THE STRUCTURE OF SCIENTIFIC DISCIPLINES. SOME NOTES ON THE EPISTEMOLOGY OF ALGORITHMIC REPRESENTATIONS OF DISCIPLINES

A PREPRINT

Eugenio Petrovich
Department of Economics and Statistics
University of Siena
Siena, Italy
eugenio.petrovich@gmail.com

June 9, 2019

ABSTRACT

Representations of research fields are devices that are used to grasp the inner structure of disciplines. They play a crucial role in the disciplinary life since they ease the fulfillment of many epistemic and social tasks (from systematizing the intellectual content to organizing the cognitive labor of researchers, to manage appointments in the university). In the first part of the paper, I will distinguish between two types of field representations: top-down and bottom-up. The former are based on the assessment of the structure of research fields by experts, whereas the latter are based on computer algorithms. In the second part of the paper, I will focus on these ones and, by taking as an example co-citation maps, I will argue that they capture the latent forces that shape the intellectual action of the researchers contributing to a scientific discipline. These forces have both a negative and positive function (i.e., they are *structuration* forces). The negative function is to constraint the possible intellectual actions of the researchers, the positive is to enable the organized accumulation of knowledge by channeling information and reducing noise.

***K*eywords** Research discipline · Co-citation analysis · Structuration

1 Introduction

Any discipline is characterized by a moment in which its practitioners begin to reflect on the structure of their knowledge domain. What are the sub-areas of the field? What is their ordering principle? Are they distinguished by the research object or the research method? What do they share and what distinguish them?

The answers to questions like these are crucial in the life of a research field from many points of view. First, they are pivotal in the division of the cognitive labor within the field. Second, they promote a systematization of the intellectual content the researchers within the field produce. Third, they simplify the activity of information storage and retrieval, providing the institution demanded to this (such as university libraries) with classificatory information. Fourth, they allow to organize both the institutional layer of the discipline (the university chairs), and its communicative infrastructure (the journals). Fifth, they provide a guide for programming the university curricula which are needed to train and recruit new members in the field. Lastly, they are crucial in the building of the self-identity of the researchers belonging to the field, promoting in this way the academic respectability of the field itself. Now, one of the most important outcomes of this reflexive activity are representations of the field.

Representations are synthetic devices by which the inner structure of the field is made visible and thus it can be discussed or used for practical purposes (e.g., organizing a library catalog). Representations appear in different formats and shapes, as well as in diverse domains (from theory to management). They can consist in a simple, mono-dimensional list of sub-disciplines or appear as an organized tree-like hierarchical structure in which sub-areas are nested one into

the other to form a complex tree. They can be one-page long or reach the extension of encyclopedias or "summae" of the discipline. They can be a list of classificatory labels to pigeonhole books in a library catalog or a system of "scientific areas" ("settori scientifico-disciplinari") for regulating appointments and promotions at the university.

2 Top-down and bottom-up representations

From a very general point of view, all these different types of representation can be considered as attempts to solve the same problem, namely: how can we arrange the "units" that compose the field (being them ideas, concepts, theories) into distinct, relative homogeneous groups (the sub-disciplines or sub-areas)? How can we organize the raw "material" of our field into an ordered structure?

All the traditional representation devices we mentioned above adopt a top-down strategy to solve this problem. They design, based on a variety of considerations, a classificatory scheme (e.g., a sub-disciplinary taxonomy) and then proceed to classifying the units according to the scheme's categories. Depending on the domain of application of the representation device, the consideration on which the taxonomy is based can be either highly theoretical and argumentatively justified (for instance, in the case of expert encyclopedic disciplinary synthesis), or merely based on common sense and institutional inertia (as it happens probably with university management schemes).

In the last years, the development of digital tools and computer algorithms has allowed an entirely new kind of strategy to answer the grouping problem. Instead of a top-down approach, digital methods allow for the first time to implement a bottom-up strategy.

In a bottom-up strategy, we let the units, so to say, organize by themselves. We choose one or more property of the units we want to group, and the algorithm calculates their mutual similarity and dissimilarity relative to those properties. In computer science and statistics, the techniques designed for this purpose are collectively called "clustering techniques". They serve to individuate the underlying structure of a set of "data-points" (i.e., the units we want to classify, considered only in some of their properties). Their output is, classically, of two kinds: dendrogram or maps. In a dendrogram, the structure is visualized as a tree, in which the leaves of the tree correspond to the unities and the different branches correspond to the clusters in which the units can be grouped. The dendrogram communicates clearly a hierarchical structure. In a map, on the other hand, the structure is visualized by placing the items in different point of a two-dimensional plane. Similar items should be placed closer in the map, as in geographical maps. The clusters appear as denser regions of the map. The hierarchical element is less visible than in the dendrogram.

3 Differences between top-down and bottom-up representations

Now, both classical top-down and computer-based bottom-up representations share the same objective: representing the inner structure of a research field. Thus, both produce a similar outcome: a representation of this structure. One can think that they are just different ways to reach the same goal. I will argue that this is not the case: the two representations are very different. Specifically, I think that the bottom-up representations unveil a latent organization of the field that top-down ones rarely grasp.

To understand this, let us focus on the bottom-up representations produced by co-citation analysis. In co-citation analysis, the similarity between two documents (classically, scientific papers) is calculated based on the number of times they are cited together in a certain set of other documents (their "co-citation strength"). To produce a co-citation map, the co-citation strengths of each pair of documents are reported in a square N x N matrix, called adjacency matrix, then their raw values are elaborated statistically (for instance, they are translated into cosine measures of similarity), and lastly, the co-citation matrix is visualized on a 2- or 3-d plane to grasp its structure intuitively. Computer scientists and statisticians have developed different techniques for doing this (e.g., Multi-Dimensional Scaling, MDS, or the VOS algorithm).

Now, co-citation maps succeed in revealing the structure of the field because the co-citation data (the matrix of the co-citation strength) are not randomly distributed. Given the average co-citation rate of all the documents, there are some of them that are cited together more frequently (i.e., above the average), and some of them that are almost never cited together. In this way, the different distribution of the co-citation strengths among the documents confers to the whole matrix a structure, i.e., an organization in terms of denser regions (the clusters) and sparser regions. The visualization algorithms aim to reproduce these distributions by placing closer in the plane frequently co-cited documents and far the seldom co-cited documents.

The crucial thing about co-citation maps is that they are based on thousands of micro-actions by the researchers of the field. These micro-actions consist in the selection of references that each author does when she is writing a paper. Indeed, the raw material from which the co-citation strengths are calculated are exactly the reference lists that each

author produces. Note that each reference list can be considered as a (very) partial representation of the field to which the paper belongs. The paper reflects the entire field in so far as each new contribution must be related, by references, to the existing body of knowledge (the field). A paper without references, by contrast, would not reflect the field since it would not be connected to any other documents.

Now, when we produce the co-citation matrix, what we are doing is to aggregate the thousands of partial perspectives contained in the reference lists of each single paper. It is as if we were collecting and aggregating photographs taken by thousands of different points of view. What may be the outcome of such a process? If the photographs were not coherent one with the other, the result would be an incomprehensible patchwork. We will conclude that they do not represent anything, or perhaps that they represent too different subjects. By contrast, if they coalesce into a coherent picture, we will conclude that they represent an object taken from different perspectives. The coherent picture, obtained by their aggregation, is a faithful representation of the object itself.

Something similar happens with the co-citation analysis. If there is a coherent structure behind the reference lists, this is captured by the co-citation matrix since the scores in the matrix will not be randomly distributed, but they will exhibit a structure. Such a structure consists in the clusters of frequently co-cited documents.

When we map big sets of documents (or scientific journals) with co-citation analysis, it turns out that, indeed, they exhibit a structure. They are not randomly distributed. This may appear obvious: at the end of the day, we would not even start to map something if we would not expect to find something to be mapped (a structure). On the other hand, this is a result bringing deeper implications.

To understand why, let us compare the co-citation map with a classic bottom-up representation, such as a list of sub-disciplines in a research field. It is very interesting that, often, the latter can be mapped onto the second, that is we can label the clusters of the map with the categories of the list. This process is sometimes called the "validation" of the map and the fact that it often succeeds warrants that the structure revealed by the map is not alien to the field experts. However, I think it would be a big mistake to confuse the labels with the regions of the map, the categories with the cluster. They are different both from the epistemological and ontological point of view.

From an epistemological point of view, let us examine the role that our knowledge plays in the two cases (the bottom-up and the top-down representations). In the top-down, we start by knowing the categories, the list of sub-disciplines composing the field. In the bottom-up, by contrast, we do not know the structure before the mapping (even if we assume that a structure exists). We start by knowing only the properties of the documents (in our case, the reference lists). The clusters are unknown to us before the mapping takes place. Therefore, the categories of the top-down list are an aid to interpret the clusters, they are not the clusters themselves! They are labels we use to give a name to the latent structuring forces that are behind the clusters.

My claim is the following: bottom-up representations allow us to capture the latent structuring forces that organize, i.e., provide a structure, to our units. Top-down representations, by contrast, are external to the units. They are imposed by an external observer to the set. It may happen that the top-down representation grasps the same latent forces that are present within the set: however, the way in which the top-down representations are known remains very different from the procedure that generates the bottom-up representations. Top-down are a priori, bottom-up are a-posteriori. Their epistemological difference reflects a deeper difference of ontological kind. Top-down representations capture explicit mental processes of the actors of the filed, bottom-up representations capture the latent forces shaping the (intellectual) actions of the actors.

Bottom-up representations results from an explicit mental process, in which the actors belonging to a field start to think reflexively about the structure of their own field. As we said above, this auto-reflexivity can be stimulated by different circumstances, from highly theoretical issues to managerial needs. Even if top-down representations can be developed collectively by several actors (as it is the case of disciplinary encyclopedias), they will never be able to include the activity of all actors. They will remain, thus, always partial representations (even if they can be very good partial representations!). Bottom-up representations, by contrast, start by selecting a type of actions that the actors perform without the explicit aim of representing the field (e.g., the action of citing references) and then consider the actions of all the actors (at least, when data are available). The analysis reveals that these actions are not random, but they are organized in patterns. The patterns are the effect of latent forces shaping these actions. The structure revealed is thus the force field in which the actors are immersed.

4 Structuration forces

Here we face a key theoretical problem: how to understand the relationship between the actors and the structure resulting from their actions? This is a central problem in social theory (theoretical sociology). In the context of disciplines, this turns out as the problem of the relationship between the research field and the researchers that collectively produce it.

From a certain point of view, the research field is no more than the sum of the intellectual actions of the researchers. However, it would be wrong to deduce from this that the individual actor is completely free to shape the field as she wants. If it is true that the collectivity of researchers produces the structure of the field, we must recognize that, from the viewpoint of the individual researchers, the field has an independent existence which limits the individual's degrees of intellectual freedom. In order to produce the field, the individual must reproduce, at least to a certain extent, the structure of the field itself. The degree in which the individual can alter, by her action, the structure depends on different variables (for instance, the dimension of the fields in terms of publications and researchers). In general, we should say that these variables determine together the inertia of the field. The actors perceive the structure of their discipline as a resistance to change.

In terms of the bottom-up representations, the inertia should be attributed both to the clusters of documents and the overall configuration of the map. The latent forces that structure the field, and that are visualized as clusters in the map, exert a sort of gravitational force towards the actors. In other terms, escaping the clusters, like leaving the attractional force of the Earth, has a cost. By the same token, creating a new cluster requires "energy" since a new attracting pole must be established. Lastly, changing the overall configuration of the map is an action probably beyond the possibility of any individual actor. The overall inertia is too large.

Until now, we focused on the "negative" side of the structure, framing its impact on the actors in terms of constrain. However, it would be wrong to attribute only this function to structures. A cluster is at the same time a constraint and the context which enables action. Imagine a field without a structure: no standard literature would be available, it would not be clear where a research should start, no common results would be shared in the research community. This is the situation of the Kuhnian pre-paradigmatic science: the knowledge that is produced in this environment is unstable, does not accumulate, and is constantly disputed. The presence of a structure, i.e., of latent forces structuring the field, by contrast, allows the actors to find a state of the art, or, in Kuhnian terms, a paradigm or, in Lakatosian terms, a research programmes. The key positive function of structures is twofold: on the one hand, to reduce uncertainty, on the other hand, to focus the attention of the researchers. In general, the structure allows to manage and channel information by reducing the noise. Specifically, the intermediate sub-structures, i.e. the sub-disciplines which show up in the map as clusters play a central role in this respect. By channeling the information into relatively homogeneous sets (the sub-disciplinary literatures), they allow the individual actors to cognitively dominate the information and thus to contribute to the production of new knowledge.