

# A Reductionist Approach to Hypothesis-Catching for the Analysis of Self-Organizing Decision-Making Systems

Heiko Hamann  
Department of Computer Science  
University of Paderborn  
Paderborn, Germany  
heiko.hamann@uni-paderborn.de

July 23, 2013

## Abstract

A difficulty in analyzing self-organizing decision-making systems is their high dimensionality which needs to be reduced to allow for deep insights. Following the hypothesis that such a dimensionality reduction can only be usefully determined in an act of a low-scale scientific discovery, a recipe for a data-driven, iterative process for determining, testing, and refining hypotheses about how the system operates is presented. This recipe relies on the definition of Markov chains and their analysis based on an urn model. Positive and negative feedback loops operating on global features of the system are detected by this analysis. The workflow of this analysis process is shown in two case studies investigating the BEECLUST algorithm and collective motion in locusts. The reported recipe has the potential to be generally applicable to self-organizing collective systems and is efficient due to an incremental approach.

## 1 Introduction

Self-organizing systems (SOS) are characterized by nonlinear dynamics, for example, due to positive and negative feedbacks and to the multitude of interactions between their components [1]. As a consequence of nonlinearity and of the large quantity of microscopic details, they are difficult to analyze. The challenging endeavor of trying to understand complex entities is common to science in its entirety to a certain extent. However, SOS may be different from other complex entities concerning the methods of reductionism because they are thought to be difficult to reduce, that is, cause and effect seem difficult to distinguish. There is the well-known controversy whether reductionism has the potential to explain everything [2] or maybe ‘more is different’ [3], that is, novel phenomena emerge on higher levels which cannot be reduced [4].

In the following we deliberately take the reductionist perspective on our subject of analyzing SOS. We report a method that focuses on the known properties of these systems and incrementally divides the configuration space into structured subsections. However, we also consider the fundamental difficulties one faces in the analysis of these systems. We assume that the effective analysis of a particular SOS is similar to a low-scale scientific discovery in terms of the applied methods. Hence, the analysis of SOS would include a certain degree of creativity similar to what is commonly thought of a scientific discovery. This is supported by a study that identified undecidable macroscopic properties in Ising models [5, 6]: “The development of macroscopic laws from first principles may involve more than just systematic logic, and could require conjectures suggested by experiments, simulations or insight.” Consequently, the ideal of an algorithm that automatically analyzes self-organizing processes seems unlikely to exist (cf. automatic abductive scientists [7, 8]). Instead, the hypothesis is that the analysis of self-organizing processes has to be a scientific endeavor and therefore cannot be efficiently automatized unless we find an algorithm

for scientific discovery itself. Hence, for our approach we need to transfer methods of scientific discovery to our domain. We do this following Hanson [9] and assume that scientific discovery is done by reasoning from data to hypothesis (cf. abductive reasoning and pragmatism [10]). We assume that the process of ‘hypothesis-catching’ [9] is driven by empirical data and an iterative process of catching a hypothesis, testing it, followed by refining or discarding it and starting over.

A difficulty in analyzing SOS, for example a swarm consisting of  $N$  moving individuals, is the high-dimensional configurations  $\gamma = (\mathbf{x}_0, \dots, \mathbf{x}_N, \mathbf{v}_0, \dots, \mathbf{v}_N, s_0, \dots, s_N)$ , for position  $\mathbf{x}_i$ , velocity  $\mathbf{v}_i$ , and state  $s_i$  of individual  $i$ . The task is to find a mapping  $g$  that projects the high-dimensional configuration space  $\Gamma$  ( $\gamma \in \Gamma$ ) to a low-dimensional space  $\Phi$  ( $g : \Gamma \rightarrow \Phi$  with  $\dim(\Phi) \ll \dim(\Gamma)$ ) while  $\Phi$  still represents the key features of the original system, in particular its evolution in time [11] and allows to assign causes and effects to regions of configuration space. The dimension reduction could, for example, be done with principal component analysis (PCA). The time series of configurations  $\gamma_t, \gamma_{t+1}, \dots$  could also be seen as a hidden Markov model inference problem [12] or one could use  $\epsilon$ -machines [13]. With these statistical methods one tries to find system states that have statistically equivalent futures, that is, one searches for the most predictive states. Although general features of the resulting Markov chains, such as the overall number of identified states, are helpful for the analysis of the system, a particular state of such a Markov chain will typically have not much explanatory power (e.g., the state does not represent a typical behavior or a comprehensible feature of the system). That would probably be asking for too much to have an automatic process that gives novel insights about a complex system.

Instead most reported methods to model SOS rely on non-automatic processes. For example, in swarm robotics one can start with a finite state machine that describes an individual robot’s behavior from which probabilistic finite state machines (Markov chains) are derived that describe the state of an individual robot and its environment or even fractions of the swarm [14]. In the latter case each state represents the average number of robots in a particular state at a given time.

In this paper, we report not an automatic, statistical analysis tool, but a recipe of how SOS can incrementally be modeled and how cause and effect can be identified and separated in order to achieve models with high explanatory power. Instead of finding the most predictive states, we structure the configuration space by defining subsets that are dominated by the same kind of feedback (positive or negative). The underlying hypothesis is that a valuable theory should assign different causes to positive feedback processes on the one hand and negative feedback processes on the other hand. These two kinds of processes should not be mixed. The reported method could optionally be transferred into a (semi-)automatic brute-force approach that tests a number of pre-defined hypotheses which would, however, limit the range of possible solutions.

In the following we introduce the two self-organizing decision-making systems that we take as examples to apply our method of analysis.

## 1.1 Case study 1: BEECLUST


The first case study is based on the BEECLUST algorithm which is a model algorithm for robot swarms. It is based on observations of young honeybees and was analyzed in many studies (e.g., see [15, 16, 17]). The algorithm implements the aggregation of a swarm at a maximum of a potential field without the application of a greedy gradient ascent. Controlled by this algorithm 3 agents stop (this threshold is varied in other studies) when they approach each other, they measure the local value of the potential field, and wait for some time proportionally to this measurement. Clusters start to form and eventually the majority of the swarm aggregates close to the global optimum of the potential field. See Fig. 1a for a definition of the BEECLUST algorithm.

The setting used in the following experiments is shown in Fig. 1b. The potential field is symmetric with peaks at the two short sides of the rectangular arena. Initially all  $N = 25$  robots are moving and scattered according to a random uniform distribution. There is noise in the agents’ sensors: only 90 percent of agent-agent approaches are detected by the agents.

The difficulty in analyzing BEECLUST is the complexity of microscopic processes in clusters of aggregated robots. Small clusters form and disappear fast. Big clusters persist much longer once they have formed because, for example, robots can get trapped within. Similarly to condensation

- 1) Each agent moves straight until it perceives an obstacle  $O$  within sensor range.
- 2) If  $O$  is a wall the agent turns away and continues with step 1.
- 3) If  $O$  is another agent and there is a third agent as well, the agent measures the local potential field value. The higher the potential field value the longer the agent stays still. After this waiting period, the agent turns away from the other agent and continues with step 1.

(a) algorithm (stop threshold of 3)



(b) setting, positions of stopped agents (circles) and moving agents (triangles) with trajectories of the last 20 time steps, contours show levels of the potential field

Figure 1: The BEECLUST algorithm and the used setting.

processes, there is a break-even in the average expected in- and outflow of robots depending on the cluster size. A particular difficulty of this setting is the symmetric potential that provokes a symmetry break due to initial fluctuations. Hence, the macroscopic feature of interest is whether the majority of agents gather at the left half or at the right half of the arena.

## 1.2 Case study 2: Collective motion in locusts

The desert locust, *Schistocerca gregaria*, shows collective motion, often called ‘marching bands’, in the growth stage of a wingless nymph [18]. The collective motion is expressed in the directional alignment of a majority of locusts, it is density-dependent, and individuals seem to change their direction as a response to neighbors [18]. In experiments the complexity of the collective motion is reduced to a pseudo-1-d setting by using a ring-shaped arena. Microscopic [19] and macroscopic models [20] of this behavior have been reported. Here we use a microscopic model similar to the self-propelled particles model of Cziráok et al. [19]. The system is defined in 1-d space, hence a particle  $i$  has coordinate  $x_i \in [0, 1)$  and discrete, dimensionless velocity  $u_i \in \{-1, 1\}$ . The dynamics is defined by

$$x_i(t+1) = x_i(t) + v_0 u_i(t), \quad (1)$$

$$u_i(t+1) = \begin{cases} G(L_i(t), R_i(t)), & \text{with probability } P_d \\ -u_i(t), & \text{with probability } P_n, \\ u_i(t), & \text{else} \end{cases} \quad (2)$$

where  $L_i$  and  $R_i$  are the numbers of neighboring particles  $j$  located in the interval  $[x_i - \Delta r, x_i + \Delta r]$  with  $u_j = -1$  (neighbor is moving to the left) or  $u_j = 1$  (neighbor is moving to the right) respectively, perception range of a particle  $\Delta r = 0.045$ , and nominal velocity  $v_0 = 0.001$ .  $P_d = 0.1$  is the particle’s probability of reconsidering its direction of movement.  $P_n = 0.015$  is the particle’s probability of inverting its direction of movement spontaneously, hence, it implements noise. With  $G$  we implement a local majority decision

$$G(L, R) = \begin{cases} -1, & L > R \\ +1, & R > L, \\ u_r, & R = L \end{cases} \quad (3)$$

where  $u_r \in \{-1, +1\}$  is a random tie breaker choosing  $-1$  or  $+1$  with equal probability. The initial condition is a random uniform distribution for both the particles’ coordinates  $x_i$  and their velocities  $u_i$ . The following experiments were done with a population of  $N = 45$  particles.

## 2 Hypothesis-catching recipe

The recipe is based on two main concepts, Markov chains and an urn model. The aim of this procedure is to discover the main causes that trigger the observed behavior and that have as much explanatory power as possible.

### 2.1 Markov chains

In this work the investigated collective systems are modeled as Markov chains. The challenge is to determine a set of Markov states  $S = \{s_0, s_1, \dots\}$  that usefully structure the above mentioned configuration space  $\Gamma$  into subsets  $\Gamma_{s_i} \subset \Gamma$ . The assumed Markov property is only an approximation because not necessarily each configuration  $\gamma \in \Gamma_{s_i}$  shows the exact same dynamics and there might be correlations with the past dynamics. The states in  $S$  are associated with agents that have certain local and global properties, that is, the states of the Markov chain combine microscopic and macroscopic properties.  $S$  itself is structured according to a feature  $f$  that describes the key feature of the observed phenomenon (e.g., ‘in favor of option 1’ in opinion dynamics) and

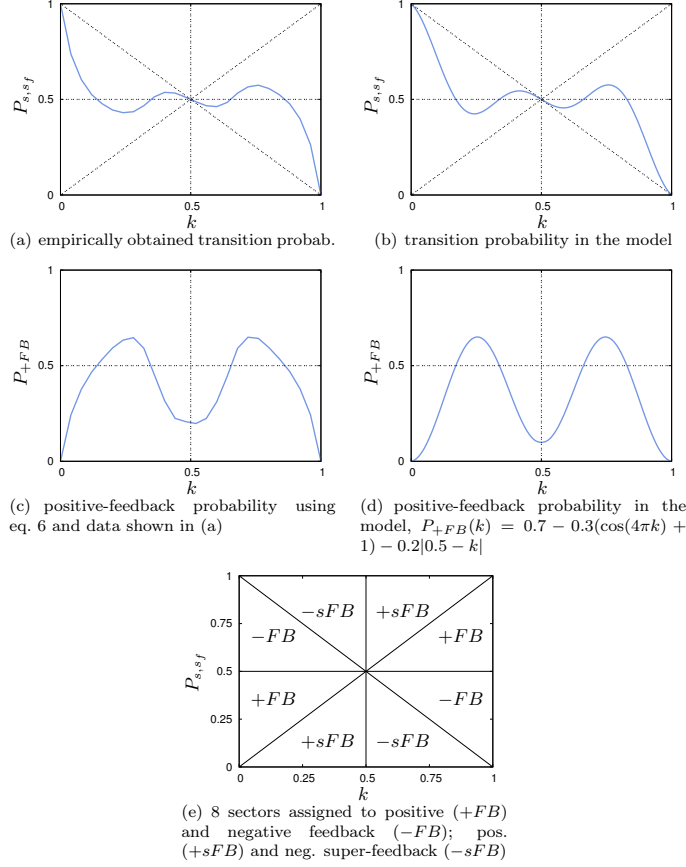


Figure 2: Transition probabilities  $P_{s,s_f}$  between feature-states and the respective positive-feedback probabilities  $P_{+FB}$ .

can be assigned to each agent  $a$  out of the set of all agents  $A$ . For simplicity we restrict this study to discrete, binary key features (on or off, left or right, etc.) but the proposed method can be generalized to more complex cases (e.g., by discretizing continuous features). The Boolean function  $f$  determines whether an agent  $a$  has the key feature. The subset of agents with key feature at time  $t$  is given by  $A_f(t) = \{a|f(a,t)\}$  (e.g., agents that move to the left) and the key feature ratio  $k$  is the fraction of agents with key feature:  $k = |A_f(t)|/|A|$ . The set of states  $S$  is separated into two symmetrical halves according to the key feature:  $S_f$  and  $S \setminus S_f$ . We call the two sets  $S_f$  and  $S \setminus S_f$  *feature-states* and transitions  $s_f \rightarrow s$  and  $s \rightarrow s_f$  between states  $s_f \in S_f$  and  $s \in S \setminus S_f$  *feature-transitions* because these are of relevance to the macroscopically defined observed phenomenon.

Each feature-transition between states  $s_f$  and  $s$  can take one of two types: those pointing towards key feature states ( $s \rightarrow s_f$ ) and those pointing away from them ( $s_f \rightarrow s$ ). The probability  $P(s \rightarrow s_f, k)$  that a certain feature-transition occurs for a given key feature ratio  $k$  is measured empirically using simulations of the actual collective system. By normalizing

$$P_{s,s_f}(k) = \frac{P(s \rightarrow s_f, k)}{P(s \rightarrow s_f, k) + P(s_f \rightarrow s, k)} \quad (4)$$

we get the probability that one out of these two feature-transitions occurs for a given  $k$  (notice the symmetry  $P_{s,s_f}(k) = 1 - P_{s_f,s}(k)$ ). A typical example of such transition probabilities is given in Fig. 2a.

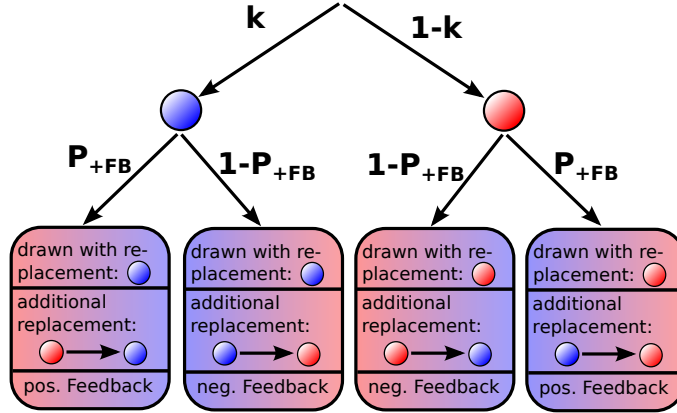


Figure 3: Swarm urn model; having the key feature is represented by blue and lacking the key feature by red. The diagram shows the 4 cases: drawing an agent with key feature followed by positive or negative feedback and drawing an agent without key feature followed by pos. or neg. feedback.

## 2.2 Urn models

The above defined transition probability  $P_{s,s_f}(k)$  can be modeled and understood by an urn model [21, 22] (see Fig. 3 for a schematic overview) similar to common urn models such as the Ehrenfest urn model [23] and the Pólya urn model [24]. An urn is filled with  $|A|$  marbles of two types:  $|A_f|$  marbles representing agents with key property and  $|A \setminus A_f|$  marbles representing agents without key property. An initial draw of a marble from the urn with replacement has hence linear dependence on the current key feature ratio  $k$  and resembles the frequency of agents with key feature  $|A_f|$ . In a second step, a second marble from the urn is replaced which either implements positive feedback by reinforcing the agent group represented by the first marble (say a blue marble was drawn first, then positive feedback is implemented by replacing a red marble from the urn with a blue one) or negative feedback by reinforcing the other agent group (say a blue marble was drawn first, then negative feedback is implemented by replacing a blue marble from the urn with a red one). The feedback is implemented by a probability of positive feedback  $P_{+FB}$  and hence the subsequent respective increase with probability  $P_{+FB}$  and decrease with probability  $1 - P_{+FB}$  of the two agent groups.

In former works, this ‘swarm urn model’ was applied to simple Markov chains (mostly 2-state Markov chains) but it can also be applied to more complex Markov chains by focusing on pairs of states at a time. For such a pair of states it is possible to determine a transition probability  $P_{s,s_f}$  for a given positive-feedback probability  $P_{+FB}$  by summing over the two respective cases

$$P_{s,s_f}(k) = kP_{+FB}(k) + (1 - k)(1 - P_{+FB}(1 - k)), \quad (5)$$

and also the inverse procedure is possible, that is, to determine positive-feedback probability  $P_{+FB}$  based on an empirically obtained probability of the corresponding transition  $P_{s,s_f}(k)$ . Rearranging eq. 5 and using the assumed symmetry  $P_{+FB}(k) = P_{+FB}(1 - k)$  yields

$$P_{+FB}(k) = \frac{P_{s,s_f}(k) - 1 + k}{2k - 1}, \quad (6)$$

see [21, 22] for details. Fig. 2c shows the positive-feedback probability obtained from the measurements shown in Fig. 2a using eq. 6. Fig. 2d shows a produced positive-feedback probability that generates the transition probability  $P_{s,s_f}$  shown in Fig. 2b using eq. 5 which is similar to that shown in Fig. 2a.

Two properties noticed in the diagrams shown in Fig. 2 are of importance here. First, a positive-feedback probability of  $P_{+FB}(k) > 0.5$  means that positive feedback is predominant for

the given  $k$ . If we have both predominant pos. feedback ( $\exists k_0 : P_{+FB}(k_0) > 0.5$ ) and predominant neg. feedback ( $\exists k_1 : P_{+FB}(k_1) < 0.5$ ) within one transition, then we call the transition ‘mixed’ (as in Fig. 2c and d). Second, we have accordingly several local maxima in  $P_{s,s_f}$  with  $P_{s,s_f} > 0.5$ , that is, sub-intervals  $k \in I_f \subset [0, 1]$  with a majority of transitions to  $s_f$  and sub-intervals  $I \subset [0, 1]$  with a majority of transitions to  $s$ . In analogy to multi-modal probability distributions, we interpret these local maxima as indicators for having mixed several causes of different origin within a pair of states  $s_f$  and  $s$ . Our reasoning is that this mixing of causes has to be resolved either by subdividing these two states into at least four states or by choosing two alternative states, that is, choosing a different structure of configuration space. We say that we are only satisfied by states that do not mix causes, that is, we have either only positive feedback ( $\forall k \in (0, 1) : P_{+FB}(k) \geq 0.5$ ,  $\forall k, 0 < k < 0.5 : P_{s,s_f} \leq 0.5$ ,  $\forall k, 0.5 < k < 1 : P_{s,s_f} \geq 0.5$ ) or only negative feedback. The ability to detect this quality of state pairs in Markov chains implements the main guidance in the reported recipe for the analysis of self-organizing decision-making systems.

In this study we discover a different kind of feedback that was not found in the former studies of these models and systems [21, 22]. Therefore, we distinguish between *standard feedback*, which cannot increase above  $P_{+FB}(k) = 1$  with consequent transition probability  $P_{s,s_f}(k) = k$  and cannot decrease below  $P_{+FB}(k) = 0$  with consequent  $P_{s,s_f}(k) = 1 - k$ , and *super-feedback* that is nonlinear in  $k$  (e.g.,  $P_{s,s_f}(k) < k$  for  $k < 0.5$ ). The nonlinearity in  $k$  means that drawing a marble with key property is more likely than a mere random draw with probability  $k$  in the urn model. When doing measurements based on trivial 2-state Markov chains, super-feedback seems not to be common in natural and artificial systems [22]. Hence, its occurrence seems to depend on the selective definition of states based on local and global features. In order to model super-feedback the before reported swarm urn model [21, 22] needs to be extended. This is done by substituting the initial draw of a single marble with a majority or minority decision of a randomly arranged group of a given size. That way the model gains similarity to voter models by Galam [25].

Positive super-feedback is modeled by a majority decision. For a given group size  $G$ ,  $G$  marbles are drawn in the first step (to avoid ties we assume  $G$  is odd). The majority of this group determines which agent type is reinforced by positive feedback in the second step. Say the majority of that group has key property, then in the subsequent step positive feedback would be the reinforcement of the key property. The group decision introduces the above mentioned nonlinearity in  $k$  because for  $k > 0.5$  the probability of a majority in a group of  $G > 1$  is bigger than  $k$ . The transition probability for positive super-feedback for a group size of  $G$  is a generalization of eq. 5 and is defined by

$$\begin{aligned}
P_{s,s_f}^{+sFB}(k) = & \\
& \left( \sum_{n \in \{\lceil G/2 \rceil, \dots, G\}} \binom{G}{n} k^n (1-k)^{G-n} \right) (1 - P_{-FB}(k)) \\
& + \left( \sum_{n \in \{\lceil G/2 \rceil, \dots, G\}} \binom{G}{n} (1-k)^n k^{G-n} \right) P_{-FB}(1-k), \tag{7}
\end{aligned}$$

which is obtained by the appropriate combinatorics of summing over all possible compositions of majorities and the two cases out of four in the swarm urn model that contribute positively (see Fig. 3). Hence, eqs. 5 and 6 are only special cases of eq. 7 for  $G = 1$ . The qualitative influence of group size  $G$ , how to determine it, and whether it is directly connected to actual microscopic processes is discussed below. Note that differently from eq. 5, we explicitly model the negative-feedback probability because for positive super-feedback it is the negative feedback probability  $P_{-FB}$  that takes shapes similar to the positive feedback probabilities  $P_{+FB}$  in standard feedback.

Negative super-feedback is just the opposite of the above  $P_{s,s_f}^{-sFB}(k) = 1 - P_{s,s_f}^{+sFB}(k)$ . It can be viewed as a minority decision to illustrate the antagonism. The minority of the group determines which group is reinforced by positive feedback in the second step. Say the minority of that group has key property, then in the subsequent step positive feedback would be reinforcement of the key

property. The transition probability for negative super-feedback and a group size of  $G$  is defined by

$$\begin{aligned}
P_{s, s_f}^{-sFB}(k) = & \\
& \left( \sum_{n \in \{1, \dots, \lfloor G/2 \rfloor\}} \binom{G}{n} k^{G-n} (1-k)^n \right) (1 - P_{+FB}(k)) \\
& + \left( \sum_{n \in \{1, \dots, \lfloor G/2 \rfloor\}} \binom{G}{n} (1-k)^{G-n} k^n \right) P_{+FB}(1-k). \tag{8}
\end{aligned}$$

According to this extended swarm urn model we define 8 feedback sectors in the transition–probability diagrams as shown in Fig. 2e. Keeping these sectors in mind, the transition probabilities can be classified directly depending on the sectors into which they extend (see Fig. 5 for examples): positive feedback (only  $+FB$  sectors involved), negative feedback (only  $-FB$  sectors), positive super-feedback (only  $+sFB$  sectors), negative super-feedback (only  $-sFB$  sectors), and mixed (sectors with different signs).

### 2.3 Recipe for analysis

The recipe is a simple 3-step instruction and implements an iterative approach to probabilistic abduction [7, 8]. It is driven by two dominant features: the hypothesis-catching abilities of the human mind and the guiding method of feedback analysis using the swarm urn model.

#### Step 1: Initial design of a Markov chain

This is done along two dimensions: the key feature dimension which is mirror symmetric (here: horizontally) and a mix of other global and local features (here: vertically).

- a) Start with the micro-level Markov chain. If micro-states already address the key property, then arrange these states horizontally, otherwise vertically.
- b) If the micro-states do not address the key property, then add feature-states which catch the macroscopic key property of the observed phenomenon along with the respective feature-transitions.
- c) Possibly add additional states that represent combinations of other global and local features according to the current hypothesis.

#### Step 2: Check feedbacks

Analyze pairwise the mutually symmetric feature-transition probabilities. Check whether they have non-mixed feedback and whether they are time-invariant or at least limited time-variant.

#### Step 3: Refinement or re-start

Refine your hypothesis if the results are encouraging or re-start the process with step 1. In the case of refinement, catch a hypothesis that defines additional/alternative states for each pair of states that belong to the feature-transitions with inhomogeneous feedback. Continue with step 2.

In the following case studies the key feature is *left* in both cases. In the BEECLUST system *left* represents the property of staying at the left half of the arena. In the locust system *left* represents the property of moving to the left which is associated with moving counter-clockwise in the ring-shaped arena. Hence, in the following illustrations of Markov chains the horizontal key feature dimension represents nicely the mirror symmetry and the left vs. right dichotomy.

## 3 Results

Before the results are reported some practical aspects of the implementation of the simulation and the analysis of the data are discussed. An option is to separate the simulator and the analysis



of the data in a process sequence which then necessitates to store the chronological sequence of all seen system configurations. This is, however, a large amount of data because SOS are high-dimensional and many samples are needed to get meaningful data (many gigabytes and samples in the order of  $10^5$  in the case of the presented case studies). However, once this data is available the complete analysis is quick because the simulator does not need to be re-run for each iteration of the recipe. A compromise is to do a 2-step dimensionality reduction: first, data from the simulator is reduced to reduce the amount of data to be stored; second, during the analysis the dimensionality is further reduced according to the defined Markov chain. The idea of the first dimensionality reduction is to optimize the tradeoff between reducing the amount of data considerably but at the same time still allowing for as many ways of later analysis as possible. In the case of the BEECLUST scenario this is easily done by storing a chronological sequence based on only four states for each agent: moving at the left/right half, stopped at the left/right half. The subsequent analysis according to the Markov chain then reassigns these sequences of simple states to the respective states of the Markov chain. In the case of the locust scenario it is more challenging to usefully reduce the data because as many graph-theoretic methods of analysis should still be applicable to the data as possible. That implies, however, that adjacency matrices or lists need to be saved. For the BEECLUST scenario the 2-step approach was implemented and for the locust scenario an all-in-one solution was preferred<sup>1</sup>.

In the following, diagrams of the empirically obtained transition probabilities and positive feedback probabilities are given. To allow the presentation of two case studies, the discussion is limited to a qualitative analysis that documents the feasibility of the reported approach. Hence, an in-depth investigation of statistical significance is out of this paper’s scope. The diagrams show multiple lines for different time steps to allow for a rough estimation of time-variance in the shown values.

### 3.1 Case study 1: BEECLUST

At the microscopic level the BEECLUST scenario is described by a 2-state Markov chain consisting of state  $M$  for moving and state  $S$  for being stopped, see Fig. 4a. In a first approach, we start with the simplest macroscopic Markov chain, that is, we duplicate each micro-state and get four states: feature-states *left* and *right* each with *moving* and *stopped* (see Fig. 4b). The resulting transition probabilities for  $M_r \rightarrow M_l$  are shown in Fig. 5a and indicate a mixed transition. In a second approach, we refine the above Markov chain by introducing states  $HS$  of agents that had stopped on the respective side but are now moving again and are still on the same side (see Fig. 4c). Such a state makes use of an agent’s history, increases the memory of the Markov chain, and introduces elements of a second-order Markov chain but does not change the approach qualitatively. The resulting transition probabilities for  $HS_r \rightarrow M_l$  are shown in Fig. 5c which indicates a mixed transition and the probabilities are also very time-variant. Also the transition  $M_r \rightarrow M_l$  is still mixed (Fig. 5d). In a third approach, we further refine the above by introducing states  $HSM$  of agents that had stopped on one side before but are now moving on the other side. The resulting transition probabilities for  $HS_r \rightarrow HS_r M_l$  are shown in Fig. 5e and indicate positive super-feedback combined with positive feedback. The probabilities are time-variant but show positive feedback for all measurements. The resulting transition probabilities for  $HS_l M_r \rightarrow HS_l$  are shown in Fig. 5f and indicate negative super-feedback which is almost time-invariant. The transition  $M_r \rightarrow M_l$  is not plotted again because it is similar to that of the second approach shown in Fig. 5d, hence, it is still mixed. However, the states  $M_l$  and  $M_r$  are drained over the course of time (no return from states  $S_l$  and  $S_r$ , see Fig. 4d) as shown in Fig. 6a which gives the relative frequency of states (due to symmetry only values for *left* are shown). For completeness, Fig. 5b shows the transition probabilities for a control experiment which used *up* and *down* as global feature (upper and lower half of arena) instead of *left* and *right*. It shows the data for the transition  $M_u \rightarrow M_d$ . As expected the transition probabilities are independent from the incorrect key feature.

<sup>1</sup>source code available online: <http://heikohamann.de/saso2013>

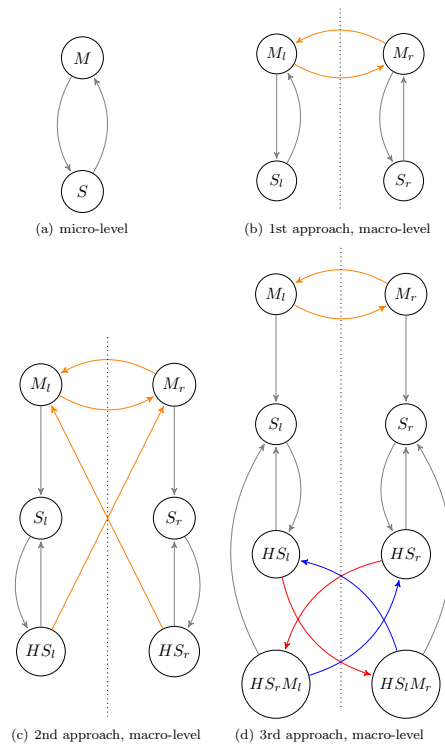


Figure 4: BEECLUST scenario, Markov chains for micro- and macro-level, feature-states are *left* and *right* separated by dotted line, micro-transitions (gray), feature-transitions: mixed feedback (orange), positive feedback (red), and negative feedback (blue).

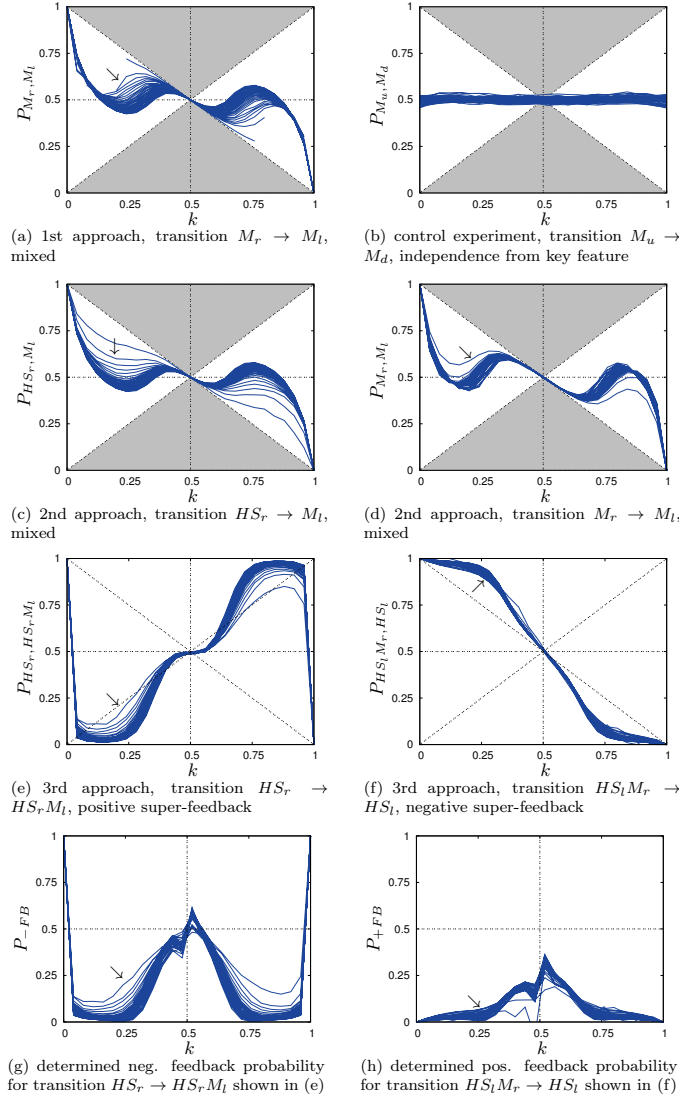


Figure 5: BEECLUST scenario, transition and feedback probabilities;  $k$  gives fraction of swarm with key feature (being in left half); arrows indicate direction towards later measurements.




Figure 6: left, BEECLUST: relative state frequencies over time; right, locusts: histogram of connected component sizes for transition  $RR \rightarrow LR$  (median of 33 and mean of 27.9 for last shown time step  $t = 2000$ ).

Using the measured transition probabilities the underlying feedback probabilities are calculated using eq. 6 in the case of standard feedback and with respective conversion using eqs. 7 and 8 in the case of super-feedback. For the probabilities of the transitions  $HS_r \rightarrow HS_r M_l$  and  $HS_l M_r \rightarrow HS_l$  this was done. Both of these transitions show super-feedback in which case we have group size  $G$  as a free parameter in eqs. 7 and 8. The appropriate group size is found in a simple procedure. In the case of super-feedback, one starts with the smallest group size of  $G = 3$  (only odd group sizes are used to avoid ties). If parts of the resulting feedback probability are negative or above one ( $\exists k : P_{sFB}(k) < 0$  or  $\exists k : P_{sFB}(k) > 1$ ), this is an indicator for a too small group size. Accordingly the next bigger group size  $G + 2$  is tested. The procedure is continued until the minimal group size with  $\forall k : 0 \leq P_{sFB}(k) \leq 1$  is found. Still, the appropriate group size is not unique because also bigger group sizes may give valid probabilities. Hence, the found group size is only a lower limit. For both of the investigated transitions the determined group size is  $G = 9$  and the feedback probabilities are given in Fig. 5g and h. Whether this result implies directly the neighborhood sizes of agents in states  $HS$  or  $HSM$  prior to their transition or possibly rather the differences between involved group sizes is left for future work but a similar case is investigated in more depth in the other case study. Also note that deviations for  $k \approx 0.5$  are expected because the equations for the feedback probability have singularities at  $k = 0.5$  for all group sizes (e.g., see denominator of eq. 6 for  $G = 1$ ).

The findings of this first application of the analysis recipe as represented by Fig. 5e-h are interesting because they exclusively focus on individual, moving robots. Without these results one was tempted to focus the analysis on the duel between competing clusters at the right and the left half. Instead the above analysis points to the positive feedback generated by individual robots that are lost to a cluster on one side when they switch sides as moving robots. The interpretation of the resulting Markov chain is non-trivial but a few direct conclusions can be drawn. The positive super-feedback in the transitions  $HS_r \rightarrow HS_r M_l$  means that each agent taking the transition from  $HS_r$  to  $HS_r M_l$  (i.e., leaving the half of the arena at which its former cluster is placed) helps to drive the system towards the formation of a majority (either  $k \ll 0.5$  or  $k \gg 0.5$ ) and especially outside the interval  $k \in [0.3, 0.7]$  this drive is intensive (indicated by low negative feedback in Fig. 5g). In turn, each agent taking the opposite transition from  $HS_l M_r$  to  $HS_l$  (i.e., returning to the half of the arena at which its former cluster is placed) slows down the progress of the system. With such analyses at hand, SOS might be optimized. For example, the system can be improved by trying to decrease the number of agents that take negative super-feedback transitions. The determined group size of  $G = 9$  is relatively high (total number of agents is 25). Thus agents taking a super-feedback transition contribute much more to the system's global dynamics than other agents taking standard feedbacks because the strong bias in the transition probability corresponds directly to big average changes in the key feature.

### 3.2 Case study 2: Collective motion in locusts

On the microscopic level the locust scenario is also described by a 2-state Markov chain ('left' and 'right' referring to counterclockwise and clockwise motion on the ring respectively, see Fig. 7a). In contrast to the BEECLUST scenario, however, the two states represent already the key feature. If this Markov chain is used, the transition probabilities for  $R \rightarrow L$  are mixed [21, 22]. The first hypothesis, that is tested, is based on the concept of the largest connected component (LCC). The distances between particles and their perception range  $\Delta r$  define neighborhoods which imply a graph (particles are vertices and neighborhood determines edges). Connected components are subsets of vertices that are connected by a single edge directly or by a path consisting of a number of edges. The rationale is that in those cases when the LCC represents a vast majority of the swarm, it possibly dominates the collective decision. Hence positive feedback could be based on whether a particle belongs to the LCC. The size  $C$  of the LCC is classified by three states of the Markov chain in dependence on the population size  $N$ :  $\frac{2N}{3} < C$  (C1),  $\frac{N}{3} < C < \frac{2N}{3}$  (C2), and  $C < \frac{N}{3}$  (C3). The total number of states doubles to 6 because each state exists in two types: left (L) and right (R). However, only one transition shows pure negative feedback and is almost time-invariant (transition  $RC1 \rightarrow LC1$ , see Fig. 8a). A second transition possibly also shows




Figure 7: Locust scenario, Markov chains, mixed feedback (orange), positive feedback (red), negative feedback (blue)

negative feedback (transition  $RC2 \rightarrow LC2$ , see Fig. 8b), and all remaining 7 feature-transitions are of mixed feedback (data not shown). We do not discuss statistical significance here because even if transition  $RC2 \rightarrow LC2$  is showing statistically significant positive feedback, there is a big interval  $k \in [0.2, 0.8]$  with neutral transition probability of 0.5 which limits this hypothesis' value. Also other separations of LCC sizes into states have been tested without improvements.

With the second hypothesis we drop the idea that the LCC is important. Instead we focus only on the connected component to which the respective particle belongs. We distinguish between components with a majority of particles in state *left* and components with a majority of particles in state *right*. As tie-breaker rule we count ties as components with a majority of state *right* to avoid dedicated 'tied states'. Hence, we get 4 states: particle state *left* and component majority *left* ( $LL$ ), particle state *right* and component majority *left* ( $RL$ ), similarly  $LR$  and  $RR$ . We get negative super-feedback for transition  $RL \rightarrow LL$  which is also almost time-invariant (see Fig. 8c), positive super-feedback for transition  $RR \rightarrow LR$  which is more time-variant (see Fig. 8d), and all other transitions are mixed. We give only data for one mixed transition that is of interest,  $RR \rightarrow LL$  (see Fig. 8e), because it is close to negative feedback except for the intervals  $k \in [0.2, 0.37]$  and  $k \in [0.63, 0.8]$ .

The third hypothesis is a refinement of the former. We add states  $L_{<3}$  and  $R_{<3}$  for components of size 1 and 2 without considering the majority within these small components (i.e.,  $L$  and  $R$  give the considered particle's state). This increases the total number of states to 6. The Markov chain with all feature-transitions is shown in Fig. 7b. By adding these two extra states, the above mentioned transition  $RR \rightarrow LL$  has improved and shows arguably negative feedback now but is still considerably time-variant (see Fig. 8f). Also note the negative super-feedback within the interval  $k \in [0.42, 0.58]$ . There are still several mixed transitions but all transitions that include  $L_{<3}$  or  $R_{<3}$  are infrequent and also transition  $LR \rightarrow RL$  is infrequent (data not shown). The transitions  $RL \rightarrow LL$  and  $RR \rightarrow LR$  do not change considerably in comparison to those shown in Fig. 8c and d (data not shown). For two transitions we determine the positive feedback probability based on eq. 7 and follow the above procedure to determine an appropriate group size  $G$ . The results are shown in Fig. 8g for  $RL \rightarrow LL$  and h for  $RR \rightarrow LL$ . Both diagrams show deviations again due to the singularity at  $k = 0.5$ . The determined group size in both cases is  $G = 31$ . In order to investigate whether this determined group size has only a theoretical meaning within the urn model or can be related to microscopic features, the distribution of connected component sizes is measured. The implementation of the locust model detects particles that take the transition  $RR \rightarrow LR$ . At that time, the size of the particle's connected component

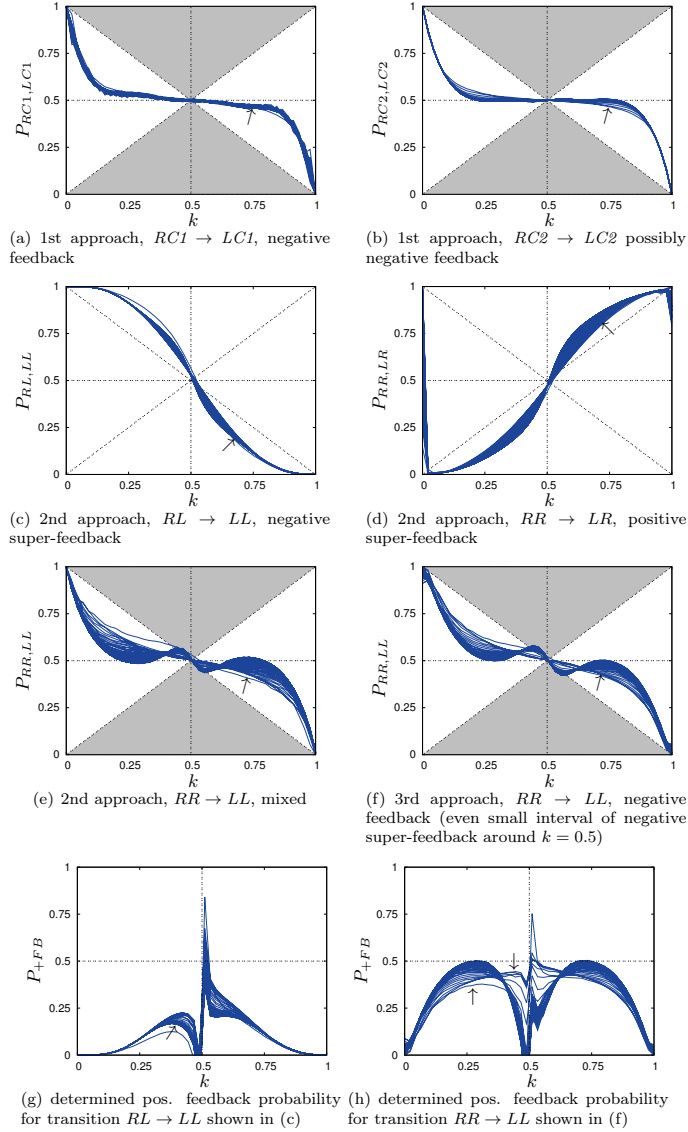


Figure 8: Locust scenario, transition and feedback probabilities;  $k$  gives fraction of swarm with key feature (going left); arrows indicate direction towards later measurements.

is determined. This data is shown as histogram in Fig. 6b. The different lines give the data for different times. Early in the simulation there is a peak at  $C = 8$  and when the system has converged there is a peak at  $C = 38$ . For the last shown time step ( $t = 2000$ ) the mean is 27.9 and the median is 33. Hence, we can speculate that the above determined group size of  $G = 31$  would correspond to the average situation. However, note that the connected component is not identical to the neighborhood of an individual and its local majority decision (eq. 3) because the component can extend over a bigger distance. A component's majority may only influence a particle over a sequence of time steps by bridging longer distances in a series of particle decisions. That way the majority decision of the urn model with group size  $G = 31$  can be interpreted as a macroscopic representation that averages microscopic behavior over space and time. Still, the investigation of group sizes is left for future work.

The findings of this second application of the analysis recipe are of interest because the analysis of the locust scenario is a difficult scientific problem itself. The found importance of majorities in connected components of sizes bigger than two reveals that the alignment process in locusts is not a competition for the dominance of the largest connected component only or for the mere overall ratio of one of the two directions but a distributed, concurrent competition in and between all bigger connected components. The importance of three-body processes in this locust model was reported before [26]. The situations when the majority in a component is counteracted seem to be of particular importance (positive super-feedback in transition  $RR \rightarrow LR$ ). However, also those moments when the majority in a component switches are relevant, especially for situations when the swarm is undecided  $k \approx 0.5$ . The negative super-feedback in transition  $RR \rightarrow LL$  for  $k \in [0.42, 0.58]$  stabilizes the system at the undecided state. However, leaving that state would actually be important for the system's effectivity. Hence, the results for transition  $RR \rightarrow LL$  raise interesting new questions for future research. For example, it should be tested whether these switches in the majority of a component rely on spontaneous switches of individuals or maybe on merging with other components with opposite majorities.

## 4 Conclusion

In this paper the former swarm urn model [21, 22] is extended to include the concept of super-feedback which is a feedback that is nonlinear in the current ratio of the considered key feature. Majority decisions in randomly arranged groups of a certain size are introduced to the swarm model to include this extended concept of feedback. Using this urn model and Markov chains, a recipe is reported that gives a guideline on how to define hypotheses and how the Markov chains should be designed following the detection of feedbacks. The recipe implements Hanson's philosophy that scientific discovery is driven by empirical data [9] and can be interpreted as an incremental abduction process [10]. The iterative analysis process defined by the recipe gives the researcher, who is analyzing self-organizing decision-making systems, a direct response to formulated hypotheses. That way it gives new impulses and triggers the creation of new approaches of how to structure the configuration space appropriately. The structuring is implemented by defining appropriate states for the Markov chain. The states represent subsets of configuration space which are characterized by mixtures of local and global properties. Purely microscopic causes would not have enough explanatory power while focusing on purely macroscopic causes would lack the links to the micro-level.

The analysis in two case studies has generated non-trivial results. In the case of the BEECLUST algorithm the results point to the importance of individual agents (e.g., agents that had stopped but then changed sides) in contrast to basing the analysis on the growth and decay of particular clusters. In the case of the locust scenario the results point to the importance of changing majorities in all connected components of size 3 or bigger and to the importance of majorities in components that are counteracted or reinforced. These findings are the result of an abductive reasoning process and hence have the status of unverified hypotheses. The required next step is their verification to determine whether the found Markov chains truly point to the most explanatory system configurations.

In future work this approach will be applied to additional scenarios and a brute-force approach will be implemented that picks the best hypothesis of how to separate configuration space out of a pre-defined list of hypotheses. Also a reversion of this method into a constructive guidance for the design of collective decision-making systems will be investigated.

## Acknowledgment

The author thanks Marco Dorigo, Payam Zahadat, and Gabriele Valentini for comments that helped improving the manuscript.

## References

- [1] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems*. Oxford Univ. Press, 1999.
- [2] S. Weinberg, “Reductionism redux,” *The New York Review of Books*, Oct. 5 1995.
- [3] P. W. Anderson, “More is different,” *Science*, vol. 177, no. 4047, pp. 393–396, Aug. 4 1972.
- [4] J. Crutchfield, “The calculi of emergence: Computation, dynamics, and induction,” *Physica D*, vol. 75, no. 1-3, pp. 11–54, Aug. 1994.
- [5] M. Gu, C. Weedbrook, Á. Perales, and M. A. Nielsen, “More really is different,” *Physica D*, vol. 238, pp. 835–839, 2009.
- [6] P.-M. Binder, “Computation: The edge of reductionism,” *Nature*, vol. 459, pp. 332–334, May 2009.
- [7] L. Magnani, *Abduction, reason, and science: Processes of discovery and explanation*. New York: Springer, 2001.
- [8] P. Thagard, *Computational philosophy of science*. MIT Press, 1988.
- [9] N. R. Hanson, *Patterns of Discovery*. Cambridge University Press, 1958.
- [10] C. S. Peirce, *Collected papers of Charles Sanders Peirce*. Belknap Press, 1974, vol. 5.
- [11] H. Hamann and H. Wörn, “A space- and time-continuous model of self-organizing robot swarms for design support,” in *First IEEE Intern. Conf. on Self-Adaptive and Self-Organizing Systems (SASO’07)*. IEEE Press, Jul. 2007, pp. 23–31.
- [12] O. Cappé, E. Moulines, and T. Rydén, *Inference in hidden Markov models*. Springer, 2005.
- [13] J. P. Crutchfield and C. R. Shalizi, “Thermodynamic depth of causal states: Objective complexity via minimal representations,” *Phys. Rev. E*, vol. 59, pp. 275–283, Jan 1999.
- [14] A. Martinoli, K. Easton, and W. Agassounon, “Modeling swarm robotic systems: A case study in collaborative distributed manipulation,” *Int. Journal of Robotics Research*, vol. 23, no. 4, pp. 415–436, 2004.
- [15] H. Hamann, T. Schmickl, H. Wörn, and K. Crailsheim, “Analysis of emergent symmetry breaking in collective decision making,” *Neural Computing & Applications*, vol. 21, no. 2, pp. 207–218, Mar. 2012.
- [16] H. Hamann, B. Meyer, T. Schmickl, and K. Crailsheim, “A model of symmetry breaking in collective decision-making,” in *From Animals to Animats 11*, ser. LNCS, S. Doncieux, et. al., Eds., vol. 6226. Springer, 2010, pp. 639–648.



- [17] T. Schmickl and H. Hamann, “BEECLUST: A swarm algorithm derived from honeybees,” in *Bio-inspired Computing and Communication Networks*, Y. Xiao, Ed. CRC Press, Mar. 2011.
- [18] J. Buhl, D. J. T. Sumpter, I. D. Couzin, J. J. Hale, E. Despland, E. R. Miller, and S. J. Simpson, “From disorder to order in marching locusts,” *Science*, vol. 312, no. 5778, pp. 1402–1406, 2006.
- [19] A. Czirók, A.-L. Barabási, and T. Vicsek, “Collective motion of self-propelled particles: Kinetic phase transition in one dimension,” *Phys. Rev. Lett.*, vol. 82, no. 1, pp. 209–212, 1999.
- [20] C. A. Yates, R. Erban, C. Escudero, I. D. Couzin, J. Buhl, I. G. Kevrekidis, P. K. Maini, and D. J. T. Sumpter, “Inherent noise can facilitate coherence in collective swarm motion,” *PNAS*, vol. 106, no. 14, pp. 5464–5469, 2009.
- [21] H. Hamann, “Towards swarm calculus: Universal properties of swarm performance and collective decisions,” in *Swarm Intelligence: 8th Intern. Conf., ANTS 2012*, ser. LNCS, M. Dorigo, et al., Eds., vol. 7461. Springer, 2012, pp. 168–179.
- [22] ———, “Towards swarm calculus: Urn models of collective decisions and universal properties of swarm performance,” *Swarm Intelligence*, 2013, in press. [Online]. Available: <http://dx.doi.org/10.1007/s11721-013-0080-0>
- [23] P. Ehrenfest and T. Ehrenfest, “Über zwei bekannte Einwände gegen das Boltzmannsche H-Theorem,” *Physikalische Zeitschrift*, vol. 8, pp. 311–314, 1907.
- [24] H. Mahmoud, *Pólya urn models*. Chapman & Hall, 2008.
- [25] S. Galam and S. Moscovici, “Towards a theory of collective phenomena: Consensus and attitude changes in groups,” *Europ. J. of Social Psychology*, vol. 21, no. 1, pp. 49–74, 1991.
- [26] C. Huepe, G. Zschaler, A.-L. Do, and T. Gross, “Adaptive-network models of swarm dynamics,” *New Journal of Physics*, vol. 13, p. 073022, 2011.