

# Agents, Simulation and Applications

---

by

Adeline Uhrmacher and Danny Weyns

# I

- 1 Stigmergic Cues and their Uses in Coordination: An Evolutionary Approach** *Luca Tummolini, Marco Mirolli, and Cristiano Castelfranchi* 1-1
- Introduction • Stigmergy: Widening the Notion but not too Much • Two Uses of Stigmergy in Coordination • Stigmergy in Cooperation and Competition • Why Pheromonal Communication is not Stigmergic • Understanding Stigmergy through Evolution • Future Work • Conclusion • Acknowledgments

# 1

## Stigmergic Cues and their Uses in Coordination: An Evolutionary Approach

---

1.1	Introduction.....	1-1
	Overview	
1.2	Stigmergy: Widening the Notion but not too Much .....	1-3
	A Tale of a Wrong Story • Is it Communication? • Stigmergic Cues as Practical Behavioral Traces	
1.3	Two Uses of Stigmergy in Coordination .....	1-6
	Coordination and Cues of Interference • Communication and Signals • Behavioral Communication and Implicit Signals • Stigmergic Self-Adjustment and Stigmergic Communication	
1.4	Stigmergy in Cooperation and Competition.....	1-9
	Cooperative and Competitive Interference • Collaborative and Conflictual Stigmergic Coordination	
1.5	Why Pheromonal Communication is not Stigmergic .....	1-10
	Pheromones are not Stigmergic Cues • The Role of Stigmergy in the Evolution of Pheromonal Communication	
1.6	Understanding Stigmergy through Evolution .....	1-11
	The Basic Model • Evolution of Practical Behavior • Evolution of Stigmergic Self-Adjustment and Indirect Coordination • Evolution of Stigmergic Communication	
1.7	Future Work .....	1-18
1.8	Conclusion .....	1-19
1.9	Acknowledgments .....	1-20
	References .....	1-20

Luca Tummolini  
*Institute of Cognitive Sciences and Technologies - CNR*

Marco Mirolli  
*Institute of Cognitive Sciences and Technologies - CNR*

Cristiano Castelfranchi  
*Institute of Cognitive Sciences and Technologies - CNR*

### 1.1 Introduction

---

Since its origins, the field of Multi-Agent Systems (MAS) has been entrenched with the notion of *direct communication* [Hew77, Smi80, BG88, Dur88].

In particular, the shift from blackboard systems [EM88], as proto-MAS, to proper MAS was characterized by abandoning a form of indirect coordination in favor of a more direct strategy. The ‘agents’ in a blackboard system used to leave their partial solutions to a

problem in a data structure shared with other ‘agents’ and each reacted to the product of the work of other agents by producing new work (i.e. new partial solutions). At the time, a single thread of control in writing in the blackboard created a bottleneck in problem solving. This constraint was overcome, at the beginning, by adopting direct communication between multiple blackboard systems (via message passing) and then, eventually, by dropping the indirect approach altogether. Though a new seminal approach was born, the baby, or at least a useful toy of his, was thrown out with the bath water.

From there on, in fact, the crucial coordination mechanism for Multi-Agent systems has been modeled on the most important communication tool available to humans: verbal language. With the help of a shared conventional language, humans in fact coordinate their activities thanks to verbal exchanges in the forms of speech acts [Sea69]. Inspired by their crucial role in human societies, speech acts have been studied and approximated also for the coordination of artificial agents in distributed systems [CL90, Sad92]. Direct communication soon began to be considered as necessary to achieve an effective MAS: in order to build autonomous agents capable of coordinating with each other, such agents must be provided with an Agent Communication Language to exchange messages (e.g. FIPA-ACL \*) on the background of shared ontologies.

However, direct communication by means of an agent language is just one possible mechanism to design an effective MAS that, together with several benefits, bears also many computational costs. To overcome these limitations many practitioners are nowadays turning their attention to approaches based on self-organizing techniques. In fact, taking inspiration from the biological sciences, agents have been moulded on social insects and their capacity to interact “indirectly” through environmental modifications (a phenomenon known as *stigmergy*, see [Gra59]). The importance of this mechanism has been investigated both in biological [Wen91, KT95, CDF<sup>+</sup>01] and artificial organisms [Par97, BHG<sup>+</sup>98, PBS05, MZ07]. Though self-organization and stigmergy are orthogonal phenomena [TB99], the narrow and unclear definition of what stigmergy is and its too close association with self-organizing approaches has also impeded the detachment of this coordination mechanism from the original domains of investigation (i.e. collaborative collective behaviors), and so has limited the impact of this approach for a wider range of problems that can be relevant for the MAS community.

Somehow ironically, indirect coordination is quite reminiscent of the blackboard system approach that was stigmergic in its very essence. By bringing the best of these two worlds together, the rehabilitation of indirect forms of coordination might be a critical step towards the next generation of agent technologies.

### 1.1.1 Overview

The general aim of this chapter is to clarify the notion of stigmergy and its uses for coordination by means of conceptual and operational models obtained with a set of evolutionary agent-based simulations.

In a series of recent papers we have argued that stigmergy is best understood as a peculiar form of *communication* [Cas06a, TCR<sup>+</sup>05, TC07]. In particular, stigmergic behavior has been analyzed as a kind of *indirect communicative behavior*, i.e. asynchronous communication enabled by a shared environment used as a repository for signals. Here we partially

---

\*See the Communicative Act Library Specification of the Foundation for Intelligent Physical Agents (FIPA): <http://www.fipa.org>

correct our previous view by distinguishing the notion of *stigmergic cue* from its uses in coordination: *stigmergic self-adjustment* and *stigmergic communication*.

In the first part of the chapter, we review current definitions adopted by leading scholars in different fields. Though the core properties of stigmergic behavior have been acknowledged, a clear-cut conception is still missing. Occasionally the phenomenon is classified as a kind of coordination, other times as communication. To better clarify what stigmergy is as well as when it is communication and when it is not, we introduce the general theory of behavioral implicit communication and we clarify what is peculiar of the stigmergic case.

In the second part of the chapter, we explore the evolution of stigmergic behavior adopting a simulative approach. In particular, we have simulated a population of artificial agents living in a virtual environment containing safe and poisonous items (fruits): eating safe fruits increases the fitness of an individual, while eating poisonous ones decreases it. The behavior of the agents is governed by artificial neural networks whose free parameters (i.e. the weights of the networks' connections) are encoded in the genome of the agents and evolve through a genetic algorithm. Agents interact with their environment and between each other through the traces that their behaviors leave in such an environment.

Biological plausibility aside, the simulations are designed to provide an operational model of stigmergic cues together with a principled way to understand their possible uses. By making explicit the transition from (1) a multi-agent system in which agents individually look for their resources to one in which (2) each agent indirectly coordinates with what the other agents do (stigmergic self-adjustment) and, finally, (3) to a situation in which each agent acts also to send a message about what kind of resources are available in a risky environment (stigmergic communication), the simulations offer a precise analysis of the difference between traces that are just signs with a behavioral content and traces that are signals with a behavioral message.

## 1.2 Stigmergy: Widening the Notion but not too Much

---

### 1.2.1 A Tale of a Wrong Story

The term *stigmergy* has been coined by the French entomologist Pierre-Paul Grassé to explain the behavior of termites during nest-costruction [Gra59]. The term comes from the combination of two greek words: “stigma”, meaning a puncture or sting (by extension, whatever provokes a reaction), and “ergon”, meaning work. Stigmergy, then, refers to the role of work in stimulating a reaction of other agents or, with Grassé: “the stimulation of the workers by the very performances they have achieved” [Gra59]. Stigmergy is considered by Grassé as a *peculiar kind of stimulation*: the consequence of previously accomplished work. It is such consequence that drives and orients additional construction: “the significant stimuli imperatively direct workers' responses” [Gra67].

The concept was introduced as the key factor to explain how complex structures such as termites' nests [Gra59] (or a complex problem-solving like rebuilding a tunnel-gallery after a sudden damage [Gra67]) can be regulated without, clearly, any complete plan being known by the termites. The complexity of a nest is astonishing and requires the collaborative coordination of a huge number of workers. Exploiting the stimulating factor of accomplished work was considered as the key to success.

Though many were able to replicate the original findings of Grassé, soon the importance of stigmergy was also somehow challenged (see [DJ88] for a review of these critiques). Several researchers remained skeptical about the fact the stigmergy could be sufficient for

the coordination of complex building tasks, especially when the construction behavior was not strictly sequential (i.e. composed of an ordered succession of phases, one causally connected with the other in a sequential fashion). Other cues were also suggested to be necessary in order to shift between different activities or to stop the whole endeavor once completed.

Curiously enough, in the very same year in which Grassé published his considerations on termites' construction behavior, another great entomologist, Edward O. Wilson, had discovered that ants, by leaving pheromones in their environment, create odor trails that attract other ants [Wil59].

It is, however, due to the work of Bruinsma [Bru79] that the role of pheromones also in coordinating termites' nest construction was finally unravelled. Bruinsma discovered that both pheromone trails and the impregnation of soil pellets with pheromone are critical factors in attracting other termites and orienting their activities. The existence of an initial deposit of soil pellets *impregnated with pheromone* stimulates workers to accumulate more material. This reaction originates a positive feedback mechanism because the accumulation of material reinforces the attractivity of deposits through the pheromones emitted by those materials.

The emergence of pillars, walls and royal chambers in termite nests can be accounted for once the *self-organizing properties* of this kind of interactions are understood. In particular it has been suggested that termites' behavior follows just a single pattern (i.e. picking up and depositing a soil pellet, if pheromone is present) [BTD<sup>+</sup>97] and this building behavior is modulated by several environmental conditions (e.g. when an air stream drives molecules of pheromones in a given direction the pattern is influenced), not last the fact that the environment is changed by work of other termites. It is self-organization that explains the emergence of a structure at the global level from interactions among its low level components, without any explicit coding of this process in the individual agents [GGT07].

The basic intuition of Grassé was in the end right: no sophisticated cognitive processing is performed by the termites in order to construct their nest. However it is the self-organizing dynamics enabled by the properties of pheromone (i.e. its additivity and its decadence rate) that regulate and modulate termites' behaviors.

Unfortunately, it seems that stigmergy, defined by Grassé himself as a particular kind of stimulation, does not play a distinct role in this process: *it is not the nest under construction that functions as significant stimulus for other termites*. Environmental modifications due to the construction activity do play a crucial role in modulating termites' behavioral patterns but not *as* stimuli for the termites.

If this fact is too embarrassing, one can, as Theraulaz and Bonabeau has done [TB99], *change* the original conception of Grassé to include also the influence of pheromone when left in the environment as a kind of stigmergy.

Indeed, as it has been explicitly recognized by Holland and Melhuish [HM99], the modern practice is to extend the definition of stigmergy by replacing the sense of "work" crucial for Grassé (stigmergy as the product of work functioning as a significant stimulus) to that of "any environmental change produced by the animal".

Though for us this apparently innocent extension is to be resisted (see Sections 1.2.3 and 1.5.2 below), it has been a fortunate one. So much that, today, the fact that there exist two main varieties of stigmergy is commonly accepted: the first being characterized by leaving special markers in the environment (i.e. pheromones), the second being the exploitation of the results of the work so far achieved by the agents [TB99, DBT00, Par06].

### 1.2.2 Is it Communication?

Though Grassé never mentioned it, stigmergy is also usually considered as a form of communication.

Wilson has been the first to explicitly link stigmergy to a communicative process and coined the name *sematectonic communication* (from the Greek “sema” for sign or token and “tekton” for craftsman, builder) [Wil75]. Reviewing different communication strategies in social animals, Wilson naturally conceived stigmergy as a particular kind of communication exploiting the “structures built by animals”.

More recently, many researchers have been explicit in considering stigmergy as a peculiar form of indirect communication, and more precisely: “communication by altering the state of the environment in a way that will affect the behaviours of others for whom the environment is a stimulus” [KES01] (see also [DBT00, Mat95, TC07]).

Though clearly, there are cases that can be defined as *stigmergic communication*, stigmergy by itself is not necessarily so (see Section 1.3.4 and the simulation results described in 1.6.4). The main reason why stigmergy is so commonly seen as a part of a communicative process lies in the fact that stigmergy has been always approached in the limited context of collaborative activities (i.e. nest construction, collective sorting, collaborative foraging, division of labor, clustering etc.), and collaborative tasks such these immediately create the conditions for communicative behaviors in which an agent aims (functionally or intentionally) to influence another one.

Though stigmergic communication is useful as a coordination mechanism in cooperative behavior, there is room for a different use which we name *stigmergic self-adjustment*. In Section 1.3.4 by providing a principled distinction between self-adjustment and communication as two ways to coordinate one’s behavior, we will fill that room.

### 1.2.3 Stigmergic Cues as Practical Behavioral Traces

Apart from his classification of stigmergy as a kind of communication, Wilson also noticed that the role of stigmergy in eliciting additional work was clearly a special case and offered the example of male crabs that build structures to attract females for reproductive purposes. He felt that Grassé’s characterization was too restricted, and wanted stigmergy “to denote the evocation of any form of behavior or physiological change by the evidences of work performed by other animals, including the special case of the guidance of additional work” [Wil75].

Though this is certainly sensible, the use of work in a strict sense is dispensable too. Why do we want to restrict the phenomenon to the construction of something? Why excluding a simple footprint?

On the other hand, Holland and Melhuish [HM99] suggestion (“any environmental change produced by the animal”) is too broad, because it would over-generalize the original intuition of the French entomologist. Stigmergy would simply refer to any kind of indirect influence, and any road sign would qualify as a case of stigmergy.

How can we reasonably generalize the notion of stigmergy?

Here is our proposal. The basic intuition of Grassé was that the traces of “work” left in the environment might become significant stimuli by themselves for an agent. A *trace* is the effect of any behavior persisting in the environment. But what behaviors are the relevant ones?

At a first approximation, we can say that relevant behaviors are those whose goal or function is a practical one: i.e. any behavior that is not adapted just to influence another agent counts as a relevant practical one. Hence, from this perspective, *stigmergy occurs whenever*

another agent's behavioral trace (a persistent effect of a practical behavior whatsoever) is used as guide for one's own future behavior.

In other words, there is a case of stigmergy whenever a behavioral trace of an agent is also a *cue* for another organism. Cue is a notion used by biologists precisely to refer to any (animate or inanimate) feature in the environment, that an agent is adapted (thanks to evolution or learning) to use as a guide for future actions [Has94]. Though many stimuli are in principle detectable by the sensory organs of an animal, only a subset of them are actually cues (i.e. are guides for action).

Very close to the spirit of Grassé's original intuition, in our approach, *stigmergy refers to a subset of cues and, in particular, to the practical behavioral 'traces' we leave around whenever we act in a physical world.*

The heuristic value of this precise definition, of course, is to be judged against its pragmatic consequences, that is, by the way it helps to identify and understand phenomena that would be, otherwise, obscured. The next two Sections are thus devoted to explore such consequences.

### 1.3 Two Uses of Stigmergy in Coordination

---

The most dramatic fact of living in the same world is that the actions of an agent can interfere with the success of those undertaken by others [Cas98]. Sociality is first of all due to this very basic condition. In fact, beyond one's own individual qualities, there is a social component behind the evolutionary and practical success of an organism. And this social component is mainly due to *interference*, and the opportunities and obstacles it creates. How do agents deal with interference? Is communication necessary to solve interference problems?

In order to set the stage for understanding two distinct uses of stigmergic cues, in this Section we clarify the relationship between coordination and communication.

#### 1.3.1 Coordination and Cues of Interference

At the simplest level *coordination* is a form of individual social action [Cas97]. Coordination is coping with positive and negative interference by means of adaptation (thanks to evolution, learning, reasoning or design). In particular, to coordinate in a social context, an individual agent takes into account the interference created by another agent in one's own action, or better, *coordination is adaptation to another agent's behavior in order to increase one's success\**.

While in negative interference the behavior is adapted to an obstacle, in positive interference the behavior is adapted to an opportunity. In cases of reciprocal coordination, the agents' behaviors are adapted to each others.

More precisely, the agents need to detect some cue in the environment which predicts such interfering conditions. As stated above, a cue is whatever one is adapted to use as a guide for future actions [Has94]. On this basis coordination can be more clearly defined as *adaptation by reading cues of interference.*

There are, however, two fundamental ways to adapt by means of these cues of interference: (1) the agent changes his own behavior to *avoid* the obstacle or to *exploit* the opportunity

---

\*In this chapter, we are interested only in "social" coordination between the behaviors of different agents. Non-social kinds of coordination such as the coordination of distinct behaviors of a single agent are therefore excluded.



(self-adjustment) or (2) the agent changes his behavior in order to influence the behavior of the interfering agent either by *impeding* or *inducing* such behavior (hetero-adjustment or influence). See Table 1.1.

**TABLE 1.1** Two strategies for solving interference problems

	Coordination by self-adjustment	Coordination by hetero-adjustment or influence
Negative Interference	Avoidance	Impediment
Positive Interference	Exploitation	Inducement

Finally, though coordination is based on such cues of interference, this does not mean that coordination can only be *direct*: via the detection of those cues that are the interfering behaviors themselves. *Indirect coordination* is possible *whenever acting on the basis of some cue other than the behavior of some agent, one is also adapting one's own behavior to those of others*. In traffic, for example, each of us individually act on the visual cues provided by the traffic lights, while at the same time, we also coordinate indirectly with all the others whose behavior we do not directly perceive.

### 1.3.2 Communication and Signals

We have seen that an agent can solve an interference problem in a more active way than by mere self-adjustment. The agent, in fact, can actively eliminate an obstacle or create an opportunity. In other words, one can *influence* the others to one's own benefit. Coercion is an obvious way to do so. Another is communication.

The most general approach to communication can be found in biology, where it is usually approached as the transfer of a *signal* from an agent to another one. When communication is defined in this way, of course, all the explanatory burden is on a clear notion of signal, which two leading scholars, Maynard-Smith and Harper, define as: “any act or structure which alters the behaviour of other organisms, which evolved because of that effect, and which is effective because the receiver's response has also evolved” [MSH03].

This way of understanding communication as signaling has the immediate advantage of making explicit that, at its core, communication is a form of influence, hence a peculiar kind of coordination. Communication, however, is distinct from coercion because the “the receiver's response has also evolved”. Being the product of evolution, learning, reasoning or design, receiver's response is in her own interests, while in coercive coordination this is not the case. If a stag pushes another one backwards, this behavior is not a signal. If the stag roars and the other animal retreats it is a signal because the retreating is evolved (or learned, decided or designed).

Moreover, the fact that the action or structure is also evolved *because* of its influencing effect makes clear that the action or structure is not only a cue. While a cue is defined only relatively to the agent using it, a signal is defined relative to its influencing effect on the behavior of others which benefit, first of all, the agent who produces or manifests it.

Given this definition of communication, it is also simple to distinguish its *direct* and *indirect* varieties. In direct communication, the signaller's behavior or structure itself is the signal which is interpreted by another agent. Differently, indirect communication exploits signals that persist in the environment and that are effective in their influencing function notwithstanding the absence of the signaller's itself.

### 1.3.3 Behavioral Communication and Implicit Signals

The above definition of communication, however, also suggests that signals are essentially products of cooperation in which both parties benefit from their exchange. Though a reliable communicative system is one that is evolutionary stable, an unstable but deceitful system, in which only the signaller benefits from the exchange, is still a communicative system after-all.

What is sufficient for a signal, for us, is just that it is defined relatively to the signaller's benefit and that it is at least a cue for some other agent who has evolved to respond to it as a cue but not necessarily as a signal: *a signal is any behavior or structure that alters the behavior of other organisms, which evolved because of that effect, and which is effective because that behavior or structure is a cue for the receiver*. The evolutionary advantages of the signaller and the receiver might diverge.

This is the kind of signal that we call *implicit*, and that characterizes a peculiar kind of communication: *behavioral implicit communication* [Cas00, Cas06b].

Usual practical behaviors in fact are meaningful and can become cues for other agents to be used as guides for future actions. Interference is the reason why agents living in a shared world will become attuned to such cues, and will become behavior readers to improve their coordination abilities. However, once agents are behavior readers, those, whose behavior is read, can start exploiting this fact to their own advantage. Though their behaviors maintain their practical aim (i.e. in traffic, a U-turn behavior), by being read and understood by other agents observing it, the practical behavior can be done also for influencing other agents by the very production of this behavioral cue (i.e. an agent can wait to see if other drivers are noticing the U-turn behavior before doing it because he intends to use this behavior also as a signal with the content: "I'm doing a U-turn" and so provoking their slowing down). As this example clarifies, beyond being a sub-case of coordination, a specific use of communication can also be a coordination mechanism. Since coordination is in general any adaptation by means of cues of interference, communication can also be an instrument for coordination by self-adjustment. In fact, when communication is focused on the production of cues of interference (e.g. by signalling an action one is going to perform), then communication is clearly one of the most powerful coordination mechanisms that we, and our agents, can take advantage of.

### 1.3.4 Stigmergic Self-Adjustment and Stigmergic Communication

It is now possible to specify two distinct uses of stigmergy as analyzed in Section 1.2.3: stigmergic self-adjustment and stigmergic communication.

*Stigmergic self-adjustment* amounts to the process of coordination merely by means of stigmergic cues. That is, *adaptation by the detection of practical behavioral traces*, i.e. the effects of practical behaviors registered in the environment. Because the cue is the effect of a behavior and not the behavior itself, stigmergic self-adjustment is a kind of indirect coordination.

Differently, *stigmergic communication* is *influence by an evolved (or learned, reasoned, designed) production of these stigmergic cues*. Given the notion of behavioral implicit communication, it is also clear that stigmergic communication is *implicit* in the sense that the stigmergic signals are effective because the effects of practical actions are cues for other agents (i.e. traces left and accessible in a shared environment that function as implicit signals). In stigmergic communication, the messages that are transmitted are *behavioral* ones, i.e. the implicit signals primarily refer to the agents behaviors or what a behavior can be a diagnostic or prognostic sign of (for an analysis of the kinds of meanings that can

be transmitted with stigmergic communication see [TC07]). Finally, because the signal is an effect of a behavior and not the behavior itself, stigmergic communication is a kind of indirect communication.

As for the more general case, stigmergic communication is also a powerful coordination mechanism for at least two reasons: firstly, because it is indirect, and, secondly, because it can exploit practical behavioral traces as cue of interference. While the former property is almost universally acknowledged, the latter has been too quickly forgotten. One of the main aims of this chapter is precisely to rectify such mistake.

## 1.4 Stigmergy in Cooperation and Competition

---

Stigmergy has been discovered in the particular context of collaborative nest construction. However, the focus on the role of stigmergy in collaboration has turned to be a bias originated from this initial domain of investigation. Limiting the attention to collaborative scenarios have impeded to notice, for example, that there are natural uses of stigmergic cues also in competitive and conflictual situations.

### 1.4.1 Cooperative and Competitive Interference

In particular, when the goals or the fitness of the agents are such that each benefits from others' success there is room for *cooperation*. More precisely, when interference is positive this entails that the behavior of an agent facilitates the success of the others such that acting for a common goal (i.e. cooperate) can be mutually advantageous. On the contrary, very often it is the case that an increase in fitness or the fulfillment of one's goals compromises the success of others (negative interference). In this case, the agents acting in ways that are reciprocally destructive are, even if unaware, in *competition*.

Of these very basic kinds of interference, there also more specific instances. Cooperation, for example, can be obtained thanks to collaboration, which means that not only the agents functionally or intentionally favor each other, but also that they adopt the goal of each other. *Goal adoption* is the process of internalizing the goal of another agent as one's own and acting *for* another agent [CC95]. Hence, *collaboration* is a specific way to achieve cooperation: it is cooperation via mutual goal adoption.

Differently, competition can originate conflict which is, again, a specific way in which agents might compete. In *conflict* in fact, agents are motivated by the fact that the others should not fulfill their goals. As in the case of collaboration, also in conflict the agents take into consideration the goals of the others but in order to frustrate them.

In-between pure cases of cooperation and competition, there are of course many situations that are partly cooperative and partly competitive but a simpler picture is enough for the aims of this chapter.

### 1.4.2 Collaborative and Conflictual Stigmergic Coordination

Notwithstanding its relevance for cooperation and collaboration, it is evident that there is plenty of competitive and conflictual situations in which stigmergy is used.

Collaborative activities, where the agents are ready to help each other to reach a common goal, are contexts in which stigmergy is obviously natural. Take for example the cooperative task in which two agents collaborate in answering customers' orders in an Italian bar. It is common, during breakfast for example, that many people ask at the same time for very different things: a coffee, a cappuccino and a sandwich for instance. How can the different

orders be correctly processed? Who has ordered what? Which is the guy who ordered the cappuccino? A very diffuse strategy to solve this problem is to employ stigmergic collaborative communication. While, for instance, a waiter executes the orders and starts preparing the coffee, the cappuccino and picking the sandwich, the other adaptively changes the environment. Given that the shapes of the dishes where the three items are delivered to the customers are distinctively different, one waiter immediately reacts to the orders by associating each customer with her appropriate dish. In this way, when the items are ready, the other waiter can easily associate each customer with the correct order. While no verbal communication is usually employed, the collaborative task is smoothly solved relying only on indirect communication through stigmergic cues.

Differently, consider the role of fences. Clearly such constructions are there to function as physical obstacles for strangers. Fences are built to avoid that strangers enter in a protected area, hence they are conflictual structures. Though they provide physical interference if somebody intends to trespass them, more than often they are not unavoidable obstacles. How can a fence be effective? The answer turns to be simple: by being also a stigmergic signal. In fact, as a practical behavioral trace left there by somebody, every fence has also a behavioral content related to the conflictual action of erecting it. The combination of its practical function and its being a stigmergic cue, for those that understand it, is often enough to achieve its intended purpose.

## 1.5 Why Pheromonal Communication is not Stigmergic

---

At the end of Section 1.2.1, we have suggested that the common practice of acknowledging the exploitation of special markers left in the environment is a misleading extension of the notion of stigmergy. Pheromones are probably the most diffuse example of stigmergy between the practitioners in fields like Artificial Life, Swarm Intelligence, Collective Robotics and, more recently, Multi-Agent Systems. Being so widely accepted as the prototypical case, our insistence on their being not stigmergic will arise much skepticism.

In this Section, we intend to clarify this issue and to explore the relevance of stigmergy in the evolution of pheromonal communication.

### 1.5.1 Pheromones are not Stigmergic Cues

The use of pheromones is not a case of stigmergy because these chemical molecules are not practical behavioral cues. A pheromone in fact is not only a cue. It is a chemical molecule whose sole function is to excite another agent (again from the Greek, “pherein”, to carry or transfer and “hormon”, to excite or stimulate), that has evolved for this influencing effect, and which is effective because the response has also evolved [KL59, Wya03]. Hence a pheromone is, by definition, an *explicit signal*.

What pheromones and stigmergic cues have in common is that they are traces, i.e. products of behavior that persist in the environment and this property can have useful consequences, not least by enabling self-organization or self-assembly. But widening the notion to include all kinds of indirect communication is, on the other hand, losing sight of what is peculiar of the stigmergic case.

Consider a familiar example. When driving in traffic, it is one thing to stop because there is a red light in front of you, it is another one to stop because the car in front of you has stopped. While the red light is an explicit signal left in the environment by somebody, the car in front of you is just a cue of a practical behavioral kind, a stigmergic kind.

Beside restoring Grassé’s first conceptualization (see Section 1.2.3), our more restrictive

notion avoids collapsing stigmergic communication into any kind of indirect communication.

### 1.5.2 The Role of Stigmergy in the Evolution of Pheromonal Communication

Understanding stigmergy from this perspective is also a contribution for a better explanation of explicit communicative systems themselves.

How in fact can an explicit communicative system evolve if not from an implicit one? As it has been suggested in the biological study of animal communication, a signalling system can evolve from a process of “ritualization” of a behavior that did not evolve for communication at the beginning [Tin52] but was already used as a cue by other individuals to gain information. And the same is true for pheromonal communication.

As stated in Section 1.5.1, pheromones are not stigmergic cues but chemical ones emitted only for communicative purposes. Pheromones are explicit signals which, most probably, *derived from stigmergic cues*. In [SS99], it is, for instance, offered a model of the evolution of pheromonal communication between goldfishes for reproductive purposes. In this model it is speculated that, at the beginning, hormonal products were released by females to water just as a side effect of ovulation. Given the relevance of these hormonal traces as predictors of the biological state of a female, the male soon became attuned to these traces as reproductive cues. Finally, by exploiting this fact, evolution has driven towards a specialization in the production of these hormones just for influencing purposes, and turning the hormones into real pheromones. Hormonal cues are in our approach stigmergic because they are traces of a process other than communication (i.e. ovulation). When specialized in pheromones, they become explicit signals.

Though the evolutionary transition from implicit to explicit signals is a matter for further research, keeping stigmergic communication clearly apart from other kinds of indirect communication is of extreme importance if we want to understand explicit communication system themselves (i.e. their evolution).

## 1.6 Understanding Stigmergy through Evolution

---

In order to clarify even further the notions so far discussed, in this Section we describe a set of agent-based simulations intended as operational models of stigmergic cues and their uses in coordination. In particular, we simulate a population of artificial agents living in a virtual environment containing safe and poisonous items (‘fruits’): eating safe fruits increases the energy (fitness) of an individual, while eating poisonous ones decreases it. If nothing in the physical properties of these items can be used for discriminating the good resources from the bad ones, how can the agents be successful in such a risky environment?

Knowledge of this properties is the prototypical kind of knowledge that is acquired by direct experience with the external environment. However when the risks associated with experience are high, an evolutionary pressure exists to learn indirectly exploiting the experience made by others. Omnivores, for example [VA00], are characterized by food neophobia, i.e. a form of hesitancy to eat novel food. Neophobia is considered as an efficient behavioral strategy to avoid the risk of ingesting poisonous substances. It is also known, however, that social context affects the acceptance of novel food by moderating such instinctive hesitancy.

In this Section, we explore a possible strategy to solve this problem by showing that, given specific circumstances, the individual capacity to categorize some natural kinds as safe or poisonous might rely on stigmergic cues left by other agents that, moreover, act in order to facilitate such individual categorization.

In our simulations, the behavior of the agents is governed by artificial neural networks whose free parameters (i.e. the weights of the networks' connections) are encoded in the genome of the agents and evolve through a genetic algorithm. Agents interact with their environment and between each other through the traces that their behaviors leave in such an environment. By way of such an agent-based model, we explore the evolution of three kinds of behaviors:

1. the basic practical behavior;
2. the use of practical behavioral traces as cues of positive interference (stigmergic self-adjustment);
3. the use of such stigmergic cues as implicit signals (stigmergic communication).

Though the simulations are offered only as operational definitions of our domain of investigation (i.e. stigmergy) and not as a plausible biological explanation of this form of social facilitation, they may also suggest that a communicative process might be discovered, whereas today biologists seem to overlook this possibility.

### 1.6.1 The Basic Model

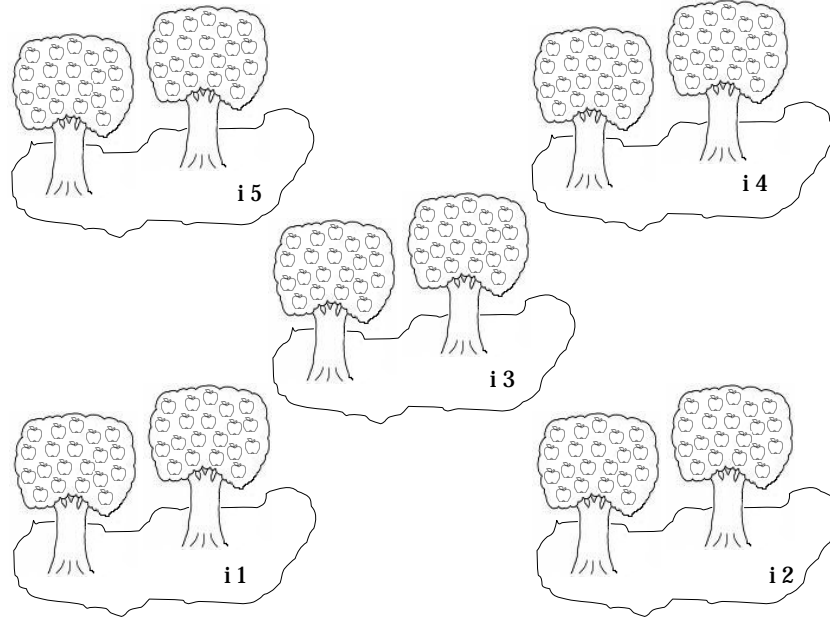
We simulate a population of 100 agents living in an environment composed of 5 'islands' (Figure 1.1). Each island contains 2 'trees' with 20 'fruits' each. One of the trees produces edible fruits, while the fruits produced by the other tree are poisonous. Eating an edible fruit increases an agent's fitness (i.e. its probability of reproduction) of 1 unit, while eating a poisonous fruit decreases an agent's fitness of 2 units. Poisonous and edible fruits (and trees) are perceptually identical, so agents cannot rely on direct perceptual cues to decide from which tree to feed. Each agent visits all the 5 islands where it can stay for at most 4 time steps. In each time step an agent can decide whether to eat a fruit from one of the two trees or to leave the island.

Agent's behavior is governed by an artificial neural network with 4 groups of input units sending connections to a single group of 5 hidden units which in turn connects to a group of three action units (Figure 1.2). The three possible actions are: (a) eat a fruit from the first tree; (b) eat a fruit from the second tree; (c) leave the island.

Actions are decided according to a winner-takes-all mechanism: in each time-step the agent performs the action corresponding to the action unit with the maximal activation. When the agent decides to eat a fruit from a given tree, a random fruit disappears from that tree and the fitness of the agent is updated according to the quality of the eaten fruit (+1 for edible fruits, -2 for poisonous ones). The first two groups of input units are composed by 20 units each, with each unit representing the presence/absence of a particular fruit in one of the two trees of an island.

Activation is binary: if the fruit is present, the corresponding unit is activated (1), otherwise it is not (0). The third group of input units is composed by 2 units representing the feedback from the agent's 'body' regarding the quality of the last eaten fruit: during the first time step in an island, no fruit has been eaten hence both units are silent (0); during the successive time steps, one of the two units is highly activated (10) depending on whether the last eaten fruit was edible or poisonous. The last group of 3 input units constitutes a memory of the last performed action: each unit corresponds to one action and it is highly activated (10) if that action has been the last to be performed and silent otherwise (0). (During the first time step in an island, no action has been performed hence all the three units are silent.)

We simulate a population of 100 individuals evolving through a genetic algorithm. Each



**FIGURE 1.1** The environment. The five islands (i 1-5) with two trees each (one with edible and one with poisonous fruits. At the beginning of a trial each tree has 20 fruits). See the main text for details.

individual lives for 5 trials, where a trial consists in visiting all the five island (eaten fruits re-grow between one trial and the other). After all individuals have lived their lives, their fitness is calculated according to the following formula:

$$f(x) = \frac{\sum_{t=1}^T \sum_{i=1}^I en_e - pn_p}{TIC} \quad (1.1)$$

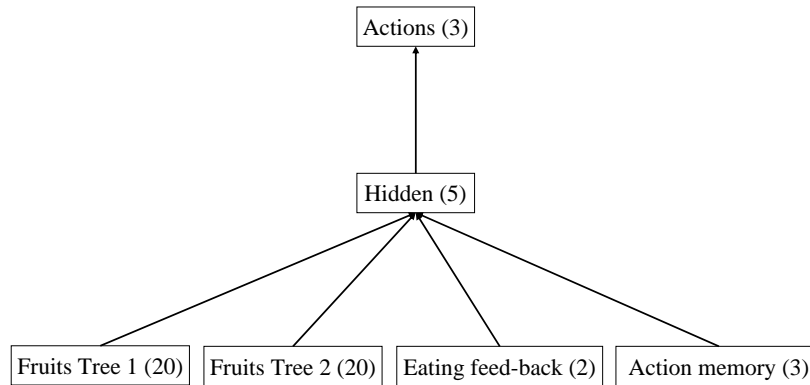
where  $n_e$  and  $n_p$  are, respectively, the number of edible and poisonous fruits eaten by  $x$ ,  $e$  and  $p$  are two constants (set to 1 and 2, respectively), and  $T$ ,  $I$ , and  $C$  are the number of trials (5), islands (5), and maximum cycles spent in an island (4), respectively.

The genome of individuals contains all the connection weights of the neural network, with each free parameter being encoded as an 8-bits string, whose value is then uniformly projected in the range  $[-5.0, +5.0]$ . The 20 best individuals of each generations are selected for reproduction and generate 5 offsprings each, which inherit their parent's genome with a 0.04 probability for each bit of being replaced with a new randomly selected value. The evolutionary process lasts 250 generations.

### 1.6.2 Evolution of Practical Behavior

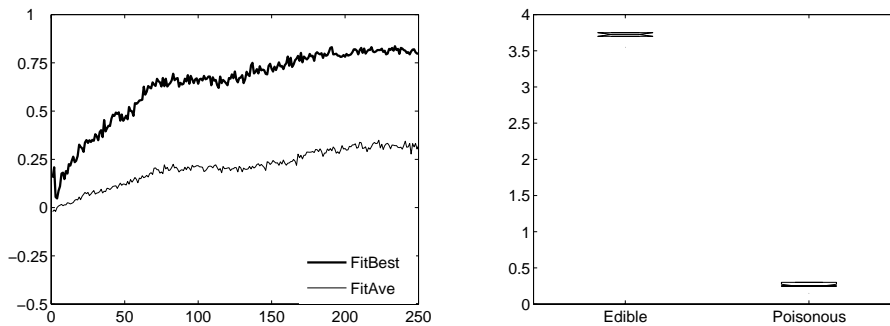
In this kind of simulation, which we will call 'base-line', individuals never interact with each other, so they cannot evolve any form of social behavior. Since at their arrival on an island the two trees are perceptually identical and contain the same number of fruits, agents cannot but choose randomly from which tree to feed.

On the other hand, once they have eaten a fruit from a tree, agents can decide whether to keep on eating from the same tree or to change tree depending on the bodily feedback about the quality of the eaten fruit: if the fruit was poisonous agents have to change tree, while if the fruit was edible agents have to keep on eating from the same tree.



**FIGURE 1.2** The neural network. Each block corresponds to a group of neurons. Numbers in parentheses correspond to the numbers of units of the group. Arrows represent all-to-all connections between groups.

Since there is no learning during individuals' life, this kind of behavioral strategy must evolve phylogenetically. This is in fact what happens in our simulations. Figure 1.3 shows the average results of 10 replications (with different random initial conditions) of this baseline simulation. Figure 1.3 (left) shows average fitness and the fitness of the best individual during the 250 generations, while Figure 1.3 (right) shows average number of edible and poisonous fruits eaten by the best individual of the last generation during each staying on an island.



**FIGURE 1.3** Results of the base-line simulation. Left: Average and best fitness along the 250 generations. Right: Average number of edible and poisonous fruits eaten by the best individual of the last generation during each staying on an island. Average results of 10 replications of the simulation.

Agents possessing the optimal behavioral strategy discussed above will eat on average 3.5 edible fruits and 0.5 poisonous ones during each visit on an island (0.5 poisonous fruits on average when they have to randomly guess, and  $0.5 + 3$  edible, since during the second, third and fourth choices the information about the previous action together with the bodily feedback can be used for eating the fruits from the right tree).

Hence, the average expected maximal fitness would be  $(3.5 + 0.5 \cdot (-2))/4 = 0.625$ .

The described optimal behavioral strategy did manage to evolve, indeed: at the end of



the 250 generations the best evolved individual has eaten on average about 3.7 edible fruits and about 0.3 poisonous ones (Figure 1.3, right), thus reaching a fitness of about 0.8 (Figure 1.3, left).

The reason why the best individual reaches a fitness which is higher than the maximal expected one is just chance: in each generation the best individual will happen to be the one which, beyond possessing an optimal behavioral strategy, is also the luckiest, having discovered the edible tree at the first choice more than the expected 50% of the times.

### 1.6.3 Evolution of Stigmergic Self-Adjustment and Indirect Coordination

Living in a solipsistic world, agents of the base-line simulation can evolve only individual practical abilities. In order to test whether our agents might evolve even social abilities, we run a second set of simulations in which individuals are allowed to live in a social world.

This ‘social’ simulation runs exactly like the base-line one but for the following modification. Each individual of the population shares its environment with 4 of its own clones (i.e. individuals possessing the same genome and hence the same neural controller). We make agents interact (through their environmental modifications) only with their own clones because we are interested in the emergence of (stigmergic) communication, but we are not interested in the problems of altruism which are typically posed by the evolution of communicative behaviors [MP05, MP08]. Allowing interactions only between clones assures us that no problem of altruism can arise, since in this way any possible altruistic behavior will favor only individuals possessing the same altruistic genes which code for the behavior, thus guaranteeing that the behavior can pass through generations (this is the strongest possible form of kin selection, [Ham64]).

At the beginning of each trial, each clone is put in a different island where it stays for a maximum of 4 cycles (less if it decides to leave the island earlier). After that, all the clones change their islands so that during each round each island is visited by a different clone and, after 5 rounds, each clone has visited each island. Like in the base-line simulation, the same process is applied for 5 trials.


What makes this simulation interesting for our purposes is that, while in-between different trials all the fruits of all the trees re-grow, within each trial they do not. The consequence is that *in this simulation the environment in which the agents live registers the traces of the behavior of other agents*.

In particular, in all but the very first round within a trial, the trees of an island will lack the fruits eaten by the agents which have already visited that island. Since, as shown by the results of the previous simulation, individuals do not eat randomly, but rather they tend to eat much more edible fruits than poisonous ones, *practical behavioral traces are meaningful*: the tree lacking more fruits is the edible one.

The results of this simulation (Figure 1.4) show that agents in the social condition are in fact able to exploit the information provided by the traces left in the environment by other individuals. In fact, in this condition the best evolved individual reaches a fitness of about 0.9 (Figure 1.4, left), which is significantly higher than the fitness reached by the best individual of the base-line, individual condition.

This fitness increase is explained by the fact that the best individual of the social condition eats an average of about 3.9 edible fruits and only about 0.1 poisonous fruits for each staying on an island (Figure 1.4, right). This is in turn due to the fact that, in the social condition, agents have to choose randomly which tree to feed to only during the first round of each trials. During the other four rounds they can eat from the tree containing less fruits thus avoiding more poisonous fruits and eating more edible fruits than the agents of the individual

condition.



**FIGURE 1.4** Results of the 'social' simulation. Left: Average and best fitness along the 250 generations. Right: Average number of edible and poisonous fruits eaten by the best individual during each staying on an island. Average results of 10 replications of the simulation.

This ability of exploiting the traces left in the environment by the actions of other agents is a paradigmatic example of what we call *stigmergic self-adjustment*. In fact, the traces left from previous eating behaviors become cues, stigmergic cues to which agents' behaviors are now sensitive. By acting on these cues, the agents indirectly coordinate their behaviors towards the edible items, and avoiding the poisonous ones.

#### 1.6.4 Evolution of Stigmergic Communication

In the social simulation just described the traces in the environment left by the practical behaviors of eating fruits become stigmergic cues because they carry valuable information which is exploited by the agents who perceive the traces so to increase their own individual fitness.

On the other hand, they are still not (implicit) *signals* in that these cues are produced *exclusively* for the practical purposes of isolated individuals, and not *also* for social ones (i.e. in order to influence the behavior of others), as required by our definition of implicit signals.

The practical behavior which produces the cues, i.e. the behavior of eating from the tree with edible fruits whenever you can, is not affected *in any way* by the fact that there are other agents which might use the traces of that behavior as cues.

In fact, agents of the social condition differ from the agents of the individual condition just for the way they react to environmental cues. When there are no environmental cues, i.e. in the first cycle of the first round of a trial, the behavior of the 'social' agents is exactly the same as the behavior of the isolated agents: i.e. choose a random tree. And agents of the two conditions behave exactly in the same way also when the information about the quality of the fruits is directly available to an agent from its own body. For instance, after a fruit has been eaten, the behavioral rule of both kinds of agents is 'change tree if the fruit was poisonous while keep on feeding from the same tree if the fruit was edible'.

To distinguish proper stigmergic *communication* from mere stigmergic *self-adjustment*, it is necessary that not only the perceiver of the trace modifies its behavior in order to exploit the information provided by the trace, but also *the trace producing behavior must, in some way, be modified* (for influencing). This does not necessarily mean that the behavior itself

will change in its morphology. It is also possible that the behavior remains the same, while it is the *conditions under which exactly the same behavior is produced* which may vary.

In order to show how this is possible we run a third set of simulations in which we slightly modify the set-up of the social simulation in the following way.


In this new simulation, the increase of fitness provided by eating an edible fruit is no more constant, but is proportional to the *hungriness* ( $h$ ) of the eating individual. Hungriness is 1 each time the individual arrives in an island and decreases of 0.5 each time the individual eats an edible fruit. Furthermore, each eating action has a cost ( $c$ ) of 0.1 on individual fitness. Hence, the fitness function used in this condition, which we will call the ‘hungriness’ simulation, is the following:

$$f(x) = \frac{\sum_{t=1}^T \sum_{i=1}^I h e n_e - p m_p - c n_a}{TIC} \quad (1.2)$$

where  $n_a$  is the number of eating actions performed by the individual.

The results of this simulation are shown in Figure 1.5. The results in terms of fitness (Figure 1.5, left) are not surprising: both average and best fitness significantly decrease with respect to the other two conditions because in this case edible fruits tend to provide less energy (depending on hungriness) and because there is always a cost to be paid when eating a fruit.

What is interesting here are the results regarding the average number of fruits eaten by the best individual during each staying on an island (Figure 1.5, right). Given the conditions of this simulation, an individual should never eat more than 2 edible fruits for each island: in fact, after it has eaten 2 edible fruits an individual hungriness has decreased to 0, and hence eating a third fruit would not increase fitness at all. On the contrary, it would in fact *decrease* fitness since the individual would pay the cost associated to the eating action. But the best individual of the last generation of this ‘hungriness’ simulation eats an average of 2.5 edible fruits for each island.



**FIGURE 1.5** Results of the ‘hungriness’ simulation. Left: Average and best fitness along the 250 generations. Right: Average number of edible and poisonous fruits eaten by the best individual during each staying on an island. Average results of 10 replications of the simulation.

This pattern of eating behavior clearly cannot be explained by referring only to individual advantages. The only function that eating more than 2 edible fruits can have is a social function: that is, *leaving more informative traces to forthcoming agents*. If during the first round on an island an individual would eat only two edible fruits before leaving the island, then the next agent would have to discriminate between trees differing only of one or two fruits (depending on whether the previous agent had chosen the edible tree first or had eaten

also a poisonous fruit). This would be of course theoretically possible, but it is evidently too difficult for the simple neural networks possessed by our agents. In fact, consider that, beyond this discrimination ability, the neural networks have also to encode all the other behavioral rules described above, and that fruits are eaten randomly from a tree. Hence the discrimination ability must be general to all the positions of fruits in the trees.

The number of poisonous fruits eaten by the best individual of the hungriness simulation is significantly less than that of the best individual of the individual simulation (i.e. about 0.5), and quite the same as that of the best individual of the social simulation (i.e. about 0.1). This means that the individuals of this condition do indeed exploit the information provided by the traces left by other agents as efficiently as the individuals in the social condition. But in this case the traces are proper *signals* and not mere *cues* because the behavior responsible for the trace production has *not only* the practical function of increasing individual fitness *but also* the social function of producing more informative traces for the individuals to come\*.

Hence, this is a paradigmatic case of *indirect behavioral implicit (i.e. stigmergic) communication*. It is indirect because it exploits the traces left in the environment; it is behavioral because these traces are the results of practical actions; it is implicit because the behavior of producing traces has not lost its practical individual function (agents still need to eat, and, in general, do eat for the individual selective advantage provided by eating); it is communication because the same practical behavior (in this case the conditions of its production) has been slightly modified for influencing purposes, thus acquiring also the communicative function of leaving more readable traces.

## 1.7 Future Work

---

Stigmergic approaches are nowadays gaining consensus in several distinct fields of computer science. Due to its relevance for situatedness, stigmergy has been straightforwardly adopted in robotics, and especially by bio-inspired approaches such as Swarm Intelligence or Behavior-based Robotics. As a general mechanism for indirect interaction, it is now raising interest also in Multi-Agent Systems where open environments and huge numbers of heterogeneous agents require new coordination mechanisms. To be fully integrated as first-level abstraction, stigmergy, however, must be more clearly understood than is today. We contend in fact that a principled approach to stigmergy will be essential for providing a new engineering framework for MAS in which both direct and indirect communicative behaviours will be possible at the same time.

In this chapter, we have contributed to a more advanced understanding of this very basic phenomenon by providing a precise definition of what stigmergy is, by distinguishing stigmergic self-adjustment from stigmergic communication, and by presenting operational definitions of these phenomena by means of a set of evolutionary agent-based simulations.

However, indirect and direct communicative systems still need further research in order to understand their specific roles and possibly their mutual advantages and disadvantages.

---

\*The communicative behavior of eating more fruits than expected from a purely individualistic point of view so to leave more informative traces for the agents to come is clearly an altruistic behavior. The evolution of such a behavior is explained by the fact that interacting individuals are clones, meaning that the producer and the receiver of the altruistic behavior always share the same (altruistic) genes. Hence, the altruistic behavior can arise just because the cost for the agent eating one more edible fruit is inferior to the advantage for the forthcoming agent of finding more informative traces.

Effective design of computational systems employing these mechanisms needs more research on how these mechanisms might co-exist together by supplementing each the limits of the other one.

Similarly, explicit and implicit kinds of communication will also be essential at the same time because each of them fits different needs as natural societies, both of human and animal agents, clearly show. In fact, natural language heavily relies upon more implicit communicative capacities to expand the kind of meaning that can be communicated (i.e. behavioral implicit communication is needed to formalize the Gricean speakers meaning). The interplay between implicit and explicit communicative behaviours is however still an open challenge that is left out for future research.

Even more than this, one of the major future challenges for the MAS community lies, in our opinion, in designing self-organizing multi- agent systems capable of evolving by themselves their explicit communicative protocols on the basis of more primitive communicative capacities like the ones we have explored in this chapter.

While in this chapter we have mainly adopted an evolutionary perspective with very simple reactive agents as prototypical models, stigmergy and its uses are available also for “cognitive” agents of the Belief-Desire-Intention kind. The relative benefits of enabling stigmergy between deliberative agents is still however an open issue.

Finally, we are convinced that in human-human, human-agent, human-robot, agent-agent, and robot-robot interaction the possibility of communicating through an action and its effects and products could be extremely relevant. In fact, behavioral communication for coordination (in particular stigmergy) has some nice properties and advantages that deserve to be stressed. It is naturally and intrinsically ‘situated’ in space and time, and thus it transmits information in a perceptual, non-mediated way without any special, arbitrary codification of this kind of information which will necessary be somewhat ‘abstract’. This information has also the very nice feature of not being discrete, digitalized, a characteristic that sometimes might be critical\*.

## 1.8 Conclusion

---

Good old commercial airplanes were flown by the captain and the first officer by means of physically connected control wheels. When the captain was in charge of conducting the plane, the first officer was aware of his steering decisions just by sensing the movements through his own wheel. Differently, modern cockpits took advantage of the advancements in electronics, and have been designed with much more nice and small control wheels very similar to contemporary joysticks. Invisible one from the other, both the captain and first officer now have independent wheels that control the airplane once the corresponding officer is in charge. An unfortunate consequence of this technological innovation has been, however, that the quality of the situation awareness in the cockpit has been compromised [Nor93]. In fact, the officers in these new cockpits need to rely much more on verbal communication to understand what is happening, while situation awareness was guaranteed before by means of a very simple case of stigmergic communication.

This story taken from aviation bears few similarities with what we have discussed in the introduction. The shift from blackboard systems to multi-agent systems has been analogously characterized by the abandonment of a form of stigmergic communication between

---

\*We thank Stefano Nolli with whom we discussed this point.

the agents to adopt the computational variety of verbal communication, that is Agent Communication Languages.

Adopting the perspective defended in this chapter, Multi-Agent Systems are, in fact, characterized by the use of *direct explicit communication* between the agents, that, as it is nowadays widely contested, are difficult to employ in open systems with an unpredictable number of heterogeneous agents. Shared and fixed ontologies and common communication protocols cannot always cope with this complex dynamics.

In this chapter, we have argued for the relevance of a different coordination mechanism: *indirect implicit communication via practical behavioral traces* (i.e. stigmergic communication). To do so, we have provided a clear definition of stigmergic cues as practical behavioral traces. We have also distinguished two varieties of stigmergic uses in self-adjustment and communication, and clarified the difference between explicit and implicit communicative systems.

Though direct and explicit communication and indirect and implicit communication seem, at a first glance, incompatible and alternative strategies, the problems in aviation teach us an important morals: they are not! Having understood the importance of stigmergic communication has helped aviation designers to exploit the relative advantages of both communicative systems without relying only on the properties of one of them. Today's cockpits are designed to take advantage both of implicit and explicit communication between the crew and exploit useful redundancies for safety reasons.

Similarly, the next generation of multi-agent systems should take advantage of all these mechanisms in order to combine them in ways that today are still not foreseeable.

## 1.9 Acknowledgments

---

This work, as part of the European Science Foundation EUROCORES Programme TECT, was supported by the funds from the Italian National Research Council (CNR) and the EC Sixth Framework Programme.

## References

---

1. A.H Bond and L. Gasser, editors. *Readings in Distributed Artificial Intelligence*. Morgan Kaufman, San Mateo, CA, 1988.
2. E. Bonabeau, F. Henaux, S. Guérin, D. Snyers, P. Kuntz, and G. Theraulaz. Routing in telecommunications networks with ant-like agents. In *IATA '98: Proceedings of the second international workshop on Intelligent agents for telecommunication applications*, pages 60–71, London, UK, 1998. Springer-Verlag.
3. O. H. Bruinsma. *An analysis of building behaviour of the termite *Macrotermes subhyalinus* (Rambur)*. PhD thesis, Landbouwhogeschool, Wageningen, 1979.
4. E. Boneabeau, G. Theraulaz, J.L. Deneubourg, N.R. Franks, O. Rafelsberger, J.L. Joly, and S. Blaco. A model for the emergence of pillars, walls and royal chambers in termite nests. *Philosophical Transactions of the Royal Society of London*, 353:1561–1576, 1997.
5. C. Castelfranchi. Principles of individual social action. In Tuomela R. and Hintikka G., editors, *Contemporary Action Theory*. Kluwer, Dordrecht, 1997.
6. C. Castelfranchi. Modelling social action for AI agents. *Artificial Intelligence*, 103:157–182, 1998.
7. C. Castelfranchi. When doing is saying. The theory of behavioral implicit communication. Technical report, ISTC-CNR, 2000.

8. C. Castelfranchi. From conversation to interaction via behavioral communication. In S. Bagnara and G. Crampton Smith, editors, *Theories and Practice in Interaction Design*, pages 157–179. Erlbaum, New Jersey (USA), 2006.
9. C. Castelfranchi. Silent agents: From observation to tacit communication. In J.S. Sichman, H. Coelho, and S.O. Rezende, editors, *Advances in Artificial Intelligence - IBERAMIA-SBIA 2006*, volume 4140 of *Lecture Notes in Computer Science*, pages 98–107. Springer, 2006.
10. R. Conte and C. Castelfranchi. *Cognitive and Social Action*. UCL Press, London, 1995.
11. S. Camazine, J. Deneubourg, N.R. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-Organization in Biological Systems*. Princeton University Press, Princeton, NJ, 2001.
12. P.R. Cohen and H.J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.
13. M. Dorigo, E. Bonabeau, and G. Theraulaz. Ant algorithms and stigmergy. *Future Generation Computer Systems*, 16:851–871, 2000.
14. H.A. Downing and R.L. Jeanne. Nest construction by the paper wasp, polistes: a test of stigmergy theory. *Animal Behavior*, 36:1729–1739, 1988.
15. E. H. Durfee. *Coordination of Distributed Problem Solvers*. Kluwer Academic, Boston, MA, 1988.
16. R. Englemore and T Morgan, editors. *Blackboard Systems*. Addison-Wesley, Reading, MA, 1988.
17. S. Garnier, J. Gautrais, and G. Theraulaz. The biological principles of swarm intelligence. *Swarm Intelligence*, 1(1):3–31, 2007.
18. P.P. Grassé. La reconstruction du nid et les coordinations inter-individuelles chez *bellicositermes natalensis* et *cubitermes*. La théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs. *Insectes Sociaux*, 6:41–81, 1959.
19. P.P. Grassé. Nouvelles expériences sur le termite de Muller (*macrotermes mulleri*) et considérations sur la théorie de la stigmergie. *Insectes Sociaux*, 14(1):73–102, 1967.
20. W. D. Hamilton. Genetic evolution of social behavior. *Journal of Theoretical Biology*, 7(1):1–52, 1964.
21. O. Hasson. Cheating signals. *Journal of Theoretical Biology*, 167:223–238, 1994.
22. C. Hewitt. Viewing control structures as patterns of message passing. *Artificial Intelligence*, 8(3):323–364, 1977.
23. O. Holland and C Melhuish. Stigmergy, self-organization, and sorting in collective robotics. *Artificial Life*, 5:173202, 1999.
24. J. Kennedy, R. C. Eberhart, and Y. Shi. *Swarm intelligence*. Morgan Kaufmann Publishers, San Francisco, CA, 2001.
25. P Karlson and M. Luscher. Pheromones: a new term for a class of biologically active substances. *Nature*, 183:55–56, 1959.
26. I. Karsai and G. Theraulaz. Nest building in a social wasp: postures and constraints. *Sociobiology*, 26:83–114, 1995.
27. M.J. Mataric. Issues and approaches in the design of collective autonomous agents. *Robotics and Autonomous Systems*, 16:321–331, 1995.
28. M. Mirolli and D. Parisi. How can we explain the emergence of a language which benefits the hearer but not the speaker? *Connection Science*, 17(3-4):325–341, 2005.
29. Marco Mirolli and Domenico Parisi. How producer biases can favour the evolution of communication: An analysis of evolutionary dynamics. *Adaptive Behavior*, 16(1):27–52, 2008.

30. J. Maynard-Smith and D. Harper. *Animal Signals*. Oxford University Press, Oxford, 2003.
31. M. Mamei and F. Zambonelli. Pervasive pheromone-based interaction with rfid tags. *ACM Transactions on Autonomous and Adaptive Systems*, 2(2):4, 2007.
32. D. Norman. *Things that make us smart*. Perseus Publishing, Cambridge, MA, 1993.
33. V. Parunak. Go to the ant: Engineering principles from natural agent systems. *Annals of Operations Research*, 75:69–101, 1997.
34. V. Parunak. A survey of environments and mechanisms for human-human stigmergy. In D. Weyns, V. Parunak, and F. Michel, editors, *Environments for Multiagent Systems II*, number 3830 in Lecture Notes in Computer Science, pages 163–186. Springer-Verlag, Berlin / Heidelberg, 2006.
35. V. Parunak, S. Brueckner, and J. Sauter. Digital pheromones for coordination of unmanned vehicles. In D. Weyns, V. Parunak, and F. Michel, editors, *Environments for Multiagent Systems I*, number 3374 in Lecture Notes in Computer Science, pages 179–183. Springer-Verlag, Berlin / Heidelberg, 2005.
36. M.D. Sadek. A study in the logic of intention. In *Proceedings of Knowledge Representation and Reasoning*, pages 462–473, 1992.
37. J. R. Searle. *Speech Acts: an Essay in Philosophy of Language*. Cambridge University Press, Cambridge, 1969.
38. R.G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers C*, 29(12):1104–1113, 1980.
39. P.W. Sorensen and N.E. Stacey. Evolution and specialization of fish hormonal pheromones. In Johnston et al., editor, *Advances in chemical signals in vertebrates*. Kluwer Academic, New York, 1999.
40. G. Theraulaz and E. Bonabeau. A brief history of stigmergy. *Artificial Life*, 5(2):97–116, 1999.
41. L. Tummolini and C. Castelfranchi. Trace signals: The meanings of stigmergy. In D. Weyns, V. Parunak, and F. Michel, editors, *Environments for Multi-Agent Systems III*, number 4389 in Lecture Notes in Artificial Intelligence, pages 141–156. Springer-Verlag, Berlin / Heidelberg, 2007.
42. L. Tummolini, C. Castelfranchi, A. Ricci, M. Viroli, and A. Omicini. Exhibitionists and voyeurs do it better: A shared environment approach for flexible coordination with tacit messages. In D. Weyns, V. Parunak, and F. Michel, editors, *Environments for Multiagent Systems I*, number 3374 in Lecture Notes in Computer Science, pages 215–231. Springer-Verlag, Berlin / Heidelberg, 2005.
43. N. Tinbergen. Derived activities: their causation, biological significance, origin and emancipation during evolution. *Quarterly Review of Biology*, 27:1–32, 1952.
44. E. Visalberghi and E. Addessi. Seeing group members eating familiar food enhances the acceptance of novel foods in capuchin monkeys. *Animal Behaviour*, 60:69–76, 2000.
45. J.W. Wenzel. Evolution of nest architecture. In K.G. Ross and R.W. Matthews, editors, *The social biology of wasps*, pages 480–519. Cornell University Press, 1991.
46. E.O. Wilson. Source and possible nature of the odor trail of fire ants. *Science*, 129(3349):643, 1959.
47. E.O. Wilson. *Sociobiology: The New Synthesis*. Harvard University Press, Cambridge, MA, 1975.
48. D. Wyatt. *Pheromones and Animal Behaviour: Communication by Smell and Taste*. Cambridge University Press, Cambridge, 2003.