

Building Optimal Macroscopic Representations of Complex Multi-agent Systems

Application to the Spatial and Temporal Analysis of International Relations Through News Aggregation

Robin Lamarche-Perrin¹(✉), Yves Demazeau², and Jean-Marc Vincent¹

¹ LIG, Université Grenoble Alpes, 38000 Grenoble, France
{Robin.Lamarche-Perrin, Jean-Marc.Vincent}@imag.fr

² LIG, CNRS, 38000 Grenoble, France
Yves.Demazeau@imag.fr

Abstract. The design and the debugging of large-scale MAS require abstraction tools in order to work at a macroscopic level of description. Agent aggregation provides such abstractions by reducing the complexity of the system's microscopic representation. Since it leads to an information loss, such a key process may be extremely harmful for the analysis if poorly executed. This paper presents measures inherited from information theory to evaluate abstractions and to provide the experts with feedback regarding the quality of generated representations. Several evaluation techniques are applied to the spatial and temporal aggregation of an agent-based model of international relations. The information from on-line newspapers constitutes a complex microscopic representation of the agent states. Our approach is able to evaluate geographical abstractions used by the domain experts in order to provide efficient and meaningful macroscopic representations of the world global state.

Keywords: Large-scale MAS · Agent aggregation · Macroscopic representation · Information theory · Geographical and news analysis

1 Introduction

Because of their increasing size, complexity, and concurrency, current multi-agent systems (MAS) can no longer be understood from a microscopic point of view. Design, debugging and optimization of such large-scale distributed applications require tools that proceed at a higher representational level by providing insightful abstractions regarding the system's dynamics. Among abstraction techniques (*e.g.*, dimension reduction, subsetting, segmentation, clustering [1]), this paper focuses on *data aggregation*. It consists in losing some information regarding the agent level to build simpler – yet meaningful – macroscopic representations.

Such a process is not trivial for the interpretation of the data by the observer. In particular, unsound aggregations may lead to critical misrepresentations of the

MAS behavior. Hence, one needs to determine what are the *good* abstractions and how to properly use them. At each stage of MAS development, aggregation processes should be carefully monitored and feedback should be provided regarding the quality of the generated macroscopic representations. A simple example can demonstrate how critical the aggregation process can be. Figure 1 shows two groups of agents that are simplified by two abstract entities with an averaged behavior. Intuitively, group A constitutes a *good* abstraction since the induced global behavior is relatively similar to the microscopic one, unlike group B. Hence, in order to scale-up, aggregation of redundant information should be encouraged to reduce the representation complexity (group A), but details regarding heterogeneous behaviors should be preserved in order to control the information loss and proceed to a sound analysis (group B).

Very little work has been done in the MAS community to quantify such aggregation properties. The main contribution of this paper consists in introducing measures from information theory (Kullback-Leibler (KL) divergence [2] and Shannon entropy [3]) to clarify the notion of *good* aggregation. From these measures, we provide generic feedback techniques and an algorithm that builds multiresolution representations out of hierarchically organized MAS. These techniques and algorithms are applied to the agent-based modeling of international relations: agents represent countries, and their behavior is extracted from on-line newspapers. Geographers exploit multilevel aggregates to build statistics regarding world areas. We show how these geographical abstractions should be used to better understand the system states and its evolution through time.

Section 2 presents the work related to the main concern of this article. Section 3 presents the agent-based model of the ANR CORPUS GEOMEDIA application. Sections 4 and 5 introduce KL divergence and the size of representations to respectively estimate *information loss* and *complexity reduction*. Section 6 shows how these measures can be combined to identify *optimal* aggregations and to build multiresolution representations of hierarchically organized MAS. Section 7 applies

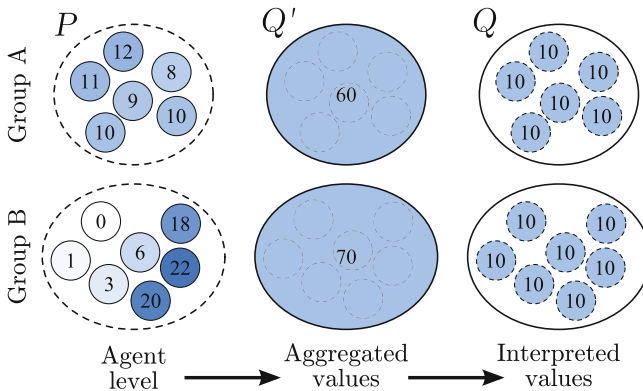


Fig. 1. Averaging the behavior of groups of agents may reduce the redundant information (group A) or lead to an undesired information loss (group B)

these aggregation techniques to the time dimension in order to provide macroscopic representations of the system’s dynamics.

2 Related Work

Aggregation can take place in every stage of the MAS development: from its design to its use. Even if abstraction techniques may differ, each stage should carefully take into consideration the quality of the provided aggregations. First, from a software perspective, this section shows that very few research efforts have been done to tackle this issue. (1) Most classical simulation platforms and monitoring systems do not even provide the user with abstraction tools; (2) some do handle the issue, but are still at an early stage of thought. Secondly, on a theoretical aspect, this section explains why classical techniques (*e.g.* data clustering, graph analysis) are not entirely satisfying to build consistent abstractions.

In a comprehensive survey of agent-based simulation platforms [4], Railsback *et al.* evaluate some tools by testing classical features of MAS modeling and analysis. Unfortunately, the abstraction problem is not tackled by this survey, thus indicating that such considerations are seldom if ever taken into account. Most platforms (Java Swarm, Repast, MASON, NetLogo and Objective-C Swarm) are limited to the microscopic simulation of agents. Railsback warns about the lack of “a complete tool for statistical output” in these platforms [4]. The provision of global views on the MAS macroscopic behavior thus constitutes an on-going research topic. Some tools for large-scale MAS monitoring however address this issue. For example, in some debugging systems, abstractions are used to reduce the information complexity of execution traces; however, they are either limited to the simplification of agents internal behavior, and do not tackled multi-agent organizational patterns [5], or they are provided without any feedback regarding their quality for the analysis [6, 7].

Some techniques from graph analysis and data clustering build groups of agents out of their *microscopic properties* (see for example [8–10]). Such considerations may meet ours from a theoretical point of view, but the approach presented in this report supports a very different philosophy: *abstractions should be built regarding some macroscopic semantics*. We claim that, to be meaningful, the aggregation process needs to rely on exogenous high-level abstractions defined by the experts. Hence, our approach should rather be related to studies on multilevel agent-based models [11]. These works openly tackle the abstraction problem by designing MAS at several organizational levels according to the expert definitions. Such approaches aim at reducing the computational cost of simulations depending on the expected level of detail. The algorithm and measures presented in this paper may provide a formal and quantitative framework to such researches.

To conclude, aggregation techniques should be more systematically implemented on MAS platforms in order to handle large-scale systems. They should combine consistent macroscopic semantics from the experts and feedback regarding the abstractions quality. In this paper, abstractions used by geographers are evaluated according to their information content.

3 Agent-Based Modeling of International Relations

This section presents the GEOMEDIA agent-based model. It consists in the microscopic representation of countries by *agents* and the macroscopic representation of world geographical areas by *groups of agents* and by *organizations*.

3.1 Microscopic Data: The Agent Level

Let A be a set of agents constituting the MAS microscopic level. Visualization tools aim at displaying and explaining the properties of these agents: their behavior and internal states, the events they are associated with, the messages they exchange, and so on. Given a variable v that expresses such properties, the set of values $\{v(a)\}_{a \in A}$ constitutes the *microscopic representation* of the system (illustrated by distribution P in Fig. 1).

The ANR CORPUS GEOMEDIA project¹ is interested in the analysis of world international relations through a media point of view. This project is conducted in collaboration with geographers and media experts from the CIST (*Collège International des Sciences du Territoire*, Paris). In that context, we make the assumption that citations or co-citations of countries, within news, are good indicators to represent and understand their political, economical and cultural relations. For example, we may assume that an often-cited country is likely to politically interact with the newspaper country. Our agent-based model has two dimensions:

- The agents of the model represent the $|A| = 193$ United Nation member states, selected by geographers depending on their significance for the analysis of international relations.
- The temporal dimension contains $|T| = 90$ weeks, from the 3rd of May 2011 to the 20th of January 2013. This preliminary aggregation to the week level aims at reducing the chaotic variations of the day level and focusing on the more significant variations related to media events.

The experiments presented in this paper focus on a very basic variable: the number of articles that cite a country during a given time period. We use the 59,234 articles published by the “world” RSS flow of *The Guardian*² during the analyzed period and stored in the GEOMEDIA database. For each article, we look for the occurrences of the country names, the country adjectives, and the inhabitants names (*e.g.*, “Spain”, “Spanish”, and “Spaniard(s)” for the **S**pain agent). Thus, for each agent a and time period t , we count the number of articles $v(a, t)$ that “cite a during t ”. A total of 138,811 citations have been found within the dataset, distributed within 77 % of the articles (3 citations/article in average if we set aside the 23 % that contain no citation at all).

¹ Founded by the French *National Agency for Research* (ANR-GUI-AAP-04). See the dedicated website for details: <http://geomedia.hypotheses.org/>.

² <http://www.theguardian.com/world>

In order to spot critical aspects of the international systems, geographers are interested in detecting significant events in the news. Such events correspond to unexpected values of the variable according to the following hypothesis: *the citation numbers of countries are homogeneous through time*. In that sense, the marginal values of the dataset give the expected citation number. For an agent a and a time period t , we thus expect the observed value $v(a, t)$ to be close to:

$$v^*(a, t) = \frac{v(a, T) v(A, t)}{v(A, T)}$$

where $v(a, T)$ is the citation number of agent a during the whole observation period T ; $v(A, t)$ is the total number of citations, regarding all agents in A , during the time period t ; and $v(A, T)$ is the total number of citations within the dataset. A media event thus correspond to a high observed value $v(a, t)$ compared to the expected value $v^*(a, t)$.

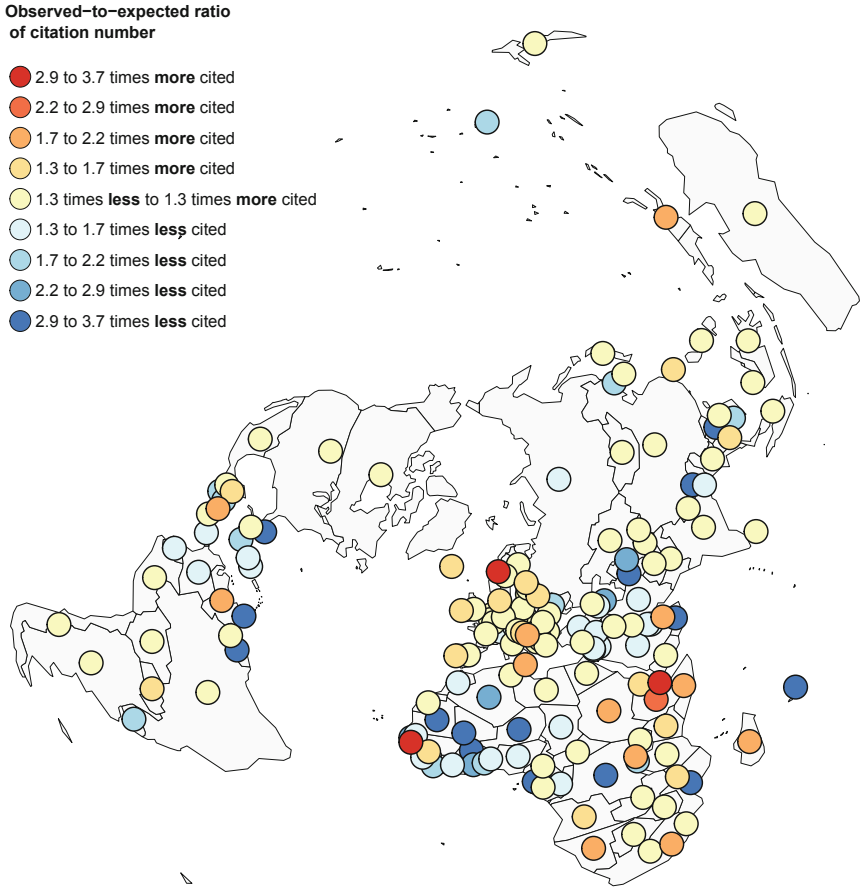


Fig. 2. Observed-to-expected ratio of countries citation within the articles published by *The Guardian* during the month of July 2011 (red circles indicate media events) (Color figure online)

Figure 2 above displays the *observed-to-expected ratio of citation number* $v(a, t)/v^*(a, t)$ for each country $a \in A$ during $t =$ “the month of July 2011”. A detailed survey of this map allows to identify geographical areas that have been unexpectedly over-cited during this time period: at the national level (*e.g.*, Norway, Djibouti, Guinea-Bissau) and at higher levels, *i.e.* for groups of countries (*e.g.*, Europe, Horn of Africa). However, the quantity of information displayed in such a microscopic representation makes it quite hard to read. In particular, the visual clutter in dense areas prevents the proper interpretation of data. To overcome this difficulty, Figs. 3 and 4 propose to focus on areas of particular interest (resp. Europe and Africa). In the following sections, we will focus on two particular events that occurred in these geographical areas:

1. The observed citation number of the Norway agent is 3.7 times higher than expected (see Fig. 3). This is explained by the terrorist attacks that occurred in Norway the 22th of July 2011³. This event belongs to the national level and thus constitutes a *microscopic event* within the system’s spatial dimension.
2. Countries of the Horn of Africa also present unexpected citation numbers (from 1.9 times to 3.4 times the expected value for Rwanda, Sudan, Somalia, Ethiopia and Djibouti, see Fig. 4). This is explained by the food crisis that

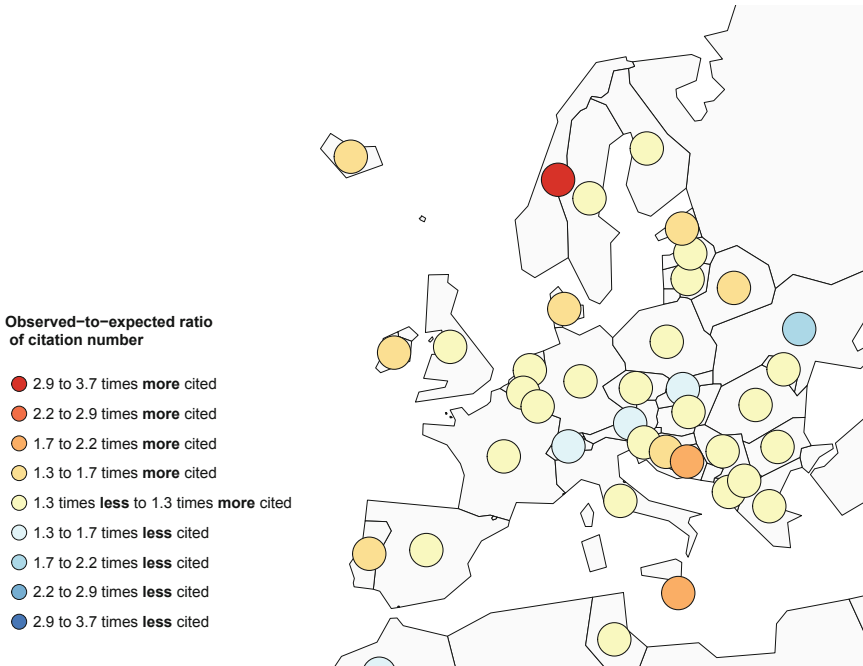


Fig. 3. Observed-to-expected ratio of citation number (zoom on European countries with a national event detected in the Norway agent)

³ http://en.wikipedia.org/wiki/2011_Norway_attacks

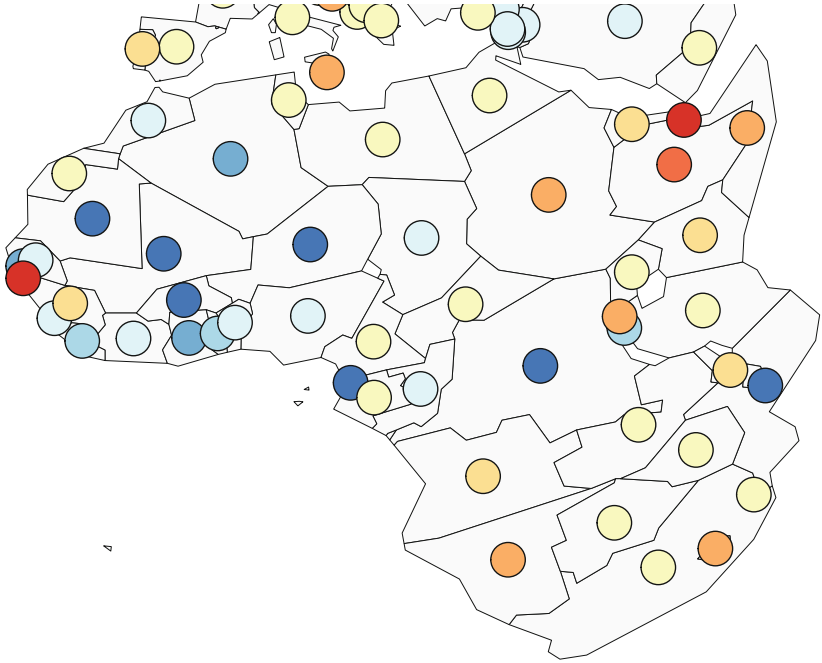


Fig. 4. Observed-to-expected ratio of citation number (zoom on African countries with an international event detected in the Horn of Africa group of agents)

has been reported in this world area starting from the beginning of July 2011⁴. Unlike the previous one, this event is not located at the national level, but regards a group of agents located in a spatially spread out area.

3.2 Macroscopic Data: Groups and Organizations

Even if the maps in Figs. 3 and 4 allow to easily spot the two events we are interested in, they do not manage to give the global overview of the world-wide system that is necessary for an informed analysis. *Data aggregation* aims at resuming the microscopic information to provide such an overview.

A *group* $G \subset A$ is subset of agents that are members of a consistent organizational pattern. It can be interpreted as an *abstract entity* that sums up the behavior of its underlying agents. Hence, groups satisfy a recursive definition: a group is either an agent or a set of groups. Quantitative variables expressing agents properties may be extended on groups according to an aggregation operator: *e.g.*, sum, mean, median, extrema [1]. In our case, since we work with extensive variables, *i.e.* variables that are proportional to the aggregate size, $v(G, t)$ is defined as the *sum* of the values of the underlying agents (see distribution Q' in Fig. 1):

⁴ http://en.wikipedia.org/wiki/2011_East_Africa_drought

$$v(G, t) = \sum_{a \in G} v(a, t)$$

We define an *organization* O as a set of groups that constitutes a *partition* of the agent set A . Thus, in the scope of this paper, each agent is always a member of one and only one group. The set of group values $\{v(G, t)\}_{G \in O}$ composes a *macroscopic representation* of the system with respect to a given organization. It simplifies the variable distribution, from the detailed microscopic representation (distribution P in Fig. 1) to an aggregated one (distribution Q').



Fig. 5. Observed-to-expected ratio of citation number for the groups of agents defined by the 3rd level of the WUTS hierarchical organization

In order to be consistent with the observer’s background knowledge, groups and organizations should be derived from the structural and semantical properties of the agent space. In our context, the world’s *social*, *political*, and *economic*

organization is used by geographers to represent and explain the data. Moreover, in this paper, we also focus on the world’s *topological* organization in order to be consistent with classic geographical representations. Groups thus aggregate adjacent territories that share a cultural and historical background. In the following experiments, we consider two hierarchical organizations of countries that meet these needs, namely WUTS [12] and UNEP [13]. Such organizations define multilevel nested groups commonly used by geographers to build global statistics regarding world areas, from the microscopic level of agents to the full aggregation (see [12] for a detailed presentation of these multiscale organizations).

As an example, Fig. 5 provides with the observed-to-expected ratio of citation numbers aggregated according to WUTS_3, the 3rd level of the WUTS hierarchy. Because of the data reduction, this map is much easier to analyze than the microscopic one (see Fig. 2). In particular, the food crisis that occurred in the **Horn of Africa** group of agents is resumed by one observed macroscopic value that is globally 1.8 times higher than the expected citation number. In that case, the aggregation macroscopically represents the corresponding event. However, most of the microscopic variations have been suppressed by the aggregation process: for example, the events that occurred in the **Norway** agent are no longer represented. We thus need to control the aggregation process in order to visualize events at different levels depending on their spatial granularity. The following sections present an aggregation technique to automatically build such multiresolution representations of MAS.

4 KL Divergence as a Measure of Organization Quality

When an observer tries to interpret the data that is contained in a macroscopic representation, she necessarily makes an assumption regarding the distribution of the aggregated values over the underlying agents. For example, in Fig. 1, the observer considers that each agent has the same weight in the group. It is thus underlined that aggregated values are *uniformly distributed* over the agents (from Q' to Q). Consequently, some groups are more suitable than others to summarize the microscopic information: using group A seems relevant since P is close to Q , unlike group B. Hence, organizations should be carefully chosen in order to provide accurate abstractions. In particular, they should only aggregate homogeneous and redundant distributions of the displayed variable.

Among classical similarity measures to compare a source distribution P with a model distribution Q , Kullback-Leibler (KL) divergence is of highest interest because of its interpretation in terms of information content. This section shows how it can be exploited to provide feedback regarding the quality of groups and organizations and to ensure their proper interpretation by the observer.

4.1 Formalization and Semantics of KL Divergence

Formally, KL divergence measures the number of bits of information that one loses by using the model distribution Q to find the optimal binary coding of

countries associated to articles, instead of using the source distribution P [2]. In other words, KL divergence estimates the information that is lost by the aggregation process. But more generally, it is a measure of dissimilarity between two probability distributions. Hence, it can be interpreted as a *fitness function* between a source P and a model Q .

In our case, the “uniform hypothesis” is not suitable to interpret an aggregated representation. Indeed, for a given group, countries *do not have* the same weight regarding citation number. For example, the observer may assume that, within the **Northern America** group, the **USA** agent usually accumulates much more citations than the **Canada** and the **Mexico** agents. The aggregated value should thus be interpreted depending on that fact. The marginal values can be used to interpret an aggregated representation Q' and to give the corresponding model Q : the citations associated to a group of agents during a time period are distributed according to the total citation numbers of the underlying agents over the whole dataset. Given an agent a in a group G and a time period t , the interpreted citation number is thus given by the following formula:

$$Q(a, t) = v(G, t) \frac{v(a, T)}{v(G, T)}$$

This interpreted value is then compared to the observed microscopic value: $P(a, t) = v(a, t)$. From the KL formula in [2], we define the *divergence* of a group G (or *information loss*, in bits) as follows:

$$\begin{aligned} \text{loss}(G, t) &= \sum_{a \in G} P(a, t) \log_2 \left(\frac{P(a, t)}{Q(a, t)} \right) \\ &= \sum_{a \in G} v(a, t) \log_2 \left(\frac{v(a, t) v(G, T)}{v(G, t) v(a, T)} \right) \end{aligned}$$

As we assume that aggregated values are thus distributed among underlying agents, a group whose internal distribution is very close to the observed distribution (as group A in Fig. 1) will have a low divergence, and conversely (as group B). Moreover, KL divergence verifies the *sum property* [14], meaning that the divergence of disjoint groups is the sum of their divergences. Therefore, for an organization O , we have:

$$\text{loss}(O, t) = \sum_{G \in O} \text{loss}(G, t)$$

4.2 Divergence is Correlated with the Source of Information

This first experiment aims at showing an essential feature of the aggregation process: its quality depends on the context of the analysis. Figure 6 presents the KL divergence of groups defined by the WUTS_3 macroscopic organization for two different newspapers (*The Guardian* and *The New York Times*) that have been observed during the month of July 2011. The *darker* a group is, the *higher* its

KL divergence is, the more *heterogeneous* its internal distribution is. Such groups should not be used for aggregation since they induce a misleading interpretation of the data. In this case, the real microscopic representation significantly diverges from the macroscopic model, making these groups unsuitable for the analysis. On the contrary, *bright* groups of countries constitutes *good* abstractions in terms of information content. The aggregated representation they provide regarding the corresponding geographical area fits with the microscopic data and can thus be properly interpreted by the observer.

In the case of *The Guardian* (cf. Fig. 6(a)), the groups with a high divergence are the location of microscopic events that can be spotted in Fig. 2. In these cases, the suppression of the corresponding microscopic variations induces a significant information loss. Divergence thus indicates heterogeneous behaviors in lower levels that should be detailed in order to reveal significant microscopic events. In the case of the *The New York Times* (cf. Fig. 6(b)), the WUTS_3 groups have not the same divergence than in the previous case. First, divergence is globally higher, thus indicating a more heterogeneous microscopic behavior. This newspaper should then be analyzed at a lower level of representation than *The Guardian*. Moreover, events are not reported in the same way, or with the same intensity, depending on the newspaper editorial policies.

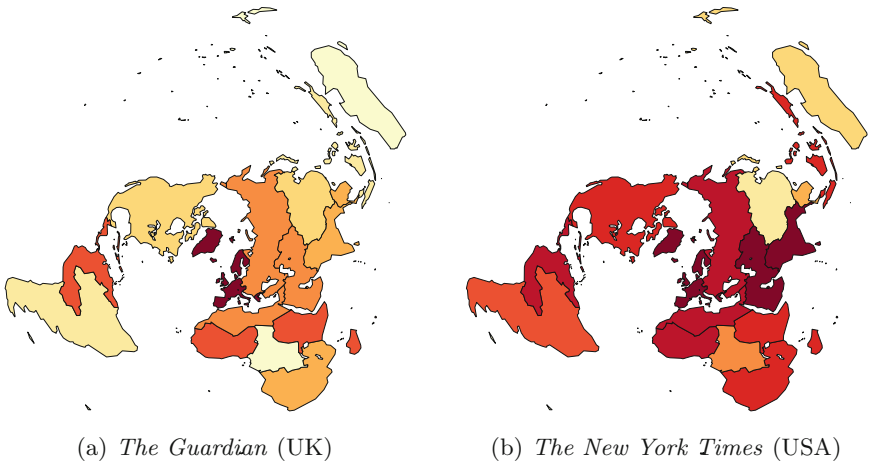


Fig. 6. Spatial variations of the KL divergence for groups of the WUTS_3 organization (the darker, the higher)

We do not aim at making explicit the various positive and negative factors explaining the citation number (*e.g.*, geographical, cultural, historical factors [15, 16]), but at showing that groups should be chosen with respect to the dataset. In our case, this is partly correlated with the source of the information. As a consequence, if an analyst uses distributed probes to observe a MAS, she does not want to use only one abstraction pattern to summarize the information.

This is consistent with the *subjectivist* account of emergence, according to which emergent phenomena strongly rely on the observation process [17].

4.3 Divergence of Groups Varies Over Time

Figure 7 presents the time variation of the KL divergence (thick red line) of the Northern America group (compounded of the USA, the Canada and the Mexico agents). Each value has been computed at the week level, by comparing the interpreted and the observed citation numbers for the countries of this group. The graph shows that a group with a globally low divergence through time can nonetheless be the source of significant information losses during specific time periods. For example, the highest peak which appears in October 2012 is explained by the US presidential elections: during this time period, the USA agent accumulated much more citations than usually, whereas the Canada and the Mexico agents did not. Consequently, the use of the Northern America group to represent events led to a massive information loss. As a result, the choice of the representation level should also fit with the analyzed time period.

Moreover, the graph in Fig. 7 shows that the divergence variations are not strictly correlated with the variations of the analyzed variable (dashed blue line): an increasing of the observed-to-expected ratio of citation number does not implies an increasing of the divergence, and conversely. Hence, the citation number is not a sufficient criterion to evaluate the information content of organizations, by contrast with divergence.

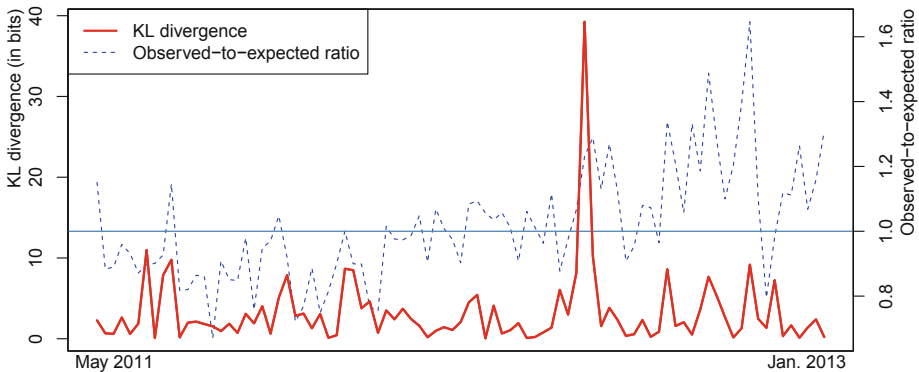


Fig. 7. Time variation of the KL Divergence and the observed-to-expected ratio of citation number of the Northern America group for *The Guardian* (Color figure online)

4.4 Divergence is Correlated with the Shape of Groups

The purpose of this third experiment is to compare two mesoscopic agent organizations: WUTS_2 and UNEP_reg (see Fig. 8). First, a global comparison indicates

which organization minimizes the KL divergence. In order to compare the results from different newspapers, the information loss induced by organizations is normalized by the total citation number of the corresponding newspaper (in the following array, “b/c” stands for “bits/citation”):

	<i>The Vancouver Sun</i>	<i>The Daily Mail</i>	<i>The Ph. Daily Inquirer</i>
WUTS_2	1.80 b/c	1.46 b/c	2.07 b/c
UNEP_reg	1.57 b/c	1.51 b/c	2.26 b/c

It appears that, both for *The Daily Mail* and *The Philippine Daily Inquirer*, divergence is slightly lower for the WUTS_2 organization than for the UNEP_reg organization. Hence, if one should choose between these two, WUTS_2 should be preferred. However, for *The Vancouver Sun*, UNEP_reg is better. Once again, abstractions should then be chosen with respect to the source of information.

We can perform a more subtle analysis in order to determine the groups optimal shape. For example, we notice in Fig. 8 that $U22 = W22 \cup \text{Mexico}$ and $W21 = U21 \cup \text{Mexico}$. Hence, one may ask “what is the best location of the Mexico agent?” Should it be aggregated with the Northern America group ($W21/U21$) or with the Latin America group ($W22/U22$)? For *The Daily Mail*, we have:

$$\text{loss}(W21) + \text{loss}(W22) = 0.048 \text{ b/c} < 0.055 \text{ b/c} = \text{loss}(U21) + \text{loss}(U22)$$

Therefore, the observed-to-expected ratio of citation number of the Mexico agent is closer to those of the Northern America group. Mexico should be aggregated accordingly. This technique allows to evaluate the geographical abstractions used by the experts in terms of information content and to choose their optimal shape for the macroscopic analysis of a given dataset.

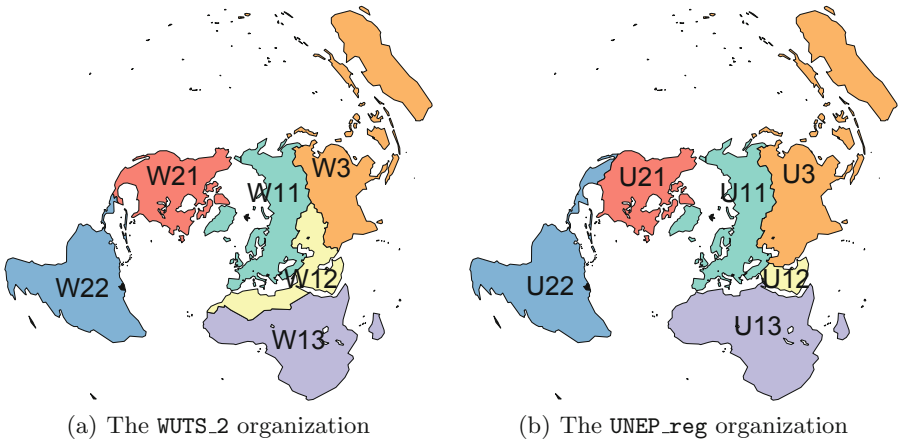


Fig. 8. Two organizations of the agent space in six similar (but not equivalent) groups: locations of the Northern Africa, the Western Asia and the Mexico subgroups differ

5 The Complexity Reduction Induced by the Aggregation

The information content is never increased by the aggregation process: for any pair of disjoint groups, we have: $\text{loss}(G_1 \cup G_2) \geq \text{loss}(G_1) + \text{loss}(G_2)$. Hence, if we only rely on KL divergence, the more detailed representation is always the best one. This is why we need a measure that also expresses what one *gains* by aggregating the microscopic data. To do so, this section presents two measures of *complexity reduction*. They estimate the information quantity that one saves by encoding a group G rather than its underlying agents:

$$\text{gain}(G) = \left(\sum_{a \in G} Q(a) \right) - Q(G)$$

where Q estimates the quantity of information needed to represent the agent a or the group G .

5.1 Number of Encoded Values

One way of measuring information quantities consists in estimating the number of bits needed to encode the values of a given representation. We may assume that it is constant for each agent or group: $Q(a) = Q(G) = q$, where q depends on the data type of the encoded values. Hence, for a group G , we have:

$$\text{gain}(G) = (|G| - 1) \times q$$

This function gives a basic complexity measure that fits well with classic visualization techniques (as for the maps in this paper) since the number of displayed values defines the granularity of the visualization. For example, according to the map expected complexity, the user can determine the number of groups that should be displayed. Figure 9 gives the organization size (number of groups) and the associated gain of each organizational level of the WUTS hierarchy.

Organization	$ O $	$\text{gain}(O)/q$
WUTS_0	1	192
WUTS_1	3	190
WUTS_2	7	186
WUTS_3	17	176
WUTS_4	36	157
WUTS_5	193	0

Fig. 9. Number of encoded values and complexity reduction of the six organizational levels of the WUTS hierarchy

However, all groups do not contain the same number of agents. Thus, Fig. 10 gives, for each level of the WUTS hierarchy, the size (number of agents) of three disjoint high-level groups of countries: **Euro-Africa**, **Americas** and **Asia-Pacific**.




Fig. 10. Number of encoded values associated to the three groups of the WUTS_1 level

The user may want to adapt the representational level of these three groups depending on the amount of detail she expects for the corresponding geographical areas. The following section presents a criterion that automatically combines KL divergence and complexity reduction to adapt the size of groups depending on their quality, thus leading to multiresolution organizations.

5.2 Shannon Entropy

The number of encoded values only depends on the groups partitioning proposed by a given organization. In contrast, Shannon entropy also depends on the *variable distribution*. It is a classical complexity measure that is consistent with KL divergence (it can be defined as “the divergence from the uniform distribution”). Briefly, entropy evaluates the quantity of information needed to encode the countries associated to *each citation* (and not to encode the citation number associated to *each agent*). Based on Shannon’s formula [3], we define the *entropy reduction* (or *gain*, in bits) of a group G as follows:

$$\text{gain}(G, t) = v(G, t) \log_2 v(G, t) - \sum_{a \in G} v(a, t) \log_2 v(a, t)$$

The choice of either one of these two complexity measures depends on the performed analysis. *Shannon entropy* should rather be used for the visualization of individuated citations, whereas *the number of encoded values* is more consistent with the visualization of aggregated values. In any case, the techniques presented in this paper are meant to be generic. They can be used with any complexity measure as long as it fits with some basic algebraic properties (see [18] for details).

6 Multiresolution Representation of Spatial Systems

As a conclusion to previous sections, finding a *good* organization relies on two aspects: complexity reduction (or *gain*), quantifying the granularity of the macroscopic representation, and KL divergence (or *loss*), quantifying the amount of information that has been lost during the aggregation process. Choosing an organization thus consists in finding a compromise between these two aspects.

6.1 Parametrized Information Criterion

We define a *parametrized Information Criterion* to express the trade-off between complexity reduction and information loss of a group G :

$$\text{pIC}(G, t) = p \times \frac{\text{gain}(G, t)}{\text{gain}(A, t)} - (1 - p) \times \frac{\text{loss}(G, t)}{\text{loss}(A, t)}$$

where $p \in [0, 1]$ is a parameter used to balance the trade-off. For $p = 0$, maximizing the pIC is equivalent to minimizing the loss: the observer wants to be as precise as possible (microscopic level). For $p = 1$, she wants to be as simple as possible (full aggregation). When p varies from 0 to 1, a whole class of nested organizations arises. The observer has to choose the ones that fulfill her requirements, between the expected amount of details and the computational resources available for the analysis.




Fig. 11. Comparison of the ratio of information loss and the ratio of complexity reduction (logarithmic scales) for the groups of the WUTS hierarchy applied to the spatial data presented in Fig. 2 (Color figure online)

Figure 11 represents the groups of agents of the WUTS hierarchy (squares) depending on the two criteria that have been previously defined: the ratio of *KL divergence* ($\text{loss}(G, t)/\text{loss}(A, t)$) and the ratio of *encoded values reduction* ($\text{gain}(G, t)/\text{gain}(A, t)$). In this plot, quality groups are easily spotted:

- The closer to the *bottom right* corner the group is (red squares), the higher the information loss is relatively to the complexity reduction. This is for example the case of the **Northern Europe** group (W1111): the unexpected citation number of the **Norway** agent makes the group very heterogeneous (see events described in Subsect. 3.1). Higher-level European groups (W111 and W11) also induces a significant information loss (they are close to the right side). Thus, avoiding such groups during the aggregation ensures that we preserve the details regarding this significant microscopic variation.
- On the contrary, the closer to the *top left* corner the group is (green squares), the more the information loss is compensated by the complexity reduction. This is the case of the **Americas** group and the **South Pacifica** group (resp. W2 and W32). This indicates that these representation levels are particularly interesting to provide with a synthetic view of the system. These groups indeed correspond to homogeneous geographical areas where no significant event has occurred during the observed time period (see map in Fig. 2).
- The **Horn of Africa** group (W132) has a better gain/loss ratio than the higher-level **Sub-Saharan Africa** group (W13). This indicates that, if some details are necessary to analyze the events occurring in Africa, the **Horn of Africa** group can however be described as a whole, without giving more details regarding this particular area. Hence, by choosing groups depending on their gain/loss ratio, the observer can represent the system with several spatial granularities in order to perfectly fit with the microscopic data.

This method allows to spot interesting groups of agent to build a synthetic but consistent macroscopic representation of the system. The rest of this section proposes an algorithm to automatize this evaluation process and to find the combinations of groups (the organizations) that jointly optimize the two criteria.

6.2 Organizations Within a Hierarchy

Given a value of the trade-off parameter p , optimal organizations are those that maximize the parametrized information criterion. Clustering techniques using *gain* and *loss* measures as distances could find such optimal partitions. However, results may have very little meaning for the MAS analysis since agents would be aggregated regardless of their location within the system. In contrast, we assume that, in most spatial MAS, there is a correlation between topology and behavior. Hence, we propose that organizations should fit with the topological constraints defined by the domain experts. In other agent-based applications, such constraints may also be derived from *semantic* properties of the system (and not necessarily *topological* properties).

In this section, we consider hierarchically organized MAS. A *hierarchy* H is a set of nested groups, defined from the microscopic level (each agent is a

group) to the whole MAS (only one group). The number of possible multiresolution organizations within such a hierarchy *exponentially* depends on the number of groups and the number of levels. For UNEP (196 groups arranged in 4 levels) and WUTS (231 groups arranged in 6 levels), we respectively have 1.3×10^6 and 3.8×10^{12} possible organizations. Finding the best one can thus be computationally expensive in case of large-scale systems. The algorithm below finds topologically-consistent organizations that maximize our parametrized information criterion. Its complexity *linearly* depends on the number of groups in the hierarchy (respectively 196 and 231 groups) by doing a depth-first search within the branches of the hierarchy. Indeed, according to the *sum property* [14] of the defined information-theoretic measures, each branch can be independently evaluated (see [19] for details).

Algorithm that linearly finds optimal organizations within a hierarchy

Require: A hierarchy H and a trade-off parameter p in $[0, 1]$.

Ensure: An organization made of groups in H that maximizes the pIC.

```

1: function FINDOPTIMALORGANIZATION( $H, p$ )
2:   if  $H$  contains only one group  $G$  then return  $\{G\}$ 
3:    $G \leftarrow$  biggest group of  $H$ 
4:   bestMicroOrganization  $\leftarrow \emptyset$ 
5:   for each direct subhierarchy  $S$  of  $H$  do
6:      $aux \leftarrow$  FINDOPTIMALORGANIZATION( $S, p$ )
7:     bestMicroOrganization  $\leftarrow$  UNION(bestMicroOrganization,  $aux$ )
8:   end for
9:   if pIC of  $\{G\} >$  pIC of bestMicroOrganization then return  $\{G\}$ 
10:  else return bestMicroOrganization
11: end function

```

6.3 Hierarchical Aggregation to Build Spatial Macro-Representations

The above algorithm has been ran on the WUTS hierarchy for the articles published by *The Guardian* during July 2011 (see Fig. 2). The plot in Fig. 12 gives the complexity reduction and the information loss associated to the organizations provided by the algorithm depending on the trade-off parameter p specified by the observer. For $p = 0$, the optimal organization corresponds to the microscopic representation (no information loss and no complexity reduction). As p increases, groups of countries are chosen within the WUTS hierarchy in order to focus on significant events. The observer can adjust the granularity of the generated representation depending on the expected level of detail. Figures 13 and 14 present two organizations respectively preserving at least 50% and 70% of the microscopic information: $\text{loss}(O, t)/\text{loss}(A, t) < 0.5$ and $\text{loss}(O, t)/\text{loss}(A, t) < 0.3$ respectively for $p < 0.86$ and $p < 0.43$ (see Fig. 12).

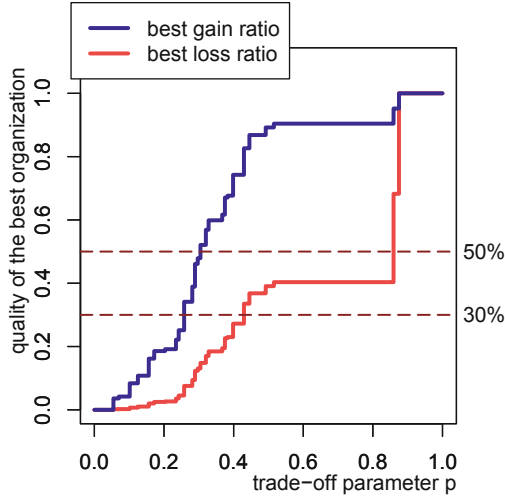


Fig. 12. Ratio of complexity reduction (*gain*) and information *loss* of the optimal organizations of the spatial data presented in Fig. 2 according to the trade-off parameter

The map in Fig. 13 represents the observed-to-expected ratio according to 17 groups of countries, two of which are high-level groups where the observed citation number is quite close to the expected value (the **Americas** group and the **Asia-Pacific** group). No significant event has been spotted at this level of detail. By contrast, two events are immediately highlighted in Europe and in Africa. They correspond to the most unexpected citation numbers for which an aggregation would induce a misleading information loss (at most 40% of information loss when $p < 0.85$ and at least 68% when $p > 0.85$, see Fig. 12).

1. The map in Fig. 13 displays the national details regarding agents of the **Northern Europe** group. Among them, the observed citation number of the **Norway** agent is 3.7 times higher than expected. The observer is thus informed that a significant event took place at the national level during July 2011 (the two terrorist attacks of the 22th, cf. Subsect. 3.1).
2. This map also displays some details regarding the **Sub-Saharan Africa** group, but only at a mesoscopic level constituted of 4 intermediary groups. Among them, the observer notices that the citation number of the **Horn of Africa** group is 1.8 times higher than expected. As the national details are not represented for this particular group, the observer may consider that this aggregated value is – at least at first glance – a good approximation of the underlying values. She concludes that the observed-to-expected ratio of citation number is *uniformly high* in the Horn of Africa. This group thus highlights an event that occurred in an extended geographical area (the food crisis that has been declared at the beginning of July 2011, cf. Subsect. 3.1).

The aggregation algorithm thus allows to highlight significant events that occur at different granularities of the system’s spatial organization by building readable multiresolution representations out of the hierarchical structure. This map thus constitutes a reasonable “first approximation” to describe the news published by *The Guardian* during the month of July 2011. However, depending on the analyzed events and the expected level of detail, the observer may adapt the granularity of such a representation by adjusting the trade-off parameter (see the map presented in Fig. 14).

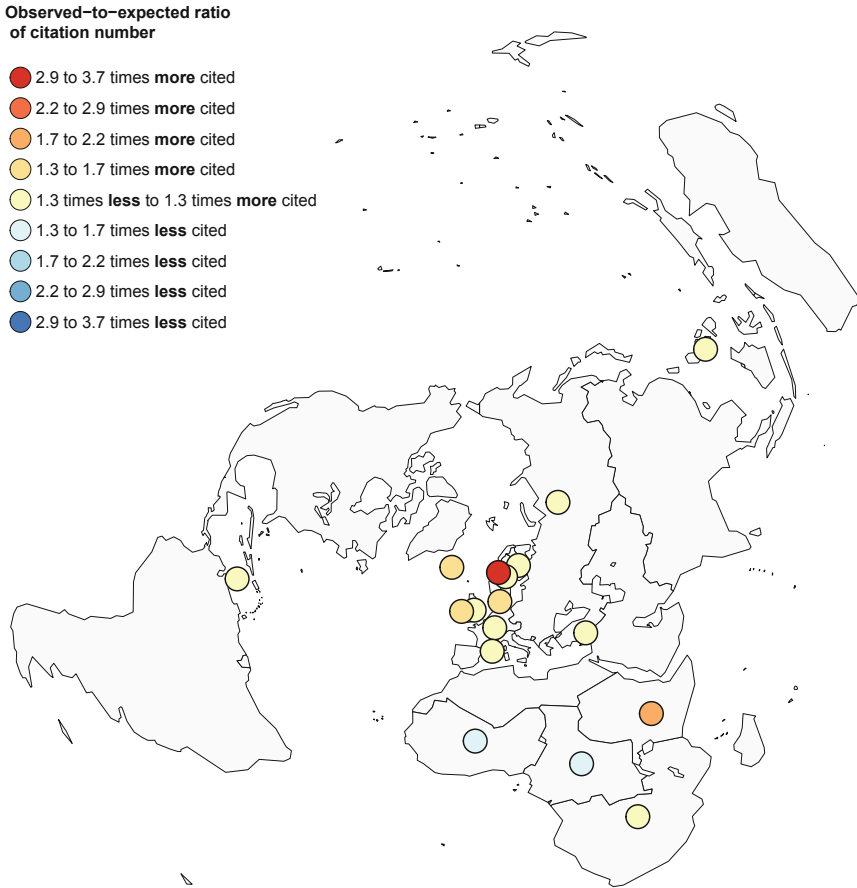


Fig. 13. Optimal geographical organization preserving at least 50 % of the microscopic information ($p < 0.86$)

The map below is a little bit more detailed than the previous one, in particular for countries of the **Asia-Pacific** group and those of the **Western Africa** group. Some other – less significant – microscopic events are thus represented:

3. The severe floods that occurred in Thailand at the end of July 2011⁵.
4. The development cooperation project that began the 16th of July between the European Union and the Republic of Guinea-Bissau.

However, in this second map, the **Horn of Africa** group of agents is still aggregated. This is only when the observer asks for at least 83% of the microscopic information ($p < 0.39$) that the algorithm provides the national details regarding this geographical area. In that way, the algorithm can adapt the generated representations to the observer's requirements.

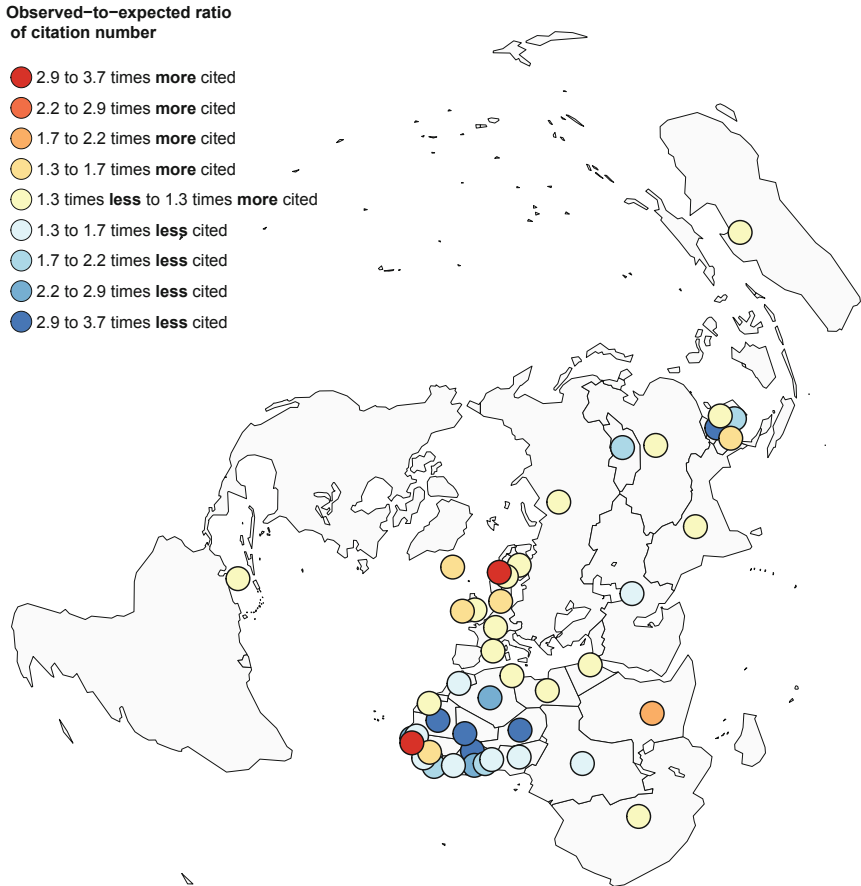


Fig. 14. Optimal geographical organization preserving at least 70% of the microscopic information ($p < 0.43$)

⁵ http://en.wikipedia.org/wiki/2011_Thailand_floods

7 Generalization to Temporal Aggregation

The time series in Fig. 15 provides with the week-level variations of the observed-to-expected ratio of citation number of the **Greece** agent by *The Guardian* between the 3rd of May 2011 and the 20th of January 2013. Peaks of unexpected citation number reveal significant events in Greece recent history. In the following, we will focus on three particular events:

1. The peak of citation that appears at the beginning of November 2011 is explained by the announcement the 31st of October of a referendum regarding the setting up of an austerity plan to reduce the Greek public debt. This announcement is widely reported by the media. On the 4th of November, the Minister of Finance announces the referendum withdrawal and the Prime Minister George Papandreou arranges, on the same evening, a vote of confidence in the Parliament that could lead to his resignation.
2. The peak that appears in the middle of May 2012 is explained by the failure of the legislative elections that are held the 6th of May and concluded the 16th by the establishment of an interim government until the organization of new elections.
3. The peak that appears at the end of June 2012 is explained by the holding of these second legislative elections on the 17th of June⁶.




Fig. 15. Microscopic time series (week level) of the observed-to-expected ratio of citation number of the **Greece** agent by *The Guardian*

⁶ See the Wikipedia page dedicated to the Greek government-debt crisis to get more details regarding the chronology of these political events:

http://en.wikipedia.org/wiki/Greek_government-debt_crisis

In this section, the aggregation technique is applied to the system's temporal dimension. In this case, macroscopic events refer to time periods during which the citation number of a given country (or group of countries) has been much higher than expected. Such periods – that breaks with the system's stable state – may also be defined at different time scales (days, weeks, month, years, and so on). When interpreting an aggregated time period, the observer can use – as for the spatial aggregation – the marginal values to distribute the aggregated citation number over the underlying microscopic time periods (*i.e.*, the week level) depending on the total number of citations on these time periods. For a microscopic time period t in an aggregated time period $T' \subset T$, the interpreted citation number is:

$$Q(a, t) = v(a, T') \frac{v(A, t)}{v(A, T')}$$

This interpreted value is then compared to the observed value: $P(a, t) = v(a, t)$, according to KL divergence:

$$\begin{aligned} \text{loss}(a, T') &= \sum_{t \in T'} P(a, t) \log_2 \left(\frac{P(a, t)}{Q(a, t)} \right) \\ &= \sum_{t \in T'} v(a, t) \log_2 \left(\frac{v(a, t) v(A, T')}{v(a, T') v(A, t)} \right) \end{aligned}$$




Fig. 16. Ratio of complexity reduction (*gain*) and information *loss* of the optimal partitions of the temporal data presented in Fig. 15 according to the trade-off parameter

The optimization of the corresponding *parametrized Information Criterion* (see Subsect. 6.1) can be achieved by the algorithm of Jackson *et al.* [20]. It consists in slicing the time series in intervals that maximize a given fitness function.

We have shown in [19] that this time-aggregation algorithm is part of a larger class of algorithms – including the space-aggregation algorithm proposed in this paper – that consist in computing the optimal partitions of a dataset under some constraints (*e.g.*, hierarchical organization for spatial aggregation, ordered dataset for temporal aggregation). The time-aggregation algorithm of Jackson *et al.* is thus exploited as previously to build multiresolution representations of the agent dynamics.

The plot in Fig. 16 gives the complexity reduction and the information loss of the time partitions provided by the algorithm depending on the trade-off parameter p . Series in Figs. 17 and 18 present two such optimal partitions respectively preserving at least 80 % and 50 % of the week-level information (resp. $p < 0.55$ and $p < 0.89$). The time series in Fig. 17 thus summarizes the microscopic data by aggregating groups of weeks for which the observed-to-expected ratio of citation number is quite homogeneous. Even if the result contains less details, it still provides significant information for the analysis of Greek news. In particular, the three significant aforementioned peaks are highlighted and easily spotted by the observer. Moreover, the variations between the peaks are synthetically represented. This allows to describe the system dynamics according to macroscopic time periods.




Fig. 17. Optimal temporal partition preserving at least 50 % of the microscopic information ($p < 0.55$)

The time series in Fig. 18 provides with an even more aggregated representation of time variations. Only the most significant events are represented:

1. The first peak, corresponding to the announcement the 31st of October of a Greek referendum, is strongly highlighted by the aggregation process.




Fig. 18. Optimal temporal partition preserving at least 80% of the microscopic information ($p < 0.89$)

2. The two peaks of May and June 2012, respectively corresponding to the legislative elections of the 6th of May and the 17th of June, are aggregated together, now constituting a unique homogeneous time period of 7 weeks. The corresponding event is interpreted as “the Greek legislative elections of 2012”, without giving the week-level details of this month-scale event. The aggregation process thus provides consistent temporal abstractions to describe and analyze Greek events at a higher level of representation.
3. This time series also gives an interesting macroscopic information: the citation number of the **Greece** agent has been globally decreasing during the observation period. This could be explained by the declining media interest regarding the Greek crisis after the arrival in news of other economical crises concerning European countries such as Spain and Italy. The aggregation process thus allows to represent the system’s temporal dynamics at several time-scales by highlighting significant macroscopic variations. If we had the sufficient temporal depth, we would be able to identify a year-long time period in Greek history corresponding to “the government-dept crisis” described as a whole.

8 Conclusion and Perspectives

The design and debugging of complex MAS need abstraction tools to work at a higher level of representation. However, such tools have to be developed and exploited with the greatest precaution in order to preserve useful information regarding the system behavior and to guarantee that generated representations are not misleading for the observer. To that extent, this paper focuses on aggregation techniques for large-scale MAS and gives clues to estimate their quality in

term of complexity and information content. They are applied to the geographical and temporal aggregation of international relations through the point of view of on-line newspapers. We show that, by combining information-theoretic measures, one can give interesting feedback regarding the use of abstractions and build multiresolution representations of the dataset that adapt to the effective information content and to the observer's requirements.

We believe that these measures and algorithms can be generalized to a large class of MAS, provided that:

- one can observe and describe the agents microscopic behavior according to several discretized microscopic dimensions (here: space and time);
- one can define measures to express the descriptions *quality* (here: complexity and information content);
- these measures have the *sum property* [14];
- the semantic and topological properties of the aggregated dimensions can be used to provide meaningful abstractions for the domain experts (here: hierarchical organizations and order of time).

Future work will apply these techniques to other dimensions of the analysis: *e.g.* for aggregation of newspapers, thematic aggregation, multidimensional aggregation [19]. Besides this work, we are currently exploiting these techniques for performance visualization of large-scale distributed systems [21]. This kind of application shows that our techniques can be scaled up to one million agents.

Acknowledgement. This work was partially funded by the ANR CORPUS GEOMEDIA project (ANR-GUI-AAP-04). We would like to thank Claude Grasland, Timothée Giraud and Marta Severo for their work on this project; and Lucas M. Schnorr for his close participation to previous work.

References

1. Elmqvist, N., Fekete, J.: Hierarchical aggregation for information visualization: overview, techniques, and design guidelines. *IEEE Trans. Visual Comput. Graphics* **16**(3), 439–454 (2010)
2. Kullback, S., Leibler, R.: On information and sufficiency. *Ann. Math. Stat.* **22**(1), 79–86 (1951)
3. Shannon, C.: A mathematical theory of communication. *Bell Syst. Tech. J.* **27**(379–423), 623–656 (1948)
4. Railsback, S.F., Lytinen, S.L., Jackson, S.K.: Agent-based simulation platforms: review and development recommendations. *Simulation* **82**, 609–623 (2006)
5. Van Liedekerke, M.H., Avouris, N.M.: Debugging multi-agent systems. *Inf. Softw. Technol.* **37**, 103–112 (1995)
6. Búrdalo, L., Terrasa, A., Julián, V., García-Fornes, A.: A tracing system architecture for self-adaptive multiagent systems. In: Demazeau, Y., Dignum, F., Corchado, J.M., Pérez, J.B. (eds.) *Advances in PAAMS. AISC*, vol. 70, pp. 205–210. Springer, Heidelberg (2010)

7. Tonn, J., Kaiser, S.: ASGARD – a graphical monitoring tool for distributed agent infrastructures. In: Demazeau, Y., Dignum, F., Corchado, J.M., Pérez, J.B. (eds.) *Advances in PAAMS. AISC*, vol. 70, pp. 163–173. Springer, Heidelberg (2010)
8. Sharpanskykh, A., Treur, J.: Group abstraction for large-scale agent-based social diffusion models with unaffected agents. In: Kinny, D., Hsu, J.Y., Governatori, G., Ghose, A.K. (eds.) *PRIMA 2011. LNCS*, vol. 7047, pp. 129–142. Springer, Heidelberg (2011)
9. Irvani, P.: Multi-level network analysis of multi-agent systems. In: Iocchi, L., Matsubara, H., Weitzenfeld, A., Zhou, C. (eds.) *RoboCup 2008. LNCS (LNAI)*, vol. 5399, pp. 495–506. Springer, Heidelberg (2009)
10. Peng, W., Grushin, A., Manikonda, V., Krueger, W., Carlos, P., Santos, M.: Graph-based methods for the analysis of large-scale multiagent systems. In: *AAMAS'09, IFAAMAS*, pp. 545–552 (2009)
11. Gil-Quijano, J., Louail, T., Hutzler, G.: From biological to Urban cells: lessons from three multilevel agent-based models. In: Desai, N., Liu, A., Winikoff, M. (eds.) *PRIMA 2010. LNCS*, vol. 7057, pp. 620–635. Springer, Heidelberg (2012)
12. Grasland, C., Didelon, C.: *Europe in the World - Final Report. Volume 1, ESPON Project 3.4.1* (2007)
13. United Nations Environment Programme: *Global Environmental Outlook: environment for development. Volume 4, Nairobi* (2007)
14. Csiszár, I.: Axiomatic characterizations of information measures. *Entropy* **10**(3), 261–273 (2008)
15. Galtung, J., Ruge, M.H.: The structure of foreign news: the presentation of the Congo, Cuba and Cyprus crises in four Norwegian newspapers. *J. Peace Res.* **2**(1), 64–91 (1965)
16. Koopmans, R., Vliegthart, R.: Media attention as the outcome of a diffusion process—a theoretical framework and cross-national evidence on earthquake coverage. *Eur. Sociol. Rev.* **27**(5), 636–653 (2011)
17. Deguet, J., Demazeau, Y., Magnin, L.: Element about the emergence issue: a survey of emergence definitions. *ComPlexUs* **3**, 24–31 (2006)
18. Lamarche-Perrin, R., Vincent, J.M., Demazeau, Y.: Informational measures of aggregation for complex systems analysis. Technical report RR-LIG-026, Laboratoire d'Informatique de Grenoble, France (2012)
19. Lamarche-Perrin, R., Demazeau, Y., Vincent, J.M.: The best-partitions problem: how to build meaningful aggregations. In: *Proceedings of the 2013 IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT'13)*, Atlanta, GA, USA, pp. 399–404. IEEE Computer Society (2013)
20. Jackson, B., Scargle, J.D., Barnes, D., Arabhi, S., Alt, A., Gioumousis, P., Gwin, E., Sangtrakulcharoen, P., et al.: An algorithm for optimal partitioning of data on an interval. *IEEE Signal Process. Lett.* **12**(2), 105–108 (2005)
21. Lamarche-Perrin, R., Schnorr, L.M., Vincent, J.M., Demazeau, Y.: Evaluating trace aggregation for performance visualization of large distributed systems. In: *Proceedings of the 2014 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS'14)*, Monterey, CA, USA (2014)