Superstition in the cognitive model: Modelling ritualised behaviour as error management *

Mariusz Rybnik[1], Ivan Puga-Gonzalez[2], F. LeRon Shults[2], Paweł Łowicki[3], Ewa Dąbrowska-Prokopowska[4], Andrew Atkinson[5], and Konrad Talmont-Kaminski[5]

[1] Institute of Computer Science, University of Bialystok, Bialystok, Poland
[2] NORCE Center for Modeling Social Systems, Kristiansand, Norway
[3] Faculty of Psychology, University of Warsaw, Warsaw, Poland
[4] Institute of Sociology, University of Bialystok, Bialystok, Poland
[5] Society and Cognition Unit, University of Bialystok, Bialystok, Poland
m.rybnik@uwb.edu.pl

Abstract. Skinner found that superstitious behaviour in pigeons results from accidental operant conditioning. We use a simple cognitive model based upon reinforcement learning to show that ritualisation of behaviour arises in analogous conditions. This makes it possible to model the creation of ritual traditions using minimal means in agent-based models, thereby opening a novel and potentially highly fruitful approach to the study of this highly significant human behaviour.

Keywords: Error Management, Cognitive Model, superstition, ritualised behaviour

1 Introduction

Most research into ritualised behaviour tends towards high-level explanations that involve, among other things, the cultural transmission of supernatural beliefs. The explanation for the spontaneous appearance of superstitious behaviours in pigeons provided by the behaviourist psychologist B.F. Skinner could not be any more different [9]. It sought to explain superstitious behaviour merely by reference to accidental operant conditioning. Understood as showing that superstitious behaviour arises spontaneously where an agent - artificial or not - attempt to control an environment that is unpredictable, this line of research has been expanded upon greatly by researchers such as Ono, who showed the same effect in humans [8], as well as as Killeen, who showed that pigeons placed in analogous conditions vary their rate of superstitious behaviour in a way that maximises their pay-off given uncertainty [7]. Killeen's observation has been generalised by Haselton and Buss, who showed that

* The research leading to these results has received funding from the Norwegian Financial Mechanism 2014-2021 (Project number 2019/34/H/HS1/00654).

human and other evolved cognitive mechanisms will be biased in favour of committing the least costly errors when operating under uncertainty, i.e., they will exhibit error management [5]. Given the need to detect potential threats/opportunities in the environment, people will avoid false negative errors at the cost of making more false positive ones. This includes overdetecting causal connections and agency. Further research has expanded the scope of this type of explanation even more by showing that the kind of behaviour Skinner observed is a normal by-product even for formal or computing mechanisms that attempt to predict system behaviour under conditions that are often met with in natural settings [2] [4]. In effect, researchers working in the Skinner tradition have shown that superstitious behaviour of all agents - be they artificial or natural - is to be explained as a by-product of the epistemic circumstances encountered when attempting to control any unpredictable system.
The aim of the line of research pursued here is to see whether ritualised behaviour can be explained in the same general terms, i.e., as the product of false positive errors resulting from error management, and to explore how changes in conditions affect spontaneous ritualisation in silico. Our ultimate goal is to create an agent-based model, to be used to study the emergence of stable rituals at the group level. In other words, we are seeking to show that rituals, including group rituals, are also to be ultimately explained in terms of a fundamental phenomenon that affects all agents given the right set of basic conditions - with most aspects of rituals as normally discussed (including supernatural beliefs or the role of anxiety) being ancillary (see [10]). The definition of ritualised behaviour used by us is drawn from Boyer and Lienard [3] where they identify five basic characteristics common to ritualised behaviours. Three are investigated here: redundancy - including elements in one's behaviour that are not necessary to achieve the apparent aim; goal-demotion - including elements in one's behaviour that have no apparent effect; rigidity - performing all of the elements, including the redundant and goal-demoted ones, in a strict order. The modelled behaviour is studied with the use of the PathGame - a (pseudo)game with clear and easily manipulable rules that makes operationalisation and analysis straight forward.

2 The PathGame

PathGame is a novel methodology developed in order to empirically investigate the conditions under which humans, as well as computational cognitive models, spontaneously ritualise their behaviour. It is based upon Stuart Vyse's earlier research [6] [12] into superstitious behaviour. Vyse's work was explicitly pursued in the Skinnerian tradition described above, so expanding upon helps to make the connection between superstitious and ritualised behaviours explicit. Getting humans and artificial agents to deal with the same problem allows direct comparison of their behaviours.
The game takes place on a four-by-four matrix. The basic goal is to move the avatar from the starting position in the top-left corner to the exit situated two blocks down and to the right. Five buttons are available:

four directional ones and a fifth, pressing which does not lead to any apparent effect. Players are allowed to play the game fifty times, sometimes receiving points upon reaching the exit cell, with the aim being to get the maximum number of points. The interactive game's interface (implemented as a TypeScript web application for the purpose of human studies not presented here) is presented in Fig. 1. **Redundancy** is operationalised as pressing the up or left buttons - pressing them is unnecessary to reach the exit. **Goal-demotion** is operationalised as pressing the fifth 'mystery' button - pressing it has no effect apart from being recorded. **Rigidity** is operationalised as repeating (within ten attempts) the same non-minimal path, i.e., one that includes redundancy or goal-demotion. Rigidity is distinguished from **automaticity**, which is operationalised as repeating any minimal path within ten attempts, and which is understood to result in humans from minimisation of cognitive effort rather than ritualisation.

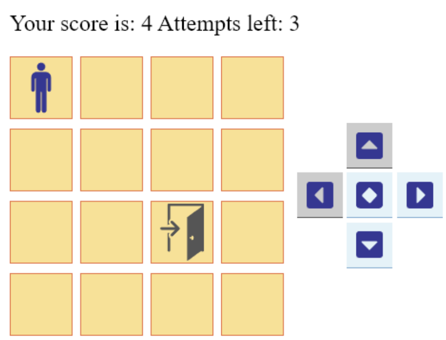

Fig. 1. PathGame implemented as web application using TypeScript

When faced with Vyse's simpler set-up (only the right and down buttons), human players tended to correctly represent the game's winning conditions when they were predictable but generated complex and completely fictitious hypotheses when points were awarded randomly. In our own pilot studies with human subjects, we found that redundancy, goal-demotion and rigidity all increased when the probability of obtaining points at each attempt was low, but automaticity grew as success was more likely. In effect, people appeared to be learning to form incorrect associations between the ritualised aspects of their behaviour and obtaining points when points were obtained randomly and rarely. As Vyse noted, "operant conditioning is not just for rats and pigeons." ([11])

3 The Cognitive Model (CM)

In order to provide an artificial model of spontaneous ritualisation using the *PathGame* paradigm, a computational CM was implemented using the Anylogic environment [1]. The goal was to determine whether spontaneous ritualisation of behaviour qualitatively akin to that exhibited by humans could be generated using a simple artificial system designed to identify non-random patterns, thereby showing a common underlying basis for the phenomenon. The CM works on the basis of a set of weights on each of the cells in the matrix, which determine the probability that the CM will make a particular move when the avatar is in that cell. The avatar moves around the matrix, with the each move determined stochastically on the basis of the weights at the currently-occupied cell, until the exit cell is reached. Depending upon whether a point was obtained at the end of the walk, the weights of the moves made are either weakened or reinforced.

3.1 Terms and parameters

A *Move* is regarded here as a movement to another valid cell or the *Goal-Demotion* pseudo-activity. *Moves* are denoted as the four cardinal geographical directions on a map: N, E, S, W, with Goal-Demotion denoted as the question mark: ?. A *Walk* is defined as a sequence of *Moves* leading from the starting position to the exit cell and is represented symbolically as a *Path*, such as: $SSEW?EE$. A single *Game* consists of a given number of *Walks*, which for the purpose of this research was always 50. CM try to gain and store the knowledge about rewards along the *Game*.
The game matrix with initial and final move weights is presented in Fig. 2. A *Move weight* is a floating-point number associated with the specific *move* at the specific cell. The *move weights* are depicted in each cell (with the obvious exception of the exit cell marked as a circle). The numbers indicate weights for each direction, with the number in the center being the *Goal-Demotion* weight. The *Move probability* of move X at a cell (i, j) can be obtained by dividing the corresponding weight t_{Xij} by the sum of all weights $\sum t_{xij}$, where $x \in \{N, E, S, W, ?\}$. Initial weights are set to ensure the CM-controlled avatar is likely to reach the exit cell in a small number of moves (just as is the case with human players). Reinforcement/attenuation during the game changes these weights significantly leading to very different behaviour depending upon the reinforcement schedule.

3.2 State diagram

The CM completed 50 *walks*. After each *walk*, *Reinforcement Learning* or *Attenuate Learning* is performed according to a random or non-random schedule. The state diagram of the CM implemented in Anylogic is presented in Fig. 3.
During the game the CM performs the following steps:
Step 1. Initialise move weights
Step 2. Make random weight-determined moves till end cell reached
Step 3. Schedule determines whether to reinforce/attenuate the walk
Step 4A. Reinforcement learning

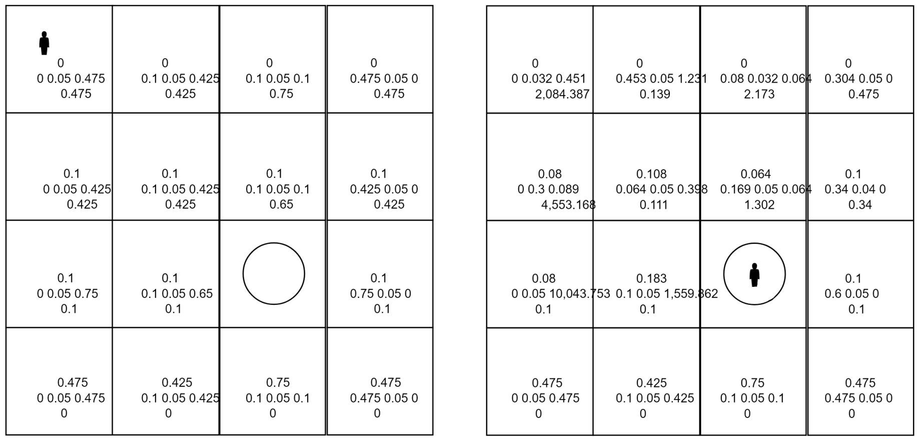

Fig. 2. The game matrix with starting move weights (on the left), and the game matrix with final move weights (on the right)

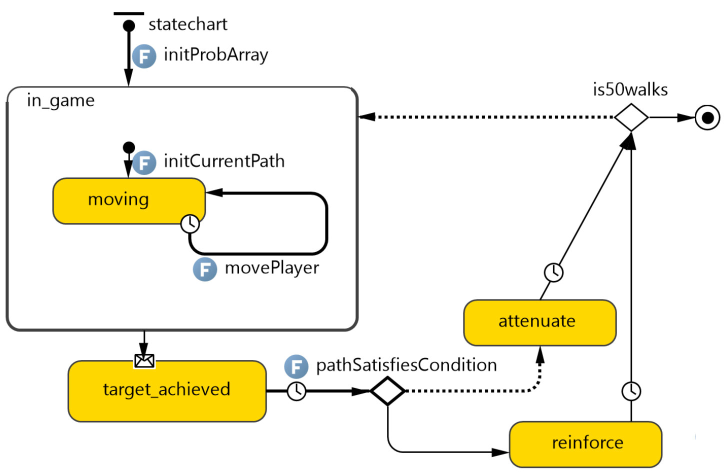

Fig. 3. State diagram of CM (in Anylogic)

OR
Step 4B. Attenuate learning
Step 5. While less than 50 walks loop back to step 2

3.3 Reinforcement vs. Attenuation

If a *walk* is rewarded, reinforcement occurs - the *moves* in the *walk* have their *move weights* increased as follows: $t_{Xij} \rightarrow t_{Xij} \times ReinforceRatio$, where $ReinforceRatio > 1$. In effect, future *walks* are more likely to

include the same *moves*. Due to stochastic nature of the CM however, some degree of indeterminacy is maintained, thereby allowing experimentation.
If a *walk* is not rewarded, attenuation occurs - the *move weights* contained in the specific *path* are attenuated: $t_{Xij} \rightarrow t_{Xij} \times AttenuateRatio$, where $AttenuateRatio \in (0;1)$. In effect, in future future *walks*, the *CM* is less likely to perform the same *moves*, as they were not beneficial.
In addition, to avoid the CM generating very long paths, path length mitigation is introduced by making reinforcement/attenuation proportional to path length, with the minimal path length of 4 being regarded as the base value.

4 Testing methodology

Given that on the Skinnerian approach being pursued here, ritualised behaviour is understood as arising due to accidental operant conditioning, it is important to test the CM in two different kinds of scenarios - random and non-random. In the non-random scenarios, the CM should be able to identify the winning paths while in the random scenarios it should generate ritualised behaviour.

4.1 Non-random Scenarios

The non-random scenarios are based on a set of predefined, alternative rules that reward specific CM behaviour. These rules are based on cells the path has to go through or avoid or specific buttons being pressed or not pressed. To test the CM in more complex non-random scenarios, some include pairs of these rules.
In these scenarios, a path increases the score by one and is reinforced if it satisfies the rule in force in the game being played. Otherwise, the path is attenuated and the score remains unchanged. Where a pair of rules is in force, both have to be satisfied by a path for it to be reinforced and for the score to increase by one.

BASIC RULES

- CELL-BASED RULES
  - R1. Avoid the bottom left quadrant
  - R2. Avoid the top right quadrant
  - R3. Walk through any 4th row cell
  - R4. Walk through any 4th column cell
- BUTTON-BASED RULES
  - R5. Press the up arrow at least once
  - R6. Press the left arrow at least once
  - R7. Never press the goal-demotion button
  - R8. Press the goal-demotion button at least once
  - R9. Press the goal-demotion button at least twice

BASIC RULE PAIRS

- **R1 and R3**: Avoid the bottom left quadrant and Walk through any 4th row cell
- **R1 and R5**: Avoid the bottom left quadrant and Press the up arrow at least once
- **R1 and R7**: Avoid the bottom left quadrant and Never press the goal-demotion button
- **R2 and R4**: Avoid the top right quadrant and Walk through any 4th column cell
- **R2 and R6**: Avoid the top right quadrant and Press the up arrow at least once
- **R2 and R7**: Avoid the top right quadrant and Never press the goal-demotion button
- **R3 and R4**: Walk through any 4th row cell and through any 4th column cell
- **R3 and R6**: Walk through any 4th row cell and Press the left arrow at least once
- **R3 and R7**: Walk through any 4th row cell and Never press the goal-demotion button
- **R4 and R5**: Walk through any 4th column cell and Press the up arrow at least once
- **R4 and R7**: Walk through any 4th column cell and Never press the goal-demotion button
- **R5 and R7**: Press the up arrow at least once and Never press the goal-demotion button
- **R6 and R7**: Press the left arrow at least once and Never press the goal-demotion button

4.2 Random Scenarios

In these scenarios, reward and the associated reinforcement/attentuation schedule are independent of the path used and depend merely upon the *ReinforceProbability* parameter, which gives the stochastic probability of reward/reinforcement.

5 Results

5.1 Results for Non-random Scenarios

The goal of these scenarios was to test whether and to what degree the CM was able to extract rules knowledge from the received rewards and to then use that knowledge to increase the number of points obtained over the length of the game. As such, the scores obtained by the CM have to be compared to the scores obtained without any reinforcement - these show how difficult the rules are to satisfy purely randomly and differ between the different non-random rules tested.
The scores obtained by the CM for the basic rules and their pairs are presented in the Fig. 4. These are the mean scores for 1000 *games* of 50 walks. Each *game* is initialized in the same way (including the same initial move weights) and with the following parameter values:

– *ReinforceRatio* = 9
– *AttenuateRatio* = 0.8
– *Path Size Mitigation* applied linearly (see section 3.3 for details).

The rules tested proved to vary greatly in terms of how easy they were to happen upon randomly, with the rule pairs generally proving more difficult. This pattern also held for cases where reinforcement was in place. However, in every case tested, reinforcement allowed the CM to obtain a clearly increased average score. In fact, it was with rule pairs that the effect of reinforcement was the most striking. Of course, the difficulty of individual rules could be modified by altering the initial weights - for example, pressing the goal-demotion button could be made much more likely, thereby making R9 much more likely to be satisfied randomly. So, not much can be read into individual rules. However, it is the overall pattern of reinforcement increasing scores that shows that for the set of rules we tested, the reinforcement learning was successful and that this CM is a satisfactory model of pattern-seeking behaviour. Having observed this behaviour, we could go on to test how this CM behaved when presented not with predictable scenarios but with purely stochastic ones.

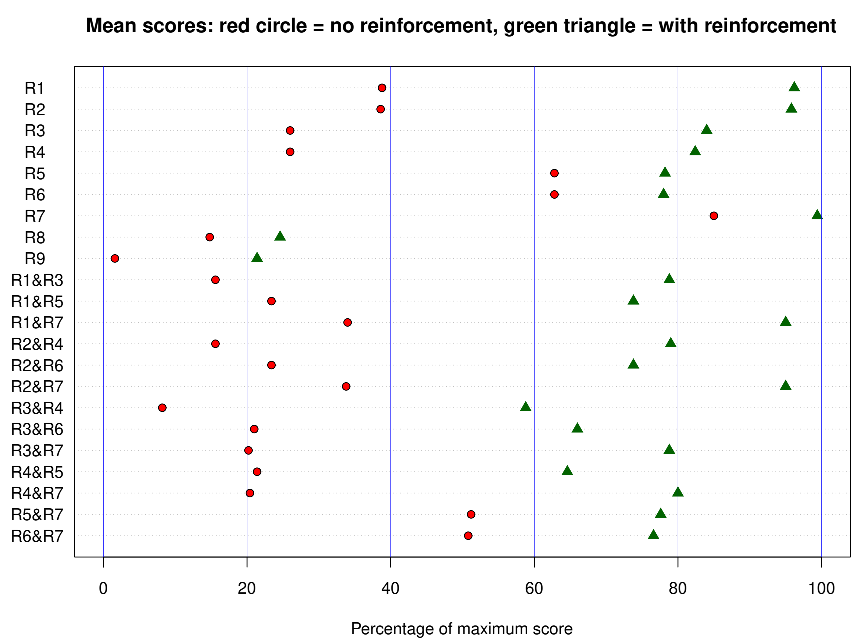

Fig. 4. CM mean scores for the individual rules and selected rule pairs, comparing reinforcement with no reinforcement

5.2 Results for Random Scenarios

Having established that the CM is capable of learning non-random rules, it was necessary to check whether this was sufficient for the CM to generate ritualised behaviour when presented with random scenarios, as operationalised in terms of redundancy, goal-demotion and rigidity (see section 2).

CM configuration is as follows:

- parameter *ReinforcementProbability* $\in < 0 : 0.01 : 1 >$ - 101 values
- 1000 *Games* per each *ReinforcementProbability* value.
- $ReinforceRatio = 9$
- $AttenuateRatio = 0.8$
- *Path Size Mitigation* applied linearly (see section 3.3 for details)

A total of 101 000 *Games* were run. Each simulation consisted of 50 time steps (i.e., 50 *Walks* or *Paths*). Rewards were awarded stochastically based on the value of *ReinforcementProbability* and independently of the path used.

The metrics for the four characteristics obtained from the experiments are presented in Fig. 5.

Goal-Demotion - as shown in the upper left plot in Figure 5 - occurs in the majority of paths when reinforcement is rare but drops away and is almost never seen in cases where the probability of reinforcement is above 0.3. This is in line with human behaviour, where it has been observed that goal-demotion also is much more common when paths are rarely reinforced.

Redundancy - as shown in the upper right plot in Figure 5 - is very high at low *ReinforcementProbability* values - occurring in the majority of walks - but then drops towards zero as the probability of reward increases, similarly to goal-demotion. This pattern has also been observed in human players.

Automaticity - shown in the bottom left plot in Figure 5 - is not a necessary characteristic of ritualised behaviour but was of interest to us given the methodology used. It behaved in the opposite manner to goal-demotion and redundancy, in that it was almost non-existent when reinforcement was low and quickly came to dominate at higher reinforcement probabilities. This pattern is also analogous to the one observed with human players.

Rigidity - shown in the bottom right plot in Figure 5 - proved to be the problematic measure. As can be seen, it is quite rare at high *ReinforcementProbability* values but is completely absent when reinforcement drops to zero. As such it will be discussed at some length in the final section. Rigidity is regarded here as the repetition of a non-minimal path within the ten most recent walks. It is low with low *ReinforcementProbability* values (bottom right plot in Figure 5). Due to rare learning the paths tend to be quite long and very stochastic, so they rarely repeat subsequently at all. Rigidity rises quickly with the increase of *ReinforcementProbability* and stabilizes obtaining about 6-7 occurrences out of 50 with *ReinforcementProbability* $\in (0.25, 1)$. It could be interpreted that the paths tend to be more focused and repetitive, both minimal ones that conform to *Automaticity* and non-minimal ones that conform to

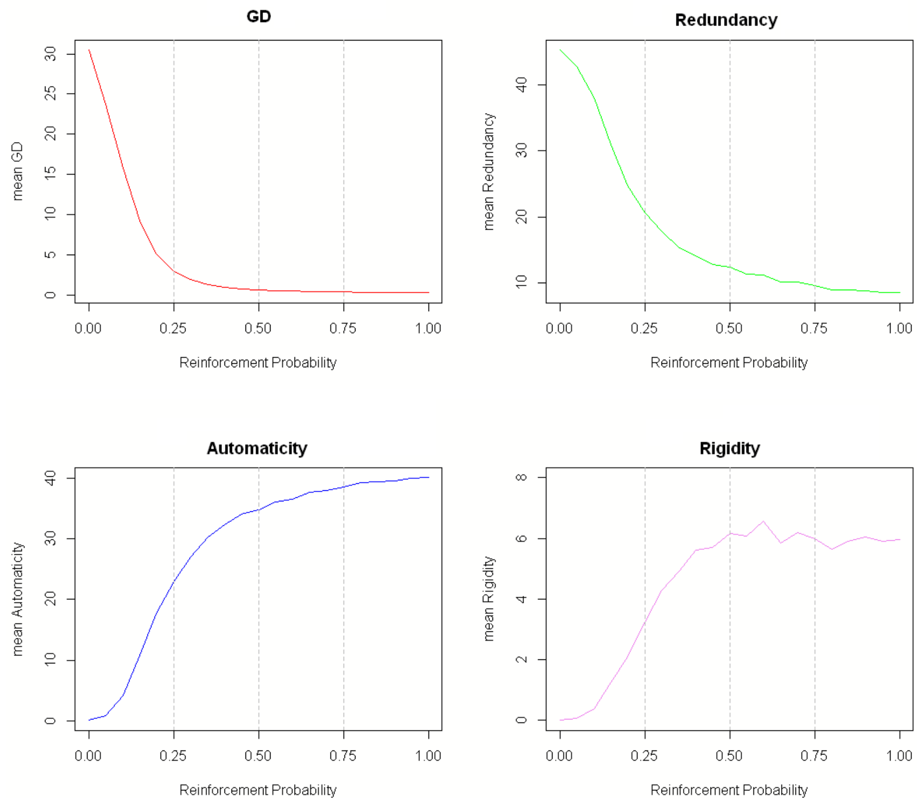

Fig. 5. Average Error Management metrics versus *ReinforcementProbability*

Rigidity. It is also worth noting that *Rigidity* drops slightly within the mentioned range, probably due to *Automaticity* raising, that overtake also non-minimal repetitions.

6 Discussion

The cognitive model was able to learn simple rules in the non-random scenarios, thereby showing that we developed a simple learning model that could be usefully compared to the behaviour generated by humans. When presented with random scenarios, it spontaneously generated redundant and goal-demoted behaviour characteristic of ritualised behaviour, with both redundancy and goal-demotion occurring more commonly in conditions that also favour that behaviour in humans. Rigidity proved more elusive. While random scenarios did indeed generate rigidity, they did not do so in the pattern met with in humans, where rigidity tends to co-vary with goal-demotion and redundancy. However, this difference in results is to be understood in terms of the limitations of the learning model used rather than as a problem with the main hypothesis. The reason is that, when tested, humans tend to rigidify their behaviour in a manner that was not open to the CM tested here. Specifically, they tend to form series of paths that they then use repeatedly, believing that the rules determining success change from walk to walk in a set pattern. In effect, they rigidify not over individual paths but over sets of them. This was not possible for the CM we created, as it has no memory of paths used as its behaviour is merely determined on the basis of the move weights.

Keeping in mind this significant limitation, it is possible to conclude that the study has been successful in showing that ritualisation of behaviour, or at least some aspects of it, is simply caused by accidental operant conditioning working in a scenario that is unpredictable - in line with the explanation that Skinner gave for superstition. As such, superstition and ritual must be seen as having at least in part a common epistemic basis that means that we should expect such behaviour in any system capable of learning when it is placed in the relevant set of conditions. This means in particular, that supernatural beliefs that commonly are connected to rituals are not basic to them, nor is anxiety necessary to generate ritualised behaviour.

The behaviour of the CM is also interesting in view of error management theory. When dealing with relatively easy random scenarios in which success was highly likely, the CM rapidly ended up automatising its behaviour while low success rates led to much greater innovation. To fully explore this aspect of ritualisation, however, it would be necessary to make each move costly thereby putting the CM in a position where it has to decide whether to use a more costly non-minimal path that may in some scenarios be required to earn points.

In future modelling work we will explore spontaneous ritualisation of behaviour using more complex cognitive models capable of rigidifying over sets of paths and then create an agent-based model to explore the social conditions necessary to sustain a ritual-tradition.

Bibliography

[1] AnyLogic: Simulation Modeling Software Tools & Solutions (2023), https://www.anylogic.com/

[2] Beck, J., Forstmeier, W.: Superstition and Belief as Inevitable By-Products of an Adaptive Learning Strategy. Human Nature 18, 35–46 (03 2007)

[3] Boyer, P., Liénard, P.: Why ritualized behavior? Precaution Systems and action parsing in developmental, pathological and cultural rituals. Behavioral and Brain Sciences 29(6), 595 – 613 (2006)

[4] Foster, K.L., Kokko, H.: The Evolution of Superstitious and Superstition-Like Behavior. Proceedings of The Royal Society B: Biological sciences 276, 31–37 (10 2009)

[5] Haselton, M., Buss, D.: Error management theory: A new perspective on biases in cross-sex mind reading. Journal of personality and social psychology 78, 81–91 (02 2000)

[6] Heltzer, R., Vyse, S.: Intermittent Consequences and Problem Solving: The Experimental Control of "Superstitious" Beliefs. The Psychological record 44 (04 1994)

[7] Killeen, P.: Superstition: A Matter of Bias, Not Detectability. Science (New York, N.Y.) 199, 88–90 (02 1978)

[8] Ono, K.: Superstitious behavior in humans. Journal of the experimental analysis of behavior 47, 261–71 (06 1987)

[9] Skinner, B.F.: Superstition in the pigeon. Journal of experimental psychology 38 2, 168–72 (1948)

[10] Talmont-Kaminski, K.: Malinowski's magic and skinner's superstition: Reconciling explanations of magical practices, chap. Mental culture: Classical social theory and the cognitive science of religion, pp. 98–109. Acumen Publishing (01 2012)

[11] Vyse, S.A.: Believing in magic: the psychology of superstition. Updated edition. New York, Oxford University Press. (2014)

[12] Vyse, S.: Behavioral variability and rule generation: General, restricted, and superstitious contingency statements. The Psychological Record pp. 487–506 (01 1991)