A cybernut to crack: an agent based model of actions copiers and results copiers

Alberto Acerbi[1], Heinz Gretscher[1] and Claudio Tennie[1]

[1]Max Planck Institute for Evolutionary Anthropology, Deutscher Platz, 6, 04103 Leipzig, Germany
alberto.acerbi@eva.mpg.de

Abstract

We propose a simple agent based model in which individuals can learn from a model by “copying” its actions or by “copying” its results. Our goal is twofold: firstly, we show how these different strategies, together with the strength of the coupling between actions and corresponding results, can impact, in our model, the final distribution of actions in the population. In particular, we demonstrate how populations of actions copiers tend to produce homogeneous action ‘cultures’, while result copiers are more dependent on the strength of the coupling between actions and results. Secondly, we propose some hypothesis on how our outcomes can be compared to observed real-world scenarios and how this may sometimes help field work researchers to infer, from observed population behaviors, the learning mode that underpin them.

Introduction

Social transmission of behavior can be realized through numerous, distinct mechanisms [22]. Here, we concentrate on a cognitive advanced subset of social learning, subsumed under the term ‘copying’ [10, 22]. Humans rely heavily on copying [19]. Copying also plays some role outside the human world in primates [21, 24], and possibly in other species (for example, for dolphins, see [7]; for birds, see [25]).

A growing interest in formal models of social learning is gradually spreading to animal behavior [6], evolutionary anthropology [14], robotics and scholars generally interested in the simulation of adaptive behavior (see chapters in [11]). Modeling strategies are particularly appealing here, as social learning typically depend on several processes at different scales as individual processes of learning (shaped by genetic evolution), characteristics of the environment, and characteristics of the population [9, 12]. For such reasons social learning is often difficult to study in controlled empirical conditions as well as in natural observation.

Models of ethologists and anthropologists are usually inspired by evolutionary biology in their re-adapting the mathematical tools being used in evolutionary genetics to fit social learning dynamics [14]. Results from these kinds of models are extremely important for developing a study of social learning compatible with the natural sciences. Those researches are shedding light on central topics such as the relative effectiveness of social- versus individual learning in variable environmental conditions as well on the different possible strategies in the choice of models to copy from (general rules of thumb like "copy from the majority" or "copy from the best individual" and so on) [6, 14]. However, these studies generally assume that a behavior is simply a quantitative variable passed on from one individual to another, with some probability to acquire it correctly or not (see examples in [14]).

Recent empirical studies in comparative psychology have instead shown that social learning is not one single phenomenon, that the transmission of behaviors can happen in different ways, and that the specific mechanisms involved in social learning can impact the resulting diffusion dynamics at population level [18, 22].

In particular, current primate social learning studies [5, 8] make a distinction between two forms of extracting information from a model [3]: copying its *actions* (body movements and body part relations during demonstrations), and copying the consequential *results* (changes in the environment brought about by these actions). This distinction is crucial because these two types of copying are very different in their propensity to allow for the accumulation of culture [19]. Recently it has been shown that human children can copy both types of information, which demonstrates that the distinction is a valid one [16]. [16] was conducted in a lab situation in which this problem can be adequately controlled (see also [8]); however, such controls are not available to field workers. Thus, field workers studying social learning are forced to undertake very labor intensive long-term studies in order to try to pinpoint the learning mechanisms shown by their studied populations [13].

Here, we imagine a population of animals who open a certain type of nuts (“cybernuts”). Further, we assume that these cybernuts can be potentially crackable in two distinct ways (e.g. twisting its two halves into different directions vs. crushing the whole nut between the hands). Such types of problems are now near-standard in social learning studies where they are referred to as the two-target method [4]. In our example, if one observes a population using only the twisting action, for which social learning mechanisms can one conclude? The answer is tricky: subjects might have either copied each others actions or the consequential results (i.e. the rotation of its halves against each other or the cracking of the shell). There is also (at least) a third and troublesome possibility: there can sometimes be a genetic propensity (or some affordances of the task) to learn only one specific way, even though potentially there are two. Even if single individuals never stumble across this single solution on their own, they might all converge on this method simply by having their *attention* directed after having observed others (this is known as local or stimulus enhancement [15, 17])

Our experiment is an initial attempt to introduce these mechanisms into a simple and abstract computational model (i.e., actions- or results-copiers plus individual learners). We then evaluate the different action-outcomes produced. In designing the experiment we used the two-target method. In other words, our agents are dealing with "cybernuts" which can be potentially opened with two different actions, leading to two different results. Importantly, while our model is abstract, it makes a very realistic assumption that was previously neglected in social learning studies: it introduces and varies a coupling factor between actions and results.

Method

Simulated populations are composed by 100 agents. Each agent is defined by three parameters: *A* (its current action, that varies through the simulation), *R* (its current result, that varies through the simulation), and *L* (its learning "mode", that is fixed). There are two possible actions, *A0* and *A1*, leading to two corresponding results, *R0* and *R1*, respectively. There are three possible learning modes, *IL* (individual learners), *AC* (actions copiers), and *RC* (results copiers). Populations are always homogeneous in respect to the learning mode, that is, given a single population, *all* agents are actions copiers, results copiers, or individual learners.

A main assumption underlying our model is that the coupling between actions and results, and vice versa, is not perfectly reliable. Indeed, depending on the level of an individual's proficiency and on the manipulated object's properties, employing the same action can potentially lead to a different outcome (e.g. a fragile cybernut-shell could be crushed while a subject tries to twist it open). Therefore we included two parameters (*A0→R0*, *A1→R1*), varying into the range [0; 100], that define the probability that an action has as outcome the corresponding result (i.e. that *A0* leads to *R0* and that *A1* leads to *R1*). If, for example, in a particular experimental condition, *A0→R0* is equal to 75, this means that an agent that performs A0 has the 75% of probability to obtain R0, and the 25% of probability to obtain R1. In the simulations, as we will explain later, we systematically vary the values of the two parameters from 0 to 100, with increases of 1 unit.

Accordingly, we also expected that observing a result can lead *RC* subjects to produce the "alternative" action. Whether an individual will end up using the corresponding action largely depends on the degrees of freedom in its motor system (behavioral flexibility) and on its personal pre-experience with manipulating similar objects (e.g. an individual used to twist open objects will be more likely to infer a twisting motion from two opened nutshell halves than an individual totally inexperienced in twisting). Therefore we also introduced the idea of inaccurateness for the process of "action acquisition". In an analogous way as described above for actions, two parameters (*R0→A0*, *R1→A1*), varying into the range [0; 100], define the probability that an agent, if it "sees" a result, will perform the correspondent action (i.e. observing *R0* it will perform *A0* and observing *R1* he will perform *A1*). As for *A0→R0* and *A1→R0*, we systematically vary the values of the two parameters from 0 to 100, with increases of 1 unit.

At the beginning of each simulation run all agents of the population are initialized with the same learning mode. For each agent an action is randomly chosen (so that at the beginning the two actions are approximately equally distributed in the population), and, depending from the parameters *A0→R0* and *A1→R1*, the results are assigned to each agent.

A run of the simulation consists in 100000 steps. For each step, one agent is randomly picked up in the population. Depending form the learning mode of the agents, learning takes place in three possible ways:

1. IL populations. The agent performs one of the two actions with the same probability. Depending on the value of parameters *A0→R0* and *A1→R1* it obtains a result.

2. AC populations. A model is also randomly chosen in the population; the agent copies the action of the model and performs it. Depending on the value of parameters *A0→R0* and *A1→R1*, it obtains a result.

3. RC populations. A model is also randomly chosen in the population; the agent copies the result of the model, and, depending on the value if parameters *R0→A0* and *R1→A1* it performs one of the two actions. As above, depending on the value of parameters *A0→R0* and *A1→R1*, it obtains a result. The rationale underlying this is (1) that regardless of the learning mode, individuals will always learn motor routines and (2) that result copiers will not keep trying to bring about the observed results if they are (accidentally) successful with a different one.

We analyze the final composition of the population, in terms of the distribution of actions, after having systematically varied the values of the four parameters *A0→R0*, *A1→R1*, *R0→A0*, and *R1→A1*. We focus on actions outcomes, since actions are, in real world studies, easier to observe and since they are at the focus of current social learning research [20, 26]. In varying parameters, we concentrate on some possible combinations, varying two parameters for each condition (leaving the other two equal to 100, i.e. with a perfectly reliable coupling). We have a total of four experimental conditions:

A. varying the probability to obtain results from both actions (we varied *A0→R0* and *A1→R1*).

B. varying the probability to obtain the result performing one action and the probability to infer, from the *corresponding* result, *that* action (we varied *A0→R0* and *R0→A0*).

C. varying the probability to obtain the result performing one action and the probability to infer, from the *opposite* result, the *other* action (we varied *A0→R0* and *R1→A1*).

D. varying the probability to infer the results from the two actions (we varied *R0→A0* and *R1→A1*).

Hence, we have four experimental conditions for each of the three homogeneous populations (individual learners, actions copiers, and results copiers), for a total of 12 experiments. For each experiments, as we varied at the same time two parameters we have a total of 10000 (100 X 100) different outcomes.

For each condition we present a graph showing the final distribution of the actions, (averaged on 100 simulation runs) and the corresponding standard deviation between populations.

Results

Individual learner populations

In the case of individual learners the outcomes of the simulations show that the variation of actions/result coupling and vice versa does not affect the outcomes. Regardless of any values of the parameters, the final composition, in terms of actions, of the populations, is the same. Therefore we present here the plot of just one condition (A), in which we varied *A0→R0* and *A1→R1*: the results of conditions B, C, and D, are exactly the same.

In the individual learner case, both action strategies survive in each population and they show a 50% distribution. The standard deviation is close to zero for all populations, which means that the final outcome of the different populations is fairly similar. This indicates that the 50% distribution is a sign of a mixture of action strategies *in the same* population. This outcome is easy to explain, as both actions are equally adaptive and each agent has no preference for either action or result.

Actions copiers populations

As for populations of individual learners, the variation of actions/results coupling and vice versa does not affect the outcomes. Again, we present here only the plot of condition A.

Like in the individual learners' case, the final average distribution of the two actions is close to 50%. However, the standard deviation plot is different, showing a higher standard deviation. This indicates that the 50% distribution, in the case of action copiers populations, is due to the fact that, starting from the same initial conditions *different* populations converge in one or in the other action strategy, depending on random initial effects. In other words, as soon as one action type slightly prevails on the other, it quickly "invades" the population.

This means that action copiers populations, unlike the individual learners ones, tend to become homogeneous in terms of the action used, but, unlike the result copiers (see below), the coupling actions / results does not affect toward which of the actions they tend to converge.

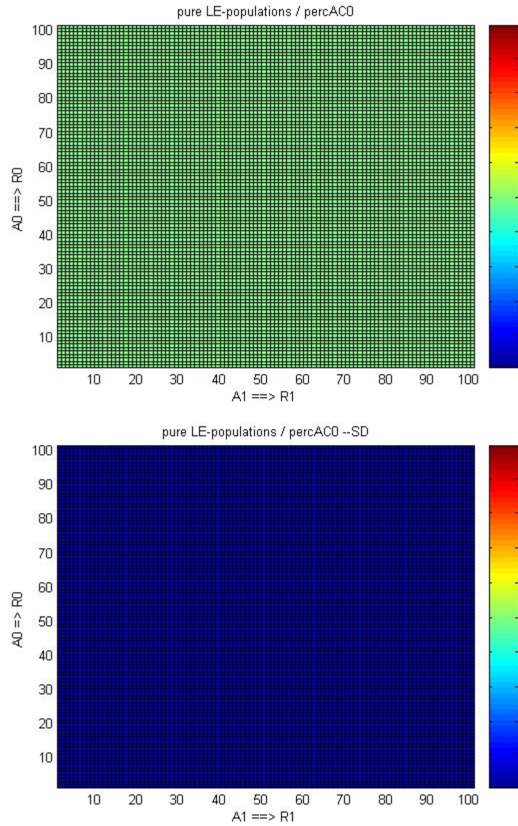

Fig. 1. Average final distribution of actions, expressed as a percentage of A0 agent at the end of simulations (*top*) and standard deviation of the same dataset (*bottom*), for populations of individual learners (see details in the text). X-axis shows the value of parameter *A1→R1*, and Y-axis the value of parameter *A0→R0* (experimental condition 1A)

Results copiers populations

3A. Outcomes of results copiers populations are the most complex. When varying the coupling between the two actions and the two results we can observe a fairly linear distribution of the final outcomes of actions in the populations.

When one action is highly correlated with the correspondent result and the other is not, all the populations converge on solving the problem using the correlated action (see high-left and low-right corners in the distribution graph). The convergence linearly decreases reaching a 50% distribution when the coupling is the same for both action/result couplings. Interestingly, however, the standard deviation is not the same in the high-right corner of the graph and in the low-left: in the latter, indeed, standard deviation is, as in the

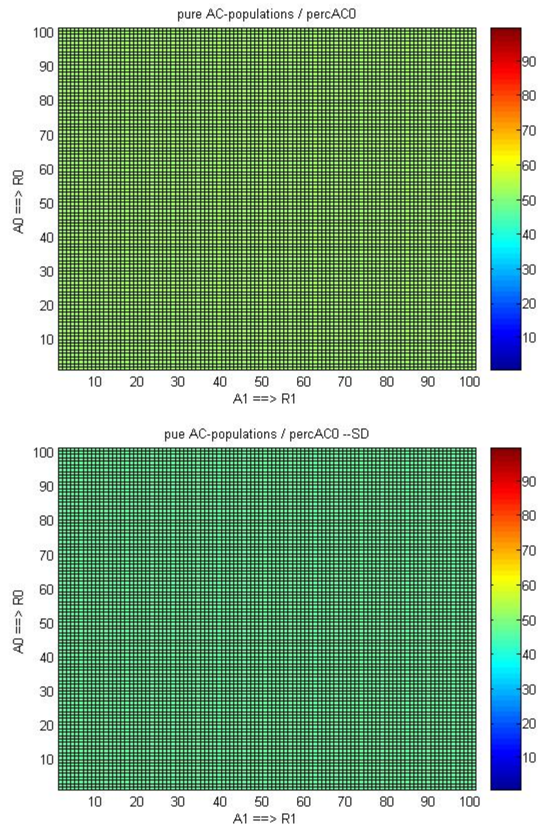

Fig. 2. Average final distribution of actions, expressed as a percentage of A0 agent at the end of simulations (*top*) and standard deviation of the same dataset (*bottom*), for populations of actions copiers (see details in the text). X-axis shows the value of parameter *A1→R1*, and Y-axis the value of parameter *A0→R0* (experimental condition 2A)

rest of the graph, close to zero, signaling, as in the individual learning case, a mixture of action strategies in the same population, when the two actions are both poorly correlated with the corresponding results. In the high-right corner, instead, where both the actions are highly correlated with the corresponding results, the standard deviation is higher as was in the action copier case, hence producing homogeneous populations very sensitive to initial random effects.

3B. When varying the coupling between one action and the result and the coupling with the *corresponding* result with the same action, independently from the entity of the variation, all populations converge on using the other action.

3C. When varying the coupling between one action and its result and between the *other* result and its action, the outcome is, in some respect, similar to condition 3A.

However, an important difference can be noticed. Here, starting from the two "convergent" corners, the decreasing is not linear, but one of the two actions, namely the action for which the variation is related to the difficulty in producing the correct result (opposite to the action that is difficult to *infer* starting from the result that this action should produce), is more robust than the other.

In other words, also if the correlation between one action and its result is low (low part of graph), as long as there is a tendency to perform that action also when observing the other result (that is, when the correlation between the other result and its action is low, left part of the graph) this action becomes stable in the population.

3D. When varying the coupling between the two results and the two actions we obtain a result that is basically the same of condition 3A.

Discussion

The main outcome of our model is that the actions used by populations of results copiers are, in general, most dependent on the degree with which actions are coupled with results, in comparison to the two other populations. In our model this is related to the fact that the outcomes of population of individual learners and of actions copiers are not dependent on the values of the two parameters *R0→A0* and *R1→A1*. Hence populations of results copiers are subjected to a "double loop" in each learning act that produces more variation at the population level. In the real world, great apes are now generally thought to be results copiers more than actions copiers [22]. However, this is still debated, and even studies using lab situations produce mixed results [16, 8]. Our result suggest that, if it is true that great apes are (predominantly or exclusively) results copiers, then the resulting extra complexity makes it hard to pinpoint the learning mechanism a particular population might use just by observing the spread of action techniques. In real life situations, it would be very difficult to get a quick estimate for the coupling strength of specific actions and results for a two-target task, thus field workers, in all likelihood, will not be able per se, and without the help of lab studies or modeling, to at least distinguish between possible different "worlds" of social learners.

Individual learner populations in our example started, in terms of actions, mixed and ended mixed. This is not surprising given the assumptions we put in our model. For real world field workers, this individual learners data mean that if a trait is very mixed and no consistencies are found between models and observers, then copying is unlikely to account for the trait in question. This is already well known. However, under special circumstances, also result copiers can show a mixed pattern, namely result copiers with *low* levels of coupling correlation involving both actions or both results (see low-left corners in the plots of conditions 3A and 3C). We believe the reason for this similarity is that the more we relax the coupling in result copiers the more they become like individual learners in the following sense: since result copiers start working their way towards their own produced results in a backward fashion (i.e. observed result leads to own action selection leads to own result produced) more copying effects are being lost on this longer path than is true for action copiers. This will then result in a mixed use of actions over

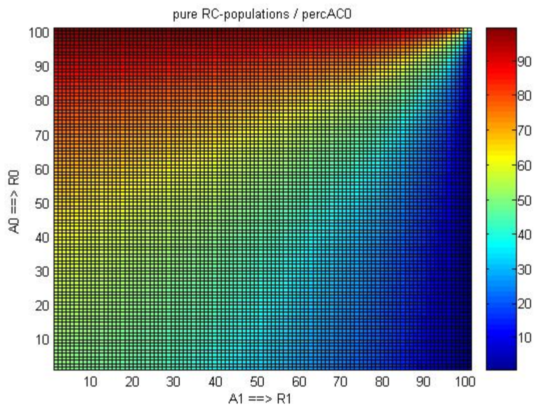
pure RC-populations / percAC0
A0 ==> R0
A1 ==> R1

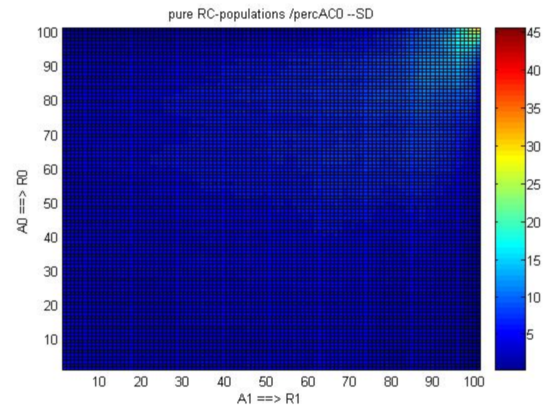
pure RC-populations /percAC0 --SD
A0 ==> R0
A1 ==> R1

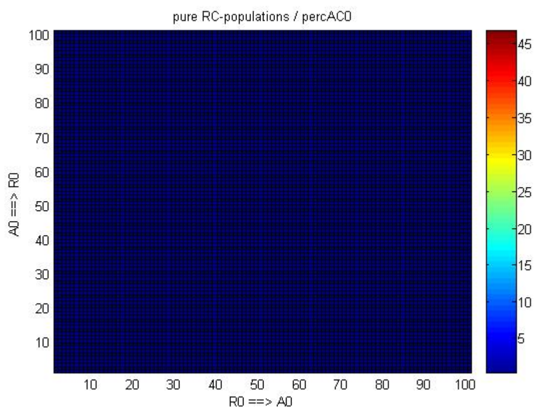
pure RC-populations / percAC0
A0 ==> R0
R0 ==> A0

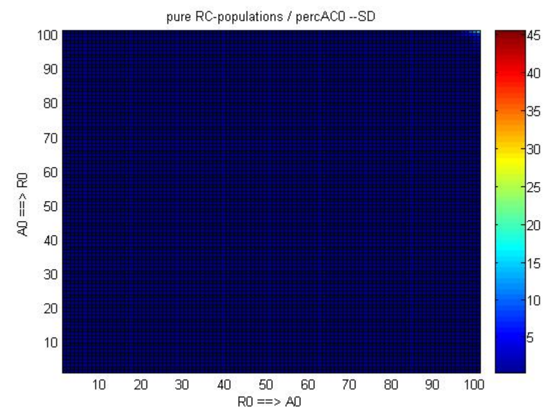
pure RC-populations / percAC0 --SD
A0 ==> R0
R0 ==> A0

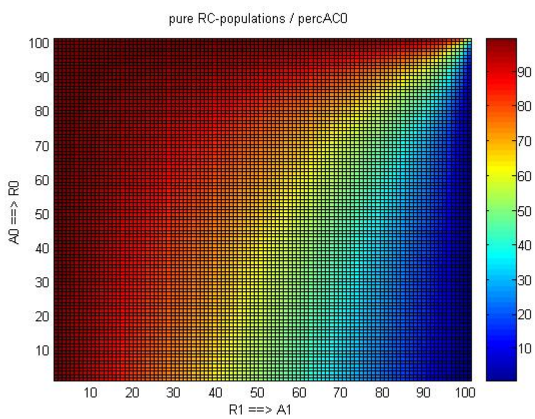
pure RC-populations / percAC0
A0 ==> R0
R1 ==> A1

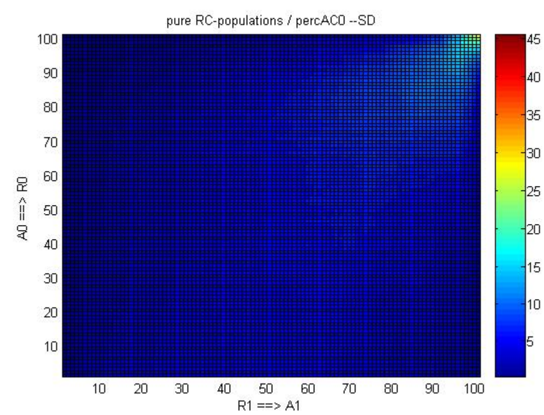
pure RC-populations / percAC0 --SD
A0 ==> R0
R1 ==> A1

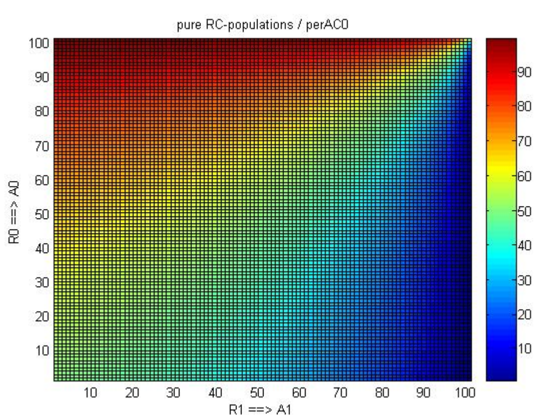
pure RC-populations / perAC0
R0 ==> A0
R1 ==> A1

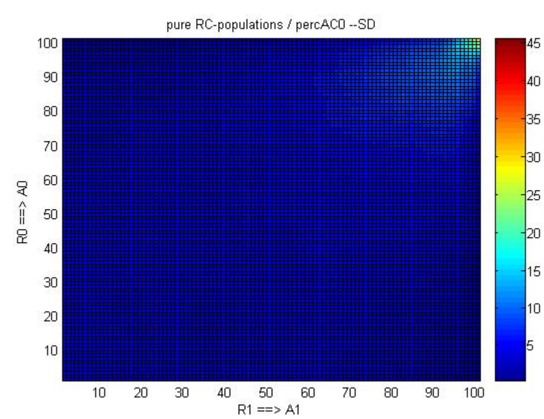
pure RC-populations / percAC0 --SD
R0 ==> A0
R1 ==> A1

Fig. 3. Average final distribution of actions, expressed as a percentage of A0 agent at the end of simulations (*left*) and standard deviation of the same dataset (*right*), for populations of results copiers (see details in the text). Outcomes of conditions 3A, 3B, 3C, and 3D are shown from high to low. 3A: X-axis shows the value of parameter $A1 \rightarrow R1$, and Y-axis the value of parameter $A0 \rightarrow R0$. 3B: X-axis shows the value of parameter $R1 \rightarrow A1$, and Y-axis the value of parameter $A0 \rightarrow R0$. 3C: X-axis shows the value of parameter $A1 \rightarrow R1$, and Y-axis the value of parameter $A0 \rightarrow R0$. 3D: X-axis shows the value of parameter $R1 \rightarrow A1$, and Y-axis the value of parameter $R0 \rightarrow A0$.

time. This pattern would mask copying by looking like having come about through enhancement effects (see Introduction).

We must also note that individual learners might not always produce a pattern like the one we observed in experiment 1A. Like we mentioned in the Introduction, the observed pattern can also depend on whether a genetically pre-channeled trial-and-error acquisition (possibly via affordance learning) tends to favor one particular method over the other (and in turn one particular action). In an extreme case, the whole individual learner population might in fact be induced to use only one method. These cases might then be mistaken for population of true action copiers. As a heavily copying species we instinctively recognize sameness in action patterns (compare [2]). In fact, this is what renders studying social learning so hard: the human eye anthropomorphizes matching behavior and instantly attributes it to copying, while in fact many other mechanisms can bring about matching behavior between model and observer as well (see Introduction). However, we believe this scenario is rarely the case for individual learners, or when it is it should be relatively easy to identify an imbalance of the task in question: when a real "nut" naturally gives rise to just one method, then it should not be concluded for copying. Such a "natural drive" can be easily tested in baseline lab studies, where naïve individuals are presented with the task (though, for motivational reasons, some enhancement neutral with regard to the two methods might still be necessary). If subjects tested under these conditions drastically prefer one of the two possibilities, then the behavior of their wild conspecifics is unlikely to be transmitted via copying.

Maybe the most important finding of our work is that actions copiers robustly produce homogeneous action populations in relation to any coupling strength of actions and results. Thus, all other things being equal and both target methods equally likely to be developed in baseline situations, if one observes a population where all subjects solely use the very same action then this population is inherently more likely to be an actions copier population than a results copier or an individual learner population. Thus, field researcher should, as they do, watch out for such populations, and also establish in lab situations that really two outcomes to the task are a) possible and b) realized by naïve individuals (i.e., that there is no "natural drive" attached to the task). Of course, under special circumstances, even a homogeneous action population might not actually be an action copier population. Namely with result copiers with *high* levels of coupling correlation involving one of the actions (or one of the results) one will also observe the same pattern of homogeny. Future empirical research will need to directly address the question of coupling strength between actions and results.

Conclusion

With an abstract agent – based model we were able to show that, in contrast to both actions copiers and individual learners, the results copier populations are highly sensitive to changes in coupling strength between actions and results. We believe varying coupling strength is a realistic assumption for two-target tasks. Our analysis further shows an inherent difference between actions copiers and results copiers. While the former robustly give rise to homogeneous action 'cultures', the latter depend very much on accuracy before they, too, would produce such type of 'culture'. Thus, if such cultural traits have to evolve, then either actions copiers or very accurate result copiers should be favored. Individual learners in our analysis and with our setting were unable to produce such patterns.

In future we plan to extend the model, in particular analyzing dynamics for non homogeneous populations, and using evolutionary computation techniques to let the social learning strategies evolve. More "computationally" complex action chains could be also employed, making our assumptions more realistic and permitting a better comparison with real world dynamics.

Acknowledgments

We would like to thank Joanna Bryson, Mathias Franz, and Charlie Nunn for discussion on related issues.

References

1. Alissandrakis, A., Nehaniv, C. L., Dautenhahn, K. (2004). Towards Robot Cultures? - Learning to Imitate in a Robotic Arm Test-bed with Dissimilar Embodied Agent, *Interaction Studies,* 5(1): 3-44.
2. Asendorpf, J. B., Warkentin, V., Baudonnière, P.-M. (1996). Self-Awareness and Other-Awareness II: Mirror Self-Recognition, Social Contingency Awareness, and Synchronic Imitation, *Developmental Psychology*, 32(2): 313-321.
3. Call, J., Carpenter, M. (2002). Three Sources of Information in Social Learning. In Dautenhahn, K., Nehaniv, C. L. (eds.), *Imitation in Animals and Artifacts*, pp. 211-228. MIT Press. Cambridge, MA.
4. Dawson, B. V., Foss, B. M. (1965) Observational Learning in Budgerigars. *Animal Behaviour* 13(4): 470-474.
5. Dindo, M., Thierry, B., Whiten, A. (2008). Social Diffusion of Novel Foraging Methods in Brown Capuchin Monkeys (Cebus Apella). *Proceedings of the Royal Society Biological Sciences Series B* 275(1631): 187-193.
6. Galef, B. G., Laland. K. N. (2005). Social Learning in Animals: Empirical Studies and Theoretical Models. *Bioscience* 55(6): 489-500.
7. Herman, L. M.: Vocal, Social, and Self-Imitation by Bottlenosed Dolphins (2002). In Dautenhahn, K., Nehaniv, C. L. (eds.), *Imitation in Animals and Artifacts,* pp. 63-108. MIT Press. Cambridge, MA.
8. Hopper, L. M., Spiteri, A., Lambeth, S. P., Schapiro, S. J., Horner, V., Whiten, A. (2007). Experimental Studies of Traditions and Underlying Transmission Processes in Chimpanzees. *Animal Behaviour* 73: 1021-1032.
9. Marocco, D., Acerbi, A. (2007). Adaptation and Social Facilitation in a Population of Autonomous Robots. In Berthouze et al. (eds.)

Proceedings of the 7th International Conference on Epigenetic Robotics, Lund, Lund University Cognitive Series.
10. Miklósi, Á. (1999). The Ethological Analysis of Imitation. *Biological Reviews of the Cambridge Philosophical Society* 74(3): 347-374.
11. Nehaniv, C. L , Dautenhahn, K. (2007) *Imitation and Social Learning in robots, Humans, and Animals.* Cambridge University Press. New York.
12. Noble, J., Todd, P. M., Tuci, E. (2001).. Explaining social learning of food preferences without aversions: an evolutionary simulation model of Norway Rats. *Proceedings of the Royal Society B: Biological Sciences*, 268(1463): 141-149.
13. Perry, S., Panger, M., Rose, L., Baker, M., Gros-Louis, J., Jack, K., (2003). Traditions in Wild White-Faced Capuchin Monkeys. In Fragaszy, D. M., Perry, S. (eds.), *The Biology of Traditions: Models and Evidence*, pp. 391-425. Cambridge University Press. New York
14. Richerson, P. J., Boyd, R. (2004). *Not By Genes Alone: How Culture Transformed Human Evolution.* Chicago University Press, Chicago.
15. Spence, K. W. (1937). Experimental Studies of Learning and the Higher Mental Processes in Infrahuman Primates. *Psychological Bulletin* 34(10): 806-850.
16. Tennie, C., Call, J., Tomasello, M. (2006). Push or Pull: Imitation vs. Emulation in Great Apes and Human Children. *Ethology* 112 (12): 1159-1169
17. Thorpe, W. H. (1963). *Learning and Instinct in Animals* (2nd ed.). Methuen, London.
18. Tomasello, M. (1996). Do Apes Ape? In Heyes, C. M., Galef, B. G. (eds.), *Social Learning in Animals: The Roots of Culture*, pp. 319-346. Academic Press. San Diego.
19. Tomasello, M. (1999). *The Cultural Origins of Human Cognition*. Harvard University Press, Cambridge.
20. Voelkl, B., Huber, L. (2000) True Imitation in Marmosets. *Animal Behaviour*, 60: 195-202.
21. Whiten, A., Goodall, J., McGrew, W. C., Nishida, T., Reynolds, V., Sugiyama, Y., et al. (1999). Cultures in Chimpanzees. *Nature* 399(6737): 682-685 (1999)
22. Whiten, A., Horner, V., Litchfield, C. A., Marshall-Pescini, S. (2004). How Do Apes Ape? *Learning & Behavior* 32(1): 36-52.
23. Whiten, A., Horner, V., de Waal, F. B. M. (2005). Conformity to Cultural Norms of Tool Use in Chimpanzees. *Nature* 437(7059): 737-740.
24. Whiten, A., Spiteri, A., Horner, V., Bonnie, K. E., Lambeth, S. P., Schapiro, S. J., et al. (2007). Transmission of Multiple Traditions Within and Between Chimpanzee Groups. *Current Biology*, 17: 1038-1043.
25. Zentall, T. R. (2004). Action Imitation in Birds. *Learning & Behavior*, 32(1): 15-23
26. Zentall, T. R., Sutton, J. E., & Sherburne, L. M. (1996). True Imitative Learning in Pigeons. *Psychological Science* 7(6):343-346.