



Citation for published version:

Napolitano, L, Evangelou, E, Pugliese, E, Zeppini, P & Room, G 2018, 'Technology networks: the autocatalytic origins of innovation' Royal Society Open Science, vol. 5, no. 6, RSOS-172445 .
<https://doi.org/10.1098/rsos.172445>

DOI:

[10.1098/rsos.172445](https://doi.org/10.1098/rsos.172445)

Publication date:

2018

Document Version

Peer reviewed version

[Link to publication](#)

University of Bath

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Technology networks: the autocatalytic origins of innovation

Lorenzo Napolitano^{1,3}, Evangelos Evangelou¹, Emanuele Pugliese^{1,2,3},
Paolo Zeppini^{1,4} and Graham Room¹.

April 19, 2018

¹ *University of Bath, BA2 7AY Bath, UK*

² *International Finance Corporation, World Bank Group, 20433 Washington, USA¹*

³ *Istituto dei Sistemi Complessi (ISC)-CNR, 00185 Rome, Italy*

⁴ *Université Côte d'Azur, CNRS, GREDEG, 06560 Valbonne, France*

Abstract

We analyse the autocatalytic structure of technological networks and evaluate its significance for the dynamics of innovation patenting. To this aim, we define a directed network of technological fields based on the International Patents Classification, in which a source node is connected to a receiver node via a link if patenting activity in the source field anticipates patents in the receiver field in the same region more frequently than we would expect at random. We show that the evolution of the technology network is compatible with the presence of a growing autocatalytic structure, *i.e.* a portion of the network in which technological fields mutually benefit from being connected to one another. We further show that technological fields in the core of the autocatalytic set display greater fitness, *i.e.* they tend to appear in a greater number of patents, thus suggesting the presence of positive spillovers as well as positive reinforcement. Finally, we observe that *core shifts* take place whereby different groups of technology fields alternate within the autocatalytic structure; this points to the importance of recombinant innovation taking place between close as well as distant fields of the hierarchical classification of technological fields.

1 Introduction

A large body of research on complex systems – physical, biological and socio-economic as well – has focused on the relation between the structure of interactions within heterogeneous populations of agents and the dynamic properties of the aggregate system they populate. This has implied a change of perspective from *linear* narratives, in which the direction of causality connecting phenomena is unambiguous, towards processes of cumulative causation [1, 2]. This relation is particularly evident in the study of biological systems, where models of pure resource competition are unable to explain the persistent variety of ecosystems. To this aim, it is useful

¹The findings, interpretations and conclusions expressed herein are those of the authors and do not necessarily reflect the view of the World Bank Group, its Board of Directors or the governments they represent.

to introduce the idea that some species in a heterogeneous population can serve as catalysts (or inhibitors) for the survival of other species. The interaction between preys and predators is one of the best-known examples of this kind of relation, but a similar mechanism has also been observed in settings like plant-pollinator interactions, opportunistic behaviour, and symbiotic relations [3]. A particularly relevant feature of these ecological systems is the presence of *autocatalytic sets* [4] – self-sustaining subsystems, in which each species benefits directly or indirectly from its cohabitation with the others. The relevance of interactions in the above framework lends itself to a complex systems interpretation for which networks are a natural tool of analysis.

The idea of catalytic interaction can be fruitfully extended to human systems and, in particular, to the realm of technological innovation. In this setting, interactions take place between technological fields whenever existing ideas are applied to new problems or used to bridge previously unrelated fields thereby expanding the set of technological capabilities and spawning further innovation [5, 6, 7, 8, 9, 10]. Consider, for instance, the development of lasers, which have opened the way to a number of innovations in different industries, ranging from telecommunication to data storage and health care. It is also possible that the combination of multiple technological fields will give rise to a radical – possibly disruptive – novelty or to technological ‘convergence’ [11, 12, 13]. An example of the former is the case of opto-electronics, since optical and electrical devices lie at the root of the technological framework for modern telecommunication systems; a notable example of convergence is instead provided by smart-phones and electronic tablets, which combine functionalities that could previously be found only separately in computers, telephones, and television sets. In general, combining knowledge from previously isolated domains has become extremely relevant in several innovation-oriented domains, such as academic research projects – which often involve scientific collaborations between groups with heterogeneous backgrounds [14] – and industrial endeavours – where R&D collaborations have become common practice [15, 16, 17] especially in sectors characterized by a quick pace of technological progress (*e.g.* biotechnology [18] and IT [19]). For this reason, the network structure of both scientific [20, 21, 22] and industrial collaborations [23] has been studied in depth in the past.

This study investigates the autocatalytic structure of the network of interactions between technological fields extracted from patent data and shows how the tools of complex systems analysis are able to forecast the evolution and future relevance of individual fields based on their role in the interaction network. Following [24], we say that a technological field is catalytic for another one if the development in a region of innovations involving the former is positively associated with the future emergence of innovative expertise in the latter in the same region. We employ patents as a proxy for inventions, in line with an established body of research about the patterns of technological change, which has been pioneered by the scholarly [25, 26] and institutional effort to tap into their potential to shed light on relevant open questions concerning the drivers of technological progress, the relative importance of technological domains, and the significance of technological proximity *vis à vis* technological variety for the emergence of radical and incremental innovation [27]. One of the decisive advantages of using patents as sources of data stems from the fact that exclusive commercial rights to an invention are granted to applicants provided that they publish a complete description of the patented invention allowing it to be replicated by others once the exclusive rights expire. In order to assess innovativeness, patent offices map the claimed novel features of each invention to the technological fields it impacts through a standard classification system and collect the above information into dedicated databases. This leads both to an extensive coverage of the innovation spectrum and a high degree of standardization that allows large scale analysis.

The rest of paper is organized as follows. Section 2 defines the technological network based on patent classification codes; Section 3 summarizes the features of autocatalytic networks; Section 4 reports the results of our study of the autocatalytic structure of the technology network;

Section 5 concludes.

2 The patent network

2.1 Connecting regions and technological fields

Our analysis relies on the patent data contained in PATSTAT [28], a comprehensive database collecting information about applications filed at national and regional patent offices all around the world. PATSTAT contains several tables linking over fifty million patent applications to information such as the filing date of applications, the patent families² they belong to, and their technological content as encoded by the International Patent Classification (IPC) codes attributed to the patent claims by patent office examiners. IPC codes define a hierarchical classification consisting of five levels (*sections, classes, subclasses, groups, subgroups*), which includes 8 codes at the coarsest level (sections) and over 70 thousand codes at the bottom of the classification tree.

We associate patents to the location of their assignees through Orbis, a commercial database of firm level data maintained by the Bureau van Dijk, which collects the list of patent PATSTAT identifiers of the applications filed by companies that have been active in patenting at some point in time. We match the technological codes associated to patent families with the firm-level data so to unambiguously localize firms geographically through their country of residence and their postal code. This allows us to construct a geographical matrix attributing IPC codes to regions through the patent portfolios and the geographical locations of the companies³.

We observe different regions worldwide across time to uncover the effects that innovation in a field produces on other fields within the same region. In order to build comparable regions across different countries, we need a spatial identification system. For European Countries we connect postal codes of patenting companies to the associated NUTS 3⁴ regions of the standard European classification (corresponding *e.g.* to Provinces in Italy and Districts in the United Kingdom). Since extra-European countries are not included in the NUTS classification, we resort to national classifications when necessary and align them to achieve a broad overall accordance with the employed European classification and hence assure the consistency of the geographical tree. Our data define a spacial hierarchy comprising over 3000 regions in 39 countries and aggregating information about the location of around 500 thousand patenting firms.

As mentioned above, the basic units of observation for the construction of the data matrices ($\mathbf{W}(t)$) are individual firms and the patents they own. In particular, we group patents into families and consider the latter as individual inventions because of the strong contiguity between documents they group together. In building the matrices, we assume that every family containing patents filed in year t counts as one unit and weighs accordingly within $\mathbf{W}(t)$. Moreover, we make the hypothesis that the technologies expressed within a patent family can be reasonably accounted for by considering the set of unique IPC codes they contain, thus avoiding to double-count codes that appear in patent applications filed in the same year that belong to the same family. We evenly split the unit of weight attributed to each family that is active in a given year between all unique combinations of technology codes and regions it maps to, therefore defining

²A patent family is a collection of patent applications addressing the same invention, usually in different patent offices.

³Note in passing that companies (or groups of companies) based in multiple locations do not necessarily assign the patents they own to the address of their corporate headquarters, but sometimes attribute inventions to one of their regional offices or a subsidiary.

⁴Third level of the hierarchical *Nomenclature des Unités Territoriales Statistiques*, the standard classification of European sub-national regions established by the European statistical office. The acronym NUTS is also commonly associated to the English name of the classification, *Nomenclature of Territorial Units for Statistics*.

element $W_{r,i}(t)$ as the sum, for all families containing patents filed in year t , of the shares of each family attributed to field i and region r .

For the purpose of the analysis, we need to transform each $\mathbf{W}(t)$ into a presence-absence matrix. In line with the literature, we assign a value of 1 to a location-technology pair if the value within the corresponding cell is compatible with a measure of revealed advantage. This allows us to reduce noise and avoid overstating the relevance of technological fields in those regions in which they play only a marginal role. In particular, we use revealed comparative advantage [29] to produce a matrix $\mathbf{M}(t)$ in which $M_{r,i}(t)$ is recorded as a presence ($M_{r,i}(t) = 1$) if in the corresponding $\mathbf{W}(t)$ we have that

$$\frac{W_{r,i}}{\sum_i W_{r,i}} > \frac{\sum_r W_{r,i}}{\sum_{r,i} W_{r,i}}, \quad (1)$$

and an absence ($M_{r,i}(t) = 0$) otherwise.

Due to time lags between patent filing and data publication by patent offices, the version of PATSTAT we employ (2014a) contains reliable data up to 2011, after which coverage falls sharply. For this reason, 2011 is the most recent year we include in the analysis. As for the left extreme of the time interval, coverage is not an issue especially for the second half of the 20th century, even though the number of patents filed decreases quickly going backwards. We stop at 1980 because it strikes a balance between having a long time interval for the analysis and dealing with yearly $\mathbf{M}(t)$ matrices that are not excessively sparse.

2.2 Directed network between technological fields

The aim of this work is to measure the relationship between the patenting activity taking place within a geographical region (r) involving technological field i at time t and the patenting activity performed in r involving a possibly different field (j) at time $t + \delta$ ⁵. To this end, we count how often patents in field i are present at time t in regions that produce patents in field j at time $t + \delta$. We discount for regional diversification $d_r(t)$ – *i.e.* the number of fields in which region r is active at time t – and the ubiquity of different fields $u_i(t)$ – *i.e.* the number of regions in which each field is represented at time t – to establish a measure of the excess probability that innovation in a technological field precedes in another field in the same place. Applying the procedure proposed by [24] to the conceptual framework proposed by [30, 31, 32], we obtain:

$$B_{i,j}(t, \delta) = \frac{1}{u_i(t)} \sum_r \frac{M_{r,i}(t)M_{r,j}(t + \delta)}{d_r(t + \delta)}, \quad (2)$$

where $d_r(t) = \sum_j M_{r,j}(t)$ and $u_i(t) = \sum_r M_{r,i}(t)$. $B_{i,j}(t, \delta)$ can be interpreted as the probability that a region displaying a revealed competitive advantage at time t in the field i will reveal a competitive advantage at time $t + \delta$ in field j

$$B_{i,j}(t, \delta) = \text{Probability}(j, t + \delta | i, t) = \sum_r \text{Probability}(j, t + \delta | r) \text{Probability}(r | i, t), \quad (3)$$

assuming that the information about the capabilities linking pairs of technological fields is fully captured by their co-occurrence within each region, *i.e.* that $\text{Probability}(j, t + \delta | r, i, t) = \text{Probability}(j, t + \delta | r)$.

An equivalent way of interpreting equation 3 is illustrated in Figure 1, which depicts the tripartite directed network connecting (1) technological field i at time t to (2) regions, and regions to (3) technological field j at time $t + \delta$. In this framework, $B_{i,j}(t, \delta)$ is equivalent to the probability that a random walk on the network starting from technology i at time t will reach technology j by time $t + \delta$. It is worth noting that the choice of the scale of analysis is not neutral,

⁵The results presented in this paper refer to $\delta = 1$

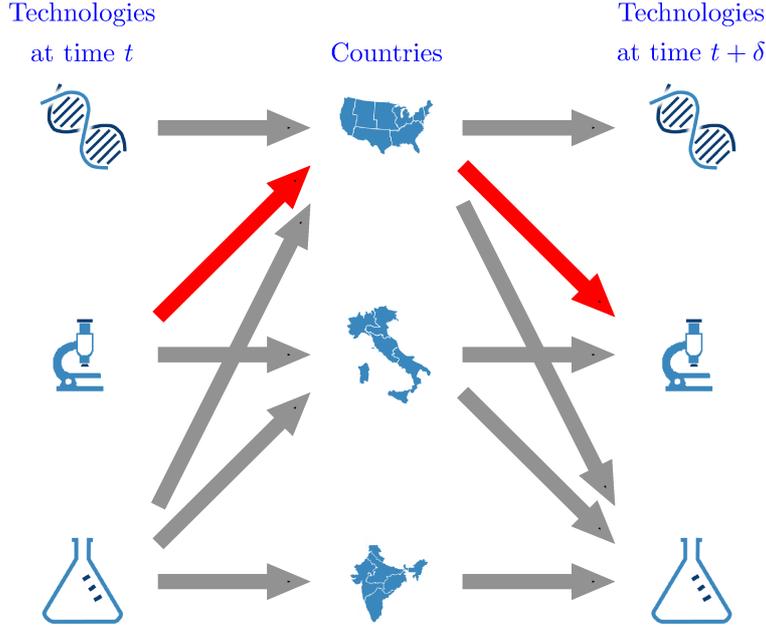


Figure 1: Model representation of the three-layer technology-country-technology.

since – like any socio-economic process – also innovation has different characteristics depending on the resolution at which it is observed. However, equation 3 can be naturally applied to observe the system at different technological granularities. In what follows, we employ IPC classes for the analysis presented in the main text and later check the robustness of the results on the more disaggregated IPC subclasses (see Supplementary Information). Classes and subclasses are two nested levels of the IPC hierarchy that divide the spectrum of technological fields associated to patents respectively into 121 and 640 fields.

2.3 Assessing link significance

Matrix $\mathbf{B}(t, \delta)$ of equation 2 represents a directed weighted network connecting all technologies⁶. However, to decide whether a link between two technologies is statistically significant, a *null model* is required to account for the fact that some links could seem relevant due to the properties of the graph without, however, being the product of any actual catalytic effect. For example, very advanced technologies can be developed only by a minority of regions around the World and, for this reason, the associated codes might often appear in the same region at different points in time though no catalytic relation connects them. Following [24], we use the bipartite configuration model [33] to test the empirical network against a randomly generated counterpart that displays, on average, the same degree distribution. Thus, in the random graphs we generate, each region has the same expected diversification in terms of technological codes as the empirical data and each technological code has the empirically observed expected ubiquity. The degree is thus the only information we extract from the empirical matrices to generate the null models. Generating a large number of null matrices with the same null model (in our case, 1000 null matrices for each pair of years), it is possible to establish the significance of each link between technologies and generate a matrix \mathbf{P} with element $P_{i,j}$ representing the percentile of the null distribution containing the link associated to $B_{i,j}$. We can then define the statistical significance of each individual link.

⁶ For ease of presentation, in the following we omit the time indexes and refer to $\mathbf{B}(t, \delta)$ and its elements as \mathbf{B} and $B_{i,j}$ respectively. The same simplification applies to all the matrices derived from \mathbf{B} .

The final step consists in constructing the unweighed directed adjacency matrix \mathbf{C} that contains only the significant links in the network. A link from field i to field j is included in \mathbf{C} if the corresponding $P_{i,j}$ is larger than a fixed threshold, say $1 - p$. However, if this comparison is performed separately for each pair of fields by looking at individual significance levels, by definition, we would expect a share p of false positive links to be retained. Effectively, the probability of false positive links for the whole network will be higher than the desired level p . This is known as the multiple comparisons problem in the statistics literature [34]. A method to accurately control for the proportion of false positive links simultaneously for all pairwise link tests, and thus maintaining the overall significance to p , is the false-discovery-rate procedure of [35]. The network used in our analysis was constructed using this procedure by retaining significant links in the constructed network from the matrix \mathbf{P} with overall significance level $p = 5\%$.

3 Autocatalytic networks

The matrix \mathbf{C} defined in the Section 2.3 is a directed adjacency matrix representing the links between the N nodes that represent technological fields: a directed link from a technological field i (source) to another technological field j (receiver) is present ($C_{i,j} = 1$) if there is a significant signal suggesting that patenting activity in i promoted patenting activity in j in the same region. Within this representation, directed technological links either exist or not, and the technology network is completely specified by a binary *adjacency matrix* of size $N \times N$ with elements $C_{i,j}$.

Inspired by [1], [36, 4] have proposed a model of network evolution for biological systems based on a catalytic interpretation of the relationship between species, which we apply to the case of interactions between technologies. To describe the fundamental analytical properties of catalytic systems, we use the simplest dynamical model

$$\dot{y}_i = \sum_j C_{i,j} y_j, \quad (4)$$

where y_i is the intensity of activity in technological class i . In this simple example, activity in field i pushes the innovation activities in every field j such that $C_{i,j} = 1$. Notice that the class of models proposed by [36] is vast and the following results are still valid for more realistic models of the dynamic behaviour of the innovation system.

Consider for instance an empty network ($C_{i,j} = 0 \forall i, j$). In this toy example, $\dot{y}_i = 0 \forall i$ and any initial condition is an equilibrium. Let us now add a link from i to j . In the model, while all the other nodes still experience a constant activity, we have $\dot{y}_j = y_i$, which implies $y_j(t) = y_j(0) + y_i(0)t$. When more links are present, we can have different polynomial behaviours for the activity in different technology classes. However, an interesting case arises when there is a direct or indirect reciprocal influence between i and j . In this case, i and j form an autocatalytic set, which, in terms of the technology network implies that the more patents employ technology i , the more will also employ technology j , and *vice versa*. The catalytic cycle creates an exponential dynamics involving the innovation activities acting on fields in the set. This exponential behaviour is in line with the empirical observation of the innovation system, both in terms of patents and in terms of productivity growth. More in general, we can see the same exponential dynamics for any arbitrarily long cycle, *i.e.* for any closed path connecting a subset of the nodes in the network.

To this end, it is useful to introduce the notion of an *autocatalytic set* (ACS), which is defined as a “subgraph, each of whose nodes has at least one incoming link from a node belonging to the same subgraph” [37]. In what follows, any set of nodes connected through one or more cycles will be called the *core* of its corresponding ACS, which, in addition to the nodes in the core, also includes the *periphery*, *i.e.* the set of nodes that are catalysed by the core but have

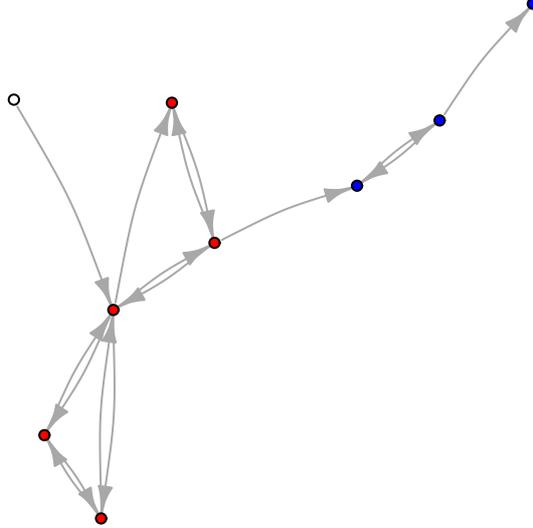


Figure 2: Example of a graph containing an *autocatalytic set* (ACS). The red nodes belong to the *core* of the ACS, the blue nodes belong to the *periphery* of the ACS, while the white node does not belong to the ACS, because it has no incoming link from the ACS itself.

no outgoing links feeding into the closed path. Because of this configuration, peripheral nodes have a passive role in the ACS and thus are not a part of the core; however, they still benefit from the boost provided by their incoming links. Figure 2 depicts a simple network containing an ACS and highlights the relevant subsets in which its nodes can be partitioned. By expressing equation 4 in its matrix form, $\dot{\mathbf{y}} = \mathbf{C}\mathbf{y}$, it is easy to show the relationship between the presence of cycles in the network and the presence of positive *eigenvalues* of the adjacency matrix \mathbf{C} . Since the adjacency matrix is *non-negative*, the Perron-Frobenius theorem guarantees the existence of a real eigenvalue which is the largest of all other eigenvalues. This one is called the Perron-Frobenius eigenvalue (PFe), λ_1 . It is possible to prove [38, 39] that a cycle exists in the graph if the PFe is greater than 0. A formal proof of the theorem is beyond the scope of this paper, but it can be shown that, if the PFe is greater than 0, any innovation activity corresponding to a positive element of the Perron-Frobenius Eigenvector (PFE), \mathbf{y}_1 , experiences an exponential growth because $\dot{\mathbf{y}}_1 = \lambda_1 \mathbf{y}_1$. Indeed, the PFe is informative of the presence or absence of closed directed paths (loops) in the graph, and its corresponding Eigenvector has elements different from zero corresponding to the nodes of the autocatalytic set. A larger PFe indicates a faster exponential growth driven by a higher connectivity in the core of the ACS. Notice that, more in general, a matrix can have more than one PFe (and corresponding PFE) if it has more than one ACS (two ACS are distinct if there is no path connecting the two cores).

In what follows, we show how this very simple model can both i) give us a novel understanding of the innovation system as a process of cumulative causation, and ii) identify the core technologies in the evolving technological landscape.

4 Autocatalytic structure of the technology network

4.1 Mapping the network

Figure 3 reports a selection of technological networks in which the nodes are the IPC classes (the interested reader is referred to Section A in the appendix for the analysis of the network of IPC subclasses). In these plots, nodes are linked whenever there is a statistically significant

relation between them. After generating the network of technology classes, we map the network searching for an autocatalytic structure as defined in Section 3. A unique ACS is present in the network between 1980 and 2010, the size of which increases over time. At the beginning of the sample, multiple clusters are present that do not form an autocatalytic set because no closed path between them is present. By 1998 most clusters have been ‘captured’ by the ACS, which in 2010 spans approximatively half of the technology network. In particular, the core occupies the largest portion of the ACS and peripheral nodes are only a minority.

Figure 4 reports the time evolution of the relevant statistical indicators of the autocatalytic network configuration, namely the maximum eigenvalue (upper-left panel), the size (number of classes) of the core (upper-right panel), the size of the periphery of the ACS (lower-left panel) and the size of the entire ACS (lower-right panel). The plots clearly show that the autocatalytic character of the network becomes stronger in time and that its growth seems to occur in two phases: a first phase in which the periphery expands (bottom left panel) and a second one where most of technologies in the ACS transition to the core (see bottom left and bottom right panels).

4.2 Fitness

Having observed the emergence of an autocatalytic structure in the technology network of classes, we want to understand how this structure affects the *fitness*⁷ of the technological fields. In this article, we define fitness of a technological field the number of patent applications filed in a certain year featuring said field. The idea is that a higher patenting rate is indicative of a higher technological productivity, and a proxy of the ‘innovativeness’ of a technology class. By comparing class fitness in different parts of the network – namely the core, the ACS, and the rest of the nodes – we test our hypothesis that technology classes inside the ACS display higher fitness than other classes.

Figure 5 shows the time series of the fitness of the classes belonging to the core (red), to the whole ACS (blue), and to the rest of the technology network. Figure 5a reports the average fitness of nodes in each of the above mentioned subsets. It clearly emerges that the nodes in the ACS, and especially those in the core, show an increasing average fitness, while the fitness of the nodes in the rest of the network fluctuates around a constant value. This value is twice smaller than the fitness of the nodes in the ACS at the beginning of the sample and is ten times smaller by 2011. The average fitness in the ACS and in its core is largely of the same order of magnitude, although for nodes of the core it is almost always larger (it is lower only in 1990 and in 1996). These figures lend support to the hypothesis that technology classes belonging to the ACS benefit from an autocatalytic advantage. Figure 5b proxies an absolute measure of fitness by reporting the time series of the total number of patents in the ACS, its core, and the rest of the network. The plot shows that there were many more patents outside the ACS until the early 1990’s, when only a few nodes were part of the ACS (Figure 3). However, as soon as the size of the ACS started increasing substantially, this measure of absolute fitness quickly surpassed the total fitness of the rest of the network by a large measure. Although less meaningful than the average fitness time series, this figure is evidence of the transition experienced by the network of technological fields, with the emergence of a large autocatalytic structure around the end of the century. Our analysis suggests that technology classes belonging to the ACS are more innovative due to the positive effect of catalytic links on knowledge flows. In particular, classes in the core of the ACS experience the positive feedback of self-reinforcing cyclical catalytic structures, which give rise to cumulative processes of innovation. Figure 5c reinforces the above intuition by showing the evolution of the distribution of total fitness between the core of the ACS, the periphery, and the rest of the nodes. It shows that in the early 1980’s the autocatalytic

⁷ Note that we employ fitness in its original biological interpretation, *i.e.* the ability of a population to reproduce. There is no relationship between fitness as we use it in this paper and the use of the term in the related literature of Economic Complexity [40].

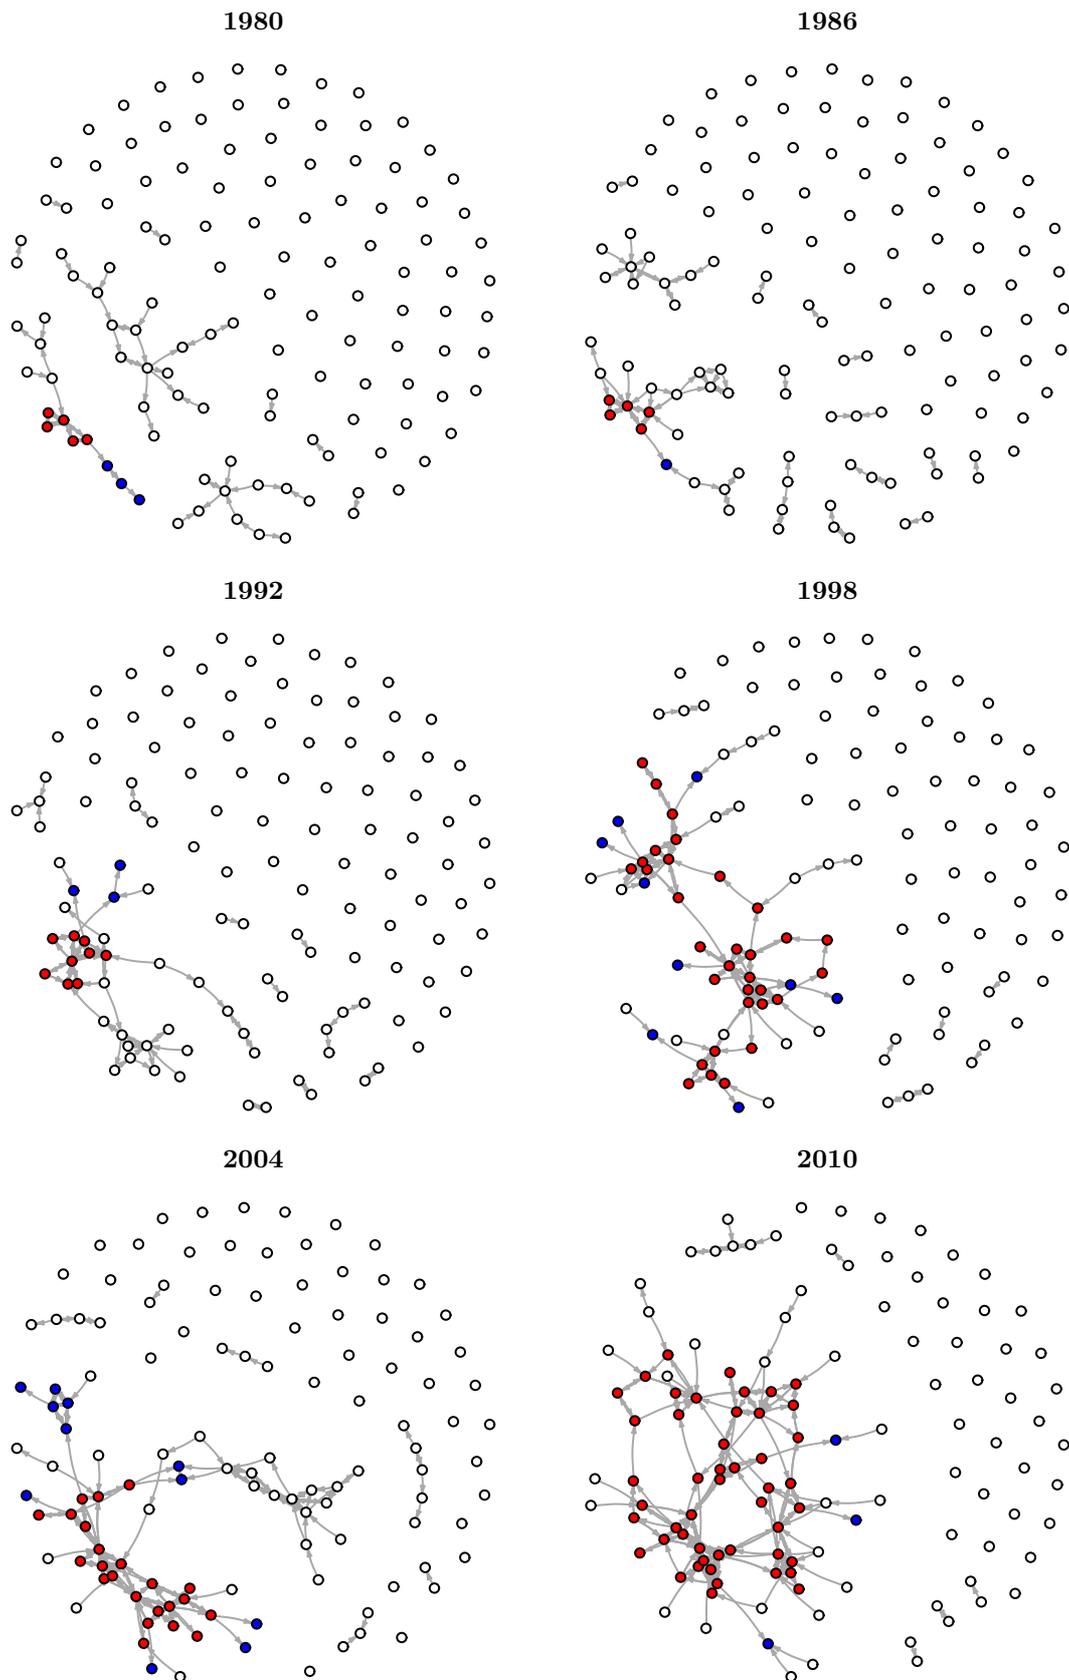


Figure 3: The technology network in six different years. The nodes are technology classes (IPC classification). Links account for statistically significant relatedness. Red nodes represent the *core* of the Autocatalytic Set (ACS) of the network, blue nodes are the *periphery* of the ACS, white nodes are outside the ACS.

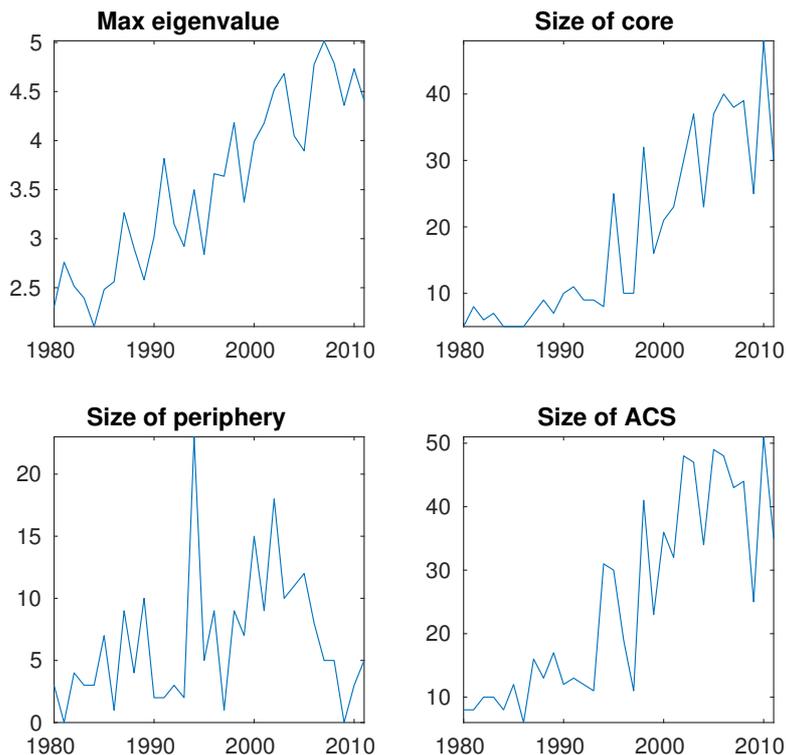


Figure 4: Statistical indicators of the autocatalytic structure in the network of technology classes from 1980 until 2011. Upper-left panel: largest eigenvalue (Perron-Frobenius) of the network adjacency matrix. Upper-right panel: size of the core. Lower-left panel: size of the periphery. Lower-right panel: size of the ACS.

structure was rather marginal in the network of technology classes, since it included only around 10% of the total number of patents. After 1985, an increasing trend started leading the ACS to gain weight in the network and become clearly predominant to the point that it concentrated almost 80% of total fitness in the last decade of the sample in the face of a more contained increase in the share of nodes it includes, as shown in Figure 5d. Considering that just above half of the classes in the network become part of the ACS in the same period, there is evidence of a strong correlation between the prominence of the ACS of the technology network and the intensity of patenting activity in the fields comprising it.

The above is a strong indication of the cumulative causation process behind innovation as identified by patenting activity. If technology classes in the ACS grow while classes outside the ACS do not, then catalytic inter-linkages between classes are a relevant driver of innovation. Taking patenting intensity as a meaningful proxy for innovation, the empirical evidence presented in Figure 5 suggests that the substantial growth of a technology class is linked to its connection to other classes ‘feeding’ into it with incoming links. This is to say that a flow of knowledge stemming from the source of the directed link can provide new knowledge and thus the basis for a new patent in the target class. The fact that stronger growth for a technology class comes from belonging to the core shows that innovation is fostered not only by technology spillovers from other classes, but especially from membership of a cycle of mutually reinforcing spillovers.

4.3 Autocatalytic structure and database hierarchy

In this section, we turn to the hierarchical structure of the database to investigate its role in the autocatalytic structure of the technology network. As mentioned in Section 2.1, PATSTAT

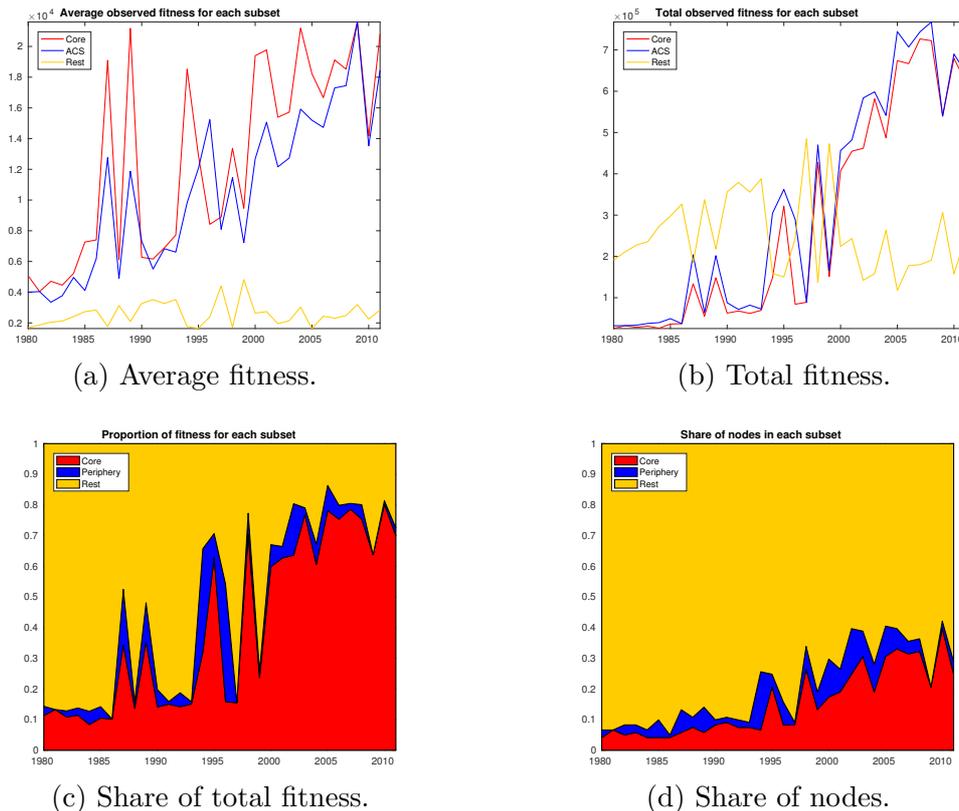


Figure 5: Total and average fitness, *i.e.* total and average number of patents per IPC class in different portions of the technology network (ACS, core, and the rest of the network).

adopts the International Patents Classification (IPC)⁸, which has a tree-like structure consisting of eight *sections* at the root, which branch out into the progressively finer-grained *classes*, *subclasses*, and so on. For the present analysis, we use classes and subclasses⁹ as nodes of the technology network, but we are also interested in understanding how the hierarchy induced by the IPC maps to the ACS. To this end, we employ the IPC sections

- HUMAN NECESSITIES (A)
- PERFORMING OPERATIONS; TRANSPORTING (B)
- CHEMISTRY; METALLURGY (C)
- TEXTILES; PAPER (D)
- FIXED CONSTRUCTIONS (E)
- MECHANICAL ENGINEERING; LIGHTING; HEATING; WEAPONS; BLASTING (F)
- PHYSICS (G)
- ELECTRICITY (H)

to cluster the nodes and see how sections map onto the structure of the technology network. In particular, we ask if the ACS discriminates sections and if significant links cut across section borders. These are not just questions of topological nature, because the answers shed light on whether recombinant innovation in general, and more specifically the cumulative causation process of autocatalytic structures take place mainly within classes or if, instead, it also involves broader connections between the coarse technological areas identified by IPC sections.

⁸ Note in passing that the IPC is not the only existing classification of technological fields. However, the IPC is by far the most common and one of the few that allows to perform a broad international comparison.

⁹The results for IPC subclasses are similar to those for classes; the interested reader is referred to appendix A for further details.

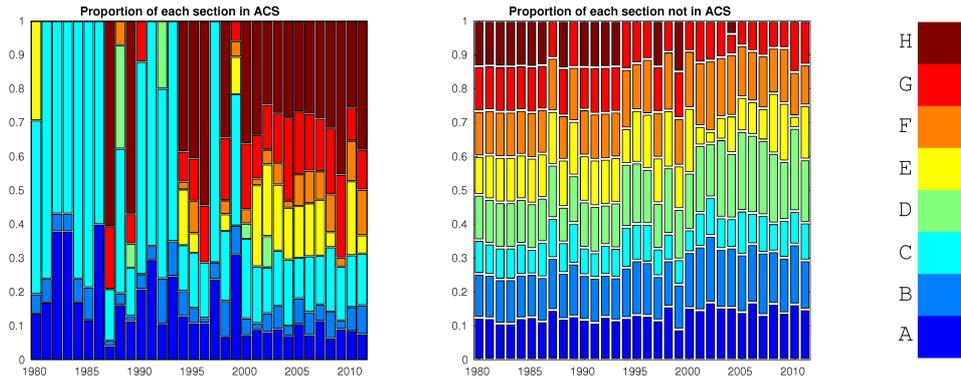


Figure 6: Relative share of nodes in the ACS (left panel) and outside the ACS (right panel) for the eight technology sections of the database.

Figure 6 reports of the share of each IPC section within and outside the ACS thus showing how relevant each section is for the two subsets of nodes. The main result, displayed in the left panel, is that the ACS is not at all static. At the beginning of the sample, the ACS consists almost entirely of the classes contained in section C (Chemistry, metallurgy), the only exception being one node of section A (Human necessity). However, over time a more variegated picture unfolds and, by 2011, the ACS spans all sections but one – D (Textiles; paper) – which remains outside the ACS consistently and only makes a few sporadic appearances. A further observation suggested by Figure 6 concerns the different composition of the ACS and the rest of the network; while the latter presents a quite uniform distribution of sections along the whole period 1980-2011, the ACS is characterised by less uniformity and a richer dynamics.

Concerning the above non-uniformity, it is also possible to quantitatively assess the non-uniform distribution of sections in the growing autocatalytic structure of the technology network. To this end, we can imagine having an urn filled with marbles of different colours, each corresponding to a distinct section, in which every marble represents an IPC class. There is a different number of marbles of each colour in the urn and we want to sample as many marbles as there are classes in the ACS. The null hypothesis is that the sampling is random, meaning that marbles are sampled blindly from the urn. The alternative hypothesis is that the sampling process is preferential and tends to privilege a specific subset of colours. We test the null hypothesis of random sampling against the alternative hypothesis of biased sampling from Fisher’s non-central hypergeometric distribution [41] and find that in every year from 1980 until 2011, the statistic is above the critical value for a significance of 5% (see appendix B for further details). This indicates a significant bias in favour of some sections in terms of occupancy of the ACS over the period considered and shows that, though the expansion of the ACS brings more sections in the ACS and its core, a non-uniform distribution remains.

The pattern observed in Figure 7 is evident also if we look at the share of nodes of each section belonging to the ACS. For instance, Figure 7 shows that one section, H (Electricity), is entirely contained in the core of the ACS in the more recent years of the sample, while the situation of class E (Fixed construction) is less clear, also due to large fluctuations in the final decade. Sections F and G seem to display a somewhat growing trend in the share of classes they contribute to the ACS. The remaining sections, have a relatively stable share of classes in the ACS. The data relative to the last decade highlights the strong transversal character of the autocatalytic structure of the technology network.

Up to now we have seen that the expansion of the ACS has brought more sections in the ACS and its core, though a non-uniform distribution has remained, with some sections that are almost absent from the ACS throughout the whole period. Moreover, there is a shift from A (Human necessities) and C (Chemistry and Metallurgy) towards H (Electricity) as the most important sections of the ACS, suggesting that there has been a transition from Chemistry and

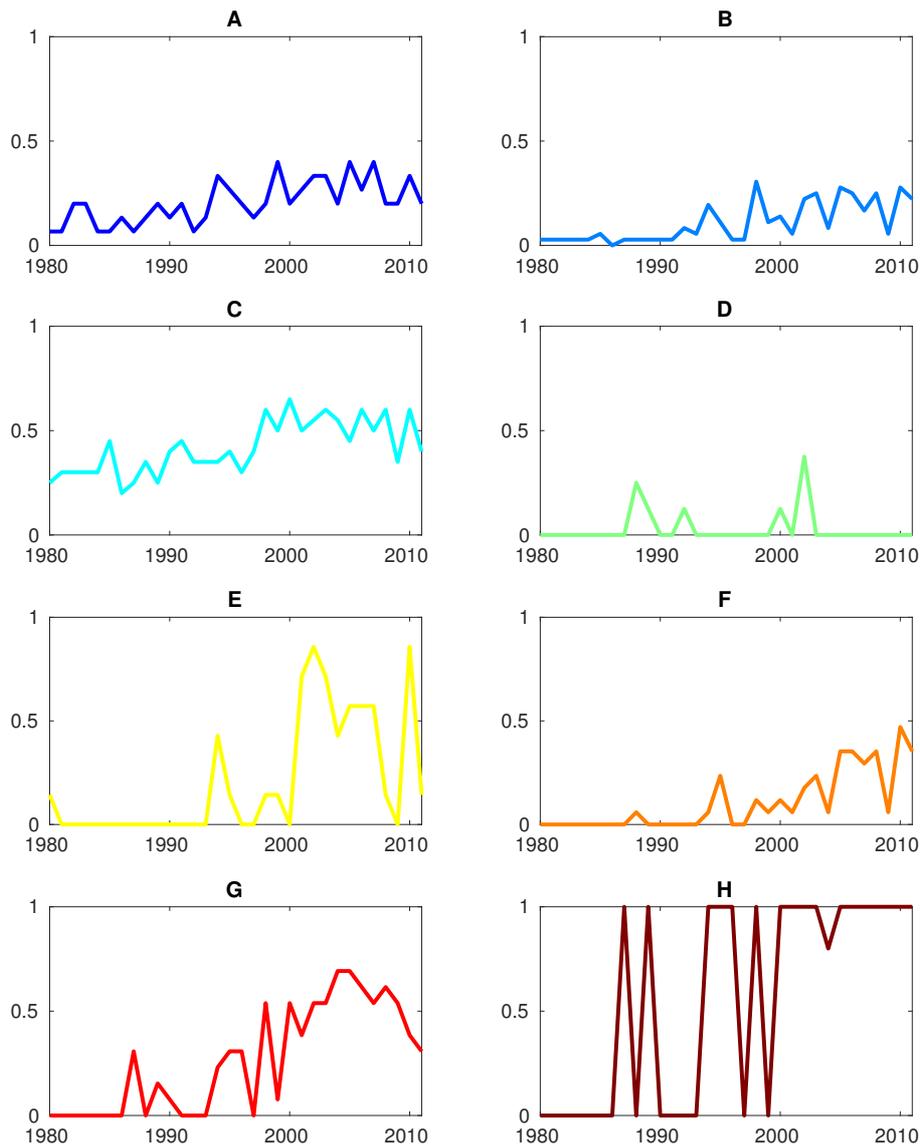
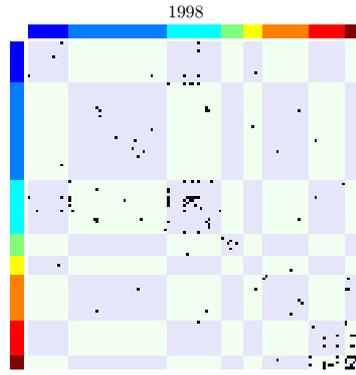


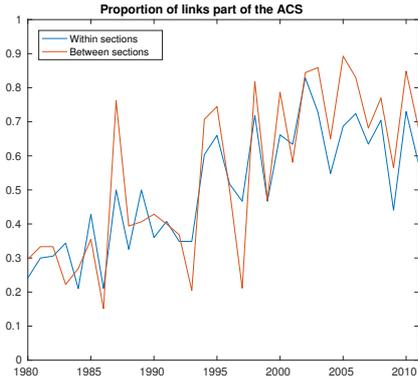
Figure 7: Time series of the proportion of nodes in the ACS for each of the eight technology sections. The plotted values are the ratio, for each section, of the number of classes belonging to the ACS to the total size of the section.

Metallurgy in the 1980's to Electricity in the 21st century as driving forces of innovation. The relevance of individual sections can be further clarified by looking at the adjacency matrix of the technology network displayed in Figure 8a and the relative distribution of links within and between sections plotted in Figure 8b. Notice that the matrix of Figure 8a is sorted according to a lexicographic ordering of the IPC class codes, so that nodes are grouped together based on the section they belong to. The presence of a block-diagonal structure would indicate that most links occur between classes of the same section. However, evidence of such a pattern appears to be mixed both in Figure 8a, which depicts the adjacency matrix of the technology network for 1998 (the central year of the sample¹⁰) and Figure 8b, which summarizes the evidence for the entire period. Note that sections in the technology network are not all equal in terms of *within* and *between* connectivity. For example, sections B and D seem to have mostly self-contained significant links, suggesting that the hierarchical structure of the IPC captures the extent of

¹⁰ The interested reader is referred to Section C information for a depiction of the adjacency matrix in a broader selection of years.



(a) Technology network for 1998.



(b) Within- and between-section links.

Figure 8: In the left panel, the 121 IPC technology classes are ordered based on sections, which go from A (Human necessities – dark blue – in the top-left corner), to H (Electricity – brown – in the bottom right corner). Black elements represent statistically significant links between classes. The right panel plots the time series of within section and between section links over time.

knowledge spillovers for technologies related to transporting, textiles, and paper. On the other hand, sections C and A, and likewise sections G and H, appear to share many links that cut the section border. This suggests that relevant spillovers can take place between ‘distant’ (and versatile) technologies and that a relevant role is played by the subsets of the autocatalytic structures in affecting the distribution of links. In other words, innovation, as measured by the production of patents, can potentially spawn as many connections between sections as it does within sections. The autocatalytic structure thus has at least comparable importance to that of sections in defining the boundaries of the drivers of technological progress.

5 Conclusions

This study is a first step to uncover the cumulative causation processes driving technological change by detecting an autocatalytic structure of patent databases. Our results can be summarized in three main points. First, the technology landscape described by the network of patent technology codes is characterized by a clear autocatalytic structure that has grown over the years to encompass most technology classes. Second, the classes that are involved in the autocatalytic set perform better in terms of innovativeness as measured by the rate of growth in the number of patents containing them. Finally, the autocatalytic structure of the technology network is more evident and stronger than the hierarchical structure of the database, in that as many links connect classes from different sections as they do classes from the same section. This implies that recombinant innovation arising from interdisciplinary technological interactions is a stylized fact of technological change.

We believe that our approach based on detecting autocatalytic structures can be successfully extended in more fundamental ways. We owe at least part of the inspiration to study autocatalysis in technological systems to the work of [1, 36, 4], who proposed autocatalytic networks as a model of self-organization for interlinked biological species. However, a distinctive characteristic of those models is the interplay of a *fast dynamics* driving species evolution and a *slow dynamics* ‘reshuffling’ the links of the network through species replacement or mutation. In this work, we instead assume that the equivalent of the fast dynamics acts on the population of technological codes and observe the changing network without providing a model for its evolution or its relation to the slow dynamics. A reason for this is that, while it is reasonable to assume that

population and network evolution take place at very different time scales in biological networks, this seems less plausible in the domain of innovation, where intuition suggests that success of individual technological fields and their mutual interactions might change at more similar speeds.

In principle technological codes are a powerful device to explore not only the different time scales of the dynamics shaping technological progress, but also the different scales of technological definition and categorization. For example, recent studies have observed that about 60% of new patents use novel combinations of codes [42] taken from the most recent version of the technological classification. Moreover, an interesting property of codes is the fact that the classifications they are drawn from are not static, but rather change over time to keep up with the pace of technological change. In fact, recombination of existing knowledge appears as a distinguishing feature of innovation, a stylized fact that can be directly observed through changes of the classification system [43, 44]. We envisage two main avenues for future research stemming from our study: first, an empirical analysis of autocatalytic sets at different scales of technological classification, to uncover possible fractal structures; second, a modelling framework that can reproduce and explain the statistical features of the empirical network of technology classes. From a methodological viewpoint, it would be interesting to explore the same questions using fundamentally different data about patents, such as citation networks [45, 46] or co-occurrences of technology codes within patent documents or families.

This article is the initial phase of a broader research project aimed at empirically assessing and understanding cumulative causation in social systems [47]. Within the realm of technological change, the results of this line of research bears potentially relevant implications for technological investment strategies at the corporate and institutional level as well as for innovation policy. More specifically, understanding the role of individual technological fields in the evolution of the wider technology system can be of the utmost importance for designing policies that address the different challenges of present times, from economic development and inequality to energy, security, and climate change.

References

- [1] Kauffman SA. Autocatalytic sets of proteins. *Journal of theoretical biology*. 1986;119(1):1–24.
- [2] Arthur B. Competing technologies, increasing returns, and lock-in by historical events. *Econom J*. 1989;99:116–131.
- [3] Domínguez-García V, Muñoz MA. Ranking species in mutualistic networks. *Scientific reports*. 2015;5.
- [4] Jain S, Krishna S. Large extinctions in an evolutionary model: the role of innovation and keystone species. *Proceedings of the National Academy of Sciences*. 2002;99(4):2055–2060.
- [5] Schumpeter JA. *The Theory of Economic Development*. Cambridge, MA: Harvard University Press; 1934.
- [6] Nelson RR, Winter SG. *An Evolutionary Theory of Economic Change*. Cambridge, MA: Harvard University Press; 1982.
- [7] Frenken K, Izquierdo LR, Zeppini P. Branching innovation, recombinant innovation, and endogenous technological transitions. *Environmental Innovation and Societal Transitions*. 2012;4:25–35.
- [8] Strumsky D, Lobo J, Van der Leeuw S. Using patent technology codes to study technological change. *Economics of Innovation and New technology*. 2012;21(3):267–286.
- [9] Loreto V, Servedio VD, Strogatz SH, Tria F. Dynamics on expanding spaces: modeling the emergence of novelties. In: *Creativity and universality in language*. Springer; 2016. p. 59–83.
- [10] Iacopini I, Milojević S, Latora V. Network dynamics of innovation processes. *Physical review letters*. 2018;120(4):048301.
- [11] Patel P, Pavitt K. The technological competencies of the world’s largest firms: complex and path-dependent, but not much variety. *Research Policy*. 1997;26(2):141–156.
- [12] Gambardella A, Torrisi S. Does technological convergence implies convergence in markets? Evidence from the electronics industry. *Research Policy*. 1998;27(5):445–463.
- [13] Korzinov V, Savin I. General Purpose Technologies as an emergent property. *Technological Forecasting and Social Change*. 2017;.
- [14] Powell WW, White DR, Kopur KW, Owen-Smith J. Network dynamics and field evolution: the growth of inter-organizational collaboration in the life sciences. *American Journal of Sociology*. 2005;110(4):1132–1205.
- [15] Ahuja G. Collaboration networks, structural holes, and innovation: a longitudinal study. *Administrative Science Quarterly*. 2000;45(3):425–455.
- [16] Hagedoorn J. Inter-firm R&D partnerships: an overview of major trends and patterns since 1960. *Research Policy*. 2002;31:477–492.
- [17] König MD, Battiston S, Napoletano M, Schweitzer F. The efficiency and stability of R&D networks. *Games and Economic Behaviour*. 2012;75(2):694–713.

- [18] Roijackers N, Hagedoorn J. Inter-firm R&D partnering in pharmaceutical biotechnology since 1975: Trends, patterns, and networks. *Research Policy*. 2006;35(3):431–446.
- [19] Hanaki N, Nakajima R, Ogura Y. The dynamics of R&D collaboration in the IT industry. *Research Policy*. 2010;39(3):386–399.
- [20] Newman MEJ. The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences*. 2001;98(2):404–409.
- [21] Barabási AL, Jeong H, Néda Z, Ravasz E, Schubert A, Vicsek T. Evolution of the social network of scientific collaborations. *Physica A: Statistical mechanics and its applications*. 2002;311(3):590–614.
- [22] Goyal S, Leij MJVD, Moraga-González JL. Economics: An emerging small world. *Journal of political economy*. 2006;114(2):403–412.
- [23] Letterie W, Hagedoorn J, van Kranenburg H, Palm F. Information gathering through alliances. *Journal of Economic Behaviour & Organization*. 2008;66(2):176–194.
- [24] Pugliese E, Cimini G, Patelli A, Zaccaria A, Pietronero L, Gabrielli A. Unfolding the innovation system for the development of countries: co-evolution of Science, Technology and Production. *arXiv preprint arXiv:1707.05146*; 2017.
- [25] Griliches Z. *Patents and Technology*. Cambridge, MA: Harvard University Press; 1998.
- [26] Hall BH, Jaffe AB, Trajtenberg M. *The NBER patent citation data file: Lessons, insights and methodological tools*. National Bureau of Economic Research; 2001.
- [27] Castaldi C, Frenken K, Los B. Related variety, unrelated variety and technological breakthroughs: an analysis of US state-level patenting. *Regional studies*. 2015;49(5):767–781.
- [28] PATSTAT. *EPO Worldwide Patent Statistical Database - 2014 Spring Edition*. European Patents Office; 2014.
- [29] Balassa B. Trade liberalisation and “revealed” comparative advantage. *The Manchester School*. 1965;33(2):99–123.
- [30] Teece DJ, Rumelt R, Dosi G, Winter S. Understanding corporate coherence: Theory and evidence. *Journal of Economic Behaviour & Organization*. 1994;23(1):1–30.
- [31] Zaccaria A, Cristelli M, Tacchella A, Pietronero L. How the taxonomy of products drives the economic development of countries. *PloS one*. 2014;9(12):e113770.
- [32] Yan B, Luo J. Measuring technological distance for patent mapping. *Journal of the Association for Information Science and Technology*. 2017;68(2):423–437.
- [33] Saracco F, Di Clemente R, Gabrielli A, Squartini T. Randomizing bipartite networks: the case of the World Trade Web. *Scientific reports*. 2015;5.
- [34] Colquhoun D. An investigation of the false discovery rate and the misinterpretation of p-values. *Royal Society open science*. 2014;1(3):140216.
- [35] Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the royal statistical society Series B (Methodological)*. 1995;p. 289–300.

- [36] Jain S, Krishna S. A model for the emergence of cooperation, interdependence and structure in evolving networks. *Proceedings of the National Academy of Sciences*. 2001;98(2):543–547.
- [37] Jain S, Krishna S. Graph theory and the evolution of autocatalytic networks. arXiv preprint nlin/0210070. 2002;.
- [38] Rothblum UG. Algebraic eigenspaces of nonnegative matrices. *Linear Algebra and its Applications*. 1975;12(3):281–292.
- [39] Jain S, Krishna S. 16 Graph theory and the evolution of autocatalytic networks. *Handbook of Graphs and Networks: From the Genome to the Internet*. 2006;.
- [40] Tacchella A, Cristelli M, Caldarelli G, Gabrielli A, Pietronero L. A new metrics for countries' fitness and products' complexity. *Scientific reports*. 2012;2.
- [41] Fog A. Sampling Methods for Wallenius' and Fisher's Noncentral Hypergeometric Distributions. *Communications in Statistics - Simulation and Computation*. 2008;37(2):241–257.
- [42] Youn H, Strumsky D, Bettencourt L, Lobo J. Invention as a combinatorial process: evidence from US patents. *J of the Royal Society Interface*. 2015;12(106).
- [43] Gkotsis P, Vezzani A. Technological diffusion as a recombinant process. Directorate Growth & Innovation and JRC-Seville, Joint Research Centre; 2016.
- [44] Lafond F, Kim D. Long-run dynamics of the US patent classification system. Available at SSRN: <https://ssrn.com/abstract=2924387>; 2017.
- [45] Hall BH, Jaffe A, Trajtenberg M. Market value and patent citations. *RAND Journal of economics*. 2005;p. 16–38.
- [46] Érdi P, Makovi K, Somogyvári Z, Strandburg K, Tobochnik J, Volf P, et al. Prediction of emerging technologies based on analysis of the US patent citation network. *Scientometrics*. 2013;95(1):225–242.
- [47] Room G, Spence A, Evangelou E, Zeppini P, Pugliese E, Napolitano L. Dynamics of Cumulative Innovation in Complex Social Systems (DCICSS). UK: University of Bath. Available at; 2017.

A Results for IPC subclasses

In this section, we examine the network of technologies defined by IPC subclasses. Figure 9 displays a selection of snapshots of the ACS at the subclass level, while Figures 10a and 10b respectively show the time series of the average fitness and the total fitness of subclasses in the ACS, its core and the rest of the network. As with IPC classes (see Figure 5), we find a clear positive trend of fitness for the ACS and its core over time. In particular, the average fitness of subclasses in these two sets grows from about 1000 patents per subclass in the ACS to about 2000 patents per subclass, and the average fitness for the core increases even more, to about 3000 patents per subclass (Figure 10a). At the same time, the average fitness in the rest of the network only barely increases reaching around 500 patents per subclass. Also the figures for the total fitness are indicative (Figure 10b) and show a lower level of fitness for ACS and core compared to the rest of the network in the first two decades of the investigated time period. The rest of the network is later surpassed by the fast growing ACS around the turn of the century. If we discount for the fact that subclasses are still less numerous in the ACS (and we normalise by the total number of subclasses in a set) we obtain the picture presented by the average fitness in Figure 10a.

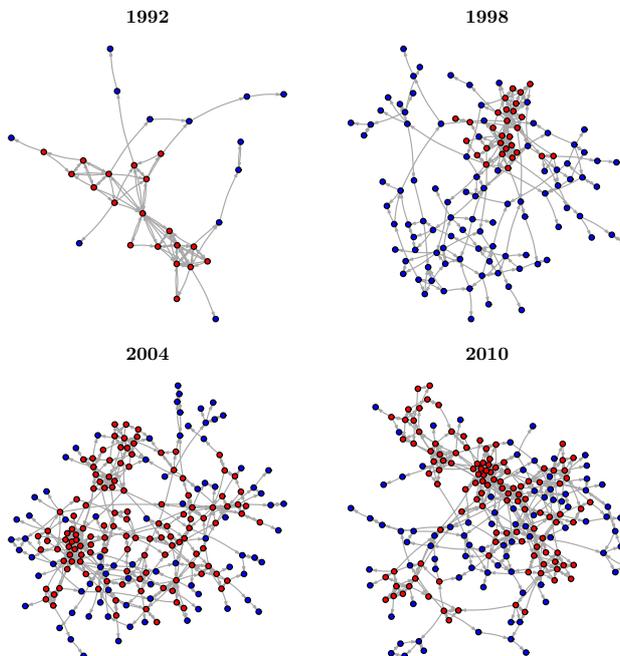


Figure 9: Network of technology subclasses: snapshots of the ACS.

Figure 10c shows the relative share of the three subsets of the subclasses network in terms of absolute fitness (number of patents). This figure presents a similar trend to the fitness in the network of IPC classes (Figure 5c): the ACS increases in relative terms, while the rest of the network shrinks. This time, the periphery of the ACS maintains a sizeable share (about 10%). In the last year (2011), the aggregate share of the ACS is about 60% for subclasses, compared to about 80% in the network of classes. However, considering that for subclasses the aggregate number of nodes in the various Autocatalytic sets is roughly one third of the population, the autocatalytic structure emerges as a dominant pattern also in the network of subclasses.

The evidence of the analysis performed confirms our hypothesis at the level of subclasses, that autocatalytic structures are fostering innovation as measured by the number of filed patents. Our hypothesis is thus verified also at the level of IPC subclasses as it is at the level of classes: like for technology classes, also the finer structure defined by subclasses clearly shows that technology fields benefit from the inclusion in the ACS. This implies that the self-reinforcing

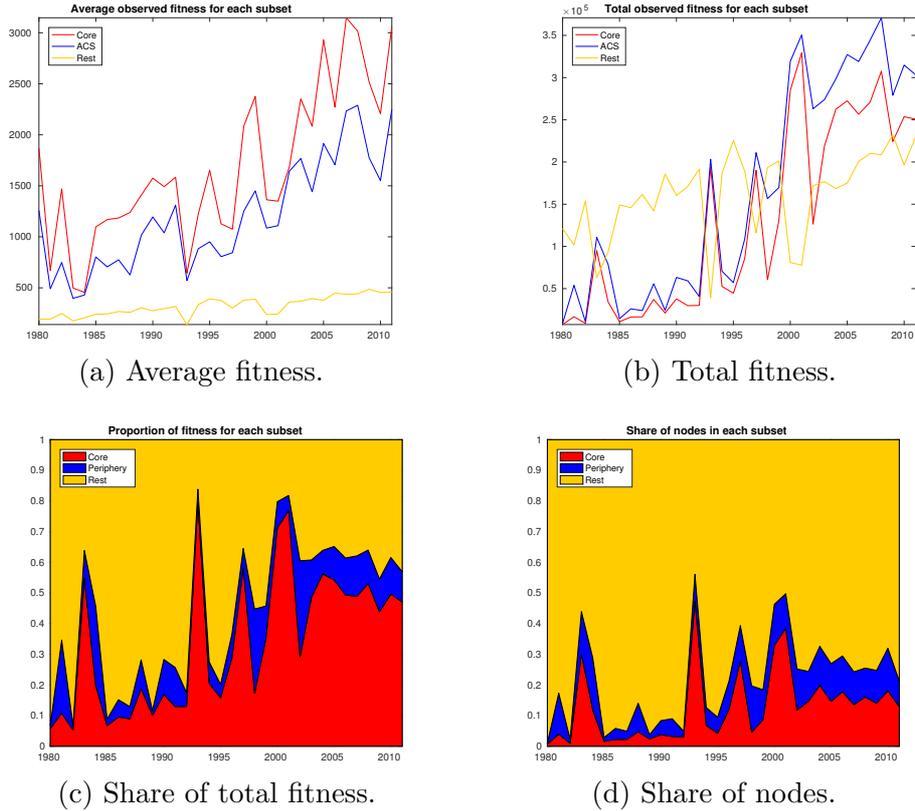


Figure 10: Total and average fitness, *i.e.* total and average number of patents per subclass in different portions of the technology network (ACS, core, and rest).

process of cumulative innovation giving rise to the same quicker pace of technological progress that we observed in Section 4.2 is detected also in a more disaggregated representation technology network.

B Variety in the ACS: evolution and significance

In this section, we investigate the composition of the ACS by running a statistical test of the uniformity of the contribution of IPC sections to the set. The null hypothesis is that the occurrence of each section within the ACS is compatible with an unbiased sampling without replacement according to the generalised hypergeometric distribution. The alternative hypothesis is sampling from Fisher’s non-central hypergeometric distribution [41]. The log-likelihood ratio statistic for each year is calculated and plotted in Figure 11, with a lower value indicating evidence in favour of the null hypothesis. This statistic can be used as a measure of variety for the population of the ACS. The asymptotic distribution of the log-likelihood ratio statistic is the chi-square distribution with 7 degrees of freedom. Its critical value corresponding to a 5% significance level is 14.07. As shown in Figure 11, in every year from 1980 until 2011 the statistics is above the threshold value, indicating a significant bias in terms of occupancy of the ACS by sections over the period considered, with an increasing trend.

C Links within and across IPC sections

This section focuses on whether links in the ACS more frequently connect classes belonging to the same IPC section or to different sections. Figure 12 shows a selection of snapshots of the

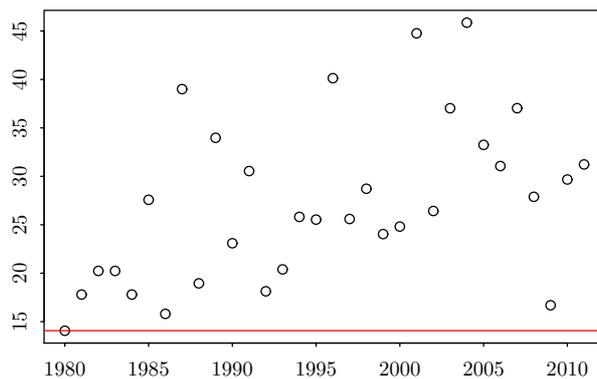


Figure 11: Variety measure of sections with respect to autocatalytic structures. Values are the log-likelihood ratio test statistic for each year contrasting the null hypothesis that occurrence in the ACS of each section happens with unbiased sampling without replacement for a fixed ACS size against the alternative hypothesis of biased sampling. The horizontal red line corresponds to a 5% significance level.

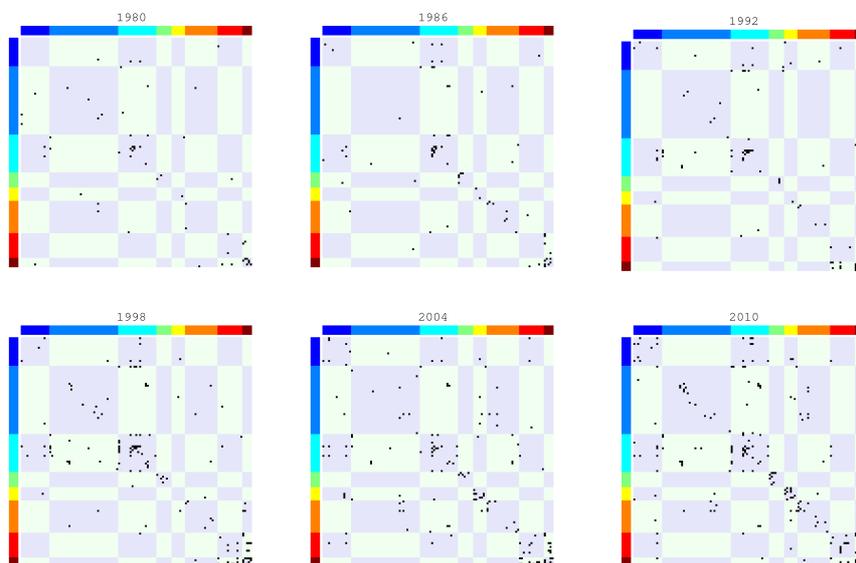


Figure 12: Adjacency matrix of the network of technology classes for a selection of uniformly sampled years. The 121 classes are ordered on the basis of their sections, which go from A (Human necessities – dark blue – in the top-left corner of each matrix), up to H (Electricity – brown – in the bottom right corner). Black dots indicate that a link is present between two classes.

technology network adjacency matrix sampling the analysed time period at regular intervals from 1980 to 2010. Comparing Figure 12 with Figure 8a of Section 4.3, it clearly emerges that the structure of the matrix is robust over time and that the point made in the main text concerning the structure of links in the technology network is generally valid. In particular, it can be seen that over time the number of links in the network increases but a significant proportion of them lies both within and across sections. This indicates that important knowledge flows connect classes of different sections, and that recombinant innovation is a stylised feature of technological change. Put differently, sections do not affect the distribution of links between classes alone, since between-section co-occurrences also appear to play a relevant role in some sections.