Binary Consensus via Exponential Smoothing

Marco A. Montes de Oca[1], Eliseo Ferrante[2,3], Alexander Scheidler[4], and Louis F. Rossi[1]

[1] Department of Mathematical Sciences, University of Delaware, USA
{mmontes,rossi}@math.udel.edu,
[2] IRIDIA, CoDE, Université Libre de Bruxelles, Belgium
[3] Socioecology and Social Evolution Lab, Katholieke Universiteit Leuven, Belgium
eferrant@ulb.ac.be,
[4] Fraunhofer IWES, Germany
alexander.scheidler@iwes.fraunhofer.de

Abstract. In this paper, we reinterpret the most basic exponential smoothing equation as a model of social influence. Exponential smoothing has been used as a time-series forecasting and data filtering technique since the 1950s. The most basic exponential smoothing equation, $S^{t+1} = (1-\alpha)S^t + \alpha X^t$, is used to estimate the value of a series at time $t+1$, denoted by S^{t+1}, as a convex combination of the current estimate S^t and the actual observation of the time series X^t. In our work, we interpret the variable S^t as an agent's tendency to imitate the observed behavior of another agent, which is represented by a binary variable X^t. We study the dynamics of the resulting system when the agents exhibit a behavior for a period of time, called latency, whose duration is stochastic. Latency allows us to model real-life situations such as product adoption, or action execution. When different latencies are associated with the two different behaviors, a bias is produced. This bias makes all the agents in a population adopt one specific behavior. This phenomenon has applications in domains where it is desired that a group of agents adopts a behavior by consensus.

Key words: Consensus, Collective Decision-Making, Self-Organization

1 Introduction

The old adage "When in Rome, do as the Romans do" summarizes the intuition that it is sometimes wise to imitate the behavior of seemingly more knowledgeable individuals, especially when we are in a new environment. Recent research has provided evidence that imitating is indeed the best thing to do even in situations previously thought to require individuals to rely more on themselves [13]. When agents[1] are under pressure to choose an adaptive action, that is, an action that provides them with some benefit, imitation can be seen as a filtering process that allows agents to collectively discard actions that provide the lowest

[1] We use the term *agent* to refer to an entity, be it an animal or an artifact, such as a robot or a piece of software, capable of autonomous perception and action.

rewards [13]. This phenomenon may explain why in nature doing what others do is a strategy frequently used by different groups of animals. For example, ants choose the same paths other ants choose by following pheromone trails [7], sheep move in the same direction other sheep move [12], and we humans tend to cross the road whenever we see other people do so [4].

Given that imitation is a strategy used by animal groups as different as insect colonies and human crowds, we wonder to what extent individual imitation induces a good collective decision-making mechanism for groups of artificial agents. In particular, we would like to know whether imitation is an individual strategy that allows a large group of agents to make good collective decisions by consensus. As a step toward answering this question, in this paper we introduce and study the dynamics of an agent-based model in which individual agents tend to perform the action most commonly performed by other agents. In our model, presented in detail in Section 2, the rule that each individual uses to determine whether to perform a commonly observed action is similar to the basic exponential smoothing equation used for data filtering and time series forecasting [5, 8]:

$$S^{t+1} = (1 - \alpha)S^t + \alpha X^t \tag{1}$$

where S^t is the estimate of a time series at time t, X^t is the actual value of the time series at time t, and α is a parameter that determines the strength of the error correction.

In this paper, we interpret the variable S^t as the tendency of an agent to imitate an observed behavior, perform an observed action, or adopt a certain opinion. The variable X^t encodes the behavior, action, or opinion of another agent. The parameter $\alpha \in [0, 1]$ controls the strength of the imitation tendency. At the two extremes, if α is equal to zero, an agent does not change its tendency to imitate other agents; if α is equal to one, an agent copies whatever action another agent performs. Note that since the behavior of all the agents is governed by the same rule, the collective dynamics of the group is determined by the dynamics of a system of coupled exponential smoothing equations.

In Section 3, we study the dynamics of the system in the cases where agents may or may not influence other agents at all times. An example of a situation where an agent may influence other agents at all times is when someone decides to buy a product that is visible to others (e.g., someone buys a tablet computer that is frequently used). The very fact that the bought item is visible, gives information to an observer agent (e.g., that the tablet is useful) that may increase its tendency to buy the same product. There are also situations where agents may not influence others at all times. For example, in a robotics application, a robot that decides to go from one point to another is only visible to other robots while it is in their close vicinity. Thus, robot actions or choices may be observed by other robots only in a confined region of the robots' operating environment. The time that passes between two product adoptions or two executions of an action, is modeled as *latency*, which is a period of time of stochastic duration during which agents do not update their tendency or state. In our examples, a person who is evaluating a product or a robot executing an action are modeled as latent agents.

Our results are organized into two parts that correspond to the two cases described above (visible *vs.* not visible latent agents). We show that in both cases a population of agents adopts the same behavior, that is, they reach a consensus. However, the actual consensus state when latent agents are visible is the opposite than when latent agents are not. These results are discussed in Section 4 in which we also describe the similarities and differences that exist between our work and previously published works. Final conclusions are given in Section 5.

2 Exponential Smoothing as a Social Influence Model

The model proposed in this paper is based on the notion that the actions and opinions of others influence our own actions and opinions. When an agent is exposed to the behavior or the opinion of another agent, the observing/listening agent may be more likely to exhibit the observed behavior, or adopt the opinion of the other agent. In this paper, we focus on the case where there are only two observable behaviors/opinions. This model captures real-life scenarios where two options are available to the members of a population, but only one can be chosen. For example, a person has to choose between an Android or an iOS phone, a French voter must choose one of the two candidates in the second round of the elections, and in some cases, ants have to choose one of two paths between their nest and a food source.

Our model consists of a set of agents, each of which is in one of two possible states, which represent the agent's current behavior or opinion. We use the binary variable $X_i \in \{0, 1\}$ to represent an agent i's state. This variable is in turn governed by an internal real-valued variable S_i and a threshold θ. The variable S_i can be thought of as the tendency of agent i to be in one of the two possible states (hereafter, we refer to S_i simply as agent i's tendency). The threshold θ is constant and common to all agents, while S_i is variable and private to each agent.

At each time step t of the system's evolution, an agent i might be able to observe the state of another random agent $j \neq i$. When an agent observes the state of another agent, the observing agent updates its tendency as follows:

$$S_i^{t+1} = (1 - \alpha)S_i^t + \alpha X_j^t \,, \tag{2}$$

where $\alpha \in [0, 1]$ determines how much importance is given to the agent's latest observation (X_j^t) as opposed to the agent's accumulated experience (S_i^t). After updating its tendency, an agent updates its state as follows:

$$X_i^{t+1} = \begin{cases} 1, & \text{if } S_i^{t+1} > \theta \\ 0, & \text{if } S_i^{t+1} < 1 - \theta \\ X_i^t, & \text{if } 1 - \theta \leq S_i^{t+1} \leq \theta \,, \end{cases} \tag{3}$$

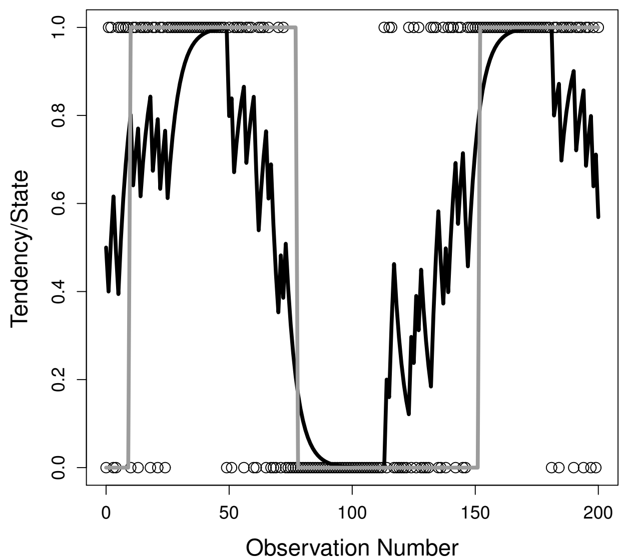

Fig. 1. Single agent behavior. Starting with an initialization of $S^0 = 0.5$ and $X^0 = 0$, an agent observes a stream of state values plotted as dots with values 0 or 1. The black line shows the evolution of the agent's tendency and the gray line shows the evolution of the agent's state. In this simulation, $\alpha = 0.2$ and $\theta = 0.8$.

where $\theta \geq \frac{1}{2}$ due to the symmetry of the actual threshold value that triggers the adoption of one or another state. Thus, an agent's state is a function of its tendency and the threshold θ.

In Fig. 1, we show an example of the behavior of a single agent controlled by the rules defined in Equations 2 and 3. In this example, the agent is exposed to a controlled stream of state observations that repeatedly switches between 1 and 0. Through this example, we observe that an agent's tendency follows the stream of observations and the threshold rule acts as a low-pass filter of the agent's tendency. Combined, these two rules make an agent adopt a state that agrees with the most commonly observed state at any point in time. This conclusion can be reached if one expands Eq. 2 iteratively as

$$S_i^n = (1-\alpha)^n S_i^0 + \alpha \sum_{k=0}^{n-1} (1-\alpha)^{n-k-1} X_{\mathbf{I}(k)}^k \,, \tag{4}$$

where S_i^n is the value of agent i's tendency after n updates, S_i^0 is agent i's initial tendency, and $X_{\mathbf{I}(k)}^k$ is the state of the agent observed at time step k. A multivariate random variable $\mathbf{I}$ is used to represent the fact that observed agents are chosen at random, and thus $\mathbf{I}(k)$ is the index of the agent observed at time step k. When $0 < \alpha < 1$, an agent's tendency is a weighted moving average of the agent's observations with exponentially decreasing weights. Any possible bias introduced by an initial tendency $S_i^0 \neq \frac{1}{2}$ vanishes when an agent performs a sufficiently large number of observations. Similarly, the weight of old observations approaches zero as the number of observations increases. This means that only the most recent observations have a significant influence on the observing agent's state. Thus, if the majority of the most recently observed states, that is, if the most commonly observed behaviors, are encoded by $X_{\mathbf{I}(k)}^k = 1$, the observing agent will have a state $X_i^{t+1} = 1$. The same reasoning applies if the majority of the most recently observed states are encoded by $X_{\mathbf{I}(k)}^k = 0$.

The collective behavior of a population of agents controlled by the rules given in Equations 2 and 3 is much richer and difficult to predict than the behavior of a single agent. In this paper, we explore this system's dynamics through Monte Carlo simulations. The analytical study of the system of coupled stochastic exponential smoothing equations that describes the system's dynamics is ongoing work. The simulation study, presented in the next Section, is focused on the effects that agent visibility, as explained in Section 1, has on the system's collective dynamics.

3 Simulations

We perform two sets of simulations. In both sets, when an agent changes state from 0 to 1, or *vice versa*, it does not update its tendency or state for a period of time of stochastic duration, called *latency* period [10, 11]. During this time, we say that an agent is *latent*. In the first set of simulations, latent agents are always visible to other agents, that is, they can influence other agents while being latent. This scenario models situations similar the iOS vs. Android example mentioned earlier. For example, after adopting one of the two competing brands, a buyer will not change immediately to the the other brand, but rather evaluate the newly-bought product for some time (which we model with the latency period). During this evaluation time, other persons will observe the buyer and may be more inclined to adopt the same product as a result of our natural tendency to do what others do. In the second set of simulations, latent agents are not visible to other agents, and therefore, cannot influence them. This scenario models situations like the robotics example mentioned in Section 1. In a robotics application, states may be interpreted as robot actions. The latency period models the duration of an action execution during which the robot may not be able to interact with other robots. Consequently, in this second case, latent robots cannot be observed by other robots.

An extra element in our simulations is known as *differential latency* [11]. This scheme associates a different average duration of the latency period to each of the two states of an agent. This is done in order to model situations where two products have different sets of features, causing a buyer to spend more time evaluating one of the competing products, or actions which have similar results but that take different times to perform. With differential latency we expect to alter the frequency with which one state is observed by the population of agents, inducing a consensus on one specific state. The simulation of differential latency proceeds as follows: While an agent is latent, it does not change state or tendency. The difference between the two sets of simulations that are reported in this section lies in the actions taken by the agents when their latency period ends. In the observable latent agents case (see Section 3.1), an agent chooses, from the whole population, a random agent to observe after its latency period ends. Then, the agent updates its tendency and becomes latent again for a period whose duration depends on its (possibly new) state. In the case of nonobservable latent agents (see Section 3.2), an agent whose latency period ends observes the state of a randomly chosen non-latent agent. If there is none available, that is, if all agents are latent, it waits until one becomes observable. This waiting time also allows other agents that switch from latent to non-latent state to observe this agent. As before, a state observation is used for updating an agent's tendency and state. The process of observing, updating tendency, updating state, and becoming latent is repeated until some stopping criterion is satisfied.

In our simulations, the duration of latency periods are normally distributed. The mean duration and standard deviation of the latency period associated with state 1, denoted by μ_1 and σ_1 respectively, are equal to 100 and 10 time steps. We keep the standard deviation of the latency period associated with the state 0 constant ($\sigma_0 = \sigma 1$) and test three different values for the mean duration of the latency period $\mu_0 = \{\mu_1, 1.2\mu_1, 2\mu_1\}$. A ratio $\mu_0/\mu_1 = 1$ is used to observe the system's dynamics when both states are associated with equal latency periods. The ratios $\mu_0/\mu_1 = 1.2$ and $\mu_0/\mu_1 = 2$ are used to observe the system's dynamics with high and low degrees of latency overlap respectively. We also explore the effects of different combinations of values of α and θ on the system's dynamics. A summary of our results is presented next.

3.1 Observable Latent Agents

Let us first present the results obtained when latent agents are always visible. In Fig. 2, we present a typical example of the evolution of the average tendency across a population of 100 agents under the three tested latency period conditions.

When the latency periods associated with the two states have equal average duration, the final average tendency is approximately 0.5. This is the result of approximately half of our simulations converging to a consensus on state 0 and the other half on state 1. Thus, we conclude that independently of the values of α, when $\mu_0/\mu_1 = 1$ the population of agents reaches a consensus on one of the two states with equal probability.

When the latency periods have different average duration, agents whose state is associated with the longest latency period are more likely to spread their state to the rest of the population because their state remains "frozen", and therefore visible, for longer periods of time. This phenomenon induces a positive feedback process whereby the state associated with the longest latency period is observed more often, and therefore copied more rapidly than the other state until every agent adopts the same state. In our simulations, state 0 is associated with the

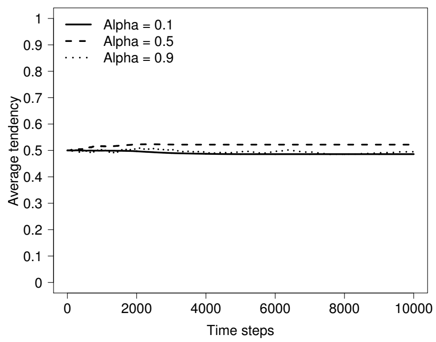

(a) $\mu_0/\mu_1 = 1.0$

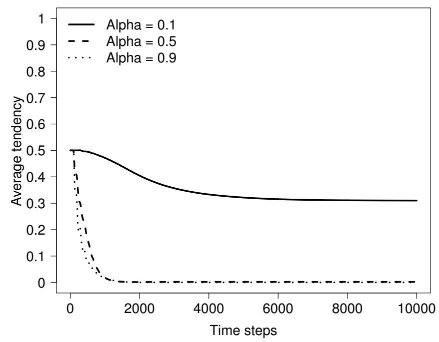

(b) $\mu_0/\mu_1 = 1.2$

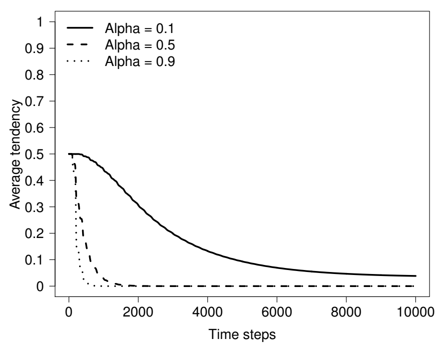

(c) $\mu_0/\mu_1 = 2.0$

Fig. 2. Average tendency over the whole population (100 agents). In these simulations, $\mu_1 = 100$, $\sigma_1 = \sigma_0 = 10$ time steps. The acceptance threshold θ is equal to 0.6 in all cases. Averages obtained through 500 independent runs of a Monte Carlo Simulation.

longest average latency period. In Fig. 2 parts (b) and (c), we can see that state 0 is the state that is more often adopted by the population. The effect of the parameter α depends on the overlap of the duration distribution of latency periods. With large overlaps (in our simulations, this is modeled with the case $\mu_0/\mu_1 = 1.2$), larger values of α have larger effects on the system's dynamics. In our simulations, with $\mu_0/\mu_1 = 1.2$ and $\alpha = 0.1$, in 69% of the cases the population reaches consensus on state 0, while with $\mu_0/\mu_1 = 1.2$ and $\alpha = 0.9$, the population reaches consensus on state 0 in 100% of the cases. With low overlaps (in our simulations, when $\mu_0/\mu_1 = 2$), the parameter α affects more the speed of convergence to consensus than the actual state on which the consensus is reached.

3.2 Nonobservable Latent Agents

We now present the results obtained when latent agents are not visible to other agents. In Fig. 3, we show the evolution of the average tendency accross a population of 100 agents under the three tested latency period distributions.

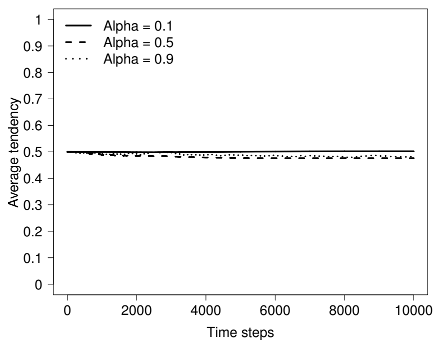

(a) $\mu_0/\mu_1 = 1.0$

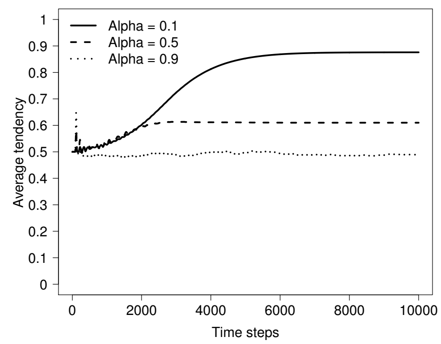

(b) $\mu_0/\mu_1 = 1.2$

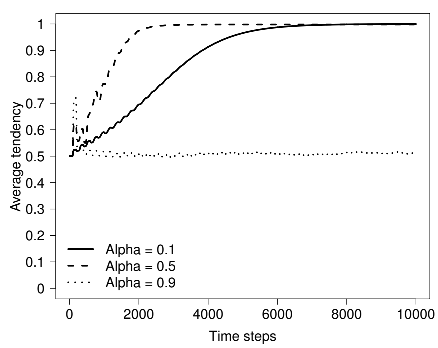

(c) $\mu_0/\mu_1 = 2.0$

Fig. 3. Average tendency over the whole population (100 agents). In these simulations, $\mu_1 = 100$, $\sigma_1 = \sigma_0 = 10$ time steps. The acceptance threshold θ is equal to 0.6 in all cases. Averages obtained through 500 independent runs of a Monte Carlo Simulation.

In Fig. 3 part (a), it is apparent that when the average duration of the latency periods associated with each of the two states is equal ($\mu_0/\mu_1 = 1$), the probability of the population to reach consensus on any of the two states is 0.5. Thus, irrespective of whether latent agents are visible all the time or only when they are in a nonlatent state, when the ratio $\mu_0/\mu_1 = 1$, a population of agents governed by Eqs. 2 and 3 initialized with $S_i^0 = 0.5$ will reach a consensus on any of the two states with equal probability.

In Fig. 3 parts (b) and (c), we see the results when $\mu_0/\mu_1 = 1.2$ and $\mu_0/\mu_1 = 2$, respectively. When $\mu_0/\mu_1 = 1.2$, the smaller the value of the parameter α, the higher the probability of the population to reach consensus on state 1, which is associated with the shorter average latency period. This phenomenon occurs because large values of α make agents favor copying over using their past experience. Therefore, when the average duration of the latency periods is similar, the agents switch from one state to the other more rapidly with large values of α than with small ones. The result is that the frequency with which different states are observed is increased and thus, the probability of reaching consensus on only one state is decreased. The results in part (c) of Fig. 3 can be explained using the same argument. In this case, the duration of the latency periods is sufficiently different to induce a strong bias toward the state associated with the shorter average latency period (in our case, state 1) even for $\alpha = 0.5$. However, just as in the previous case, a large value of α means that agents copy any observed state, which produces switches that lead the system to consensus on any of the two states with equal probability.

4 Discussion

In some situations, doing what others do may be beneficial to some agents but not necessarily to the whole group. For example, if food is clustered in patches and all the members of a group copy each other and exploit only one discovered patch, very soon food in that patch will be depleted. A better strategy would be to switch between exploitation and exploration behaviors to ensure everyone has enough food [6]. However, there are also cases where imitation results in the group reaching a state of consensus that benefits everyone in the group. For example, the aforementioned trail-laying and following behavior of ants allows them to find the shortest route from their nest to a food source [7]. The shortest route is not only more energetically efficient, it also reduces the exposure to predators. This last example also shows us that consensus can emerge even in very large groups, which means that in these cases, consensus is the result of multiple agent-to-agent interactions and not of agents knowing what everyone in the group does. Therefore, the ability of large groups to reach consensus on an action that benefits everyone in the group may be regarded as a form of collective intelligence.

In this paper, we explored some circumstances under which imitation makes a population of agents reach a consensus. In particular, we studied the effects of latency periods of different duration on the system's dynamics. We saw how imitation and latency induce a positive feedback process that eventually leads the population to a consensus state. Imitation and positive feedback have been already used successfully in some software-based swarm intelligence systems [1] for optimization. For example, in ant colony optimization (ACO) [3] algorithms, very simple agents simulate ants following pheromone trails laid by other ants.

In ACO, these artificial ants move on a graph that represents adjacency relations between solution components to a combinatorial optimization problem. At each decision point, ants are attracted toward nodes with higher levels of "pheromone", which is a numerical variable associated with each pair of solution components i and j, denoted by τ_{ij}. The result is an "emergent" optimization process. At each iteration of an ACO algorithm, these pheromone variables are updated using a rule that resembles the following equation:

$$\tau_{ij}^{t+1} = (1-\rho)\tau_{ij}^{t} + \Delta\tau \,, \tag{5}$$

where $0 \leq \rho \leq 1$ is a parameter that simulates "pheromone evaporation", and $\Delta\tau$ is a quantity that reinforces a pheromone value. The similarity between Equations 2 and 5 is apparent. However, in ACO algorithms, the quantity $\Delta\tau$ is different across problems and ACO variants. Such heterogeneity makes a unified approach to the study of the imitation dynamics that occur in ACO algorithms practically impossible. Nevertheless, our model enables the study of the abstract process of imitation and thus, any insights gained from its study may shed light onto the operation of ACO algorithms for specific problems.

Imitation is also an important idea behind the particle swarm optimization (PSO) [9] algorithm, which is another successful swarm intelligence algorithm for optimization. In a PSO algorithm, agents move in an n-dimensional space. Their positions represent candidate solutions to a continuous optimization problem. The stochastic rules used to update the particles' position make particles move toward the position of (i.e., imitate) their most successful neighbors. It can be shown that the expected value of the i-th particle's position obeys the equation:

$$E(\mathbf{x}_i^{t+1}) = w\mathbf{x}_i^{t} + \frac{\mathbf{L}+\mathbf{G}}{2} - w\mathbf{x}_i^{t-1} \,, \tag{6}$$

where $0 \leq w \leq 1$ is a parameter called inertia weight in the PSO jargon, $\frac{\mathbf{L}+\mathbf{G}}{2}$ is the midpoint between the position of the particle's best local neighbor, denoted by $\mathbf{L}$, and the position of the best particle of the swarm, denoted by $\mathbf{G}$. Despite the fact that the PSO algorithm is clearly a second order system, we can write $w = (1-\alpha)$ for some $0 < \alpha < 1$, which makes Equation 6 resemble Equation 2. Thus, our model seems to capture certain commonalities present in several swarm intelligence systems with the advantage that our model is not focused on any specific application. Future work should be aimed at determining the extent to which the higher level of abstraction of Equation 2 is useful for the design and analysis of swarm intelligence systems.

As a model aimed at understanding the dynamics of collective decision and action, our model is not unique. In fact, our model is a member of a class of models known in the statistical physics literature as *opinion formation models* [2]. In the large body of literature dealing with these kinds of models, we can find two that are closely related to the model proposed in this paper. One of these models was proposed by Scheidler *et al.* [14]. Their model consists of a population of agents each of which has a memory of fixed size, which stores the values of the last k observed states. If all of these states are all equal to, say state 1, then the agent adopts state 1. The authors called this state update mechanism the *k-unanimity* rule. Our model and the k-unanimity model are similar in that the k-unanimity rule can be seen as a special case of our tendency mechanism. In our case, we sum the values of the observed states with exponentially decreasing weights tunable via the parameter α. The k-unanimity rule, on the other hand, can be interpreted as a sum of the values of the observed states with equal weights. In both cases, a threshold determines whether an agent changes state or not. The other model that is related to ours, is the so-called majority rule model [11]. In the majority rule model, teams of three agents are repeatedly formed, each time with different agents. When three agents form a team, they exchange their states and the state of the team's majority (that is, the state of at least two of these agents) is adopted by all the agents in the team. The majority rule may be seen as a distributed implementation of the k-unanimity rule with $k = 2$. Another common element between our model and the majority rule model is the concept of differential latency, which was first introduced in [11].

The model presented here may be seen as a more general model than the aforementioned models because in many cases it is possible to find a value for α that reduces Equation 2 to a simple average of the value of the last k observed states. Therefore, it is possible to reduce our model to the k-unanimity model. The majority model is also subsumed within our model because, as we discussed earlier, the majority rule model is a distributed implementation of the k-unanimity model. Thus, even though the dynamics at the individual level are different, one can find a configuration of the model presented here that mimics the dynamics of the majority rule at the collective level.

5 Conclusions

When an animal is part of a group, its behavior is influenced by the behavior of other animals that are also members of the same group. The fact that this phenomenon occurs across very diverse animal groups (including human groups), seems to indicate that there are intrinsic benefits to imitation, or at least action exploiting socially acquired information. In this paper, we introduce a simple social influence model that can be seen as a system of coupled exponential smoothing equations. The model is a binary decision model in which the tendency of an agent to adopt one of the two available behaviors/opinions is reinforced by the observation of another agent exhibiting a particular behavior or holding a certain opinion. We explored the dynamics of the system when agents can always be observed or when agents are observable only at certain times. Both of these circumstances model real-life situations. We observe that consensus is reached in both cases; however, the consensus state in one case is the opposite of the consensus state in the other case.

Future work includes a more through study of the system's dynamics both in simulation and analytically. An application area that benefits from our work is swarm intelligence. We discussed how our model captures the essential dynamics of two families of swarm intelligence algorithms for optimization, ACO and PSO, as well as collective decision-making processes in swarms of robots.

Acknowledgment

Part of the work presented in this paper was carried out while the first author was affiliated with IRIDIA, Université libre de Bruxelles. Eliseo Ferrante acknowledges support from the Research Foundation Flanders (FWO) of the Flemish Community of Belgium through the H2Swarm project.

References

1. Bonabeau, E., Dorigo, M., Theraulaz, G.: Swarm Intelligence: From Natural to Artificial Systems. Santa Fe Institute Studies on the Sciences of Complexity, Oxford University Press, New York (1999)
2. Castellano, C., Fortunato, S., Loreto, V.: Statistical physics of social dynamics. Reviews of Modern Physics 81(2), 591–646 (2009)
3. Dorigo, M., Stützle, T.: Ant Colony Optimization. Bradford Books, MIT Press, Cambridge, MA (2004)
4. Faria, J.J., Krause, S., Krause, J.: Collective behavior in road crossing pedestrians: the role of social information. Behavioral Ecology 21(6), 1236–1242 (2010)
5. Gardner Jr., E.S.: Exponential smoothing: The state of the art–Part II. International Journal of Forecasting 22(4), 637–666 (2006)
6. Giraldeau, L.A., Valone, T.J., Templeton, J.J.: Potential disadvantages of using socially acquired information. Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences 357(1427), 1559–1566 (2002)
7. Goss, S., Aron, S., Deneubourg, J.L., Pasteels, J.M.: Self-organized shortcuts in the argentine ant. Naturwissenschaften 76(12), 579–581 (1989)
8. Hyndman, R., Koehler, A., Ord, K., Snyder, R.: Forecasting with Exponential Smoothing. The State Space Approach. Springer, Berlin, Germany (2008)
9. Kennedy, J., Eberhart, R.: Particle swarm optimization. In: Proceedings of IEEE International Conference on Neural Networks, pp. 1942–1948. IEEE Press, Piscataway, NJ (1995)
10. Lambiotte, R., Saramäki, J., Blondel, V.D.: Dynamics of latent voters. Physical Review E 79(4), 046107, 6 pages (2009)
11. Montes de Oca, M.A., Ferrante, E., Scheidler, A., Pinciroli, C., Birattari, M., Dorigo, M.: Majority-rule opinion dynamics with differential latency: A mechanism for self-organized collective decision-making. Swarm Intelligence 5(3-4), 305–327 (2011)
12. Pillot, M.H., Gautrais, J., Gouello, J., Michelena, P., Sibbald, A., Bon, R.: Moving together: Incidental leaders and naïve followers. Behavioural Processes 83(3), 235–241 (2010)
13. Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M.W., Fogarty, L., Ghirlanda, S., Lillicrap, T., Laland, K.N.: Why copy others? Insights from the social learning strategies tournament. Science 328(5975), 208 – 213 (2010)
14. Scheidler, A., Brutschy, A., Ferrante, E., Dorigo, M.: The k-unanimity rule for self-organized decision making in swarms of robots. Tech. rep., IRIDIA, CoDE, Université Libre de Bruxelles, Brussels, Belgium (October 2011)