



ÉCOLE POLYTECHNIQUE

Thèse pour l'Obtention du Titre de Docteur
Domaine: Sciences Humaines et Sociales
Spécialité: Sciences Sociales et Sciences Cognitives

CAMILLE ROTH

Co-evolution in Epistemic Networks Reconstructing Social Complex Systems

—

Co-évolution dans les réseaux épistémiques
Un exemple de reconstruction en sciences sociales

soutenue le 19 novembre 2005

Jury

HENRI BERESTYCKI	CAMS, EHESS	Président
PAUL BOURGINE	CREA, CNRS & Ecole Polytechnique	Directeur de thèse
DAVID A. LANE	Université de Modène, Italie	Examineur
MICHEL MORVAN	ENS-Lyon & EHESS	Rapporteur
DOUGLAS R. WHITE	Université de Californie-Irvine, Etats-Unis	Rapporteur

“L’Ecole Polytechnique n’entend donner aucune approbation, ni improbation, aux opinions émises dans cette thèse, ces opinions doivent être considérées comme propres à leur auteur”

Acknowledgements

I wish to express my deepest gratitude to my advisor Paul Bourguine for having directed this research work, especially for the challenging discussions we had and his ever-rigorous mathematical views.

I wish to thank Michel Morvan and Douglas White for having accepted to be reviewers (“*rapporteurs*”) of my work, and for the relevant advices they gave me towards the completion of the present manuscript. I also wish to thank Henri Berestycki and David Lane for serving as members of the jury.

This work has been carried at the CREA (Centre de Recherche en Epistémologie Appliquée) of the Ecole Polytechnique: I would like to thank its director, Jean Petitot, and its members, researchers, graduate students, assistants, for their conviviality, thoughtful advices and intellectual enlightenment. The lab, in particular, always provided me the material means I needed — this tremendously facilitated the achievement of my work. Thanks also to the CNRS, for being confident in my research proposal and the subsequent 3-year funding they were kind enough to provide me.

I had the occasion to interact with many people during my thesis, some I even had the pleasure to collaborate with, yet all of them have closely or loosely helped me and contributed to the advancement of my research. As such, I cannot envisage to comprehensively *and* fairly acknowledge all of them — I must nonetheless thank in particular Michel Bitbol, David Chavalarias, Jean-Philippe Cointet, Matthieu Latapy, Clémence Magnien, Sergei Obiedkov, Nadine Peyriéras, Thierry Rayna, Richard Topol and Douglas White. I also had many interesting interactions with several members of the EU-funded ISCOM project (“Information Society as a COMplex system”) coordinated by David Lane, and the CNRS-funded PERSI project (“Programme d’Etude des Réseaux Sociaux et de l’Internet”) coordinated by Matthieu Latapy — I thank both of them for involving me into these projects.

Special thanks go to my parents & my friends, for supporting me — mind the gallicism...

Contents

General introduction	9
I Knowledge Community Structure	15
Introduction	17
1 Epistemic communities	21
1.1 Context	21
1.2 Definitions	23
1.3 Formal framework	25
2 Building taxonomies	31
2.1 Taxonomies and lattices	31
2.2 Galois lattices	32
2.3 GLs and categorization	34
2.3.1 About relevant categorization	36
2.3.2 Assumptions on EC structure	37
2.3.3 GLs and selective categorization	38
2.4 Comparison with different approaches	39
3 Empirical results	43
3.1 Experimental protocol	43
3.2 Results and comparison with random relations	45
3.2.1 Empirical versus random	46
3.2.2 Rebuilding the structure	47
4 Community selection	51
4.1 Rationale	51
4.2 Selection methodology	53

5	Taxonomy evolution	57
5.1	Empirical protocol	59
5.2	Case study, dataset description	60
5.3	Rebuilding history	61
5.3.1	Evolution description	61
5.3.2	Inference of an history	63
5.3.3	Comparison with real taxonomies	64
6	Discussion and conclusion	67
II	Micro-foundations of epistemic networks	73
	Introduction	75
7	Networks	77
7.1	Global overview	77
7.2	A brief survey of growth models	79
7.3	Epistemic networks	81
8	High-level features	85
8.1	Empirical investigation	85
8.2	Degree distributions	85
8.3	Clustering	89
8.4	Epistemic community structure	93
9	Low-level dynamics	97
9.1	Measuring interaction behavior	97
9.1.1	Monadic PA	99
9.1.2	Dyadic PA	100
9.1.3	Interpreting interaction propensions	100
9.1.4	Activity and events	101
9.2	Empirical PA	103
9.2.1	Degree-related PA	103
9.2.2	Homophilic PA	105
9.2.3	Other properties	107
9.2.4	Concept-related PA	108
9.3	Growth- and event-related parameters	109
9.3.1	Network growth	109
9.3.2	Size of events	110
9.3.3	Exchange of concepts	112

10 Towards a rebuilding model	115
10.1 Outline	115
10.2 Design	117
10.3 Results	120
10.4 Discussion	123
Conclusion	125
III Coevolution, Emergence, Stigmergence	129
Introduction	131
11 Appraising levels	133
11.1 Accounting for levels	134
11.2 Emergentism	134
11.3 What levels are not	137
11.4 Observational reality of levels	138
11.4.1 Different modes of access	138
11.4.2 Illustrations	140
12 Complex system modeling	143
12.1 Complexity and reconstruction	143
12.1.1 Objectives	143
12.1.2 Commutative decomposition	144
12.1.3 Reductionism failure	145
12.1.4 Emergentism	146
12.2 A multiple mode of access	146
12.2.1 The observational viewpoint	146
12.2.2 Introducing new levels	149
12.2.3 Rethinking levels	150
13 Reintroducing retroaction	153
13.1 Differentiating objects	153
13.2 Agent behavior, semantic space	154
13.3 Coevolution of objects	156
13.4 "Stigmergence"	156
Conclusion	159

General conclusion	161
List of figures	168
References	169
Index	189

General introduction

Agents producing, manipulating, exchanging knowledge are forming as a whole a socio-semantic complex system: a complex system made of agents who work on and are influenced by semantic content, by flows of information in which they are fully immersed but, at the same time, on which they can have an impact and leave their footprints. Social psychologists and epistemologists, *inter alia*, have already a long history in studying the properties of such knowledge communities. Yet, the massive availability of informational content and the potential for extensive interactivity has made the focus slip from single “groups of knowledge” to the entire “society of knowledge”. Simultaneously, the change in scale has called for the use of new methods, as well as the characterization of new phenomena, with knowledge being distributed and appraised on a more horizontal basis — in a networked fashion. On the other hand, many different “sub-societies” of knowledge co-exist, possibly overlapping and interwoven, although usually easily distinguished by their means, methods, and people.

Reconstruction issues Therefore, the research community has taken a renewed and unprecedented interest in studying these communities, in both a theoretical and a practical perspective:

- theoretically, it conveys the hope of naturalizing further social sciences.
- practically, it entails several potential applications — as regards research policy in particular, since scientists themselves form a knowledge community; but also as a means for political planning, innovation diffusion improvement, to cite a few.

The present thesis lies within the framework of this research program. Specifically, we aim to know and be able to model the behavior and the dynamics of such knowledge communities. Alongside, we address more broadly the question of reconstruction in social science, and notably the reconstruction of the evolution of a social complex system. Reconstruction is a reverse problem consisting fundamentally in *successfully reproducing several stylized facts observed in the original empirical*

system. To this end, we distinguish the *lower level* of microscopic objects (including agents, agent-based interactions, etc.), and the *higher level* of macroscopic descriptions (communities, global structures). Thus, we wish to know whether it is possible to:

- (i) deduce high-level observations of such a system from strictly low-level phenomena; and
- (ii) reconstruct the evolution of high-level observations *from* the dynamics of lower-level objects.

For instance, social scientists are using more and more frequently social network analysis to infer high-level phenomena which would have traditionally undergone a strictly high-level description: qualifying the cohesion of a community, finding the roots of a crisis, explaining how roles are distributed, etc. By doing so, they are clearly carrying an analysis related to the *first issue*, “(i)”: they exhibit a formal relationship between higher and lower level objects — they reconstruct the “social structure” (Freeman, 1989), benchmarked against classically proven high-level descriptions. In this respect they make the assumption that the chosen lower level (for instance a social network) yields enough information about the phenomenon; the benefit being often that low-level information is easier to collect and entails more robust descriptions. **In formal terms, the first issue is equivalent to the following question: given a high-level phenomenon H , and low-level objects L , is there a P such that $P(L) = H$, for any empirically valid pair L and H ? — then, how to find it?** This approach must be accurate in an evolutionary framework as well: given empirical dynamics λ^e and η^e on L and H respectively, such that for any time t :

$$\begin{cases} \lambda^e(L_t) = L_{t+\Delta t} \\ \eta^e(H_t) = H_{t+\Delta t} \end{cases} \quad (1)$$

we must find a P such that:

$$P \circ \lambda^e = \eta^e \circ P \quad (2)$$

In other words, we must have $P(L_{t+\Delta t}) = H_{t+\Delta t}$: it must be possible to describe the final observation on H from the evolution of L . The reconstruction scheme is detailed on Fig. 1, the commutative diagram in particular is encountered in the context of dynamical systems — see (Rueger, 2000) and references herein, and (?; Turner & Stepney, 2005).

Thereafter, once P is defined, the *second issue*, “(ii)”, is to show that a low-level dynamics enables the reconstruction of the higher level dynamics. This approach is generally a traditional problem of modeling, although in our framework we insist on the constraint that low-level objects, not high-level descriptions, play a




Figure 1: The *reconstruction problem* comes to find (i) a valid P (the projection P from L onto H is valid if, knowing the empirical dynamics η^e and λ^e , the above diagram commutes, i.e. $P \circ \lambda^e = \eta^e \circ P$) and (ii) a satisfying λ (i.e. such that $P \circ \lambda = P \circ \lambda^e$). See (Rueger, 2000; ?) for comprehensive discussions on this kind of diagrams.

central role (Bonabeau, 2002). **Thus, the *second issue* comes to find a dynamics λ such that it correctly reproduces the empirical high-level dynamics η^e , through P .** As such, the model objectives are restricted to rebuilding high-level phenomena. Indeed, the point is not necessarily to find a dynamics λ yielding empirically valid low-level phenomena (i.e. such that we have $\lambda(L_t) = L_{t+\Delta t}$), but simply to find λ such that the desired high-level objects are correctly described (i.e. *only* $P \circ \lambda(L_t) = H_{t+\Delta t}$ must hold). Thus, the fact that $\lambda \neq \lambda^e$ or that $L_{t+\Delta t} \neq \lambda(L_t)$ is not problematic, *as long as* $P \circ \lambda = P \circ \lambda^e$: λ needs not be a model of λ^e , and the knowledge of L_t needs not be perfect; it only needs to be valid “through P .” This allows successful reconstruction even when it is not possible to describe λ^e comprehensively, or when L is imperfectly known — only reconstructed high-level descriptions have to be accurate. For instance, being unable to predict the actual number of friends of a given agent (a specific fact on L) should not prevent us from rebuilding the fact that the distribution of acquaintances follows a power-law (a specific fact on H).

Reconstructing a knowledge community We may now focus on the above-mentioned social complex system, a knowledge community, for which our thesis solves a reconstruction problem. We will indeed rebuild several aspects of the structure of such a community — these are high-level phenomena. Foremost among these aspects is the description of the community in smaller, more precise sub-communities. Here an “epistemic community” is understood as a descriptive instance only, not as a coalition of people who have some interest to stay in the community: it is a set of agents who simply share the same knowledge concerns.

Epistemologists traditionally describe a whole field of knowledge by characterizing and ordering its various epistemic communities, and they basically achieve this task by gathering communities in a hypergraph, which we call *epistemic hypergraph*. A *hypergraph* is a graph where edges can connect groups containing more than two nodes.

We thus support the following *thesis*: the structure of a knowledge community, and in particular its epistemic hypergraph, is primarily produced by the co-evolution of agents and concepts.

In the first part, we will propose a method for exhibiting a hierarchical epistemic hypergraph for any given community. More precisely, we will exhibit a P that yields H (the community structure) from L (agent and concept-based descriptions) — this corresponds to the first issue. Given the assumptions, an adequate and efficient method for achieving this task consists in using Galois lattices. By checking the adequation between the resulting hypergraph and an empirical high-level epistemological description of the knowledge community — i.e. of the kind epistemologists would produce and work on — we will confirm the validity of the projection. Better, for any time t , P will yield H_t from L_t , and as such, given the empirical low-level dynamics λ^e , we will reproduce the empirical high-level dynamics η^e . This provides subsequently a formal way of partially defining the field of “scientometrics”, which consists in describing scientific field and paradigm evolution from low-level quantitative data.

Further, in the second part, we will *micro-found* the high-level phenomena in the dynamics of the lower level of agents and concepts — this addresses the second issue. More precisely, we will introduce a co-evolutionary framework based on a social network, a semantic network and a socio-semantic network; as such an epistemic network made of agents, concepts, and relationships between all of them. We will then show that dynamics at the level of this epistemic network are sufficient to reproduce several stylized facts of interest. Given H and the empirical dynamics η^e on H , we will therefore propose methods to design λ from low-level empirical data on L such that $P \circ \lambda(L) = \eta^e \circ P(L)$. Since the dynamics will be based on the co-evolution at the the lower level L of the epistemic network, we will substantiate our claim that epistemic communities are produced by the co-evolution of agents and concepts.

It is nonetheless worth noting that the co-evolution occurs at the lower level of the three networks only. We are thus within the framework of “*simple emergence*”: the high-level is deduced from the lower level, but the lower level is to be influenced by low-level phenomena only. In addition, we will underscore the fact that exogeneous phenomena may also account for the social complex system evolution (including for instance ‘strength’ of concepts, external policies, etc.). We will consequently moderate the thesis, arguing eventually that reconstructing epistemic

communities involves *at least* the dynamic co-evolution of agents and concepts.

In the third and last part, we will defend a more general epistemological point on the methods and achievements of this kind of reconstruction. We will notably situate our effort within the whole apparatus of complex system appraisal. In this respect, we will suggest in particular that a successful rebuilding is no more than a claim that some particular high-level stylized facts, observed with high-level instruments (epistemologists and experts in our case) can be *fully deduced* from low-level objects (here, the epistemic network). As such, reduction of a high-level to a lower level should be understood as the successful full deduction of the higher-level from a relevantly chosen lower level. This remark will eventually support our choice of a co-evolutionary framework.

Part I

Knowledge Community Structure

Summary of Part I

In this part, we introduce a formal framework based on Galois lattices that categorizes epistemic communities automatically and hierarchically, rebuilding a whole community taxonomy in the form of a *hypergraph of significant sub-communities*. The longitudinal study of these static pictures makes historical description possible, by capturing stylized facts such as field emergence, decline, specialization and interaction (merging or splitting). The method is applied to empirical data and successfully validated by categories and *histories* given by domain experts. We thus design a valid projection function P from a low-level defined by links between agents and concepts to the high-level of epistemological descriptions.

Introduction of Part I

Scientists, journalists, political activist groups, socio-cultural communities with common references are various instances of the so-called *society of knowledge*. They are in all respects smaller, embedded “sub-societies” of knowledge, with their own norms, methods, and specific topics; as such independent to some extent, though possibly partially overlapping. Yet, it is remarkable that any knowledge community, whatever its level of generality — the whole society, the scientific community, biologists, embryologists, embryologists working on a particular model-animal — appears to be structured in turn in various implicit subcommunities, with each subgroup contributing to knowledge creation in a *distributed* and *complementary* manner. Expertise seems indeed to be heterogeneously distributed over all agents, with different levels of specificity and distinct areas of competence: there are very few topics that all agents are able to deal with. As specialization occurs, knowledge communities become subsequently more structured: boundaries appear between subgroups, both *horizontally*, with the appearance of several branches, and *vertically*, with different levels of generality for appraising a given topic.

In this part of our thesis, we propose a method for **building, ordering and appraising the epistemic hypergraph of a given knowledge community**, which as a result can be compared to high-level descriptions of the knowledge community structure. The epistemic hypergraph is a graph of knowledge communities, where each community gathers both agents and concepts. At first sight, we denote by knowledge community, or *epistemic community*, any kind of group of agents who are interested in some common knowledge issues: a group of research for instance investigating a precise topic, a whole field of research, a larger scientific field, a paradigm; besides, the notion is also not necessarily restricted to academic groups. A knowledge community needs not be a community of practice (Lave & Wenger, 1991; Wenger & Snyder, 2000) because its agents need not be acquainted or involved in a common practical task; although a community of practice is certainly a special type of knowledge community. On the whole, agents involved in a same epistemic community interact using shared paradigms, meanings, judgments, opinions (Haas, 1992; Cowan *et al.*, 2000), all of which being to a certain

extent publicly available concepts, especially in larger scale communities. Therefore, in itself, an “epistemic complex system” achieves widespread social cognition: new concepts are being introduced by some agents, others work on them, build upon them, refine, falsify, improve, etc. This phenomenon has even been recently sensibly boldened by the fact that the whole process of knowledge elaboration has slipped from a rather centralized, well-recognized organization to a mainly decentralized, collectively interactive and networked system. Thus, while agents can *potentially* have access and be synchronized with a large part of the knowledge produced by the whole epistemic community, they *actually* have access only to a small portion of it, prominently because of cognitive and physical limitations. In this respect, it should be of utmost interest to have tools **enabling agents to understand the structure and the activity of their knowledge community, at any level of specificity or generality.**

More precisely, in any kind of epistemic community, agents have an implicit knowledge of the structure of the larger global community they are participating in. Embryologists know what molecular biology, biology, and science in general are about. Their knowledge is thus meta-knowledge: it is knowledge on the structure of their own knowledge communities. They can name several other fields, issues they know are close, related to their knowledge concerns, or not. Agents can distinguish various levels of specificity as well, pragmatically knowing that a given set of topics is usually a subfield of another larger field, or has affiliations with several fields, roughly knowing when knowledge communities intersect in what appears to be interdisciplinary, cross-domain enterprises.

Yet, as a matter of scalability agents have a limited and *subjective* knowledge of the extent of the community they are evolving in. As such their meta-knowledge resembles that of a folk taxonomy, in the anthropological sense, that is, a taxonomy proper to an individual (or shared by a small-sized group) and made of its own experience, as opposed to scientific taxonomies, deemed objective and systematic (Berlin, 1992). Hence, epistemologists often have the last word in elaborating and validating credible meta-knowledge. Expert-made taxonomies are prodigiously more reliable than folk taxonomies, in particular because of their tangible methodology. However, again because of scalability, elaborating this meta-knowledge still lacks precision, takes an enormous amount of work, and rarely focuses on precise groups of agents nor investigates comprehensively the whole community; in addition, the result may be biased by a particular approach on the field.

Here, we will thus study the large-scale structure of epistemic complex systems. *In fine*, we wish to introduce a method for creating automatically a taxonomy of knowledge fields — in other words, for producing a hierarchic epistemic hypergraph of the community structure (a high-level description $P(L)$ from low-level empirical data L). This hypergraph should make clear (i) which fields, disciplines,

trends, schools of thought are to be found in such an epistemic network, and (ii) what kind of relationships they entertain. In turn, the resulting taxonomy should prove consistent with the already-existing intersubjective perception of the field, which will thus be the benchmark of our procedure (the empirical H , to compare to the $P(L)$ produced by the method). Eventually, knowing the taxonomy at any given time, we should be able to describe the evolution of the system; and as such achieve a reconstruction of the history of the community *on objective grounds*.

The outline of this part is as follows: after having presented the context and introduced the formal framework (Chap. 1), we describe how to categorize epistemic communities in an hierarchically structured fashion using Galois lattices (Barbut & Monjardet, 1970) (Chap. 2) and produce a lattice-based representation of the whole knowledge community. We then apply it to empirical data, successfully comparing our results with the expected categories given by domain experts (Chap. 3). Chapter 4 details the way we build reduced taxonomies, or community hypergraphs, and Chapter 5 addresses their evolution. In particular, field progress or decline, field scope enrichment or impoverishment, and field interaction (merging or splitting) are observed in a dynamic case study. Settled both in applied epistemology and scientometrics, this approach would ultimately provide agents with processes enabling them to know dynamically their community structure.

Our main source of data is MedLine, a database maintained by the US National Library of Medicine and containing more than 11 million references to health sciences articles published in about 3,700 journals worldwide. We narrow our study to articles dealing with the “zebrafish,” a fish whose embryo is translucent and developing fast, therefore widely used as a model animal by embryologists.¹

¹Portions of this part can be found in more details in (Roth & Bourguine, 2005; Roth & Bourguine, 2006; Roth & Bourguine, 2003).

Chapter 1

Epistemic communities

In this chapter, we present the existing works concerning epistemic community appraisal and representation, and we introduce a formal framework along with various definitions.

1.1 Context

Several works ranging from social epistemology to political science and economics have given an account of the collaboration of agents within the same epistemic framework and towards a given knowledge-related goal, namely knowledge creation or validation. For social epistemologists, it is a scientist group, or *epistemic community*, producing knowledge and recognizing a given set of conceptual tools and representations — the “paradigm,” according to Kuhn (1970) — possibly working in a distributed manner on specialized tasks (Schmitt, 1995; Giere, 2002). Considering a whole knowledge field as a huge epistemic community (e.g. biology, linguistics), one can see subdisciplines as smaller, embedded, and more specific epistemic communities — subfields within a paradigm. Haas (1992) introduced the notion of epistemic community as “*a network of knowledge-based experts (...) with an authoritative claim to policy-relevant knowledge within the domain of their expertise.*” Cowan, David and Foray (2000) added that an epistemic community must *share* a subset of concepts. To them, an epistemic community is “*a group of agents working on a commonly acknowledged subset of knowledge issues and who at the very least accept a commonly understood procedural authority as essential to the success of their knowledge activities.*” The “common concern” aspect has been emphasized by Dupouet, Cohendet and Creplet (2001) who define an epistemic community as “*a group of agents sharing a common goal of knowledge creation and a common framework allowing to understand this trend.*” These authors nevertheless acknowledge the need of a notion of authority and deference.

On the other hand, scientists have shown an increasing interest for methods of knowledge community structure analysis. Several conceptual frameworks and automated processes have been proposed for finding groups of agents or documents related by common concepts or concerns, notably in knowledge discovery in databases (KDD) (Rocha, 2002; Hopcroft *et al.*, 2003) and scientometrics (Leydesdorff, 1991a; Lelu *et al.*, 2004). Dealing with and ordering categories automatically has indeed become central in data mining and related fields (Jain *et al.*, 1999), along with the massive development of informational content. Besides, since a large amount of data is freely and electronically available, the study of scientific communities in particular has attracted a large share of the interest — especially biologist communities: biology is a domain where the need for such techniques is also the most pressing because article production is so high that it becomes hard for scientists to figure out the evolution of their own community.

Yet, existing approaches in community finding are often either based on social relationships only, with community extraction methods stemming from graph theory applied to social networks (Wasserman & Faust, 1994), or on semantic similarity only, namely clustering methods applied to document databases where each document is considered as a vector in a semantic space (Salton *et al.*, 1975). There have been few attempts to link social and semantic aspects, although the various characterizations of an epistemic community insist on its duality, i.e. the fact that such a community is on one side a group of agents who, on the other side, share common interests and work on a given subset of concepts. By contrast, only scientometrics have developed a whole set of methods for characterizing specifically such communities, working on both scientists and the concepts they use. Categorization has been notably applied to scientific community representation, using *inter alia* multidimensional scaling in association with co-citation data (McCain, 1986; Kreuzman, 2001) or other co-occurrence data (Callon *et al.*, 1986; Noyons & van Raan, 1998), in order to produce two-dimensional cluster mappings and track the evolution of paradigms (Chen *et al.*, 2002).

Along with this profusion of community-finding methods, often leaning towards AI-oriented clustering, an interesting issue concerns the representation of communities in an ordered fashion. On the whole, many different techniques have been proposed for producing *and* representing categorical structures including, to cite a few, hierarchical clustering (Johnson, 1967), Q-analysis (Atkin, 1974), formal concept analysis (Wille, 1982), information theory (Leydesdorff, 1991b), block-modeling (White *et al.*, 1976; Moody & White, 2003; Batagelj *et al.*, 2004), graph theory-based techniques (Newman, 2004; Radicchi *et al.*, 2004), neural networks (Kohonen, 2000), association mining (Srikant & Agrawal, 1995), and dynamic exploration of taxonomies (Sacco, 2000). Here, the notion of *taxonomy* is particularly relevant with respect to communities of knowledge. A taxonomy is a hierarchi-

cal structuration of things into categories, as such an ordered set of categories (or *taxons*), and is a fundamental tool for representing groups of items sharing some properties. Taxonomies are useful in many different disciplinary fields: in biology for instance, where classification of living beings has been a recurring task (Whittaker, 1969; Simpson & Roger, 2004); in cognitive psychology for modeling categorical reasoning (Rosch & Lloyd, 1978; Barthélemy *et al.*, 1996); as well as in ethnography and anthropology with folk taxonomies (Berlin, 1992; Lopez *et al.*, 1997; Atran, 1998). While taxonomies have initially been built using a subjective approach, the focus has moved to formal and statistical methods (Sokal & Sneath, 1963; Benzécri, 1973).

However, taxonomy building itself is generally poorly investigated; arguably, taxonomy evolution during time has been fairly neglected. Our intent here is to address both topics: build a taxonomy of epistemic communities, then monitor its evolution — as such a work which shares the aims of history of science. At the same time while taxonomies have long been represented using tree-based structures, we wish to produce taxonomies which deal with sub-communities affiliated with multiple communities (such as interdisciplinary groups) or of diverse paradigmatic statuses (i.e., rendering equally communities centered around methods, processes, fields of application, given objects, etc.); therefore introducing lattice-based structures.

1.2 Definitions

Basically, we are first trying to know (i) which agents share the same concerns and work on the same concepts, and (ii) which these concerns or concepts are. We are thus farther from the epistemological point of view and need not characterize authoritative groups and their role. Hence, the definitions of an “epistemic community” introduced in the previous section seem to be too precise with respect to authoritative and normative properties, while they lack the ability to formalize community boundaries and extents accurately. Obviously, an epistemic community that is simply characterized by common knowledge concerns should not necessarily be a social community, with agents of the same community enjoying some sort of social link: it is neither a department nor a group of research. In addition, we want a definition that allows some flexibility in the sense that an agent or a semantic item (or *concept*) can belong to several communities. Therefore, we adopt the following definition, keeping the notion of common “knowledge issues”, to which we add *maximality*:

Definition EC-1 (Epistemic community). *Given a set of agents S , we consider the concepts they have in common and we call epistemic community of S the largest set of*

agents who also use these concepts.

In other words, taking the epistemic community (EC) of a given agent set extends it to the largest community sharing its concepts. This notion is to be compared with the *structural equivalence* introduced in sociology by F. Lorrain and H. White (1971). Structural equivalence describes a community as a group of people related in an identical manner to a set of other people. When extending this concept to a group of people related identically to the same concept set, ECs are groups of agents related in an *equivalent* manner to some concepts.

Definition EC-1 is based on an agent set, and we could define correspondingly an epistemic community as the largest set of concepts commonly used by agents who share a given concept set. We will at first focus on agent-based epistemic communities, keeping in mind that concept-based notions are defined strictly equivalently and in a dual manner. In order to set up a comprehensive framework allowing to work on these notions, we now introduce a few basic definitions:

Definition 1 (Intension). *The intension of a set of agents S is the set of concepts which are used by every agent in S .*

Definition 2 (Epistemic group). *An epistemic group is a set of agents provided with its intension, i.e. a group of agents and the concepts they have in common.*

Consider for instance that some given agents s_1 , s_2 and s_3 work on “linguistics” (Lng), while “neuroscience” (NS) is being used by s_2 , s_3 and s_4 (Fig. 1.1). Therefore, the intension of $\{s_1, s_2, s_3\}$ is {Lng}, that of $\{s_2, s_3, s_4\}$ is {NS} and that of $\{s_2, s_3\}$ is {Lng, NS}. Some epistemic groups of this example are thus $(\{s_1, s_2, s_3\}; \{\text{Lng}\})$, $(\{s_2, s_3\}; \{\text{Lng, NS}\})$ and $(\{s_1, s_4\}; \{\emptyset\})$.

For a given set of agents S , knowing its epistemic community comes to identifying the largest group of people who share the same knowledge issues as those of agents of S (this largest group thereby includes S) — notably, for a group of agents *prototypic* of a field, this amounts to know the whole set of agents of the field.

Definition 3 (Hierarchy, maximality). *An epistemic group is larger than another epistemic group if and only if (i) their intensions are the same and (ii) the agent set of the former contains that of the latter.*

An epistemic group is said maximal if there exists no larger epistemic group.

This statement enables us not only to compare epistemic groups but also and more significantly to expand a given epistemic group to its maximal social size. Interpreting definition EC-1 within this framework leads to the following reformulation:

Definition EC-2 (Epistemic community). *The epistemic community based on a given agent set is the corresponding maximal epistemic group.*




Figure 1.1: Sample community, and relationships between agents s_1, s_2, s_3, s_4 and concepts “linguistics” (Lng), “neuroscience” (NS) and “prosody” (Prs) (dashed lines).

The epistemic community based on $\{s_4\}$, for instance, is thus $(\{s_2, s_3, s_4\}; \{NS\})$, and the one based on either $\{s_1\}$ or $\{s_1, s_2\}$ is $(\{s_1, s_2\}; \{Prs, Lng\})$.¹

Notice that we can similarly define an *EC based on a concept set* as the largest set of concepts sharing a given agent set. We introduce the concept-based notions, defined symmetrically to the agent-based notions, and thus, in the remainder of the thesis we will equivalently denote an EC by its agent set S , its concept set C or the couple (S, C) .

Definition 4 (Extension, concept-based notions). *The extension of a set of concepts C is the set of agents using every concept in C . A concept-based epistemic group is a set of concepts provided with its extension. A concept-based epistemic group is larger than another one if and only if (i) their extension are the same and (ii) the concept set of the former contains that of the latter. A concept-based epistemic community is a maximal concept-based epistemic group.*

1.3 Formal framework

In order to work formally on these notions, we need to bind agents to concepts through a binary relation \mathcal{R} between the whole agent set S and the whole concept set C . \mathcal{R} expresses any kind of relationship between an agent s and a concept c . The nature of the relationship depends on the hypotheses and the empirical data. In our case, the relationship represents the fact that s used c (e.g. in some article).

¹The epistemic community based on $\{s_2\}$ is however $(\{s_2\}; \{Prs, Lng, NS\})$; this accounts notably for the fact that s_2 can belong *both* to a generic community and to a more specific or multidisciplinary community: $(\{s_2\}; \{Prs, Lng, NS\})$ vs. $(\{s_1, s_2\}; \{Prs, Lng\})$ — see section 2.3.2 for more details.

Sets and relations Let us consider $\mathcal{R} \subseteq \mathbf{S} \times \mathbf{C}$ binding \mathbf{S} to \mathbf{C} . We introduce the operation “ \wedge ” such that for any element $s \in \mathbf{S}$, s^\wedge is the set of elements of \mathbf{C} which are \mathcal{R} -related to s . Extending this definition to subsets $S \subseteq \mathbf{S}$, we denote by S^\wedge the set of elements of \mathbf{C} \mathcal{R} -related to *every* element of S , namely:

$$s^\wedge = \{c \in \mathbf{C} \mid s\mathcal{R}c\} \quad (1.1a)$$

$$S^\wedge = \{c \in \mathbf{C} \mid \forall s \in S, s\mathcal{R}c\} \quad (1.1b)$$

Similarly, “ \star ” is the dual operation so that $\forall c \in \mathbf{C}, \forall C \subseteq \mathbf{C}$,

$$c^\star = \{s \in \mathbf{S} \mid s\mathcal{R}c\} \quad (1.2a)$$

$$C^\star = \{s \in \mathbf{S} \mid \forall c \in C, s\mathcal{R}c\} \quad (1.2b)$$

By definition we set $(\emptyset)^\wedge = \mathbf{C}$ and $(\emptyset)^\star = \mathbf{S}$.

Definitions 1, 2 and 4 mean that if S is a set of agents, S^\wedge denotes its intension, the set of concepts used by every agent in S (“ $\forall s \in S$ ”). Similarly if C is a concept set, C^\star is its extension, the set of agents who use every concept in C . Thus, epistemic groups are couples of kind (S, S^\wedge) or (C^\star, C) . On the sample community described on Fig. 1.1, we have for instance $\{s_1, s_3\}^\wedge = \{\text{Lng}\}$ and $\{\text{NS}, \text{prs}\}^\star = \{s_3\}$. As Wille (1997) points out, this formalism constitutes a robust and rigorous way of dealing with abstract notions (in a philosophical sense), characterized by their *extension* (physical implementation) and their *intension* (properties or internal content). Here, concepts are *properties* of authors who use them (they are skills in scientific fields, i.e. cognitive properties) and authors are *loci* of concepts (concepts are implemented in authors).

Properties These operations enjoy the following properties:

$$S \subseteq S' \Rightarrow S'^\wedge \subseteq S^\wedge \quad (1.3a)$$

$$C \subseteq C' \Rightarrow C'^\star \subseteq C^\star \quad (1.3b)$$

which means that the intension of a larger agent set is smaller, because more agents share less. We also have:

$$(S \cup S')^\wedge = S^\wedge \cap S'^\wedge \quad (1.4a)$$

$$(C \cup C')^\star = C^\star \cap C'^\star \quad (1.4b)$$

In other words, the intension of two agent sets is the intersection of their respective intensions because a group of agents has in common what its individuals share. Moreover, we can easily derive from (1.4) the words used by a community $S \cup S'$

by taking the intersection $S^\wedge \cap S'^\wedge$, or the authors corresponding to the union of any two sets of concepts $C \cup C'$ by taking $C^* \cap C'^*$. Accordingly,

$$S^\wedge = \left(\bigcup_{s \in S} \{s\} \right)^\wedge = \bigcap_{s \in S} s^\wedge \quad (1.5a)$$

$$C^* = \left(\bigcup_{c \in C} \{c\} \right)^* = \bigcap_{c \in C} c^* \quad (1.5b)$$

We can also conveniently read s_i^\wedge on rows and c_j^* on columns of a matrix R representing relation \mathcal{R} , as follows:

$$R = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}$$

where $R_{i,j}$ is non-zero when $s_i \mathcal{R} c_j$. For instance, $s_4^\wedge = \{NS\}$ and $\{Lng, NS\}^* = \{s_2, s_3\}$ (see Fig. 1.1).

Closure operation More important, the following property holds:

$$S \subseteq S^{\wedge*} \quad (1.6a)$$

$$C \subseteq C^{*\wedge} \quad (1.6b)$$

And thus:

Proposition 1.

$$((S^\wedge)^*)^\wedge = S^\wedge \text{ and } ((C^*)^\wedge)^* = C^* \quad (1.7)$$

Proof. Indeed, (1.3a) applied to (1.6a) leads to $(S^{\wedge*})^\wedge \subseteq S^\wedge$, while (1.6b) applied to S^\wedge gives $(S^\wedge) \subseteq (S^\wedge)^{\wedge*}$ \square

It is therefore possible to define the operation “ \wedge^* ” as a *closure operation* (Birkhoff, 1948), in that it is:

$$\text{extensive,} \quad S \subseteq S^{\wedge*} \quad (1.8a)$$

$$\text{idempotent} \quad (S^{\wedge*})^{\wedge*} = S^{\wedge*} \quad (1.8b)$$

$$\text{and increasing.} \quad S \subseteq S' \Rightarrow S^{\wedge*} \subseteq S'^{\wedge*} \quad (1.8c)$$

$S^{\wedge*}$ is called the closure of S . Extensivity means that the closure is never smaller, while idempotence implies that applying \wedge^* more than once does not change the closure. Finally, that \wedge^* is increasing corresponds to the idea that the closure of a larger set is larger.

Given two subsets $S \subseteq \mathbf{S}$ and $C \subseteq \mathbf{C}$, a couple (S, C) is said to be *closed* (or *complete*) if and only if $C = S^\wedge$ and $S = C^\star$. Yet such a closed couple is actually an epistemic group (S, S^\wedge) where $S^{\wedge\star} = S$. Closed couples correspond obviously to epistemic groups closed under $\wedge\star$, and therefore “ $\wedge\star$ ” is an operation yielding a set which cannot be enlarged further (extensivity and idempotence). It expands an epistemic group to its boundary: the largest possible set which is still based on a given agent set.²

Since the EC based on an agent set S is the largest agent set with the same intension as S , it becomes obvious that this largest set is the extension of the intension of S , or $S^{\wedge\star}$: applying $\wedge\star$ to S returns *all the agents* who use the same concepts that were common to the agents of S , hence the largest agent set — once and for all from (1.8b). Thus, the operator “ $\wedge\star$ ” yields the EC of any agent set, and according to definitions EC-1 and EC-2 we have:

Proposition 2. $(S^{\wedge\star}, S^\wedge)$ is the epistemic community based on S .

Proof. Indeed, (i) $S^{\wedge\star}$ has the same intension as S from $((S^\wedge)^\star)^\wedge = S^\wedge$ and (ii) it is the largest agent set enjoying this property: consider S' such that $S' \supset S^{\wedge\star}$ and $S'^\wedge = S^{\wedge\star\wedge}$, then $\forall \{s\} \subset S' \Rightarrow \{s\}^\wedge \supset S'^\wedge \Rightarrow \{s\}^\wedge \supset S^{\wedge\star\wedge} \Rightarrow \{s\}^{\wedge\star} \subset S^{\wedge\star}$, but $\{s\} \subset \{s\}^{\wedge\star} \Rightarrow \{s\} \subset S^{\wedge\star}$, hence $S' \subset S^{\wedge\star}$ \square

Subsequently,

Proposition 3. Any closed couple is an epistemic community.

Note that all these properties are similar and in fact dual if we consider an epistemic community based on C , subset of \mathbf{C} , and operators \star and $\star\wedge$. We may now define formally what an epistemic hypergraph is:

Definition 5 (Graph, hypergraph). A graph \mathbf{G} is a couple (V, E) where V is a set of vertices and $E \subset V \times V$ a set of edges binding pairs of vertices. A hypergraph \mathbf{hG} is a couple (V, hE) where V is a set of vertices and hE a set of hyperedges connecting set of vertices. hE is thus fundamentally a subset of $\mathcal{P}(V)$, the power set of V .

Definition 6 (Epistemic hypergraph). An epistemic hypergraph is a hypergraph of epistemic communities, $(\mathbf{S}, \{S^{\wedge\star} | S \subset \mathbf{S}\})$ with hyperedges binding groups of agents belonging to a same EC.

²Note that given $S^\wedge = \{c_1, \dots, c_n, c\}$ and $S'^\wedge = \{c_1, \dots, c_n, c'\}$, $c' \neq c$, we have $S' \notin S^{\wedge\star}$, S' is not in the closure of S . This might look strange for a human eye who would have said their domains of interest to be similar. S and S' anyway belong together to $(S \cup S')^{\wedge\star}$, or $\{c_1, \dots, c_n\}^\star$. Another property may help understand better what this closure actually corresponds to: given $S^\wedge = \{c_1, \dots, c_n\}$ and $S'^\wedge = \{c'_1, \dots, c'_n\}$ such that $\forall (i, j) \in \{1, \dots, n\}^2, c_i \neq c'_j$, we have $(S \cup S')^{\wedge\star} = \mathbf{S}$: the closure of two sets of scientists working on totally different issues is the whole community \mathbf{S} .

Each hyperedge can be labelled with the concept set corresponding to the agent set it binds, S^\wedge . For instance, $(\{s_2, s_3, s_4\}, NS)$ is an EC, so the hyperedge $\{s_2, s_3, s_4\}$ belongs to the epistemic hypergraph, and may be labelled “NS”. Note that equivalently an epistemic hypergraph could be based on concepts: $(\mathbf{C}, \{C^\wedge \mid C \subset \mathbf{C}\})$, with hyperedges binding concepts of a same EC.

Cultural background Interestingly, S^\wedge represents the concepts the whole community shares — as such, the “*cultural background*”. By contrast, C^\star contains authors who have used *every* word in the whole concept set \mathbf{C} — in the real world, it should be very rare to have $C^\star \neq \emptyset$.

Chapter 2

Building taxonomies

A relationship between the set of agents and the set of concepts is thus sufficient to capture the underlying epistemic hypergraph of a given scientific field. However, we still need to hierarchize the raw set of all ECs to build a taxonomy of the whole knowledge community, assuming that they are structured into fields and subfields. By introducing Galois lattices particularly appropriate for this purpose, we will represent ECs hierarchically. GLs are suitable for representing and ordering abstract categories relying on such a binary relation, and have been therefore widely used in conceptual knowledge systems, formal concept classification, as well as mathematical social science (Wille, 1982; Freeman & White, 1993; Godin *et al.*, 1995; Monjardet, 2003). More broadly, GLs can also be considered as hierarchically ordered epistemic hypergraphs — as such, GLs are both a categorization tool and a taxonomy building method.

2.1 Taxonomies and lattices

The canonical approach for representing and ordering categories consists of trees, which render Aristotelian taxonomies. In a tree, categories are nodes, and subcategories are child nodes of their unique parent category. A major drawback of such a taxonomy lies in its ability to deal with objects belonging to multiple categories. In this respect, the *platypus* is a famous example: it is a mammal and a bird at the same time. Within a tree, it has to be placed either under the branch “*mammal*,” or the branch “*bird*.” Another problem is that trees make the representation of paradigmatic categories extremely unpractical. Paradigmatic classes are categories based on exclusive (or orthogonal) rather than hierarchical features (Vogel, 1988): for instance *urban* vs. *rural*, *Italy* vs. *Germany*. In a tree, “*rural Italy*” has to be a subcategory of either *rural* or *Italy*, whereas there may well be no reason to assume an order on the hierarchy and a redundancy in the differentiation.

A straightforward way to improve the classical tree-based structure is a lattice-based structure, which allows *category overlap* representation. Technically, a lattice is a partially-ordered set such that given any two elements l_1 and l_2 , the set $\{l_1, l_2\}$ has a least upper bound (denoted by $l_1 \sqcup l_2$ and called “join”) and a greatest lower bound (denoted by $l_1 \sqcap l_2$ and called “meet”):

Definition 7 (Lattice). *A set $(L, \sqsubseteq, \sqcup, \sqcap)$ is a lattice if every finite subset $H \subseteq L$ has a least upper bound in L noted $\sqcup H$ and a greatest lower bound in L noted $\sqcap H$ under the partial-ordering relation \sqsubseteq .¹*

In a lattice, the *platypus* may simply be the sole member of the joint category “mammal-bird,” with the two parent categories “mammal” and “bird.” The “mammal-bird” category is “mammal” \sqcap “bird,” i.e. “mammal”-meet-“bird.” The parent category (“animal”) is “mammal” \sqcup “bird”, or “mammal”-join-“bird”. Besides, lattices may also contain different kinds of paradigmatic categories at the same level — see Fig. 2.1. Note that such an algebraic lattice is *not to be confused* with what the term “lattice” traditionally covers in physics: a mesh, a regular grid, a periodic configuration of points whose structure has nothing to do with our lattices.

2.2 Galois lattices

We hence argue that a lattice replaces efficiently and conveniently trees for describing taxonomies.² In order to create a lattice-based taxonomy of ECs, we first need to provide a *partial order* between ECs. Namely, we say that an EC is a *subfield* of a field if its intension is more precise than that of the field; in other words, if the concept set of the subfield contains that of the field. Formally, we define the strict partial order \sqsubset such that $(S, S^\wedge) \sqsubset (S', S'^\wedge)$ means that (S, S^\wedge) is a *subfield* of (S', S'^\wedge) , with:

$$(S, S^\wedge) \sqsubset (S', S'^\wedge) \Leftrightarrow S \subset S' \quad (2.1)$$

Hence (S, S^\wedge) can be seen as a specification of (S', S'^\wedge) , since its concept set is larger ($S^\wedge \supset S'^\wedge$) thus defining (S, S^\wedge) more precisely, while less agents belong to its extension ($S \subset S'$). Conversely, (S', S'^\wedge) is a “*superfield*” or a generalization of (S, S^\wedge) . We can thus render both generalization and specification of closed couples (Wille, 1992). For instance, if we consider (S, S^\wedge) as a school of thought, a subfield $(S', S'^\wedge) \sqsubset (S, S^\wedge)$ can be seen as a trend inside the school.

¹In this respect the power set of a set X provided with the usual inclusion, union and intersection, $(\mathcal{P}(X), \subseteq, \cup, \cap)$, is a lattice.

²We will not consider graded categories like fuzzy categories (Zadeh, 1965) and thick categories, such as locologies (De Glas, 1992).

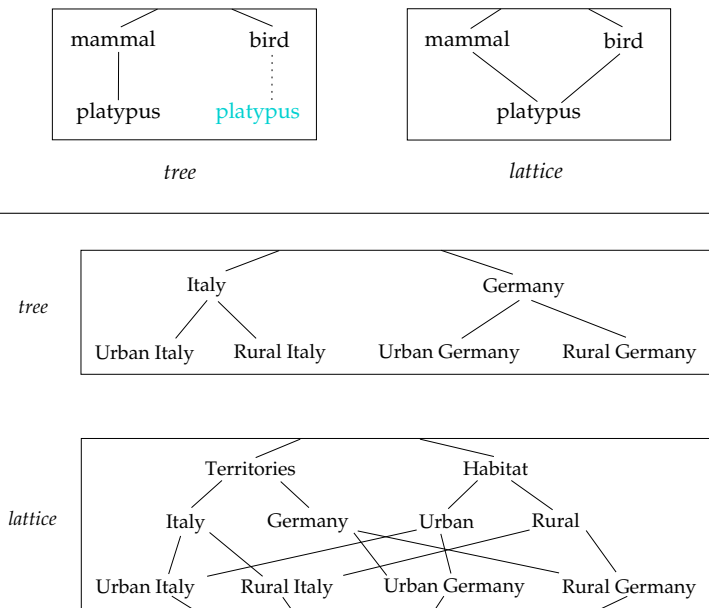


Figure 2.1: Trees *vs.* lattices. *Top:* Multiple categories: in a tree, the *platypus* needs either to be affiliated with *mammal* or *bird*, or to be duplicated in each category — in a lattice, this multiple ascendancy is effortless. *Bottom:* Paradigmatic taxonomies: in a tree, a paradigmatic distinction (e.g. territories *vs.* habitat types) must lead to two different levels and cannot be represented as a single category — in a lattice, the two paradigmatic notions may well be on the same level, leading to mixed sub-categories.

Now, using the natural partial order \sqsubseteq , gathering the set of ECs allows us to define a lattice that hierarchically orders all ECs. The *Galois lattice* (Birkhoff, 1948) is exactly the ordered set of all epistemic communities built from \mathbf{S} , \mathbf{C} and \mathcal{R} :

Definition 8 (Galois lattice). *Given a binary relation \mathcal{R} between two finite sets \mathbf{S} and \mathbf{C} , the Galois lattice $\mathcal{G}_{\mathbf{S},\mathbf{C},\mathcal{R}}$ is the set of every complete couple $(S, C) \subseteq \mathbf{S} \times \mathbf{C}$ under relation \mathcal{R} . Thus,*

$$\mathcal{G}_{\mathbf{S},\mathbf{C},\mathcal{R}} = \{(S^{\wedge*}, S^{\wedge}) \mid S \subseteq \mathbf{S}\} \quad (2.2)$$

Proposition 4. *($\mathcal{G}_{\mathbf{S},\mathbf{C},\mathcal{R}}, \sqsubseteq, \sqcup, \sqcap$) is a lattice, with \sqcup and \sqcap such that $\forall (S, C), (S', C') \in \mathcal{G}_{\mathbf{S},\mathbf{C},\mathcal{R}}$,*

$$\begin{cases} (S, C) \sqcup (S', C') = ((C \cap C')^*, C \cap C') \\ (S, C) \sqcap (S', C') = (S \cap S', (S \cap S')^{\wedge}) \end{cases}$$

Proof. Indeed, $((C \cap C')^*, C \cap C')$ is closed and belongs to $\mathcal{G}_{\mathbf{S},\mathbf{C},\mathcal{R}}$: $(C \cap C')^{\wedge*} = (S^{\wedge} \cap S'^{\wedge})^{\wedge*} = (S \cup S')^{\wedge*} = (S \cup S')^{\wedge} = C \cap C'$, from (1.4) & (1.7). Suppose now $(\sigma, \sigma^{\wedge})$ closed such that $S \subset \sigma$, $S' \subset \sigma$, so $(S \cup S') \subset \sigma$, $(S \cup S')^{\wedge*} \subset \sigma^{\wedge*} = \sigma$, i.e. $(C \cap C')^* \subset \sigma$, thus $(C \cap C')^*$ is the smallest closed σ such that $S \subset \sigma$ and $S' \subset \sigma$. The same goes for $(S \cap S', (S \cap S')^{\wedge})$. \square

A graphical representation³ of a GL is drawn on Fig. 2.2 from the sample community of Fig. 1.1: an EC closer to the top is more general: the hierarchy reproduces the generalization/specialization relationship induced by \sqsubseteq . It is straightforward to see that a GL can be seen as an epistemic hypergraph. Note that Galois lattices are also called “concept lattices” in other contexts (Wille, 1992; Stumme, 2002) — in other epistemic communities...⁴

2.3 GLs and categorization

Galois lattice theory offers a convenient way to group agents with respect to concepts they share, and as such it is yet another clustering method (CM). Nonetheless, if a GL contains all epistemic communities, ordered in a lattice-based taxonomy, we need to show why this tool is relevant as regards a *community description*

³We represent the GL using the Hasse diagram, which is a general method for rendering partially-ordered sets. In a Hasse diagram, an element is linked by a line to its covers (the smallest greater elements), and no element can be geometrically over another one if it is not greater (Davey & Priestley, 2002).

⁴Let us also mention Q -analysis (Atkin, 1974), whose principles strongly recall GLs. Again, given a relation \mathcal{R} between two sets, Q -analysis introduces polyhedra such that for each object s of the first set, the associated “polyhedron” is made of vertices c such that $s\mathcal{R}c$. The notion of “maximal hub / maximal star” replaces that of closed couple (Johnson, 1986). However, while Galois lattices focus on the hierarchy between closed couples, Q -analysis is more interested in connected paths between polyhedra, by making an extensive use of equivalence classes of Q -connected components. In particular, two polyhedra sharing at least $Q+1$ vertices are Q -near, and polyhedra between which there is a chain of Q -near polyhedra are said to be Q -connected.




Figure 2.2: Creating the Galois lattice corresponding to the sample community of Fig. 1.1. The GL contains 6 ECs. Solid lines indicate hierarchic relationships, from top (most general) to bottom (most specific); ECs are represented as a pair (extension, intension) = (S, C) with $S^\wedge = C$ and $C^* = S$.

task. Is a GL able to capture and reveal a *meaningful* structure of a given community? There are several stylized facts we would like GLs to rebuild, primarily the existence of subfields and *significant groups* of agents working within those subfields. Assuming a certain organization of scientific communities, the justification for this method will lie (i) in the fact that it partitions a field into smaller subfields corresponding to scientific communities, and (ii) in the *agreement* between epistemic communities rebuilt and extracted using GLs and those explicitly given by domain experts.

2.3.1 About relevant categorization

Let us first examine what clustering methods reveal about data: from any input set of objects provided with attributes, CMs are designed to produce an output, namely clusters of objects. CMs regroup the data even when the objects have no attribute in common, where any clustering would in fact be meaningless. In sorting objects from their size and value, clustering algorithms give results which are unlikely to represent, say, functional categories. To be relevant, CMs need to be guided by assumptions on the data structure: an obvious necessary assumption is that it does at least exhibit a clustered structure. It is necessary to inquire and specify what a given CM aims to rebuild: it would be unwise to trust its output without having checked its adequacy to data and defined what constitutes a cluster or a community. Both the choice of the CM and the choice of attributes (labelling of data) are decisive.⁵

The same holds for Galois lattices: one can draw a GL from any two sets of objects and a given relationship between them, but there is no reason *a priori* why the lattice should reveal a remarkable structure, even if it is built, represented or managed efficiently. There should exist a lot of data for which this categorization is just irrelevant. In order to know whether and why GL is an appropriate CM for producing a taxonomy of knowledge communities, it is necessary to investigate the nature and organization of these communities.

⁵One might thus distinguish (i) labelling irrelevant for the kind of data studied, while using a relevant CM; from (ii) CM irrelevant for the kind of data studied, however labelled relevantly. Take for instance a linguist who would like to group the words *light*, *dark*, *holy* and *evil* as regards their semantic field. He might consider two criteria: *brightness* and *goodness*, and select e.g. the following numerical representations: *light*: +5 (brightness), +1 (goodness); *dark*: -5, -1; *holy*: +1, +5; *evil*: -1, -5. For sure an irrelevant labelling, i.e. a bad choice in the previous criteria (say, choosing the number of vowels and the number of consonants) would obviously give him a meaningless result. But an irrelevant clustering method, e.g. based on Euclidian distances, would also give him inconsistent output in grouping *light* with *holy*, and *dark* with *evil*, while he wanted *light* with *dark*, and *holy* with *evil*.

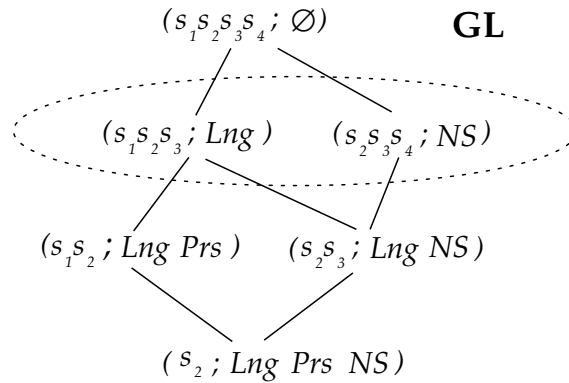


Figure 2.3: Galois lattice of the sample community (hierarchical structure drawn in solid lines relatively to \sqsubseteq , i.e. “bottom” \sqsubseteq “top”). The medium level (dashed ellipse) contains closed couples $(\{s_1, s_2, s_3\}; \{Lng\})$ and $(\{s_2, s_3, s_4\}; \{NS\})$ obviously corresponding to major fields (linguistics and neuroscience). Hierarchy yields just below interesting subcommunities like $(\{s_1, s_2\}; \{Lng, Prs\})$ or $(\{s_2, s_3\}; \{Lng, NS\})$, possibly prototypical of more specific subfields.

2.3.2 Assumptions on EC structure

Our main assumption is that there are fields of knowledge which can be described by concept lists (relevant labelling), and which are being implemented by sets of agents. Taking again the first example, some people are obviously linguists: among them, some deal with a given aspect, say prosody; some other scientists deal with neuroscience, while a few of them are interdisciplinary and use both concepts. Knowledge fields and their corresponding agent sets are epistemic communities, which are precisely what GLs consist of (see Prop. 3). Moreover and also crucial, these fields are hierarchically organized: (i) a general field can be divided into many subfields, themselves possibly having subcategories or belonging to various general fields, and (ii) some fields can be *multi-disciplinary* or *interdisciplinary* in that they respectively involve or integrate two or more subfields (Klein, 1990). For instance, cognitive science is a general field gathering various subfields such as cognitive linguistics and cognitive neuroscience, thus being multidisciplinary. But the subfield “cognitive neurolinguistics” is interdisciplinary because it mixes both parent disciplines.

GL relevance as regards these properties results from its natural partial order \sqsubseteq , which reflects a generalization/specialization relationship between fields and subfields as discussed previously (see also Fig. 2.3), as well as multidisciplinary and interdisciplinary through particular patterns called *diamonds* (see Fig. 2.4).




Figure 2.4: Zoom on Fig. 2.3 showing one possible *diamond*. A multidisciplinary field is at the top of the diamond (here “ \emptyset ”, which can be considered as “cognitive science”) and covers “cognitive linguistics” and “cognitive neuroscience”, which themselves, when combined, define an interdisciplinary subfield, “cognitive neurolinguistics”.

2.3.3 GLs and selective categorization

Thus, GLs are a relevant tool for building taxonomic lattices from simply \mathcal{R} , \mathbf{S} and \mathbf{C} . More generally, it is worth noting that we can replace *authors* with *objects*, and *concepts* with *properties*. This yields a generic method for producing a comprehensive taxonomy of any field where categories can be described as a set of items sharing equivalently some property set. This has been indeed a useful application of GLs in artificial intelligence (as “Formal Concept Analysis”) (Wille, 1982; Ganter, 1984; Wille, 1997; Godin *et al.*, 1998), and has been investigated as well in mathematical sociology recently (Wasserman & Faust, 1994; Batagelj *et al.*, 2004), as well as mathematical social science in general (Freeman & White, 1993; Monjardet, 2003; Duquenne *et al.*, 2003).

However, a serious caveat of GLs is that they may grow extremely large and therefore become very unwieldy. Even for a small number of agents and concepts, GLs contain often significantly more than several thousands of ECs. Thus, it is still unclear why a GL would produce a *useful* and *usable* categorization of the community under study. Indeed, by definition a GL contains all epistemic communities. This property is already restrictive: sets of agents or sets of concepts which have nothing or nobody in common (i.e. their intension or extension is \emptyset) or more generally which are not “closed”, are not epistemic communities and hence do not appear in the GL. Yet $\mathcal{G}_{\mathbf{S}, \mathbf{C}, \mathcal{R}}$ contains all ECs: this includes naturally most singletons ($s^{\wedge*}$, s^{\wedge}) as well as $(\mathbf{S}, \mathbf{S}^{\wedge})$, but also and especially all the intermediary ECs. Among those, many do not correspond to an existing or relevant field of knowledge, because they are too small or too specific. For a single scientist $\{s\}$, the

closure $\{s\}^{\wedge*}$ will admittedly be equal to $\{s\}$, because no other scientist than s is likely to use every concept in $\{s\}^{\wedge}$ (there are strong chances that $\forall s' \in \mathbf{S}, \exists w \in s^{\wedge}$ and $\notin s'^{\wedge}$). Agent s is “original”.

Consider the agents working on an actual knowledge field F (e.g. a real discipline). If we consider only a few of these agents, there is a strong chance that they share some original concepts other than those of F . These few agents S will thus constitute a small EC, $(S^{\wedge*}, S^{\wedge} \supsetneq F)$. However, the more agents working on F in S , the less likely they are to share concepts other than those of F , and the more likely the decreasing intension S^{\wedge} reaches F . For any agent set S whose intension S^{\wedge} reaches F , the corresponding epistemic community $S^{\wedge*}$ is the whole community working on F . This induces a gap between (i) small ECs using F plus some additional original concepts, and (ii) the suddenly emerging EC $(S^{\wedge*}, S^{\wedge} = F)$ — “emerging” because it suddenly gathers many more agents than S . We conjecture that there is a relevant level for which closed sets $S^{\wedge*}$, and identically $C^{*\wedge}$, are representative of a field or a trend. This also means that some epistemic communities listed by GLs are deemed to be prototypical of these fields. They are located between the whole agent set, too general, and too specific communities, that is, at a medium level of size and generality which is to be compared to the *basic-level* of categorization introduced by Rosch and Lloyd (1978).⁶ This medium level shall constitute our basic-level of epistemic categorization, in such a way that the field would be too general above it (“superordinate categories”), and too precise under it (“subordinate categories”).

Given these assumptions, $\mathcal{G}_{S,C,R}$ is expected to exhibit significant structural properties which could help design criteria for detecting major trends (basic-level categories) within a more general field, in a somewhat automated manner. In particular, in the light of the present remarks populated ECs should be remarkable ECs. We will bring empirical evidence to support this conjecture in Chap. 3. More broadly, our objective is to use GLs in order to extract a *significant epistemic hypergraph of relevant ECs*, which is *in fine* a taxonomy matching empirical expert-based descriptions of the community structure.

2.4 Comparison with different approaches

Community and group detection have been investigated in both computer science (graph theory as well as artificial intelligence) and sociology. Clustering methods originating from computer science rely on graph theory and then on algorithms

⁶Basic levels obey in particular to two principles (Barthélemy *et al.*, 1996): (i) a principle of minimal cognitive cost (which suggests for instance to look at largest communities), and (ii) a principle of reality (which requires to check that reality fits the assumptions on category structure).

that partition graphs in a number of clusters, fixed *a priori* or not (such as spectral bisection or Kernighan-Lin algorithm (Newman, 2004)), or on object properties viewed as a multi-dimensional vector, where objects are grouped according to their relative similarity (such as *k-means* (Hartigan, 1975), probabilistic neural networks (Specht, 1990), Kohonen maps (Kohonen, 2000)), similarity measures being mostly based on Euclidian distance. The main drawback of these methods is their relevance for social science: they eventually infer communities with no particular assumption on the nature of the social groups that these CMs are supposed to extract from data. Thus, produced clusters have an unclear connection with what social scientists would call communities.

Sociologists by contrast introduce hypotheses and tools proper to social networks — such as cohesion and strong ties (Burt, 1978; Wellman *et al.*, 1988), centrality (Freeman, 1977; Friedkin, 1991) or structural equivalence (Lorrain & White, 1971) — which yield CMs more adequate to social group detection than generic computer science methods, including for instance hierarchical clustering (Johnson, 1967), structural balance (Doreian & Mrvar, 1996), blockmodeling (Batagelj *et al.*, 1999) or, more recently, structural cohesion and *k-components* (Moody & White, 2003), and the Girvan-Newman algorithm (Girvan & Newman, 2002) and its improvement by Radicchi *et al.* (2004).

In addition, most of these methods produce hierarchically structured clusters which are in fact more or less *dendrograms*. Yet a dendrogram is a cluster tree, and ascendancies cannot be multiple: a community is bound to be embedded into a lineage of increasing communities. It cannot have ascendancies in various “directions,” and an agent cannot be part of many non-embedded, overlapping communities.

In any case, methods relying only on single networks of social relationships (e.g. co-authorship) may prove to be insufficient and inefficient in order to find epistemic communities which, as we said before, are not necessarily socially linked. One-mode data (or projection of two-mode data onto one-mode data) also entails a loss of crucial *structural* information (see Fig. 2.5). Consider for instance a one-mode concept network where links arise between two concepts whenever they share some authors: there would be no way, here, to distinguish a triangle of concepts sharing the same set of authors, from a triangle of concepts linked through pairs of totally different author sets; this distinction is however central in our case. Data duality brought by the reciprocal linkage of agents to concepts and the corresponding symmetry between agent-based and concept-based notions (definitions 1, 2, 3 and EC-2, and definition 4) is moreover well rendered by a GL, being a hierarchy of closed couples considered equivalently as agent sets or as concept sets.




Figure 2.5: Two significantly different two-mode datasets (*left*) yield an identical one-mode projection (*right*), when linking pairs of agents sharing at least one concept. s_1, s_2, s_3 are agents, c, c_1, c_2, c_3 are concepts.

Chapter 3

Empirical results

In this chapter, (i) we present a first experimental protocol, enabling us to create a static taxonomy from bibliographic data, and (ii) we validate a basic stylized fact, the presence of ECs having a large agent set — *a feature which cannot be explained only by the popularity of some concepts*, as we will show.

3.1 Experimental protocol

To conduct our experiments on scientific communities, we need data stipulating which agents use which concepts. We consider article collections, assuming that articles are a faithful account of what their authors are working on. However, an important point is to define what a *concept* is, such that it appears in an article. Is it a paradigm such as “*universal gravitation*” or a simple word like “*operon*”? For instance, authors provide their articles with keywords: considering these keywords as concepts might constitute a relevant level of categorization while being a convenient idea. Yet, keywords are poor indicators, for authors often omit important keywords. Depending on the database, keywords for a same article may differ.

Word groups as concepts Getting concepts through words and nominal groups (terms) from the title, abstract or body is safer. At first we considered that *each word or nominal group is a concept*, even if we were still hampered by linguistic phenomena such as homonymy, polysemia, synonymy (Jackendoff, 2002), syllepsis (Jacquelinet *et al.*, 2000), and the fact that different authors may have different definitions of the same word or understand different concepts under an identical nominal group (Lavie, 2003). Some techniques (Wang *et al.*, 2000) could be used to determine the contextual meaning of nominal groups, but we assumed that nominal groups represent sufficiently distinguishable and homogenous references to concepts — we also ignored the fact that their meaning possibly evolves with time

(Leydesdorff, 1997). This definition does not prevent us from observing higher-level concepts such as theories or even paradigms, because we can refer to these concepts *a posteriori* by considering sets of words, for example interpreting {"cell," "DNA," "gene," "genetics," "molecular"} as *molecular biology*.

We proceeded with title and abstract words only, because complete article contents are seldom available. While apparently rough, these minimal assumptions yielded significant results anyway.

Data processing We treated the data according to the following methodology:

1. Collect and automatically process article data (title, abstract, authors) for a given community and period of time. As regards abstract and title, we apply a basic linguistic processing consisting in:
 - Excluding insignificant words (*stop-words*), such as common and rhetorical English words ("often," "then," "we," etc.) and irrelevant words with respect to the domain ("demonstrate," "postulate," "specimen," "study," etc.), using a list of more than 2,500 words, to which we add non-words such as figures, percentages, dates, etc.
 - Excluding rare words, i.e. words appearing n times or less in the whole corpus (such as words appearing only once, also called *hapax legomena* or *hapaxes*). We took $n = 4$.
 - Stemming the remaining words, i.e. reducing morphological variants of words to their stem (root form) using a slightly improved version of Porter's stemming algorithm (Porter, 1980), and then creating the corresponding word classes (for example, "genetic" and "genetics" both reduce to "genet").
2. Identify unique authors and unique words, and then create the weighted matrix R of links between authors and words, where R_{ij} is equal to the number of articles where author i used concept j (see Fig.3.1).
3. Consider a representative sample of the whole community by extracting randomly and uniformly some lines from matrix R . We chose to keep each line with probability .25 (this step aims at reducing GL computation cost by a factor 40).
4. Make R a binary matrix with respect to a given threshold α , i.e. replace R_{ij} by 1 if $R_{ij} > \alpha$, otherwise by 0: this means that an author will not be related to a concept he used less than α times. We used a threshold of 0. Increasing the threshold would critically reduce both computation costs and results significance.




Figure 3.1: Experimental protocol: step 1 and 2 help create the core network, and the corresponding relationship weighted matrix shown here (authors on rows, concepts on columns). Some agents are removed through step 3 (hence some little used concepts disappear). The GL is then computed from the binary relation matrix obtained after step 4.

5. Calculate the Galois lattice for the binary relation \mathcal{R} built upon matrix R , using an implementation of Ganter’s algorithm (Ganter, 1984; Lindig, 1998).

3.2 Results and comparison with random relations

We ran the process on articles published between 1990 and 1995 obtained through a search for “zebrafish” in publicly available bibliographic data from the MedLine database, totalizing 418 articles, 797 authors and 2129 words after step 2 of the protocol.¹ After step 3, only 218 authors and 1817 concepts remained in R . This is the matrix we used for computing the GL (steps 4 and 5).

¹This community was chosen in part because we are sure that scientists working on the zebrafish explicitly mention the name of the animal, at least in the abstract. This would be less certain if we were looking for scientists working on molecular biology, or quantum mechanics for instance. Of course, restricting the data to articles present in MedLine could induce a bias, yet this database is also one of the most comprehensive for the field.

Some authors and concepts appeared more frequently than others. There is a characteristic distribution of links from agents to concepts and from concepts to agents: a lot of agents (resp. concepts) are linked to few concepts (resp. agents), a small number of agents are related to many concepts, few concepts are related to many agents. We could fear GL artefacts because frequent authors or frequent concepts are more likely to share or be shared by more concepts or agents. Being part of bigger closed sets and increasing the number of these big sets, they modify the GL structure, especially high-size closed sets. We could compare our results with those from GLs calculated with random-generated relationships where this exact property of the empirical data was kept. We kept the distributions of links on rows and columns in the relationship matrix from step 3 while we reshuffled the links themselves, using an algorithm introduced by Molloy and Reed (1995). This algorithm consists in assigning a number of outgoing links to concepts to each author, according to the desired distribution, and identically assigning a number of outgoing links to authors to each concept; then matching randomly the dangling links between authors and concepts. We call “*random case*” the results obtained from computations on 40 such randomly rewired relationship matrices. We also considered two other random cases: (i) keep the same density in the relationship (same proportion of real links in respect of possible links), which is approximately one link out of 30; and (ii) keep only the distribution of links from agents to concepts. Interestingly, the corresponding GLs are dramatically small, with 16,000 epistemic communities whose sizes do not exceed 5% of the whole community (see Fig. 3.2). Therefore, these cases were not investigated further.

3.2.1 Empirical versus random

Fig. 3.2 represents the total number of epistemic communities versus the size of their agent set. The empirical GL contains 214,000 closed couples, with communities ranging from 1 to 196 agents, except the epistemic community (S, \emptyset) containing all of the 218 agents under study. The random case contains an average of around 207,000 closed couples in the random case (standard deviation $\sigma \simeq 64,700$), with agent set sizes ranging only from 1 to 60 ($\sigma \simeq 5$). While the empirical GL is approximately of the same size as random GLs, it contains more high-size epistemic communities (371 communities representing more than a fifth of the whole agent set, against a dozen communities for the random case). There is a quite perfect fit on low-size closed couples, yet the empirical GL is denser on high-size couples. Cumulated densities, the proportions of closed couples containing at least a given number of agents, are shown on Fig. 3.3: 1% of the GL in the empirical case is made of epistemic communities containing 30 agents or more, against 0.05% in the random case. This proportion is one thousandth against one thirty-thousandth for




Figure 3.2: Raw distributions of agent set sizes.

communities with 50 agents or over. In the empirical case, we thus have a strongly significant discrepancy of *at least one order of magnitude more* populated ECs with more than 10% of the whole agent set.

3.2.2 Rebuilding the structure

The presence of large groups of structurally equivalent agents pointing to the same groups of concepts supports therefore the conjecture outlined in section 2.3: high-size epistemic communities are thus a remarkable stylized fact of our empirical data. It is also of interest to know whether these communities are significant and relevant, and if they help partition a field into smaller subfields corresponding to real epistemic communities.

Our zebrafish expert, Nadine Peyri eras, showed that it was the case:

- (i) The first and biggest community is unsurprisingly centered around the word "zebrafish" and contains 196 agents (90% of the whole). The fact that it does




Figure 3.3: Cumulated densities of agent set sizes.

not reach 100% of the community reflects the imperfection of the empirical data collection and processing.

- (ii) Then, a lot of large epistemic communities use a small set of words, namely “gene,” “expression,” “pattern,” “embryo,” “develop” and “vertebrate.” A majority of the 218 agents are present in at least one of these communities. This word set seems accordingly to characterize the core paradigm of zebrafish researchers, even if each agent does not use it entirely. According to our expert and to Grunwald and Eisen (2002), the zebrafish is used as a vertebrate animal model for the study of gene expression and function during embryonic development.

Similarly, another word subset of interest is made of “cloning,” “stage,” “transcription,” “sequence,” “protein,” “region,” “encode,” which constitute the intentions of large epistemic communities (50 agents). According to our expert, these words are proper to molecular biology or developmental studies, including zebrafish study, which consists in isolating the mutated genes from a large number of mutant fish lines then in investigating their effect on biological processes.

- (iii) Thereafter, two major groups emerge: (i) one with the epistemic community

based on “*growth*” (39 agents), and (ii) the other around three epistemic communities whose intensions are “*neuron*” (70 agents), “*brain*” (36 agents) and {“*nervous*”, “*system*”} (28 agents), with many agents in common and which altogether makes a group of 84 single agents. With only 15 agents in common, communities (i) and (ii) represent two distinct groups totalizing 108 agents. These groups correspond exactly to what the literature describes as significant subfields.²

Smaller communities help structure the field: the epistemic community based on {“*toxicity*”} is made of 23 agents with 9 shared with “*growth*” and only 3 with “*brain*”. This latter group might be related to the study of the toxic effect of growth factors. The epistemic community based on words “*acid*” (45 agents) has an interesting descent, {“*acid*”, “*amino*”} (22 agents) and {“*acid*”, “*retino*”} (21 agents), with only 3 agents in common in the extension of {“*acid*”, “*amino*”, “*retino*”}, so this is a diamond with no relationship between people working on amino acid and retinoic acid. Also, the closed couple with intension {“*spinal*”, “*cord*”} (28 agents) includes the one based on {“*spinal*”, “*cord*”, “*neural*”, “*ventral*”} (20 agents) with almost as many agents, suggesting that (i) “*spinal*” and “*cord*” cannot be dissociated and (ii) people working on spinal cord are also very familiar with concepts “*neural*” and “*ventral*”.

These findings summed up on Fig. 3.4 show that GLs are efficient both for determining the community paradigm (or common background) and for finding prevailing communities as well as basic-level subcommunities. This first partition is made from data of the period 1990-1995 and is supposed to be a *static* picture of the community structure in December 1995. Methods for studying the community evolution through the dynamics of the GL will be described in section 5.

These results also show the usefulness of binding agents to concepts networks and taking into account data of both types, since detected communities here are not necessarily socially grounded: agents who belong to the same EC are likely for example to have never collaborated. It would have been certainly uneasy, if not impossible, to detect them with single-network based methods. Moreover, distributions of links between agents and concepts do not account alone for the particular clustered structure of ECs. There is more structure in the empirical network than distributions of links would suggest.

²At the beginning of the 90’s, according to Grunwald and Eisen (2002), “among the first mutants to be isolated was one that was later discovered to be deficient in a *growth factor* needed for axis determination, a second deficient in myofibril organization, and a third in which a specific portion of its *nervous system* failed to form”.

According to the program of the first conference on zebrafish development and genetics at the CSH Laboratory in 1994, there were seven theme-based sessions, including two on nervous system and one on growth control. Approximately, these two fields represented half the sessions and half the community.

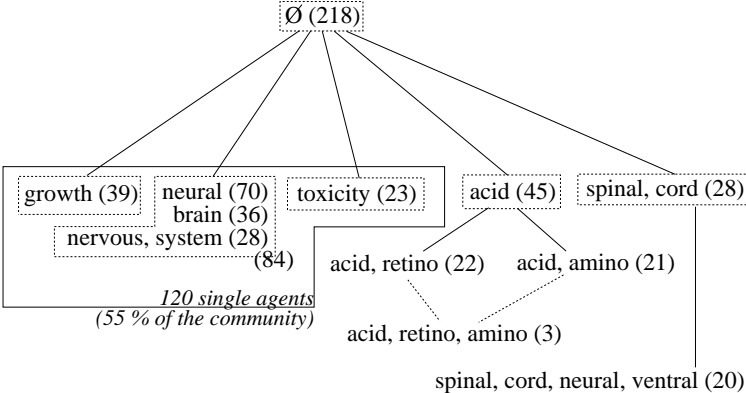


Figure 3.4: Partial view of the actual GL, which contains more than 200,000 closed couples. It shows intension and extension sizes in brackets of selected epistemic communities. There are various possible partitions of the whole agent set, depending on what one is looking at: objects, processes, methods. Note that on this figure we ignored communities containing paradigmatic words (“develop,” “gene,” etc.), thus focusing on more discriminating ECs.

Chapter 4

Community selection

So far, from a low-level L made of a relation \mathcal{R} between agents and concepts, Galois lattices helped us define a projection $P(L)$ which matches two high-level phenomena: (i) the presence of ECs gathering many agents, and (ii) an expert-based description of the community. Now, we would like to improve taxonomies produced by GLs, so that we are also able to provide an history of the field that matches an expert-based history.

To this end, a critical issue relates to the design of better criteria for distinguishing *basic-level* epistemic communities: what makes an epistemic community be a “basic-level” community? Which ECs should we extract from the GL to build a reduced and meaningful hypergraph of ECs? The property of gathering an important proportion of agents is a good yet insufficient first estimate. This quite simple criterion bears some major drawbacks, such as the fact that small communities are ignored, even if they correspond to well-defined but isolated fields. In this respect taking communities close to the top is more relevant.¹ These communities are indeed just more specific than the whole community. Hence, a more detailed set of selection properties may include distance from the top epistemic community, distance from the empty epistemic community (\emptyset, C) , and concept set size. In this section we explore the reduction of the GL to a *manageable* taxonomy.

4.1 Rationale

As we previously noticed GLs are usually very large, thus, considering only useful and meaningful patterns instead of manipulating whole lattices becomes crucial (in particular in an epistemological thus dynamic perspective, it would be signif-

¹In other words, those belonging to the *maximal antichain*, which is the subset of the ECs of $\mathcal{G}_{\mathcal{S}, \mathcal{C}, \mathcal{R}}$ which are not comparable one to each other, and which are maximal (each one of them is not included in any other EC).




Figure 4.1: From the original GL to a selected poset, or partial epistemic hypergraph.

icantly harder to track a series of GLs than just examining a static lattice). This means selecting from a possibly huge GL which ECs are relevant to taxonomy rebuilding, and excluding a large number of irrelevant ECs that could blur the picture of the community. In other words, we consider a partial, manageable view of the whole GL which we choose in order to reflect the most significant part and patterns of the taxonomy. Formally, the partial view is not anymore a lattice as defined previously: it is a partially-ordered set, or *poset*; nonetheless it overlays on the lattice structure and still enjoys the taxonomical properties we are interested in (see Fig. 4.1). For the sake of clarity, we will name “*partial epistemic hypergraph*” such a poset.

Selection preferences This selection process has so far been an underestimated topic in the study of GLs, with an important part of the effort focused on GL computation and representation (Dicky *et al.*, 1995; Godin *et al.*, 1998; Ferré & Ridoux, 2000; Kuznetsov & Obiedkov, 2002). Nevertheless, some authors insist on the need for semantic interpretations and approximation theories in order to cope with GL combinatorial complexity (Van Der Merwe & Kourie, 2002; Duquenne *et al.*, 2003). In our case, we need to specify selection preferences, i.e. which kind of ECs are relevant for a concise taxonomy description.

At first, we would certainly focus on the largest ECs while ignoring either too small or too specific closed sets, as we did so far: if a set of properties, attributes or concepts corresponds to a field, one can expect that the corresponding extension is

of a significant size. Since fields tend to be made of large groups of agents, and also because a GL mostly consists of small communities, size proved to be a segregating and efficient criterion, categorizing a large portion of the whole community — however still an insufficient criterion. Indeed, using only this criterion may be over-selective or under-selective, notably in the following cases:

- *Small yet significant sets.* One should not pay attention to very small closed sets, for instance those of size one or two: in general they cannot be considered representative of any particular EC. There is thus a pertinent threshold for the size criterion. However, this may still exclude some small ECs that could actually be relevant, notably those prototypical of a minority community. If so, some other criteria might apply as well:
 - (i) such ECs indeed, while being small, are unlikely to be subsets of other ECs and are more likely to be located in the surroundings of the lattice top;
 - (ii) alternatively, they may be unusually specific with respect to their position in the lattice;
 - (iii) finally, being outside the mainstream may make them less likely to mix with other ECs, thus having fewer descendants.
- *Large yet less significant sets.* Large contingent ECs may augment the GL uselessly. This is the case:
 - (i) when two ECs are large: it is likely that their intersection exists and has fortuitously a significant size — we could discriminate ECs whose size is not significant enough with respect to their smallest ascendant.
 - (ii) when empirical data fails to mention that some agents are linked to some properties: two or more very similar ECs appear where only one exists in the real world² — we could avoid this duplicity by excluding ECs whose size is too close to that of their smallest ascendant.

4.2 Selection methodology

Extending preferences and criteria Hence, agent set size does not matter alone and selection preferences cannot be based on size only. For instance, small ECs distant from the top are likely to be irrelevant, and certainly the most uninteresting ECs are the both smaller and less generic ones. To keep small meaningful ECs and

²Indeed, let s_1, s_2, s_3, s_4 and s_5 work on c_1, c_2, c_3, c_4 and c_5 , in reality. Suppose now that some data for s_5 is missing and that we are ignorant of the fact that s_5 works on c_5 . Then there will be two distinct communities: $(\{s_1, s_2, s_3, s_4\}, \{c_1, c_2, c_3, c_4, c_5\})$ and $(\{s_1, s_2, s_3, s_4, s_5\}, \{c_1, c_2, c_3, c_4\})$, which cover a single real EC.

to exclude large insignificant ones, some more criteria are required to design the above preferences. For a given epistemic community (S, C) , we may propose the following criteria:

1. size (agent set size), $|S|$;
2. level (shortest distance to the top³), d ;
3. specificity (concept set size), $|C|$;
4. sub-communities (number of descendants), n_d ;
5. contingency / relative size (ratio between the agent set and its smallest ascendant), λ .

Selection heuristics Then, we design several simple selection heuristics adequately rendering selection preferences. Selection heuristics are functions attributing a score to each EC by combining these criteria, so that we only keep the top scoring ECs. We may not necessarily be able to express all preferences through a unique heuristic. Therefore, the selection process involves several heuristics: for instance one function could select large communities, while another is best suited for minority communities. We ultimately keep the best nodes selected by each heuristic (e.g. the 20 top scoring ones).

Notice that agent set size $|S|$ remains a major criterion and should take part in every heuristic. Indeed, a heuristic that does not take size into account could assign the same score, for example, to a very small EC with few descendants (like those at the lattice bottom) and to a larger EC with as many few descendants (possibly a worthy heterodox community). In other words, given an identical size, heuristics will favor ECs closer to the top, having less descendants, etc. In general we need heuristics that keep the significant upper part of the lattice. Hence distance to the top d is important as well and should be used in many heuristics.

While we can possibly think of many more criteria and heuristics, we must yet make a selection among the possible selection heuristics, and pick out some of the most convenient and relevant ones. In this respect, the following heuristics are a *possible* choice:

1. $|S|$: select large ECs,
2. $\frac{|S|}{d}$: select large ECs close to the top,

³We take here the shortest length of all paths leading to the top EC (S, \emptyset) (the whole community). Indeed, paths from a node to the top are not unique in a lattice; we could also have chosen, for instance, the average lengths of all paths.

3. $|S| \frac{|C|}{d}$: select large ECs unusually specific,
4. $\frac{|S|}{dn_d}$: select large ECs close to the top and having few descendants,
5. $\frac{|S|}{d}(\lambda - \lambda^+)(\lambda^- - \lambda)$: select large non-contingent ECs close to the top.⁴

Fine tuning these heuristics eventually requires an active feedback from empirical data. For instance, one could prefer to consider only the first heuristics, and accordingly to focus on taxonomies including only large, populated, dominant ECs. Exploring further the adequacy and optimality of the choice and design of these heuristics would also be an interesting task — heuristics yielding e.g. a maximum number of agents for a minimal number of ECs — however unfortunately far beyond reach in the present effort. We will thus authoritatively keep and combine these few heuristics to build the partial epistemic hypergraph from the original GL, as shown on Fig. 4.1. In any case, correct empirical results with respect to the rebuilding task will acknowledge the validity of this choice.

⁴That is, of a moderate size relatively to their parents: $\lambda \in [\lambda^-; \lambda^+]$ — we could thus expect to exclude fortuitous EC intersections when $\lambda < \lambda^-$, and duplicate ECs when $\lambda > \lambda^+$.

Chapter 5

Taxonomy evolution

To monitor taxonomy evolution we monitor partial epistemic hypergraph evolution. To this end, we create a series of partial epistemic hypergraphs from GLs corresponding to each period, and we capture some *patterns* reflecting epistemic evolution by comparing successive static pictures. In other words, we proceed to a longitudinal study of this series.

Interesting patterns include in particular:

- *progress or decline of a field*: a burst or a lack of interest in a given field;
- *enrichment or impoverishment of a field*: the reduction or the extension of the set of concepts related to a field;
- *reunion or scission of fields*: the merging of several existing fields into a more specific subfield or the scission of various fields previously mixed.

In terms of changes between successive partial epistemic hypergraphs, the first pattern simply translates into a variation in the population of a given EC: the agent set size increases or decreases.

The second pattern reduces in fact to the same phenomenon. Indeed, suppose “*linguistics*” is enriched by “*prosody*”, i.e. $\{Lng\}$ is enriched by $\{Prs\}$, thus becoming $\{Lng, Prs\}$. This means that the population of $\{Lng, Prs\}$ is increasing. Since this EC is still a subfield of $\{Lng\}$, the enrichment of $\{Lng\}$ by $\{Prs\}$ translates into an increase of its subfield. Similarly, the decrease of $\{Lng, Prs\}$ would indicate an impoverishment of the superfield $\{Lng\}$.¹

¹More formally, say a field (S, C_1) is enriched by a concept c , becoming $(S', C_1 \cup c)$. This means that the subfield $(S', C_1 \cup c)$ is increasing — as it is a subfield of (S, C_1) , it is a subfield increase. In the limit case, when all agents working on C_1 are also working on c , the superfield (S, C_1) becomes exactly $(S, C_1 \cup c)$. In all other cases, it is $(S', C_1 \cup c)$, a strictly smaller subfield of (S, C_1) , with $S' \subset S$. Conversely, if a field $(S', C_1 \cup c)$ is to lose a specific concept c , the subcategory $(S', C_1 \cup c)$ is going to decrease relatively to (S, C_1) .

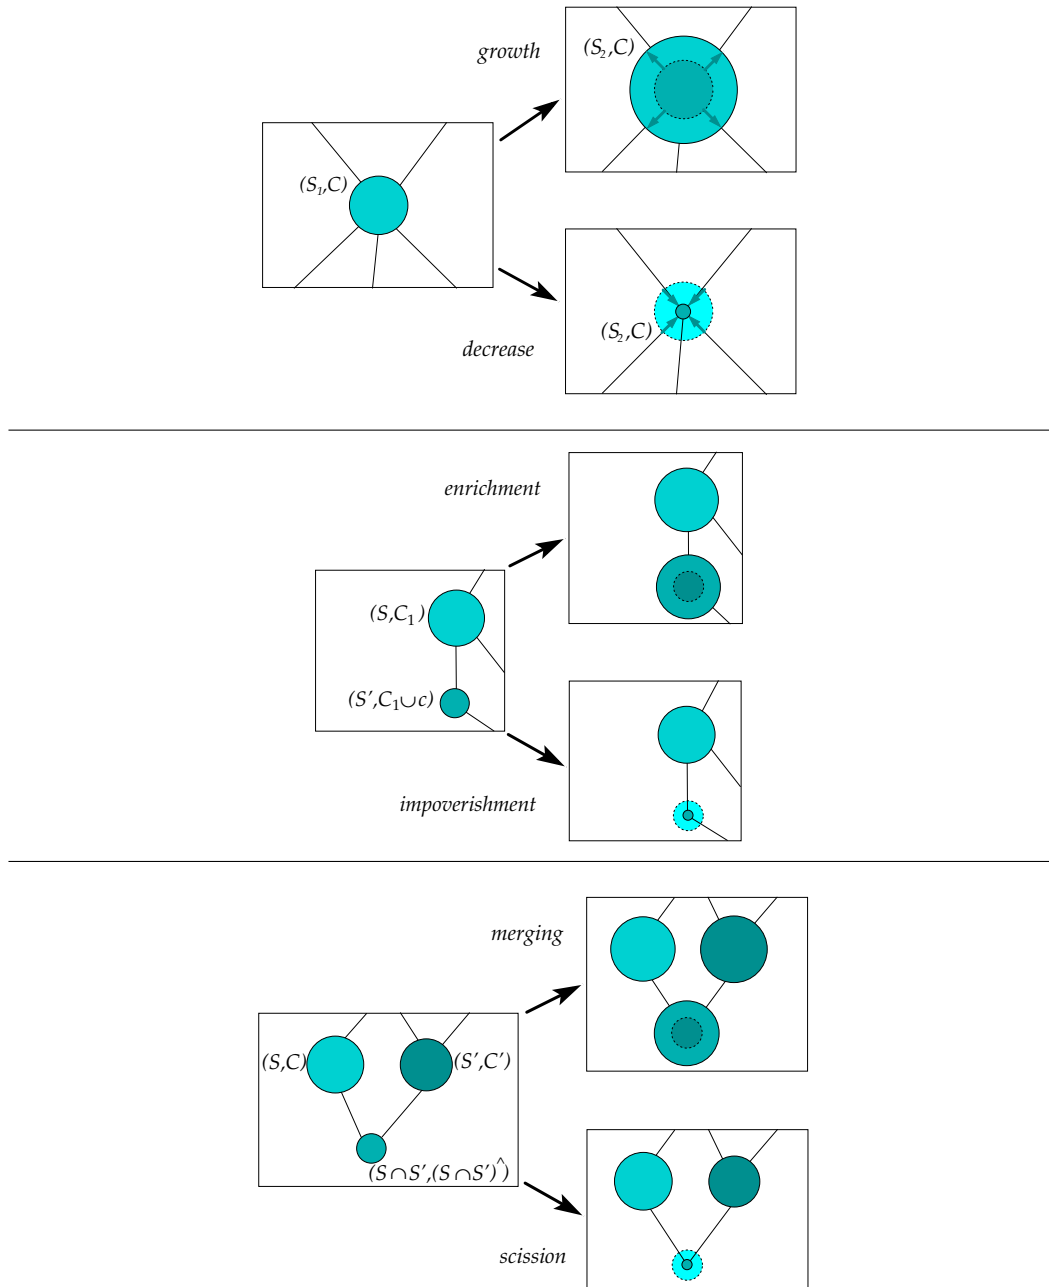


Figure 5.1: *Top*: progress or decline of a given EC (S_1, C) , whose agent set is growing (above) or decreasing (below) to S_2 . *Middle*: enrichment or impoverishment of (S, C_1) by a concept c , through a population change of the subfield $(S', C_1 \cup c)$. *Bottom*: emergence or disappearance of a joint community (diamond bottom) based on two more general ECs, (S, C) and (S', C') . Disk sizes represent agent set sizes.

Finally, the union of various fields into an interdisciplinary subfield as well as the scission of this interdisciplinary field comes in fact to an increase or a decrease of a joint subfield — geometrically, this means that a diamond bottom is emerging or disappearing (see Fig. 5.1–bottom). Obviously a merging (respectively a scission) is also an enrichment (resp. impoverishment) of each of the superfields.

Hence, each of these three kinds of patterns corresponds to a growth or a decrease in agent set size. The interpretation of the population change ultimately depends on the EC position in the partial epistemic hypergraph, and should vary according to whether (i) there is simply a change in population, (ii) the change occurs for a subfield and (iii) this subfield is in fact a joint subfield. These patterns, summarized on Fig. 5.1, describe epistemic evolution with an increasing precision. More precise patterns could naturally be proposed, but as we shall see, these ones are nevertheless sufficiently relevant for the purpose of our case study.

5.1 Empirical protocol

We complete here the empirical protocol presented in Chap. 3 to make it suitable for this method. To describe the community evolution over several periods of time, as previously we use data telling us *when* an agent s uses a concept c . Accordingly, we divide the database into several time-slices, and build a series of relation matrices aggregating all events of each corresponding period. Before doing so, we need to specify the way we choose the *time-slice width* (size of a period), the *time-step* (increment of time between two periods) and the way we *attribute a concept* to an agent, thus to an article.

Time-slice width We must choose a sufficiently wide time-slice in order to take into account minority communities (who publish less) and to get enough information for each author (especially those who publish in multiple fields).² Doing so also smoothes the data by reducing noise and singularities due to small sample sizes.

However, when taking a longer sample size, we take the risk of merging several periods of evolution into a single time-slice. There is arguably a tradeoff between short but too insignificant time-slices, and long but too aggregating ones. This parameter must be empirically adapted to the data: depending on the case, it might be relevant to talk in terms of months, years or decades.

²For instance, extremely few authors publish more than one paper during a 6-month period, so obviously 6-month time-slices are not sufficient.




Figure 5.2: Series of overlapping periods P_1 , P_2 and P_3 .

Time-step The time-step is the increment between two time-slices, so it defines the pace of observation. We need to consider overlapping time-slices, since we do not want to miss developments and events covering the end of a period and the beginning of the next one. Therefore, we need to choose a time-step strictly shorter than the time-slice width, as shown on Fig. 5.2.

Moreover, the time-step is strongly related to the community *time-scale*: seeing almost no change between two periods would indicate that we are below this time-scale. We need to pick out a time-step such that successive periods exhibit sensible changes.³

5.2 Case study, dataset description

We considered the same particular community of embryologists working on the model animal “zebrafish”, but extended the set of articles to the whole period 1990–2003. Thus, we covered what experts of the field call the beginning of the major growth of this community, up to recent times. As such, this timespan corresponds to a recent and important period of expansion for this community, which gathered approximately 1,000 agents at the end of 1995, and reached nearly 10,000 people by end-2003. We chose a time-slice width of 6 years, with a time-step of 4 years — that is, a 2 years overlap between two successive periods. We thus splitted the database in three periods: 1990-1995, 1994-1999 and 1998-2003.

To limit computation costs, we restricted the dictionary to the 70 most used and

³We may nevertheless suggest a more objective method for choosing time-step and overlap sizes. Consider indeed the density of evolution patterns “ $d(i) = \# \text{patterns during } i / \text{time-slice width}$ ”, for a given time-slice i . To this end we need to define clearly when a pattern is present: we have to define a threshold μ such that we consider a pattern to be present as soon as a given EC size changes by $\mu\%$ between two periods. The goal is thus to get the maximum uniformity in time-slice significance, which is equivalent to have the smallest variance for d . We could finally draw the variance σ_d for various values of time-step and overlap, and select values that yield the smallest variance.

significant words in the community, selected with the help of our expert. We also considered for each period a random sample of 255 authors. Besides, we used a fixed-size author sample so as to distinguish taxonomic evolutions from the trend of the whole community. Indeed, as the community was growing extremely fast, an EC could become more populated because of the community growth, while it was in fact becoming less attractive. With a fixed-sized sample, we could compare the relative importance of each field with respect to others within the evolving taxonomy.

5.3 Rebuilding history

5.3.1 Evolution description

Few changes occurred between the first and the second period, and between the second and the third period: the second period is a transitory period between the two extreme periods. This seems to indicate that a 4-year time-step is slightly below the time-scale of the community, while 8 years can be considered a more significant time-scale.⁴

We hence focus on two periods: the first one, 1990-1995, and the third one, 1998-2003. The two corresponding partial epistemic hypergraphs are drawn on Fig. 5.3 (page 50). We observe that:

- *First period (1990-1995), first partial epistemic hypergraph:* $\{develop\}$ and $\{pattern\}$ strongly structure the field: they are both large communities and present in many subfields.

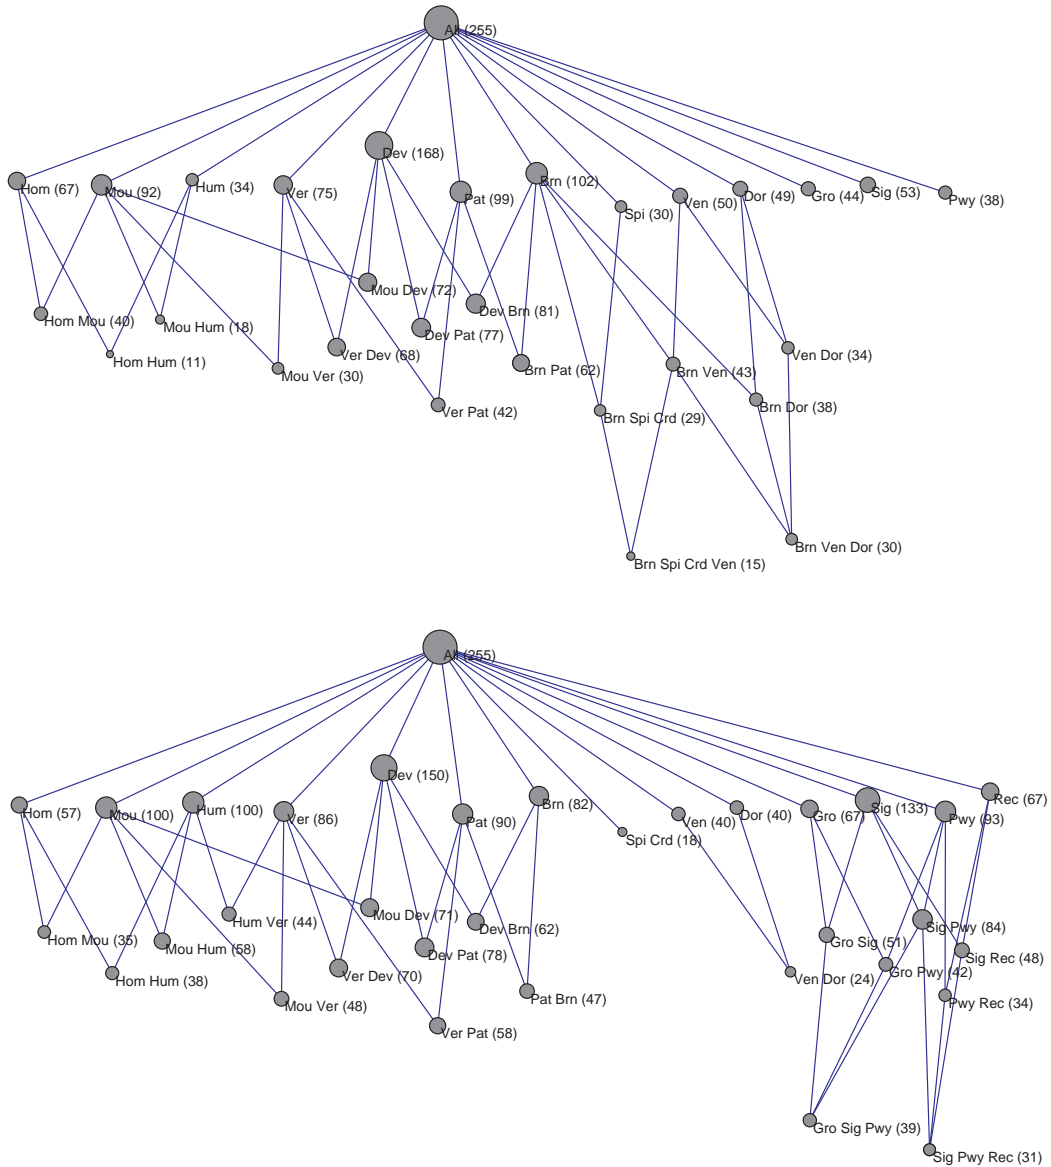
Then, slightly to the right of the partial hypergraph, a large field is structured around *brain*⁵ and *ventral* along with *dorsal*. Excepting one agent, the terms *spinal* and *cord* form a community with *brain*; this dependance suggests that the EC $\{spinal, cord\}$ is necessarily linked to the study of *brain*. Subfields of $\{brain\}$ also involve *ventral* and *dorsal*. In the same view, $\{brain, ventral\}$ has a common subfield with $\{spinal, cord\}$.

To the left, another set of ECs is structured around $\{homologous\}$, $\{mouse\}$ and $\{vertebrate\}$, and $\{human\}$, but significantly less.

- *Third period (1998-2003), second partial epistemic hypergraph:* We still observe a strong structuration around $\{develop\}$ and $\{pattern\}$, suggesting that the core

⁴Kuhn (1970) asserts that old ideas die with old scientists — equivalently new ideas rise with new scientists. In this community, 8 years could represent the time required for a new generation of scientists to appear and define new topics; e.g. the time between an agent graduation and his first students graduation.

⁵We actually grouped *brain*, *nerve*, *neural* and *neuron* under this term.



Legend: All: the whole community, Hom: *homologue/homologous*, Mou: *mouse*, Hum: *human*, Ver: *vertebrate*, Dev: *development*, Pat: *pattern*, Brn: *brain/neural/nervous/neuron*, Spi: *spinal*, Crd: *cord*, Ven: *ventral*, Dor: *dorsal*, Gro: *growth*, Sig: *signal*, Pwy: *pathway*, Rec: *receptor*.

Figure 5.3: Two partial epistemic hypergraphs representing the community at the end of 1995 (*top*) and at the end of 2003 (*bottom*). Figures in parentheses indicate the number of agents per EC. Lattices established from a sample of 255 agents (out of 1,000 for the first period vs. 9,700 for the third one).

topics of the field did not evolve.

However, we notice the strong emergence of three communities, *{signal}*, *{pathway}* and *{growth}*, and the appearance of a new EC, *{receptor}*. These communities form many joint subcommunities together, as we can see on the right of this lattice, indicating a convergence of interests.

Also, there is a slight decrease of *{brain}*. More interestingly, there is no joint community anymore with *{ventral}* nor *{dorsal}*. The interest in *{spinal cord}* has decreased too, in a larger proportion.

Finally, *{human}* has grown a lot, not *{mouse}*. These two communities are both linked to *{homologous}* on one side, *{vertebrate}* on the other. While the importance of *{homologous}* is roughly the same, the joint community with *{human}* has increased a lot. The same goes with *{vertebrate}*: this EC, which is almost stable in size, has a significantly increased role with *{mouse}* and especially *{human}* (a new EC *{vertebrate, human}* just appeared).

5.3.2 Inference of an history

To summarize in terms of dynamic patterns: some communities were stable (e.g. *{pattern}*, *{develop}*, *{vertebrate, develop}*, *{homologous, mouse}*, etc.), some enjoyed a burst of interest (*{growth}*, *{signal}*, *{pathway}*, *{receptor}*, *{human}*) or suffered less interest (*{brain}* and *{spinal cord}*). Also, some ECs merged (*{signal}*, *{pathway}*, *{receptor}* and *{growth}* altogether; and *{human}* both with *{vertebrate}* and *{homologous}*), some splitted (*{ventral-dorsal}* separated from *{brain}*). We did not see any *strict* enrichment or impoverishment — even if, as we noted earlier, merging and splitting can be interpreted as such.

We can consequently suggest the following story: (i) research on brain and spinal cord depreciated, weakened their link with ventral/dorsal aspects (in particular the relationship between ventral aspects and the spinal cord), (ii) the community started to enquire relationships between signal, pathway, and receptors (all actually related to biochemical messaging), together with growth (suggesting a messaging oriented towards growth processes), indicating new very interrelated concepts prototypical of an emerging field, and finally (iii) while mouse-related research is stable, there has been a significant stress on human-related topics, together with a new relationship to the study of homologous genes and vertebrates, underlining the increasing role of *{human}* in these differential studies and their growing focus on human-zebrafish comparisons (leading to a new “interdisciplinary” field).

Point (ii) entails more than the mere emergence of numerous joint subcommunities: all pairs of concepts in the set *{growth, pathway, receptor, signal}* are in-

volved in a joint subfield. Put differently these concepts form a clique of joint communities, a pattern which may be interpreted as *paradigm emergence* (see Fig. 5.3–bottom).

5.3.3 Comparison with real taxonomies

We compared these findings with empirical taxonomical data, coming both from:

1. Expert feedback: Our expert, Nadine Peyri ras, confirms that points (i), (ii) and (iii) in the previous paragraph are an accurate description of the field evolution. For instance, according to her, the human genome sequencing in the early 2000s (International Human Genome Sequencing Consortium, 2001) opened the path to zebrafish genome sequencing, which made possible a systematic comparison between zebrafish and humans, and consequently led to the development described in point (iii). In addition, the existence of a subcommunity with *brain*, *spinal cord* and *ventral* but not *dorsal* reminded her the initial curiosity around the ventral aspects of the spinal cord study, due to the linking of the ventral spinal cord to the mesoderm (notochord), i.e. the rest of the body.
2. Literature: The only article yet dealing specifically with the history of this field seems to be that of Grunwald & Eisen (2002). This paper presents a detailed chronology of the major breakthroughs and steps of the field, from the early beginnings in the late 1960s to the date of the article (2002). While it is hard to infer the taxonomic evolution until the third period of our analysis, part of their investigation confirms some of our most salient patterns: “Late 1990s to early 2000s: Mutations are cloned and several genes that affect common processes are woven into molecular pathways” — here, point (ii). Note that some other papers address and underline specific concerns of the third period, such as the development of comparative studies (Bradbury, 2004; Dooley & Zon, 2000).
3. Conference proceedings: Finally, some insight could be gained from analyzing the evolution of the session breakdown for the major conference of this community, “Zebrafish Development & Genetics” (Cold Spring Harbor Laboratory, 1994, 1996, 1998, 2000, 2001, 2002, 2003). Topic distribution depends on the set of contributions, which reflects the current community interests; yet it may be uneasy for organizers to label sessions with a faithful and comprehensive name — “*organogenesis*” for instance covers many diverse subjects. Reviewing the proceedings roughly suggests that comparative and sequencing-related studies are an emerging novelty starting in 1998, at the

beginning of the third period, which agrees with our analysis. On the contrary, the importance of issues related to the brain & the nervous system, as well as signaling, seem to be constant between the first and the third period, which diverges from our conclusions.

The expert feedback here is obviously the most valuable, as it is the most exhaustive and the most detailed as regards the evolving taxonomy — the other sources of empirical validation are more subject to interpretation and therefore more questionable. A more comprehensive empirical protocol would consist in including a larger set of experts, which would yield more details as well as a more intersubjective viewpoint, thus objective.

Chapter 6

Discussion and conclusion

We presented here a method for extracting a meaningful taxonomy of any knowledge community, in the form of hypergraphs, and successfully validated it with empirical expert-based descriptions for a given scientific community. In other words, we designed a valid projection function P from the low-level of relations between agents and concepts to the high-level of epistemological descriptions. In particular, in Sec. 5.3, the two partial epistemic hypergraphs can be seen as $P(L_{1995})$ and $P(L_{2003})$, which match expert-based H_{1995} and H_{2003} . More, the transition from H_{1995} to H_{2003} (η^e) is also reproduced: we provide a valid high-level dynamics η by describing the taxonomy evolution description.

The computer programs we created to achieve data processing, empirical experiments and Galois lattice computations will also be made available shortly, as open source software. It will thus be possible to reuse them in potentially any other similar case. We are hopeful that the process can be widely used for representing and analyzing static and dynamic taxonomies: in the first place, it could be helpful to historians of science, in domains where historical data is lacking — notably when examining the recent past. Studies such as the recent history of the zebrafish community, written by scientists themselves from this community (Grunwald & Eisen, 2002), could profit from such non-subjective analysis. In this particular case the present study might be considered the second historical study of the “zebrafish” community. At the same time, with the growing number of publications, some fields produce thousands of articles per year. It is more and more difficult for scientists to identify the extent of their own community: they need efficient representation methods to understand their community structure and activity.

More generally, unlike many categorization techniques, community labelling here is straightforward, as agents are automatically bound to a semantic content. Additionally, these categories would have been hard to detect using single-network-

based methods, for instance because agents of a same EC are not necessarily socially linked. Moreover, projection of such two-mode data onto single-mode data often implies massive information loss (see Sec. 2.3). Finally, the question of overlapping categories — hardly addressed when dealing with dendrograms — is easily solved when observing communities through lattices.

Also, using this method is possible in at least any practical case involving a relationship between agents and semantic items. As stated by Cohendet, Kirman & Zimmermann (2003), “*a representation of the organization as a community of communities, through a system of collective beliefs (...), makes it possible to understand how a global order (organization) emerges from diverging interests (individuals and communities).*”¹ In addition to epistemology, scientometrics and sociology, other fields of application and validation include economics (start-ups dealing with technologies, through contracts), linguistics (words and their context, through co-appearance within a corpus), marketing (companies dealing with ethical values, through customers cross-preferences), and history in general (e.g. evolution of industrial patterns linked to urban centers (White & Spufford, 2006)). Having significant results in many distinct fields would support the overall robustness of GL-based taxonomy building.

Lattice manipulation On the other hand, our method could enjoy several improvements. Practically, note that computing the whole GL then selecting a partial epistemic hypergraph is certainly not the most efficient option. Rather, computing the upper part and its “valuable” descendance (computing a fixed number of ECs, starting from the top) should perform better — similarly to what is done with “iceberg lattices” (Stumme *et al.*, 2002). Thus GL computation complexity, which is theoretically exponential, is limited upfront by the number of ECs which should be computed. This requires however to use monotonic selection heuristics, i.e. heuristics respecting the lattice partial order: if $(S, N) \sqsubset (S', N')$, then $h(S, N) < h(S', N')$. Similarly, selection heuristics must allow for significant child nodes to appear. Indeed, when two fields do not seem to form a joint subfield in the partial hypergraph, it is hard to know whether they actually form a joint subfield but are below the threshold. In the second lattice for instance, although of similar importance as *{spinal cord}* (17 vs. 18 agents), the EC *{brain, spinal cord}* is excluded by the selection threshold and does not appear, possibly leading us to wrongly deduce that *{brain}* does not mix with *{spinal cord}*.

In the same direction, we could endeavor to exclude false positives such as fortuitous intersections (as discussed in section 4.1) and merge clusters of ECs

¹“*Une représentation de l’organisation comme une communauté des communautés, à travers un système de croyances collectives (...), permet (...) de comprendre comment émerge un ordre global (organisation) à partir d’intérêts divergents (individus et communautés).*”

into single multidisciplinary ECs (like for instance “*signal*,” “*pathway*,” “*receptor*”). This would lead to reduced partial hypergraphs containing merged sublattices. Questions arise however regarding the best way to define a cluster of ECs without destroying overlapping communities, one of the most interesting feature of GLs. Accordingly, it could also be profitable to disambiguate and regroup terms in the lattice using for instance Natural Language Processing (NLP) tools (Ide & Véronis, 1998): certainly not everyone assigns the same meaning to “*pattern*,” we would thus have to introduce “*pattern-1*,” “*pattern-2*,” etc.

More generally, improving linguistic processing could be very informative, and could first include the use of:

- Lemmatizers: algorithms giving the root of a word, instead of using a stemmer like the one used here (the “Porter stemmer,” though it is also a quite simple yet efficient lemmatizer);
- Taggers: algorithms detecting word grammatical status in context, e.g. “subject,” “verb,” etc.;
- Morphological analyzers: algorithms recognizing the shape of a word actually composed of two or more words, like “molecular biology,” “positron emission tomography,” etc.;
- Dictionaries: ontologies of the domain, returning classes of words considered as equivalent (as stated in Chap. 3), like “zebrafish” and “rerio brachydanio,” the former being the common name of the latter;
- Disambiguators: algorithms determining the meaning of words by examining the context in which they are used (Wang *et al.*, 2000).

Most of these tools already exist, although their joint use would require a judicious work of integration. Alternatively, it could be useful to compare these results with those from data processed by human experts, where all linguistic processing problems become quite obsolete. For instance, (i) by providing them with a fixed list of concepts and making them classify agents according to this list, or (ii) by making them identify a restricted list of words they know to be sufficiently descriptive for a given set of articles (e.g. protein nomenclature consisting of very specific names (Lelu *et al.*, 2004)).

Lastly, considering that some authors are more or less strongly related to some concepts, the binary relationship may seem too restrictive. To this end, we could use a weighted relation matrix together with fuzzy GLs (Belohlavek, 2000).

Dynamics study Another major class of improvements is related to the study of the dynamics. Indeed, we are now able to represent an evolving taxonomy but we ignore whether individual agents have fixed roles or not. In particular, the stability of the size of an EC does not imply the stability of its agent set. Fortunately, even if our random agent samples are not consistent across periods, it would be easy to rebuild the whole community taxonomy by filling the partial ECs with their corresponding full agent sets. In this case, field scope enrichment or impoverishment could be described in a better way: by monitoring an identical agent set, and by watching whether its intension increases or not.

More generally, we could address this topic by considering the lattice *dynamics*, instead of adopting a *longitudinal* approach. A dynamic study would yield a better representation of field evolution at smaller scales, nevertheless saving us the empirical discussion about the right time-step.

Conclusion of Part I

In this part, we proposed a method for describing and categorizing knowledge communities as well as capturing essential stylized facts regarding their structure. After having reviewed the definitions in use in social science for knowledge communities, or “epistemic communities,” we formally defined an epistemic community as the largest group of agents who share and work on the same concepts — as such, a conception close to structural equivalence. We showed next that the Galois lattice structure was an adequate clustering method with respect to this definition. Assuming that such communities are structured in fields and subfields of common concerns, a GL faithfully represents epistemic community taxonomies by automatically partitioning the community into hierarchic fields and subfields. In addition, it accurately renders overlaps among epistemic communities, commonly called interdisciplinary fields. Finally, because it relies on the very duality of epistemic communities (agents having common interests), our method diverges from single-network-based methods using for instance relationships or semantic proximity.

Yet, it was unclear whether this was sufficient to make it a useful method for appraising so-produced taxonomies, because the set of all epistemic communities could possibly prove really huge and intractable. *GLs organize the data but they do not reduce it much.* To this end, we conjectured the existence of criteria enabling us to discriminate within the lattice between “uninteresting” communities and interesting ones; among which EC size and position in the lattice were of particular interest. With respect to heuristics based on these criteria, selecting the most relevant epistemic communities produced a partial epistemic hypergraph providing a

manageable representation of the hierarchical structure.

Empirical results on an embryologist community centered around the model animal zebrafish confirmed this expectation even with imperfect data quality, mostly because of an approximative linguistic processing. More generally, we managed to reproduce a partition of the community assessed by domain experts. Consequently, the longitudinal study of such partial taxonomies made possible an historical description. In particular, we proposed to capture stylized facts related to epistemic evolution such as field progress, decline and interaction (merging or splitting). We ultimately applied our method to the subcommunity of embryologists working on the “zebrafish” between 1990 and 2003, and successfully compared the results with taxonomies given by domain experts.

Part II

Micro-foundations of epistemic networks

Summary of Part II

The main purpose of this part is to *micro-found* the high-level features we observed in the Part I — exhibit L and λ such that $P \circ \lambda(L) = \eta^e(H)$. In particular, we aim to know which processes at the level of agents may account for the emergence of epistemic community structure. To achieve a morphogenesis model reproducing this phenomenon, we first need to build tools that enable the estimation of interaction and growth mechanisms from past empirical data. Then, assuming that agents and concepts are co-evolving, we successfully reconstruct a real-world scientific community structure for a relevant *selection* of high-level stylized facts.

Introduction

“Des Esseintes (...) faisait l'exégèse de ces textes; il se complaisait à jouer pour sa satisfaction personnelle, le rôle d'un psychologue, à démonter et à remonter les rouages d'une œuvre”²

A rebours, J.-K. Huysmans.

In the preceding part, we characterized EC structure as a high-level stylized fact for a socio-semantic complex system. Here, we will endeavor to “micro-found” these features. In other words, we would like to rebuild this phenomenon from a lower-level perspective, starting from the local behavior of agents immersed in such an epistemic network. This task is threefold:

- First, define formally the framework of epistemic networks,
- Second, design measurement tools and proceed with the observation of relevant empirical facts of the networks, both high- and low-level,
- Third, reconstruct the real-world structure with the help of a dynamic network morphogenesis model.

On the whole, this amounts to find the solution of a reverse problem: given an evolving epistemic network, what kind of (possibly minimal) dynamics allow to rebuild its structure? To bind this problem to our general reconstruction framework, this comes to find λ such that given η^e and P , we have $P \circ \lambda = \eta^e \circ P$.

We make the following assumption: **modeling interactions at the level of agents who co-evolve with the concepts they manipulate is sufficient to carry the micro-founded reconstruction of this social complex system.** This question relates more broadly to a current issue in structural social science. Modeling social network formation has indeed constituted a recent challenge for this area of research. Social networks are usually interaction networks — nodes are agents and links between nodes represent interactions between agents. In this respect,

²“Des Esseintes (...) expounded these texts; he took a delight, for his own personal satisfaction, in playing the part of psychologist, in unmounting and remounting the machinery of a work” (Huysmans: *Against the Grain*).

proposing morphogenesis models for these networks has involved several disciplines linked both to mathematical sociology, graph theory (computer science and statistical physics) and economics (Skyrms & Pemantle, 2000; Albert & Barabási, 2002; Cohendet *et al.*, 2003). Most of the recent interest in this topic has stemmed from the universal empirical observation that the structure of real networks — including social networks — strongly differ from that of uniform random graphs *a la* Erdős-Rényi (1959), where links between agents are present with a constant probability p . The discrepancy is particularly sensible with respect to two particular statistical parameters: the local topological structure, which has been found to be abnormally clustered and dense in real networks (Watts & Strogatz, 1998), and the node connectivity distribution (or *degree* distribution), which empirically follows a power-law (Barabási & Albert, 1999) instead of a Poisson law in Erdős-Rényi's model (ER). These phenomena suggested that link formation does not occur randomly but rather depends on node and network properties — that is, agents do not interact at random but instead according to *heterogenous* preferences for other nodes. While this fact was already well-documented in social science (Lazarsfeld & Merton, 1954; Touhey, 1974; McPherson & Smith-Lovin, 2001), general network models had been limited for long to ER-like random graphs (May, 1972; Barbour & Mollison, 1990; Wasserman & Faust, 1994; Zegura *et al.*, 1996).

Subsequently, much work has been focused on novel non-uniform interaction and growth mechanisms, in order to determine processes explaining and reconstructing complex network structures consistent with those observed in the real world (Dorogovtsev & Mendes, 2003). The consistency, in turn, has been validated through a rich set of statistical parameters measured on empirical networks, and not limited to degree distributions and clustering coefficients.

After a brief overview of existing network growth models — and particularly in relation with social networks — the goal of this part is twofold. Firstly, we design tools for measuring empirically micro-level phenomena at work in evolving networks, in order to infer and design the interaction behavior of agents. Indeed, even when cognitively, sociologically or anthropologically credible, most of the hypotheses driving these models are mathematical abstractions whose empirical measurement and justification are dubious, if any. We hence apply these instruments to the epistemic network of scientists working on the zebrafish, and eventually suggest significant implications for morphogenesis models. Secondly, we use this knowledge to introduce a model that successfully rebuilds relevant stylized facts observed in this epistemic network.³

³Some portions of this part, concerning in particular the epistemic network framework and the measurement of interaction propensities, can be found in more details in (Roth & Bourguine, 2003; Roth, 2005; ?). Besides, Sec. 9.3 is linked to a preliminary study of basic dynamic parameters published in (Latapy *et al.*, 2005).

Chapter 7

Networks

7.1 Global overview

Measuring and modeling Formally, as noted in Ch. 1, a network (or equivalently a graph) is a set of nodes (or vertices) with connections between them: links (or edges), possibly directed (going explicitly from a node to another node) or undirected (symmetric, without any orientation). Networks are omnipresent in the real world: from the lowest levels of physical interaction, in the study of mean fields and spin glasses for instance (Parisi, 1992; Fischer & Hertz, 1993), to higher levels of description such as biology (Yuh *et al.*, 1998; D’Haeseleer *et al.*, 2000; Hasty *et al.*, 2001), sociology (White *et al.*, 1976; Granovetter, 1985; Wasserman & Faust, 1994; Degenne & Forse, 1999; Pattison *et al.*, 2000; Doreian *et al.*, 2005), economics (Kirman, 1997; Cowan *et al.*, 2002; Deroian, 2002; Goyal, 2003; Carayol & Roux, 2004) and linguistics (Quillian, 1968; Fellbaum, 1998). Along with the empirical investigation of real-world networks, scientists need models for both descriptive and explanatory purposes — either to study processes immersed in a network structure, or to exhibit network creation processes deemed key for the explanation or reproduction of several stylized facts observed in the real world.

For long however, the appraisal of networks had been restricted to theoretical approaches in graph theory and small scale empirical studies on a case-by-case basis. In this respect, network models were mostly limited to the seminal work of Erdős-Rényi (1959) and their “random network model”, based on a random wiring process where each pair of nodes has a constant probability p to be bound by a link. Random networks generated by the Erdős-Rényi (ER) model are often denoted by $G_{N,p}$, because the only parameters of their model are p and the number of nodes N .

The assumption that the ER model was an accurate description of reality had remained unchallenged for a long time. Yet, the empirical study of networks is a

sibling task of the design of models: new measurement tools reveal caveats of former models, thus pushing towards the introduction of new, more accurate models. In this respect, the recent availability of increasingly larger computational capabilities has made possible the use of quantitative methods on large networks, which yielded surprising results and consequently precipitated an unprecedented interest in networks (Barabási, 2002; Dorogovtsev & Mendes, 2003; Newman, 2003). Three statistical parameters in particular appeared to provide an enormous insight on the topological structure of networks:

- the clustering coefficient — that is, the proportion of neighbors of a node who are also connected to each other, averaged over the whole network;
- the average distance — i.e. the length of the shortest path between two nodes, averaged over all pairs of nodes;
- the degree distribution — the degree (or the connectivity) of a node is basically the number of nodes this node is connected to.¹

A new turn These novel instruments opened the way to the distrust of the ER model. In 1998 indeed, Watts and Strogatz (1998) discovered that clustering coefficients for many real-world networks were in flagrant contradiction with those predicted by the ER model. They subsequently introduced a new model, “the small-world network” model, consisting of a *ring* of nodes each connected to their closest neighbors, with a proportion p of these links being randomly rewired (p is thus a rewiring probability). Empirical values for the clustering coefficient were in close adequation with those of the Watts-Strogatz model (WS), which like the ER model respects a realistic shortest path length. The “small-world” metaphor was striking and compelling, as these two features recalled intuitions about real-world networks, especially social networks. A high clustering coefficient suggests that many agents are forming dense, local areas of strongly connected nodes; in sociology, this relates to the concept of transitivity (Wasserman & Faust, 1994). On the other hand, a low shortest length path indicates that a node is generally not “far” from any other node in the network, when considering the number of intermediate agents needed to travel from a given node to another one — a feature observed in real social networks as well (Milgram, 1967; Dodds *et al.*, 2003).

At about the same time, Redner (1998) empirically measured the distribution of degrees in a citation network and found it to be scale-free — that is, it follows a power law with $P(\text{degree} = k) \propto k^{-\alpha}$. This fact contradicted the expectations of both ER and WS models: with ER, the degree distribution can be approximated

¹In a directed network, we have to distinguish the number of outgoing links from the number of incoming links, respectively denoted by outcoming degree vs. incoming degree.

by a Poisson law ($P(k) \propto \exp(-\alpha k)/k!$) (Bollobás, 1985), with an exponentially low probability of finding high-degree nodes. Nearly the same goes for WS (Barabási *et al.*, 1999). Shortly thereafter, Faloutsos *et al.* (1999) discovered that the physical topology of the Internet network was nothing but a scale-free network and Barabasi & Albert (1999) discovered the same feature in the world wide web, and collaboration networks. At this point, the ER model had been totally discredited as a way to render the topology of real-world networks. Simultaneously, dynamical processes were highlighted as an efficient feature for designing accurate models, yielding at the same time a significant and realistic insight on the self-organizing processes at work during morphogenesis.

7.2 A brief survey of growth models

History More specifically, Barabasi & Albert (BA) insisted on the point that such topology could be due to two very particular phenomena that models were so far unable to take into account: network growth, and preferential attachment of nodes to other nodes. They thus pioneered the use of these two features to successfully rebuild a scale-free degree distribution. In their network formation model, new nodes arrive at a constant rate and attach to already-existing nodes with a likelihood linearly proportional to their degree. This model was a great success and has been widely spread and reused. As a consequence, the term “preferential attachment” has been often understood as degree-related preferential attachment only, in reference to BA’s work.

Since then, many other authors introduced network morphogenesis models with diverse modes of preferential link creation depending on various node properties (attractiveness (Dorogovtsev *et al.*, 2000; Krapivsky *et al.*, 2000), age (Dorogovtsev & Mendes, 2000), common neighbors (Jin *et al.*, 2001), fitness (Caldarelli *et al.*, 2002), centrality, euclidian distance (Manna & Sen, 2002; Fabrikant *et al.*, 2002), hidden variables and “types” (Boguna & Pastor-Satorras, 2003; Söderberg, 2003), bipartite structure (Peltomaki & Alava, 2005), etc.) and various linking mechanisms (stochastic copying of links (Kumar *et al.*, 2000), competitive trade-off and optimization heuristics (Fabrikant *et al.*, 2002; Berger *et al.*, 2004; Colizza *et al.*, 2004), payoff-biased network reconfiguration (Carayol & Roux, 2004), two-steps node choice (Stefancic & Zlatic, 2005), group formation (Ramasco *et al.*, 2004; Guimera *et al.*, 2005), Yule processes (Morris, 2005), to cite a few). On the other side, growth processes (if any) were often reduced to the regular addition of nodes which attach to older nodes — sometimes growth is absent and studies are focused on the evolution of links only.

Following BA’s initial model, most of these studies aimed first and before all at

reproducing degree distributions, which had obviously to be scale-free.² Depending on the application field of the model — WWW (Kumar *et al.*, 2000), protein networks (Eisenberg & Levanon, 2003), social networks (Newman, 2001d), citation networks (Vázquez, 2001), etc. — various other stylized facts can be selected, used and compared with real-world values. Statistical parameters include notably clustering coefficient, mean distance (shortest path length), largest connex component size (giant component), assortative mixing,³ existence of feedback circuits (or cycles), number of second neighbors, and one-mode community structure (Pattison *et al.*, 2000; Newman, 2001d; Caldarelli *et al.*, 2002; Watts *et al.*, 2002; Guelzim *et al.*, 2002; Girvan & Newman, 2002; Latapy & Pons, 2004; Boguna *et al.*, 2004; Guimera *et al.*, 2005).

Methodology In such approaches, the idea is generally to exhibit high-level statistical parameters and to suggest low-level network processes, such that the former could be deduced, or recreated, from the latter. Obviously, after having selected a set of relevant stylized facts to be explained or reconstructed, designing network morphogenesis models consists of two subtasks: it requires to define the way agents are bound to interact with each other, as well as to specify how the network grows. However and even in recent papers, hypotheses on such mechanisms are often arbitrary and at best supported by qualitative intuitions. This is particularly true for the definition of the preferential attachment (PA) which rarely enjoys empirical verification, in spite of the rich diversity of propositions. While this attitude is still convenient for normative models, this is clearly insufficient for descriptive models — although even normative models should be able to suggest means to reach the “norm” they introduce.

In the remainder of this part, we will thus endeavor to (i) exhibit high-level stylized facts characteristic of epistemic networks, notably the EC structure observed in the previous part, (ii) point out relevant low-level features that may account for these high-level facts, (iii) design measurement tools to appraise these low-level features, and (iv) design a reconstruction model based on the observed low-level dynamics that rebuilds the high-level one. *In fine*, the goal of this model is to reproduce the morphogenesis of epistemic networks, and to show consequently that these networks are produced by the dynamic co-evolution of agents and concepts.

²There is a long history of models generating all sorts of power-law distributions (size of cities, incomes, etc.), dating back to the early twentieth century (from Pareto, Lotka, Zipf and Yule, to Simon and Mandelbrot) (Mitzenmacher, 2003; Newman, 2005). The significant difference in this “network-based paradigm” is that present network models are node-based (agent-based), not anymore relying on global differential equations (Bonabeau, 2002).

³This term denotes the fact that neighbors of a node have a similar degree or not: high-degree nodes connected to high-degree ones (like in social networks) or to low-degree ones (like in other kinds of networks) (Newman, 2002).

Before that, we formally introduce the objects we deal with.

7.3 Epistemic networks

In the first part, we studied ECs with the help of a single relation linking agents to concepts — as such creating a bipartite graph: a socio-semantic network. A bipartite graph (or two-mode network) is a graph whose vertices can be decomposed into two disjoint sets, such that no link exists between pairs of vertices belonging to the same set (as opposed to a monopartite graph, also called one-mode network). In addition to the socio-semantic network, we introduce two related networks: a *social network*, involving links between agents, and a *semantic network*, with links between concepts. As a result, an *epistemic network* is made of these three networks.

Definitions

Definition 9 (Social network). *The nodes in the social network \mathcal{S} are agents, and links represent the joint appearance of two agents in an event.*

Thus $\mathcal{S} = (\mathbf{S}, E_{\mathbf{S}})$, where \mathbf{S} denotes the set of agents and $E_{\mathbf{S}}$ denotes the set of undirected links. As time evolves, new events occur (e.g., new articles are published), new nodes are possibly added to \mathbf{S} and new links are created between each pair of interacting agents. We actually consider the temporal series of networks \mathcal{S}_t with $t \in \mathbb{N}$ (events are dated with an integer), in order to observe the dynamics of the network.

The semantic network is very similar to the social network:

Definition 10 (Semantic network). *The semantic network \mathcal{C} is the network of joint appearances of concepts within events, where nodes are concepts and links are co-occurrences.*

Identically to \mathcal{S} , we have $\mathcal{C} = (\mathbf{C}, E_{\mathbf{C}})$. When a new event occurs, new concepts are possibly added to the network, and new links are added between co-appearing concepts. As the social network is the network of joint appearances of agents, so is the semantic network with concepts. In the same way we did with the previous networks, we link scientists to the words they use, i.e. we add a link whenever an author and a concept co-appear within an event, establishing an obvious duality between the two networks. This duality has been exploited in the previous part for the sole purpose of describing epistemic communities, yet it is also key for explaining the reciprocal influence and co-evolution of authors and concepts.

Definition 11 (Socio-semantic network). *The socio-semantic network $\mathbf{G}_{\mathbf{SC}}$ is made of agents of \mathbf{S} , concepts of \mathbf{C} , and links between them, $E_{\mathbf{SC}}$, representing the usage of concepts by agents.*

Weighted networks An important issue relative to networks in general concerns the nature of links. Depending on the model goals and the desired precision, we may want to take into account the fact that two nodes have interacted more than once (thus introducing *link strength*), or that their interactions are more or less recent (thus introducing *link age*). Relationships should consequently be different according to whether agents have interacted only once and a long time ago, or they have recently interacted on many occasions. An easy and practical way for dealing with these notions is to use a weighted network:

- in a *non-weighted network*, we say that two nodes are linked as soon as they interact, i.e. they jointly appear in at least one event. Links can only be active or inactive.
- in a *weighted network*, links are provided with a weight $w \in \mathbb{R}^+$, possibly evolving in time. We can therefore easily represent multiple interactions by increasing the weight of a link, as well as render the age of a relationship by decreasing this weight — for instance by applying an aging function.

This latter framework is more general as it makes it possible to model a non-weighted network (by assigning weights of 1 or 0 respectively to active or inactive links), while it also leaves room for creating *ex post* a non-weighted network from a weighted network by setting a threshold on link weight (such that a link is active when its weight exceeds the threshold, otherwise inactive). Besides, the design and choice of w depends on the objectives of the modeling.

Relations Considering the three networks \mathcal{S} , \mathcal{C} and $\mathbf{G}_{\mathcal{SC}}$, we deal with three kinds of similar links: (i) between pairs of agents $E_{\mathcal{S}}$, (ii) between pairs of concepts $E_{\mathcal{C}}$, and (iii) between concepts and agents $E_{\mathcal{SC}}$; we thus set up three kinds of binary relations:

- (i) a set of binary symmetrical relations $\mathcal{R}_{\alpha}^{\mathcal{S}} \subset \mathbf{S} \times \mathbf{S}$ from the set of agents to the set of agents, and such that given $\alpha \in \mathbb{R}$ and two agents s and s' , we have $s \mathcal{R}_{\alpha}^{\mathcal{S}} s'$ iff the link between s and s' has a weight w strictly greater than the threshold α .
- (ii) a set of binary symmetrical relations $\mathcal{R}_{\alpha}^{\mathcal{C}} \subset \mathbf{C} \times \mathbf{C}$ from the set of concepts to the set of concepts, and such that given $\alpha \in \mathbb{R}$ and two concepts c and c' , $c \mathcal{R}_{\alpha}^{\mathcal{C}} c'$ iff the link between c and c' has a weight $w > \alpha$.
- (iii) a set of binary relations $\mathcal{R}_{\alpha} \subset \mathbf{S} \times \mathbf{C}$ from the set of agents to the set of concepts, and such that given $\alpha \in \mathbb{R}$, an agent s and concept c , $s \mathcal{R}_{\alpha} c$ iff the link between s and c has a weight $w > \alpha$.




Figure 7.1: Sample epistemic network with $\mathbf{S} = \{s, s', s''\}$, $\mathbf{C} = \{c, c', c''\}$, and relations $\mathcal{R}^{\mathbf{S}}$, $\mathcal{R}^{\mathbf{C}}$ (solid lines) and \mathcal{R} (dashed lines).

Noticing that $\alpha < \alpha' \Rightarrow \mathcal{R}_{\alpha'}^{(\cdot)} \subset \mathcal{R}_{\alpha}^{(\cdot)}$, thus giving $\forall \alpha > 0, \mathcal{R}_{\alpha}^{(\cdot)} \subset \mathcal{R}_0^{(\cdot)}$, we infer that the relations $\mathcal{R}_0^{(\cdot)}$ are maximal: two nodes are related whenever there exists a link binding them, whatever its weight.

In the remainder of this part, to make the things simpler we choose to assign weights equal to the number of interactions, with no aging; and we focus on the special case $\alpha = 0$, which corresponds to non-weighted networks. Consequently, we do not pay attention to weights and related phenomena: as long as there has been *any* interaction, a link is established between two nodes. More details on weighted networks can nonetheless be found in e.g. (Barrat *et al.*, 2004). In addition, we only consider *growing networks*, that is, neither nodes nor links may disappear. \mathcal{R}_0 is identical to what \mathcal{R} designates in Part I. To ease the notation, we will denote $\mathcal{R}_0^{\mathbf{S}}$ and $\mathcal{R}_0^{\mathbf{C}}$ by $\mathcal{R}^{\mathbf{S}}$ and $\mathcal{R}^{\mathbf{C}}$, respectively. Note that social, semantic and socio-semantic networks are fully characterized by \mathbf{S} , \mathbf{C} and $\mathcal{R}^{\mathbf{S}}$, $\mathcal{R}^{\mathbf{C}}$ and \mathcal{R} — see Fig. 7.1.

Chapter 8

High-level features

In this chapter, we endeavor to describe a few high-level statistical parameters particularly appropriate for epistemic networks. We thus enrich the high-level description of Part I, consisting in the epistemic hypergraph, with these new features. Translated in the above framework, events are articles, agents are their authors, and concepts are made of expert-selected abstract words.

8.1 Empirical investigation

While we could have looked at many single-network parameters (such as assortativity (Newman & Park, 2003), giant component size (Guimera *et al.*, 2005), single-network communities (Girvan & Newman, 2002; Latapy & Pons, 2004), etc.), we focused instead on features specific to this epistemic network (thus, mostly bipartite parameters) — many results and models are already available for most traditional statistical features.

As previously, empirical data comes from the bibliographical database Medline concerning the well-defined community of embryologists working on the *zebrafish*, this time during the period 1997-2004. The dataset contains around 10,000 authors, 6,000 articles and 70 concepts. The 70 concepts are the same as those selected for Part I — in addition, we consider this set to be given *a priori*: in the semantic network, only links appear, not nodes. The rationale is twofold: first, this is consistent with assumptions used for the preceding dynamic taxonomy study; second, it dramatically reduces computational complexity.

8.2 Degree distributions

In an epistemic network, ties appear in the social, semantic, and socio-semantic networks; hence, four degree distributions are of interest:

1. *The degree distribution for the social network of coauthorship, $P(k)$, shown on Fig. 8.1.* This distribution has been extensively studied in the literature, notably by Newman (2001b; 2001c; 2001d) and Barabasi et al. (2002), among others. It is traditionally said to follow a power law, although often only the tail of the distribution actually follows a power-law. It is indeed easy to see that the distribution shape is not constant: for low degrees, the distribution is sensibly flatter. Instead of a power-law, some may suggest that this distribution follows a log-normal law (Redner, 2005). This observation is very natural as the log-log plot exhibits a parabolic shape, for which the best fitting function is of a log-normal kind.¹

Note that various other shapes may address this fitting problem equally well, such as q -exponential functions (White *et al.*, 2006). In any case, it appears that a strict power-law is not the most accurate description of this degree distribution.

2. *The distribution of degrees k_{concepts} for the semantic network.* Since there are only 70 concepts the data are very sparse, we considered cumulated distributions (plotted on Fig. 8.2 for all eight periods). Obviously all concepts are progressively connected to each other, with almost every concept having a degree of 69 at the end of the last period.
3. *The distribution of degrees from agents to concepts ($k_{\text{agents} \rightarrow \text{concepts}}$).* It follows a power-law: few agents use many concepts, many agents use few concepts. The exponent is similar to that of the social network and constant across periods as well (see Fig. 8.3 — a detailed report on similar phenomena can be found in (Latapy *et al.*, 2005)).
4. *The degree distribution for links from concepts to agents ($k_{\text{concepts} \rightarrow \text{agents}}$).* Again, cumulated distributions were considered to bridge data sparsity. With time, more and more concepts are becoming popular (used by numerous agents), yet the repartition is still heterogeneous, with few concepts being used by a lot of agents, and most concepts being used by an average number of agents (see Fig. 8.3).

Considerations on bipartite graphs The socio-semantic network is obviously a bipartite graph, with agents on one side and concepts on the other. It is also possible to consider the social network itself as a bipartite graph (Wilson, 1982; Wasserman & Faust, 1994; Ramasco *et al.*, 2004; Kossinets, 2005), made of agents on one

¹The interested reader may find in (Mitzenmacher, 2003) a comprehensive comparison of processes underlying the emergence of power-law and log-normal distributions.




Figure 8.1: Degree distribution for the social network. *Dots*: $N(k)$, proportional to $P(k) = \frac{N(k)}{\sum_{k'} N(k')}$. *Solid line*: power-law fit of $P(k)$ with k^γ , here $\gamma = -3.39$. *Inset*: evolution of the exponent γ for 8 periods (mean exponent is $-3.19 \pm .10$). *Dashed line*: Lognormal fit — indeed, the distribution has a parabolic shape: this suggests that $\log N(k) = p_2(\log k)^2 + p_1 \log k + p_0$, thus $P(k) \propto k^{p_2 \log k + p_1}$. This deviates from a strict power law because of the term in $k^{p_2 \log k}$ (here, $p_2 = -0.61 \pm .06$, $p_1 = 1.45 \pm .22$).




Figure 8.2: Cumulated degree distribution for the semantic network, for all 8 periods — from top (1997, light blue) to bottom (2004, black).

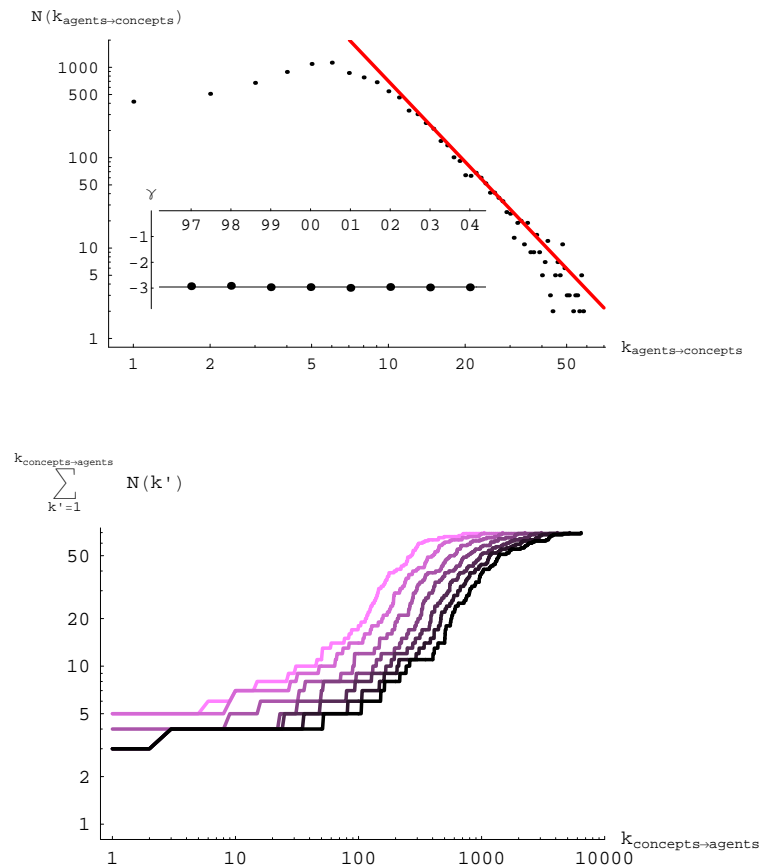


Figure 8.3: Degree distributions for the socio-semantic network. *Top*: Degree distribution from agents to concepts (dots), power-law fit (solid line), and evolution of the exponent γ for all 8 periods (from 1997 to 2004), mean γ is $-2.96 \pm .02$ (see inset). *Bottom*: Cumulated degree distribution from concepts to agents, for 8 periods (1997-2004, from light blue to black).

side, events on the other, and links from agents to events they participate in. Projecting this two-mode graph on a one-mode network (such that two agents are linked in the one-mode network *iff* they are linked to the same event in the two-mode network) yields in turn the classical social network. In this respect, it can be expected that some properties of the bipartite graph and the one-mode projection are strongly correlated: Guillaume and Latapy (2004b) for instance showed that the one-mode projection of a bipartite network preserves scale-free degree distributions. In other words, if the degree distribution from one side of a bipartite graph to the other side follows a power-law, then the projection follows a power-law of the same exponent.

Yet, such bipartite graphs “agents–events” are another (richer) way of considering the social network, by keeping events apart instead of losing some of the information embedded in events. For instance, by doing so the fact that some agents participated in the *same* event is not lost. More generally, any one-mode network can be considered bipartite, if one expands the underlying event structure to a new network of events — to this end, Guillaume & Latapy (2004a) even try to recompose events from a one-mode network.

Nonetheless, this bipartite graph is special: events are bound to appear only once, agents cannot attach to old events; as such, the side of events is merely *historical*. Here, the social network is not the one-mode projection of the socio-semantic network. Agents can bind to old concepts, so can concepts to old agents. In spite of this, social and semantic networks could enjoy some of the properties of a one-mode projection from a bipartite graph, if we consider that these networks are created by using the co-appearance of agents and concepts in *common* events. Thus, there are two underlying bi-partite graphs made of events: agents and events, and concepts and events. The social and semantic networks are respectively one-mode projections of each of these bipartite graphs. Because of their strictly historical structure, we nonetheless discard the ‘artificial’ networks of events.

8.3 Clustering

The clustering coefficient is another valuable parameter, introduced by Watts & Strogatz (1998). It is basically a measure of the transitivity in one-mode networks: in other words, it expresses the extent to which neighbors of a given node are also connected — the sociological metaphor translates into: “friends of friends are friends”. This coefficient is usually found to be abnormally high in social networks, when compared to random networks such as those produced by ER, BA models. By contrast, it is successfully reconstructed by the WS model. Along with degree distribution, this stylized fact has been the target of many more recent

models (Jin *et al.*, 2001; Ebel *et al.*, 2002; Ravasz & Barabási, 2003; Newman & Park, 2003).

Two competing formal definitions have been proposed, potentially yielding significantly different values (Ramasco *et al.*, 2004):

- either a local coefficient, $c_3(i)$, measuring the proportion of neighbors of node i who are connected together,

$$c_3(i) = \frac{[\text{number of pairs of connected neighbors}]}{k_i \cdot (k_i - 1)/2} \quad (8.1a)$$

where k_i is the degree of node i .

- or a global measure C_3 (proportion of connected triangles in the whole network with respect to connected triplets),

$$C_3 = \frac{3 \cdot [\text{number of triangles}]}{[\text{number of broken triangles}]} \quad (8.1b)$$

The factor three comes from the fact that for each triangle there are three “broken triangles” (triplets where only two pairs are connected, see Fig. 8.4).

We focus on the local coefficient for it makes it possible to examine the clustering structure with respect to node properties, in particular node degrees. Here, each article adds complete subgraphs of authors, or *cliques*, to the social network: all authors of a given article are linked to each other. In a network where events are addition of cliques, the clustering coefficient is very likely to be close to one, since each event adds an overwhelming quantity of triangles. Therefore, only nodes participating in multiple events can have neighbors who are not themselves connected to each other. Empirically, the local clustering coefficient is close to 1 and decreases rather slowly with node degree (Fig. 8.5).

As such, in the case of event-based networks, c_3 seems to be a trivial, very poorly informative criterion as regards the clustering structure. Indeed, c_3 is virtually bound *by definition* to be high. More generally, networks built with an underlying event structure are shown to naturally exhibit a high c_3 (Guillaume & Latapy, 2004b; Ramasco *et al.*, 2004).²

Bipartite clustering Very recently, bipartite clustering coefficients have been proposed as a means to have a meaningful clustering measure in spite of this caveat.

²Assuming that the number of agents per event is higher than 2 — otherwise events reduce to simple dyadic interactions, and we fall back onto classical models of single links addition (Catanzaro *et al.*, 2004). This may also explain why many dyadic-interaction models fail to reproduce real-world high clustering coefficients.




Figure 8.4: *Left*: Comparison between a transitive triplet, or *triangle* (top), and a broken triangle, or simply connected triplet (bottom). One-mode clustering coefficients measure the proportion of triangles vs. broken triangles, either globally (C_3) or locally (c_3). *Right*: Comparison between a *diamond* and a broken diamond, with pairs (s', s'') both connected to (c', c'') (top) or not (bottom). Similarly, C_4 and c_4 provide a measure of the proportion of diamonds with respect to broken diamonds.

In a strictly bipartite graph, clearly triangles are impossible: the bipartite socio-semantic network does not render links between agents. To bridge this, a sensible idea consists in measuring the proportion of diamonds; that is, measuring how many pairs of nodes from one side, who are connected *together* to a node of the other side, are also connected to another node of the other side (see Fig. 8.4).³ In other words, are two agents connected to a same concept likely to be connected to other concepts? Like for the monopartite clustering coefficient, there exists both a global version C_4 (Robins & Alexander, 2004) and, latterly, a local one c_4 (Lind *et al.*, 2005):

- locally, c_4 is the proportion of common neighbors among the neighbors of a node:

$$c_4(i) = \frac{\sum_{i_1=1}^{k_i} \sum_{i_2=i_1+1}^{k_i} \kappa_{i_1, i_2}}{\sum_{i_1=1}^{k_i} \sum_{i_2=i_1+1}^{k_i} [(k_{i_1} - \kappa_{i_1, i_2})(k_{i_2} - \kappa_{i_1, i_2}) + \kappa_{i_1, i_2}]} \quad (8.2a)$$

where κ_{j_1, j_2} is the number of nodes which the j_1 -th & j_2 -th neighbors of i have in common (leaving out i).

- globally, C_4 evaluates the proportion of diamonds with respect to potential diamonds:

$$C_4 = \frac{4 \cdot [\text{number of diamonds}]}{[\text{number of broken diamonds}]} \quad (8.2b)$$

For one diamond there are four broken diamonds (i.e., couples of connected pairs of nodes where one node from one side is not connected to one node of the other side).

Again we focus on the local coefficient c_4 , which appears to be one order of magnitude larger compared to that measured in random networks with a power-law degree-distribution — see Fig. 8.5. Therefore, the real socio-semantic network enjoys an abnormally high level of bipartite clustering: many pairs of agents linking together to certain concepts are more likely to share other concepts than in a random network. Note that, as such, the bipartite coefficient is a measure of a very local kind of structural equivalence (quantifying a “limited structural equivalence” restricted to groups of size 2).

³Obviously, many other shapes could also be worth considering; we focused on this one because it is very basic yet insightful.




Figure 8.5: *Left*: $c_3(k)$ as a function of node degree — c_3 is close to 1 and slightly decreasing. *Right*: $c_4(k)$, very slightly decreasing, with an average value of $3.7 \cdot 10^{-4}$, to be compared to $\simeq 3 \cdot 10^{-5}$ in random scale-free networks (Lind *et al.*, 2005).

8.4 Epistemic community structure

A key high-level stylized fact characteristic of epistemic networks is the particular distribution of ECs obtained through GLs, as presented in the previous part. An adequate epistemic network model should ultimately yield the same EC profile as in the real-world, which shows a significantly larger proportion of high-size ECs — see Fig. 8.6.

Semantic distances Besides, just as we observed the bipartite clustering between agents and concepts, we may want to know whether agents in the network are semantically close to each other. Likewise, and more specifically, in which manner are they semantically close to their social neighborhood? To this end, we need to introduce a semantic distance. By semantic distance we mean a function of a dyad of agents that enjoys the following properties: (i) decreasing with the number of shared concepts between the two agents, (ii) increasing with the number of distinct concepts, (iii) equal to 1 when agents have no concept in common, and to 0 when they are linked to identical concepts. Given $(s, s') \in \mathbf{S}^2$, we build a semantic distance $\delta(s, s') \in [0; 1]$ satisfying the previous properties:⁴

$$\delta(s, s') = \frac{|(s^\wedge \setminus s'^\wedge) \cup (s'^\wedge \setminus s^\wedge)|}{|s^\wedge \cup s'^\wedge|} \quad (8.3)$$

Note that this kind of distance, based on the Jaccard coefficient (Batagelj & Bren, 1995), has been extensively used in Information Retrieval, as well as recently for link formation prediction in (Liben-Nowell & Kleinberg, 2003) — however, we

⁴Recall that s^\wedge denotes the set of concepts s is linked to (cf. Part I).




Figure 8.6: Raw distribution of epistemic community sizes, in an empirical GL calculated for a relationship between a random sample of 250 agents, and 70 concepts.

need not focus on this particular similarity measure.

Discretizing δ Written in a more explicit manner, with $s^\wedge = \{c_1, \dots, c_n, c_{n+1}, \dots, c_{n+p}\}$ and $s'^\wedge = \{c_1, \dots, c_n, c'_{n+1}, \dots, c'_{n+q}\}$, we have $\delta(s, s') = \frac{p+q}{p+q+n}$; n and p, q representing respectively the number of elements s^\wedge and s'^\wedge have in common and have in proper. We also verify that if $n = 0$ (disjoint sets), $\delta(s, s') = 1$; if $n \neq 0, p = q = 0$ (same sets), $\delta(s, s) = 0$; and if $s^\wedge \subset s'^\wedge$ (included sets), $\delta(s, s') = \frac{q}{q+n}$. It is moreover easy though cumbersome to show that $\delta(\cdot, \cdot)$ is also a *metric* distance.

As δ takes real values in $[0, 1]$ we need to discretize δ . To this end, we use a uniform partition of $[0, 1[$ in $I - 1$ intervals, to which we add the singleton $\{1\}$. We thus define a new discrete distance d taking values in $\mathcal{D} = \{d_1, d_2, \dots, d_I\}$ such that: $\mathcal{D} = \left\{ \left[0, \frac{1}{I-1}[, \left[\frac{1}{I-1}, \frac{2}{I-1}[, \dots, \left[\frac{I-2}{I-1}, 1[, \{1\} \right\}$.

Then, we look at the distribution of semantic distances in the network, both on a global scale (by computing the distribution for all pairs of agents) and on a more local scale (by carrying the computation for pairs of already-connected agents only). Results are shown on Fig. 8.7, and suggest that while similar nodes are usually rare in the network, the picture is radically different when considering the social neighborhood: acquaintances are at a strongly closer distance.⁵

⁵Although part of the phenomenon is biased by the fact that co-authors receive *by definition* the same concepts when they write an article (especially for distance 1, which is obviously over-represented because of, at first, co-authors who write only one paper), this fact alone is not sufficient to explain the distribution of distances restricted to the social neighborhood.




Figure 8.7: *Left:* Distribution of semantic distances on the whole graph. *Right:* Distribution of semantic distance for the social neighborhood of agents only.

Chapter 9

Low-level dynamics

Designing a credible social network morphogenesis model requires to understand both low-level interaction and growing mechanisms, as noted earlier in Sec. 7.2. The aim of the present chapter is thus to show how we design such low-level dynamics λ from empirical data.

9.1 Measuring interaction behavior

Formally, the preferential attachment (PA) is the likeliness for a node to be involved in an interaction with another node with respect to node properties. Existing *quantitative* estimations of PA and subsequent validations of modeling assumptions are quite rare, and are either:

- related to the classical degree-related PA (Barabási *et al.*, 2002; Eisenberg & Levanon, 2003; Jeong *et al.*, 2003; Redner, 2005), sometimes extended to a selected network property, like common acquaintances (Newman, 2001a); or
- reducing PA to a scalar quantity: for instance using direct mean calculation (Guimera *et al.*, 2005), econometric estimation approaches (Powell *et al.*, 2005) or Markovian models (Lazega & van Duijn, 1997; Snijders, 2001).¹

In addition, the extent to which distinct properties correlatively influence PA is widely ignored. Thus, while of great interest in approaching the underlying interactional behavioral reality of social networks, these works may not be able to provide a sufficient empirical basis and support for designing trustworthy PA mechanisms. Yet in this view we argue that the following points are key:

¹Let us also mention link prediction from similarity features based on various strictly structural properties (Liben-Nowell & Kleinberg, 2003), obviously somewhat related to PA.

1. Node degree does not make it all — and even the popular degree-related PA (a linear “rich-get-richer” heuristics) seems to be inaccurate for some types of real networks (Barabási *et al.*, 2002), and possibly based on flawed behavioral foundations, as we will suggest below in Sec. 9.2.1.
2. Strict social network topology and derived properties may not be sufficient to account for complex social phenomena — as several above-cited works insinuate, introducing “external” properties (such as e.g. node types) may influence interaction; explaining for instance homophily-related PA (McPherson & Smith-Lovin, 2001) requires at least to *qualify* nodes with the help of non-structural data. In reference networks, the probability for citing a paper decreases with time, since papers are gradually forgotten or obsolete (Redner, 1998; Dorogovtsev & Mendes, 2000).
3. Single scalar quantities cannot express the rich heterogeneity of interaction behavior — for instance, when assigning a unique constant parameter to preferential interaction with closer nodes, one misses the fact that such interaction could be significantly more frequent for very close nodes than for loosely close nodes, or discover that for instance it might be quadratic instead of linear with respect to the distance, etc.
4. Often models assume properties to be uncorrelated which, when it is not the case, would amount to count twice a similar effect;² knowing correlations between distinct properties is necessary to correctly determine their proper influence on PA.

To summarize, it is crucial to conceive PA in such a way that (i) it is a flexible and general mechanism, depending on relevant parameters based on both topological and non-topological properties; (ii) it is an empirically valid function describing the whole scope of possible interactions; and (iii) it takes into account overlapping influences of different properties.

In order to measure PA, we now have to distinguish between (i) single node properties, or *monadic* properties (such as degree, age, etc.) and (ii) node dyad properties, or *dyadic* properties (social distance, dissimilarity, etc.). When dealing with monadic properties indeed, we seek to know the propension of some kinds of nodes to be involved in an interaction. On the contrary when dealing with dyads, we seek to know the propension for an interaction to occur preferentially with some kinds of couples. Note that a couple of monadic properties can be considered dyadic; for instance, a couple of nodes of degrees k_1 and k_2 considered as a dyad

²Like for instance in (Jin *et al.*, 2001) where effects related to degree and common acquaintances are combined in an independent way.

(k_1, k_2) . This makes the former case a refinement, not always possible, of the latter case.

9.1.1 Monadic PA

Suppose we want to measure the influence on PA of a given monadic property m taking values in $\mathcal{M} = \{m_1, \dots, m_n\}$. We assume this influence can be described by a function f of m , independent of the distribution of agents of kind m . Denoting by “L” the event “attachment of a new link”, $f(m)$ is simply the conditional probability $P(L|m)$ that an agent of kind m is involved into an interaction.

Thus, it is $f(m)$ times more probable that an agent of kind m receives a link. We call f the *interaction propensity* with respect to m . For instance, the classical degree-based PA used in BA and subsequent models — links attach proportionally to node degrees (Barabási & Albert, 1999; Barabási *et al.*, 2002; Catanzaro *et al.*, 2004) — is an assumption on f equivalent to $f(k) \propto k$.

$P(m)$ typically denotes the distribution of nodes of type m . The probability $P(m|L)$ for a new link extremity to be attached to an agent of kind m is therefore proportional to $f(m)P(m)$, or $P(L|m)P(m)$. Applying the Bayes formula yields indeed:

$$P(m|L) = \frac{f(m)P(m)}{P(L)} \quad (9.1)$$

with $P(L) = \sum_{m' \in \mathcal{M}} f(m')P(m')$.

Empirically, during a given period of time ν new interactions occur and 2ν new link extremities appear. Note that a repeated interaction between two already-linked nodes is not considered a new link, for it incurs acquaintance bias. The expectancy of new link extremities attached to nodes of property m along a period is thus:

$$\nu(m) = P(m|L) \cdot 2\nu \quad (9.2)$$

As $\frac{2\nu}{P(L)}$ is a constant of m we may estimate f through \hat{f} such that:

$$\begin{cases} \hat{f}(m) = \frac{\nu(m)}{P(m)} & \text{if } P(m) > 0 \\ \hat{f}(m) = 0 & \text{if } P(m) = 0 \end{cases} \quad (9.3)$$

Thus $\mathbf{1}_P(m)f(m) \propto \hat{f}(m)$, where $\mathbf{1}_P(m) = 1$ when $P(m) > 0$, 0 otherwise.

9.1.2 Dyadic PA

Adopting a dyadic viewpoint is required whenever a property has no meaning for a single node, which is mostly the case for properties such as proximity, similarity — or distances in general. We therefore intend to measure interaction propension for a dyad of agents which fulfills a given property d taking values in $\mathcal{D} = \{d_1, d_2, \dots, d_n\}$. Similarly, we assume the existence of an essential dyadic interaction behavior embedded into g , a strictly positive function of d ; correspondingly the conditional probability $P(L|d)$. Again, interaction of a dyad satisfying property d is $g(d)$ times more probable. In this respect, the probability for a link to appear between two such agents is:

$$P(d|L) = \frac{g(d)P(d)}{P(L)} \quad (9.4)$$

with $P(L) = \sum_{d' \in \mathcal{D}} g(d')P(d')$.

Here, the expectancy of new links between dyads of kind d is $\nu(d) = P(d|L)\nu$. Since $\frac{\nu}{P(L)}$ is a constant of d we may estimate g with \hat{g} :

$$\begin{cases} \hat{g}(d) = \frac{\nu(d)}{P(d)} & \text{if } P(d) > 0 \\ \hat{g}(d) = 0 & \text{if } P(d) = 0 \end{cases} \quad (9.5)$$

Likewise, we have $\mathbf{1}_{P(d)}g(d) \propto \hat{g}(d)$.

9.1.3 Interpreting interaction propensions

Shaping hypotheses The PA behavior embedded in \hat{f} (or \hat{g}) for a given monadic (or dyadic) property can be reintroduced as such in modeling assumptions, either (i) by reusing the exact empirically calculated function, or (ii) by stylizing the trend of \hat{f} (or \hat{g}) and approximating f (or g) by more regular functions, thus making possible analytic solutions.

Still, an acute precision when carrying this step is often critical, for a slight modification in the hypotheses (e.g. non-linearity instead of linearity) makes some models unsolvable or strongly shakes up their conclusions. For this reason, when considering a property *for which there is an underlying natural order*, it may also be useful to examine the cumulative propension $\hat{F}(m_i) = \sum_{m'=m_1}^{m_i} \hat{f}(m')$ as an estimation of the integral of f , especially when the data are noisy (the same goes with \hat{G} and \hat{g}).

Correlations between properties Besides, if modelers want to consider PA with respect to a collection of properties, they have to make sure that the properties are uncorrelated or that they take into account the correlation between properties: evidence suggests indeed that for instance node degrees depend on age. If two distinct properties p and p' are independent, the distribution of nodes of kind p in the subset of nodes of kind p' does not depend on p' , i.e. the quantity $\frac{P(p|p')}{P(p)}$ must theoretically be equal to 1, $\forall p, \forall p'$. Empirically, it is possible to estimate it through:

$$\begin{cases} \widehat{c}_{p'}(p) = \frac{P(p|p')}{P(p)} & \text{if } P(p) > 0 \\ \widehat{c}_{p'}(p) = 0 & \text{if } P(p) = 0 \end{cases} \quad (9.6)$$

in the same manner as previously. For computing the correlation between a monadic and a dyadic property, it is easy to interpret $P(p|d)$ as the distribution of p -nodes being part of a dyad d .

Essential behavior As such, calculated propensions do not depend on the distribution of nodes of a given type at a given time. In other words, if for example physicists prefer to interact twice more with physicists than with sociologists but there are three times more sociologists around, physicists may well be apparently interacting more with sociologists. Nevertheless, \hat{f} remains free of such biases and yields the “baseline” preferential interaction behavior of physicists.

However, \hat{f} could still depend on global network properties, e.g. its size, or its average shortest path length. Validating the assumption that \hat{f} is independent of *any* global property of the network — i.e., that it is an *entirely essential* property of nodes of kind p — would require to compare different values of \hat{f} for various periods and network configurations. Put differently, this entails checking whether the shape of \hat{f} itself is a function of global network parameters.

9.1.4 Activity and events

Additionally, as regards monadic PA, \hat{f} represents equivalently an attractivity or an activity. Indeed, if interactions occur preferentially with some kinds of agents, it could as well mean that these agents are more attractive or that they are more active. If more attractive, the agent will be interacting more, thus being apparently more active. To distinguish between the two effects, it is sometimes possible to measure independently agent activity, notably when interactions occur during *events*, or when interaction initiatives are traceable (e.g. in a directed network).

In such cases, the distinction is far from neutral for modeling. Indeed, when considering evolution mechanisms focused not on agents creating links, but in-

stead on events gathering agents (Ramasco *et al.*, 2004; Guimera *et al.*, 2005), modelers have to be careful when integrating back into models the observed PA as a behavioral hypothesis. Some categories of agents might in fact be more active and accordingly involved in more events, not enjoying more attractivity. This would eventually lead the modeler to refine agent interaction behavior by including both the participation in events and the number of interactions per event, rather than just preferential interactions.

Detailing interaction propensions In other words, for a given property m , this means breaking down interaction propensions into:

- (i) *activity* $a(m)$: the conditional probability of taking part in an event:

$$a(m) = P(E|m) \quad (9.7)$$

where “E” denotes “involvement in an event”;

- (ii) *interactivity* $\iota(m, \cdot)$: the conditional distribution of the number of links during an event, such that:

$$\iota(m, l) = P(L^E = l|m) \quad (9.8)$$

where “ L^E ” denotes the random variable “number of link extremities received in an event”. The interactivity is thus directly linked to the distribution of the size of events in which agents of kind m participate. We denote by $\bar{\iota}(m)$ the mean of $\iota(m, \cdot)$:

$$\bar{\iota}(m) = \sum_{l \in \mathbb{N}} (\iota(m, l) \cdot l) \quad (9.9)$$

Hence, we now have:

Proposition 5. *f is fully decomposable into $\bar{\iota}$ and a:*

$$f(m) \propto a(m)\bar{\iota}(m) \quad (9.10)$$

Proof. $\nu(m)$ is the product of (i) the mean number of link extremities received by a node of kind m per event, and (ii) the number of nodes of kind m involved in events:

$$\nu(m) = \bar{\iota}(m) \cdot P(m|E)\nu^E \quad (9.11)$$

where ν^E is the number of events for a period. Recall from (9.1) & (9.2) that $\nu(m) =$

$2\nu \frac{f(m)P(L)}{P(m)}$, then Eq. 9.11 yields:

$$f(m) = \frac{\nu^E P(L)}{2\nu P(E)} \bar{l}(m) \cdot a(m) \quad (9.12)$$

As ν , ν^E , $P(L)$ and $P(E)$ are constants of m , we have $f(m) \propto a(m)\bar{l}(m)$. \square

For instance, very active agents (large $a(m)$) involved in events with few participants (small $\bar{l}(m)$) could appear to have the same interaction propensity f as moderately active agents (mean $a(m)$) with a moderate number of co-participants (mean $\bar{l}(m)$). Consequently, when considering monadic PA, event-based modeling requires the knowledge of both a and \bar{l} , for f alone would *not* be in general a sufficient characterization of agent interaction behavior.

9.2 Empirical PA

We now apply the above tools to the study of the epistemic network. We examine therein particularly two kinds of PA: (i) PA related to a monadic property: the node degree; and (ii) PA linked to a dyadic property: semantic distance d , rendering homophily, i.e. the propensity of individuals to interact more with similar agents. In order to have a non-empty and statistically significant network for computing propensities, we first build the network on an initialization period of 7 years (from 1997 to end-2003), then carry the calculation on new links appearing during the last year; 1, 000 new articles appear during the last year.

9.2.1 Degree-related PA

We use Eq. 9.3 and consider the node degree k as property m (thus $\mathcal{M} = \mathbb{N}$): in this manner, we intend to compute the real slope $\hat{f}(k)$ of the degree-related PA and compare it with the assumption “ $f(k) \propto k$ ”. This hypothesis classically relates to the preferential linking of new nodes to old nodes. To ease the comparison, we considered the subset of interactions between a new and an old node.

Empirical results are shown on Fig. 9.1. Seemingly, the best linear fit corroborates the data and tends to confirm that $f(k) \propto k$. The best non-linear fit however deviates from this hypothesis, suggesting that $f(k) \propto k^{0.97}$. However, the confidence interval on this exponent is $[0.6, 1.34]$ thus dramatically too wide to determine the precise exponent, which may be critical. When the data is noisy like in the present situation, since there is a natural order on k it is very instructive to plot the cumulated propensity $F(k) = \sum_{k'=1}^k \hat{f}(k')$ on Fig. 9.1. In this case, the best non-linear fit for \hat{F} is $\hat{F}(k) \propto k^{1.83} \pm 0.05$, confirming the slight deviation from a strictly linear preference which would yield k^2 .




Figure 9.1: *Left:* Degree-related interaction propensity \hat{f} , computed on a one-year period, for $k < 25$ (confidence intervals are given for $p < .05$); the solid line represents the best linear fit. *Right:* Cumulated propensity \hat{F} . Dots represent empirical values, the solid color line is the best non-linear fit for $\hat{F} \sim k^{1.83}$, and the gray area is the confidence interval.




Figure 9.2: *Left:* Activity $a(k)$ during the same period, in terms of articles per period (events per period) with respect to agent degree; solid line: best linear fit. *Right:* Cumulated activity $A(k) = \sum_{k'=1}^k a(k')$, best non-linear fit is $k^{1.88} \pm 0.09$.

Rich-work-harder. This precise result is not new and tallies with existing studies on degree-related PA (Newman, 2001a; Jeong *et al.*, 2003). Nevertheless, we wish to stress a more fundamental point concerning this kind of PA. Indeed, considerations on agent activity lead us to question the usual underpinnings and justifications of PA related to a monadic property. Regarding in particular degree-related PA, we question the “*rich-get-richer*” metaphor describing rich, or well-connected agents as more attractive than poorly connected agents, thus receiving more connections and becoming even more connected.³

When considering the activity of agents with respect to k , that is, the number of events in which they participate (here, the number of articles they co-author), “rich” agents are proportionally more active than “poor” agents (Fig. 9.2), and thus obviously encounter more interactions. It might thus well simply be that richer agents work harder, not are more attractive; the underlying behavior linked to preferential interaction being simply “proportional activity.”⁴

While formally equivalent from the viewpoint of PA measurement, the “*rich-get-richer*” and “*rich-work-harder*” metaphors are not behaviorally equivalent. One could choose to be blind to this phenomenon and keep an interaction propensity proportional to node degree. On the other hand, one could also prefer to consider higher-degree nodes as more active, assuming instead that the number of links per event is degree-independent and that agents do neither *prefer*, nor *decide* to interact with famous, highly connected nodes; a hypothesis supported by the present empirical results. These two viewpoints, while both consistent with the observed PA, bear distinct implications for modeling — especially in event-based models. More generally, such feature supports the idea that events, not links, are the right level of modeling for social networks (Sec. 9.1.4) — with events reducing in some cases to a dyadic interaction.

9.2.2 Homophilic PA

Homophily conveys the idea that agents prefer to interact with other resembling agents. Here, we assess the extent to which agents are “homophilic” by using the inter-agent semantic distance introduced in Sec. 8.4, thus using the socio-semantic network. As we previously underlined, the point is not to focus on this particular similarity measure: rather, we wish to show that simple properties non-related to the strict social structure may also strongly influence interaction behavior in the social network.

³“(…) the probability that a new actor will be cast with an established one is much higher than that the new actor will be cast with other less-known actors” (Barabási & Albert, 1999).

⁴Moreover, if we assume that k is an accurate proxy for agent activity (i.e. a behavioral feature), and if the number of coauthors does not depend on k (which is actually roughly the case in this data, see Fig. 9.8), then observing a quasi-linear degree-related PA should not be surprising.




Figure 9.3: *Left*: Homophilic interaction propensity \hat{g} with respect to $d \in \mathcal{D} = \{d_1, \dots, d_{15}\}$ (thick solid line) and confidence interval for $p < .05$ (thin lines). The y-axis is in log-scale. *Right*: Because of the two extremas it seems natural to try to fit the graph using a third-degree polynomial: $\log(g(d)) = 4.7 \cdot 10^{-3}d^3 - 9.6 \cdot 10^{-2}d^2 + 2.2 \cdot 10^{-1}d - 1.76$ (dashed line). Simpler is a linear fit on the log-log graph: $\log(g(d)) = -0.29d$ (solid line). The original empirical data is plotted here with dots — obviously, many other fitting functions are conceivable.

We obtain an empirical estimation of homophily with respect to this distance by applying Eq. 9.5 on d , with $I = 15$. The results for \hat{g} are gathered on Fig. 9.3 and show that while agents favor interactions with slightly different agents (as the initial increase suggests), they still very strongly prefer similar agents, as the clearly decreasing trend indicates (sharp decrease from d_4 to d_{13} , with d_4 being one order of magnitude larger than d_{13} — note also that $\hat{g}(d_1) = 0$ because no new link appears for this distance value). Agents thus display semantic homophily, a fact that fiercely advocates the necessity of taking semantic content into account in the perspective of modeling such networks.

Correlation between degree and semantic distance In other words, the exponential trend of \hat{g} suggests that scientists seem to choose collaborators most importantly because they are sharing interests, and less because they are attracted to well-connected colleagues, which besides actually seems to reflect agent activity. As underlined in Sec. 9.1.3, when building a model of such network based on degree-related and homophilic PA, one has to check whether the two properties are independent, i.e. whether or not a node of low degree is more or less likely to be at a larger semantic distance of other nodes. It appears here that there is no correlation between degree and semantic distance: for a given semantic distance d , the probability of finding a couple of nodes including a node of degree k is the same as it is for any value of d — see Fig. 9.4.




Figure 9.4: Degree and semantic distance correlation estimated through $\hat{c}_d(k) = \frac{P(k|d)}{P(k)}$, plotted here for three different values of d : $d \in \{d_5, d_8, d_{11}\}$, along with $y = 1$.

9.2.3 Other properties

Specifying the list of properties is nevertheless a process driven by the real-world situation *and* by the stylized facts the modeler aims at rebuilding and considers relevant for morphogenesis. While we examined a reduced example of two significant properties (node degree and semantic distance), measuring PA relatively to other parameters could actually be very relevant as well — such as PA based on social distance, common acquaintances, etc. However, the goal is also to exhibit behaviorally credible as well as non-overlapping, non-correlated properties, if possible. In this respect, neither common acquaintances nor social distance seem to be good candidates.

Let us nonetheless examine social distance in more details. The social distance l between two agents is the length of the shortest path linking them in the social network, with $l = \infty$ when no path exists.⁵ Obviously, l is also a dyadic parameter. The rationale for considering this property is that one may expect that agents at a short social distance are more likely to interact. The shorter the distance, the more likely two agents are to get gathered in a common event: if they have at least one common acquaintance (distance 2), if there is a pair of acquaintances of each agent who know each other (distance 3), etc. Notice that agents at distance 1 are already neighbors so, as regards our definition of a “new link”, there are no new links between pairs at distance one.

The interaction propension h with respect to social distance is plotted on Fig. 9.5, and reveals a strong PA towards “closer” agents. However, social distance is corre-

⁵The algorithm to compute shortest path length in an unweighted graph principally consists in taking the first vertex, assigning it distance 0, then assigning distance 1 to all neighbors, taking the list of all neighbors, assigning them a distance 2, etc. — this is a special version of Dijkstra’s algorithm (1959) on an unweighted network.




Figure 9.5: Social distance-related interaction propensity \hat{h} with respect to $l \in \mathcal{L} = \{1, 2, \dots, 7, 8, \infty\}$ (thick solid line) and confidence interval for $p < .05$ (thin lines). The y-axis is in log-scale. *Inset:* Fit of \hat{h} (empirical data, *dots*), using either an affine function ($\log(\hat{h}(l)) = -.65 - .60l$, *solid line*) or an inverse function ($\log(\hat{h}(l)) = -4.7 + 4.6/l$, *dashed line*). This second function, apparently better, suggests that there is a limit in the decrease of the propensity: after some distance, the preference is the same for everybody.

lated at least to degree (Newman, 2001c) (nodes of degree 0 for instance are always at an infinite distance of everyone in the social network) and in this respect a reductive parameter: two agents at distance 2 are certainly more likely to interact if they have a lot of common acquaintances than just one, and social distance does not distinguish between the two phenomena.⁶ By contrast, we are sure from Sec. 9.2.2 that degree and semantic distance are independent.

9.2.4 Concept-related PA

Yet, we may also wonder how concepts are chosen: for instance, like for social interactions, are well-connected concepts used more often in articles, thus ‘interacting’ with even more authors? It turns out that concepts are present with a frequency proportional to their socio-semantic degree, which is the number of agents who use them, therefore reflecting their popularity — see Fig. 9.6.

⁶In this respect, distances based on random walks could be a good compromise (Gaume, 2004), as this takes into account the fact that two agents are connected through a more or less dense web of common acquaintances in the broad sens (“proxemy”).




Figure 9.6: Cumulated activity of concepts, with respect to their socio-semantic degree $k_{\text{concepts} \rightarrow \text{agents}}$. A non-linear fit yields $A_{\text{concepts}}(k_{\text{c} \rightarrow \text{a}}) \propto k_{\text{c} \rightarrow \text{a}}^{2.19}$, implying a slightly supra-linear activity $a_{\text{concepts}}(k_{\text{c} \rightarrow \text{a}})$, i.e. $\propto k_{\text{c} \rightarrow \text{a}}^{1.19}$.

9.3 Growth- and event-related parameters

These features yield an essential insight on how local interactions occur. Now, in order to complete the description of the way the network grows, studying how events are structured in terms of both authors and concepts is also a crucial information. Regularly, new articles are produced, involving on one side a certain number of authors who have already authored a paper (old nodes) and possibly a fraction of new authors (new nodes), and on the other side, concepts that the authors bring in as well as new concepts.

9.3.1 Network growth

The first step is to determine the raw network growth, in terms of new nodes. How many new events appear, how many new articles are written during each period? Articles gather existing authors as well as new authors around concepts. Since we consider the set of concepts to be fixed *a priori*, new nodes appear in the social network only. The evolution of the size of the social network N_t depends on the number of new nodes per period ΔN_t , with $N_{t+1} = N_t + \Delta N_t$. In turn, there is a strong link between ΔN_t and the number of articles n_t , depending on the fraction of new authors per article.

As we can see on Fig. 9.7, the growth of both ΔN_t and n_t is roughly linear with time. For instance, we can approximate the evolution of n by $n_{t+1} = n_t + n_+$, for a given arithmetic growth rate of n_+ ; every period the number of new articles increases by n_+ . In our case, $n_+ \simeq 96$ ($\sigma \simeq 28$). ΔN and n seem to be linearly correlated, suggesting that the proportion of new authors in all articles is stable




Figure 9.7: For each period, number of articles n_t (blue triangles), number of new agents ΔN_t (red stars), and total size of the social network at the beginning of the period N_t (dark boxes). *Inset:* Comparison functions $(\Delta N_t)^2/N_t$ (dark boxes), n_t^2/N_t (red stars) and $\Delta N_t/n_t$ (blue triangles), modulo a multiplicative constant. All quantities appear to be constant, and linear fits yield respectively $(\Delta N_t)^2 \simeq 490N_t$, $n_t^2 \simeq 96.8N_t$ and $\Delta N_t \simeq 2.25n_t$.

across periods.

9.3.2 Size of events

This leads us to study how articles are structured: in particular, how many agents are gathered in an event, and how many of them are new nodes? As shown on Fig. 9.8, the distribution of the number of agents per article appears to follow roughly a geometric distribution.⁷ On the other hand, the weight of new authors within articles obeys a distribution centered around three modes $\{0, 0.5, 1\}$, suggesting that in most cases either (i) authors are all new, (ii) they are all old, or (iii) half are new & half are old. Since this proportion is stable across periods, n_t is a good indicator of network growth: new articles appear and pull new authors into the network — on average, articles gather 4.4 authors, among which 55% are new, thus $.55 \times 4.4 = 2.42$ new authors, which is close to the coefficient of the best linear fit of ΔN with respect to n : $\Delta N \sim 2.25n$.

Since the size of the network is increased by ΔN in a period, and ΔN here shows a linear behavior, N should exhibit a quadratic growth; which is confirmed by comparing $(\Delta N)^2$ to N as shown on Fig. 9.7 (the same goes for n^2 vs. N). The fact that the number of articles per period linearly increases is however proper to

⁷In addition, the number of coauthors does not depend on node degree, suggesting that more active agents are not working with a different number of collaborators when coauthoring an article (see inset on Fig. 9.8-top): agent interactivity is independent of degree, $\bar{r}(k) = \bar{r}$.

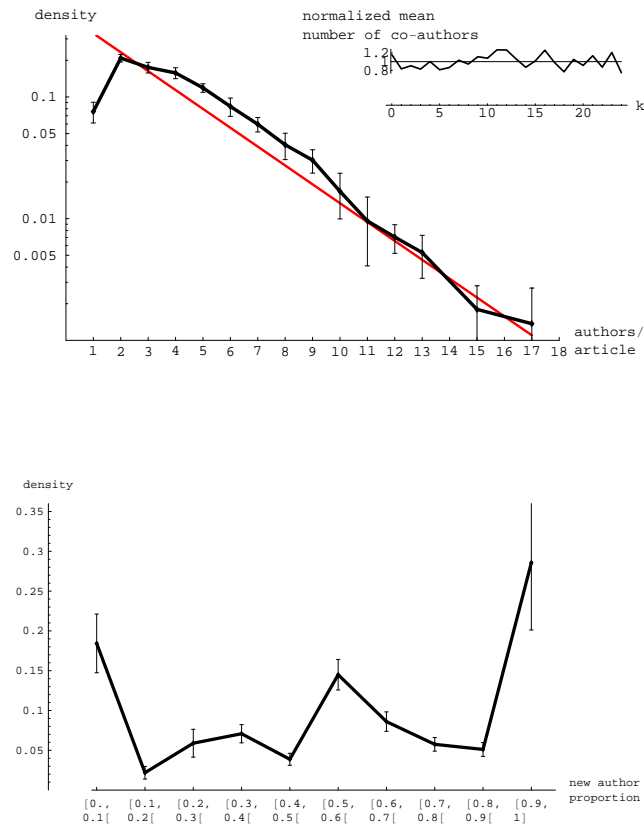


Figure 9.8: *Top*: Distribution of the size of events (black line), averaged on 8 periods 97-04, with confidence intervals for $p < .05$. The mean number of authors is 4.4 ($\sigma = 3.1$), and the best non-linear fit is $\propto \exp^{-\mu n}$ with $\mu = .36 \pm .06$ (red line). The inset shows the mean number of coauthors with respect to degree k , relatively to the global mean number of co-authors: in case of independence, this ratio equals 1. *Bottom*: Proportion of new authors with respect to total authors, averaged on 7 periods (98-04) — the mean proportion is 0.55, but $\sigma = .33$ because of the tri-modal distribution.

the evolution of *this* empirical situation. The evolution of n and N is a consequence of this — this is obviously not the case for all networks: if for instance this field of research were to be abandoned, we would have a decrease of articles, not a linear growth.

9.3.3 Exchange of concepts

Knowing the structure of articles, and how authors are gathered, we now investigate how concepts are chosen. The distribution of the number of concepts is plotted on Fig. 9.9, and could be accurately approximated by a geometric distribution. Besides, while old authors bring a certain proportion of their concepts, some concepts are used for the first time: they do not belong to the intension of authors. The distribution of the proportion of new concepts — *new to the authors* — also shown on Fig. 9.9, makes it possible to distinguish concepts chosen within the intension of authors, from new, unused ones. It has a single mode 0, but is on the whole relatively flat.

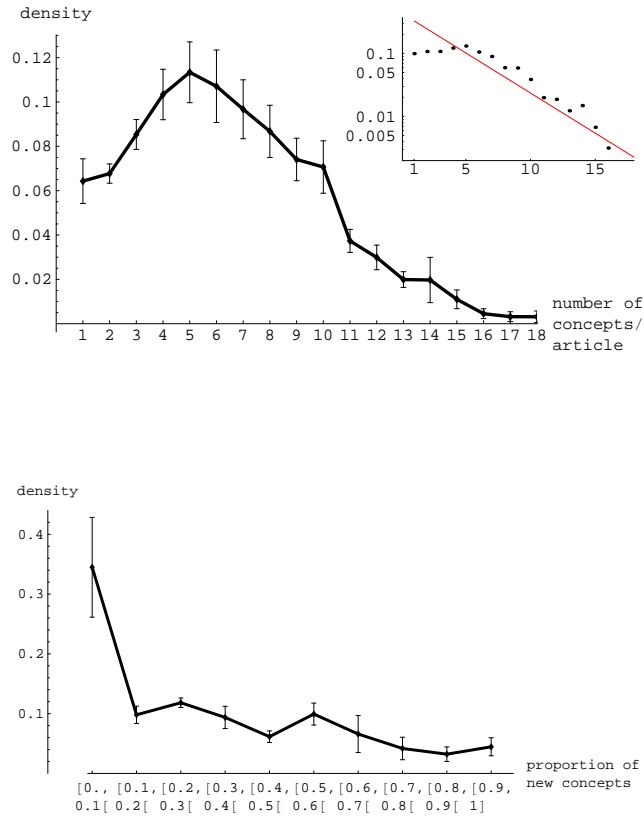


Figure 9.9: *Top*: Distributions of concepts per article — mean: 6.5, $\sigma = 3.6$. In the inset, the solid line represent the best exponential fit, $\propto e^{-\mu n}$ with $\mu = 0.29$. *Bottom*: Distribution of the proportion of new concepts that none of the agents anteriorly used — only for articles where there is at least one old agent. The mean is .32, with $\sigma = .28$.

Chapter 10

Towards a rebuilding model

10.1 Outline

To sum up, the empirical epistemic network of the field “zebrafish” could be described as follows:

- power-law degree distributions from agents to agents and from agents to concepts;
- a high-level of structurally equivalent groups, both because of a high bipartite clustering coefficient and because of a particular EC structure observed through GLs;
- a particular distribution of semantic distances;
- interaction behavior characterized by a preference to interact with similar, well-connected agents (or, equivalently, who are more active), and to use well-connected, popular concepts (or, equivalently, which are more ‘suitable’), in the precise manner outlined in Sec. 9.2;
- a quadratically growing social network because of a constant growth rate of new authors and articles;
- quasi-geometrically distributed numbers of agents per article and concepts per article, with a trimodal distribution for the proportion of new authors, and a unimodal distribution for the proportion of new concepts.

In short, using the empirically-measured low-level parameters (composition of articles and interaction preferences) we aim at designing a reconstruction model able to reconstruct a high-level structure compatible with real-world stylized facts (degree and semantic distance distributions, bipartite clustering and EC structure).

To this end, three crucial modeling features are implemented: (i) event-based network growth, (ii) co-evolution between agents and concepts, and (iii) realistic low-level descriptions, especially regarding interactions.

Respecting PA in n -adic interactions Yet, event-based modeling introduces serious challenges towards accurately implementing PA. In classical dyadic-interaction-based models, where events involve only two agents, it is utmost easy to choose pairs of agents with respect to PA based on a set of uncorrelated properties, monadic or dyadic. This category also covers models where agents make links to a certain number of other agents on a peer-to-peer basis — for instance in the BA model, where new nodes arrive and attach to a given number n of old nodes; this can actually be considered as n dyadic interactions, not a n -adic interaction; at no time sets of more than 2 nodes have to be composed to create links.

On the contrary in n -adic-interaction-based models, where interactions involve n agents *altogether* and thus induce the addition of n -cliques (with links between all pairs of agents), composing the set of agents *while at the same time* respecting interaction propensities for all $[n(n-1)/2]$ links could be an extremely tricky puzzle. In any case, it now appears very dubious to base network growth on simple dyadic interactions: n -adic interactions are simply everywhere. So, how to proceed in this case? Two situations are to be distinguished:

- as regards PA based on a monadic property m , the picture is still easy if \bar{v} is independent of m , since choosing agents with respect to $f(m)$ or $a(m)$ is equivalent. Then agents can be chosen proportionally to $a(m)$, which is nothing else than $P(E|m)$ and PA is obviously respected for all links between pairs of agents.¹ Otherwise, if \bar{v} depends on m , it would be hard to randomly form events which respect both activities and interactivities for all kinds of nodes.

In our case, we observed on Fig. 9.8 that the number of co-authors does not depend on degree, i.e. $\bar{v}(k)$ is a constant. In other words, agents make the same number of links for every event they participate in, whatever their degree is. This is consistent with the previous observation that the degree-based propensity $f(k)$ has the same shape as the activity $a(k)$ (Sec. 9.2.1).

- as regards PA based on a dyadic property d , the picture is quite different: agents must be chosen so that all links between all pairs of agents respect the

¹In particular, this is what *necessarily* happens with dyadic-interaction-based models (where events always gather 2 agents), which still constitute the core of network growth models (cf. detailed list in Sec. 7.2). Such models are credible in networks where events are by definition of size two (e.g. peer-to-peer networks, Internet transmissions, phone calls). Then $\bar{v}(m)$ always equals 1, and agents can be indifferently chosen with respect to a propensity (which is traditionally the case) or to an activity, *because* $\bar{v}(m) = 1$.

alleged dyadic PA. To make it simpler, our answer is to introduce an initial node i (an “initiator”) which in turn chooses all other nodes with respect to a dyadic PA.² The choice of the initiator must obey criteria consistent with interaction behavior; for instance, it needs to be chosen proportionally to agent activity. Then, other nodes are chosen according to (i) activity and (ii) dyadic PA with respect to the initiator.

Still, without any further assumption there is no guarantee that dyadic propensions are respected for links between these other nodes, i.e. between nodes that do not involve the initiator — between agents *around* the initiator. In our case, the fact that δ is a metric distance nonetheless warrants that the semantic distance between any pair of nodes (x, y) remains similar to their respective distance to i : $\delta(x, y) \leq \delta(i, x) + \delta(i, y)$.

10.2 Design

We may now introduce a minimal event-based model of a coevolving epistemic network. Events are articles, made of (i) agents, who are more or less active depending on their degree k , and gather preferentially with respect to their interests — the former being entirely independent of the latter, and (ii) concepts, which are more or less popular, depending on their degree $k_{\text{concepts} \rightarrow \text{agents}}$. The low-level dynamics is thus as follows:

1. *Creating events.* n_t articles are created at each period:

$$n_{t+1} = n_t + n_+ \quad (10.1a)$$

n_+ fixed to 100.³ This makes the number of events close to that of the real network. The set of articles is denoted by \mathcal{A}_t such that:

$$\begin{aligned} \mathcal{A}_t &= \{\mathcal{A}_t(i) \mid i \in \{1, \dots, n_t\}\} \\ \mathcal{A}_t(i) &= (\mathcal{S}_t(i), \mathcal{C}_t(i)) \end{aligned} \quad (10.1b)$$

where $\mathcal{S}_t(i)$ is the author set of the i -th article, and $\mathcal{C}_t(i)$ the concept set.

²Another solution could consist in quantifying propensions of n -adic interaction between n members of a given event with respect to a n -dimensional vector of parameters — that is, a n -adic PA, generalizing further the framework presented hitherto. Yet, this kind of measurement would really not be convenient. On top of that, for most networks — even large ones — it may be rare to get statistically significant estimations for a decent number of n -adic configurations.

³We have to keep in mind that n_+ remains an exogenous parameter of the model, adapted to the situation of a growing network for a growing community.

2. *Defining event sizes.* Author set and concept set sizes follow geometric laws respecting means observed on Fig. 9.8 and Fig. 9.9, respectively, i.e.:

$$\begin{aligned} |\mathcal{S}_t(i)| &\rightsquigarrow \mathcal{G}(1/m_s) \\ |\mathcal{C}_t(i)| &\rightsquigarrow \mathcal{G}(1/m_c) \end{aligned} \quad (10.1c)$$

where m_s (resp. m_c) is the mean number of authors (resp. concepts) per article.

3. *Choosing authors.* New agents within author sets are denoted by $\mathcal{S}'_t(i) \subset \mathcal{S}_t(i)$. Because of the tri-modal distribution (Fig. 9.8), $\mathcal{S}_t(i)$ contains either only new authors, either only old authors, or equally old and new authors, equiprobably. Thus,

$$|\mathcal{S}'_t(i)| = \begin{cases} [P = \frac{1}{3}] & |\mathcal{S}_t(i)| \\ [P = \frac{1}{3}] & \lfloor \frac{1}{2} |\mathcal{S}_t(i)| \rfloor \\ [P = \frac{1}{3}] & 0 \end{cases} \quad (10.1d)$$

If $\mathcal{S}_t(i) > \mathcal{S}'_t(i)$, there is at least one old agent, and the initiator is randomly chosen proportionally to her social network degree k . Then, other old agents of $\mathcal{S}_t(i) \setminus \mathcal{S}'_t(i)$ are picked according to probability $P(L|k, d)$, where k is the degree of the agent to be chosen, and d the semantic distance between her and the initiator — in accordance with empirical measurements, we have:⁴

$$\begin{cases} P(L|k, d) = P(L|k)P(L|d) \\ P(L|k) \propto k \\ P(L|d) \propto \exp(\mu d) \end{cases} \quad (10.1e)$$

with $\mu = -.29$.

Finally, $|\mathcal{S}'_t(i)|$ new nodes are created, and ultimately added to \mathcal{S} .

4. *Choosing concepts.* New concepts are denoted by $\mathcal{C}'_t(i) \subset \mathcal{C}_t(i)$. By new, we mean concepts that no old agent of $\mathcal{S}_t(i)$ uses. These concepts represent a fixed proportion of the article concept set, that is,

$$|\mathcal{C}'_t(i)| = \mu_c |\mathcal{C}_t(i)| \quad (10.1f)$$

where μ_c is the mean proportion of new concepts (see Fig. 9.9).

Thus, concepts are chosen:

⁴We consider that $P(L|k = 0) = P(L|k = 1)$, which is in reasonable agreement with the data (certainly choosing $P(L|k = 0) = 0$ would doom single agents to remain single for their whole life).

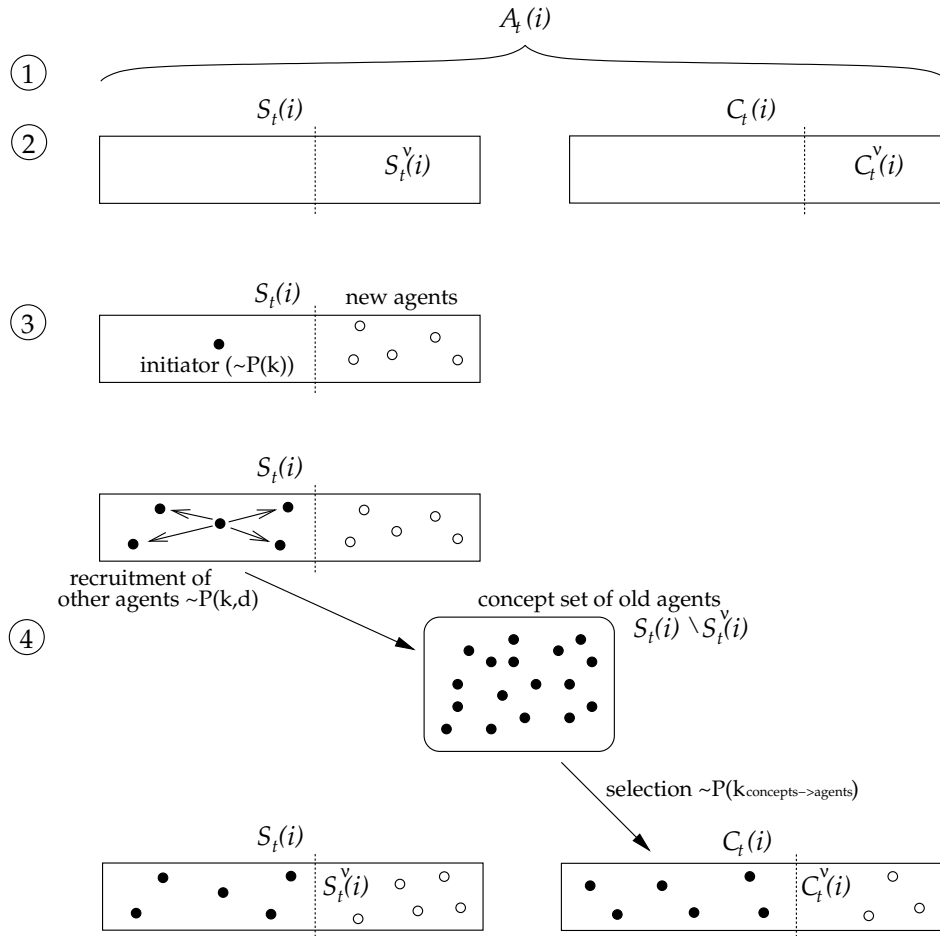


Figure 10.1: Modeling an event by specifying article contents.

- (i) for $C_t(i) \setminus C_t^v(i)$, from the concept set of authors $(\cup_{s \in S_t(i)} s^\wedge)$;
- (ii) for $C_t^v(i)$, from the whole concept set;
- (iii) and for all, randomly proportionally to their degree $k_{concepts \rightarrow agents}$ (stylization of Fig. 9.6).

5. *Updating the network.* When author and concept sets are defined (Fig. 10.1),

the whole network is updated:

$$\left\{ \begin{array}{l} \mathbf{S}_{t+1} = \mathbf{S}_t \cup \bigcup_{i \in \{1, \dots, n_t\}} \mathcal{S}_t^\nu(i) \\ \mathcal{R}_{t+1}^{\mathbf{S}} = \mathcal{R}_t^{\mathbf{S}} \cup \bigcup_{i \in \{1, \dots, n_t\}} \{\mathcal{S}_t(i) \times \mathcal{S}_t(i)\} \\ \mathcal{R}_{t+1}^{\mathbf{C}} = \mathcal{R}_t^{\mathbf{C}} \cup \bigcup_{i \in \{1, \dots, n_t\}} \{\mathcal{C}_t(i) \times \mathcal{C}_t(i)\} \\ \mathcal{R}_{t+1} = \mathcal{R}_t \cup \bigcup_{i \in \{1, \dots, n_t\}} \{\mathcal{S}_t(i) \times \mathcal{C}_t(i)\} \end{array} \right. \quad (10.1g)$$

10.3 Results

We ran the model for 8 periods $t \in \{1, \dots, 8\}$, starting with an empty epistemic network — in other words, the morphogenesis starts *from scratch*. Obviously, periods correspond to years. One hundred new articles were to appear during the first period, with a growth rate of 100 articles per period per period: $n_1 = 100$, $n_+ = 100$. We focus on networks obtained after simulations are completed for 8 periods, and we have a satisfying adequation for every stylized fact, both in shape and in magnitude:

- *Rebuilding network size.* Simulated networks contain 10982 agents on average ($\sigma = 215$, for fifteen runs), agreeing with empirical data.
- *Rebuilding degree distributions.* Results for all four degree distributions are shown on Fig. 10.2, indicating a very good fit — in particular, power-law tails have a similar exponent, with a shape which fits a log-normal distribution similar to that of the empirical case.
- *Rebuilding clustering coefficients.* Clustering coefficients are accurately reproduced, as shown on Fig. 10.3.
- *Rebuilding epistemic community structure.* GLs have been computed for 250-agents samples (see Fig. 10.4), following the protocol of Part I: distributions of EC sizes are close to those of the real network, and exhibit the same effect when compared to the “random case”.⁵ Semantic distances are also correctly rebuilt, see Fig. 10.5.

⁵There is a slight deviation for high-size ECs, which are found in lower number in the simulations than in the real network. This could actually be due to a selection bias where empirical data are *ex post* selected data on a given community (the zebrafish field), where high-size communities are gathered around paradigmatic words (“*develop*”) which the model only partly reproduces.




Figure 10.2: Social, semantic and socio-semantic degree distributions. Simulation results (black dots or thick line) globally fit the empirical data (blue thin line). For instance, the exponent of a power-law fit for social network degree distribution is $\gamma = -3.10 \pm .04$, on average (empirical fit was $\gamma = -3.39$).




Figure 10.3: *Left*: Simulated $c_3(k)$ (dots) compared to the empirical value (blue solid line). *Right*: The same, for $c_4(k)$.




Figure 10.4: Number of ECs with respect to agent set sizes, in GLs computed for samples of 250 agents. Simulation results (thick black line) fit the empirical data (thin blue line). We also computed random “rewired” cases, as we did in Part I (keeping degree distributions on both sides, from agents to concepts and from concepts to agents): as expected, they contain significantly less ECs, by one order of magnitude (thin red line).




Figure 10.5: *Left*: Simulated mean distribution of semantic distances on the whole graph (dots) compared to original empirical data (blue line). *Right*: Same quantities, but computed only for the social neighborhood of each agent. Note the red thin solid line, representing simulations not using homophily.

10.4 Discussion

Hence, *epistemic communities are produced by the co-evolution of agents and concepts*. Not only is the high-level structure accurately reconstructed by our model, but low-level dynamics are consistent as well — this is a not a minor point: rebuilding high-level phenomena remains dubious if the low-level dynamics is incorrect. Truthfulness of descriptions must reach the higher level *as well as* the lower level. In any case, we may still wonder what weight some of our hypotheses bear towards the apparition of high-level phenomena: is our model a minimal model *as regards the stylized facts we selected*?

In particular, consider basic event-based models for social networks — which have become popular very recently among a few other authors as well (Ramasco *et al.*, 2004; Guimera *et al.*, 2005; Peltomaki & Alava, 2005) — that simply rest on n -adic events instead of dyadic interactions and that do not even specify any kind of PA. Yet, these models lead to scale-free distributions and high one-mode clustering coefficients. These results suggest that PA is not required to rebuild degree distributions and c_3 , by contrast to dyadic-interaction-based models (such as BA model).

Recall that our model features (i) event-based modeling, (ii-a) degree-related preferential attachment (or activity) for the choice of agents and (ii-b) for concepts, and (iii) homophily of agents. Are the high-level stylized facts still reproduced if we loosen some of these hypotheses? Since many combinations of simplified models are envisageable, we only examine what happens when relaxing one hypothesis at a time; and sum up the results hereafter.

1. *Relaxing social-degree-based PA*. Only agent degree distributions change (from agents to agents and from agents to concepts), with a different power-law fit exponent ($\gamma = 2.48$ for the social network without this kind of PA, vs. 3.39 with it — the degree distribution is thus “flatter”, which is consistent with the suppression of the accumulative effect of this PA).
2. *Relaxing semantic-degree-based PA*. Here, reconstruction of both EC structure and semantic distance distribution fails. The effect of concept popularity seems central to the emergence of epistemic communities.
3. *Relaxing homophily-based PA*. This is certainly the most surprising result: the only change concerns the semantic distance distribution for the social neighborhood (see Fig. 10.5-right) — yet, this change is slim, especially as regards a feature that has such a heterogeneous impact (recall that the homophilic propensity is exponential).
4. *Relaxing event-based modeling*. This hypothesis is at the core of the model, so

revisiting it may require to strongly reshape the whole model. Let us only fix the fact that $|\mathcal{S}_t(i)| = 2$, which amounts to classical dyadic interactions — all other mechanisms remain unchanged. Then, degree distributions do not enjoy the log-normal shape and are only scale-free; which is unsurprising from (Barabási & Albert, 1999).⁶ Also, clustering coefficients are not reproduced (which is also unsurprising (Ramasco *et al.*, 2004) and consistent with the fact that a high c_3 is simply due to clique addition). Thus, relaxing event-based modeling creates empirical inconsistencies *even for the simplest* topological criteria.

⁶Yet, any constant number of authors per article ($|\mathcal{S}_t(i)| = c$) also leads to a very particular degree distribution, contrarily to what (Guimera *et al.*, 2005) found. For other values of $c > 2$, by definition social network degree distributions are likely to be biased around multiples of $(c - 1)$ — especially for low degrees.

Conclusion of Part II

The main achievement of this part has been to micro-found the particular community structure that we highlighted in Part I. We investigated the formation of an emerging scientific community, that of the “zebrafish”, considered as a social process of knowledge building and community organization. Using real-world observations, we asked whether we could in turn reconstruct artificially the evolution of this scientific field, through the lens of selected stylized facts deemed relevant for this epistemological task.

We assumed that modeling agents co-evolving with concepts was enough to micro-found the evolution of this social complex system. In other words, the social constitution, arrangement, configuration, manipulation and reconfiguration of concepts was assumed to account for most of the scientific field structure. We had thus to design a low-level dynamics λ consistent with empirical data, and adequately rebuilding η^e , through P . To this end, after outlining the kind of stylized facts to be reconstructed, we needed to create tools enabling the estimation, from past data, of the interaction and growth processes at work in the epistemic network. Only thereafter could we hope for a realistic, descriptive model of the dynamic co-evolution of agents and concepts, and the resulting structure.

We have thus argued for an empirical stance in designing model hypotheses, although this attitude can often prohibit analytical solutions and compel to the use of simulation-based proofs. *In fine*, introducing credible empirically-based hypotheses would help attract really more social scientists into this promising field. Social scientists are usually not seeking *normative* models. More specifically, in the search for hypotheses eager to explain a given “high-level” phenomenon, scientists have to make inductions on low-level features which reconstruct the phenomenon. We suggest that it is eventually essential to know whether the alleged low-level dynamics is empirically grounded too — even if the model reproduces the desired stylized facts, and even if the hypotheses do not look *ad-hoc* (like for instance introducing scale-free preferences to rebuild scale-free networks). Normative models are certainly nice, but not necessarily useful towards a descriptive task.

In particular, quantifying interaction processes plays here a crucial role — heterogeneous interaction behaviors are indeed the cornerstone of many recent social network formation models. Preferential attachment (PA), which is the common way of designating this heterogeneity, is obviously a robust method to avoid the classical random graph model. PA was established by the success of a pioneer model (Barabási & Albert, 1999) rebuilding a major stylized fact of empirical networks, the scale-free degree distribution. However, while it has subsequently been widely used, generally few authors attempt to check or quantify the rather arbitrary assumptions on PA. Therefore, we designed measurement tools yielding a comprehensive description of interaction behaviors with respect to any kind of property, structural or not. In addition to epistemic networks, this framework could also be easily applied to any other kind of network, especially non-growing networks — likewise, a whole class of empirically-based morphogenesis models can be designed (Boguna & Pastor-Satorras, 2003; Cohendet *et al.*, 2003). This kind of hindsight on the notion and status of PA should be useful even for normative models.

The final success of the reconstruction gives full credit to the claim of the present thesis: **the structure of knowledge communities is at least produced by the co-evolution of agents and concepts**. Yet, we also argue that such co-evolution may still depend on exogenous parameters. We can indeed imagine that various low-level measurements (size of groups, interaction behavior, growth rate, etc.) would be different in other research groups, other epistemic areas, or other eras. Take for instance the growth of the field: how comes that there is such an interest in the zebrafish? Practical reasons can be put forward: it is a translucent vertebrate, quickly developing, sufficiently close to human, very helpful for many more fields other than embryology. But all of this is proper to the *contingent* nature of the zebrafish. Later, a cure for cancer could be found from the study of the zebrafish, likely to pull in a large number of scientists; *or not*: this discovery depends on *unpredictable* properties of the zebrafish itself. We strongly doubt that these features could be endogenized in any model.

More generally, the uncertainty on novelty and new knowledge (new concepts as well as new usage of old concepts) appearing in the social complex system is not truth-related uncertainty: it is not something which is already-known, which may happen or not, and which is easily substitutable by a probability. Rather, it is a radically different uncertainty, one on the *ontology* (Lane & Maxfield, 2005): “what ontology will agents dispose of in the future?” Epistemologists have long been interested in exploring the justification of new ideas, but few attempted to explain how discoveries occur. In such cases, random intuition (“lucky guesses”) and induction are often called on. Some authors on the contrary argue that the discovery of new knowledge is rooted in already-existing knowledge (Gigeren-

zer, 2003): novel reinterpretations of existing notions and tools have an innovative feedback onto theories and concepts. But here too, we cannot predict the way tools will be reinterpreted. In both situations, we still have to cope with ontological irreducibility: a model cannot express and yield anything *newer* than what is already specified by the language and the grammar of the model, which are closed (Chavalarias, 2004, p.257).

In any case, we must therefore keep in mind that real-world epistemic networks are not closed. In our model, we decided to keep some things exogenous: we had for instance a fixed growth rate n_+ and a fixed set of *a priori* equivalent concepts C . In reality, new topics can arrive in the system — either through items that are not represented in the model (like conferences, news (Gruhl *et al.*, 2004)), underlining the problem of boundary specification (Laumann *et al.*, 1989); or from phenomena that are simply unpredictable (like the cure for cancer, cf. *supra*), for which modeling is most likely to fail. Let us mention in particular two modeling methods that could be proposed to account for new knowledge creation: (i) innovation is modeled by a random probabilistic increase in the amount of knowledge, which is thereby assumed to be quantifiable, monotonic, and whose nature is fixed (e.g. in (Cowan *et al.*, 2002)); (ii) innovation is a generative process, producing new items from already-existing items; for instance Lane (1993) proposed λ -calculus as a way to generate truly novel objects, generally thanks to a chaotic process — such generative processes however could hardly be considered realistic, even if they are indeed undecidable and unpredictable, hence compatible with ontological uncertainty (which probabilistic models are not).

Hence and more broadly, the potential dependence on undecidable exogenous parameters leads us to moderate the claim of our thesis: whereas the reconstruction has obviously proven to be a success, within a given time-period and all its particularities, it is nonetheless likely that other processes *in which the epistemic network is immersed* could also play a significant role. As such, under the provision that such parameters are stable for the considered time-scale, we clearly demonstrated that the reconstruction of the dynamics of a social complex system is within reach.

Part III

Coevolution, Emergence, Stigmergence

Summary of Part III

In this part, we make an epistemological point that provides a significant insight on how to rebuild a social complex system. After detailing different attitudes towards appraising the relationships between levels of description, we argue that distinct levels are merely *distinct observations on a process*. We then present implications on reconstruction methodology and complex system modeling, and particularly emphasize the role of level design in making sound distinctions among objects. We distinguish the special case of systems of agents producing artefacts which in turn have an effect onto them, a feature shared by many social systems.

Introduction of Part III

"(...) because I know that you are a part of Humanity, of which I am also a part, and that you partly take part in the part of something which is also a part and of which I am also in part a part, together with all the particles and parts of parts, of parts, of parts, of parts, of parts... Help! Oh, confounded parts! Oh, bloodthirsty, nightmarish parts, you've grabbed me once again, is there no escaping you, hah, where can I find shelter, what am I to do?"

Ferdydurke, Witold Gombrowicz.

In this final part, we wish to make an epistemological point that should provide a crucial methodological insight on social complex system modeling. So far, we have proven that epistemic networks are the result of low-level interactions of agents co-evolving with concepts. To do so, we have appraised this socio-semantic complex system both (i) starting from disciplines & community structure, and looking at how this may be expressed in terms of agents and concepts, exhibiting a valid "*P*" (Part I); and (ii) using low-level dynamics of epistemic networks to reconstruct high-level phenomena (Part II). As such, we filled the explanatory gap between the lower level of agents & concepts and the higher level of epistemological descriptions. We now wish to investigate the epistemology of our approach, and suggest broader implications on social complex system modeling. In order to do so, we will focus on the status of the different levels of description, the subsequent relationships they may entertain, and the modeling methodology required to give an account of these relationships. **We will argue that modeling social complex systems tends to require the introduction of co-evolutive frameworks at the lower level** of the kind we presented here. More generally, we argue that some high-level phenomena cannot be explained without a fundamental *viewpoint change* in not only low-level dynamics but also in the *design* of low-level objects themselves. In other words, it may be important to reconsider (and sometimes differentiate) objects at a given level in order to achieve a successful reconstruction. Emphasizing *level design* is particularly insightful in situations where structures created by a level exhibit an efficient causal feedback on this level. Surprisingly, these cases do not involve downward causation, but simply relate to causation of *a priori* distinct objects onto each other, or *coevolution* of phenomena.

The outline of this part is as follows: in Chap. 11 we suggest that distinct lev-

els, considered as phenomena of a unique underlying process, only exist to the observer and as such may still yield overlapping, redundant and thus correlated information about the process (Bonabeau & Dessalles, 1997; Gershenson & Heylighen, 2003; Bitbol, 2005). Chapter 12 presents meaningful implications on modeling, and highlights a few yet essential methodological points required for complex system modeling. In Chapter 13, we support the idea that while levels are often simply different aspects of a process, objects could still be usefully differentiated to describe certain kinds of causality between phenomena: for instance, agents produce artifacts that in turn influence them, with no downward causation. The notion of “emergence” is consequently enriched by the concept of “*stigmergence*” of artifacts. We conclude that co-evolution is a central feature of socio-semantic complex systems.

Chapter 11

Appraising levels

The concern of any scientific field is to describe certain kinds of objects, along with the regularities that govern them. The global picture of scientific research is subsequently made of disciplines focused on particular levels of description: physics is concerned with fields and particles, biology with cells and living organisms, social sciences with agents and institutions. Often, a level can be considered to “*rely on*” more fundamental levels — for instance, agents are living organisms, organisms are “made of” cells, cells are “made of” molecules. These notions usually translate in terms of “whole/part” relationships.

Modern science, and complex system science in particular, has also been taking this conception in a reverse, compositionalist direction: items at some level are organized systemically and *compose* higher-level objects — higher in size, because they are made of at least one entity and, often, higher in inertia (i.e. slower time-scale). For example, molecules build up cells, cells build up organisms, which build up agents, and so on. Like our epistemic network model, an important associated challenge is the reconstruction of high-level phenomena through the iterated, cumulated interplay of low-level objects: complex scientists dream to rebuild high-level descriptions from low-level ones. Thus they would bridge explanatory gaps between levels and cancel out separations between scientific fields. To this end, investigating the nature of levels of description becomes a crucial topic — especially addressing the two following key questions: (i) how to appraise different levels? (ii) how to assess their links and potential mutual influence upon each other? We also indicate why this attitude leads to reconsider the notions of upward and downward causation — namely, a level having a causally efficient influence on other levels.

11.1 Accounting for levels

In order to appraise the nature of levels, as mentioned above, several attitudes are available. Classical answers include dualism, reductionism and, as a tentative bridge between these two extremes, emergentism, where higher levels are supposed to emerge from lower levels. Here, we review these stances and present their caveats, notably dismissing the idea that levels exist as entities, and suggesting instead that they are merely observations of a single process — as such, distinct aspects, various phenomena of a same underlying “*x*.”

Let us recall the two most classical positions that could be first suggested:

Definition 12 (Dualism). *Dualism is a position for which different levels correspond to different entities, and have a proper reality by themselves.*

Thus in the dualist position, different levels must be appraised through different means and enjoy distinct realms. Causality happens at all levels. Even if one can for instance describe the cells that compose the body, the body is supposed to enjoy a substantial reality by itself that cannot be explained in terms of the lower level, and accordingly a proper causal efficiency — this amounts, for instance, to vitalism.

Definition 13 (Reductionism). *Reductionism states that all phenomena can be explained, computed and rebuilt from the lower level, up to higher levels.*

Opposite to dualism, the reductionist viewpoint denies that higher levels exist by themselves: they are at best convenient macroscopic descriptions. Here, only the lower level enjoys reality and causal efficiency. This eventually amounts to physicalism: physical entities and laws are sufficient to explain the entire world, at least in theory.¹

11.2 Emergentism

These two conflicting positions nevertheless exhibit some weaknesses. Apart from its unconvincing non-materialistic aspects (Papineau, 2001), the dualist viewpoint eventually amounts to pluralism, with as many ontologies as there are levels. Worse, it is in fact a *subjective* pluralism, because conceptions of levels mostly depend on a quite subjective if not arbitrary ontology.² How could levels created by

¹(Bickhard & Campbell, 2000) “Everything else is epiphenomenal to that, and can be eliminatively reduced to it — perhaps with the caveat of the cognitive limitations of human beings to handle the complexities required. In this cognitive view, higher levels are necessary considerations only because of their relative cognitive simplicity for humans, not for any metaphysical or even physical reasons.”

²As Emmeche et al. (2000) observe, “Our methods for making such distinctions [of primary levels] are of course dependent on the historical development of scientific theories and disciplines.”

scientists be real entities, especially when considering the multiplicity of levels at stake (physical, chemical, biological, individual, social, etc.)?

On the other hand, it is unclear whether reductionism allows the rebuilding of the whole world and its different levels. In this respect, it appears sometimes unlikely that theories on a given level could be reduced to an applied, iterated version of lower-level theories (Anderson, 1972; Laughlin & Pines, 2000; Lane, 2005). Practical reasons (computing the behavior of more than a handful of particles proves quickly to be impossible) as well as less practical reasons (such as Anderson's example of nuclei whose spherical shape is due to an infinite approximation of lower-level particle properties) suggest that "*the Theory of Everything is not even remotely a theory of every thing*" (Laughlin & Pines, 2000).

While the dualist position is based on the *a priori* existence of several levels, the reductionist position actually *eliminates* the higher levels to the benefit of the lowest level.³ These two stances are strikingly contradictory, and the tension is particularly disturbing when one dismisses dualism but still wants to consider higher levels to be irreducible, granting them some reality.

Bridging the gap The emergentist position is an attempt to reconcile both views, by assuming *emergence*. The point is to bridge the possible failures of reductionism: the higher level is not reducible, the whole is more than the sum of its parts, even in theory; but it is physically grounded so it needs to *emerge* from the lower level. No dualism is supposed *a priori*, but the cumulated, aggregated action of small objects somehow leads to the *emergence of novel* higher-level objects that are not reducible to lower-level objects. To make things clearer, we adopt the following definition of emergentism:

Definition 14 (Emergentism). *Emergentism assumes that low-level phenomena are the cause of high-level phenomena, yet in turn not necessarily reducible to low-level phenomena.*

The resulting high-level and low-level phenomena then come to influence each other through *causally efficient* mechanisms. This classical picture of emergence distinguishes the interacting objects (physical phenomena at the *lower-level*) from the emerging objects (emergent structures at the *higher-level*). Yet providing the lower level with causally efficient properties onto the higher level induces two possibly unsatisfactory consequences: either the higher-level is an epiphenomenon (a mere consequence of low-level phenomena, which cannot cause anything itself), or it enjoys causal properties as well (which amounts to downward causation).

³Some call this "eliminativist physicalism", because processes are supposed to be fully characterized by the lowest physical level *only*.

In the first case indeed when causation goes only upwards, some authors underline the epiphenomenality of higher-level phenomena (Kim, 1999; Campbell & Bickhard, 2001). The argument is fundamentally as follows: denoting lower-level states by “ L ” and higher-level states by “ H ”, at the lower level L causes L' , however at the same time L causes H and L' causes H' ; so why would we need H and H' for? These two properties seem in fact merely epiphenomenal. Thus, “[i]f emergent properties exist, they are causally, and hence explanatorily, inert and therefore largely useless for the purposes of causal/explanatory theories” (Kim, 1999).

But then, epiphenomenality does not differ much from reductionism, and according to Bitbol (2005), “*emergentists are inclined to require productive causal powers of the emergent properties on the basic properties.*” In other words, the whole may impose constraints onto the parts. In such a framework, where both upward and downward causations are present, interactions of low-level items (in L) create a higher-level object (in H), which in turn, is supposed to have an influence on the lower-level items ($L \rightarrow H \rightarrow L'$). Hence causation goes downwards too, and H adds something to the lower-level. To Donald Campbell, who introduced the term ‘downward causation’, “*All processes at the lower levels of a hierarchy are restrained by and act in conformity to the laws of the higher levels*” (Campbell, 1974a).⁴ In other words, the whole influences the part through top-down constraints.

Definition 15 (Downward causation). *Downward causation corresponds to the fact that a system of objects which integrates a larger whole is in turn affected by the larger whole.*

For instance, cell interactions produce some emergent psychological feature (e.g. stress) which in turn induces biological changes (blood pressure increase). Similarly, consciousness is considered causally efficacious on the activity of the body (Thompson & Varela, 2001).

Although widely spread, this conception could be surprising: indeed, can a lower level *create* a higher level which in turn influences the lower level? Accordingly, detractors of downward causation argue essentially that it is redundant and, even worse, that it violates the causal rules defining the lower level; hence, they suggest, a critically erroneous principle — see e.g. (Emmeche *et al.*, 2000).

⁴More precisely, Campbell illustrates this idea as follows: “*The organisational levels of molecule, cell, tissue, organ, organism, breeding population, species, in some instances social system (...) are accepted as factual realities rather than as arbitrary conveniences of classification, with each of the higher orders organising the real units of the lower level.*”

11.3 What levels are not

Basically, each one of the three positions posits different assumptions on the status of levels, considering higher levels to exist:

- (i) *a priori* — dualism;
- (ii) *a posteriori* — emergentism;
- (iii) *only at the bottom* — reductionism.

The two first options assume the objective existence of the higher level. Let us not elaborate on strict dualism. So what about emergent levels? Often, emergent properties are called on when a system exhibits highly unexpected and/or unpredictable high-level properties.⁵ Emergentism here underscores the potential failure of reductionism in manipulating high-level properties. Granting an independent objective status to the higher level makes it possible to develop assertions and predictions on it (and particularly on what is considered irreducible or unpredictable) while still grounding the system into low-level objects. Using downward causation, it is even possible to cast back the higher level into the lower level.

But as Emmeche *et al.* (2000) put it, “*it is unclear what the ramifications are of assuming that a physical cause could have an effect which was not physical.*” Arguing that emergent properties are hard to predict from underlying properties is not a reason to abandon a strictly reductionist viewpoint. The reason why the reductionist approach still fails in practice could simply be that we miss tools, cognitive or formal, to observe and predict high-level phenomena from the low-level ones. One must tell whether there is a real emergence of irreducible novel objects or not — not only that these new properties are a convenient descriptive and predictive tool. In other words, emergentists must explain why the fact that “*each level can require a whole new conceptual structure*” (Anderson, 1972) is not simply *epistemological*. In this respect, considering temperature, which is simply an instrument and enjoys no reality by itself, Bitbol (2005) notices that “[*it*] looks as if it were a new and autonomous property, but it is only relative to the thermometric technique”. Yet, he underlines that even in the particular case of property fusion in quantum mechanics — low-level properties merge to yield an upper-level property, which in turn

⁵A common definition of ‘emergent’ is precisely “unpredictable from the basic laws”. As Shalizi (2001) notes, “*to call something emergent is therefore not to say anything about the property at all, but merely to make a confession of scientific and mathematical incompetence.*” Similarly, an easily deducible macroscopic phenomenon is rarely considered “emergent”: if the low-level mechanism at the origin of the high-level property is clearly explainable (with linear dynamic systems being the limit case), its status as an emergent feature is often weakened or considered trivial (again, particularly in the case of linearity (Bickhard & Campbell, 2000)).

forms different lower-level properties — there is no objective reality of the higher-level: *“in the upward direction, fusion of potential experimental information occurs; not fusion of actual property.”*

Now, the assumption of the existence of a lowest level, which makes the core of reductionism, is problematic as well. This point has been indeed recently challenged by Bickhard & Campbell (2000) who deny any supremacy to the lower level: *“there is no ‘bottoming out’ level in quantum field theory — it is patterns of process all the way down, and all the way up.”* For reductionism lies on the hypothesis that only higher levels are decomposable into smaller objects, a decomposition which ultimately reaches physical items governed by physical laws; yet what happens if patterning occurs at *all* levels? If we cannot consider the lowest level to involve elementary properties, then Bitbol suggests that *“no level can claim for itself the privilege of being for sure the ultimate one; ultimate and monadic.”*⁶

11.4 Observational reality of levels

11.4.1 Different modes of access

To summarize, all levels, both higher and lower, seem to vanish as substantial objects — as Bitbol puts it, *“the physical process may have no substantial roof of emergent properties, it has no substantial ground of elementary properties either.”* This apparently yields a tricky paradoxical situation, where objects and hence causality are bound to have no shelter anymore, while things still happen. To solve this, suggesting instead that properties *at any level* are the result of an observational operation proves to be a unifying and compelling answer (Bonabeau & Dessalles, 1997; Gershenson & Heylighen, 2003; Bitbol, 2005). Notably, focusing on quantum property fusion, Bitbol stresses the fact that *“[w]hat emerges is only a new mode of possible cognitive relation between the microscopic environment and the available range of experimental devices.”*

This remark is crucial and can obviously be extended to any kind of phenomenon. The whole point is to see that properties are defined only under a given instrumental apparatus, and that even lowest-level properties are always appraised through an *“instrumental intervention.”* Thus, we have to consider that there are different modes of access to a same process, not different levels that co-exist. In other words, there is a dual mode of instrumental access, not a duality of entities. In this view, we can have different kinds of properties (microscopic or

⁶This viewpoint is already present in (Campbell, 1974b): *“For a weak microscope, we assume that the homogeneous texture provided at its limit of resolution is a function of those limits, not an attribute of reality. We do this because through more powerful scopes this homogeneity becomes differentiated. By analogy, we extend this assumption even to the most powerful scope.”*




Figure 11.1: Distinct, partially overlapping aspects of an underlying process x .

macroscopic, monadic or relational) leading to the introduction (by the observer) of several kinds of related objects and phenomena — and accordingly have different modes of access to a real process, by operating on any level. Thus different ways to appraise properties emerge, *not* levels.

Therefore, Bitbol stresses out that “[t]here may be emergence without emergent properties. Not asymmetric emergence of high-level properties out of basic properties, but symmetrical co-emergence of microscopic low-level features and high level behavior.” As such, considering the co-emergence of several modes of observation is not a physicalist position, for it does not assume a lowest physical level, yet it is not dualist as well, because it does not imply dualist entities but simply the simultaneous observation of a unique process at different levels. Here levels have no consistence, rather they are *observational*: in this respect, one may say that they exist *a observatori*. By contrast with the other trends presented so far, we will call this position “observationism.”

An underlying process “ x ” is thus appraised through observations, which are phenomena in the etymological sense: things that *appear*. Each of the observed aspects of a process can be considered as a partial projection $p_i(x)$ of the underlying “ x .” Each $p_i(x)$ yields possibly overlapping information on x : the mean kinetic energy of a perfect gas gives indeed the *same* information as does a thermometer. But the thermometer is able to provide the temperature of fluids and solids as well — the thermometer, as a high-level observation instrument, yields information which obviously the mean kinetic energy cannot render. More generally, it is dubious that we could exhibit a set of instruments $\{p_1, p_2, \dots\}$ that would *wholly* characterize the process x , in the sense that any observation concerning x could be *deduced* from this minimal set of instruments, even infinite (i.e., we suggest it is impossible to find a covering of x with p_i , see Fig. 11.1).

11.4.2 Illustrations

This conception is instructive in situations involving iterated actions producing an emergent structure that in turn influences individual action, where downward causation is often supposed to play a key role. Let us consider first waves “emerging” from water: in this case water molecules move by obeying strictly mechanical laws at the lower level. Yet at a higher level a wave emerges, which in turn *like an independent object* seems to have a downward causal effect on the molecules that participate in the wave by draining them into a high-level dynamics that individual molecules cannot resist. Rather, it is a phenomenon which lends itself to dual-mode appraisal, either at the high-level of the wave or at the lower-level of molecules. Local laws applying to the lower-level are not to be modified, and molecule positions are consistent with what is to be observed at a higher-level. Looking at the wave however provides only *information* about low-level phenomena (position, movement of water molecules).

The same goes with Schelling’s (1971) celebrated model of segregated neighborhood formation. In this model, agents are placed on a grid and assigned a random color, blue or red. They behave according to a simple and unique rule consisting in changing locations in order to be surrounded by at least a certain fraction α of same-color agents. When running the model, for a sufficient value of α , large areas of same-color agents appear, as such a global pattern emerging from strictly local rules. Downward causation seems at work when “emerging” patterns in turn influence agents who join segregated neighborhoods. But this is simply apparent: the agent does not choose ‘consciously’ to join segregated neighborhoods. Her behavioral and causal rules are the same as before and need not be changed to observe an emergent macro-level behavior consisting of “agents going to same-color neighborhood.”

In the case of epistemic networks, the fact that higher-level epistemic communities appear bears no influence as such on agents: agents are still characterized by their low-level behavior. Appraising differently the process through a high-level instrument — Galois lattices — reveals high-level patterns. Agents could even appear to join epistemic communities. But in the definition of our model, agents are not explicitly influenced by epistemic communities. Other examples include norm emergence from repeated games between agents (Epstein & Axtell, 1996; Axtell *et al.*, 2001), network formation from repeated agent-based interactions (Skyrms & Pemantle, 2000), to cite a few. For every of these cases, high-level phenomena may appear to have a backward effect on the behavior of lower-level objects. Instead, the higher level simply yields large-scale information on the lower-level, but it does not induce a modification of the behavior itself, which remains unchanged. In other words, observing the higher-level provides us with knowledge on the *out-*

come of low-level behavior. Therefore, with respect to lower levels, higher levels are often macroscopic and partially informative observations — possibly expressible as a “pattern” of low-level items.

Chapter 12

Complex system modeling

Even when adopting such an observational position, the way of linking levels remains an open question — at least for the modeller. What are the implications of these philosophical considerations on modeling phenomena? How should models deal with different levels of access? Before suggesting answers, we need first to detail more extensively the operational motives of reductionists and emergentists and, by doing so, recall some goals and methods of complex system science.

12.1 Complexity and reconstruction

12.1.1 Objectives

Basically, complex system science craves for explaining high-level phenomena by playing with lower level objects. More precisely, with the help of low-level descriptions, it aims at (i) checking whether some already-known high-level descriptions are properly reconstructed (validation of higher-level phenomena), or (ii) discovering new high-level descriptions (new unexpected and potentially counterintuitive phenomena).

This attitude has two main epistemological advantages over strictly high-level descriptions: it follows Occam's razor law and, subsequently and more importantly, it works with *simpler* and, often, *more reliable* mechanisms. Simplicity means that objects are governed by more simple laws, while reliability here qualifies mechanisms that enjoy a more accurate and stable experimental validation.¹ This is most of the motto of complex system science: rebuild complex high-level behavior based on simple and well-understood "atoms."

¹Some other epistemological benefits of this approach can be found in more details in (Bonabeau, 2002) for example.

12.1.2 Commutative decomposition

In order to win the challenge of reconstruction, one could first adopt a reductionist version of the paradigm of complexity, modeling *only* low-level items. This approach discards theories of the higher level to the benefit of “micro-founded” science — as such, it discards all impermeability between scientific fields. For instance, instead of using laws and theories of psychology, one may be willing to *rebuild* them by iterating the activity of neurons, which compose here the lower level, governed by biological laws — and this is a current issue in computational neuroscience, e.g. for explaining adaptive change capabilities from neural plasticity (Destexhe & Marder, 2004).

Here, it is necessary to characterize how lower-level properties translate into higher-level properties by a projection function P (or composition function) expressing the higher-level H from the lower-level L ; that is, $P(L) = H$. Without P , how would somebody playing with low-level items expect to say anything about high-level phenomena H ? The definition of P is however not sufficient to achieve successful reconstruction: low-level dynamics observed through P must also be consistent with higher-level dynamics. Dynamical consistence means that a sequence of low-level states projected by P corresponds to a valid sequence of high-level states. More formally,² if we denote by λ (resp. η) the transfer function of a low-level state L (resp. high-level state H) to another one L' (resp. H') — in short, $\lambda(L) = L'$, $\eta(H) = H'$ — this means that P must form a *commutative* diagram with λ and η so that, as suggested in the general introduction (Rueger, 2000; Nilsson, 2004; Turner & Stepney, 2005):

$$P \circ \lambda = \eta \circ P \tag{12.1}$$

Indeed, the left side of Eq. 12.1 is the high-level result of a low-level dynamics, while the right side yields the outcome of a high-level dynamics. The aim of the reconstruction is to equate the latter with the former.

Hence commutativity is the cornerstone of the process; should this property not be verified, reconstruction would fail. How to check it? Since P is a definition and λ is designed by the modeler, η is truly the benchmark of the reconstruction. There are nevertheless two ways of considering η : (i) either η stems from *a priori* knowledge of higher-level theories (e.g., “can we rebuild these Zipf laws arising in that context?”); (ii) or η is discovered *a posteriori* from the model (e.g. “what unexpected phenomena may emerge? are they empirically valid?”). Verifying Eq. 12.1 in the first case refers to a successful reduction, while in second case it induces

²Although formulated in a specific way, this formalism could be easily transposed for a wide range of kinds of dynamics, discrete or continuous.

new knowledge for the scientist, because the challenge is to exhibit a solution $\bar{\eta}$ of Eq. 12.1, then to test this theoretical solution against reality.³

12.1.3 Reductionism failure

Nevertheless, Eq. 12.1 should hold in any case. Sometimes verifying it works perfectly, thanks to an analytical proof — such as in the famous case of temperature of gases: “*Physics can make it intelligible that mean kinetic energy of the molecules of a gas plays exactly [the] causal role [that temperature plays]*” (Beckermann, 2001); the causal role of gas temperature has been reduced to physical phenomena (molecular interactions). Sometimes it works less perfectly, because analytical resolution is hardly tractable; here only proofs on statistically sufficient simulation sets are available, using several initial states L . This is a somewhat positivist attitude, but as Epstein (2005) notices, each simulation is nonetheless a proof on a particular case, so the reconstruction may be considered a success as long as Eq. 12.1 holds true for statistically enough particular cases.

But sometimes it just doesn’t work: commutativity does not hold. For we assume η to be empirically fixed, the failure must be due either to λ or to P . Suppose that we stick to the fact that H is always correctly described by $P(L)$.⁴ Then λ must be jeopardized. In this case the fact that the low-level dynamics entails, through P , a high-level dynamics different from that given by η means that λ misses something: $\lambda(L)$ is invalid, otherwise $P(L')$ would equate H' . Solutions consist in improving the description of the low-level dynamics. In this paradigm, reductionism could fail only for practical reasons, for instance if λ has to be too complicated for commutativity to hold.⁵

³In more details: in the first case, consider an example where one already knows the empirical dynamics η^e of a given law of city size distribution ($\eta^e(H) = H'$, where both H and H' follow Zipf laws) (Pumain, 2004). The high-level state H is composed by P of low-level objects (cities and their populations) whose dynamics is deemed to be λ . Initially, $P(L) = H$. Suppose now that $P \circ \lambda(L) = H''$: if $H'' = H'$, $P \circ \lambda = \eta^e \circ P$, the reconstruction succeeded, otherwise it failed. In the second case, consider an example where one wants to observe the adoption rate of an innovation (a high-level dynamics) from low-level agent interactions (Deroian, 2002). Here also, P and λ are defined by the modeller, only η^e is induced by assuming the commutativity, i.e. find a η that satisfies Eq. 12.1. Often, this approach stops here: it rests on the *stylized* high-level dynamics η deduced from the interplay of P and λ . But at this point it should be straightforward to try to measure the empirical η^e , which comes down to the kind of empirical validations carried out in the first case: “does $\eta(H) = \eta^e(H)$?”

⁴I.e., $P(L) = H$ for all empirically valid couple of low- and high-level states (L, H) . Note that this is necessarily the case when H describes higher-level patterns on L . This is what some authors seem to call second-order properties (Kim, 1998).

⁵For the sake of instrumental practicality then, it is even possible to say that λ depends also on H , but only because $P(L) = H$, which amounts to no more than repeat that λ depends on L , through the instrumental “simplifier” P .

12.1.4 Emergentism

In spite of that, it may also be that reductionism fails for ontological reasons: P is incorrect and, more generally, it is impossible to define P . This is for example what Anderson (1972) suggests in his famous quote: “*Psychology is not applied biology.*” In other words, even with an ideally perfect knowledge of λ , reconstruction attempts would fail from the beginning because of the *inobservability* of H from L . Here the whole is more than its parts, and the higher level enjoys some sort of independence, even when acknowledging that in reality everything is physically grounded. Obviously, this is the emergentist position. H is substantially independent, and causation relationships between both levels are necessary to expect that L and λ explain something about H and η — and possibly reciprocally when assuming downward causation. In other terms, η is enriched to take L into account, and λ may be enriched to take H into account: $\lambda(L, H) = L'$, $\eta(L, H) = H'$; with possibly both levels exerting a *causally efficient* influence on each level dynamics. *In fine*, the modeller wants both λ and η to be empirically correct. So far, this is not formally different from what a “pure” dualism would yield.

Yet when considering that it is the lower-level that causes the emergence of the higher-level, most problems underlined in Sec. 11.2 & 11.3 emerge as well. Still, reductionism is uneasy to trust, because of its conception of a lowest level where all causality happens and for which projection functions P onto any level do exist (at least in theory). So, in many cases where reductionism actually fails in spite of a “solid” λ , complex system methodology nonetheless agrees with the emergentist stance.

12.2 A multiple mode of access

12.2.1 The observational viewpoint

This dilemma appears to be easily solved from an observational viewpoint. Within this framework levels are only a different way to access a same process, and L and H are “observation” functions: the high-level and the low-level are simply two simultaneous manifestations of the same process. Nonetheless, this is still a monist conception of reality: there is a single ontology, that of the process.

When levels themselves are merely informations, links between levels are thus bound to be only informational. The higher level may yield sufficient information about the underlying process, so that we can have an idea of what happens *and* what does not happen at the lower-level, and vice-versa. For example, when some individual expresses some stress (a psychological observation), one could guess that the blood pressure is higher (a biological observation). There is top-down




Figure 12.1: Relationships between levels and their dynamics in the case of (1) reductionism, (2) emergentism or dualism, and (3) observationism.

as well as bottom-up informational constraining, because information from some level specifies the dynamics of another level. To clarify this, dynamics could be rewritten as $\lambda(L|H) = L'$ and $\eta(H|L) = H'$ — see Fig. 12.1. Here again the success of the model will be measured by the empirical correctness of both λ and η .⁶ If for instance there is *ideally* enough information in the lower level about the higher level, then sufficiently valid models of the lower level bear hopes that the higher level could be rebuilt.

In case the reconstruction fails, there are two alternatives: either, as before, λ and/or η are not precise enough. Or, the chosen decomposition in levels is not informative enough about the phenomenon, and we have to check whether we are not missing something crucial when designing levels. Lane (2005) underlines this effect with a striking metaphor about “details”: there is basically no use trying to explain crises from dynamics on social classes, when the relevant item that is informative of the high-level crisis is actually at a very lower level concerning individual action. In other words, sometimes there are details that may *account for* the high-level dynamics such that the chosen decomposition into a lower-level

⁶One can introduce useful modeling approximations that seemingly give some thickness to the higher-level, but are clearly not to be confused in any way with substantial independence or downward causation. A frequent knack consists indeed in considering that the high-level is evolving slowly comparatively to low-level objects (which sometimes are considered low-level precisely because their timescale is faster), therefore being somewhat fixed and apparently independent. In this respect, some distinguish the “emergence” of higher-level items (characterized by larger, slower quantities) from the “immergence” of lower-level items in a stable, fixed high-level environment — such as boundaries (Bourgine & Stewart, 2004). This is not far from what Rueger (2000) calls “robust supervenience,” in case a high-level phenomenon enjoys some temporal stability.

dynamics is *essentially* unefficient for high-level prediction. Here, it may simply be that observing L will *never* yield enough information about H , and this bears identical consequences for modeling.

On the whole, this is a strong change in viewpoint:

- First, there is no “substantial” reality of levels, but an *observational reality* only (Sec. 11.4).
- Second, and consequently, there is no reciprocal causation of higher- and lower-level, but simply *informational links*: high- and low-levels are distinct but simultaneous observations of a same underlying process, through an instrumental “equipment” defined by the observer/scientist, that *may or may not* yield information about other levels.
- Third, and most importantly, for some phenomena it is hopeless to expect to rebuild them from some given lower-level descriptions — not because there is something irreducible in the higher level, that provides it with thickness, but because the *lower level of description itself is essentially maladapted*. Thus improving dynamics is not sufficient, and rethinking levels is mandatory.
- Lastly, the conception of “higher” and “lower” levels becomes simply a notion of *different* levels, because of a *distinct instrumental apparatus*. Therefore, problems regarding the specification of why the “higher” level is truly *above* the lower level (timescale? size? inertia?) vanish.

In this respect, both reductionism and emergentism are inadequate conceptions for appraising and modeling complex systems. Reductionism works in particular cases where the low-level description yields enough information about the high-level, giving the impression that the high-level is *reducible*, while in fact it is simply *fully deducible*. Therefore, reductionism makes the bet that physical interactions yield enough information about *any* other “higher” level, at least in principle. This is a intuitive yet very audacious bet. Emergentism on the other hand bears serious causality problems. Dualism is consistent theoretically, but clearly lacks plausibility (especially if it leads to subjective pluralism).

Application to epistemic network reconstruction In Part II we have adopted an apparent reductionist stance, starting from low-level description (epistemic networks) to rebuild high-level phenomena (epistemic communities, *inter alia*). But being reductionist would amount to say here that everything could be caused by networks built on agents and concepts. Obviously, this is not the case: only for the H we exhibited in Part I do we have a valid reconstruction from the L suggested in Part II. In other words, we showed that this L yields enough information about the

stylized facts H we selected: we could define a P such that $P(L) = H$, thanks, *inter alia*, to Galois lattices. To compare with the case of temperature, the high-level information we had through experts is like the temperature of a perfect gas obtained through a thermometer: there are low-level phenomena (epistemic network and molecular activity alike) from which we can deduce the high-level information.

More broadly, the claim is thus the following: given a high-level phenomena, it *may* be possible to find a finite set of low-level observations (potentially only one) that yield enough information to *fully deduce* the given higher level. But there is no set of finite low-level descriptors such that *any* (high-level) phenomenon can be fully deduced, even in theory — and not even at the physical level of atoms and molecules.

12.2.2 Introducing new levels

By contrast, observationism is both consistent and potentially efficient to rebuild any given complex phenomenon as long as levels are relevantly defined.⁷ In this respect, explaining phenomena at some level may require more than one level. A quite frequent need is that of a third level, intermediary between higher and lower levels: a “meso-level” deemed more informative than the macro-level while more assessable than the micro-level; sometimes crucial to understand some types of phenomena (Laughlin *et al.*, 2000). A triad of macro-, meso- and micro-levels seems rather arbitrary, and one may well imagine that some research topics involve even more levels (such as e.g. studying a (i) system of (ii) cities made of (iii) coalitions of (iv) agents who are (v) learning neural networks). While in some cases new levels are necessary (because the basic levels are essentially deficient), introducing a few levels may also be just more convenient. Here, there is no trouble using as many levels as desired, since there is only one unique and simultaneous process producing to *all* levels — and many ways to look at it. At this point activity-based modeling is a precious modeling feature, for it enables a multi-level appraisal but also yields a *natural* insight on level-specific properties (Bonabeau, 2002).

Now, how to design new levels? Various authors support the idea that introducing a new level is interesting insofar as it makes possible a better understanding and/or prediction of the system (Crutchfield, 1994; Clark, 1996; Shalizi, 2001; Gershenson & Heylighen, 2003). More precisely, the argument is essentially that emergent properties are high-level properties that “*are ‘easier to follow,’ or ‘simplify the description,’ or otherwise make our life, as creatures attempting to understand the world around us, at least a little easier*” (Shalizi, 2001). This calls clearly for choosing an observation level that provides easily key information on a given phenomenon.

⁷It is also compatible with reductionism which is a particular case where a level is “fully-informative” about another level (generally higher).

Here, instead of considering (emergent) high-level properties as something complicated, impossible to understand, or even irreducible — a negative and slippery definition — this informational attitude looks the high-level as something that must enable a more convenient understanding and prediction of the phenomenon — a positive definition.

This stance is very enlightening theoretically: to give meaning to complex systems we design new observational instruments and description grammars that help reduce reality dimensions and complexity. Going further operationally, compelling methods (Crutchfield, 1994) and effective algorithms (Shalizi & Shalizi, 2004) have been proposed to find and build automatically & endogeneously a new level of observation (i) based on low-level phenomena and (ii) simplifying their description. In any case, these tools appear to be powerful for detecting higher-order properties and informative, relevant patterns, for it yields an immediate description of H and, if the grammar is simultaneously built, a valid η too (at least statistically). However, as Shalizi (2001) notes, “*the variables describing emergent properties must be fully determined by lower-level variables.*” It becomes clear then that the new simplified “high-level” description is a clever projection function P of the lower level.

12.2.3 Rethinking levels

More generally, such methods produce relevant “high-level” description grammars, possibly hierarchically ordered, which are still based on an initial lower level (Bonabeau & Dessalles, 1997). In addition, while simpler, the newly created levels are not necessarily (i) more natural and intuitive or (ii) more importantly, complete: their efficiency is indeed limited in case the reductionist approach fails, i.e. when the chosen lower levels are not informative enough about the considered phenomenon. What happens for instance when creating high-levels from neural activity in order to describe some psychological phenomenon, while in fact there are crucial data in glial cells (Pfrieger & Barres, 1996)? What new descriptions extracted from neural activity could be effective when glial cells do a key part of the job? Consider indeed someone trying to make learning emerge from neurons and failing to do so: she could conclude that learning is a irreducible high-level description that emerges from neurons, yet models of such a thing would be irretrievably unsuccessful, if not reconsidering lower level design. Neurons are simply not sufficiently informative about learning processes. As such, emergentism could also be a dangerous pathway.

Also, the question here goes deeper: can an automatic (bottom-up) process yield an *essentially new* vision on things? This sounds as if a deterministic machine could address the problem of ontological uncertainty. In short, it may be hopeless

to expect a machine to yield a truly innovative insight starting from already deficient levels. Coming back to the central problem of rebuilding efficiently a given phenomenon through a “complex system” approach, this means that mistakes are not to be found necessarily in the dynamics λ , η , etc. nor in putative projection functions P , Q , etc.; but rather in the definition itself of levels L , H , etc. In other words, a successful reconstruction may require not only to find a valid and efficient grammar, but also to rethink the very bricks that constitute any potential grammar.

Chapter 13

Reintroducing retroaction

13.1 Differentiating objects

In the previous chapter, we detailed consequences on modeling methodology of the idea that different levels are simply different manifestations of a same process. By denying them any substantial reality and by dismissing any causal efficiency from a level to another, downward causation should be interpreted as informational dependence of low-level phenomena on high-level phenomena.¹

Yet, of course, causality may still occur between distinct objects at a same level: for instance, agents have a causal influence upon other agents. Causality may also happen between different levels, as long as it happens between different items: a hand can move the molecules that constitute a stick. A given wave moves molecules other than those that constitute this wave. Here, there is simultaneity in the movement of the hand and of its molecules, while there is causality of the hand on the stick or, equivalently, on stick molecules. In this respect, when defining a level one must describe the objects it contains as well as the causal links between these objects.

To illustrate this, consider that a neuron can *interact* with another neuron and at the same time, at a higher-level of observation, a bunch of neurons is able to *affect* other bunches of neurons. Observing a bunch of neurons provides partial *information* on the state of each individual neuron, whereas *causality* happens between different bunches of neurons and, simultaneously, between neurons of these different bunches; depending on whether one looks high-level or low-level. Therefore, if one acknowledges that there are also glial cells on the playground, causal relationships are to be expected between neurons and glial cells. At the level of the brain,

¹The modeler may yet overlook the question of the status of levels, as long as equations correctly render inter-level links/dependencies (Bourgine, personal communication). It is however really important to know where the error comes from when reconstruction fails — this is why a particular attention must be paid to level design itself.

one may consider low-level observation of neurons and high-level observation of psychological facts. Suppose now that refining the picture leads to consider the nervous system as a set of both neurons and glial cells. From there, high-level observation instruments can be designed for neurons and, separately, for glial cells. Causation occurs between neurons and glial cells (as it occurs between two neurons too), and there is a real efficient causation when glial cells observed from a high-level standpoint induce a change on individual neurons. This shall not be downward causation.

13.2 Agent behavior, semantic space

This point however helps understanding an intriguing objection that may be raised when considering intentional systems: in social systems notably, agents are able to observe what happens at a higher level, and modify their behavior accordingly. Large-scale artefacts created by agents, such as semantic items or institutions, seem to interfere with laws at the agent level. Does this induce some kind of downward causation? As we will show below, such causal influence of the higher level actually corresponds to coevolution of different kinds of objects — thus accentuating the need for accurate level descriptions, and for accurate distinction between objects.

Consider again Schelling’s model outlined in Sec. 11.4: one could be tempted to say that the higher level exerts a causal influence on the lower level: agents *decide to join* same-color neighborhoods. As we noted, it is simply a two-mode access to a same phenomenon, where agents go increasingly to places where they are surrounded by same-color agents. Eventually, using “neighborhoods” as a new high-level of description, agents *appear to join* same-color neighborhoods.

In the real world however, it seems that agents do not stick to their alleged low-level behavior (i.e. going where they are surrounded by at least $\alpha\%$ of same-color neighbors). Instead, they actually adopt another kind of behavior by really deciding to move to neighborhoods, not only to places verifying local properties. Thus, their local, low-level behavior *itself* is modified by this high-level feature. Believing in this case that this is downward causation would require to ignore that the agent behavior has been *enriched*. More precisely, the low-level description has been modified by adding a new capability to the cognitive equipment of agents: agents are now equipped with the notion of neighborhood.

Thus, what used to exist only in the eye of the modeler/observer — the presence or not of neighborhoods — has been introduced within the model, under the form of a high-level representation available to agents: agents are observers and they can access high-level descriptions. In the original Schelling model, the fact




Figure 13.1: Differentiating several kinds of objects restores the discrimination between causal links (solid lines) and informational links (dashed lines). The general picture (*top*) is applied to the two examples of this section (*below*).

that there is a neighborhood *does not* change agent behavior: neighbor colors not neighborhoods have a causal impact on agents. In the modified model, which is more realistic,² neighborhoods have a causal impact on agents *in addition to* local features such as neighbor colors. In both models, agent moves can be provoked by color-based (semantic) features; in the new one, they are furthermore affected by neighborhoods. There is still no downward causation, but a richer causal impact of *other neighbors*, both low- and high-level (local neighbors, and neighborhoods).³

²With agents more sensible to considerations on the neighborhood than to a low-level scrutiny of each location.

³High- and low- level semantic features are two observations of a same process, so there may also exist an informational overlap of both levels (e.g., the existence a blue neighborhood bears low-level information on neighbor colors).

13.3 Coevolution of objects

Here, agent behavior is causally linked to a semantic space, appraised through representational capacities, either low-level (“color of closest neighbors”) or *possibly* high-level (“belonging to a neighborhood”). Therefore, we may more generally discern two kinds of influence:

- (i) upward/downward informational dependence of a level on another, through different observation levels of a same phenomenon. Water molecules are not meant to take the wave into account, and there are two modes of access: informational links clarify the classical picture of downward causation (Bitbol, 2005).
- (ii) co-evolution of objects, through an efficient explicit causality between two different kinds of objects given *a priori*. Obviously, this remains a classical causation.

The global picture is summarized on Fig. 13.1 — put this way, it should also be possible to address tangled hierarchies explicitly without having to deal with causation violations.

To take another example, suppose we try to model the way agents create a semantic structure and paradigms through concept associations, which themselves in turn influence agents by what seems at first sight to be downward causation. This sounds like an enriched version of the model of Part II, where agent behavior has been extended to take into account high-level phenomena; as such, we get off the framework of the simple emergence of *H*. We must then distinguish: (i) the two-mode access to different features or phenomena of epistemic networks (agents and concepts, vs. social semantic and epistemic communities), and (ii) the co-evolution between objects belonging to the three kinds of networks.

Introducing co-evolutionary objects the way we did is thus crucially linked to level design. Indeed, accounting for the morphogenesis of epistemic networks using social data only may be essentially insufficient. This compels the modeler to modify the description: adding a semantic space (containing *concepts*) is required to explain the formation of such networks and the appearance of patterns (communities of agents).

13.4 “Stigmergence”

A co-evolutionary framework also yields an insight on why high-level artifacts (such as institutions) may have a proper influence on agents. Here social acts are actually “immersed” in an *environment* which influences social behavior and on

which agents may act. For instance, when an agent arrives in an epistemic network links between concepts are already present — a portion of the bibliography has already been written — but she may act upon them and make semantic associations vary and influence other agents (and herself).

In a more abstract manner, institutions are produced by agents, yet have a causal effect on agents because they can take them into account — they are equipped to recognize them. When agents build artifacts, create institutions, they produce something that is not ascribed to the particular social situation being modelled. Artifacts do exist *outer of* agents, they are *stigmergic* — in the sense Karsai & Penzes (1993) use when they describe wasps building their comb and being influenced by it, generalized in (Bonabeau *et al.*, 2000) with agents producing external, stigmergic three-dimensional structures that influence them. Thus we may talk of “*stigmergence*” of institutions or artifacts, not emergence; inducing in this case (diachronic) co-evolution, not downward causation.

Conclusion of Part III

In most scientific disciplines, levels of description can be considered to rely on objects which are themselves the focus of lower-level disciplines. In this picture, complex system science has been the cornerstone of a recent and natural effort to try to explain higher level phenomena with the help of lower-level descriptions. As an interdisciplinary area of research, this new field attempts to bridge levels by binding both lower and higher levels into a systemic framework, in order to eventually *rebuild* phenomena through the interplay of both high- and low-level objects.

This also requires considerations on how relationships between levels should be appraised. After reviewing several possible attitudes towards the status of levels (dualism, reductionism, and emergentism) we supported the idea that these three stances were possibly unsatisfactory — either because of plausibility, successfulness or consistency. Rather, noting that even the lowest level could not be the “ultimate and monadic level”, we built upon recent suggestions that levels were simply different modes of access to a process. This led us to present and adopt a viewpoint inducing only one ontology, that of the process, and many ways to look at it. In this framework, levels are instrumental apparatus created by scientists to partially access reality: they are distinct but simultaneous observations of a same underlying process. Thus, what appeared to be upward or downward causation can be reduced to informational dependence.

We then detailed the implications for modeling methodology. Indeed, a given description level may only yield (partial) information about other levels. In some cases, this information is insufficient to rebuild a given phenomenon, and new levels may be required. In the perspective of reconstruction, because some given levels may be essentially insufficiently informative for explaining a given phenomenon, we hence insisted on the idea that designing levels was as crucial as designing the dynamics. In particular, in the case of network morphogenesis the fact that, say, clustering coefficient reconstruction from the strict social network fails may be due to a wrong low-level dynamics λ . Yet, as regards epistemic community structure reconstruction, there is simply *no* P that may yield H from the strict

social network of collaborations. We are compelled to enrich the description of *L*, introducing epistemic networks.

Dismissing the possibility of retroaction could nevertheless be puzzling in several cases, in particular in artefactual systems. For instance, when studying innovation and social change, innovation is obviously not only a question of increasing production with no influence on the production processes: agents modify the production processes with respect to what they produce – hence, retroaction often happens. Putting forward level design helps reintroducing the possibility of causally efficient actions between levels, *through distinct objects*. Indeed, this kind of retroaction must not be confused with alleged downward causation; it only follows from objective differentiation, entailing causation on a “horizontal” basis. Agents produce something that remains *external*, then influences their actions. Instead of emergence, we suggest that this notion of reciprocal action of an external item should be denoted by the new term *stigmergence*.

General conclusion

“Explaining the distribution of cultural representations would be isolating the causes (...) of the capacity for some representations to propagate until becoming precisely cultural, that is, revealing the reasons of their contagiousity.”⁴ (Lenclud, 1998)

The present dissertation provides a theoretical overview of the purposes of complex system reconstruction along with an empirical achievement on a particular case study of knowledge community rebuilding. We have argued that epistemic communities are mostly produced by the co-evolution between agents and concepts. More precisely,

- in Part I, we proposed a method for describing and categorizing knowledge communities as well as capturing essential stylized facts regarding their structure. In particular, **we rebuilt the taxonomy of a whole epistemic community using a formal framework based on Galois lattices**. Then, studying the evolution of these taxonomies made possible an historical description of knowledge fields, describing *inter alia* field progress, decline, specialization, interaction (merging or splitting).
- in Part II, we *micro-founded* the particular structure observed in Part I: which processes at the level of agents may account for the emergence of epistemic community structure? To achieve a morphogenesis model of this phenomenon, and thus of epistemic networks, we needed to build tools enabling the empirical estimation of interaction and growth processes. Then, assuming that agents and concepts are co-evolving, **we successfully reconstructed the structure of a real-world scientific community on a selection of relevant high-level stylized facts**.

⁴*“Expliquer la distribution des représentations culturelles, ce serait isoler les causes (...) du pouvoir détenu par certaines représentations de se propager jusqu’à devenir justement culturelles, c’est-à-dire déceler les facteurs de leur contagiosité.”*

- in Part III, we argued that modeling social complex systems tends to require the introduction of co-evolutive frameworks of the kind presented in the preceding parts. More generally, investigating the methodology of complex system science, we **suggested that some high-level phenomena cannot be explained without a fundamental *viewpoint change* in *not only* low-level dynamics *but also* in the design of low-level objects themselves.**

Naturalizing cultural anthropology As such, this thesis also makes a preliminary to the study of knowledge diffusion and cultural pattern formation. Indeed, three canonical explanations are available to account for *cultural similarity* (Aunger, 2000): (i) genetics (i.e. convergent biological evolution), (ii) individual learning (through convergent cultural evolution), and (iii) social learning (through transmission and adoption of knowledge). It is easy to dismiss genes as an appropriate explanation: culture evolves on a dramatically shorter time-scale than that of genetic evolution. The second point alone, because it assumes the existence of cultural attractors for mankind, lacks credibility: here, cultural diversity confronts cultural similarity. On the contrary, social epistemology underlines the fact that knowledge construction is only marginally individual-based. Kornblith (1995) for instance insists on the influence of society from birth: we are immersed from the beginning in a cultural and conceptual bath, “*Language is not reinvented by each individual in social isolation, nor could it be.*”

The third argument, social learning, or *social cognition*, is thus a convincing account — Bloch (2000) summarizes the point: “*One generation may have no idea about electricity, while the next may be innovating a new computer program under Windows. This is not due to a speeding up of ‘cultural evolution’ but the result of a totally different process: the fact that humans can communicate knowledge to each other.*” Subsequently, the co-evolutionary morphogenesis model presented here is an important step for explaining cultural similarity through a naturalistic approach (Sperber, 1996): the structure and dynamics of epistemic networks has indeed a crucial impact on processes taking place on it, such as, precisely, knowledge propagation. In this respect, Pastor-Satorras & Vespignani (2001) for instance show that even with a very simplistic epidemiologic model, disease propagation follows very different paths depending on network structure.

Yet, our morphogenesis model nevertheless dismissed important considerations regarding in particular:

1. agent behavior enrichment, following the way cognitive economics improve ‘classical’ economics (Bourguine, 2004). For instance, agent behavior could be enriched to use knowledge on epistemic communities — high-level phenomena — so that it is closer to reality. This is credible at least in scientific

networks: agents refer to themselves and their work using e.g. disciplines, they do not only interact on the basis of individual properties.

2. endogenization of additional phenomena which, as suggested at the end of Part II, is strongly linked to modeling novelty and induces ontological uncertainty. Here, it is likely that we could not dismiss purely historical features: we certainly reach the boundaries of any reconstruction model in social science.

Bridging these caveats, when possible, and assessing their impact on the structure of epistemic networks — especially on features that precisely influence knowledge propagation and transmission — would be a first improvement. Besides studying cultural similarity on a social basis, including homophily, we should also investigate why cultural similarity relates to *conceptual similarity*, on an individual and cognitive basis. How comes that concepts cover identical representations among several agents of a same (epistemic) community? Working on the notion of “concept” appears to be decisive in order to depart from a strict memeticist point of view, and especially to take into account critics of memetics by cultural anthropology (Kuper, 2000; Atran, 2003). On one hand indeed, memetics could appear as a seducing program with respect to social learning, for it offers three significant features: a unit of cultural transmission (memes), a process of transmission (imitation) and characteristics of the transmission (survival of fitter ideas). Yet, memetics also entails three major drawbacks: (i) the atomistic assumption that there are bits of knowledge is very controversial; as is (ii) the assumption that there is high-fidelity transmission (imitation), when there is in most cases contextual reformulation, or reproduction; finally memetics does not address (iii) what a fitness function is, and what makes a meme be selected. In this thesis, we nevertheless assumed that using the same *term* was identical to sharing the same representation, and agents gathering in an event were *exchanging* concepts, without alteration or reinterpretation — a viewpoint that memetics would not deny. Hence, acknowledging the weaknesses of this position, we should also improve the cognitive description of processes at work in epistemic networks.⁵

⁵In particular, several authors argue that concepts are patterns in a semantic space (Colby, 2003). Empirical evidence suggests that e.g. kinship concepts are roughly located in the same area of a multidimensional semantic representation (Romney *et al.*, 1996). In other words, people of a same “culture”, using the same language could be *almost* in agreement on the meaning of concepts. Henrich & Boyd (2002) explain such aggregation by assuming that there are cognitive attractors: then, a concept is a pattern of “versions” that resemble each other. As Sperber notices, “a myth is the set of its versions.” This position does not deny that concepts are “continuously graded entities,” but it suggests that these entities aggregate around alleged attractors. Eventually, classes of equivalences of patterns might thus be of great use to model concepts.

Towards an autonomous society In any case, the work presented in this dissertation is a first brick towards enabling agents to understand the dynamics of the global social system they are participating in, and more broadly towards the achievement of a truly *autonomous* society, in Castoriadis' (1983) sense: a society which, knowing its own structure, organization, and representations, is able to determine its own *laws*. Then, what would indeed be a society which knows its own dynamics, and which precisely adapts its behavior with respect to the knowledge of its own dynamics?

List of Figures

1	The reconstruction problem	11
1.1	Sample community with s_1, s_2, s_3, s_4 and Lng, NS and Prs	25
2.1	Comparison of trees vs. lattices	33
2.2	Creating the Galois lattice	35
2.3	Galois lattice and hierarchy	37
2.4	Zoom on a diamond in a Galois lattice	38
2.5	Loss of information in one-mode projections	41
3.1	Experimental protocol: steps 1–5	45
3.2	Raw distributions of agent set sizes.	47
3.3	Cumulated densities of agent set sizes.	48
3.4	Partial view of the empirical GL, static case	50
4.1	From the original GL to a selected poset, or partial epistemic hypergraph.	52
5.1	Dynamic patterns: progress, decline, enrichment, impoverishment, merging, scission	58
5.2	Series of overlapping periods P_1, P_2 and P_3	60
5.3	Two partial epistemic hypergraphs, 1995 and 2003	62
7.1	Sample epistemic network $S, C, \mathcal{R}, \mathcal{R}^S, \mathcal{R}^C$	83
8.1	Empirical degree distribution for the social network	87
8.2	Empirical degree distribution for the semantic network	87
8.3	Empirical degree distributions for the socio-semantic network	88
8.4	Description of monopartite and bipartite clustering coefficients	91
8.5	Empirical clustering coefficients	93
8.6	Raw distribution of EC sizes, GL computed with 70 concepts	94
8.7	Distribution of empirical semantic distances	95
9.1	Degree-related interaction propensity	104
9.2	Degree-based activity	104
9.3	Homophilic interaction propensity	106

9.4	Degree and semantic distance correlations	107
9.5	Social distance-related interaction propension	108
9.6	Cumulated activity of concepts with respect to $k_{\text{concepts} \rightarrow \text{agents}}$	109
9.7	Network growth: number of old and new agents, number of articles	110
9.8	Distribution of the size of events, and composition	111
9.9	Distributions and composition of concepts per article	113
10.1	Modeling an event by specifying article contents	119
10.2	Simulated social, semantic and socio-semantic degree distributions	121
10.3	Simulated distribution of c_3 and c_4	121
10.4	Simulated distribution of EC sizes	122
10.5	Simulated distributions of semantic distances	122
11.1	Distinct, partially overlapping aspects of an underlying process x	139
12.1	Reductionism, emergentism, observationism	147
13.1	Differentiating several kinds of objects	155

References

- R. Albert and A.-L. Barabási (2002). Statistical mechanics of complex networks. *Reviews of modern physics*, **74**, 47–97.
- P. Anderson (1972). More is different. *Science*, **177**, 393–396.
- R. Atkin (1974). *Mathematical structure in human affairs*. London: Heinemann Educational Books.
- S. Atran (1998). Folk biology and the anthropology of science: Cognitive universals and cognitive particulars. *Behavioral and Brain Sciences*, **21**, 547–609.
- S. Atran (2003). Théorie cognitive de la culture, une alternative évolutionniste à la sociobiologie et à la sélection collective. *L'homme*, **166**.
- R. Aunger (ed) (2000). *Darwinizing culture: The status of memetics as a science*. Oxford: Oxford University Press.
- R. Axtell, J. M. Epstein, and H. P. Young (2001). The emergence of classes in a multi-agent bargaining model. *Pages 191–211 of: H. P. Young and S. Durlauf (eds), Social dynamics*. Cambridge: MIT Press.
- A.-L. Barabási, H. Jeong, R. Ravasz, Z. Neda, T. Vicsek, and T. Schubert (2002). Evolution of the social network of scientific collaborations. *Physica A*, **311**, 590–614.
- A.-L. Barabási (2002). *Linked: The new science of networks*. Cambridge, Mass.: Perseus Publishing.
- A.-L. Barabási and R. Albert (1999). Emergence of scaling in random networks. *Science*, **286**, 509–512.
- A.-L. Barabási, R. Albert, and H. Jeong (1999). Mean-field theory for scale-free random networks. *Physica A*, **272**, 173–187.
- A. Barbour and D. Mollison (1990). Epidemics and random graphs. *Pages 86–89 of: J.-P. Gabriel, C. Lefevre, and P. Picard (eds), Stochastic processes in epidemic theory*. Lecture Notes in Biomaths, 86. Springer.

- M. Barbut and B. Monjardet (1970). *Algèbre et combinatoire*. Vol. II. Paris: Hachette.
- A. Barrat, M. Barthélemy, R. Pastor-Satorras, and A. Vespignani (2004). The architecture of complex weighted networks. *PNAS*, **101**(11), 3747–3752.
- J.-P. Barthélemy, M. De Glas, J.-P. Desclés, and J. Petitot (1996). Logique et dynamique de la cognition. *Intellectica*, **23**, 219–301.
- V. Batagelj and M. Bren (1995). Comparing resemblance measures. *Journal of classification*, **12**(1), 73–90.
- V. Batagelj, A. Ferligoj, and P. Doreian (1999). Generalized blockmodeling. *Informatica*, **23**, 501–506.
- V. Batagelj, A. Ferligoj, and P. Doreian (2004). Generalized blockmodeling of two-mode networks. *Social networks*, **26**(1), 29–54.
- A. Beckermann (2001). Physicalism and new-wave-reductionism. *Grazer Philosophische Studien*, **61**, 257–261.
- R. Belohlavek (2000). Fuzzy Galois connections and fuzzy concept lattices: From binary relations to conceptual structures. *Pages 462–494 of: V. Novak and I. Perfilova (eds), Discovering the world with fuzzy logic*. Heidelberg: Physica-Verlag.
- J.-P. Benzécri (1973). *L'analyse des données. Tome 1. La taxinomie*. Paris: Dunod.
- N. Berger, C. Borgs, J. Chayes, R. D'Souza, and R. Kleinberg (2004). Competition-induced preferential attachment. *Pages 208–221 of: Proceedings of the 31st international colloquium on automata, languages and programming*.
- B. Berlin (1992). *Ethnobiological classification - principles of categorization of plants and animals in traditional societies*. Princeton: Princeton University Press.
- M. Bickhard and D. T. Campbell (2000). Emergence. *Pages 322–348 of: P. B. Andersen, C. Emmeche, N. O. Finnemann, and P. V. Christiansen (eds), Downward causation. minds, bodies and matter*. Aarhus: Aarhus University Press.
- G. Birkhoff (1948). *Lattice theory*. Providence, RI: American Mathematical Society.
- M. Bitbol (2005). Ontology, matter and emergence. *Phenomenology and the cognitive science*.
- M. Bloch (2000). A well-disposed social anthropologist's problem with memes. *In: R. Aunger (ed), Darwinizing culture: The status of memetics as a science*. Oxford: Oxford University Press.
- M. Boguna and R. Pastor-Satorras (2003). Class of correlated random networks with hidden variables. *Physical Review E*, **68**, 036112.

- M. Boguna, R. Pastor-Satorras, A. Diaz-Guilera, and A. Arenas (2004). Models of social networks based on social distance attachment. *Physical Review E*, **70**, 056122.
- B. Bollobás (1985). *Random graphs*. London: Academic Press.
- E. Bonabeau (2002). Agent-based modeling: Methods and techniques for simulating human systems. *PNAS*, **99**(3), 7280–7287.
- E. Bonabeau and J.-L. Dessalles (1997). Detection and emergence. *Intellectica*, **25**(2), 85–94.
- E. Bonabeau, S. Guérin, D. Snyers, P. Kuntz, and G. Theraulaz (2000). Three-dimensional architectures grown by simple 'stigmergic' agents. *Biosystems*, **56**, 13–32.
- P. Bourgin (2004). What is cognitive economics? *Pages 1–12 of: P. Bourgin and J.-P. Nadal (eds), Cognitive economics – An interdisciplinary approach*. Berlin: Springer.
- P. Bourgin and J. Stewart (2004). Autopoiesis and cognition. *Artificial life*, **10**, 327–345.
- J. Bradbury (2004). Small fish, big science. *PLoS Biology*, **2**(5), 568–572.
- R. S. Burt (1978). Cohesion versus structural equivalence as a basis for network subgroups. *Sociological methods and research*, **7**, 189–212.
- G. Caldarelli, A. Capocci, P. D. L. Rios, and M. A. Munoz (2002). Scale-free networks from varying vertex intrinsic fitness. *Physical review letters*, **89**(25), 258702.
- M. Callon, J. Law, and A. Rip (1986). *Mapping the dynamics of science and technology*. London: MacMillan Press.
- D. T. Campbell (1974a). 'Downward causation' in Hierarchically Organized Biological Systems. *Pages 179–186 of: F. Ayala and T. Dobzhansky (eds), Studies in the philosophy of biology*. Macmillan Press.
- D. T. Campbell (1974b). Evolutionary epistemology. *Pages 413–463 of: P. A. Schilpp (ed), The philosophy of Karl Popper*. La Salle, Ill.: Open Court.
- R. J. Campbell and M. H. Bickhard (2001). *Physicalism, emergence and downward causation*. <http://eprints.anu.edu.au/archive/00000029>.
- N. Carayol and P. Roux (2004). Micro-grounded models of complex network formation. *Cahiers d'interactions localisées*, **1**, 49–69.

- C. Castoriadis (1983). La logique des magmas et la question de l'autonomie. *Pages 421–443 of: P. Dumouchel and J.-P. Dupuy (eds), L'auto-organisation. De la physique au politique.* Paris: Seuil.
- M. Catanzaro, G. Caldarelli, and L. Pietronero (2004). Assortative model for social networks. *Physical Review E*, **70**, 037101.
- D. Chavalarias (2004). *Métadynamiques en cognition sociale*. Ph.D. thesis, Ecole Polytechnique, Paris, France. Part III.
- C. Chen, T. Cribbin, R. Macredie, and S. Morar (2002). Visualizing and tracking the growth of competing paradigms: Two case studies. *Journal of the american society for information science and technology*, **53**(8), 678–689.
- A. Clark (1996). *Being there: Putting brain, body, and world together again*. Cambridge: MIT Press. Chap. 6, Emergence and Explanation, pages 103–128.
- P. Cohendet, A. Kirman, and J.-B. Zimmermann (2003). Emergence, formation et dynamique des réseaux – modèles de la morphogenèse. *Revue d'Economie Industrielle*, **103**(2-3), 15–42.
- B. N. Colby (2003). Toward a theory of culture and adaptive potential. *Mathematical anthropology and cultural theory*, **1**(3).
- Cold Spring Harbor Laboratory (1994, 1996, 1998, 2000, 2001, 2002, 2003). *Zebrafish development & genetics*. Cold Spring Harbor, NY.
- V. Colizza, J. R. Banavar, A. Maritan, and A. Rinaldo (2004). Network structures from selection principles. *Physical review letters*, **92**(19), 198701.
- R. Cowan, P. A. David, and D. Foray (2000). The explicit economics of knowledge codification and tacitness. *Industrial & corporate change*, **9**(2), 212–253.
- R. Cowan, N. Jonard, and J.-B. Zimmermann (2002, July). *The joint dynamics of networks and knowledge*. Computing in Economics and Finance 2002 354. Society for Computational Economics.
- J. P. Crutchfield (1994). The calculi of emergence: Computation, dynamics, and induction. *Physica D*, **75**, 11–54.
- B. A. Davey and H. A. Priestley (2002). *Introduction to lattices and order*. 2nd edn. Cambridge, UK: Cambridge University Press.
- M. De Glas (1992). A local intensional logic. In: *International conference on algebraic logic and their computer science applications*. Warsaw: Stefan Banach Mathematical Institute.
- A. Degenne and M. Forse (1999). *Introducing social networks*. Sage Publications Inc.

- F. Deroian (2002). Formation of social networks and diffusion of innovations. *Research policy*, **31**, 835–846.
- A. Destexhe and E. Marder (2004). Plasticity in single neuron and circuit computations. *Nature*, **431**, 789–795.
- P. D’Haeseleer, S. Liang, and R. Somogyi (2000). Genetic network inference: from co-expression clustering to reverse engineering. *Bioinformatics*, **16**(8), 707–726.
- H. Dicky, C. Dony, M. Huchard, and T. Libourel (1995). ARES, Adding a class and REStructuring inheritance hierarchies. *Pages 25–42 of: Actes de BDA’95 (Bases de Données Avancées), Nancy*.
- E. W. Dijkstra (1959). A note on two problems in connexion with graphs. *Numerische Mathematik*, 269–271.
- P. S. Dodds, R. Muhamad, and D. J. Watts (2003). An experimental study of search in global social networks. *Science*, **301**, 827–829.
- K. Dooley and L. I. Zon (2000). Zebrafish: a model system for the study of human disease. *Current opinion in genetics & development*, **10**(3), 252–256.
- P. Doreian and A. Mrvar (1996). A partitioning approach to structural balance. *Social networks*, **18**(2), 149–168.
- P. Doreian, V. Bategelj, and A. Ferligoj (2005). *Generalized blockmodelling*. Cambridge: Cambridge University Press.
- S. N. Dorogovtsev and J. F. F. Mendes (2000). Evolution of networks with aging of sites. *Physical Review E*, **62**, 1842–1845.
- S. N. Dorogovtsev and J. F. F. Mendes (2003). *Evolution of networks — From biological nets to the Internet and WWW*. Oxford: Oxford University Press.
- S. N. Dorogovtsev, J. F. F. Mendes, and A. N. Samukhin (2000). Structure of growing networks with preferential linking. *Physical Review Letters*, **85**(21), 4633–4636.
- O. Dupouet, P. Cohendet, and F. Creplet (2001). *Economics with heterogenous agents*. Berlin: Springer. Chap. Organisational innovation, communities of practice and epistemic communities: the case of Linux.
- V. Duquenne, C. Chabert, A. Cherfouh, A.-L. Doyen, J.-M. Delabar, and D. Pickering (2003). Structuration of phenotypes and genotypes through Galois lattices and implications. *Applied artificial intelligence*, **17**(3), 243–256.
- H. Ebel, J. Davidsen, and S. Bornholdt (2002). Dynamics of social networks. *Complexity*, **8**(2), 24–27.

- E. Eisenberg and E. Y. Levanon (2003). Preferential attachment in the protein network evolution. *Physical review letters*, **91**(13), 138701.
- C. Emmeche, S. Koppe, and F. Stjernfelt (2000). Levels, emergence, and three versions of downward causation. *Pages 13–34 of: P. B. Andersen, C. Emmeche, N. O. Finnemann, and P. V. Christiansen (eds), Downward causation. minds, bodies and matter.* Aarhus: Aarhus University Press.
- J. M. Epstein (2005). *Remarks on the foundations of agent-based generative social science.* Tech. rept. 00506024. Santa Fe Institute.
- J. M. Epstein and R. Axtell (1996). *Growing artificial societies: social science from the bottom up.* Washington, DC, USA: The Brookings Institution.
- P. Erdős and A. Rényi (1959). On random graphs. *Publicationes mathematicae*, **6**, 290–297.
- A. Fabrikant, E. Koutsoupias, and C. H. Papadimitriou (2002). Heuristically optimized trade-offs: A new paradigm for power laws in the internet. *Pages 110–122 of: Icalp '02: Proceedings of the 29th international colloquium on automata, languages and programming.* London, UK: Springer-Verlag.
- M. Faloutsos, P. Faloutsos, and C. Faloutsos (1999). On power-law relationships of the Internet topology. *Computer communication review*, **29**(4), 251–262.
- C. Fellbaum (ed) (1998). *WordNet: An electronic lexical database.* Cambridge, Mass: MIT Press.
- S. Ferré and O. Ridoux (2000). A file system based on concept analysis. *Pages 1033–1047 of: J. W. Lloyd, V. Dahl, U. Furbach, M. Kerber, K.-K. Lau, C. Palamidessi, L. M. Pereira, Y. Sagiv, and P. J. Stuckey (eds), Computational logic.* Lecture Notes in Computer Science, vol. 1861. Springer.
- K. H. Fischer and J. A. Hertz (1993). *Spin glasses.* Cambridge: Cambridge University Press.
- L. C. Freeman and D. R. White (1993). Using Galois lattices to represent network data. *Sociological methodology*, **23**, 127–146.
- L. C. Freeman (1977). A set of measures of centrality based on betweenness. *Sociometry*, **40**, 35–41.
- L. C. Freeman (1989). Social networks and the structure experiment. *Pages 11–40 of: L. C. Freeman, D. R. White, and A. K. Romney (eds), Research methods in social network analysis.* Fairfax, Va.: George Mason University Press.
- N. E. Friedkin (1991). Theoretical foundations for centrality measures. *American journal of sociology*, **96**(6), 1478–1504.

- B. Ganter (1984). *Two basic algorithms in concept analysis*. Tech. rept. preprint #831. TH-Darmstadt.
- B. Gaume (2004). Balades aléatoires dans les petits mondes lexicaux. *I3 Information Interaction Intelligence*, 4(2).
- C. Gershenson and F. Heylighen (2003). When can we call a system self-organizing? *Pages 606–614 of: W. Banzhaf, T. Christaller, P. Dittrich, J. T. Kim, and J. Ziegler (eds), Advances in artificial life, 7th european conference, ECAL 2003 LNAI 2801*. Springer-Verlag.
- R. Giere (2002). Scientific cognition as distributed cognition. *Pages 285–299 of: P. Carruthers, S. Stich, and M. Siegal (eds), The cognitive basis of science*. Cambridge University Press.
- G. Gigerenzer (2003). Where do new ideas come from? A heuristics of discovery in the cognitive sciences. *Pages 99–139 of: M. Galavotti (ed), Observation and experiment in the natural and social sciences*. Amsterdam: Kluwer Academic Publishers.
- M. Girvan and M. E. J. Newman (2002). Community structure in social and biological networks. *PNAS*, 99, 7821–7826.
- R. Godin, G. Mineau, R. Missaoui, and H. Mili (1995). Méthodes de classification conceptuelle basées sur les treillis de Galois et applications. *Revue d'intelligence artificielle*, 9(2), 105–137.
- R. Godin, H. Mili, G. W. Mineau, R. Missaoui, A. Arfi, and T.-T. Chau (1998). Design of class hierarchies based on concept (Galois) lattices. *Theory and practice of object systems (TAPOS)*, 4(2), 117–134.
- S. Goyal (2003). Learning in networks: A survey. *In: G. Demange and M. Wooders (eds), Group formation in economics: Networks, clubs, and coalitions*. Cambridge: Cambridge University Press.
- M. Granovetter (1985). Economic action and social structure: The problem of embeddedness. *American journal of sociology*, 91(3), 481–510.
- D. Gruhl, R. Guha, D. Liben-Nowell, and A. Tomkins (2004, May 17-22). Information diffusion through blogspace. *In: Proceedings of WWW2004*.
- D. J. Grunwald and J. S. Eisen (2002). Headwaters of the zebrafish – emergence of a new model vertebrate. *Nature rev. genetics*, 3(9), 717–724.
- N. Guelzim, S. Bottani, P. Bourguine, and F. Képès (2002). Topological and causal structure of the yeast transcriptional regulatory network. *Nature genetics*, 31(5), 60–63.

- J.-L. Guillaume and M. Latapy (2004a). Bipartite graphs as models of complex networks. *In: Lecture notes in computer science (LNCS), proceedings of the international workshop on combinatorial and algorithmic aspects of networking, Banff, Canada.*
- J.-L. Guillaume and M. Latapy (2004b). Bipartite structure of all complex networks. *Information processing letters*, **90**(5), 215–221.
- R. Guimera, B. Uzzi, J. Spiro, and L. A. N. Amaral (2005). Team assembly mechanisms determine collaboration network structure and team performance. *Science*, **308**, 697–702.
- P. Haas (1992). Introduction: epistemic communities and international policy coordination. *International organization*, **46**(1), 1–35.
- J. A. Hartigan (1975). *Clustering algorithms*. Wiley, New York, NY.
- J. Hasty, D. McMillen, F. Isaacs, and J. J. Collins (2001). Computational studies of gene regulatory networks: *in numero* molecular biology. *Nature reviews genetics*, **2**, 268–279.
- J. Henrich and R. Boyd (2002). *Five misunderstandings about cultural evolution*. forthcoming in *The Epidemiology of Ideas*, D. Sperber ed., London: Open Court Publishing.
- J. E. Hopcroft, O. Khan, B. Kulis, and B. Selman (2003). Natural communities in large linked networks. *Pages 541–546 of: KDD '03: Proceedings of the ninth ACM SIGKDD international conference on knowledge discovery and data mining*. Washington, D.C.: ACM Press.
- N. Ide and J. Véronis (1998). Word sense disambiguation: The state of the art. *Computational linguistics*, **24**(1), 1–40.
- International Human Genome Sequencing Consortium (2001). Initial sequencing and analysis of the human genome. *Nature*, **409**, 860–921.
- R. Jackendoff (2002). *Foundations of language: Brain, meaning, grammar, evolution*. Oxford: Oxford University Press.
- C. Jacquelinet, O. Bodenreider, and A. Burgun (2000). Modelling syllepse in medical knowledge bases with application in the domain of organ failure and transplantation. *In: Proceedings of OntoLex 2000, Workshop on ontologies and lexical knowledge bases, Sozopol, Bulgaria.*
- A. K. Jain, M. N. Murty, and P. J. Flynn (1999). Data clustering: a review. *ACM computing surveys*, **31**(3), 264–323.
- H. Jeong, Z. Néda, and A.-L. Barabási (2003). Measuring preferential attachment for evolving networks. *Europhysics letters*, **61**(4), 567–572.

- E. M. Jin, M. Girvan, and M. E. J. Newman (2001). The structure of growing social networks. *Physical Review E*, **64**(4), 046132.
- J. H. Johnson (1986). Stars, maximal rectangles, lattices: A new perspective on q-analysis. *International journal of man-machine studies*, **24**(3), 293–299.
- S. C. Johnson (1967). Hierarchical clustering schemes. *Psychometrika*, **2**, 241–254.
- I. Karsai and Z. Penzes (1993). Comb building in social wasps: Self-organization and stigmergic script. *Journal of theoretical biology*, **161**(4), 505–525.
- J. Kim (1998). *Mind in a physical world*. Cambridge: MIT Press.
- J. Kim (1999). Making sense of emergence. *Philosophical studies*, **95**, 3–36.
- A. Kirman (1997). The economy as an evolving network. *Journal of evolutionary economics*, **7**(4), 339–353.
- J. T. Klein (1990). *Interdisciplinarity: History, theory, and practice*. Detroit, MI: Wayne State University Press.
- T. Kohonen (2000). *Self-organizing maps*. 3rd edn. Berlin: Springer.
- H. Kornblith (1995). A conservative approach to social epistemology. In: F. Schmitt (ed), *Socializing epistemology: The social dimensions of knowledge*. Lanham, MD: Rowman and Littlefield.
- G. Kossinets (2005). Effects of missing data in social networks. *Social networks*. To appear.
- P. L. Krapivsky, S. Redner, and F. Leyvraz (2000). Connectivity of growing random networks. *Physical Review Letters*, **85**, 4629–4632.
- H. Kreuzman (2001). A co-citation analysis of representative authors in philosophy: Examining the relationship between epistemologists and philosophers of science. *Scientometrics*, **51**(3), 525–539.
- T. S. Kuhn (1970). *The structure of scientific revolutions*. 2nd edn. Chicago, IL: University of Chicago Press.
- R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tomkins, and E. Upfal (2000). Stochastic models for the web graph. *Page 57 of: IEEE 41st annual symposium on Foundations of Computer Science (FOCS)*.
- A. Kuper (2000). If memes are the answer, what is the question? In: R. Aunger (ed), *Darwinizing culture: The status of memetics as a science*. Oxford: Oxford University Press.

- S. O. Kuznetsov and S. A. Obiedkov (2002). Comparing performance of algorithms for generating concept lattices. *Journal of experimental and theoretical artificial intelligence*, **14**(2-3), 189–216.
- D. A. Lane (2005). *Hierarchy, complexity, society*. Working paper.
- D. A. Lane and R. R. Maxfield (2005). Ontological uncertainty and innovation. *Journal of Evolutionary Economics*, **15**(1), 3–50.
- D. A. Lane (1993). Artificial worlds and economics, part I. *Journal of Evolutionary Economics*, **3**, 89–107.
- M. Latapy and P. Pons (2004). Computing communities in large networks using random walks. *arXiv e-print archive*, 0412568.
- M. Latapy, C. Magnien, M. Mariadassou, and C. Roth (2005). A basic toolbox for the analysis of dynamics of growing networks. In: *Proceedings of the 7th "rencontres francophones sur l'algorithmique des télécommunications" – Algotel*.
- R. B. Laughlin and D. Pines (2000). The theory of everything. *PNAS*, **97**(1), 28–31.
- R. B. Laughlin, D. Pines, J. Schmalian, B. P. Stojkovic, and P. Wolynes (2000). The middle way. *PNAS*, **97**(1), 32–37.
- E. O. Laumann, P. V. Marsden, and D. Prensky (1989). The boundary specification problem in network analysis. *Pages 61–87 of: L. C. Freeman, D. R. White, and A. K. Romney (eds), Research methods in social network analysis*. Fairfax, Va.: George Mason University Press.
- J. Lave and E. Wenger (1991). *Situated learning: Legitimate peripheral participation*. Cambridge: Cambridge University Press.
- R.-J. Lavie (2003). Systemic productivity must complement structural productivity. In: *Proceedings of language, culture and cognition: An international conference on cognitive linguistics*. Braga, Portugal, July 2003.
- P. F. Lazarsfeld and R. K. Merton (1954). Friendship as a social process: a substantive and methodological analysis. *Pages 18–66 of: M. Berger (ed), Freedom and control in modern society*. New York: Van Nostrand.
- E. Lazega and M. van Duijn (1997). Position in formal structure, personal characteristics and choices of advisors in a law firm: a logistic regression model for dyadic network data. *Social networks*, **19**, 375–397.
- A. Lelu, P. Bessières, A. Zasadzinski, and D. Besagni (2004). Extraction de processus fonctionnels en génétique des microbes à partir de résumés medline. In: *Proceedings of the journées francophones d'extraction et de gestion des connaissances, egc 2004*. Clermont-Ferrand, France.

- G. Lenclud (1998). La culture s'attrape-t-elle ? *Communications, EHESS, Centre d'études transdisciplinaires*, **66**, 165–183.
- L. Leydesdorff (1991a). In search of epistemic networks. *Social studies of science*, **21**, 75–110.
- L. Leydesdorff (1991b). The static and dynamic analysis of network data using information theory. *Social networks*, **13**, 301–345.
- L. Leydesdorff (1997). Why words and co-words cannot map the development of the sciences. *Journal of the American society for information science*, **48**(5), 418–427.
- D. Liben-Nowell and J. Kleinberg (2003). The link prediction problem for social networks. *Pages 556–559 of: Cikm '03: Proceedings of the 12th international conference on information and knowledge management*. New York, NY, USA: ACM Press.
- P. G. Lind, M. C. Gonzalez, and H. J. Herrmann (2005). Cycles and clustering in bipartite networks. *Physical Review E*, **72**, 056127.
- C. Lindig (1998). *Concepts, a free and portable implementation of concept analysis in C*. Open source software package available on <http://www.st.cs.uni-sb.de/~lindig/src/concepts-0.3f.tar.gz>.
- A. Lopez, S. Atran, J. D. Coley, D. L. Medin, and E. E. Smith (1997). The tree of life: Universal and cultural features of folkbiological taxonomies and inductions. *Cognitive psychology*, **32**(3), 251–295.
- F. Lorrain and H. C. White (1971). Structural equivalence of individuals in social networks. *Journal of mathematical sociology*, **1**(49–80).
- S. S. Manna and P. Sen (2002). Modulated scale-free network in euclidean space. *Physical Review E*, **66**, 066114.
- R. K. May (1972). Will a large complex system be stable? *Nature*, **238**(413–414).
- K. W. McCain (1986). Cocited author mapping as a valid representation of intellectual structure. *Journal of the american society for information science*, **37**(3), 111–122.
- M. McPherson and L. Smith-Lovin (2001). Birds of a feather: Homophily in social networks. *Annual review of sociology*, **27**, 415–440.
- S. Milgram (1967). The small world problem. *Psychology today*, **2**, 60–67.
- M. Mitzenmacher (2003). A brief history of generative models for power law and lognormal distributions. *Internet mathematics*, **1**(2), 226–251.

- M. Molloy and B. Reed (1995). A critical point for random graphs with a given degree sequence. *Random structures and algorithms*, **161**(6), 161–179.
- B. Monjardet (2003). The presence of lattice theory in discrete problems of mathematical social sciences. *Why. Mathematical social sciences*, **46**(2), 103–144.
- J. Moody and D. R. White (2003). Structural cohesion and embeddedness: a hierarchical conception of social groups. *American sociological review*, **68**(103–127).
- S. A. Morris (2005). Bipartite yule processes in collections of journal papers. *In: 10th International Conference of the International Society for Scientometrics and Informetrics, Stockholm, Sweden, July 24-28*.
- M. E. J. Newman (2001a). Clustering and preferential attachment in growing networks. *Physical review letters E*, **64**(025102).
- M. E. J. Newman (2001b). Scientific collaboration networks. I. Network construction and fundamental results. *Physical Review E*, **64**, 016131.
- M. E. J. Newman (2001c). Scientific collaboration networks. II. Shortest paths, weighted networks, and centrality. *Physical Review E*, **64**, 016132.
- M. E. J. Newman (2001d). The structure of scientific collaboration networks. *PNAS*, **98**(2), 404–409.
- M. E. J. Newman (2002). Assortative mixing in networks. *Physical review letters*, **89**, 208701.
- M. E. J. Newman (2003). The structure and function of complex networks. *SIAM review*, **45**(2), 167–256.
- M. E. J. Newman (2004). Detecting community structure in networks. *European physical journal B*, **38**, 321–330.
- M. E. J. Newman (2005). Power laws, Pareto distributions and Zipf's law. *Contemporary physics*, **46**(5), 323–351.
- M. E. J. Newman and J. Park (2003). Why social networks are different from other types of networks. *Physical Review E*, **68**(036122).
- M. Nilsson (2004). *Hierarchical organization in smooth dynamical systems*. Working paper, to appear in *Artificial Life*.
- E. C. M. Noyons and A. F. J. van Raan (1998). Monitoring scientific developments from a dynamic perspective: self-organized structuring to map neural network research. *Journal of the american society for information science*, **49**(1), 68–81.
- D. Papineau (2001). The rise of physicalism. *In: B. Loewer and C. Gillet (eds), Physicalism and its discontents*. Cambridge: Cambridge University Press.

- G. Parisi (1992). *Field theory, disorder and simulations*. Singapore: World Scientific.
- R. Pastor-Satorras and A. Vespignani (2001). Epidemic spreading in scale-free networks. *Physical review letters*, **86**(14), 3200–3203.
- P. Pattison, S. Wasserman, G. Robins, and A. M. Kanfer (2000). Statistical evaluation of algebraic constraints for social networks. *Journal of mathematical psychology*, **44**, 536–568.
- M. Peltomaki and M. Alava (2005). Correlations in bipartite collaboration networks. *arXiv e-print archive*, **physics**, 0508027.
- F. W. Pfrieger and B. A. Barres (1996). New views on synapse-glia interactions. *Current opinion in neurobiology*, **6**, 615–621.
- M. F. Porter (1980). An algorithm for suffix stripping. *Program*, **14**(3), 130–137.
- W. W. Powell, D. R. White, K. W. Koput, and J. Owen-Smith (2005). Network dynamics and field evolution: The growth of interorganizational collaboration in the life sciences. *American journal of sociology*, **110**(4), 1132–1205.
- D. Pumain (2004). *Scaling laws and urban systems*. SFI Working Paper 04-02-002.
- M. R. Quillian (1968). Semantic memory. In: M. Minsky (ed), *Semantic information processing*. Cambridge: M.I.T. Press.
- F. Radicchi, C. Castellano, F. Cecconi, V. Loreto, and D. Parisi (2004). Defining and identifying communities in networks. *PNAS*, **101**(9), 2658–2663.
- J. J. Ramasco, S. N. Dorogovtsev, and R. Pastor-Satorras (2004). Self-organization of collaboration networks. *Physical review E*, **70**, 036106.
- E. Ravasz and A.-L. Barabási (2003). Hierarchical organization in complex networks. *Physical Review E*, **67**, 026112.
- S. Redner (1998). How popular is your paper? An empirical study of the citation distribution. *European Phys. Journal B*, **4**(131–134).
- S. Redner (2005). Citation statistics from 110 years of physical review. *Physics today*, **58**, 49–54.
- G. Robins and M. Alexander (2004). Small worlds among interlocking directors: Network structure and distance in bipartite graphs. *Computational and mathematical organization theory*, **10**, 69–94.
- L. M. Rocha (2002). Semi-metric behavior in document networks and its application to recommendation systems. *Pages 137–163 of: V. Loia (ed), Soft computing agents: A new perspective for dynamic information systems*. International Series Frontiers in Artificial Intelligence and Applications. Amsterdam: IOS Press.

- A. K. Romney, J. P. Boyd, C. C. Moore, W. H. Batchelder, and T. J. Brazill (1996). Culture as shared cognitive representations. *PNAS*, **93**, 4699–4705.
- E. Rosch and B. Lloyd (1978). Cognition and categorization. *American psychologist*, **44**(12), 1468–1481.
- C. Roth (2005). Generalized preferential attachment: Towards realistic social network models. In: *ISWC 4th Intl semantic web conference, Workshop on Semantic Network Analysis*.
- C. Roth and P. Bourguine (2003). Binding social and cultural networks: a model. *arXiv e-print archive*, **nlin.AO**, 0309035.
- C. Roth and P. Bourguine (2005). Epistemic communities: Description and hierarchic categorization. *Mathematical population studies*, **12**(2), 107–130.
- C. Roth and P. Bourguine (2006). Lattices for dynamic, hierarchic & overlapping categorization: the case of epistemic communities. *Scientometrics*. To appear.
- A. Rueger (2000). Robust supervenience and emergence. *Philosophy of science*, **67**(3), 466–489.
- G. M. Sacco (2000). Dynamic taxonomies: A model for large information bases. *IEEE Transactions on knowledge and data engineering*, **12**(3), 468–479.
- G. Salton, A. Wong, and C. S. Yang (1975). Vector space model for automatic indexing. *Communications of the ACM*, **18**(11), 613–620.
- T. C. Schelling (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*, **1**, 143–186.
- F. Schmitt (ed) (1995). *Socializing epistemology: The social dimensions of knowledge*. Lanham, MD: Rowman & Littlefield.
- C. R. Shalizi (2001). *Causal architecture, complexity and self-organization in time series and cellular automata*. Ph.D. thesis, University of Wisconsin at Madison, U.S.A. Chap. 11.
- C. R. Shalizi and K. L. Shalizi (2004). Blind construction of optimal non-linear recursive predictors for discrete sequences. *Pages 504–511 of: M. Chickering and J. Halpern (eds), Uncertainty in artificial intelligence: Proceedings of the 20th conference*.
- A. G. Simpson and A. J. Roger (2004). The real ‘kingdoms’ of eukaryotes. *Current biology*, **14**(17), R693–R696.
- B. Skyrms and R. Pemantle (2000). A dynamic model of social network formation. *PNAS*, **97**(16), 9340–9346.

- T. A. Snijders (2001). The statistical evaluation of social networks dynamics. *Sociological methodology*, **31**, 361–395.
- B. Söderberg (2003). A general formalism for inhomogeneous random graphs. *Physical review E*, **68**, 026107.
- R. R. Sokal and P. H. A. Sneath (1963). *Principles of numerical taxonomy*. San Francisco, CA: W.H. Freeman.
- D. F. Specht (1990). Probabilistic neural networks. *Neural networks*, **3**(1), 109–118.
- D. Sperber (1996). *Explaining culture: A naturalistic approach*. Oxford: Blackwell Publishers.
- R. Srikant and R. Agrawal (1995). Mining generalized association rules. In: *Proceedings of the 21st vldb (very large databases) conference*.
- H. Stefancic and V. Zlatic (2005). Preferential attachment with information filtering—node degree probability distribution properties. *Physica A*, **350**(2-4), 657–670.
- G. Stumme (2002). Formal concept analysis on its way from mathematics to computer science. *Pages 2–19 of: ICCS '02: Proceedings of the 10th international conference on conceptual structures*. London, UK: Springer-Verlag.
- G. Stumme, R. Taouil, Y. Bastide, N. Pasquier, and L. Lakhal (2002). Computing iceberg concept lattices with TITANIC. *Data and knowledge engineering*, **42**, 189–222.
- E. Thompson and F. J. Varela (2001). Radical embodiment: neural dynamics and consciousness. *Trends in cognitive sciences*, **5**(10), 418–425.
- J. C. Touhey (1974). Situated identities, attitude similarity, and interpersonal attraction. *Sociometry*, **37**, 363–374.
- H. Turner and S. Stepney (2005). Rule migration: Exploring a design framework for modelling emergence in CA-like systems. In: *ECAL Workshop on Unconventional Computing*. To appear in *International Journal of Unconventional Computing*.
- F. J. Van Der Merwe and D. G. Kourie (2002). Compressed pseudo-lattices. *Journal of experimental and theoretical artificial intelligence*, **14**(2-3), 229–254.
- A. Vázquez (2001). Disordered networks generated by recursive searches. *Europhysics letters*, **54**(4), 430–435.
- C. Vogel (1988). *Génie cognitif*. Paris: Masson. Chap. Les taxinomies.

- L. Wang, W. Song, and D. Cheung (2000). Using contextual semantics to automate the web document search and analysis. *In: Proceedings of the first international conference on Web Information Systems Engineering (WISE)*. Honk Kong, China, July 2000.
- S. Wasserman and K. Faust (1994). *Social network analysis: Methods and applications*. Cambridge: Cambridge University Press.
- D. J. Watts and S. H. Strogatz (1998). Collective dynamics of 'small-world' networks. *Nature*, **393**, 440–442.
- D. J. Watts, P. S. Dodds, and M. E. J. Newman (2002). Identity and search in social networks. *Science*, **296**, 1302–1305.
- B. Wellman, P. J. Carrington, and A. Hall (1988). Networks as personal communities. *Pages 130–184 of: B. Wellman and S. D. Berkowitz (eds), Social structures: A network analysis*. Cambridge, UK: Cambridge University Press.
- E. Wenger and W. M. Snyder (2000). Communities of practice: the organizational frontier. *Harvard business review*, **1**, 139–145.
- D. R. White and P. Spufford (2006). *Medieval to modern: Civilizations as dynamic networks*. Book Ms.
- D. R. White, N. Kejzar, C. Tsallis, D. Farmer, and S. D. White (2006). A generative model for feedback networks. *Physical Review E*, **73**, 016119.
- H. C. White, S. A. Boorman, and R. L. Breiger (1976). Social-structure from multiple networks. I: Blockmodels of roles and positions. *American journal of sociology*, **81**, 730–780.
- R. H. Whittaker (1969). New concepts of kingdoms of organisms. *Science*, **163**, 150–160.
- R. Wille (1982). Restructuring lattice theory: an approach based on hierarchies of concepts. *Pages 445–470 of: I. Rival (ed), Ordered sets*. Dordrecht-Boston: Reidel.
- R. Wille (1992). Concept lattices and conceptual knowledge systems. *Computers mathematics and applications*, **23**, 493.
- R. Wille (1997). Conceptual graphs and formal concept analysis. *Pages 290–303 of: Proceedings of the fourth international conference on conceptual structures*. Lecture Notes on Computer Science, no. #1257. Berlin: Springer.
- T. P. Wilson (1982). Relational networks: An extension of sociometric concepts. *Social networks*, **4**(2), 105–116.

-
- C. Yuh, H. Bolouri, and E. H. Davidson (1998). Genomic cis-regulatory logic: Experimental and computational analysis of a sea urchin gene. *Science*, **279**(5358), 1896–1902.
- L. A. Zadeh (1965). Fuzzy sets. *Information and control*, **8**, 358–353.
- E. W. Zegura, K. L. Calvert, and S. Bhattacharjee (1996, March). How to model an internetwork. *Pages 594–602 of: IEEE Infocom*, vol. 2. IEEE, San Francisco, CA.

Index

- activity, 101
- antichain, 51
- autonomous society, 166
- categorization
 - basic-level, 39
 - clustering method, 34
- closed couple, 28
- closure operation, 27
- clustering coefficient
 - bipartite, 90
 - monopartite, 89
- clustering method, *see* categorization
- concept
 - exchange, 112
 - network, *see* network, semantic terms, 43
- degree, 78
 - distribution, 78
- dendrogram, 40
- diamond
 - in a graph, 91
 - in a lattice, 37
- distance
 - semantic, 93
 - social, 107
- downward causation, 136
- dualism, 134
- dyadic, 98
- emergentism, 135
- epistemic community
 - enrichment, impoverishment, 57
 - formal definition*, 24
 - merging, scission, 57
 - natural definition*, 23
 - progress, decline, 57
 - subfield & superfield, 32
- epistemic group, 24
- exogenous, 126
- extension, 25
- field, *see* epistemic community
- Galois lattice
 - definition*, 34
 - graphical representation*, 34
 - Hasse diagram, 34
- graph, 28, *see* network
- homophily, 98, 105
- hypergraph
 - definition*, 28
 - epistemic, 28
 - partial, 52
- instrumental apparatus, 138
- intension, 24
- inter-disciplinary, 37
- interaction
 - n*-adic, 116
 - propension, 99
- interactivity, 102
- knowledge community, 17
- lattice, 32
 - Galois lattice, 34
- levels
 - definition*, 10
 - design of, 149
 - dynamics, 10
- memetics, 165

- micro-found, 75
- monadic
 - level, 159
 - property, 98
- multi-disciplinary, 37
- network
 - bipartite graph, 81
 - growth, 79, 109
 - projection, 40
 - random, *see* random graph
 - semantic, 81
 - social, 81
 - socio-semantic, 81
 - two-mode, *see* network, bipartite
 - weighted, 82
- novelty, 126

- observationism, 139

- paradigmatic category, 31
- partial order
 - subfield & superfield, 32
- partially-ordered set, 52
- Poisson law, 79
- poset, *see* partially-ordered set
- power-law, 78
- preferential attachment, 79, 97

- Q*-analysis, 34

- random graph
 - Barabasi-Albert model, 79
 - Erdős-Rényi model, 77
 - rewiring, 46
 - small-world, 78
 - Watts-Strogatz model, 78
- reconstruction
 - issues, 9
 - micro-foundation, 75
- reductionism, 134

- selection heuristics, 54
- social cognition, 18, 164
- social distance, *see* distance

- social structure, 10
- society of knowledge, 9
- stigmergence, 156
- structural equivalence, 24

- taxonomy, 22
 - Aristotelian, 31
 - evolution, 57
 - folk, 18
- transitivity, 89
- tree, 31

- zebrafish, 19

Abstract

Agents producing and exchanging knowledge are forming as a whole a socio-semantic complex system. Studying such knowledge communities offers theoretical challenges, with the perspective of naturalizing further social sciences, as well as practical challenges, with potential applications enabling agents to know the dynamics of the system they are participating in. The present thesis lies within the framework of this research program. Alongside and more broadly, we address the question of reconstruction in social science. Reconstruction is a reverse problem consisting of two issues: (i) deduce a given high-level observation for a considered system from low-level phenomena; and (ii) reconstruct the evolution of some high-level observations from the dynamics of lower-level objects.

In this respect, we argue that several significant aspects of the structure of a knowledge community are primarily produced by the co-evolution between agents and concepts, i.e. the evolution of an epistemic network. In particular, we address the first reconstruction issue by using Galois lattices to rebuild taxonomies of knowledge communities from low-level observation of relationships between agents and concepts; achieving ultimately an historical description (*inter alia* field progress, decline, specialization, interaction – merging or splitting). We then micro-found various stylized facts regarding this particular structure, by exhibiting processes at the level of agents accounting for the emergence of epistemic community structure. After assessing the empirical interaction and growth processes, and assuming that agents and concepts are co-evolving, we successfully propose a morphogenesis model rebuilding relevant high-level stylized facts. We finally defend a general epistemological point related to the methodology of complex system reconstruction, eventually supporting our choice of a co-evolutionary framework.

Keywords: *Complex systems, social cognition, reconstruction, applied epistemology, Galois lattices, taxonomies, dynamic social networks, mathematical sociology, cultural co-evolution, scientometrics, knowledge discovery in databases.*

“L’Ecole Polytechnique n’entend donner aucune approbation, ni improbation, aux opinions émises dans cette thèse, ces opinions doivent être considérées comme propres à leur auteur”