$
Positive Iota output thresholds	$+\theta_a = 0.5$
Movement threshold	$\theta_b = 0.5$

999 **A4 justifications and rationale**

1000 Hidden layer size was set based on preliminary testing - larger hidden layers
1001 didn't provide noticeably better performance on the task, but did increase com-
1002 putational time to obtain results. The consequence of increasing the maximum
1003 Iota value (and decreasing the minimum Iota value) would be to allow activity
1004 from resource of agents to propagate further within the visual field, therefore
1005 objects on the edges of visual fields would have greater influence on agent de-
1006 cisions - it is not anticipated that this would cause a large change in agent
1007 behavior. The maximum and minimum Iota values set here were found to be
1008 sufficient for allowing object activation to influence agent decisions. Adjusting
1009 thresholds would be expected to affect the likelihood of agent behaviors being
1010 applied. Current thresholds do not bias agent decisions in favor of any of the
1011 actions.

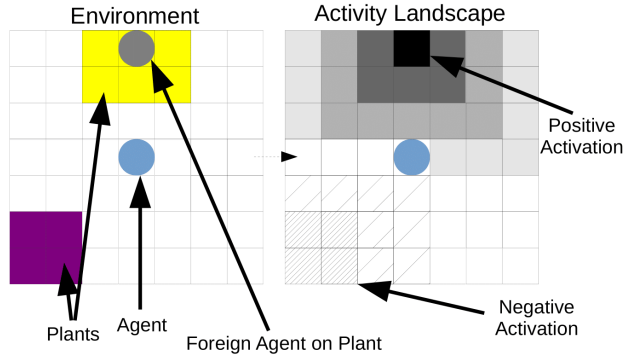


Figure 1: Mock-up transition from agent visual field to shunting network activity landscape. The left-hand grid shows the agent's visual field with two plant objects and one other agent occupying the same space as a plant. The right-hand grid shows an example activity landscape for the visual field. 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 equation 1. The central agent then chooses to move within its immediate Moore neighbourhood to the cell with highest activity value.

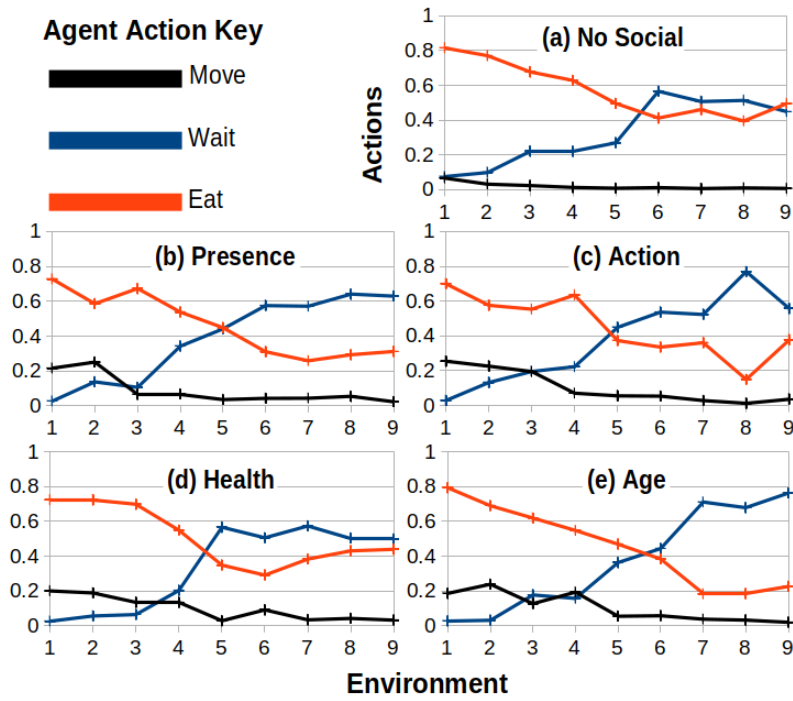


Figure 2: Median agent action profiles for each social information strategy over each environment difficulty.




Figure 3: Median **move** actions for each social information strategy over all environments. Data points on the primary y-axis represent the median proportion of the move action. Data points on the secondary y-axis represent the Z-score value from a Mann-Whitney U test comparing, for each environment, the median actions for the two social information strategies presented. Z-scores which indicate statistically significant p values are highlighted on the x-axis.

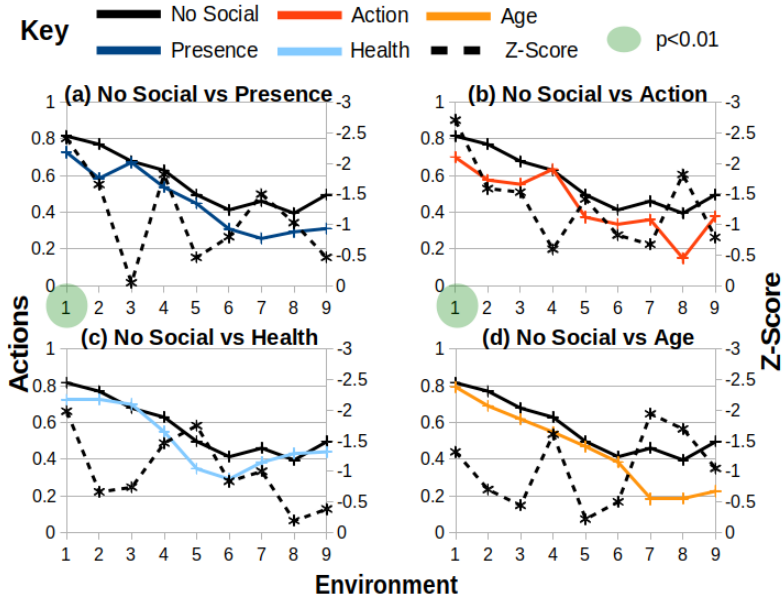


Figure 4: Median **eat** actions for each social information strategy over all environments. Data points on the primary y-axis represent the median proportion of the eat action. Data points on the secondary y-axis represent the Z-score value from a Mann-Whitney U test comparing, for each environment, the median actions for the two social information strategies presented. Z-scores which indicate statistically significant p values are highlighted on the x-axis.

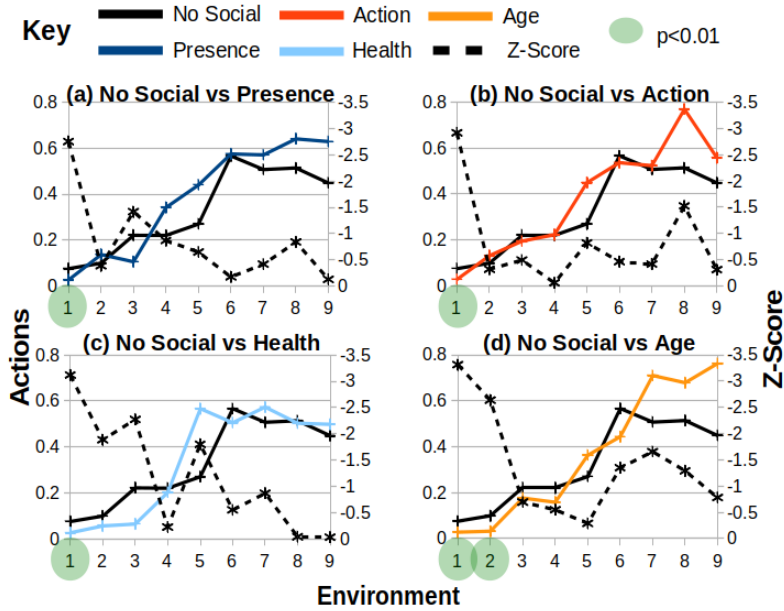


Figure 5: Median **wait** actions for each social information strategy over all environments. Data points on the primary y-axis represent the median proportion of the wait action. Data points on the secondary y-axis represent the Z-score value from a Mann-Whitney U test comparing, for each environment, the median actions for the two social information strategies presented. Z-scores which indicate statistically significant p values are highlighted on the x-axis.

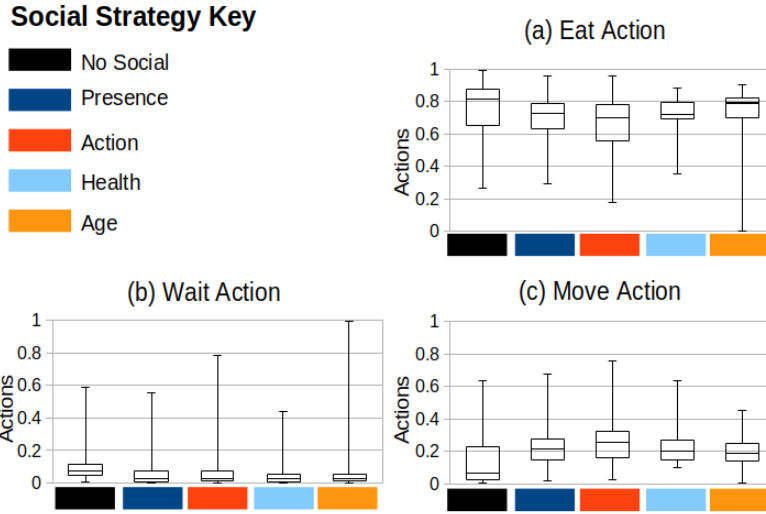


Figure 6: Action box-plots for each action, for each social information strategy in environment 1, where there is a 1:1 ratio of positive to negative food resources.

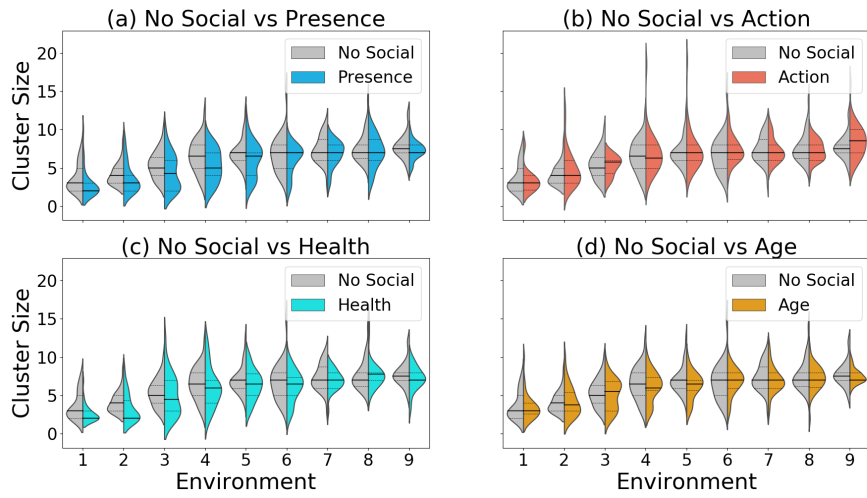


Figure 7: Distribution of cluster sizes for each social information strategy against the no social strategy over each environment difficulty.

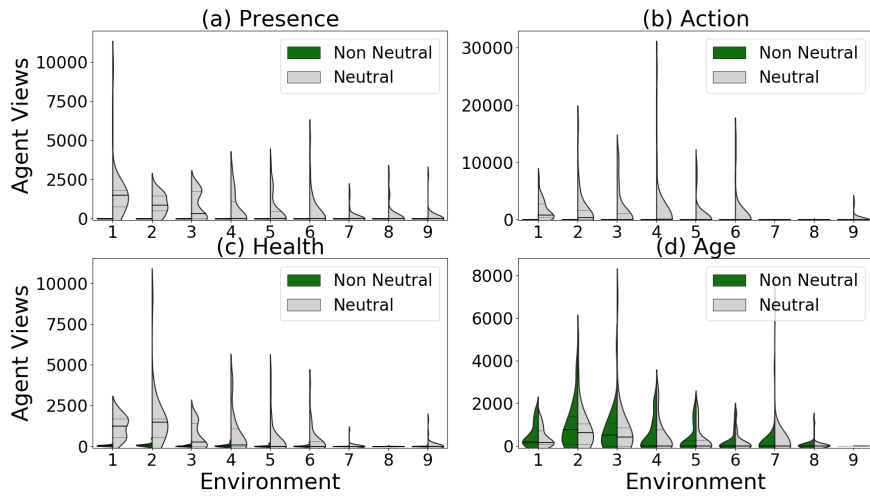


Figure 8: Distribution of neutral vs non neutral agent views for each social strategy over each environment difficulty.

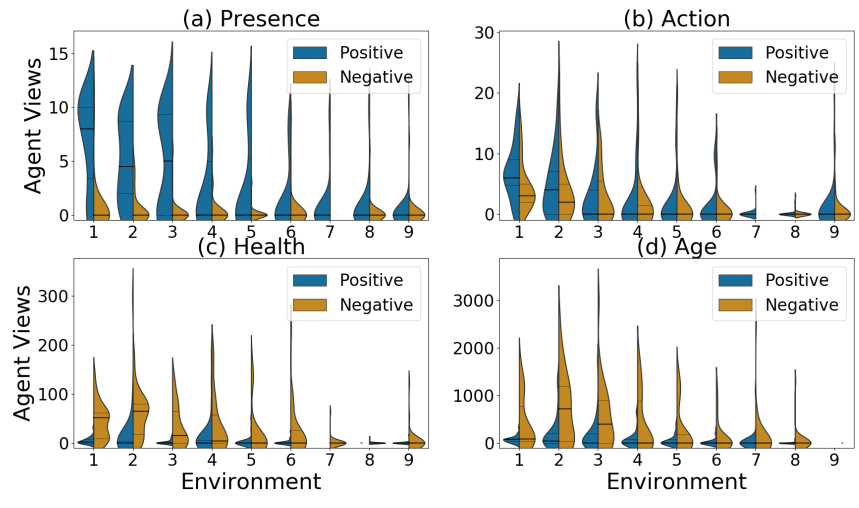


Figure 9: Distribution of positive vs negative agent views for each social strategy over each environment difficulty.

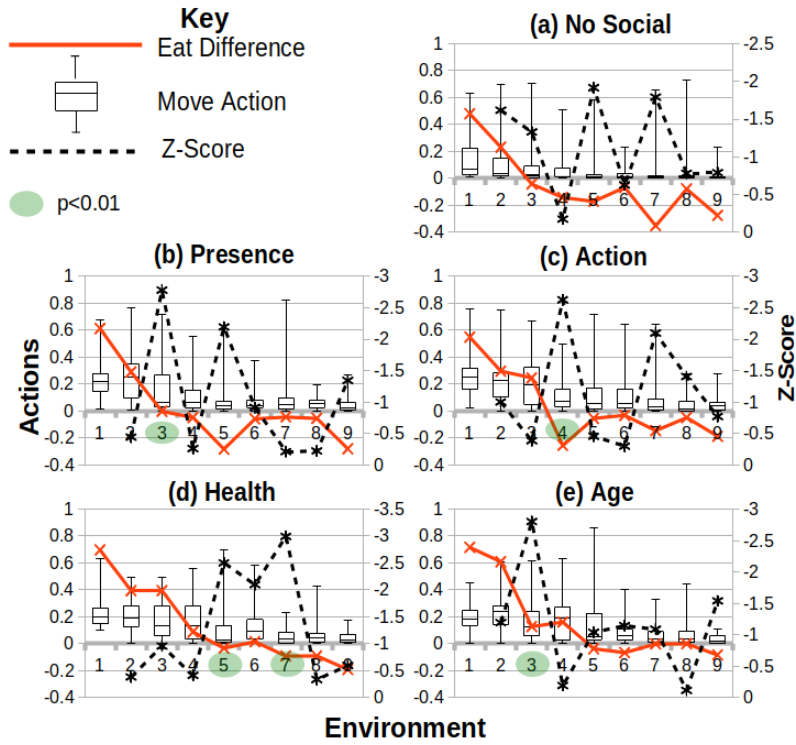


Figure 10: The median differences between successful and unsuccessful eat actions is presented on the primary y-axis along with the box plots for the move action. The Z-score from Mann-Whitney U tests, which compare the action data for the environment on which a data point falls with the previous environment, is presented on the secondary y-axis. These Z-scores indicate which transitions in action behavior between previous environments are significant. Z-scores which indicate statistically significant p values are highlighted on the x-axis.

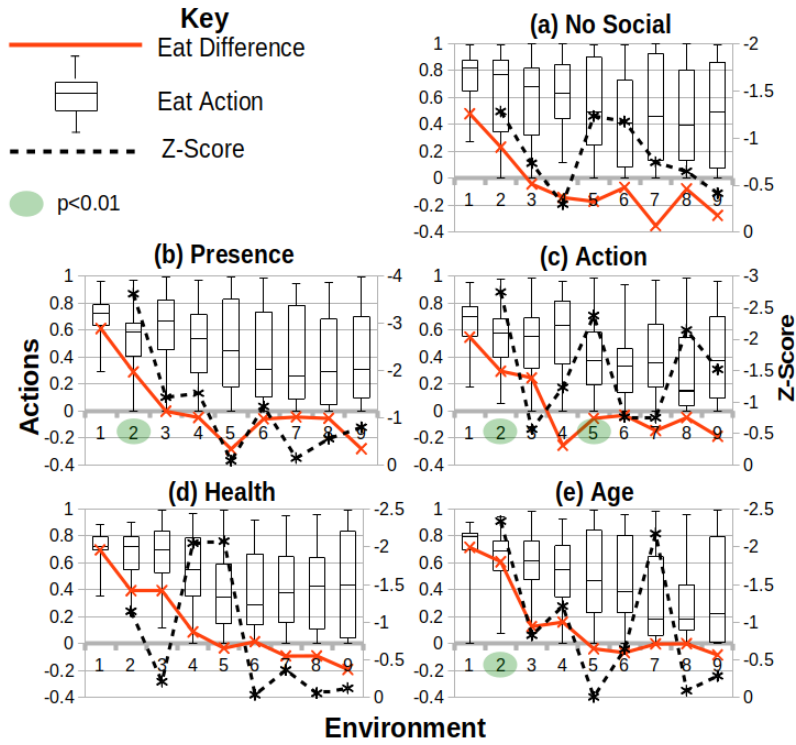


Figure 11: The median differences between successful and unsuccessful eat actions is presented on the primary y-axis along with the box plots for the eat action. The Z-score from Mann-Whitney U tests, which compare the action data for the environment on which a data point falls with the previous environment, is presented on the secondary y-axis. These Z-scores are indicate which transitions in action behavior between previous environments are significant. Z-scores which indicate statistically significant p values are highlighted on the x-axis.

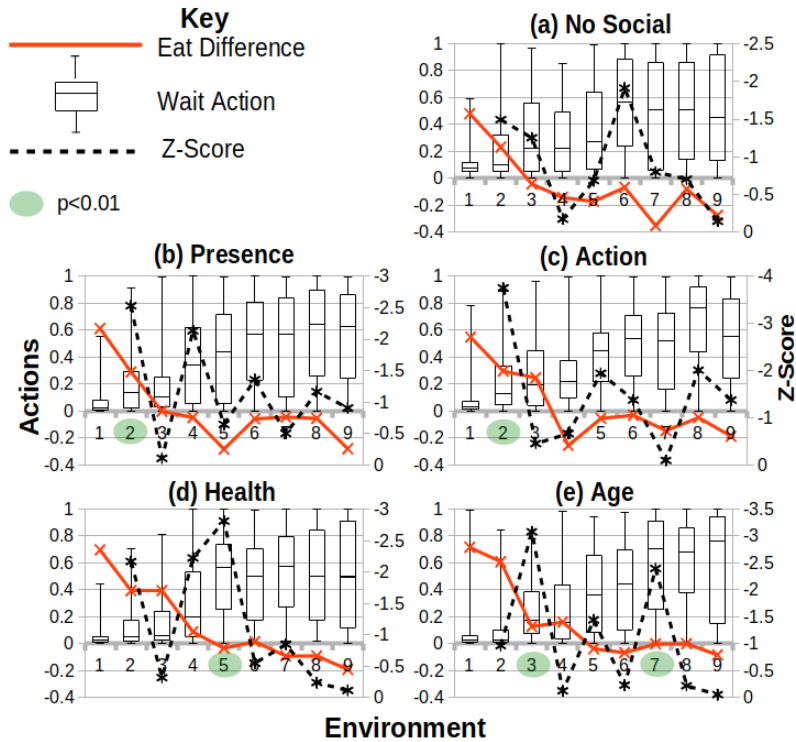


Figure 12: The median differences between successful and unsuccessful eat actions is presented on the primary y-axis along with the box plots for the wait action. The Z-score from Mann-Whitney U tests, which compare the action data for the environment on which a data point falls with the previous environment, is presented on the secondary y-axis. These Z-scores are intended to indicate which transitions in action behavior between previous environments are significant. Z-scores which indicate statistically significant p values are highlighted on the x-axis.