

RESEARCH

Open Access



Community structure in co-inventor networks affects time to first citation for patents

William Doonan¹, Kyle W. Higham^{1,2}, Michele Governale¹ and Ulrich Zülicke^{1*} 

*Correspondence:

uli.zuelicke@vuw.ac.nz

¹Te Pūnaha Matatini, School of Chemical and Physical Sciences, Victoria University of Wellington, PO Box 600, 6140 Wellington, New Zealand

Full list of author information is available at the end of the article

Abstract

We have investigated community structure in the co-inventor network of a given cohort of patents and related this structure to the dynamics of how these patents acquire their first citation. A statistically significant difference in the time lag until first citation is linked to whether or not this citation comes from a patent whose listed inventors share membership in the same communities as the inventors of the cited patent. Although the inventor-community structures identified by different community-detection algorithms differ in several aspects, including the community-size distribution, the magnitude of the difference in time to first citation is robustly exhibited. Our work is able to quantify the expected acceleration of knowledge flow within inventor communities and thereby further establishes the utility of network-analysis tools for studying innovation dynamics.

Keywords: Patent citations, Community detection, Knowledge propagation

Introduction and motivation

Inventions can be codified in patent applications that, if granted, bestow certain rights to the assignees. Patent documents contain citations to other patents as part of the requirement to acknowledge the state of prior art and to delimit the invention's legal scope. These citations are understood to represent knowledge flows between inventors and have been used widely to study various economic and social aspects of innovation dynamics (Hall et al. 2002; Breschi and Lissoni 2004; Jaffe and de Rassenfosse 2017).

The degree to which social and working relationships between inventors influence patenting and knowledge propagation in technological-innovation space has been the subject of recent interest (Stolpe 2002; Balconi et al. 2004; Singh 2005; Sorenson et al. 2006; Breschi and Lissoni 2009; Breschi and Lenzi 2016; Tóth et al. 2018). We shed new light on this topic from a network-analysis perspective by studying the effect of community structure in the co-inventor network on patent-citation dynamics. It can be intuitively expected that knowledge about inventions will be transmitted, and thus become available for utilization, faster within groups of inventors that have collaborated before. One of the main aims of our present work is to verify this expectation and quantify the acceleration of knowledge flow through inventor communities. We use the time lag until first citation (Gay et al. 2005; Lee and Sohn 2017) as a proxy measure for the fastest speed at which

knowledge about an invention can propagate. This is motivated by the fact that, as in the case of an electrical pulse whose leading edge is the real carrier of information, independent of the rest of its line shape, the time to first citation is more representative of the speed of information flow than any other, aggregate or average, citation measure.

In previous studies, inventor communities were identified using attributes such as geographical co-location (Jaffe et al. 1993; Singh 2005; Breschi and Lenzi 2016) or employment at the same firm (Kogut and Zander 1992; Stolpe 2002; Singh 2005; Breschi and Lissoni 2009). Here our approach is different. We use established community-detection algorithms (Fortunato 2010; Fortunato and Hric 2016) to identify communities of inventors in the co-inventor network that has been constructed by projection from a bipartite inventor-and-patent network. This method can directly reveal communities based on collaborative inventive activity without relying on any externally observed relationship proxies. On the other hand, the algorithmic identification of communities may introduce biases arising from idiosyncrasies inherent to particular detection methods. Our present work also investigates how robustly community effects in patent-citation dynamics can be identified using the generally different community structures obtained by various established detection algorithms.

Community structures in citation networks for scientific articles have been analyzed recently (see, for example, Refs (Lambiotte and Panzarasa 2009; Chen and Redner 2010; Liu et al. 2017)), and opportunities for similar studies on the available large amounts of rich patent-citation data are now also beginning to be realized (Nakai et al. 2018; Gao et al. 2018). Our study offers an example for the type of interesting insights that can be gained by applying network-analysis-based community-detection methods in the important area of technical innovation.

The remainder of this article is structured as follows. In the following “[Methods](#)” section, we start by discussing the data used to construct the co-inventor network and the five community-detection algorithms utilized in our work. Basic properties of communities identified on the largest connected component of the co-inventor network by the different algorithms are compared before details are given on how the citation-lag distributions are assembled. The subsequent “[Results and discussion](#)” section presents a detailed analysis of the observed distributions for the time to first citation, focusing in particular on how these depend on whether inventors from the citing and the cited patent share membership in one of the algorithmically detected communities. Our “[Conclusions](#)” are summarized in the final section. A list of abbreviations used throughout this article is given in [Table 1](#).

Methods

The availability of disambiguated US-patent data (Li et al. 2014) makes it possible to construct the co-inventor network associated with a particular cohort of patents. To be

Table 1 List of abbreviations

Abbreviation	Meaning
ARI	Adjusted Rand index
LCC	Largest connected component
USPC	US Patent Classification
USPTO	US Patent and Trademark Office

specific, we use patents granted by the US Patent and Trademark Office (USPTO) during 1995–1999 and assigned to one of the following classes from the US Patent Classification System (USPC) (US Patent Classification 2013): 257 (Active solid-state devices), 326 (Electronic digital logic circuitry), or 438 (Semiconductor device manufacturing: process)¹. The projection of the bi-partite patent-and-inventor network onto inventors yields the co-inventor network (Balconi et al. 2004) associated with this cohort. Going further than using a simple projection as was done in previous work (Balconi et al. 2004; Singh 2005), we introduce edge weights w_{ij} reflecting the frequency and intensity of the co-inventing activity (Newman 2001b) between two inventors i and j ;

$$w_{ij} = \sum_{\alpha} \frac{\delta_i^{(\alpha)} \delta_j^{(\alpha)}}{n_{\alpha} - 1} . \quad (1)$$

The sum in Eq. (1) is over all patents in the chosen cohort that list more than one inventor, and $\delta_i^{(\alpha)} = 1$ ($\delta_i^{(\alpha)} = 0$) if inventor i is (not) among the n_{α} inventors listed on the particular patent α . In the following, we restrict ourselves to considering only the largest connected component (LCC) of the thus-constructed co-inventor network, which can be expected to capture the most relevant aspects of inventors' connectedness (Newman 2001a; Breschi and Lissoni 2009). Amounting to approximately 31% of the total co-inventor network by number of nodes, the LCC comprises a diverse assortment of inventors as evidenced by the spread of firms their patents are assigned to. For example, IBM, Hitachi LTD, Motorola INC, Kabushiki Kaisha Toshiba, and Texas Instruments represent 14%, 7.2%, 4.6%, 4.1%, and 4.0% of nodes in the LCC, respectively. Other companies each account for portions smaller than 4% of the total in the component. Table 2 provides an overview of relevant summary statistics pertaining to the patent cohort and co-inventor network considered here.

We have used five established community-detection algorithms to analyze the LCC of the co-inventor network: Greedy (Newman 2004), Louvain (Blondel et al. 2008), Infomap (Rosvall and Bergstrom 2008), Random Walks (Pons and Latapy 2006), and Propagating Labels (Raghavan et al. 2007). Although different algorithms generally yield different community structures, clear similarities are exhibited between the structures obtained by conceptually related approaches such as Greedy and Louvain on the one hand, and Infomap, Random Walks, and Propagating Labels on the other. In particular, we observe the well-known resolution-limit issue (Fortunato and Barthélemy 2007) where the communities delineated by approaches that maximise modularity (Newman and Girvan 2004) (Greedy and Louvain) are typically larger in size and generally subsume the many

Table 2 Summary statistics of patent data utilized for this study

	No. of patents	No. of inventors
Full patent cohort	46 034	49 703
Patents gaining a first citation within 10 years	44 137	—
Largest connected component (LCC) of co-inventor network	—	15 525
Patents associated with inventors from the LCC	20 021	—
LCC-associated patents gaining a first citation within 10 years	19 284	—
LCC-associated patents with a first citation by LCC inventors	16 271	—
Subset of these patents whose first citation is a self-citation	2 584	—


The full patent cohort comprises the USPTO patents from USPC classes 257, 326, and 438 granted during 1995–1999. Citations are tracked for each patent during the 10-year period after its time of grant

smaller communities identified by other approaches (Infomap, Random Walks, Propagating Labels). This may be inferred graphically from visualizations of community structures, such as those shown in Fig. 1. A more quantitative comparison is possible based on the size distributions for communities obtained by application of each of the five different detection algorithms to the LCC, given in Fig. 2, and the corresponding community-structure-related summary statistics provided in Table 3. We also use the adjusted Rand index (ARI) (Hubert and Arabie 1985) to measure similarity between community structures generated by different algorithms as well as those from different runs of the same algorithm. ARI scores close to 1 are obtained when comparing the results of multiple runs of the same algorithm on the LCC of the co-inventor network, indicating the generally very good or, for Propagating Labels, at least satisfactory robustness of each method. The ARI value of 0.8 found when comparing the community structures generated from the Greedy and Louvain algorithms attests to their high degree of similarity. Similarity to a much lesser extent is exhibited between the partitionings arising from the Infomap and Propagating-Labels methods, as well as between those of Greedy/Louvain and Random Walks. All other pairwise comparisons yield a very small ARI score of 0.2 and below.

Having identified inventor communities in the LCC on the co-inventor network, we track the citations acquired by patents from the chosen cohort over a ten-year period from each individual patent's time of grant. Following the usual convention (Hall et al. 2002; Mehta et al. 2010), we assign the time of citation to be the application date of the citing patent. Hence, citations to patents in the cohort considered here can only originate from patents applied for before 2010. All citation-related data are sourced directly from the USPTO (USPTO Bulk Data Storage System 2016). For the purpose of the present work, we specifically consider the time Δt_1 elapsed after each patent's time of grant until it acquires its first citation (Gay et al. 2005; Lee and Sohn 2017) and determine this citation to be either a self-citation, an in-community citation, or an out-of-community citation based on the previously obtained community structure. Here a *self*-citation occurs if any of the inventors listed on the citing patent is also listed as an inventor on the cited patent, which is the case independently of any community structure. An *in-community* citation is a *non-self*-citation from a patent where at least one of the listed inventors shares membership in one of the algorithmically identified communities with at least one of the inventors listed on the cited patent. An *out-of-community* citation is from a patent whose listed inventors belong to different communities than the inventors of the cited patent. To focus most precisely on how community structure affects the knowledge flow through the co-inventor network, we only consider citations from patents that have at least one of their listed inventors belonging to the LCC of the co-inventor network defined via the cohort of cited patents². Table 2 provides citation-related summary statistics for the analyzed patent cohort.

Distributions of time lags Δt_1 to first citation obtained for each subset of citation type (self, in-community, and out-of-community) are juxtaposed in Fig. 3. Their distinctive properties are discussed in greater detail in the next section. Note that, because the time of citation is the *application* date of the citing patent, and the time lag Δt_1 is measured from the *grant* date of the cited patent, negative values of Δt_1 are possible and, in fact, occur quite frequently. That patents have often already accrued citations by the time of grant from other patents that were applied for before that date but granted afterwards is a well-known feature of patent-citation dynamics (Mehta et al. 2010).

Louvain



Infomap


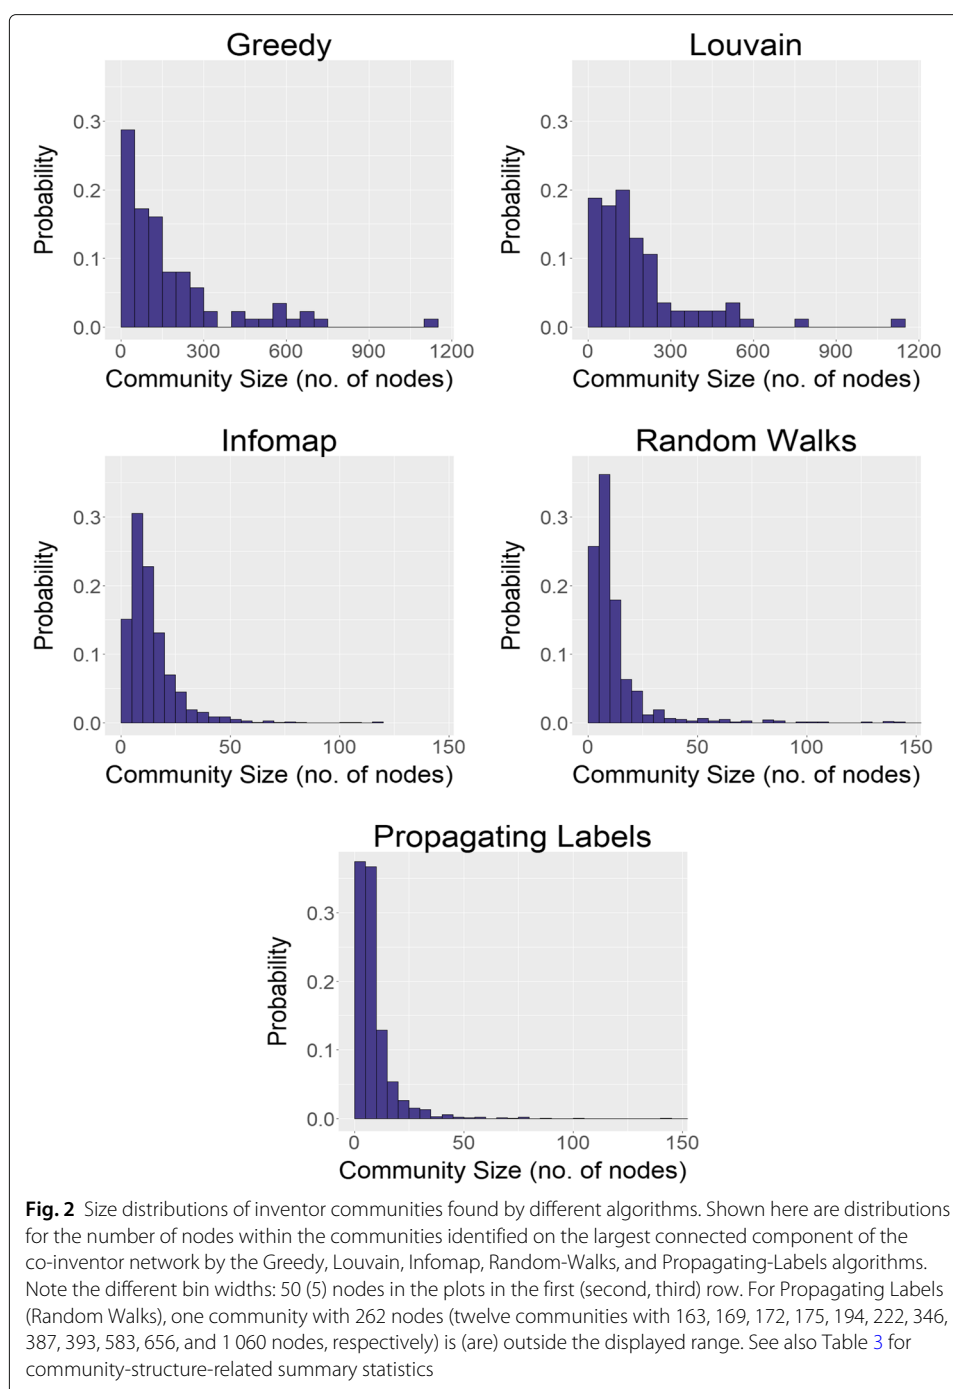


Fig. 1 Community structure within the co-inventor network. The upper (lower) panel indicates communities in the largest connected component of the co-inventor network identified by the Louvain (Infomap) algorithm using different colors (but with some colors being reused in the lower panel due to the overall large number of communities yielded by Infomap). Note how certain groups of communities identified as separate by Infomap are clustered together by Louvain. The size of the circle indicating a node is proportional to the number of edges attached to that node. Images created using Gephi (Bastian et al. 2009)



Results and discussion

As is apparent from the examples shown in Fig. 3, distributions of time lags to first citation are generally broad and skewed. We find that log-normal distributions provide a successful fit to their line shapes, except for the distinctive peak for the $\Delta t_1 = 0$ bin exhibited by the time-lag distribution of first citations that are self-citations. Further research is needed to elucidate the origin of the inflated probability for a zero time lag for self-citations; we can only speculate at this point that it could be the result of larger firms' patenting strategies.

Table 3 Comparison of inventor-community structures found by different algorithms

	Greedy	Louvain	Infomap	Rand. Wlks.	Prop. Lbls.
Total no. of communities identified	87	85	1 157	948	1 800
Size of largest community (nodes)	1 143	1 139	120	1 060	262
Similarity to Greedy (ARI)	1.0	0.8	0.1	0.5	0.1
Similarity to Louvain (ARI)	0.8	1.0	0.1	0.5	0.1
Similarity to Infomap (ARI)	0.1	0.1	0.9	0.2	0.6
Similarity to Rand. Wlks. (ARI)	0.5	0.5	0.2	1.0	0.2
Similarity to Prop. Lbls. (ARI)	0.1	0.1	0.6	0.2	0.7
No. of in-community first citations	3 179	3 031	1 415	2 772	1 398

Summary statistics provided here pertain to the communities obtained by applying each of the five indicated community-detection algorithms to the largest connected component of the co-inventor network considered in this work. Similarity between community structures is quantified in terms of the adjusted Rand index (ARI) (Hubert and Arabie 1985). Listed ARI values are averages calculated for pairs of community structures generated by multiple runs of the respective algorithms that are being compared. The total number of first citations classified as in-community citations based on each of the five community partitions is also given

The analysis of time-lag distributions for in-community vs. out-of-community first citations enables us to discuss the influence of community structure on citation dynamics, for each set of communities identified by the five different algorithms used in this work. Specifically, we consider distributional averages (mean and median values), as well as the most probable value (mode). Table 4 gives the results obtained for these by two different

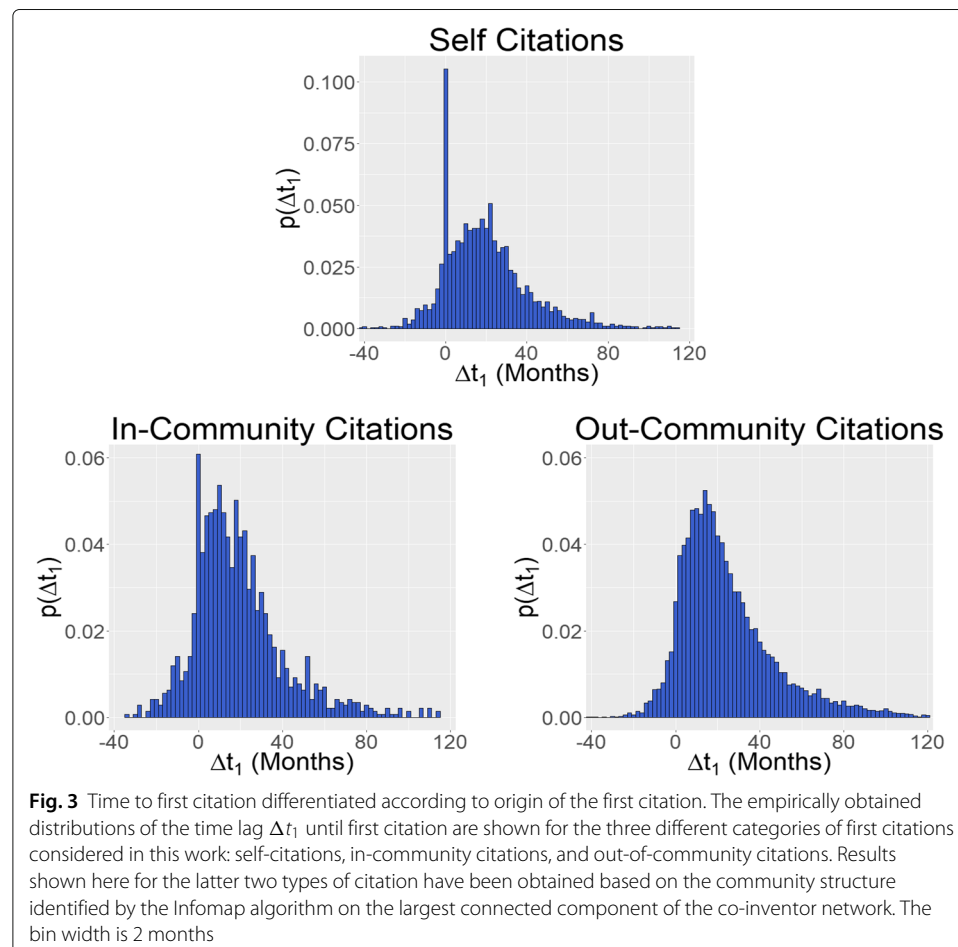


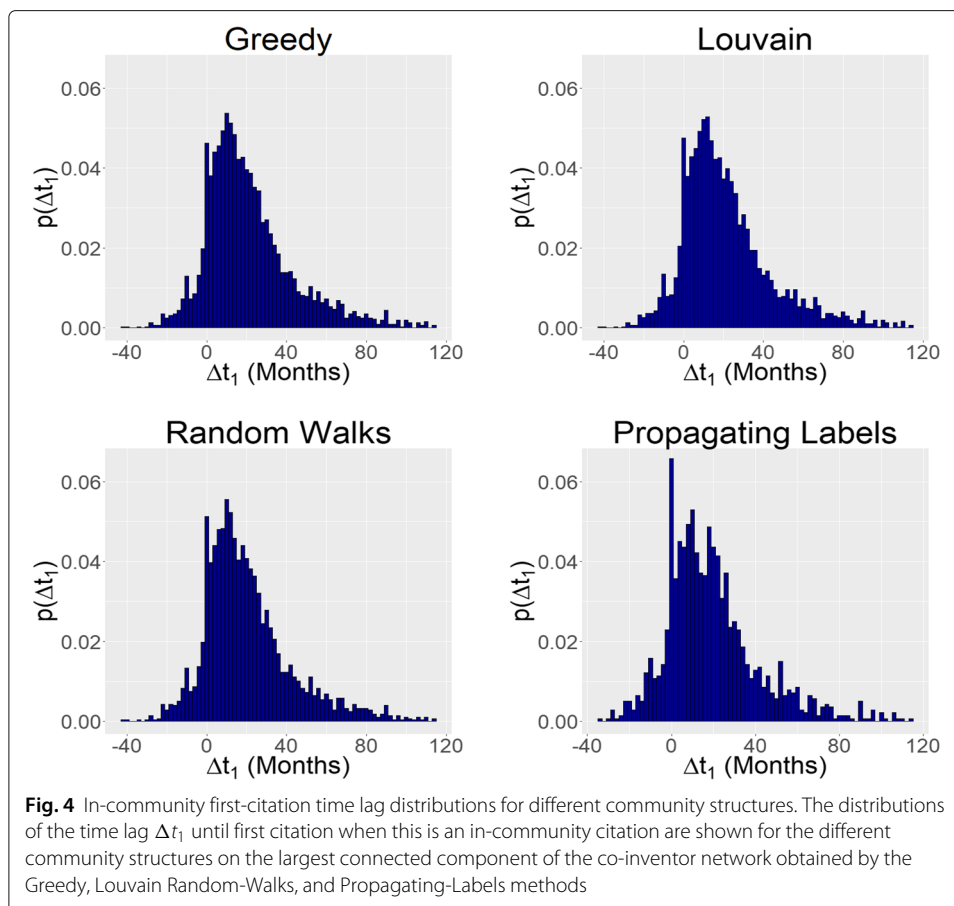
Table 4 Community-related differences in the time lag Δt_1 until first citation

		Infomap	Prop. Lbls.	Rand. Wlks.	Greedy	Louvain
In	raw data $\overline{\Delta t_1}$	19.4	19.7	20.8	21.6	21.7
	raw data $\widetilde{\Delta t_1}$	16	16	17	17	17
	raw data ³ $\text{Mo}_{\Delta t_1}$	0	0	10	10	12
	lognorm fit $\overline{\Delta t_1}$	18 ± 2	18 ± 2	18.8 ± 0.9	19.5 ± 0.9	19.5 ± 0.9
	lognorm fit $\widetilde{\Delta t_1}$	16 ± 1	16 ± 1	16.7 ± 0.7	17.4 ± 0.7	17.4 ± 0.7
	lognorm fit $\text{Mo}_{\Delta t_1}$	12 ± 2	12 ± 2	12.6 ± 0.9	12 ± 2	13.3 ± 0.9
Out	raw data $\overline{\Delta t_1}$	24.6	24.5	24.8	24.7	24.7
	raw data $\widetilde{\Delta t_1}$	19	19	20	20	20
	raw data $\text{Mo}_{\Delta t_1}$	15	15	15	15	15
	lognorm fit $\overline{\Delta t_1}$	21.5 ± 0.9	21.5 ± 0.9	21.5 ± 0.9	21.5 ± 0.9	21.5 ± 0.9
	lognorm fit $\widetilde{\Delta t_1}$	19.4 ± 0.7	19.4 ± 0.7	19.4 ± 0.7	19.4 ± 0.7	19.4 ± 0.7
	lognorm fit $\text{Mo}_{\Delta t_1}$	15.3 ± 0.9	15.4 ± 0.9	15.3 ± 0.9	15.3 ± 0.9	15.4 ± 0.9
	<i>t</i> -test full dataset	8.4	6.6	6.8	8.7	7.6
	av. <i>t</i> -score subset	2.9	1.7	1.7	2.3	2.6

The mean ($\overline{\Delta t_1}$), median ($\widetilde{\Delta t_1}$), and modal ($\text{Mo}_{\Delta t_1}$) values derived from the raw data for the empirically observed distributions of time until the first citation that is not a self-citation and either originates from within (In) or outside (Out) co-inventor communities are listed for community structures obtained using five different detection algorithms. Corresponding values derived from fits of the empirical distributions to log-normal form are also given for comparison. Welch's *t*-test scores (Welch 1947) establish the statistical significance for the observed differences between mean time lags for in-community and out-of-community citations. We give *t*-values obtained from tests applied to the full datasets of in-community and out-of-community citation lags, as well as the average *t*-score found when comparing 500 independently and randomly chosen subsets of 300 citations from the in- and out-of-community datasets

methods: direct calculation using the empirical data for the time-lag distributions, and the values derived from parameters of the fitted log-normal distributions. For the mean and median values of time lags for in-community first citations, there is excellent agreement, within uncertainties, between the two approaches for all five community structures considered. Similarly good agreement exists between the medians and modes of time lags for out-of-community first citations. The disagreement between the modes calculated directly and from the log-normal fit for in-community first-citation time lags obtained based on the Infomap and Propagating Labels algorithms is due to the slightly inflated probability for the $\Delta t_1 = 0$ bin exhibited in these cases. See Figs. 3 and 4. The magnitude of the inflated peak at $\Delta t_1 = 0$ for in-community first-citation time lags is generally comparable to the size of statistical fluctuations and overall much smaller than that observed in the time-lag distribution of first self-citations. The consistently higher mean obtained for the out-of-community first-citation time lag using the raw data as compared with that derived from the log-normal-distribution fit parameters can be traced to the fact that the fit systematically underestimates probabilities in the tail of the time-lag distribution for this type of citation. Overall, the magnitude for the means and medians found in this work for the first-citation time lag agrees very well with that reported previously in the literature (Trajtenberg 1990; Gay et al. 2005; Fisch et al. 2017; Lee and Sohn 2017).

The specific community structure determined for the co-inventor network is observed to depend sensitively on the employed algorithm. The Greedy and Louvain algorithms yield a comparatively small number of larger communities, whereas the Infomap and Propagating-Labels algorithms identify mostly much smaller communities. The community structure obtained using Random Walks lies somewhat inbetween these two extremes. See Fig. 2 and Table 3. Interestingly, the distributional properties of the time lag for in-community first citations, and even more so those pertaining to out-of-community



first citations, are found to be quite similar irrespective of the significant differences between the underlying inventor-community structures. Figure 4 shows the distributions of time lags for in-community citations based on the community structures obtained from running the Greedy, Louvain, Random-Walks, and Propagating-Labels algorithms on the LCC of the co-inventor network. Corresponding results for the community structure identified by the Infomap algorithm are given in Fig. 3. The distributions of the time lag for out-of-community first citations obtained based on the communities identified by the five different algorithms are visually barely distinguishable and therefore not shown. This suggests that algorithms producing larger communities have generally agglomerated parts of the co-inventor network that are separate communities as far as citing behavior is concerned. Hence, our study provides another real-world verification of the resolution limit (Fortunato and Barthélemy 2007) associated with algorithms based on optimizing modularity.

Comparison of the median and mean values for the time lags to first citation found for the in-community and out-of-community types, respectively, shows systematic differences. See Table 4. In particular, the median time lags to in-community first citations are about 2-3 months shorter than the corresponding medians for out-of-community first citations. The difference between the means of the time lag for in-community and out-of-community first citations is about 2-5 months. Given that we have analyzed time lags to first citation for patents whose inventors are all connected by prior co-authorship

of patents, this indicates a strong influence of community structure on patent-citation dynamics and the associated knowledge flow. Further comparison with the mean, median, and modal values found for the time lag to first citations that are self-citations is very instructive. See Table 5. The mean and median calculated using the raw data, including the inflated peak at $\Delta t_1 = 0$, are found to be essentially the same as for in-community citations. Hence, on average, in-community first citations and first-self-citations occur after similar time periods that are, again on average, shorter by about 3 months than the time lag for out-of-community first citations. However, the distributions of time lags for the in-community and self types of first citations differ significantly. This can be illustrated by analyzing the adjusted distribution of time lags for self-citations where the inflated probability at $\Delta t_1 = 0$ has been replaced with an interpolated value, or with the peak removed entirely. The mean, median and modal values found for this modified distribution agree almost exactly with those found for out-of-community citations. See again Table 5. This interesting combination of properties invites further investigation. A tentative speculation could be made that both the out-of-community type of first citation and the first-self-citations outside the $\Delta t_1 = 0$ peak originate largely from examiners, whereas the in-community type and the self-citations from the inflated $\Delta t_1 = 0$ peak are generally made by inventors themselves. As inventor and examiner citations have been separately identified on USPTO patents granted since 2001, a follow-up study utilizing a large-enough cohort of post-2001 patents with sufficiently long time window for citation accrual should soon be possible.

We employed Welch's t -test (Welch 1947) to establish whether the difference between the means of the citation lag for in-community and out-of-community citations is statistically significant. Results are summarized in Table 4. The t -scores ranging between 6.6 and 8.7 indicate that the differences in the means of citation lags are statistically extremely significant. Although the t -test is expected to yield reliable scores for non-normal distributions (Lumley et al. 2002), we also applied it to the, to a good approximation normal, distribution of the log of the time lag to first citation (after shifting the latter by a constant to eliminate negative values). Also in this case, similarly large t -scores reject the possibility of the means being equal with high confidence.

The high t -values generated in our application of the Welch test are likely a consequence of the very large sample sizes (1415 data points from in-community citations, 12 272 from out-of-community citations) that increase sensitivity to small differences. To illustrate the effect sample size has on the outcome of the t -test, we selected a random and independently chosen subset of 300 citations from each of the in-community and out-of-community datasets and performed Welch's t -test on these small samples. This was repeated 500 times to obtain averages of such t -values, which are also listed in Table 4. In the case of the Infomap-generated community partition, the mean t -value resulting

Table 5 Distributional properties of the time lag until first citation when this is a self-citation

	$\overline{\Delta t_1}$	$\widetilde{\Delta t_1}$	$\text{Mo}_{\Delta t_1}$
Using all raw data from empirical distribution	19.9	17	0
With $\Delta t_1 = 0$ peak adjusted by interpolation	21.7 ± 0.9	19.4 ± 0.7	15.0 ± 0.9
With $\Delta t_1 = 0$ peak removed	21.6 ± 0.9	19.4 ± 0.7	15.2 ± 0.9

The mean ($\overline{\Delta t_1}$), median ($\widetilde{\Delta t_1}$), and modal ($\text{Mo}_{\Delta t_1}$) values derived directly from the empirically observed distribution are shown alongside corresponding values obtained when the high peak at $\Delta t_1 = 0$ is replaced by the average of probabilities observed for $\Delta t_1 = \pm 2$ months or completely removed

from this procedure is 2.9, along with 83% of the small-sample simulations finding differences between the distributions' means to be significant to the 5% α level. This indicates that sample sizes of at least 300 first citations are needed to detect with sufficient confidence the time delay reported here between those originating from within and outside of co-inventor communities.

As further proof of the inventor-community-related cause of acceleration in patent-citation dynamics, we performed the following control experiment. Starting with the Infomap-generated community partition of the LCC, inventors were randomly re-assigned to communities within this given structure. As a result, the high-level morphology (number of communities and their individual sizes) of the partition was left unchanged, but the groupings of inventors within each community became completely random, as indicated by the ARI value of 0.0002 obtained when comparing the initial and randomized partitions. Repeating the analysis of first citations, we observe that the total number of in-community citations occurring in the randomized community structure has dropped precipitously to 129, i.e., about 10 % of the in-community citations present in the original structure. Furthermore, the mean (median) time to a first in-community citation in the randomized structure is determined to be 23.9 months (18 months), which is not significantly earlier anymore than the corresponding value of 24.4 months (19 months) found for the out-of-community first citations. The Welch's t -test score of 0.24 also indicates that there is no significant difference between the means of the distributions of times to first in-community and out-of-community citations in the situation with randomized inventor assignment to communities. The significantly reduced total number of in-community citations, and the disappearing difference between mean times to first citations that originate from within and outside of communities, both convincingly indicate that the originally established inventor communities are indeed the platforms for accelerated patent-citation dynamics.

We close this section by discussing several confounding variables whose influence may weaken our inference of a general non-trivial community-structure effect on citation dynamics. *Inventor team size*: All other things being equal, inventors working in larger teams will have a greater chance of acquiring in-community first citations to their patents, which would also be likely to occur more quickly. However, the scope for this trivial mechanism causing the observed effect is severely limited by the fact that large inventor teams are actually quite rare. See, e.g., data presented in Ref. Crescenzi et al. (2016). The average community size found in our work certainly exceeds the average inventor-team size given for the period 1995-1999 in Fig. 4 of that article. *Firm-level associations*: Certain inventor communities may just be a reflection of the inventors' employment at the same firm and, for their case, intra-firm information channels acting in parallel to prior co-invention activity may drive the observed acceleration in the in-community citation dynamics. However, for large firms, especially those with multiple geographically separated R&D centers, this alternative mechanism could be ineffective. Disentangling trivial firm-association effects from knowledge flows established via real inventor collaboration would be an interesting direction for future research. *Niche-technology associations*: The patent cohort studied in this work relates to the broad and technologically crowded semiconductor industry, and our detected communities may correlate with very specific technology types. The fact that citations are likely to come first from inventors in the same community could then arise simply because inventors that are working in the same

industrial niche are likely to build on each others' technological advances before inventors working in technologically more distant fields. To clarify this issue, the distributions of technology specializations across the respective in-community and out-of-community citing-patent cohorts would need to be studied in greater detail using a suitable proxy measure for technological similarity that could be defined either within (Trajtenberg et al. 1997) or beyond (Kay et al. 2014) existing classifications schemes.

Conclusions

We have investigated community structure on the co-inventor network associated with a particular cohort of USPTO patents where edge weights reflect the frequency of inventors' collaborative patenting activity. Five established community-detection algorithms (Greedy, Louvain, Infomap, Random Walks, and Propagating Labels) were deployed to identify communities on this network's largest connected component. The sizes and numbers of communities found by the different algorithms varied, with some similarities exhibited by algorithms using related methodology. Table 3 and Fig. 2 provide details enabling a quantitative comparison between the properties of the algorithmically detected community structures.

To investigate the effect of inventor communities on patent-citation dynamics, we analyzed the time lag to the first citation received by the patents associated with inventors from the largest connected component of the co-inventor network. Only citations originating from patents co-authored by at least one of the inventors from the largest connected component were counted for the present study. Three different types of first citation were distinguished: self-citations, in-community (non-self-)citations, and out-of-community citations. The distributions of time lags for each type of first citation were observed to have distinctive properties. Figure 3 shows results obtained based on the community structure found by the Infomap algorithm. The mean, median, and modal values for each type of distribution were determined to enable a quantitative comparison of the speeds of knowledge flow within and between different inventor communities. Results for all community structures investigated in the present work are summarized in Table 4. The median time delay to first citation for the out-of-community type turns out to be typically 3 months longer than for the in-community type. Self-citations and in-community-type citations have approximately the same median time lag, even though the distributions of time lags for these two types of first citation are markedly different. The difference between the mean time lag observed for out-of-community and in-community first citations is generally even larger than that found for the corresponding median values. Although the communities identified by the different detection algorithms utilized in this work differ in some detail, the observed influence of community structure on patent-citation dynamics was found to agree closely, even on a quantitative level. Thus our results provide a rather general quantification for the accelerated knowledge flow through inventor communities formed via collaborative patenting. Furthermore, the observation that association with distinct inventor communities based on previous co-authorship on patents results in faster citation of an inventor's later patents by members of that community provides strong further evidence in support of the fundamental importance of such collaboration-based social connections (Singh 2005; Agrawal et al. 2006; Breschi and Lissoni 2009; Breschi and Lenzi 2016) that has not always been able to be clearly observed (Stolpe 2002).

Our focus on the time lag to patents' first citation was motivated by the expectation that this quantity will likely be the best proxy measure for a real difference in time scales for knowledge propagation within and outside of co-inventor communities. It would be interesting to also compare the mean time lags between later (i.e., second, third, . . . n th) in-community and out-of-community citations. Such a study should be able to observe how the in-community advantage for knowing earlier about an invention diminishes over its repeated utilization.

The results of our work point the way to other interesting directions for future research. For example, analyzing the characteristics of the algorithmically identified inventor communities could yield useful information regarding the structure of effective innovation teams, extending previous work that focused only on direct collaborations between inventors (Crescenzi et al. 2016). Studying the dependence of the observed acceleration of information flow through inventor communities on the field and type of inventions may yield a measure for the speed at which the knowledge frontier moves in different parts of innovation space. Community-detection methods could also be deployed to elucidate relationships shaping innovation activity beyond the network of inventors, e.g., on the level of firms (Schilling and Phelps 2007) or other organizations. Thus opportunities abound for the useful application of modern network-analysis tools to innovation economics (Acemoglu et al. 2016) and related social-science studies (Hidalgo 2016).

Endnotes

¹The objectives of a particular investigation will generally determine the choice of cohort. Ours is motivated by three main needs: (i) to leave a sufficiently long time window for citation accrual, (ii) to have a sufficiently large data set to reduce statistical noise, and (iii) to include patents from similar fields that are granted around the same time to minimize variations in the inventors' work environment.

²Excluding citations by inventors outside the LCC amounts to neglecting trivial out-of-community citations that would likely influence the time-lag distribution for this type of citation mostly by shifting weight to larger Δt_1 (Singh 2005; Sorenson et al. 2006; Breschi and Lissoni 2009). If at all relevant, this could only further accentuate the differences between characteristic values (mean, median, and mode) for the time lag to first citation found for in-community and out-of-community citations.

³binned into 2-months-wide bins as shown in Figs. 3 and 4

Acknowledgements

The authors obtained useful insights into companies' patenting practices from conversations with L. Colombo and F. Natali, discussed community-detection algorithms with S. Marsland and D. O'Neale, and received advice on statistical testing from I. Doonan.

Funding

The work of WD was supported by a Victoria University of Wellington Faculty of Science Strategic Grant (GMS grant no. 217470). Funding from the New Zealand Tertiary Education Commission's Centres of Research Excellence Fund via Te Pūnaha Matatini is gratefully acknowledged by KWH, MG and UZ.

Availability of data and materials

The datasets supporting the conclusions of this article are available in the Zenodo repository, <https://doi.org/10.5281/zenodo.2574342>.

Authors' contributions

KWH, MG and UZ designed the research. WD constructed the co-inventor network and performed the data analysis, with assistance from KWH. All authors contributed to analyzing the results. UZ wrote the paper with input from all authors. All authors read and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Author details

¹Te Pūnaha Matatini, School of Chemical and Physical Sciences, Victoria University of Wellington, PO Box 600, 6140 Wellington, New Zealand. ²Present Address: College of Management of Technology, EPFL, Odyssea, Station 5, 1015 Lausanne, Switzerland.

Received: 12 November 2018 Accepted: 2 April 2019

Published online: 18 April 2019

References

- Acemoglu D, Akcigit U, Kerr WR (2016) Innovation network. *Proc Natl Acad Sci USA* 113:11483–11488. <https://doi.org/10.1073/pnas.1613559113>
- Agrawal A, Cockburn I, McHale J (2006) Gone but not forgotten: knowledge flows, labor mobility, and enduring social relationships. *J Econ Geogr* 6:571–591. <https://doi.org/10.1093/jeg/lbl016>
- Balconi M, Breschi S, Lissoni F (2004) Networks of inventors and the role of academia: an exploration of Italian patent data. *Res Policy* 33:127–145. [https://doi.org/10.1016/S0048-7333\(03\)00108-2](https://doi.org/10.1016/S0048-7333(03)00108-2)
- Bastian M, Heymann S, Jacomy M (2009) Gephi: An open source software for exploring and manipulating networks. In: Adar E, Hurst M, Finin T, Glance N, Nicolov N, Tseng B (eds). *Proceedings of the Third International AAAI Conference on Weblogs and Social Media*. AAAI Press, pp 361–362
- Blondel VD, Guillaume J, Lambiotte R, Lefebvre E (2008) Fast unfolding of communities in large networks. *J Stat Mech*:10008. <https://doi.org/10.1088/1742-5468/2008/10/P10008>
- Breschi S, Lenzi C (2016) Co-invention networks and inventive productivity in US cities. *J Urban Econ* 92:66–75. <https://doi.org/10.1016/j.jue.2015.12.003>
- Breschi S, Lissoni F (2004) Knowledge networks from patent data: Methodological issues and research targets. In: Moed HF, Glänzel W, Schmoch U (eds). *Handbook of Quantitative Science and Technology Research*. Kluwer, Dordrecht, pp 613–643
- Breschi S, Lissoni F (2009) Mobility of skilled workers and co-invention networks: an anatomy of localized knowledge flows. *J Econ Geogr* 9:439–468. <https://doi.org/10.1093/jeg/lbp008>
- Chen P, Redner S (2010) Community structure of the physical review citation network. *J Informetrics* 4:278–290. <https://doi.org/10.1016/j.joi.2010.01.001>
- Crescenzi R, Nathan M, Rodríguez-Pose A (2016) Do inventors talk to strangers? On proximity and collaborative knowledge creation. *Res Policy* 45:177–194. <https://doi.org/10.1016/j.respol.2015.07.003>
- Fisch C, Sandner P, Regner L (2017) The value of Chinese patents: An empirical investigation of citation lags. *China Econ Rev* 45. <https://doi.org/10.1016/j.chieco.2017.05.011>
- Fortunato S (2010) Community detection in graphs. *Phys Rep* 486:75–174. <https://doi.org/10.1016/j.physrep.2009.11.002>
- Fortunato S, Barthélemy M (2007) Resolution limit in community detection. *Proc Natl Acad Sci USA* 104:36–41. <https://doi.org/10.1073/pnas.0605965104>
- Fortunato S, Hric D (2016) Community detection in networks: A user guide. *Phys Rep* 659:1–44. <https://doi.org/10.1016/j.physrep.2016.09.002>
- Gao Y, Zhu Z, Kali R, Riccaboni M (2018) Community evolution in patent networks: technological change and network dynamics. *Appl Netw Sci* 3:26. <https://doi.org/10.1007/s41109-018-0090-3>
- Gay C, Le Bas C, Patel P, Touach K (2005) The determinants of patent citations: an empirical analysis of French and British patents in the US. *Econ Innov New Technol* 14:339–350. <https://doi.org/10.1080/1040859042000307329>
- Hall BH, Jaffe AB, Trajtenberg M (2002) The NBER patent citations data file: Lessons, insights and methodological tools. In: Jaffe AB, Trajtenberg M (eds). *Patents, Citations, and Innovations: A Window on the Knowledge Economy*. MIT Press, Cambridge, MA, pp 403–460
- Hidalgo CA (2016) Disconnected, fragmented, or united? a trans-disciplinary review of network science. *Appl Netw Sci* 1:6. <https://doi.org/10.1007/s41109-016-0010-3>
- Hubert L, Arabie P (1985) Comparing partitions. *J Classif* 2:193–218. <https://doi.org/10.1007/BF01908075>
- Jaffe AB, de Rassenfosse G (2017) Patent citation data in social science research: Overview and best practices. *J Assoc Inf Sci Technol* 68:1360–1374. <https://doi.org/10.1002/asi.23731>
- Jaffe AB, Trajtenberg M, Henderson R (1993) Geographic localization of knowledge spillovers as evidenced by patent citations. *Quart J Econ* 108:577–598. <https://doi.org/10.2307/2118401>
- Kay L, Newman N, Youtie J, Porter AL, Rafols I (2014) Patent overlay mapping: Visualizing technological distance. *J Assoc Inf Sci Technol* 65:2432–2443. <https://doi.org/10.1002/asi.23146>
- Kogut B, Zander U (1992) Knowledge of the firm, combinative capabilities, and the replication of technology. *Organ Sci* 3:383–397. <https://doi.org/10.1287/orsc.3.3.383>
- Lambiotte R, Panzarasa P (2009) Communities, knowledge creation, and information diffusion. *J Informetrics* 3:180–190. <https://doi.org/10.1016/j.joi.2009.03.007>
- Lee J, Sohn SY (2017) What makes the first forward citation of a patent occur earlier?. *Scientometrics* 113:279–298. <https://doi.org/10.1007/s11192-017-2480-1>
- Li G-C, Lai R, D'Amour A, Doolin DM, Sun Y, Torvik VI, Yu AZ, Fleming L (2014) Disambiguation and co-authorship networks of the US patent inventor database (1975–2010). *Res Policy* 43:941–955. <https://doi.org/10.1016/j.respol.2014.01.012>
- Liu W, Nanetti A, Cheong SA (2017) Knowledge evolution in physics research: An analysis of bibliographic coupling networks. *PLoS ONE* 12:0184821. <https://doi.org/10.1371/journal.pone.0184821>
- Lumley T, Diehr P, Emerson S, Chen L (2002) The importance of the normality assumption in large public health data sets. *Annu Rev Public Health* 23:151–169. <https://doi.org/10.1146/annurev.publhealth.23.100901.140546>
- Mehta A, Rysman M, Simcoe T (2010) Identifying the age profile of patent citations: new estimates of knowledge diffusion. *J Appl Econ* 25:1179–1204. <https://doi.org/10.1002/jae.1086>

- Nakai K, Nonaka H, Hentona A, Kanai Y, Sakumoto T, Kataoka S, Carreón ECA, Hiraoka T (2018) Community detection and growth potential prediction using the Stochastic Block Model and the long short-term memory from patent citation networks. In: 2018 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM). pp 1884–1888. <https://doi.org/10.1109/IEEM.2018.8607487>
- Newman MEJ (2001a) Scientific collaboration networks. I. Network construction and fundamental results. *Phys Rev E* 64:016131. <https://doi.org/10.1103/PhysRevE.64.016131>
- Newman MEJ (2001b) Scientific collaboration networks. II. Shortest paths, weighted networks, and centrality. *Phys Rev E* 64:016132. <https://doi.org/10.1103/PhysRevE.64.016132>
- Newman MEJ (2004) Fast algorithm for detecting community structure in networks. *Phys Rev E* 69:066133. <https://doi.org/10.1103/PhysRevE.69.066133>
- Newman MEJ, Girvan M (2004) Finding and evaluating community structure in networks. *Phys Rev E* 69:026113. <https://doi.org/10.1103/PhysRevE.69.026113>
- Pons P, Latapy M (2006) Computing communities in large networks using random walks. *J Graph Algorithms Appl* 10:191–218. <https://doi.org/10.7155/jgaa.00124>
- Raghavan UN, Albert R, Kumara S (2007) Near linear time algorithm to detect community structures in large-scale networks. *Phys Rev E* 76:036106. <https://doi.org/10.1103/PhysRevE.76.036106>
- Rosvall M, Bergstrom CT (2008) Maps of random walks on complex networks reveal community structure. *Proc Natl Acad Sci USA* 105:1118–1123. <https://doi.org/10.1073/pnas.0706851105>
- Schilling MA, Phelps CC (2007) Interfirm collaboration networks: The impact of large-scale network structure on firm innovation. *Manage Sci* 53:1113–1126. <https://doi.org/10.1287/mnsc.1060.0624>
- Singh J (2005) Collaborative networks as determinants of knowledge diffusion patterns. *Manage Sci* 51:756–770. <https://doi.org/10.1287/mnsc.1040.0349>
- Sorenson O, Rivkin JW, Fleming L (2006) Complexity, networks and knowledge flow. *Res Policy* 35:994–1017. <https://doi.org/10.1016/j.respol.2006.05.002>
- Stolpe M (2002) Determinants of knowledge diffusion as evidenced in patent data: the case of liquid crystal display technology. *Res Policy* 31:1181–1198. [https://doi.org/10.1016/S0048-7333\(01\)00192-5](https://doi.org/10.1016/S0048-7333(01)00192-5)
- Tóth G, Juhász S, Elekes Z, Lengyel B (2018) Inventor collaboration and its persistence across European regions. preprint arXiv:1807.07637. <https://arxiv.org/abs/1807.07637>
- Trajtenberg M (1990) A penny for your quotes: patent citations and the value of innovations. *RAND J Econom* 21:172–187. <https://doi.org/10.2307/2555502>
- Trajtenberg M, Henderson R, Jaffe A (1997) University versus corporate patents: A window on the basicness of invention. *Econ Innov New Technol* 5:19–50. <https://doi.org/10.1080/10438599700000006>
- US Patent Classification (2013). <https://www.uspto.gov/web/patents/classification/selectnumwithtitle.htm>. Accessed 8 Nov 2018
- USPTO Bulk Data Storage System (2016) Patent Technology Monitoring Team - Custom Bibliographic Patent Data Extract DVD-ROM Available for Download (2014). Version of 2016-09-21. <https://bulkdata.uspto.gov/data/patent/ptmtdvd/>. Accessed 8 Nov 2018
- Welch BL (1947) The generalization of 'Student's' problem when several different population variances are involved. *Biometrika* 34:28–35. <https://doi.org/10.1093/biomet/34.1-2.28>

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)
